

Imitation learning for structured prediction in natural language processing

Andreas Vlachos^{*}, Gerasimos Lampouras^{*} and Sebastian Riedel[#]

^{*}University of Sheffield
Regent Court, 211 Portobello,
Sheffield, S1 4DP, UK

[#]University College London
One Euston Square, 40 Melton Street,
London, NW1 2FD, UK

Tel: (+44) 114 22 21960, (+44) 114 22 21831, (+44) 203 54 95254
vlachos@sheffield.ac.uk, g.lampouras@sheffield.ac.uk,
s.riedel@live.ucl.ac.uk

1 Tutorial Objectives

Imitation learning is a learning paradigm originally developed to learn robotic controllers from demonstrations by humans, e.g. autonomous flight from pilot demonstrations. Recently, algorithms for structured prediction were proposed under this paradigm and have been applied successfully to a number of tasks including syntactic dependency parsing, information extraction, coreference resolution, dynamic feature selection, semantic parsing and natural language generation. Key advantages are the ability to handle large output search spaces and to learn with non-decomposable loss functions. Our aim in this tutorial is to have a unified presentation of the various imitation algorithms for structure prediction, and show how they can be applied to a variety of NLP tasks.

All material associated with the tutorial will be made available through <https://sheffieldnlp.github.io/ImitationLearningTutorialEACL2017/>.

2 Part I: Imitation Learning

In the first part, we will give a unified presentation of imitation learning for structured prediction focusing on the intuition behind the framework. We will then delve into the details of the different algorithms that have been proposed so

far under the imitation learning paradigm, including SEARN (Daumé III et al., 2009), DAGGER (Ross et al., 2011), v-DAGGER (Vlachos and Clark, 2014), and LOLS (Chang et al., 2015). Furthermore, we will give a comparative overview of the various algorithms presented that will expose their differences and their practical implications. We will conclude by discussing the relation of imitation learning to recurrent neural networks, bandit learning, adversarial learning, and reinforcement learning.

3 Part II: Applications in NLP

In the second part, we will discuss recent work applying imitation learning methods in the context of NLP. We will begin by reviewing the application of imitation learning to syntactic dependency parsing and discuss how to create expert policies, also known as dynamic oracles (Goldberg and Nivre, 2013). Furthermore, we will review how imitation learning is applied to semantic parsing (Goodman et al., 2016), and how it can benefit natural language generation (Lampouras and Vlachos, 2016), where the search space is all English sentences. In this process we will explain techniques that can enhance imitation learning and solve problems that arise in each practice, such as focused costing, noise reduction, targeted exploration, sequence correction, and change propagation. Finally, we will briefly present applications on biomedical event extraction (Vlachos and Craven, 2011), feature selection (He et al., 2013), coreference resolution (Clark and Manning, 2015), autonomous driving learning (Zhang and Cho, 2016), and pruning policies for syntactic parsing (Vieira and Eisner, 2016).

4 Structure

Part I: Imitation Learning Algorithms (90 minutes)

- Introduction, basic concepts and intuition: Structured prediction basics, per-action and end-to-end supervision, and decomposability.
- Detailed algorithm descriptions: Exact imitation, DAGGER, roll outs and non-decomposable losses, latent variables, and cost-sensitive classification.
- Advanced topics: Imitation learning by coaching, imitation learning for recurrent neural network training, bandit learning, adversarial learning, reinforcement learning.

Break (15 minutes)

Part II: Applications (90 minutes)

- Syntactic dependency parsers: Expert policy specification.
- Natural language generation: Sequence correction.
- Semantic parsing: Abstract meaning representation, focused costing, noise reduction, and targeted exploration.

- Brief presentation of other applications: Biomedical event extraction, feature selection, coreference resolution, autonomous driving learning, pruning policies for syntactic parsing.

5 About the Speakers

Andreas Vlachos (<http://andreasvlachos.github.io/>) is a lecturer at the University of Sheffield, working on the intersection of Natural Language Processing and Machine Learning. Current projects include natural language generation, automated fact-checking and imitation learning. He has also published research on semantic parsing, language modelling, information extraction, active learning, clustering and biomedical text mining.

Gerassimos Lampouras (<http://glampouras.github.io/>) is a postdoc at the University of Sheffield, working with Andreas Vlachos on developing domain-independent Natural Language Generation frameworks using Imitation Learning. Previously he was a postdoc at the Machine Reading group at University College London working with Andreas Vlachos and Sebastian Riedel. Before that he received his PhD from Athens University of Economics and Business while supervised by Ion Androutsopoulos. His research interests revolve around natural language generation, machine learning and global optimization approaches for natural language processing. His previous work also includes text summarization and question answering.

Sebastian Riedel (<http://www.riedelcastro.org/>) is a Reader at University College London and leading the Machine Reading Lab. He is an Allen Distinguished Investigator, Marie Curie fellow and received a Google Focused Research award. Sebastian is generally interested in the intersection of Natural Language Processing and Machine Learning, and particularly interested in teaching machines to read and to reason with what was read. He has published various papers on this topic, covering questions related to scalable inference, probabilistic programming and representation learning.

References

- Chang, K.-W., Krishnamurthy, A., Agarwal, A., Daumé III, H., and Langford, J. (2015). Learning to search better than your teacher. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Clark, K. and Manning, C. D. (2015). Entity-centric coreference resolution with model stacking. In *Association for Computational Linguistics (ACL)*.
- Daumé III, H., Langford, J., and Marcu, D. (2009). Search-based structured prediction. *Machine Learning*, 75:297–325.
- Goldberg, Y. and Nivre, J. (2013). Training deterministic parsers with non-deterministic oracles. *Transactions of the Association for Computational Linguistics*, 1(Oct):403–414.

- Goodman, J., Vlachos, A., and Naradowsky, J. (2016). Noise reduction and targeted exploration in imitation learning for abstract meaning representation parsing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- He, H., Daumé III, H., and Eisner, J. (2013). Dynamic feature selection for dependency parsing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1455–1464, Seattle.
- Lampouras, G. and Vlachos, A. (2016). Imitation learning for language generation from unaligned data. In *Proceedings of COLING 2016*.
- Ross, S., Gordon, G. J., and Bagnell, D. (2011). A reduction of imitation learning and structured prediction to no-regret online learning. In *14th International Conference on Artificial Intelligence and Statistics*, pages 627–635.
- Vieira, T. and Eisner, J. (2016). Learning to prune: Pushing the frontier of fast and accurate parsing. *Transactions of the Association for Computational Linguistics (TACL)*, 4.
- Vlachos, A. and Clark, S. (2014). A new corpus and imitation learning framework for context-dependent semantic parsing. *Transactions of the Association for Computational Linguistics*, 2:547–559.
- Vlachos, A. and Craven, M. (2011). Search-based structured prediction applied to biomedical event extraction. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*, pages 49–57, Portland, Oregon, USA. Association for Computational Linguistics.
- Zhang, J. and Cho, K. (2016). Query-efficient imitation learning for end-to-end autonomous driving. *CoRR*, abs/1605.06450.