Towards Identifying for Evidence of Drain Brain from Web Search Results using Reinforcement Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Brain drain is the phenomenon in which experts on a field abandon their origin country to practise their profession in a different country. It forms part of the migration patterns around the world. However *brain drain* can have damaging effects on the source of origin when it happens at large scale. This has been problem in several countries of latinamerica, particularly at the postgraduate level. The correct characterisation of this phenomena is vital to outline polices that keep or attract the talent needed in these countries. In this research, we propose a methodology to identify evidence of *drain brain* through results of web search engines which commonly contains links to career information pages given a seed name, however it could be very time consuming explore and analyse all resulting pages. For this reason, in this research we propose to exploit a Reinforcement Learning setting to learn to navigate and extract significant information from the snippet results. In this work we outline the main architecture based on the Dopamine RL framework.

1 Motivation

Figuring out the career path of the highly qualified diaspora is a major challenge for many nations OECD (2012); Meyer (2015). Sociologist rely on several tools that have developed to measure and quantify the *brain drain* phenomena such as: population censuses, labor force surveys and administrative data. However, these tools have several shortcomings in a more global, mobile and interconnected working force Turner et al. (2015). In face of these challenges, mining the web it could be an alternative approach to better understand the dynamics expert migration flow **?**Lepinay et al. (2014); Auriol et al. (2010).

However, directly mining the web for experts could not be feasible at large scale, since identifying the evidence that a certain professional remained, abandoned or returned to his original country is not trivial. For this reason, we propose to use a Deep Reinforcement Learning (DRL) based method for the extraction of evidence of mobility from snippets of web search results.

We suppose that we have a seed list of names together with a certain time and place in which our potential *brain drain* was in a particular place. This information could come from research institutions which provide scholarships, diasporas organisations such as clubs or societies or from governmental institutions such embassies. Track each member of the list to verify if they continue in the same place could be time consuming, however on one hand to google them could be straight forward but on the other hand to analyse the big amount of returned links also could be out of the question. In this work, we propose the DRL methodology as a semi-supervised method to learn to navigate the web-search snippets. Remember, that snippets contains a relevant segment of the webpage to which its links points to. Our hypothesis is that there is enough information on the snippet so that we could used it to identify evidence of mobility which latter could be used for quantifying *brain drain*.

Submitted to 33rd Conference on Neural Information Processing Systems (NeurIPS 2019). Do not distribute.

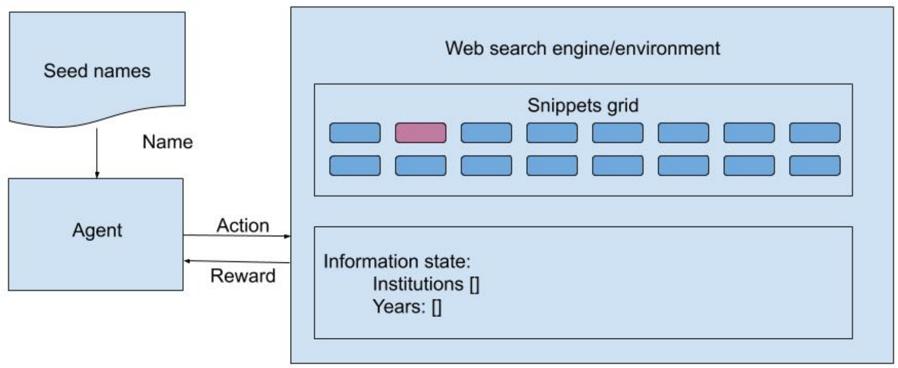

Figure 1: Architecture of DRL for identifying evidence of mobility in web search results.

2 DRL for Information Extraction

We follow the work of Narasimhan et al. Narasimhan et al. (2016) on information extraction from document collections. We model the navigation of the snippets as an optimal policy search in Reinforcement Learning, where we consider a Deep Q-network and various neural network NN models to learn Q-value functions and approximate the sequence of snippets which contain relevant information to the expert mobility. In our setting, there are four possible actions: go to the next snippet, return to the previous snippet, perform a new search and stop. At each step the RL agent looks at the state of the world which consists of *information state* based of what had learnt so far from seeing other snippets and the current snippets.

We had explored the use of a MLP and a LSTM architecture for the Deep Q-network. In the first case, information state and snippet are codified in a one-hot fashion. In the case of the LSTM, information state is codified as a one-hot vector, but the snippet text is passed directly to the LSTM. We hypnotize that the LSTM could be better to approximate the Q-function since it has access to a expression in natural language. The information state consists of bits of Name Entities extracted from the snippets and in our case they are related to the career of a professional. Such as organization names and years. The semi-supervised aspect of our proposal comes from the fact that the place and time information from the seed are used by the DRL as a factor to quantify the reward function, in this way the agent can learn to collect evidence from the snippets directed from this information which is certain for a particular person. This arrangement requires to add new actions related to the management of the information state, such as: add institution, add year, remove institution or remove year. So the agent, not only learns to navigate but to administer the information state. Figure 1 summarises the proposed architecture.

In the current work, we will present our implementation of this framework using the Gym framework Brockman et al. (2016) and the Dopamine Castro et al. (2018). Gym is a toolkit for developing and comparing reinforcement learning algorithms. Dopamine is a fast prototyping of reinforcement learning algorithms. While both have been focused on videogame applications, we have created our agents for the retrieval of textual information.

References

L. Auriol, B. Felix, and M. Schaaper. 2010. Mapping careers and mobility of doctorate holders: Draft guidelines, model questionnaire and indicators. In *OECD Science, Technology and Industry Working Papers*.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. Openai gym. *arXiv preprint arXiv:1606.01540*.

Pablo Samuel Castro, Subhodeep Moitra, Carles Gelada, Saurabh Kumar, and Marc G. Bellemare. 2018. Dopamine: A Research Framework for Deep Reinforcement Learning.

Vincent Lepinay, Andrei Mogoutov, Jean Philippe Cointet, and Lionel Villard. 2014. Russian computer scientists, local and abroad: Mobility and collaboration. In *Proceedings of the 10th Central and Eastern European Software Engineering Conference in Russia*, CEE-SECR '14, pages 18:1–18:6, New York, NY, USA. ACM.

Jean-Baptiste Meyer. 2015. *Diaspora: towards the new frontier*. IRD, Institut de recherche pour le développement, Marseille.

Karthik Narasimhan, Adam Yala, and Regina Barzilay. 2016. Improving information extraction by acquiring external evidence with reinforcement learning. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2355–2365.

OECD. 2012. *Connecting with emigrants: A global profile of diasporas*. Organization for Economic Cooperation and Development, Paris.

William Turner, Jorge Garcia Flores, and Mathilde de Saint Leger. 2015. Computer supporting diaspora knowledge networks: a case study in managing distributed collective practices. In Jean-Baptiste Meyer, editor, *Diaspora: towards the new frontier*, pages 213–243. IRD, Institut de recherche pour le développement, Marseille.