



Joint Learning of Interactive Spoken Content Retrieval and Trainable User Simulator

Pei-Hung Chung^{1*}, Kuan Tung^{2*}, Ching-Lun Tai², Hung-Yi Lee¹

¹Graduate Institute of Communication Engineering, National Taiwan University

²Department of Electrical Engineering, National Taiwan University

{chung95191, dinotuku, vincent001217, tlkagkb93901106}@gmail.com

Abstract

User-machine interaction is crucial for information retrieval, especially for spoken content retrieval, because spoken content is difficult to browse, and speech recognition has a high degree of uncertainty. In interactive retrieval, the machine takes different actions to interact with the user to obtain better retrieval results; here it is critical to select the most efficient action. In previous work, deep Q-learning techniques were proposed to train an interactive retrieval system but rely on a hand-crafted user simulator; building a reliable user simulator is difficult. In this paper, we further improve the interactive spoken content retrieval framework by proposing a learnable user simulator which is jointly trained with interactive retrieval system, making the hand-crafted user simulator unnecessary. The experimental results show that the learned simulated users not only achieve larger rewards than the hand-crafted ones but act more like real users.

Index Terms: Interactive Retrieval, Human-Computer Interaction, Deep Reinforcement Learning

1. Introduction

Interactive spoken content retrieval (SCR) system enhances a retrieval system by incorporating user-system interaction [1, 2] into the spoken content retrieval process. The primary task of SCR is to retrieve the spoken content [3] or multimedia [4, 5] desired by the user. The necessity of user-system interaction in an SCR system is shown in many aspects. In SCR, the spoken content is usually transcribed into one-best transcriptions or lattices [6, 7, 8, 9, 10]. As ASR errors are inevitable, the retrieved results often contain ambiguities. Additionally, it is difficult to display the retrieved multimedia or spoken information items to users [11].

In previous work, the Markov decision process (MDP) has been used to model spoken-content interactive information retrieval (IIR) [12, 13, 14]. Wu et al. [15] propose using deep reinforcement learning (RL) in interactive SCR. Approaches to learning interaction policy include training a system to learn against a user simulator [16, 17, 18]. In spoken-content IIR, the agent learns through interaction; however, it is not feasible to learn this using real users in user-system interaction. Thus a user simulator is used to train spoken-content IIR. The construction of the user simulator is often the key to a satisfactory system. Policies learned with a poor user simulator can be worse than heuristic hand-crafted policies [19]. In previous approaches, simulated users choose actions based on hand-crafted rules without randomness [12, 13, 14, 15]. These hand-crafted rules are not sufficient to simulate the behavior of real users. Moreover, simulated users are

unrealistic in that they always generate the same responses given the same situations.

To create a reliable user simulator, we propose a new framework for interactive SCR. We design a trainable user simulator for use in joint learning with the dialogue manager, making hand-crafted simulated users unnecessary. Joint optimization of the dialog agent and the user simulator with deep RL using simulated dialog between the two agents has been proposed for task-oriented dialogue [20], and leads to promising improvements after pre-training by supervised learning. Learning two agents jointly has also been considered in object detection dialogue [21] and negotiation dialogue [22]. In this paper, we consider interactive SCR, which is quite different from previous work. Here both the user simulator and the IIR system are learned jointly by deep Q-network (DQN) [23] from scratch without any labeled data for pre-training. We use advanced DQN algorithms [24] to improve the performance of both the interactive SCR system and the user simulator. The learned simulator displays more sophisticated behavior and acts more like a real user.

2. Proposed Approach

Figure 1 illustrates the framework of the proposed interactive SCR, in which the interactive SCR system and the user simulator are jointly learned. In the proposed framework, the user simulator is given a set of queries and their corresponding relevant spoken documents. Although the system developers must provide data to the user simulator, they do not need to recruit real users to interact with the system, which greatly facilitates system development. During training, the user simulator enters a query (e.g., “US president”) into the SCR system. The system retrieves the search results, and to improve the results, it requests more information from the user (e.g., “Please provide more information.”). Then the user simulator responds (e.g., “Trump”). The response is determined based on the current search results and the relevant documents. Given the response, the SCR system updates the search results and poses another question. This interaction between the SCR system and the user simulator continues until the simulator decides to terminate (the user is satisfied with the results, or gives up).

During this interaction, the retrieval system does not know which documents are relevant. It must learn to ask the most efficient question given the situation to obtain the information with least interaction. For the user simulator, though, although it knows which documents are relevant, it cannot directly access the database. It can only update the search results by providing useful responses to the SCR system. The DQN parameters in the dialogue manager and the user simulator are jointly learned so the SCR finds the relevant documents in the most efficient way. During interaction, real users also do their best to help the SCR system. Thus we believe that a user simulator learned in this fash-

*The two first authors made equal contributions.

This work was supported by the Ministry of Science and Technology of Taiwan

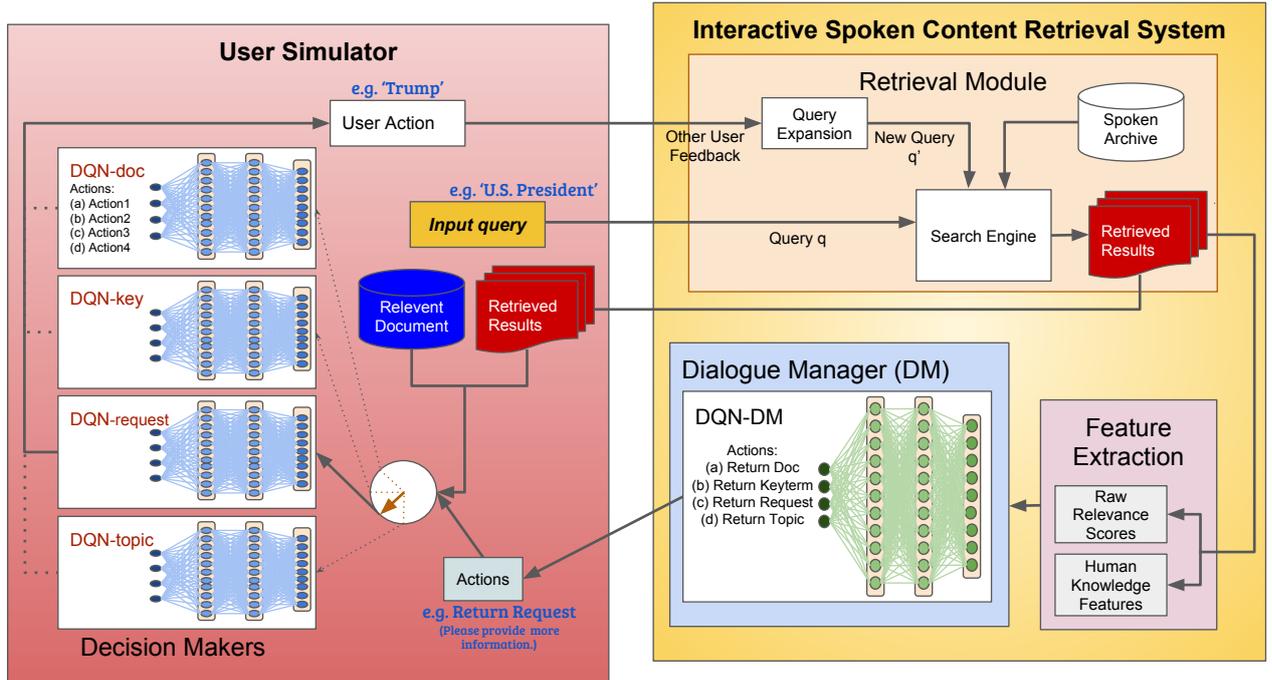


Figure 1: Block diagram of jointly learned user simulator and interactive spoken content retrieval system

ion is better than one based on hand-crafted rules.

The interactive SCR system is described in Section 2.1, and the user simulator is described in Section 2.2. The DQN is used to determine the actions of both the interactive SCR system and the user simulator. The relevant DQN algorithms are described in Section 2.3.

2.1. Interactive SCR System

The interactive SCR system comprises the retrieval module, the feature extraction module, and the dialogue manager [15]. The user first enters a query q into the system. With this user query, and potentially with other extra feedback information from the user during further user-system interaction [25, 26, 27], the retrieval module generates a list of retrieval results. The retrieved results returned from the retrieval module are fed into the feature extraction module, after which they are represented as a feature vector [15, 14]. The dialogue manager determines the action based on the extracted features. There are four possible actions: return documents¹, return key term², return request³, and return topic⁴. By taking an action and showing the retrieved results based on the current information, the user's response provides the system with more information, enabling the system to update the retrieved results. Given the updated results, the dialogue manager goes on to the next dialogue round until the user decides to end the interaction.

For the system reward, we use a tailored negative reward for each action because no matter what action is taken by the dialogue manager, all feedback actions represent an additional burden to

¹The system says "Please view the list and select one item relevant to your need".

²The system says "Is it related to t^* ?", where t^* is a key term.

³The system says "Please provide more information".

⁴The system shows a list of topics and says "Which topic is related?".

the user. To evaluate the retrieval result, we utilize the mean average precision (MAP) as an indicator: the better the MAP, the higher reward obtained by the system. The return (total reward) R of a dialogue is defined as

$$R = \sum_{k=1}^T c_k + \lambda r_k, \quad (1)$$

where c_k is the negative reward from the action taken at the k -th turn, and r_k is the MAP improvement after the interaction. λ is a constant set by the system developer.

2.2. User Simulator

Given the action of the SCR system, the user simulator offers a response. In the user simulator, for each action the corresponding decision maker decides on the suitable response. The input of the decision maker is a K -dimensional feature vector. The feature vector is the relevance of the top- K documents in the retrieved list from the SCR system, which represents the retrieval quality of the retrieval results at the current stage. If the document ranked at the k -th position is relevant, the k -th dimension is 1; otherwise, the value is zero (the relevance of the returned documents is known only by the user simulator, not the SCR system).

Because the SCR system has four possible actions, there are four decision makers in the user simulator. The four decision makers are described as follows:

- **Return Documents:** Given a retrieved list, we rank the relevant documents based on their relevance scores from the retrieval system. The response is the relevant documents ranked at a specific position. The decision maker determines the reply by document rankings.
- **Return Key Term:** If the key term appears in more than 50% of the relevant documents, the simulated user replies

“YES” with probabilities of 100%, 95%, 90%, or 85%. The probability is decided by the decision maker.

- *Return Request*: Each term t has a score $S(t) = \sum_{d \in \mathcal{R}_{\mathcal{T}}} N(t, d) \ln(1 + idf(t))$, where $N(t, d)$ is the term frequency of term t in the manual transcription of document d , $idf(t)$ the inverse document frequency of term t computed from the manual transcriptions of the document collection, and $\mathcal{R}_{\mathcal{T}}$ the real relevant document set provided by the annotators. All the terms are ranked according to $S(t)$. Again, the decision maker determines the reply by term rankings.
- *Return Topic*: Given a query, a set of topics is ranked according to their relevance to the query (the relevance is provided by the annotators). The decision maker determines the reply by topic rankings.

If the MAP of the retrieved result is higher than a threshold (thus meeting the user’s information need) or the interaction turns exceeds a maximum number (the user gives up), the interaction stops. If the interaction ends because the MAP is higher than a threshold, the dialogue is deemed a success and the user simulator receives a positive reward; otherwise, the dialogue fails, and a negative reward is given as punishment⁵. In this study the threshold is set to 0.6, and the maximum number of turns is 4.

2.3. Reinforcement Learning

The dialogue manager in the interactive SCR system as well as all the decision makers in the user simulator are deep Q-networks (DQNs) [23]. The input of each DQN is a feature representation s , known also as a *state* per reinforcement learning. The definitions of this feature representation differ between the dialogue manager and decision maker. The DQN output dimension equals the number of possible actions a in the action set. The output corresponding to each action is the state-action value $Q(s, a)$. As the possible actions for the DQN in the dialogue manager are *Return Documents*, *Return Key Term*, *Return Request*, and *Return Topic*, the output dimension for the DQN in dialogue manager is four. For the user simulator decision makers, the DQN for the Return Key Term decision maker has four outputs (corresponding to 100%, 95%, 90%, and 85%). For the remaining three decision makers, each DQN output dimension corresponds to a ranking position (1st, 2nd, 3rd, etc). All the DQNs (one in the SCR system and four in the user simulator) are trained to maximize the expected return of the system they belong to via the interaction between the interactive SCR system and the user simulator. To ameliorate the risk of instability when jointly learning the user simulator and retrieval system [20], we update the DQNs in the two systems iteratively. That is, we first fix the DQNs in the user simulator, and update the DQN in the SCR system C times. Then we switch the training agent: We update the DQNs in the user simulator another C times, while fixing the SCR system. The above procedure is conducted iteratively.

In addition to the standard DQN algorithm, we use the extended versions. As the standard DQN algorithm has been observed to overestimate the Q-values [28, 29], double DQN is also used in the following experiment, which decouples the selection from the evaluation [28, 29]. Dueling DQN [30] is further applied here. Compared to typical DQN, dueling DQN uses an altered network structure, splitting it into two streams: one learns to provide an estimate of the state value $V(s)$ for every state, and the

⁵In contrast to the dialogue manager, the user simulator does not receive a negative reward for each action.

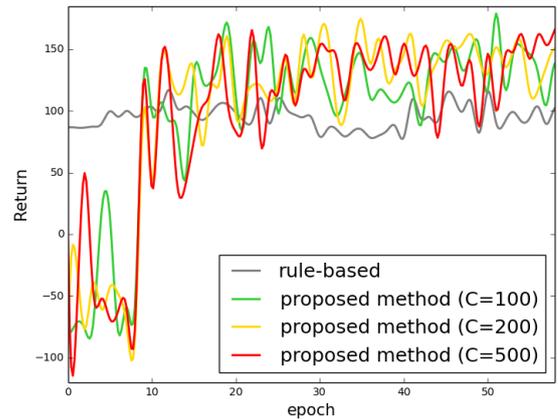


Figure 2: Learning curve of different methods

other calculates the potential advantages of each action at a given state.

3. Experiments

3.1. Data Set

In the experiments, we used a Mandarin Chinese broadcast news corpus as the spoken document collection from which to retrieve news. The news stories were recorded from radio or TV stations in Taipei from 2001 to 2003, and comprised a total of 5047 news documents covering a total length of 198 hours. For speech recognition, both one-best transcriptions and lattices were generated for spoken content retrieval. Twenty-two graduate students were recruited as the annotators, providing 163 text queries and their relevant spoken documents (not necessarily including the query terms).

3.2. Experiment Settings

Our retrieval module is based on language modeling [31, 32]. After receiving the user feedback in the dialogue, we use the query-regularized mixture model [25, 26, 27] to generate the new query q' . The input dimension K of the decision makers was set to 49. The output dimensions for the four DQNs in the user simulator were all set to four.

All the DQNs had the same hyper-parameters. We used networks with two hidden layers of 1024 nodes. We used relu as the activation function and set the batch size to 256. The initial learning rate was set to $8e-4$. MAP was selected as our retrieval evaluation metric. Ten-fold cross validation was performed in all experiments; that is, for each trial, 8 out of 10 query folds were used for training, 1 for parameter tuning (validation set), and the remaining fold for testing.

3.3. Results and Discussion

In this subsection, the user simulator is used to both train and test the SCR system. Figure 2 shows the learning curves of different methods on the training set. The rule-based system is identical to previous work [15]. The proposed approach is based on a typical DQN with the C iterative steps from Section 2.3. We find that training with the rule-based system is more stable than the proposed approach, and results in better performance in the first

Table 1: MAP and Return for different approaches evaluated on both one-best transcription and lattices

Approaches			One-best		Lattices	
User simulator	Dialogue manager	Input feature	MAP	Return	MAP	Return
(a) Rule-based	(a-1) DQN	Raw	0.5641	99.31	0.5847	112.24
		Human+Raw	0.5655	107.02	0.5790	101.99
	(a-2) Double DQN	Raw	0.5744	107.08	0.5911	112.14
		Human+Raw	0.5846	115.70	0.5959	102.52
	(a-3) Dueling DQN	Raw	0.5562	96.27	0.5860	113.70
		Human+Raw	0.5665	105.35	0.5784	96.37
(b-1) DQN	DQN	Raw	0.5758	113.02	0.6041	162.85
		Human+Raw	0.5820	130.83	0.6027	162.93
(b-2) Double DQN	Double DQN	Raw	0.6063	190.39	0.6414	224.42
		Human+Raw	0.6066	199.17	0.6375	222.66
(b-3) Dueling DQN	Dueling DQN	Raw	0.5780	108.72	0.5864	157.92
		Human+Raw	0.5757	127.49	0.5723	158.23
(c-1) Double DQN	Dueling DQN	Raw	0.5733	134.20	0.5678	143.83
		Human+Raw	0.5573	111.84	0.5905	146.83
(c-2) Dueling DQN	Double DQN	Raw	0.5899	164.00	0.6494	238.91
		Human+Raw	0.6139	208.25	0.6420	225.72

few epochs. This is reasonable because it is more difficult to train two systems together. However, with more epochs, the proposed approach eventually outperforms the previous rule-based method. C was set to 500 in the following experiments.

Table 1 shows the MAP and Return results for the SCR system based on one-best transcription and lattices. For each approach, the dialogue manager DQN utilizes two kinds of input features, which were also used in previous work [15]. *Raw* represents the relevance scores of the top-N documents in the retrieved results, and *Human* indicates the use of hand-crafted features based on human knowledge such as clarity score [33], ambiguity score [34], and weighted information gain (WIG) [35]. The table is divided into three blocks. In part (a), the rule-based user simulator was used, which is identical to Wu et al. [15]. The results of the learnable user simulator are in parts (b) and (c). In part (b), we used the same DQN algorithm for the user simulator and dialogue manager, while in part (c), we used different algorithms.

In part (a), we compare the different DQN algorithms with the rule-based user. The results show that in some cases, both double and dueling DQN outperform the typical DQN. Using the same DQN algorithm, the trainable user simulator always outperforms the rule-based user simulator (rows (b-1) v.s. (a-1), rows (b-2) v.s. (a-2), rows (b-3) v.s. (a-3)). For example, compare (b-1) with the baseline (a-1): the trainable user simulator yields improvements of 35.2% in terms of Return. We further achieve our best performance on (c-2), by utilizing dueling DQN for the user simulator and double DQN for the dialogue manager. Why this combination yields the best performance is still under investigation. It shows significant improvement, outperforming the rule-based baseline ((c-2) v.s. (a-1)) by 88.4% and joint learning using the same DQN algorithm ((b-2) v.s. (a-1)) by 99.0%. However, it is not enough to simply show that the trainable user simulator results in higher returns for the SCR system: it is possible that the user simulator is learning a very unnatural way to provide information to the SCR system. We address this in the next subsection.

3.4. Human evaluation

To verify the correlation and similarity between the proposed user simulator and real users, we performed a subjective human evaluation with 21 subjects, all of whom were graduate students. Given

Table 2: KL-divergence between rule-based simulator, learnable simulator and human results, as well as each system’s entropy

KL-divergence	Rule/human	Ours/human	Rule/ours
	5.9899	2.1503	6.1859
Entropy	Rule	Ours	Human
	0	0.6006	0.8886

a sampled query, the subject was asked to identify the most relevant document in the top-4 retrieved relevant documents⁶. Due to space limitations, we show only the results of *Return Document*; the results for the other actions are similar. The choices of human, rule-based user, and decision maker are considered random variables (with 1st, 2nd, 3rd, and 4th as the outcomes). We calculated the KL-divergence between the rule-based user simulator, our proposed trainable user simulator, and the human evaluation result. In Table 2, which lists the average results, the abbreviations are Rule, Ours, and Human respectively. The proposed user simulator and human have the smallest KL divergence (Ours/Human). This shows that the proposed joint training of the user simulator was more closely correlated to the distribution of human evaluations than the rule-based approach. We further calculated the entropy of each distribution. The entropy of the rule-based system was zero because given the same input, it always produces the same response. The result shows that the decisions made by the proposed user simulator were more diverse and also closer to human evaluation than the rule-based simulator.

4. Conclusions

User-system interaction is highly desired for SCR. Interactive SCR is usually learned using simulated users, but designing simulated users using hand-crafted rules is difficult. In this paper, we propose a new framework in which the user simulator and the interactive SCR system are jointly learned. The experimental results show that the proposed method leads to promising improvements on return. Human evaluation further shows that a trainable user simulator much more closely reflects human behavior.

⁶Because the output dimension of the decision maker for *Return Document* is 4, it returns only relevant documents ranked in the top four.

5. References

- [1] D. Robins, "Interactive information retrieval: Context and basic notions," *Informing Sci. J.*, vol. 3, pp. 57–62, 2000.
- [2] I. Ruthven, "Interactive information retrieval," *Annual Review of Information Science and Technology*, vol. 42, no. 1, pp. 43–91, 2008.
- [3] J. S. Garofolo, C. G. Auzanne, and E. M. Voorhees, "The TREC spoken document retrieval track: A success story," in *Content-Based Multimedia Information Access-Volume 1*. LE CENTRE DE HAUTES ETUDES INTERNATIONALES D'INFORMATIQUE DOCUMENTAIRE, 2000, pp. 1–20.
- [4] D. L. Patton, P. R. Ashe, and J. A. Manico, "Interactive image storage, indexing and retrieval system," Jun. 18 2002, US Patent 6,408,301.
- [5] S.-B. Cho, "Emotional image and musical information retrieval with interactive genetic algorithm," *Proceedings of the IEEE*, vol. 92, no. 4, pp. 702–711, 2004.
- [6] M. Saraclar and R. Sproat, "Lattice-based search for spoken utterance retrieval," in *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, 2004.
- [7] D. Hakkani-Tür, F. Béchet, G. Riccardi, and G. Tur, "Beyond ASR 1-best: Using word confusion networks in spoken language understanding," *Computer Speech & Language*, vol. 20, no. 4, pp. 495–514, 2006.
- [8] L.-S. Lee, J. Glass, H.-Y. Lee, and C.-A. Chan, "Spoken content retrieval; beyond cascading speech recognition with text retrieval," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 9, pp. 1389–1420, Sept 2015.
- [9] D. Can and M. Saraclar, "Lattice indexing for spoken term detection," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 8, pp. 2338–2347, 2011.
- [10] S. Parlak and M. Saraclar, "Spoken term detection for Turkish broadcast news," in *IEEE International Conference on Acoustics, Speech and Signal Processing, 2008*. IEEE, 2008, pp. 5244–5247.
- [11] Y. Zhang and C. Zhai, "Information retrieval as card playing: A formal model for optimizing interactive retrieval interface," in *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR '15. New York, NY, USA: ACM, 2015, pp. 685–694. [Online]. Available: <http://doi.acm.org/10.1145/2766462.2767761>
- [12] Y.-C. Pan, H.-Y. Lee, and L.-S. Lee, "Interactive spoken document retrieval with suggested key terms ranked by a Markov decision process," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 2, pp. 632–645, Feb 2012.
- [13] T.-H. Wen, H.-Y. Lee, and L.-S. Lee, "Interactive spoken content retrieval with different types of actions optimized by a Markov decision process," in *Thirteenth Annual Conference of the International Speech Communication Association*, 2012.
- [14] T.-H. Wen, H.-Y. Lee, P.-h. Su, and L.-S. Lee, "Interactive spoken content retrieval by extended query model and continuous state space Markov decision process," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, 2013, pp. 8510–8514.
- [15] Y.-C. Wu, T.-H. Lin, Y.-D. Chen, H.-Y. Lee, and L.-S. Lee, "Interactive spoken content retrieval by deep reinforcement learning," *arXiv preprint arXiv:1609.05234*, 2016.
- [16] J. Schatzmann, K. Weilhammer, M. Stuttle, and S. Young, "A survey of statistical user simulation techniques for reinforcement-learning of dialogue management strategies," *The Knowledge Engineering Review*, vol. 21, no. 2, pp. 97–126, 2006.
- [17] L. E. Asri, J. He, and K. Suleman, "A sequence-to-sequence model for user simulation in spoken dialogue systems," *arXiv preprint arXiv:1607.00070*, 2016.
- [18] X. Li, Z. C. Lipton, B. Dhingra, L. Li, J. Gao, and Y.-N. Chen, "A user simulator for task-completion dialogues," *arXiv preprint arXiv:1612.05688*, 2016.
- [19] B. Thomson and S. Young, "Bayesian update of dialogue state: A POMDP framework for spoken dialogue systems," *Computer Speech & Language*, vol. 24, no. 4, pp. 562–588, 2010.
- [20] B. Liu and I. Lane, "Iterative policy learning in end-to-end trainable task-oriented neural dialog models," *arXiv preprint arXiv:1709.06136*, 2017.
- [21] H. De Vries, F. Strub, S. Chandar, O. Pietquin, H. Larochelle, and A. Courville, "Guesswhat?! Visual object discovery through multi-modal dialogue," in *Proc. of CVPR*, 2017.
- [22] M. Lewis, D. Yarats, Y. N. Dauphin, D. Parikh, and D. Batra, "Deal or no deal? End-to-end learning for negotiation dialogues," *arXiv preprint arXiv:1706.05125*, 2017.
- [23] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing Atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.
- [24] M. Hessel, J. Modayil, H. Van Hasselt, T. Schaul, G. Ostrovski, W. Dabney, D. Horgan, B. Piot, M. Azar, and D. Silver, "Rainbow: Combining improvements in deep reinforcement learning," *arXiv preprint arXiv:1710.02298*, 2017.
- [25] H.-Y. Lee, T.-H. Wen, and L.-S. Lee, "Improved semantic retrieval of spoken content by language models enhanced with acoustic similarity graph," in *Spoken Language Technology Workshop (SLT), 2012 IEEE*. IEEE, 2012, pp. 182–187.
- [26] C. Zhai and J. Lafferty, "Model-based feedback in the language modeling approach to information retrieval," in *Proceedings of the Tenth International Conference on Information and Knowledge Management*. ACM, 2001, pp. 403–410.
- [27] T. Tao and C. Zhai, "Regularized estimation of mixture models for robust pseudo-relevance feedback," in *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2006, pp. 162–169.
- [28] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double Q-learning," in *AAAI*, vol. 16, 2016, pp. 2094–2100.
- [29] H. V. Hasselt, "Double Q-learning," in *Advances in Neural Information Processing Systems*, 2010, pp. 2613–2621.
- [30] Z. Wang, T. Schaul, M. Hessel, H. Van Hasselt, M. Lanctot, and N. De Freitas, "Dueling network architectures for deep reinforcement learning," *arXiv preprint arXiv:1511.06581*, 2015.
- [31] J. Lafferty and C. Zhai, "Document language models, query models, and risk minimization for information retrieval," in *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2001, pp. 111–119.
- [32] T. K. Chia, K. C. Sim, H. Li, and H. T. Ng, "Statistical lattice-based spoken document retrieval," *ACM Transactions on Information Systems (TOIS)*, vol. 28, no. 1, p. 2, 2010.
- [33] B. He and I. Ounis, "Query performance prediction," *Information Systems*, vol. 31, no. 7, pp. 585–594, 2006.
- [34] S. Cronen-Townsend, Y. Zhou, and W. B. Croft, "Predicting query performance," in *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2002, pp. 299–306.
- [35] Y. Zhou and W. B. Croft, "Query performance prediction in web search environments," in *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2007, pp. 543–550.