Accepted Manuscript

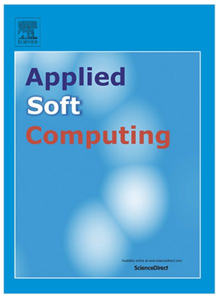

Reinforcement learning applied to Forex trading

João Carapuço, Rui Neves, Nuno Horta

PII: S1568-4946(18)30534-9
DOI: https://doi.org/10.1016/j.asoc.2018.09.017
Reference: ASOC 5096

To appear in: *Applied Soft Computing Journal*

Received date : 2 February 2018
Revised date : 20 July 2018
Accepted date : 16 September 2018

Please cite this article as: J. Carapuço, et al., Reinforcement learning applied to Forex trading, *Applied Soft Computing Journal* (2018), https://doi.org/10.1016/j.asoc.2018.09.017

Reinforcement Learning Applied to Forex Trading

João Carapuço, Rui Neves, Nuno Horta
jmcarapuco@gmail.com, rui.neves@tecnico.ulisboa.pt, nuno.horta@tecnico.ulisboa.pt

Instituto de Telecomunicações, Instituto Superior Técnico, Torre Norte, Lisboa, Portugal

Abstract

This paper describes a new system for short-term speculation in the foreign exchange market, based on recent reinforcement learning (RL) developments. Neural networks with three hidden layers of ReLU neurons are trained as RL agents under the Q-learning algorithm by a novel simulated market environment framework which consistently induces stable learning that generalizes to out-of-sample data. This framework includes new state and reward signals, and a method for more efficient use of available historical tick data that provides improved training quality and testing accuracy. In the EUR/USD market from 2010 to 2017 the system yielded, over 10 tests with varying initial conditions, an average total profit of 114.0±19.6% for an yearly average of 16.3±2.8%.
Keywords: Machine Learning, Neural networks, Reinforcement Learning, Financial Trading, Foreign Exchange

1. Introduction

Reinforcement learning (RL) is a sub-field of machine learning in which a system learns to act within a certain environment in a way that maximizes its accumulation of rewards, scalars received as feedback for actions. It has of late come into a sort of Renaissance that has made it very much cutting-edge for a variety of control problems. Some high-profile successes ushered in this new era of reinforcement learning. First, in 2013, a London-based artificial intelligence (AI) company called Deepmind, stunned the AI community with a system based on the RL paradigm that had taught itself to play 7 different Atari video-games, 3 of them at human-expert level, using simply pixel positions and game scores as input and without any changes of architecture or learning algorithm between games [1]. Deepmind was bought by Google, and by 2015 the system was achieving performances comparable to professional human game testers in over 20 different Atari games [2]. Then, that same company achieved wide mainstream exposure when its Go-playing program AlphaGo, which uses a somewhat similar approach to the Atari playing system, beat the best Go player in the world in an event that reached peaks of 60 million viewers.

This was made possible by the use of neural networks, another sub-field of machine learning. These networks consist of interconnected artificial neurons inspired by the biological brain which process information by their dynamic state response to external inputs. This combination has been used before but often proved unreliable, especially for more complex neural networks. Advances in the neural networks field such as the ReLU neuron and gradient descent algorithms with adaptive learning rate, along with contributions from Mnih et al. [1] (and many others afterwards) aimed specifically at this mixed approach, have made it much more effective. From the point of view of reinforcement learning, neural networks provide much needed generalization power to find patterns for decision making that lead to greater reward. From the point of view of neural networks, reinforcement learning is useful because it automatically generates great amounts of labeled data, even if the data is more weakly labeled than through direct human intervention, which is usually a limiting factor for neural networks [3].

In this paper our aim is to find how to adapt these new developments in RL to the creation of an algorithmic system that generates profitable trading signals in financial markets. This requires successfully accomplishing the following steps:

- Obtain a system that is able to stably, without diverging, learn and thus improve its financial performance on the dataset it is being trained on;
- Show that the system's learning generalizes to unseen data and that this generalization power can be harnessed to generate profitable decisions on a realistic simulation of live trading.

The foreign exchange market (Forex) was chosen as the testing ground for accomplishing these goals as having the largest volume of trades out of all financial markets, with roughly 25% of that volume concentrated on the EUR/USD pair [4] makes it ideal for short-term speculation. To reach the objectives above the main contributions offered in this work are: