Rebellion
RESEARCH
Ai Asset Management

Deep Reinforcement Learning & Nowcasting's Fatal Flaw

Alexander Fleiss

CEO
RebellionResearch.com

Artificial Intelligence

"The art of creating machines that perform functions that require intelligence when performed by people"

– Ray Kurzweil, 1999 U.S. National Medal of Technology recipient

Approaches to AI

- Rule-based algorithms
  - Requires the programmer to be an expert in the field
  - Difficult to adapt to other problems
- Machine Learning Algorithms
  - Creates rules and procedures using a set of relevant data
  - Uses mathematical optimization techniques

Machine Learning Advantages

- Expert knowledge is not necessary
- Computer can be “smarter” than the programmer
- Machine learning algorithms can be adapted to a wide variety of problems

AI vs. Human decision making

- Benefits of using AI
  - Incredible computation ability
  - Perfect recall
  - Dispassionate
  - Superior pattern recognition

AI vs. Human decision making

- Disadvantages of using AI
  - Framing the problem
  - Cannot come up with its own factors
  - Overfitting

Overfitting

- Spurious Correlations
  - Garbage in Garbage out
- Model complexity
  - Giving the algorithm too much leeway with performing a task

When to use AI

- Large set of relevant training data
- Complicated interconnected system
- Easy to mathematically describe the problem

- Why not use AI in financial markets?

AI For Short-term Trading

- Inefficiencies Exploited:
  - Human emotional and irrational decision-making
  - Market Making
- Large set of training data
- Easy to determine relevant factors
- Stable system

AI Investing

- Challenges with Longer Horizon investing
  - System constantly evolving
    - Market conditions constantly changing
    - Meanings of factors changes with time
  - High noise
  - System extremely complex. Many different factors contribute to stock returns.

AI Investing in Practice

- Inefficiencies Exploited:
  - Herd mentality
  - Improperly priced risk/reward characteristics
- Large set of factors chosen which correspond to different investment styles
  - Factors are stable
  - Contain a large amount of information

AI Investing in Practice

- Modified Bayesian learner to prevent overfitting
  - Bayesian learners update probabilities given new pieces of information
  - Rate of learning can be controlled
  - Ensures the algorithm has innate biases to reduce variance and does not become too precise in predictions

AI Investing in Practice

- Overcoming the LTCM/Bangladeshi Butter Problem
  - Each stock investment relies on average 30-40 factors
  - “Bad” factors expected to lead to market performance of the stock
  - The presence of several good factors still leads to expected outperformance

Deep Reinforcement Learning

Deep RL

What is Reinforcement Learning?

Deep RL

- - An agent learns the optimal policies to navigate an environment (i.e. learning optimal indicators to buy/sell stock, optimal moves to solve a puzzle, etc…)
- - State-action pairs are stored in a discrete data structure (i.e. state = stock price, action = sell)

Deep RL

- Agent makes decisions, and a reward scheme punishes/rewards that decision and updates the value of that state-action pair, with potential discounting if necessary

Deep RL

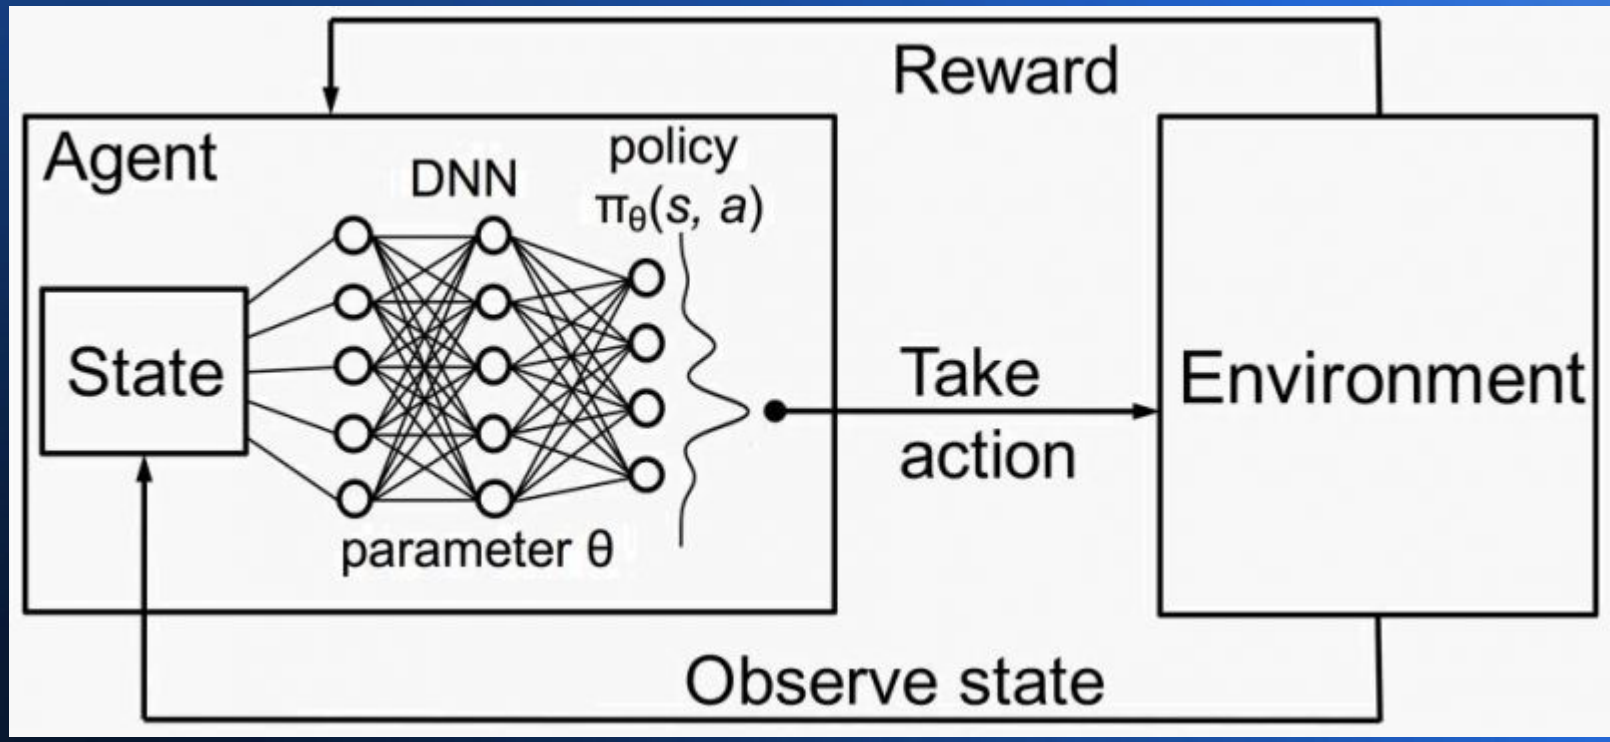
Reward
Agent
DNN
policy
$\pi_\theta(s, a)$
State
Take
action
Environment
parameter θ
Observe state

Deep RL

- When training, the agent either randomly picks actions or exploits the best action given state (epsilon-greedy)

- When testing, the agent picks the best action given state

Deep RL

- Reinforcement Learning needs exposure to lots of data, and training ends upon convergence (very computationally demanding)

Drawbacks of Reinforcement Learning

Deep RL

- **Tendency to overfit**: the agent can identify ideal state-action pairs for training environments but there is a great chance of overfitting and poor performance while testing

Deep RL

-**Very data demanding:** The agent must be trained on a very large set of data to ensure it knows how to deal with a range of situations

Deep RL

- **Must have discrete states:** Storing state-action values in a matrix or similar data structure is only feasible with limited, discrete state values

Deep RL

Why Deep Reinforcement Learning?

Deep RL

- Rather than storing state-action values in a discrete data structure such as a matrix, through training, a neural network is trained with rewards

Deep RL

- Example: input is a 1-D vector of state information and output is a gradient vector containing the estimated rewards of taking each possible action

Deep RL

- Deep Reinforcement Learning is a better choice than traditional reinforcement learning for environments that do not have discrete states

Nowcasting

The idea of predicting the future of the stock market through current economic data

Nowcasting

Why does it fail?

Nowcasting

The market is not a perfect mechanism: functionality of long term investing is not based on daily data.

Nowcasting

Nowcasting sometimes waits for too long and by then, the changes have already been incorporated in the price.

eg. Shocks due to news are incorporated almost immediately. When we finish web scraping the weeks data and take a position based on it, we are already behind the crowd.

Nowcasting

Nowcasting can only work when using specific types of data and has been shown to not work using most of the available data out there.

Rebellion
RESEARCH
Ai Asset Management

Questions

Alexander Fleiss
CEO
RebellionResearch.com