

State of the Art of User Simulation approaches for conversational information retrieval

Pierre Erbacher¹, Laure Soulier¹ and Ludovic Denoyer²

¹Sorbonne Université, CNRS, LIP6, F-75005 Paris - France

²Facebook AI research

Abstract

Conversational Information Retrieval (CIR) is an emerging field of Information Retrieval (IR) at the intersection of interactive IR and dialogue systems for open domain information needs. In order to optimize these interactions and enhance the user experience, it is necessary to improve IR models by taking into account sequential heterogeneous user-system interactions. Reinforcement learning has emerged as a paradigm particularly suited to optimize sequential decision making in many domains and has recently appeared in IR. However, training these systems by reinforcement learning on users is not feasible. One solution is to train IR systems on user simulations that model the behavior of real users. Our contribution is twofold: 1) reviewing the literature on user modeling and user simulation for information access, and 2) discussing the different research perspectives for user simulations in the context of CIR.

Keywords

User simulation, conversational information retrieval, Evaluation, Reinforcement learning

1. Introduction

Conversational information retrieval (CIR) is a well-established research field in the Information Retrieval (IR) community [1]. It is at the intersection of many application domains (interactive IR, conversational systems, question answering) but differs in many aspects. 1) Unlike interactive IR [2] that is also interested in user feedback, CIR's objective is to find the relevant information in a more natural way (i.e., interactions in natural language). 2) It is also different from task-oriented conversation systems [3] which are guided in a particular application area (i.e., travel booking). A CIR system is a specific application case of task-oriented conversation systems, but is different because the search is performed in an open domain. 3) Finally, very close to question answering systems [4] in which questions are expressed in natural language, CIR focuses on exploratory or complex information needs, on open fields, not specifically centered on a specific entity, question or document as in Q&A systems.

The objective of a CIR system is to help the user in his/her interactions with a search engine through a conversational tool. This whole search engine/conversational system will aim to: 1) respond to a complex information need expressed in natural language and/or through various feedbacks (e.g., sequences

of reformulation of requests, clicks on documents, conversations in natural language) and 2) anticipate and guide the user in his/her sequence of search actions using, e.g., query clarification or document ranking [5, 6]. This requires CIR systems to interact with users to better understand and/or refine their information needs in a long term objective.

Most recent models focus on the interactive IR aspect that builds a user profile or session modeling [7, 8, 9]. However, they do not support on a long-term, sequential and incremental decision strategy. The latter has been outlined as essential for conversational recommendation systems [10]. A growing interest in the community is to design IR systems that embed strategies to best optimize these search sequences [11, 12], in particular through interactions in natural language with users [13]. Reinforcement learning has been shown to be particularly effective in optimizing sequential decision making [14] because it can learn a policy that maximizes a long term cumulative reward using successive interactions with its environment. The idea of using reinforcement learning to optimize an IR policy has become a hot topic [13, 15]. For IR, policies would be learned to optimize in the long term the sequences of user interactions with the CIR system according to different criteria : reducing/increasing the number of interactions with the user, increasing the diversity of retrieved information, predicting future needs, guiding users actions, etc. Ideally, CIR system policies should train with real users. In practice, this is impossible because reinforcement learning algorithms are currently very inefficient: it would take millions of interactions with users to see the emergence of a correct policy. Training directly on data limits the exploration to scenarios collected from real users, so CIR systems can only visit

Causality in Search and Recommendation (CSR) and Simulation of Information Retrieval Evaluation (Sim4IR) workshops at SIGIR, 2021
Editors of the proceedings (editors): Krisztian Balog, Xianjie Chen, Xu Chen, David Maxwell, Paul Thomas, Shuo Zhang, Yi Zhang, and Yongfeng Zhang

✉ pierre.erbacher@lip6.fr (P. Erbacher); laure.soulier@lip6.fr (L. Soulier); denoyer@fb.com (L. Denoyer)

© 2021 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)

a very small part of the state space. In addition, even if a lot of search trajectories (sequence of interactions between user and system) are collected, only a few would be optimal [16].

One of the solutions is to train CIR systems with user simulations (US). US generate synthetic data close to data generated by real users. This makes it possible to : 1) estimate the most probable actions of users, particularly in new scenarios and 2) best optimize CIR systems policies by exploring trajectories not contained in collected scenarios. In this article we are interested in the problem of user modeling and simulation in different contexts of access to information (e.g., IR, recommendation, conversational systems).

1. We review the different approaches for modeling and simulating user for information access and goal oriented conversational systems (Section 2). Modeling and simulation are often confused in the literature. For a source system, in our case the user whose behavior is unknown, the associated model is a set of instructions, rules, markov chain or equations constructed from the observed data coming from this user [17]. Simulation is an agent that generates behaviors by following the rules of a model. Together, the goal is to emulate user behavior [17].
2. We discuss state-of-the-art SU evaluation, the difficulty being to measure whether the simulated behavior is realistic (Section 3).
3. Finally, we discuss the research perspectives to build user simulations for CIR (Section 4).

2. Model and simulate user behavior

The IR community produced lot of work in studying and modelling users behaviors either based on cognitive or statistical approaches. First user modeling approaches in IR were based on cognitive approaches. For instance, Belkin [18] produces a description of the users characterized by an objective or a problem as well as a knowledge of the world. He also describes the conditions necessary for an IR system to be effective with ambiguous queries. Many derivatives were then proposed (e.g., ISP [19], IS&R [20], Ellis [21]). The second line of work mainly remains on statistical models describing users' click behavior and satisfaction during search sessions [22, 23, 24, 25]. Beyond being used for building predictive models (from the IR system point of view), these user models serve also as basis of user simulations. In what follows, we will focus on the literature review of user simulation models. Please note that since CIR is a young research domain, most of the cited work below belongs to orthogonal and related fields, namely recommendations and dialog.

Historically, US have been extensively studied for conversational systems [26, 27, 28, 29, 30]. One of the first attempts is to use statistical user models as a simulator [26]. This simulates the users from the appearance probabilities of a sentence according to the previous sentence. This simulation is limited, because it only takes into account the previous response of the system, the probabilities suffer from the lack of data and moreover it is not guided by a dialogue objective. To add consistency to users' responses, Scheffler et al. [30] introduces into its simulator a dialogue objective, drawn from a list of possible objectives. The model follows an imposed dialogue objective: the probabilities are conditioned by the objectives of the real users in the corpus.

Komatani et al. [27] brings diversity by conditioning the probability of utterance on the user's ability to use a search system as well as its degree of urgency. From annotated data, a decision tree is used to classify the sentences of the corpus according to the urgency and the competence of the users. For the simulation, a degree of urgency is drawn and imposed on the model for the generation of sentences. One way to implement rule based US is to use agenda, as in [31]. In this work, the authors build a conversational US using constraints and rules which make it possible to explicitly encode the user's state as well as his objective. In addition to generating sentences according to a certain user model, the simulator is assigned a dialogue objective as well as a dialogue strategy (Agenda) which conditions the sentences generated by the model.

Agenda-based approaches are widely used to train conversational recommendation systems because they allow to build dialogue strategies for close domain using simple rules [32, 33]. Zhang et al. [34] show that US Agendas can be particularly realistic: when recommending movies, 36% of human reviewers mistook an US for a real user and 23% of reviewers were undecided.

Li et al. [33] propose a more realistic hybrid model that relies on the agenda to maintain user dialogue strategy but which takes advantage of the power of language models *seq2seq* to improve the natural language generation and language understanding. Asri et al. [35] and Kreyssig et al. [16] propose a neural model that breaks free from the agenda by using *seq2seq* language models to generate and interpret natural language. On the other hand, they keep the generation of objectives (constraints / requests) at the semantic level. To simulate the dynamic behavior of the user, the constraints are drawn independently according to a probability observed in the dataset.

Some work relies on deep learning to simulate users from data [36, 37, 38]. For example, Zhao et al. [38] learn a user simulation which provides information on the relevance of the recommendations made by the system. They use an adversarial generative approach to converge to a US. The discriminator must differentiate between

simulated and real feedback based on user search history. The generator learns to create feedbacks based on a history of interactions. By relying on reinforcement learning approaches, Chen et al. [37] propose to jointly model the user click policy and the reward function using a generative approach. Indeed, the user policy is optimized to generate the actions which maximize the reward signal. In parallel, the reward function has the role of the discriminator, it is optimized to minimize the difference in reward between the action taken by the user policy and the action taken by a real user observed in the database. A data driven approach is based on reverse reinforcement learning, the objective of which is to learn the reward function from the behavioral observations of another (often human) agent [39]. This approach contrasts with previous models which relied on explicitly defined reward functions. In this context, Chandramohan et al. [36] propose to use inverse reinforcement learning to find the reward function describing user policies observed in data. An experiment is carried out using conversational data in the field of tourist information.

Some US relies on active learning. Because learning a user model requiring a large amount of data. For this, Zhang et al. [40] proposes to include sparsely human feedback in the learning loop. In this model, the user model is learned jointly with the recommendation system. The goal is to optimize training to limit human intervention. For this, the user model is allocated a limited budget. At each turn, the model can choose how to allocate her credits to respond to the recommendation system: 0 credit to respond using the learned policy, 1 credit to access a human response, 2 credit to collect a human discussion - human. The user model learns from the new data collected. At the start of each conversation, a goal generator pulls in unsuccessful or unexplored goals. Despite a parsimonious use of human interventions, including humans in the learning loop remains costly.

3. Evaluation of User Simulation

As explained earlier, one of the objectives of US is to evaluate [26] information retrieval systems. It is therefore needed to ensure the quality of the US used for the evaluation. According to [41], a good metric, to evaluate US, should be able to measure US quality independently of the task and should be correlated with real user evaluations. Following [41], we distinguish two types of evaluation.

Direct evaluation directly measures US through metrics. For example, an evaluation approach consists in measuring whether the policy has achieved its objective, the quality of sentences generated in a dialogue system thanks to BLEU / RED / METEOR measurements [42, 43, 44, 45], likely click behavior measured by perplexity [24]). However, in the context of dialogue sys-

tems, [46] demonstrates that automatic metrics that assess word overlap with reference sentences like BLUE, RED and METEOR have very little correlation with judgment human, in particular because there is a great diversity of valid answers for each round of conversation [46, 34]. Human evaluation is also used to evaluate the quality of the responses of a user simulation. For example, Zhang et al. [34] observe the number of times human evaluators cannot differentiate between real user responses and US-generated dialogues to compare the quality between US.

Indirect evaluation measures the impact of a US on the performance of the retrieval system [29]. Shi et al. [47] assess the realism of its user model thanks to direct and indirect human annotation. For direct evaluation, they directly judge the quality of sentences from different US through different criteria: the quality of the language (such as grammar), the consistency of the sentences between the different turns as well as the diversity of the sentences. The evaluators also give an overall rating of the simulations. The indirect evaluation of the US is measured through the performance of the recommender systems that have interacted with the US. Human reviewers analyze the quality of the recommendation. They rate their satisfaction, the naturalness of the conversations and the effectiveness of the recommendations.

This section summarizes how human annotations and automatic metrics are used in the literature to assess US quality via direct and indirect evaluations. [34] shows that assessing IR systems using US is possible. There are still lot of work to be done to have reliable and realist user simulation in order to benchmark IR systems.

4. Research perspectives

In this section, we discuss the research perspectives on user simulation models for CIR.

Towards multi-modal user simulation models for open domains (CIR). To generate more realistic search sessions, US should be able to express complex queries on many different domains and simulate document selection (click). Currently, conversational US do not embed user click models. However, clicks are present in many IR scenarios for example when the system is returning keywords for disambiguating the query or selecting document. Generating complete user behaviors in the form of sequences of clicks and requests or dialogue seems essential to have realistic evaluation frameworks for CIR systems. To overcome the lack of data, one idea would be to incorporate existing click models [24] into US [31]. Another idea is to jointly learn the language model and the user click model. Simulating both conversations and clicks seems essential for learning CIR systems that return a list of

documents. Also, CIR imposes to be able to evolve on a variety of domains (open domains). The vast majority of US are built for conversational recommendation systems (CRS) for a specific domain, for example, (e.g., restaurant suggestion). CIR systems need to offer services covering a growing number of areas. A first step towards this open domain would be to benefit from the research in multi-domain CRS [48], with the publication of several multi-domain datasets [48, 49].

Towards diversity-driven user simulation

One of the difficulties brought by neural models is to generate diversity in behavior. Indeed, such models learn to satisfy the average behaviors observed in the data. However, there is a strong variability in user behavior [50]. Some are more persistent or curious in their research, while others give up quickly. In addition, the context and world knowledge is different between users. Learning a single user model on data does not reflect this diversity of users. Recently proposed US, model one behavior from data generated by many real users. One of the ideas would be to generate realistic user policies but with different styles. One method is to regroup and model by hand the different users according to a certain number of criteria (curiosity, perseverance, world knowledge, expertise, ...) such as [27]. A second method is to extract automatically these informations from a corpus of conversations. Some works propose to use US with attributes and characteristics (persona) [51, 52] to obtain both diversity and consistency in requests and dialogues. Li et al. [51] used user embedding encoding speaker specific information (persona) such as age, dialect, context, gender and other personal information that influences generated responses. Balog et al. [52] proposed a more convenient way to encode information about users persona using knowledge graph. Ideally, CIR should use knowledge they have about the user in order to optimize this user's search session. In a cold start scenario, one possibility for the CIR is to probe users to extract a maximum number of information with a limited probing budget. We believe that using diversity driven US is essential to learn such probing policies.

Towards the establishment of benchmarks for CIR systems. Currently, human annotators are the most reliable benchmark to assess CIR systems performances. But, human evaluation is too time-consuming to test the robustness of models for various information needs and behaviors. One of the solutions is to establish US to measure the performance of the various proposed CIR models. Note that [34] proposed this evaluation framework for CR systems. In addition, CIR systems should also be evaluated to on their generalization capacity:

- To open domain conversation

- To various level of query complexity
- To various user behavior

For this, we believe that the benchmark should be a population of unseen US having different persona and behavior.

Learning users behaviors in CIR context directly from collected data might be limited because: 1) Collecting conversational search session data is computationally expensive. 2) Collected data is entangled with used the IR system. One idea is to build US with different level of abstractions. From the intention level to the natural language generation. In order to facilitate future works, the IR community should have a common platform for building and testing US for sequential IR systems. In this direction, some environments simulating users have already been developed for learning recommendation system using sequential learning such as [53, 54].

5. Conclusion

This article focuses on CIR and above all on setting up a training and evaluation framework for user-centered approaches. Although human assessment is the most suitable, it is however too time-consuming for models based on deep learning or reinforcement learning. The objective of this paper is to synthesize the state of the art of user modeling and simulation approaches in various information access tasks to open up prospects for the implementation of CIR systems.

6. Acknowledgments

This work is supported by the ANR JCJC project SESAMS (Projet-ANR-18-CE23-0001).

References

- [1] J. S. Culpepper, F. Diaz, M. D. Smucker, Research Frontiers in Information Retrieval: Report from the Third Strategic Workshop on Information Retrieval in Lorne (SWIRL 2018), SIGIR Forum 52 (2018) 34–90.
- [2] C. Zhai, Interactive Information Retrieval: Models, Algorithms, and Evaluation, in: SIGIR, 2020, p. 2444–2447.
- [3] A. Bordes, Y.-L. Boureau, J. Weston, Learning End-to-End Goal-Oriented Dialog, in: ICLR, 2017.
- [4] A. Bordes, S. Chopra, J. Weston, Question Answering with Subgraph Embeddings, in: EMNLP, 2014, pp. 615–620.
- [5] L. Azzopardi, M. Dubiel, M. Halvey, J. Dalton, Conceptualizing agent-human interactions during the conversational search process, 2018.

- [6] F. Radlinski, N. Craswell, A Theoretical Framework for Conversational Search, in: SIGIR, 2017, p. 117–126.
- [7] A. Sordoni, Y. Bengio, H. Vahabi, C. Lioma, J. G. Simonsen, J. Y. Nie, A hierarchical recurrent encoder-decoder for generative context-aware query suggestion, in: CIKM, 2015, pp. 553–562.
- [8] L. Zou, L. Xia, Y. Gu, X. Zhao, W. Liu, J. X. Huang, D. Yin, Neural interactive collaborative filtering, in: ACM SIGIR, SIGIR '20, ACM, 2020, p. 749–758.
- [9] M. Mitsui, C. Shah, Bridging gaps: Predicting user and task characteristics from partial user information, in: ACM SIGIR, SIGIR'19, ACM, 2019, p. 415–424.
- [10] Y. Sun, Y. Zhang, Conversational Recommender System, SIGIR (2018) 235–244.
- [11] Z. Tang, G. H. Yang, Dynamic search - optimizing the game of information seeking, CoRR (2019). [arXiv:1909.12425](https://arxiv.org/abs/1909.12425).
- [12] H. Yang, M. Sloan, J. Wang, Dynamic information retrieval modeling, in: ACM SIGIR, SIGIR '14, 2014, p. 1290.
- [13] W. Zhang, X. Zhao, L. Zhao, D. Yin, G. H. Yang, A. Beutel, Deep Reinforcement Learning for Information Retrieval: Fundamentals and Advances, in: SIGIR, 2020, pp. 2468–2471.
- [14] R. S. Sutton, A. G. Barto, Reinforcement Learning: An Introduction, MIT Press, 1998.
- [15] A. MontazerAlghaem, H. Zamani, J. Allan, A Reinforcement Learning Framework for Relevance Feedback, in: SIGIR, 2020, p. 59–68.
- [16] F. L. Kreyssig, I. Nigo Casanueva, P. Budzianowski, M. G. Gašić, Neural User Simulation for Corpus-based Policy Optimisation for Spoken Dialogue Systems, Technical Report, 2018.
- [17] B. P. Zeigler, Theory of Modelling and Simulation | Guide books, Krieger Publishing Co., Inc, 1984. URL: <https://dl.acm.org/doi/10.5555/537562>.
- [18] N. J. Belkin, Cognitive models and information transfer, Social Science Information Studies 4 (1984) 111–129. doi:10.1016/0143-6236(84)90070-X.
- [19] C. C. Kuhlthau, Inside the search process: Information seeking from the user's perspective, JASIST 42 (1991) 361–371.
- [20] P. Ingwersen, K. Järvelin, The Turn: Integration of Information Seeking and Retrieval in Context (The Information Retrieval Series), Springer-Verlag New York, Inc., 2005.
- [21] D. G. Ellis, A Behavioral Approach to Information Retrieval System Design, Journal of Documentation 45 (1989) 171–212.
- [22] N. Craswell, O. Zoeter, M. Taylor, B. Ramsey, An experimental comparison of click position-bias models, in: WSDM, 2008, pp. 87–94.
- [23] O. Chapelle, Y. Zhang, A Dynamic Bayesian Network click model for web search ranking, in: WWW, ACM Press, 2009, pp. 1–10.
- [24] G. Dupret, B. Piwowarski, A user browsing model to predict search engine click data from past observations, in: ACM SIGIR, 2008, pp. 331–338.
- [25] A. Chuklin, P. Serdyukov, M. De Rijke, Modeling clicks beyond the first result page, in: International Conference on Information and Knowledge Management, Proceedings, ACM, 2013, pp. 1217–1220.
- [26] W. Eckert, E. Levin, R. Pieraccini, User modeling for spoken dialogue system evaluation, in: IEEE, IEEE, 1997, pp. 80–87.
- [27] K. Komatani, S. Ueno, T. Kawahara, H. G. Okuno, User modeling in spoken dialogue systems to generate flexible guidance, in: User Modelling and User-Adapted Interaction, volume 15, Springer, 2005, pp. 169–183.
- [28] O. Pietquin, A Framework for Unsupervised Learning of Dialogue Strategies, Ph.D. thesis, Faculté Polytechnique de Mons, TCTS Lab (Belgique), 2004. URL: <http://theses.ulb.ac.be:8002/ETD-db/collection/available/FPMSetd-11082004-163102/>.
- [29] J. Schatzmann, K. Weilhammer, M. Stuttle, S. Young, A Survey of Statistical User Simulation Techniques for Reinforcement-Learning of Dialogue Management Strategies, The Knowledge Engineering Review (2006) 0–1.
- [30] K. Scheffler, S. Young, Probabilistic simulation of human-machine dialogues, in: ICASSP, IEEE, volume 2, IEEE, 2000, pp. 1217–1220.
- [31] J. Schatzmann, B. Thomson, K. Weilhammer, H. Ye, S. Young, Agenda-Based User Simulation for Bootstrapping a POMDP Dialogue System, Technical Report, 2007.
- [32] Y. Hou, M. Fang, W. Che, T. Liu, A Corpus-free State2Seq User Simulator for Task-oriented Dialogue, Technical Report, 2019.
- [33] X. Li, Z. C. Lipton, B. Dhingra, L. Li, J. Gao, Y.-N. Chen, A User Simulator for Task-Completion Dialogues *, Technical Report, 2016. URL: <https://github.com/MiuLab/UserSimulator>.
- [34] S. Zhang, K. Balog, Evaluating Conversational Recommender Systems via User Simulation, in: ACM SIGKDD, Association for Computing Machinery, 2020, pp. 1512–1520.
- [35] L. E. Asri, J. He, K. Suleman, A sequence-to-sequence model for user simulation in spoken dialogue systems, in: ISCA, 2016.
- [36] S. Chandramohan, M. Geist, F. Lefèvre, O. Pietquin, User Simulation in Dialogue Systems using Inverse Reinforcement Learning, Technical Report, 2011.
- [37] X. Chen, S. Li, H. Li, S. Jiang, Y. Qi, L. Song, Generative adversarial user model for reinforcement learning based recommendation system, in: ICML,

- volume 97 of *PMLR*, PMLR, 2019, pp. 1052–1061.
- [38] X. Zhao, L. Xia, L. Zou, D. Yin, J. Tang, Simulating User Feedback for Reinforcement Learning Based Recommendations, Technical Report, 2019.
 - [39] S. Russell, Learning agents for uncertain environments (extended abstract), in: COLT’ 98, ACM Press, 1998, pp. 101–103.
 - [40] Z. Zhang, X. Li, J. Gao, E. Chen, Budgeted policy learning for task-oriented dialogue systems, in: ACL, 2019, pp. 3742–3751.
 - [41] O. Pietquin, H. Hastie, A survey on metrics for the evaluation of user simulations, *The Knowledge Engineering Review* (2008) 0–1.
 - [42] K. Papineni, S. Roukos, T. Ward, W.-J. Zhu, {B}leu: a Method for Automatic Evaluation of Machine Translation, in: ACL, 2002, pp. 311–318.
 - [43] C.-Y. Lin, ROUGE: A package for automatic evaluation of summaries, in: *Text Summarization Branches Out*, Association for Computational Linguistics, 2004, pp. 74–81.
 - [44] S. Banerjee, A. Lavie, METEOR: An automatic metric for MT evaluation with improved correlation with human judgments, in: ACL, Ann Arbor, Michigan, 2005, pp. 65–72.
 - [45] A. Sordoni, M. Galley, M. Auli, C. Brockett, Y. Ji, M. Mitchell, J.-Y. Nie, J. Gao, B. Dolan, A Neural Network Approach to Context-Sensitive Generation of Conversational Responses, in: NAACL, 2015, pp. 196–205.
 - [46] C.-W. Liu, R. Lowe, I. Serban, M. Noseworthy, L. Charlin, J. Pineau, How {NOT} To Evaluate Your Dialogue System: An Empirical Study of Unsupervised Evaluation Metrics for Dialogue Response Generation, in: EMNLP, 2016, pp. 2122–2132.
 - [47] W. Shi, K. Qian, X. Wang, Z. Yu, How to Build User Simulators to Train {RL}-based Dialog Systems, in: EMNLP-IJCNLP, 2019, pp. 1990–2000.
 - [48] A. Rastogi, X. Zang, S. Sunkara, R. Gupta, P. Khaitan, Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset, *AAAI* 34 (2020) 8689–8696.
 - [49] P. Budzianowski, T.-H. Wen, B.-H. Tseng, I. Casanueva, S. Ultes, O. Ramadan, M. Gašić, MultiWOZ - A Large-Scale Multi-Domain Wizard-of-Oz Dataset for Task-Oriented Dialogue Modelling, in: EMNLP, 2018, pp. 5016–5026.
 - [50] R. W. White, S. M. Drucker, Investigating behavioral variability in web search, in: WWW 2007, 2007, pp. 21–30.
 - [51] J. Li, M. Galley, C. Brockett, G. Spithourakis, J. Gao, B. Dolan, A Persona-Based Neural Conversation Model, in: ACL, 2016, pp. 994–1003.
 - [52] K. Balog, T. Kenter, Personal knowledge graphs: A research agenda, in: ICTIR, 2019, p. 217–220.
 - [53] E. Ie, C. wei Hsu, M. Mladenov, V. Jain, S. Narvekar, J. Wang, R. Wu, C. Boutilier, Recsim: A configurable simulation platform for recommender systems (2019).
 - [54] D. Rohde, S. Bonner, T. Dunlop, F. Vasile, A. Karatzoglou, Recogym: A reinforcement learning environment for the problem of product recommendation in online advertising, arXiv preprint arXiv:1808.00720 (2018).