

## RESEARCH INTERNSHIP PROPOSAL

Title	Efficient RL algorithms for online and continual dialogue policy learning
Advisor	Prof. Fabrice Lefèvre
<p><i>Topic:</i> Artificial Intelligence, machine Learning, Reinforcement Learning, Deep Learning.  <i>City and country:</i> Avignon, France  <i>Team or project in the lab:</i> LIA – Vocal Interactions Team  <i>Name and mail of the advisor:</i> Prof. Fabrice Lefèvre (@univ-avignon.fr)  <i>Name and mail of the head of the department:</i> Prof. Yannick Estève (@univ-avignon.fr)</p> <p><i>General presentation of the project</i>  Most RL policy learning algorithms are exclusively oriented towards convergence to an optimal solution. Yet in many cases dynamics of the environment might not be totally stationary. For instance in Natural Language Processing, and Dialogue Systems in particular, users’ preferences, abilities etc may greatly evolve over time (in relation to increased experience and familiarity with the interface). In such cases an initial one-for-all training is clearly sub-optimal and online and continual learning should be used instead. Many solutions have been proposed recently to improve stochastic policy learning but not necessarily implemented and tested in an NLP context such as Dialogue Modeling.</p> <p><i>Objective of the internship</i>  The objective of the project will be to study the possibility to address optimality efficiency of deep learning for RL (as instantiated in (double)-DQN [1], Soft Actor-Critic [2] and others algorithms [3]) in the context of data sparsity (sample efficiency) and online dynamic learning. Various tasks could serve as testbed during the project, yet the main overall application foreseen is the training of a vocal interaction dialogue manager, as in [4]. Showcase of the techniques studied and proposed could be done in any language (Python, R, Matlab etc).</p> <p><i>Bibliographic references</i>  [1] D. Silver, A. Huang, C. Maddison, <i>et al.</i> Mastering the game of Go with deep neural networks and tree search. <i>Nature</i> <b>529</b>, 484–489 (2016).  [2] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” <i>35th Int. Conf. Mach. Learn. ICML 2018</i>, vol. 5, pp. 2976–2989, 2018.  [3] K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, “A Brief Survey of Deep Reinforcement Learning,” <i>IEEE SIGNAL Process. Mag. Spec. ISSUE Deep Learn. IMAGE Underst.</i>, Aug. 2017.  [4] E. Ferreira and F. Lefèvre, “Reinforcement-learning based dialogue system for human-robot interactions with socially-inspired rewards,” <i>Comput. Speech Lang. Spec. issue Speech Lang. Interact. Robot.</i>, vol. 34, no. 1, pp. 256–274, 2015.</p> <p><i>Expected ability of the student</i>  Interest in mathematical foundations of machine learning applied to NLP. Good skills in math (theory and programming). Good command of programming environments (OS, CLI, SDK, virtual env etc).</p>	
Duration	6 months
Remuneration	Around 529€ / month
Topics	Artificial Intelligence, Machine Learning, Reinforcement Learning, Deep Learning.