

INSPIRED2: An Improved Dataset for Sociable Conversational Recommendation

Ahtsham Manzoor*, Dietmar Jannach

University of Klagenfurt, Universitätsstraße 65-67, Klagenfurt am Wörthersee, 9020, Austria

Abstract

Conversational recommender systems (CRS) that are able to interact with users in natural language often utilize recommendation dialogs which were previously collected with the help of paired humans, where one plays the role of a *seeker* and the other as a *recommender*. These recommendation dialogs include items and entities that indicate the users' preferences. In order to precisely model the seekers' preferences and respond consistently, CRS typically rely on item and entity annotations. A recent example of such a dataset is INSPIRED, which consists of recommendation dialogs for *sociable* conversational recommendation, where items and entities were annotated using *automatic* keyword or pattern matching techniques. An analysis of this dataset unfortunately revealed that there is a substantial number of cases where items and entities were either wrongly annotated or annotations were missing at all. This leads to the question to what extent automatic techniques for annotations are effective. Moreover, it is important to study impact of annotation quality on the overall effectiveness of a CRS in terms of the quality of the system's responses. To study these aspects, we manually fixed the annotations in INSPIRED. We then evaluated the performance of several benchmark CRS using both versions of the dataset. Our analyses suggest that the improved version of the dataset, i.e., INSPIRED2, helped increase the performance of several benchmark CRS, emphasizing the importance of data quality both for end-to-end learning and retrieval-based approaches to conversational recommendation. We release our improved dataset (INSPIRED2) publicly at <https://github.com/ahtsham58/INSPIRED2>.

Keywords

Conversational Recommender Systems, data quality, annotations, evaluation, dialog systems

1. Introduction

Sociable conversational recommender systems (CRS) aim to build rapport with users while interacting with them in natural language [1, 2]. CRS that rely on natural language processing (NLP) nowadays commonly utilize datasets of previously recorded dialogs between humans, where one plays the role of a *recommendation-seeker* and the other as *human-recommender*, see e.g., [3]. However, due to a certain lack of rich sociable interactions in such datasets [4], it can be challenging to build a sociable CRS that builds rapport with the users using such limited data.

Therefore, it is important to develop datasets like INSPIRED [1], which includes dialogs that implement rich social communication strategies. Such rich datasets represent a solid basis to develop trustable CRS that are able to engage users in a natural and user-adaptive manner. Another key factor for building high-quality CRS lies in the proper recognition of the named entities and other

concepts that appear in the dialogs. In the movies domain, for example, being able to exactly identifying the items (i.e., movies) and related entities and concepts (e.g., actors or genres) can play a pivotal role for building an effective system. Existing CRS for example arrange such entities and their relationships as graphs [5, 6], and these relationships often form the basis to model the users' preferences, e.g., [7, 8, 9]. Moreover, domain specific concepts and entities can also contribute to the generation of meaningful and coherent responses, especially in knowledge-aware CRS, see [10, 11, 12, 13].

Annotating items and entities can be a laborious and economically expensive process [14, 15]. Human costs are high and may even be prohibitive for domains where particular knowledge or expertise are required to accomplish the annotation task [16]. In that context, the quality of the resulting annotations is crucial, and factually wrong annotations can lead to errors or ambiguity for the downstream task. Automating the annotation task or at least automatically verifying the annotations [14] has therefore been in the focus of research for several years.

We note here that data quality is crucial both for recent *generation-based* CRS approaches as well as for *retrieval-based* approaches to build natural language conversational systems [17]. For both types of systems, the question arises to what extent better data quality, i.e., having correct annotations and noise-free conversations, leads to better results in terms of the quality of the responses returned by a system for a given user utterance, e.g., in terms of consistency and plausibility.

4th Edition of Knowledge-aware and Conversational Recommender Systems (KaRS) Workshop @ RecSys 2022, September 18–23 2023, Seattle, WA, USA.

*Corresponding author.

✉ ahtsham.manzoor@aau.at (A. Manzoor);

dietmar.jannach@aau.at (D. Jannach)

🌐 <https://ahtsham58.github.io/> (A. Manzoor); <https://www.aau.at/en/aics/research-groups/infsys/team/dietmar-jannach/> (D. Jannach)

🆔 0000-0001-9418-753 (A. Manzoor); 0000-0002-4698-8507

(D. Jannach)

© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

In this work, we study the recent INSPIRED dataset, in which the items and entities that were mentioned in the recorded utterances are explicitly annotated. These annotations were created with the help of automatic approaches using keyword or pattern matching methods. However, looking at the data, we observed a substantial number of cases where items and entities were either wrongly annotated or missing annotations at all, e.g., “My favorite [MOVIE_GENRE_1] are Groundhogs Day, [MOVIE_TITLE_2] and Borat”. In addition, there were several cases where the utterances included noise, e.g., “How did you like QUOTATION_MARKHustlersQUOTATION_MARK?”. Finally, we found instances where regular words were identified as being named entities. In this latter case, human annotations would in fact have been required.¹ Overall, such issues may limit the quality of any CRS that is built on top of such data.

To understand the severity of the problem and the potential effects of data issues on the quality of a CRS, we have *manually* corrected the dataset by fixing the annotations and by removing noise from the utterances. Then, we conducted offline experiments and human evaluations to compare the performance of different benchmark CRS when using the original (INSPIRED) and improved (INSPIRED2) datasets. Overall, the results of our analyses indicate that all CRS showed better performance in different dimensions when built on INSPIRED2. In order to facilitate the design and development of future sociable CRS, we release the INSPIRED2 dataset online at <https://github.com/ahtsham58/INSPIRED2>.

2. Related Work

In this section, we first discuss datasets and aspects of data quality in the context of CRS. Afterwards, we review different design paradigms for building CRS, followed by a discussion of predominant evaluation approaches for such systems.

Datasets and Data Quality Research interest in CRS has experienced a substantial growth in recent years, see [18, 19] for related surveys. Many current systems interact with users in natural language, and one important goal for such system is to enable them to engage in conversations that reflect human behavior. Since many of these recent systems are built on recorded dialogs between humans, the capabilities of the resulting CRS depend on the richness of the communication in the datasets, e.g., in terms of the user *intents* that can be found in the conversations, see [20] for an detailed analysis of such intents.

¹Consider the movie “It (2017)” as an example of a difficult case, e.g., when appearing in an utterance like “Have you seen It?”.

A number of new datasets for conversational recommendation were published in recent years, e.g., [3, 21, 22, 23]. Such datasets, which are commonly collected with the help of crowdworkers, can however have limitations and may be not fully representative in terms of what we would observe in reality. In some cases, for example, crowdworkers were instructed to mention a minimum number of movies in the conversations. This leads to mostly “instance-based” conversations, where crowdworkers rather mention individual movies they like than their preferred genres, see also, [3, 24, 25].

Another problem when creating such datasets lies in the recognition and annotation of named entities appearing in the conversations, as mentioned above. Annotating entities in textual data can be a tedious process that may require a substantial amount of manual effort and time. To overcome this challenge, researchers sometimes adapt a semi-automatic approach or rely on NLP-assisted tools that visualize the entities in a text in order to reduce the required manual effort [16, 15, 26]. Generally, some automatic approaches may experience problems to correctly create annotations because human judgments and opinions are required. An automated approach was used in the context of the INSPIRED dataset. Here, the items and entities were annotated using keyword or pattern matching approaches. However, verifying the outcomes of such automatic or semi-automatic approaches can again be laborious and require manual effort.

Today, structured annotations for items and entities mentioned in the conversations are common in recent datasets. For example, in the case of the ReDial dataset [3], the mentioned movie titles were annotated with unique IDs. However, the ReDial dataset has some limitations. Various meta-data concepts (e.g., genres, actors, or directors) were not annotated. Moreover, the recorded dialogs include limited social interactions or explanations for the made recommendations. On the other hand, the INSPIRED dataset includes rich sociable conversation and explanation strategies for the recommended items. Also, aspects like movie genres or actors were explicitly annotated too. A comparison of these differences can be found in [1]. The key statistics of the INSPIRED dataset are shown in Table 1.

As mentioned earlier, the INSPIRED dataset has some limitations. The keyword or pattern matching approach used for the annotations might for example not detect misspelled keywords or concepts in an utterance. Moreover, data anomalies such as noisy utterances or ill-formed language can deteriorate the performance of an annotating algorithm, leading to challenges for the downstream use of the dataset [15, 27, 28]. In reality, the level of noise can be substantial both in real-world applications and in purposefully created datasets. Therefore, data quality assurance is often considered a significant and important step in NLP applications.

Table 1
Main Statistics of INSPIRED

	Total
Number of dialogs (conversations)	1,001
Average turns per dialog	10.73
Average tokens per utterance	7.93
Number of human-recommender utterances	18,339
Number of seeker utterances	17,472

Building Conversational Recommender Systems

Research on CRS has made substantial progress in terms of their underlying technical approaches. Some early commercial system such as *Advisor Suite* [29] for example relied on an entirely knowledge-based approach for the development of adaptive and personalized applications. Similarly, early critiquing-based systems were based on detailed knowledge about item features and possible critiques and had limited learning capabilities [30, 31].

Technological advancements particularly in fields like NLP, speech recognition, and machine learning in general led to the design of today’s end-to-end learning-based CRS. In such approaches, recorded recommendation dialogs between paired humans are used to train the deep neural models, see, e.g., [8, 9, 10, 12]. Given the last *user* utterance and the history of the ongoing dialog history, these trained models are then used to *generate* responses in natural language. These responses can either include item recommendations, which are also computed with the help of machine learning techniques, or other types of conversational elements, e.g., greetings.

In terms of the underlying data, the DeepCRS [3] system was built on the ReDial dataset, which was created in the context of this work. Later on, systems were developed which also relied on this dataset as well but included additional information sources, e.g., from *DBPedia* or *ConceptNet* [32, 12], to build knowledge graphs that are then used to improve the generated utterances. A number of works also makes use of pretrained language models like BERT [33] and subsequently fine-tune them using the recommendation dialogs, see, e.g., [34]. A related approach was adapted by the authors of INSPIRED, in which they proposed two variants of a conversational system, *with* and *without* strategy labels.

Unlike generation-based systems, in retrieval-based CRS the idea is to retrieve and adapt suitable responses from the dataset of recorded dialogs. One main advantage of retrieval-based approaches is that the retrieved responses were genuinely made by humans and thus are grammatically usually correct and in themselves semantically meaningful [35]. Recent examples of such retrieval-based systems are RB-CRS [17] and CRB-CRS [36], which we designed and evaluated based on the ReDial dataset in our own previous work.

CRS Evaluation Evaluating a CRS is a multi-faceted and challenging problem as it requires the consideration of various quality dimensions. An in-depth discussion of evaluation approaches for CRS can be found in [37]. Like in the recommender systems literature in general, computational experiments that do not involve humans in the loop are the predominant instrument to assess the quality of a CRS. Common metrics to evaluate the quality of the recommendations include *Recall*, *Hit Rate*, or *Precision* [8, 21, 38]. Moreover, certain linguistic aspects such as fluency or diversity are often evaluated with offline experiments as well to assess the quality of the generated responses. Common metrics in this area include *Perplexity*, *distinct N-Gram*, or the *BLEU* score [3, 8, 13, 22].

Given the interactive nature of CRS, offline experiments and the corresponding metrics have their limitations. Mainly, it is not always clear if the results obtained from offline experiments are representative of the user-perceived quality of the recommendations or system responses in general [35]. For example, when using metrics like the *BLEU* score, usually a system response is compared with one particular given ground truth. Such a comparison has limitations when used to estimate the average quality of a system’s responses, because there might be many different alternative responses that might be suitable as well in an ongoing dialog. Still, offline evaluations have their place and value. They can for example be informative for assessing particular aspects such as the number of items or entities that appear in an utterance or conversation.

Overall, given the limitations of pure offline experiments, researchers often follow a mixed approach where some aspects of the system are evaluated offline and some with humans. Typical quality aspects in terms of human perceptions in such combined approaches include the assessment of the *meaningfulness* or *consistency* of the system responses [1, 8, 12, 13, 36].

3. Data Annotation Methodology

During the creation of the INSPIRED [1] dataset, items and other entities were annotated in an automated way, as described above. For example, genre keywords were annotated using a regular expression to match a set of predefined tokens. Regarding actors and directors and other entities, a pattern matching technique was used, where words starting with a capital letter were searched in the TMDb database². A similar technique was used for movie titles. However, as mentioned, we observe a large number of cases where items and entities were either wrongly annotated or missing annotations. To answer our research question on the impact of the quality of the

²<https://www.themoviedb.org>

underlying data on the quality of the responses of a CRS, we fixed the annotations as follows.

Procedure To fix the annotations, we interviewed a number of university students to assess their knowledge in the movies domain and their ability to do the correction task. Subsequently, we hired two students and instructed them on how to annotate and clean the dataset. First, they were briefed on the logical format of the original annotations and how to retain that format. Second, they were asked to read each utterance individually, to detect potential noise, and to analyze which items or entities (e.g., title, genre, actor, or director) are mentioned in it.

In case of ambiguity or obscurity, they were allowed to access online portals, e.g., IMDb³. Note that regarding the genres, a set of 27 keywords was provided to them, which we curated and used in our earlier research [36]. After the briefing, the dataset was split evenly for both annotators. On weekly basis, their performance and the accuracy of the annotations was checked by one of the authors. Finally, after annotating the complete dataset, a number of additional validation steps were applied.

First, using a Python script, we ensured that every placeholder is enclosed by '[' and ']' as was done originally, e.g., [MOVIE_TITLE_1]. Second, another thorough manual examination of the entire improved dataset was performed to fix any missing annotations or noise. In that context, we also double-checked the consistency of the format and of the annotations.

The INSPIRED2 Dataset In total, 1,851 new annotations were added to INSPIRED, leading to the INSPIRED2 dataset. The most mistakes or inconsistencies were found for the items, i.e., movie titles, which is the most pertinent information for developing a CRS. We present the statistics about *new* annotations in Table 2. Overall, we added around 20% new annotations in INSPIRED2. The number of issues that were fixed, e.g., duplicate annotations in an utterance, noise or factually wrong information in the original annotations, are not shown in the presented statistics. We release the INSPIRED2 both in the TSV and JSON format online.

Table 2
Statistics about new annotations added in INSPIRED2

	Total	% Increase
Number of movie titles	966	22.0
Number of movie genres	206	5.0
Number of actors, directors, etc.	519	49.0
Number of movie plots	160	54.6
Number of new annotations	1851	18.9

³<https://www.imdb.com/>

Observed Issues During the annotation process, we recorded the observed issues in the original annotations. Since the original annotations were created using automatic techniques, many issues were related to the limitations of the simple keyword or pattern matching techniques. Overall, we observed a number of cases where minor spelling mistakes or incomplete movie titles made the exact string matching approaches ineffective.

For example, in one of the utterances, “*I think I am waiting for Star Wars The Rise of Skywalker*”, the annotation was missing because the correct title is “*Star Wars: Episode IX – The Rise of Skywalker*”. Similarly, we observe a significant number of cases where an utterance was only partially annotated, e.g., “*ok is it scary like incidious or [MOVIE_GENRE_2] [MOVIE_TITLE_5]*”. In addition, at places where two entities were separated with ‘/’ instead of a space, the automatic technique often failed to create proper annotations, e.g., “*Since you like [MOVIE_GENRE_1] drama/mystery, I’m going to send you the trailer to the movie [MOVIE_TITLE_3]*”.

Also, the automatic approach used for INSPIRED sometimes had difficulties to deal with ambiguity. We found a number of cases where a regular word was annotated, although such a word did not belong to any item or entity. For example, in one of the cases, “*Are you interested in a current movie in the box office?*”, the utterance was annotated as “*Are you interested in a current movie in the box [MOVIE_TITLE_0]*”, where the word ‘office’ was mistakenly annotated as an item, i.e., *The Office (2005)*.

Overall, the main observed issues are the following.

1. Missing annotations for movie titles, genres, actors, movie plots, etc.
2. Partially annotated items and entities such as movie titles, or genres in an utterance.
3. Factually wrong annotations for movie titles.
4. Inconsistent indexing for the annotated items and entities.
5. Mistaken annotations for plain text, e.g., *family*, *box office*; human annotations may be required here.
6. Parts of the utterance or a few keywords were omitted during the annotation process.

4. Evaluation Methodology

We performed both offline experiments as well as a human evaluation to assess the impact of data quality on the quality of the responses of a CRS.

Offline Evaluation of Recommendation Quality We included the following recent end-to-end learning approaches in our experiments: DeepCRS [3], KGSF [12],

TG-ReDial [22], and the INSPIRED model *without strategy labels*⁴ [1]. This selection of models covers various design approaches for CRS, e.g., using an additional knowledge graph or not. We used the open-source toolkit CRSLab⁵ for our evaluations. This framework was used in earlier research as well, for example in [10, 39, 40]. For our analyses, we first trained the aforementioned CRS models using the original split ratio, i.e., 8:1:1, for each dataset. Afterwards, given the trained models and test data for each dataset, we ran three trials for each CRS and subsequently averaged the results for offline evaluation metrics. Note that the same procedure was adapted for both versions of the dataset, i.e., INSPIRED and INSPIRED2.

User Study on Linguistic Quality We conduct a user study to compare the perceived quality of system responses using either INSPIRED and INSPIRED2. Specifically, we randomly sampled *same* 50 dialog situations from each dataset. To create the dialog continuations, we used the *retrieval-based* CRS approaches, RB-CRS and CRB-CRS, which we proposed in our earlier work, see [36].

In order to obtain fine-grained assessments, three human judges⁶ were involved. The specific task of the judges was to assess (rate) the *meaningfulness* of a system response as a proxy of its *quality* and *consistency* in a dialog situation, see [3, 41, 12]. Note that in this study we did not explicitly assess the quality of the specific item recommendations. Instead, the focus of this study was to understand the impact of the improved underlying dataset on the linguistic quality and the consistency of the generated responses.

We used a 3-point scale for these ratings, from ‘*Completely meaningless (1)*’ to ‘*Somewhat meaningless and meaningful (2)*’ to ‘*Completely meaningful (3)*’. The human judges were provided with specific instructions on how to evaluate the meaningfulness of a response, e.g., they should assess if a response represents a logical dialog continuation and evaluate the overall language quality of the given response. Overall, the human judges were provided 50 dialogs (446 responses to rate) that were produced using the INSPIRED and INSPIRED2 datasets. We also explained the meanings and purpose of various placeholders contained in the responses to the human judges. Moreover, to avoid any bias in the evaluation process, the judges were not made aware which response was created for which dataset by which CRS. Also, the order of the dialogs and the system responses were randomized.

⁴The INSPIRED *with strategy labels* model was not publicly available.

⁵<https://github.com/RUCAIBox/CRSLab>

⁶These judges were PhD students and were different than the ones who fixed the annotations.

5. Results

Recommendation Quality Table 3 shows the accuracy results for the evaluated CRS models. Specifically, we provide the results for the different benchmark CRS models in terms of the performance *difference* when using the original and improved annotations. Overall, we can observe an almost consistent gain in performance for all models and on all metrics except Hit@50 when the improved dataset is used. The obtained improvements can be quite substantial, indicating that improved data quality can be helpful for CRS of different types, including (i) CRS, which do not rely on additional knowledge sources, (ii) CRS that leverage additional knowledge sources, (iii) CRS that are guided by a topic policy, and (iv) CRS that rely on pre-trained language models like BERT.

Interestingly, we see negative effects for two measurements in which Hit@50 is used as a metric. A deeper investigation of this phenomenon is needed, in particular as the other metrics at this (admittedly rather uncommon) list length, MRR@50 and NDCG@50, indicate that the improved dataset is helpful to increase recommendation accuracy. At the moment, we can only speculate that the improved annotations in the ongoing dialog histories led to more diverse or niche recommendations compared to the original dataset. We might assume that the missing annotations in many cases referred to less popular movies, so that the recommendations without the improved annotations will more often recommend popular movies, which is commonly advantageous in terms of hit rate and recall.

Linguistic Quality We recall that three human evaluators assessed the linguistic quality of the system responses (dialog continuations), which were created either based on the INSPIRED or the INSPIRED2 dataset. As underlying CRS systems, we considered the retrieval-based approaches RB-CRS and CRB-CRS, as mentioned above. For our analysis, we averaged the scores by the three evaluators. Table 4 shows the mean ratings across all dialog situations as well as the standard deviations. We find that also in the case of retrieval-based approaches, improving the quality of the underlying dataset was helpful, leading to higher mean scores, without observing larger standard deviations. A Student’s t-test reveals that the observed differences in the means are statistically significant ($p < 0.001$).⁷

Comparison of Knowledge Concepts in Responses To understand the impact of the new annotations on the responses in terms of the richness of knowledge concepts, we compute the number of items and entities that

⁷We provide the data and compiled results of our study online.

Table 3

Accuracy results obtained in the *offline* evaluation. V1 represents INSPIRED, V2 denotes INSPIRED2, and “% Change” represents the actual performance gain/loss when using INSPIRED2 compared to INSPIRED.

		Hit@1	Hit@10	Hit@50	MRR@1	MRR@10	MRR@50	NDCG@1	NDCG@10	NDCG@50
DeepCRS [3]	V1	0.0006	0.0464	0.1726	0.0065	0.0148	0.0193	0.0065	0.0220	0.0478
	V2	0.0256	0.0578	0.1222	0.0256	0.0306	0.0333	0.0256	0.0366	0.0504
	% Change	4161.11	24.50	-29.20	294.95	106.99	72.48	294.95	66.08	5.31
KGSF [12]	V1	0.0022	0.0216	0.0744	0.0032	0.0061	0.0084	0.0022	0.0097	0.0211
	V2	0.0066	0.0303	0.0587	0.0057	0.0123	0.0134	0.0066	0.0165	0.0223
	% Change	207.27	40.46	-21.11	75.58	100.97	58.40	207.27	70.36	5.78
TG-ReDial [22]	V1	0.0365	0.1149	0.2344	0.0365	0.0572	0.0626	0.0365	0.0707	0.0967
	V2	0.0511	0.1315	0.2417	0.0511	0.0742	0.0792	0.0511	0.0877	0.1118
	% Change	40.00	14.46	03.12	40.00	29.64	26.48	40.00	24.05	15.51
INSPIRED [1]	V1	0.0151	0.0550	0.1532	0.0151	0.0241	0.0286	0.0151	0.0312	0.0527
	V2	0.0194	0.0734	0.1855	0.0194	0.0293	0.0353	0.0194	0.0392	0.0650
	% Change	28.57	33.33	21.13	28.57	21.59	23.44	28.57	25.44	23.28

Table 4

Results of Human Evaluation

		INSPIRED	INSPIRED2
RB-CRS	Average score	2.30	2.46
	Std. deviation	0.62	0.59
CRB-CRS	Average score	2.31	2.46
	Std. deviation	0.55	0.55

appeared in the system responses. Specifically, we compute the number of placeholders in the responses, before they would be replaced by the recommendation component, see [36]. In Table 5, we present the statistics for RB-CRS and CRB-CRS for both dataset versions. Overall, we find that the responses for the improved dataset contain between 20% and 27% more concepts and entities. We note that an increase in concepts is expected, as INSPIRED2 has almost 20% more annotations. However, the important observation here is that the retrieval-based CRS approaches actually surfaced these richer system responses frequently.

Table 5

Number of Items and Entities included in Responses

	INSPIRED	INSPIRED2	% Increase
RB-CRS	174	222	27.6
CRB-CRS	208	251	20.7

BLEU Score Analysis Finally, in order to understand to what extent (offline) linguistic scores correlate with the perceived quality of responses as was done in [1, 8], we performed an analysis of the BLEU scores obtained for the different datasets. Specifically, given a system response and the corresponding *ground truth* response, we preprocess both sentences and compute the BLEU scores for $N = \{1, 2, 3, 4\}$ grams. We provide the results of this analysis online. In sum, the analysis shows that

the BLEU scores generally improve when the underlying data quality is higher, i.e., in the case of the INSPIRED2 dataset. These findings are thus well aligned with the outcomes of our human evaluation study, where using INSPIRED2 as an underlying dataset turned out to be favorable.

6. Conclusion

Datasets containing recorded dialogs between humans are the basis for many modern CRS. In this work, we have analyzed the recent INSPIRED dataset, which was developed to build the next generation of sociable CRS. We found that automatic entity and concept labeling has its limitations and we have improved the quality of the dataset through a manual process. We then conducted both computational experiments as well as experiments with users to analyze to what extent improved data quality impacts recommendation accuracy and the quality perception of the system’s responses by users. The analyses clearly indicate the benefits of improved data quality across different technical approaches for building CRS. We release the improved dataset publicly and hope to thereby stimulate more research in sociable conversational recommender systems in the future.

References

- [1] S. A. Hayati, D. Kang, Q. Zhu, W. Shi, Z. Yu, IN-SPIRED: Toward sociable recommendation dialog systems, in: EMNLP '20, 2020.
- [2] F. Pecune, L. Callebort, S. Marsella, A socially-aware conversational recommender system for personalized recipe recommendations, in: Proceedings of the 8th International Conference on Human-Agent Interaction, HAI '20, 2020, p. 78–86.
- [3] R. Li, S. E. Kahou, H. Schulz, V. Michalski, L. Charlin, C. Pal, Towards deep conversational recommendations, in: NIPS '18, 2018, pp. 9725–9735.
- [4] D. Jannach, L. Chen, Conversational Recommendation: A Grand AI Challenge, *AI Magazine* 43 (2022).
- [5] M. Di Bratto, M. Di Maro, A. Origlia, F. Cutugno, Dialogue analysis with graph databases: Characterising domain items usage for movie recommendations (2021).
- [6] C.-M. Wong, F. Feng, W. Zhang, C.-M. Vong, H. Chen, Y. Zhang, P. He, H. Chen, K. Zhao, H. Chen, Improving conversational recommender system by pretraining billion-scale knowledge graph, in: ICDE '21, 2021, pp. 2607–2612.
- [7] Y. Cao, X. Wang, X. He, Z. Hu, T.-S. Chua, Unifying knowledge graph learning and recommendation: Towards a better understanding of user preferences, in: WWW '19, 2019, pp. 151–161.
- [8] Q. Chen, J. Lin, Y. Zhang, M. Ding, Y. Cen, H. Yang, J. Tang, Towards knowledge-based recommender dialog system, in: EMNLP-IJCNLP '19, 2019, pp. 1803–1813.
- [9] J. Zhou, B. Wang, R. He, Y. Hou, CRFR: Improving conversational recommender systems via flexible fragments reasoning on knowledge graphs, in: EMNLP '21, 2021, pp. 4324–4334.
- [10] K. Chen, S. Sun, Knowledge-based conversational recommender systems enhanced by dialogue policy learning, in: IJCKG '21, 2021, pp. 10–18.
- [11] A. Wang, C. D. V. Hoang, M.-Y. Kan, Perspectives on crowdsourcing annotations for natural language processing, *Language resources and evaluation* 47 (2013) 9–31.
- [12] K. Zhou, W. X. Zhao, S. Bian, Y. Zhou, J.-R. Wen, J. Yu, Improving conversational recommender systems via knowledge graph based semantic fusion, in: KDD '20, 2020, pp. 1006–1014.
- [13] Y. He, L. Liao, Z. Zhang, T.-S. Chua, Towards enriching responses with crowd-sourced knowledge for task-oriented dialogue, in: MuCAI '21, 2021, p. 3–11.
- [14] T. Arjannikov, C. Sanden, J. Z. Zhang, Verifying tag annotations through association analysis, in: ISMIR '13, 2013, pp. 195–200.
- [15] J. S. Grosman, P. H. Furtado, A. M. Rodrigues, G. G. Schardong, S. D. Barbosa, H. C. Lopes, Eras: Improving the quality control in the annotation process for natural language processing tasks, *Information Systems* 93 (2020) 101553.
- [16] B. C. Benato, J. F. Gomes, A. C. Telea, A. X. Falcão, Semi-automatic data annotation guided by feature space projection, *Pattern Recognition* 109 (2021) 107612.
- [17] A. Manzoor, D. Jannach, Generation-based vs. retrieval-based conversational recommendation: A user-centric comparison, in: RecSys '21, 2021.
- [18] D. Jannach, A. Manzoor, W. Cai, L. Chen, A survey on conversational recommender systems, *ACM Computing Surveys* 54 (2021) 1–36.
- [19] C. Gao, W. Lei, X. He, M. de Rijke, T.-S. Chua, Advances and challenges in conversational recommender systems: A survey, *AI Open* 2 (2021) 100–126.
- [20] W. Cai, L. Chen, Predicting user intents and satisfaction with dialogue-based conversational recommendations, in: UMAP '20, 2020, p. 33–42.
- [21] D. Kang, A. Balakrishnan, P. Shah, P. Crook, Y.-L. Boureau, J. Weston, Recommendation as a communication game: Self-supervised bot-play for goal-oriented dialogue, in: EMNLP-IJCNLP '19, 2019, pp. 1951–1961.
- [22] K. Zhou, Y. Zhou, W. X. Zhao, X. Wang, J.-R. Wen, Towards topic-guided conversational recommender system, in: ICCL '20, 2020, pp. 4128–4139.
- [23] Z. Fu, Y. Xian, Y. Zhu, Y. Zhang, G. de Melo, COOKIE: A dataset for conversational recommendation over knowledge graphs in e-commerce, 2020. [arXiv:2008.09237](https://arxiv.org/abs/2008.09237).
- [24] K. Christakopoulou, F. Radlinski, K. Hofmann, Towards conversational recommender systems, in: KDD '16, 2016, pp. 815–824.
- [25] X. Ren, H. Yin, T. Chen, H. Wang, Z. Huang, K. Zheng, Learning to ask appropriate questions in conversational recommendation, in: SIGIR '21, 2021, pp. 808–817.
- [26] P. Stenetorp, S. Pyysalo, G. Topić, T. Ohta, S. Ananiadou, J. Tsujii, Brat: a web-based tool for nlp-assisted text annotation, in: ACL '12, 2012, pp. 102–107.
- [27] C. Zong, R. Xia, J. Zhang, Data Annotation and Preprocessing, 2021, pp. 15–31.
- [28] P. Röttger, B. Vidgen, D. Hovy, J. B. Pierrehumbert, Two contrasting data annotation paradigms for subjective nlp tasks, 2021. [arXiv:2112.07475](https://arxiv.org/abs/2112.07475).
- [29] D. Jannach, ADVISOR SUITE – A knowledge-based sales advisory system, in: ECAI '04, 2004, pp. 720–724.
- [30] K. McCarthy, Y. Salem, B. Smyth, Experience-based critiquing: Reusing critiquing experiences to im-

- prove conversational recommendation, in: ICCBR '10, 2010, pp. 480–494.
- [31] L. Chen, P. Pu, Critiquing-based recommenders: survey and emerging trends, *User Modeling and User-Adapted Interaction* 22 (2012) 125–150.
- [32] Q. Chen, J. Lin, Y. Zhang, H. Yang, J. Zhou, J. Tang, Towards knowledge-based personalized product description generation in e-commerce, in: *KDD '19*, 2019, pp. 3040–3050.
- [33] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, in: *NAACL-HLT*, 2019.
- [34] L. Wang, H. Hu, L. Sha, C. Xu, K.-F. Wong, D. Jiang, Finetuning large-scale pre-trained language models for conversational recommendation with knowledge graph, 2021. [arXiv:2110.07477](https://arxiv.org/abs/2110.07477).
- [35] A. Manzoor, D. Jannach, Conversational recommendation based on end-to-end learning: How far are we?, *Computers in Human Behavior Reports* (2021) 100139.
- [36] A. Manzoor, D. Jannach, Towards retrieval-based conversational recommendation, *Information Systems* (2022) 102083.
- [37] D. Jannach, Evaluating conversational recommender systems, *Artificial Intelligence Review* forthcoming (2022).
- [38] T. Zhang, Y. Liu, P. Zhong, C. Zhang, H. Wang, C. Miao, KECRS: Towards knowledge-enriched conversational recommendation system, 2021. [arXiv:2105.08261](https://arxiv.org/abs/2105.08261).
- [39] Y. Zhou, K. Zhou, W. X. Zhao, C. Wang, P. Jiang, H. Hu, C^2 -crs: Coarse-to-fine contrastive learning for conversational recommender system, in: *WSDM '22*, 2022, pp. 1488–1496.
- [40] Y. Li, B. Peng, Y. Shen, Y. Mao, L. Liden, Z. Yu, J. Gao, Knowledge-grounded dialogue generation with a unified knowledge representation, 2021. [arXiv:2112.07924](https://arxiv.org/abs/2112.07924).
- [41] D. Jannach, A. Manzoor, End-to-end learning for conversational recommendation: A long way to go?, in: *IntRS Workshop at RecSys '20*, Online, 2020.