

Artificial Intelligence and applications in Finance

1. Research area

The aim of the research is to create new and novel architectures for Artificially Intelligent agents with the ability to perceive their environment, learn how the environment operates, construct a model of the environment, and finally leverage this model to optimize sequential actions (decisions) under uncertainty. Whilst this architecture is sufficiently broad and can be used for a diverse set of tasks, the research is primarily focused on applications in finance, such as asset pricing, portfolio optimization and risk management.

2. Problem statement

Current AI agents are trained using reinforcement learning techniques which require a large number of trials for the agent to attain mastery in a given task via trial-and-error. In the case of many real-world tasks, such as financial trading agents, portfolio optimizer agents, autonomous cars, chatbots, etc., the trial-and-error approach is unsuitable for deployment in the real world. For example, a financial trading agent cannot risk incurring huge financial losses until it learns the correct behavior. Furthermore, AI agents typically learn to perform well on the environment they are trained on, however they do not perform well in different or entirely novel tasks. On the other hand, humans can easily learn to perform new tasks and achieve relatively good performance with just a few minutes of practice, since they are able to re-use existing knowledge to generalize across tasks and transfer their knowledge to entirely new domains.

Research student
Achilleas Bekiaris

1st Supervisor
George Magoulas
(Computer Science)

2nd Supervisor
Andrea Cali
(Computer Science)

3. Methodology

The objective of the research is to develop AI agents in a modular approach, where the complete agent will be a sum of various sub-components. The agent is assumed to operate in a POMDP (Partially Observable Markov Decision Process) framework.

To accelerate training and minimize computational requirements, the agent will have access to the output of specialized components such as pre-trained models for various tasks (e.g. language model like BERT, asset pricers etc). These pre-trained modules will serve as sensory inputs and facilitate Transfer Learning. Furthermore, the agent will be endowed with a “memory” module where past interactions, states, probability densities, transition dynamics and rewards will be stored.

Finally, the actions of the AI agent will be coordinated by a central “decision” module capable of processing the sensory inputs from the pre-trained models, evaluate these against past memory and learn to optimize decision making given the state of the environment.

The complete agent will be trained end-to-end using techniques such as Reinforcement Learning and unsupervised techniques such as Variational Autoencoders and GANs (Generative Adversarial Networks). The proposed architecture is displayed in Figure 1 below.

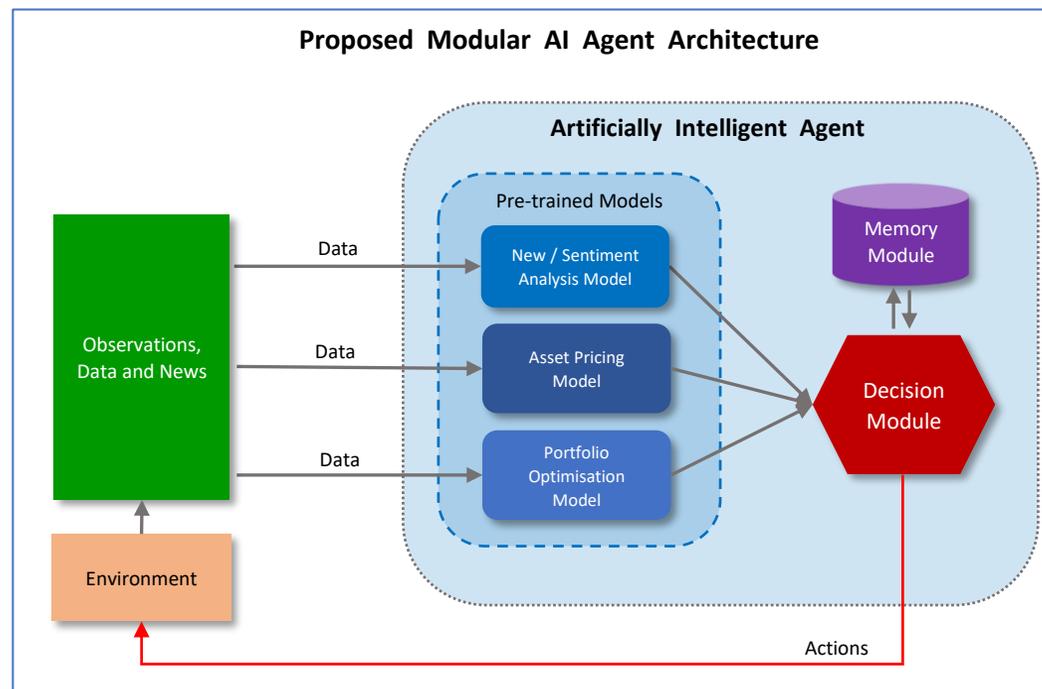


Figure 1. Proposed modular architecture for an AI agent with an ability to perceive the world via inputs (e.g. Data feeds, Language), ability to learn a set of Models and store acquired knowledge in memory. The Decision Module will be able to use the acquired knowledge of the World Models to make optimal actions and interact with the environment.