

Balancing Reinforcement Learning Training Experiences in Interactive Information Retrieval

Limin Chen, Zhiwen Tang, Grace Hui Yang

{lc1117,zt79,grace.yang}@georgetown.edu

InfoSense, Department of Computer Science, Georgetown University

ABSTRACT

Interactive Information Retrieval (IIR) and Reinforcement Learning (RL) share many commonalities, including an agent who learns while interacts, a long-term and complex goal, and an algorithm that explores and adapts. To successfully apply RL methods to IIR, one challenge is to obtain sufficient relevance labels to train the RL agents, which are infamously known as sample inefficient. However, in a text corpus annotated for a given query, it is not the relevant documents but the irrelevant documents that predominate. This would cause very unbalanced training experiences for the agent and prevent it from learning any policy that is effective. Our paper addresses this issue by using domain randomization to synthesize more relevant documents for the training. Our experimental results on the Text REtrieval Conference (TREC) Dynamic Domain (DD) 2017 Track show that the proposed method is able to boost an RL agent's learning effectiveness by 22% in dealing with unseen situations.

CCS CONCEPTS

• Information systems → Users and interactive retrieval.

KEYWORDS

Dynamic Search; Interactive IR; Deep Reinforcement Learning

ACM Reference Format:

Limin Chen, Zhiwen Tang, Grace Hui Yang. 2020. Balancing Reinforcement Learning Training Experiences in Interactive Information Retrieval. In *43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*, July 25–30, 2020, Virtual Event, China. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3397271.3401200>

1 INTRODUCTION

Reinforcement Learning (RL) fits Interactive Information Retrieval (IIR) well as both of them center on accomplishing a goal during an interactive process. In RL, a machine agent maximizes its cumulative rewards collected during the course of interactions with an environment. In IIR, a search system satisfies an information need during the course of interactions with a user and a corpus. These commonalities have inspired approaches for IIR using RL solutions

[14, 20]. In these solutions, the search system is the RL agent; and both the user and text corpus the RL environment.

Training an RL agent could be difficult. First, it is expensive to train the agent. To learn some policy useful, the agent may need to take millions of steps to interact with the environment. It becomes even more expensive when real humans are involved in the interactions, as what we have here in IIR. Simulations have therefore been proposed to replace real human users to train interactive systems [19]. Second, these simulators are usually created based on pre-annotated ground truth corpora. The corpora usually contain many irrelevant documents and a few relevant documents. Therefore the simulators would not be able to form a balanced training environment.

The unbalance in the forming of training environments would prevent the RL agent from learning good policies. It is because the rewards – relevant documents – are too few and the RL agent may not be able connect a long sequence of retrieval actions to a distant future reward, thus will never learn how to perform a task. This is known as the problem of “sparse rewards”. Trained with sparse rewards, when the RL agent is deployed in a real dynamic environment, such as in the Web, it would be likely to make wrong decisions, especially when the agent encounters documents never seen before.

In this paper, we propose a novel domain randomization method to enhance the training experiences of RL agents in interactive retrieval. Our method, *Document Environment Generation (DEG)*, proposes to automatically generate positive learning environments as many as needed during a simulated training. DEG derives a stream of synthetic environments from available relevant documents by merging relevant segments within the documents with other irrelevant segments extracted from the corpus. In addition, our method dynamically restrains the changing rate of a policy to make sure that the RL agent adapts conservatively to the synthetic environments. We experiment our method on the Text REtrieval Conference (TREC) 2017 Dynamic Domain (DD) Track and the results show that the proposed method is able to statistically significantly boost the agents' learning effectiveness.

2 RELATED WORK

Reinforcement learning has been successfully applied in several applications in text domains. They include dialogue systems [5], recommendation systems [7] and dynamic search [14]. These systems aim to find items that match with user's interests by modeling the interactions between a system and a human user. The required participation of real human users makes it quite costly to sample training trajectories for the system. Researcher have proposed to build simulators to reduce the sample complexity [19]. However, the difference between a simulation and a real environment, known as

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGIR '20, July 25–30, 2020, Virtual Event, China

© 2020 Association for Computing Machinery.
ACM ISBN 978-1-4503-8016-4/20/07...\$15.00
<https://doi.org/10.1145/3397271.3401200>

their reality gap, makes it challenging to directly deploy an RL agent into the real world [2]. Domain randomization has been proposed to alleviate this problem in robotics. For instance, [17] randomized the objects and textures when training a robotic arm to grasp in cluttered environments. [8] randomized the mass and damping of robotic arms. [10] randomized the floor plan when training a drone to fly indoor.

A similar technique, data augmentation, has been used in supervised learning to add more training data. In computer vision, it is proposed to augment image data by performing geometry transformation, such as flipping and cropping [13], and photometric transformation, such as edge enhancement and color jittering [16]. In natural language processing, it is proposed to generate new text by various language generation methods. For instances, [6] created noisy texts by substituting words with its synonyms. [3] flipped the order of characters to improve the robustness of a text classifier. [18] also augmented training samples by random insertion and deletion.

In this work, we randomize a process to generate new texts by leveraging a unique characteristic in IR. The characteristic is that whether a document is relevant largely depends on if it contains the query keywords and only the matching parts in the document would determine the relevance of the document.

Training an RL agent in a series of different environments sometimes would result in ‘catastrophic forgetting’ [4]. We handle this problem in our work too because our agent needs to learn from an enlarged and more diverse set of environments. [9] dealt with this problem by training an ensemble of neural networks, each for a separate task, and sharing weights among them. [4] augmented a loss function to slow down the learning on important parameters. [1] transferred policies learned among different tasks with network distillation. Unlike them, DEG handles this problem by restraining the policy’s changing rate based on how much the newly generated environments differ from the original.

3 SETUP

In IIR, a human user searches for relevant information within a text collection from an interactive search engine. The Text REtrieval Conference (TREC) Dynamic Domain (DD) Track [19] provides a platform to evaluate interactive search engines. In the Track, a simulated user starts an initial query. At each subsequent time step, a search system retrieves a set of documents and returns the top K ($K = 5$ in TREC DD) to the simulated user. The simulated user provides explicit feedback regarding how relevant these documents are and then the search system adjusts its search algorithm based on the feedback. The process repeats until the search stops.

CE3 [15] is a state-of-the-art IIR system. It is based on the proximal policy optimization (PPO) algorithm [12]. In this paper, we improve this state-of-the-art system by incorporating domain randomization. Following CE3’s practice, DEG also uses a corpus-level state representation. To build the states, our method first segments each document into a fixed number (M) of segments. Each segment is on the same topic and maps into a separate doc2vec vector after compression. State at time t is formed like taking a snapshot of the entire corpus by stacking together the representations of all documents at t . This global state representation at t is expressed in

the embedding function \mathcal{S} :

$$s_t = \mathcal{S}(C, \mathcal{D}_1 \cup \mathcal{D}_2 \dots \cup \mathcal{D}_{t-1}) \quad (1)$$

where C is the text corpus and \mathcal{D}_i is the set of documents retrieved at time i . The state representation’s dimension is $|C| \times M \times n$, where $|C|$ is the size of the corpus, M the number of segments per document, and n the lower dimension ($n \ll$ vocabulary size) of the doc2vec vector after compression. The retrieved documents at time t , i.e. \mathcal{D}_t , will be marked as ‘visited’ on this global representation and passed on to next run of retrieval.

Action at time t , a_t , is a weighting vector for document ranking. The ranking function calculates a relevance score between document d_i and search topic (query) q as the weighted sum over all segments $x_{ij} \in d_i$:

$$score_{i,t} = f(a_t, d_i) = \sum_{j=1}^M x_{ij} \cdot a_t \quad (2)$$

where x_{ij} is the doc2vec representation of the j^{th} segment in the i^{th} document and a_t is the action vector (weighting vector) at time t .

Reward r_t is derived from the relevance ratings annotated by human assessors. These relevance ratings are from 0 (irrelevant) to 4 (highly relevant). We define r_t as the sum of all relevance ratings for the returned documents \mathcal{D}_t after duplicated results are removed:

$$r_t = \sum_{d_i \in \mathcal{D}_t \setminus (\mathcal{D}_1 \cup \mathcal{D}_2 \dots \cup \mathcal{D}_{t-1})} rel(d_i) \quad (3)$$

where $rel()$ gets the relevance rating for d_i by looking up and summing the ratings for all passages in d_i in the ground truth.

4 PROPOSED METHOD

This section presents our method on generating more relevant documents to form more balanced training environments and using an adaptive clipping rate to constrain the policy from dramatic changes.

4.1 Generate Synthetic Environments

Information retrieval is a task driven by keyword matching. Users recognize relevant documents by recognizing query keywords that are present in the documents. As long as the keywords are kept in a document, the document is considered as relevant. We therefore propose to separate relevant segments from irrelevant segments in a relevant document and put them into different uses.

Our process to create new relevant documents is the following. First, DEG separates the corpus into three parts, relevant segments from relevant documents, irrelevant segments from relevant documents, and irrelevant segments from irrelevant documents. Note all segments in an irrelevant document are irrelevant. Second, an ‘irrelevant pool’ of segments, P , is formed by putting together all irrelevant segments from both relevant and irrelevant documents. Third, for each relevant document, DEG samples N segments, p_1, p_2, \dots, p_N , from the irrelevant pool P :

$$N \sim \mathcal{U}(0, F) \quad (4)$$

where \mathcal{U} is a uniform distribution and F is a parameter to control the maximum ‘fake’ level that the synthesized documents would

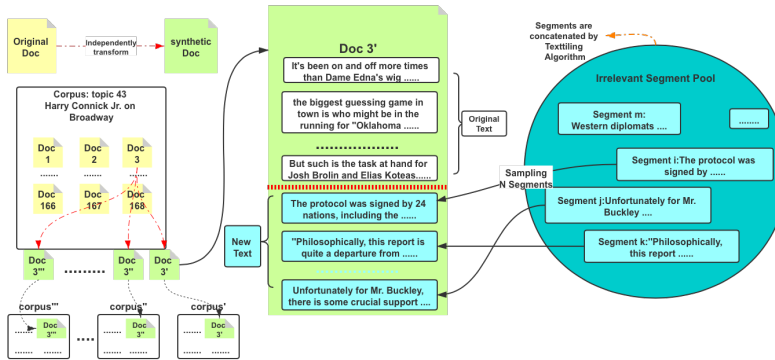


Figure 1: Generating new relevant documents $Doc3'$, $Doc3''$ and $Doc3'''$ independently from $Doc3$ (Topic DD17-43).

be. Fourth, irrelevant segments p_1, p_2, \dots, p_N replace the original irrelevant segments in the relevant document and form a synthetic relevant document. Each sampling can independently generate a new document from the same seed document. Fifth, documents synthesized from the same document share the same relevance rating as the seed document as they all contain the same set of relevant segments. Last, the process repeats for every relevant document in the corpus. DEG can generate as many new relevant documents as possible from the same original document. Figure 1 illustrates DEG’s document generation process.

These generated documents are added into the corpus to form a new training environment. With more relevant documents now in the training environment, the RL agent is able to meet a reward (relevant document) more frequently when it interacts with the environment. It is then able to learn a policy with enough reward signals. Moreover, the agent can be exposed to many different training environments when we include different synthesized documents to form a new environment. The agent is thus trained with not only more data but more diverse data, which would help with its generalizing ability.

Note that the generation process needs knowledge about passage-level ground truth relevance, which may not be available in datasets outside of TREC DD. However, such knowledge may also be obtained from matching keywords with ground truth documents. The relevance can also be derived from implicit feedback such as clicks and dwell time. We think the method proposed here is applicable to other settings for interactive retrieval. The test environments used in this work do not use DEG because it would be unfair to know ground truth in a test environment.

4.2 Adaptive Training

Training an RL agent on vastly different environments may result in a problem known as “catastrophic forgetting” of the original environment [4]. In our case, the synthetic environment could be very different from the original and causes this problem. We propose to control the policy derivations within a bound based on how different the synthesized environment is from the original. We call this strategy “adaptive training”.

The RL framework we use is based on PPO [12, 15], which optimizes the following objective: $J(\theta) = \hat{E}_t[\min(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1 -$

$\epsilon, 1 + \epsilon)\hat{A}_t)]$, where $r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{old}(a_t|s_t)}$ is a ratio denoting the change of action distribution. The $\text{clip}()$ function limits the change ratio within $[1 - \epsilon, 1 + \epsilon]$, where ϵ is the maximum change of the action distributions and usually set to a fixed value, such as 0.2. This formulation has already been able to prevent drastic policy change.

Our work goes one step further to use a dynamically defined ϵ , instead of a fixed value. In DEG, the synthetic environment differs from the original environment in terms of the documents it consists of. Thus, we calculate ϵ based on how much the synthetic documents differ from the original document. Their difference is measured by Euclidean distance in the embedding space. For the m^{th} synthetic environment, ϵ_m is calculated as

$$\epsilon_m = \alpha u^{-\|S(C) - S(C'_m)\|} \quad (5)$$

where C is the original environment, C' is the synthetic environment, S is the embedding function shown in Eq. 1, $\|\cdot\|$ is the Euclidean distance function, and α and u are hyper parameters.

For environments that differ more from the original environment, ϵ_m puts a tighter restraint on the change rate, which slows down the change of the policy. In this way, DEG prevents the agent from catastrophic forgetting when trained with more diverse environments.

5 EXPERIMENTS

We experiment on the TREC 2017 Dynamic Domain Track [19]. It uses the New York Times corpus [11]. Sixty search topics were created for the Track and each is a distinct learning environment. We compare the performance of the following IIR methods: CE3 [15], a PPO-based RL method without domain randomization. DEG, the proposed method with domain randomization and adaptive training; $\alpha = 0.2$, $u = 1.006$, and $F = 5$. DEG (fixed), a variant of DEG with a fixed ϵ ; ϵ is set to 0.2 as it works the best for many topics.

5.1 Effectiveness under Unseen Situations

To understand how well a trained agent can perform interactive retrieval under unseen situations, we split the TREC DD 2017 dataset into non-overlapping training and test sets. The test dataset is considered as unseen environments to the agents. For each search

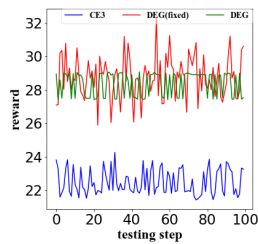


Figure 2: Average Reward




Figure 3: Case study: Topics 49 and 21

topic, we train RL agents using CE3, DEG or DEG (fixed) on the training set and test them on the test set. We compare the agents’ performance by measuring the rewards they obtain in the first 100 runs of interactions in the test set. Figure 2 reports the reward of CE3, DEG, and DEG (fixed), averaged over test search topics.

We observe that both DEG runs achieve a significant 5% absolute and 22% relative ($p < 0.0001$, double-tailed t-test) improvement over CE3. It shows that agents trained with domain randomization work much better than trained without it. We think it is because DEG generates a lot more relevant documents to form more balanced training environments and makes the agents generalize better. In addition, we found that although on average DEG and DEG (fixed) work comparable, DEG demonstrates much lower variance in the rewards that it can obtain. The proposed use of adaptive training contributes to a more conservative policy update, which results in more stable rewards. This is because DEG is less likely to take a big step to fall into a trajectory that is either too good or too bad.

5.2 Case Studies

Figure 3 shows two representative examples, topics 49 and 21. In Topic 49, CE3 gets nearly zero rewards in unseen environments, which indeed works very poor. On the other hand, DEG variants much better adjust to unseen environments and can both get impressive rewards. Domain randomization seems to help the agents develop much better generalization ability. In Topic 21, we can see that DEG (fixed)’s rewards fluctuate a lot as the interactions go on while DEG produces much more stable rewards. Although DEG may miss some high-value steps, it too avoids the wrong moves, which is desirable for a task that cares much about precision at top positions.

6 CONCLUSION

This paper proposes a domain randomization method to enhance reinforcement learning in interactive retrieval. Our method generates

new relevant documents from existing documents in a collection to increase frequency of meeting relevant documents during RL training. The experiments show that this strategy can significantly boost an RL agent’s performance in interactive retrieval under unseen situations. The proposed adaptive training method also helps the RL agent explore more steadily with stable rewards.

ACKNOWLEDGMENTS

This research was supported by U.S. National Science Foundation Grant no. IIS-145374. Any opinions, findings, conclusions, or recommendations expressed in this paper are of the authors, and do not necessarily reflect those of the sponsor.

REFERENCES

- [1] Glen Berseth, Cheng Xie, Paul Cernek, and Michiel Van de Panne. 2018. Progressive reinforcement learning with distillation for multi-skilled motion control. In *ICLR '18*.
- [2] Yevgen Chebotar, Ankur Handa, Viktor Makoviychuk, Miles Macklin, Jan Issac, Nathan Ratliff, and Dieter Fox. 2019. Closing the sim-to-real loop: Adapting simulation randomization with real world experience. In *ICRA '19*.
- [3] Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. HotFlip: White-Box Adversarial Examples for Text Classification. In *ACL '18*.
- [4] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* (2017).
- [5] Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016. Deep Reinforcement Learning for Dialogue Generation. In *EMNLP '16*.
- [6] Yitong Li, Trevor Cohn, and Timothy Baldwin. 2017. Robust training under linguistic adversity. In *EACL '17*.
- [7] Feng Liu, Ruiming Tang, Xutao Li, Weinan Zhang, Yunming Ye, Haokun Chen, Hui Feng Guo, and Yuzhou Zhang. 2018. Deep reinforcement learning based recommendation with explicit user-item interactions modeling. *arXiv preprint arXiv:1810.12027* (2018).
- [8] Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. 2018. Sim-to-real transfer of robotic control with dynamics randomization. In *ICRA '18*.
- [9] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. 2016. Progressive neural networks. *arXiv preprint arXiv:1606.04671* (2016).
- [10] Fereshteh Sadeghi and Sergey Levine. 2016. Cad2rl: Real single-image flight without a single real image. *arXiv preprint arXiv:1611.04201* (2016).
- [11] Evan Sandhaus. 2008. The new york times annotated corpus. *Linguistic Data Consortium, Philadelphia* 6, 12 (2008), e26752.
- [12] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *arXiv e-prints* (2017). arXiv:1707.06347 [cs.LG]
- [13] Connor Shorten and Taghi M Khoshgoufar. 2019. A survey on image data augmentation for deep learning. *Journal of Big Data* 6, 1 (2019), 60.
- [14] Zhiwen Tang and Grace Hui Yang. 2019. Dynamic Search—Optimizing the Game of Information Seeking. *arXiv preprint arXiv:1909.12425* (2019).
- [15] Zhiwen Tang and Grace Hui Yang. 2020. Corpus-Level End-to-End Exploration for Interactive Systems. In *AAAI '20*.
- [16] Luke Taylor and Geoff Nitschke. 2017. Improving deep learning using generic data augmentation. *arXiv preprint arXiv:1708.06020* (2017).
- [17] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. 2017. Domain randomization for transferring deep neural networks from simulation to the real world. In *IROS '17*.
- [18] Jason W Wei and Kai Zou. 2019. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. In *EMNLP '19*.
- [19] Grace Hui Yang, Zhiwen Tang, and Ian Soboroff. 2017. TREC 2017 Dynamic Domain Track Overview. In *TREC '17*.
- [20] Xiangyu Zhao, Long Xia, Jiliang Tang, and Dawei Yin. 2019. "Deep reinforcement learning for search, recommendation, and online advertising: a survey" by Xiangyu Zhao, Long Xia, Jiliang Tang, and Dawei Yin with Martin Vesely as coordinator. *ACM SIGWEB Newsletter* Spring (2019), 1–15.