

Reinforcement Learning in Finance



Reinforcement Learning in Finance

Before we can adequately explore the applications of reinforcement learning in finance, we must first define reinforcement learning and how it relates to computer science.

Reinforcement learning, supervised learning, and unsupervised learning are the three branches of machine learning methods in artificial intelligence.

Supervised learning is the approach to machine learning that uses labeled data sets. The data sets are designed to train the algorithms into classifying data or predicting accurate outcomes. With labeled inputs and outputs, the model can measure its accuracy and thereby learn over time.

Unsupervised learning relies on machine learning algorithms to analyze and cluster unlabeled data sets. The algorithm to find hidden patterns in the data without needing human intervention for dynamic programming.

Reinforcement learning is different from the other two because it is based on the idea of trial-and-error decision making which measures learning through the idea

of rewards rather than through labeled data.

Reinforcement learning is the process of training machine learning models to make a series of decisions. The model generally starts from random trials and then trains itself to use a more complicated model.

In other words, reinforcement learning is a learning process in which an algorithm interacts with its environment using trial-and-error to reach a predefined goal. The approach is designed so that the learning agent can maximize the reward while minimizing the penalties for each correct step it takes to reach the goal.

How is reinforcement learning from deep learning?

Deep learning is another one of the popular methods of machine learning which is commonly used in financial markets.

Where reinforcement learning uses a system of penalties and rewards to force the computer to solve problems by itself with limited human involvement, deep learning is a model based on the human brain. The model uses a deeper data set that consists of neural network layers to help computers learn more about abstract features of the particular data. That's why deep learning is particularly useful in forecasting in finance.

With that out of the way let's look at some key phrases that are important to know before we move on to the practical applications of reinforcement learning in finance.

- **Deep Reinforcement Learning (DRL):** Algorithms that use deep learning to approximate value or policy functions at the core of reinforcement learning.
- **Policy Gradient Reinforcement Learning Technique:** A method used in solving reinforcement learning problems. These target modeling and optimizing the policy function.
- **Deep Q Learning:** This involves using a neural network to approximate Q value function which creates an exact matrix for the working agent. The working agent can then refer to this matrix to maximize its long-term reward.
- **Gated Recurrent Unit (GRU):** This is a special type of recurrent neural

network that is implemented using a gating mechanism.

- **Gated Deep Q Learning Strategy:** This is a combination of Deep Q Learning and GRU.
- **Gated Policy Gradient Strategy:** This uses a combination of policy gradient technique and GRU
- **Deep Recurrent Q Network:** This approach combines the recurrent neural networks and the Q learning technique.

Common Reinforcement Learning Algorithms

Reinforcement learning doesn't rely on a specific algorithm but instead consists of multiple algorithms that use similar approaches. The difference is between algorithms why mainly in their strategies for exploring environments. The reinforcement learning framework is adaptive.

- **State-action-reward-state-action (SARSA):** In this reinforcement learning approach, the agent is given what's known as a policy. The optimal policy is nothing more than a probability that gives it the odds of certain actions resulting in rewards.
- **Q-learning:** With this reinforcement framework, the agent doesn't receive a policy so its exploration of the environment is self-directed.
- **Deep Q-Networks:** These reinforcement learning algorithms use neural networks in addition to reinforcement learning techniques. They use self-directed exploration while future actions are based upon a random sample of past beneficial actions the neural network learned.
- **Actor-critic:** This is a temporal difference version of policy gradient, made up of two networks - the actor and the critic. The actor decides which action to take, and the critic tells the action how good the action was and how it should adjust in the short term.

Why Use DRL for Stock Trading?

DRL does not require a large labeled training data set. This is particularly advantageous because the amount of data we have available grows exponentially every day. If we had to label add a set of its size, it becomes both very time-

consuming and labor-intensive.

Since the goal of stock trading is to maximize returns while avoiding risk, DRL solves the optimization problem by maximizing the expected return from future actions over a certain period of time. Stock trading provides a continuous process of testing new ideas, getting market feedback, and attempting to optimize trading strategies over time. It's possible to model stock trading processes like the Markov decision process which serves as the very foundation of reinforcement learning.

It's been shown that DRL algorithms can easily outperform human players in a variety of situations. By defining the reward function as the change of portfolio value, DRL maximizes portfolio value over time. Because the stock market provides sequential feedback, DRL sequentially increases the model performance throughout the training process. The exploration-exploitation technique balances out different things and taking takes advantage of what's been figured out which is different from other learning algorithms. Plus, there's no need for any skilled humans to provide labeled samples or training examples. During the exploration process, the agent is encouraged to explore areas that have been uncharted by humans.

DRL also has experience replay because it is able to overcome the correlated samples issues. it does this by sampling many batches of transitions from a pre-saved replay memory randomly. Since it uses continuous action space, it can handle large dimensional data. DRL is empowered by neural networks that are powerful enough to handle large state space and action space, unlike Q learning.

Machine learning helps in quantitative finance because, without it, it would be impossible to analyze all the large datasets. However, it is not foolproof, and for the risk-averse, is less than ideal.

Trading Bots

One of the simplest applications of reinforcement learning in finance is bots that are powered with that can learn from the stock market environment and trading simply by interacting with it. By using trial and error to optimize their learning strategy based on the characteristics of all the stocks listed in the stock market

these bots help to:

- Save time
- Diversified trading across all Industries
- Trade on a twenty-four-hour basis

Chatbots

Typically, chatbots are trained using the help of a sequence-to-sequence model. However, adding reinforcement learning to their training offers advantages for stock trading and finance.

- These chatbots can provide real-time quotes to their user operators and act as brokers.
- Conversational user interface-based chatbots can serve on a customer service team to help people solve their problems. This approach saves time and keeps the support staff from using their resources on easily repeatable tasks so they can focus on more complex issues.
- It's also possible for chatbots to provide suggestions on opening and closing sales values within trading hours.

Peer-to-Peer Lending Risk Optimization

Peer-to-peer lending has gained popularity in recent years because it's an easy way to provide both individuals and businesses with loans online. There are many online services, such as Lending Club that provide a matching service between lenders and investors.

With this kind of marketplace, reinforcement learning is particularly helpful. You can use it to:

- Analyze borrowers' credit scores to reduce risk
- Estimate the likelihood of the borrower being able to meet their debt obligations
- Predict analyzed returns. As online businesses have lower overhead, lenders can reasonably expect higher returns compared to the investment and savings products offered by traditional banks

Portfolio Management

Portfolio management refers to taking assets, putting them into the stocks, and managing them continuously to help you or your clients achieve their financial goals. Using Deep Policy Network reinforcement learning, you can optimize the allocation of assets over time for portfolio optimization. Deep reinforcement learning offers the following benefits here:

- Enhanced efficiency and success rate for human managers
- Decreased organizational risk
- Increased return on investment (ROI) for organizational profit

Price Setting Strategies

One of the most difficult parts of understanding stock prices is the complex and dynamic nature of price changes. To understand these properties, Gated Recurrent Unit (GRU) networks work well with reinforcement learning because they provide advantages including:

- Extracting the informative financial features that can represent intrinsic characters of individual stocks
- Helping to determine the stop loss and stock profit during trading to reduce transaction costs.

Recommendation Systems

With online trading services, recommendation systems that are based on reinforcement learning techniques are important. When trained well, these systems could help recommend the right stocks to users while they are trading. Reinforcement learning how to choose the best stock or mutual funds after they've been trained on a number of stocks, which leads to a better return on investment.

Maximizing Profit

Combining all of the points above, it's possible to get an automated system built with the goal of achieving high returns during financial trading while

simultaneously keeping the initial Investments as low as possible.

An agent could be trained with the help of reinforcement learning to take the minimum asset from any source and allocate it to the right stock to double the return in the future.

In today's environment, enforcement learning agents are able to learn optimal trading strategies that go beyond simple buy and sell strategies that people typically apply. You can achieve this with the help of the Markov decision process model using a deep recurrent Q Network.

While all of this is certainly impressive, many of the projects done today are essentially for fun. They're trained with past data but aren't necessarily back-tested properly. In the instance of unforeseen data, the downside risk is much larger than the model can expect.

The stock market is a complicated system and it's hard for any machine learning system to understand stocks based on historical data only. The performance of machine learning-based trading strategies can be great, but it is also possible to drain savings, so always take these projects with a grain of salt. If they were 100% accurate 100% of the time, everyone would be rich.

Challenges of Reinforcement Learning

Though reinforcement learning has high potential, it can be difficult to deploy and its application remains limited. A major barrier for deployment for this type of machine learning is that it relies on the exploration of the environment.

For instance, if you were to deploy a robot that was reliant on reinforcement learning to navigate a complex physical environment, it seeks new States and takes different actions as it moves. It is therefore difficult to consistently take the optimal action for future rewards in this real-world environment because of how frequently the environment changes.

Because it requires so much time to ensure the learning is done correctly with this method, it does put a limit on its usefulness. When you consider how insensitive it can be on computing resources when the training environment becomes too complex, the demands on time and compute resources increase. As

such, many opt for supervised or semi-supervised learning in its place. It can deliver more efficient and faster results than reinforcement learning with the proper amount of data available since it can be used with fewer resources and achieve the expected reward.

Ultimately, both deep learning and reinforcement learning belong in the financial market and can have great applications. However, we need to practice more to maximize both methods in stock trading systems on a case-by-case basis, rather than trying to choose one or the other to use a broad application.

These two approaches are not comparable as apples to apples but rather apples to oranges and thereby must be used in different applications where they make the most sense.

What's your goal today?

1. Use PLANERGY to manage purchasing and accounts payable

We've helped save billions of dollars for our clients through better spend management, process automation in purchasing and finance, and reducing financial risks. To discover how we can help grow your business:

- Read our case studies, client success stories, and testimonials.
- Visit our "Solutions" page to see the areas of your business we can help improve to see if we're a good fit for each other.
- Learn about us, and our long history of helping companies just like yours.

Book a Live Demo

2. Download our guide "Preparing Your AP Department For The Future"

Download a free copy of our guide to future proofing your accounts payable department. You'll also be subscribed to our email newsletter and notified about new articles or if have something interesting to share.

download a free copy of our guide

3. Learn best practices for purchasing, finance, and more

Browse hundreds of articles, containing an amazing number of useful tools, techniques, and best practices. Many readers tell us they would have paid consultants for the advice in these articles.

Related Posts