

Contribution ID: 37

Type: Talk (15min + 5min)

Deep Reinforcement Learning in agent-based model AgriPoliS to simulate strategic land market interactions

Thursday, March 7, 2024 9:00 AM (20 minutes)

AgriPoliS (Agricultural Policy Simulator) is an agent-based model for simulating the development of agricultural regions, focusing on the structure changes under economical, ecological and societal factors. The farms in the region are modeled as agents in AgriPoliS, which interact with each other through different markets, most importantly, land market. Every year the agents make their decisions about bidding for new land plots, stable and machine investments, and production processes through mixed linear programming (MIP) optimizations, to maximize their profit, where the focus is only for the current year with no regard for the future implication of these decisions. In this work, we enhance the agents with Deep Reinforcement Learning, giving them the ability to develop strategic instead of myopic decisions to maximize their long-term profits within the simulation period through strategic bidding behavior in the land market.

As the first step, only one agent is enhanced with Reinforcement Learning while the other agents adopt the standard behavior. As we are interested in the effects of the strategic bidding behavior of the agent, we formulate bidding prices as the action space. The state space consists of the state variables of the agents and the region under investigation, which include liquidity, current stables and machines, the distribution of remaining contract duration for rented land, rent prices, spatial distribution of free land plots in the region and the distribution of competitive agents in the neighborhood. Equity capital of the agent is considered as the reward from the environment by taking an action. Using the state variables, the agent chooses the action (bidding price) to present to the land market. Based on the success and/or failure of his bid, the agent proceeds to make investment and production decisions and obtains the results of his decisions. This continues until the end of the simulation period when the cumulative equity capital is accessed.

The learning framework consists of two parts, AgriPoliS and learning algorithm. AgriPoliS, which is implemented in C++, functions as the environment. It takes the action from the learning algorithm, which is implemented in Python, and deliver the states and rewards to the learning algorithm. The communication between the two parts is realized through the message queue system zeorMQ.

Since the action space is continuous, we implemented the DDPG (Deep Determinant Policy Gradient) algorithm with PyTorch. After tuning the learning hyper parameters, the enhanced agent could learn a stable strategy, which varies the relative bidding prices and maximizes the cumulative rewards. The first results are promising, the effect of the agent changing its bidding behavior not only affects its equity capital but also the equity capital of other farms. More algorithms like TD3 (Twin Delayed DDPG) and SAC (Soft Actor-Critic) are under work and we are also interested to resolve the learning stability issue of the algorithms.

Slot length

Primary authors: DONG, Changxing (Leibniz Institute for Agricultural Development in Transition Economies (IAMO)); NJIRU, Ruth; APPEL, Franziska; BALMANN, Alfons

Presenter: DONG, Changxing (Leibniz Institute for Agricultural Development in Transition Economies (IAMO))

Session Classification: AI/ML Research Software

Track Classification: HPC and Massive Data: AI/ML Research Software