

DEEP REINFORCEMENT LEARNING FOR CRYPTOCURRENCY TRADING: PRACTICAL APPROACH TO ADDRESS BACKTEST OVERFITTING

Anonymous authors

Paper under double-blind review

ABSTRACT

Designing profitable and reliable trading strategies is challenging in the highly volatile cryptocurrency market. Existing works applied deep reinforcement learning methods and optimistically reported increased profits in backtesting, which may suffer from the *false positive* issue due to overfitting. In this paper, we propose a practical approach to address backtest overfitting for cryptocurrency trading using deep reinforcement learning. First, we formulate the detection of backtest overfitting as a hypothesis test. Then, we train the DRL agents, estimate the probability of overfitting, and reject the overfitted agents, increasing the chance of good trading performance. Finally, on 10 cryptocurrencies over a testing period from 05/01/2022 to 06/27/2022 (during which the crypto market **crashed two times**), we show that the less overfitted deep reinforcement learning agents have a higher return than that of more overfitted agents, an equal weight strategy, and the S&P DBM Index (market benchmark), offering confidence in possible deployment to a real market.

1 INTRODUCTION

A profitable and reliable trading strategy in the cryptocurrency market is critical for hedge funds and investment banks. Deep reinforcement learning methods prove to be a promising approach (Fang et al., 2022), including crypto portfolio allocation (Jiang & Liang, 2017; Ang et al., 2022) and trade execution (Hambly et al., 2021). However, three major challenges prohibit the adoption and deployment in a real market: 1) the cryptocurrency market is highly volatile; 2) the historical market data have a low signal-to-noise ratio (Conrad et al., 2018); and 3) there are large fluctuations (e.g., market crash) in the cryptocurrency market.

Existing works may suffer from the backtest overfitting issue. Many methods (Xiong et al., 2018; Liu et al., 2021; Yang et al., 2020) adopted a walk-forward backtest method and optimistically reported increased profits in backtesting. The walk-forward method divides data in a training-validation-testing manner, but using a single validation set can easily result in overfitting. Another approach (Jiang & Liang, 2017) considered a k -fold cross-validation method (essentially leave one period out) with an assumption that the training set and the validation sets are drawn from an IID process, which does not hold in financial tasks. For example, (Liu & Tsyvinski, 2021) showed that there is a very strong time-series momentum effect in crypto returns. Finally, DRL algorithms are highly sensitive to hyperparameters, resulting in high variability of DRL algorithms' performance (Clary et al., 2019; Henderson et al., 2018; Mania et al., 2018). A researcher may get 'lucky' on the validation set and obtain a *false positive* agent. A key question from practitioners is that "*are the results reproducible under different market situations?*"

This paper proposes a practical approach to address the backtest overfitting issue. Researchers in the DRL cryptocurrency trading niche submit papers containing overfitted backtest results. The value of a quantitative metric for detecting overfitting will be significant. First, we formulate the detection of backtest overfitting as a hypothesis testing. Such a test employs an estimated probability of overfitting to determine whether a trained agent is acceptable. Second, we provide detailed steps to estimate the probability of backtest overfitting. If the probability exceeds the threshold of the hypothesis test, we reject it. Finally, on 10 cryptocurrencies over a testing period from 05/01/2022

to 06/27/2022 (during which the crypto market **crashed two times**), we show that the less overfitted deep reinforcement learning agents have a higher return than that of the more overfitted agents, an equal weight strategy, and the S&P DBM Index (market benchmark), offering confidence in possible deployment to a real market. The DRL-based strategy appears to be related to the volatility-managed strategies - buy (sell) when volatility is high (low), and is popular in the asset management industry (Moreira & Muir, 2017). We hope DRL practitioners may apply the proposed method to ensure that their chosen agents do not overfit during the training period.

The remainder of this paper is organized as follows. Section 2 reviews related works. Section 3 describes the cryptocurrency trading task. Section 4 proposes a practical approach to address the backtest overfitting issue. Section 5 presents performance evaluations. We conclude the paper in Section 6.

2 RELATED WORKS

Existing works can be classified into three categories: 1) backtest using the walk-forward method, 2) backtest using cross-validation methods, and 3) backtest with hyperparameter tuning.

2.1 BACKTEST USING WALK-FORWARD

The Walk-Forward (WF) method is the most widely applied backtest practice. WF trains a DRL agent over a training period and then evaluates its performance over a subsequent validation period. However, WF validates in one market situation, which can easily result in overfitting (Agarwal et al., 2021; Lin et al., 2021; Chan et al., 2019). WF may not be a good representative of future performance, as the validation period can be biased, e.g., a significant uptrend. Therefore, we want to train and validate under a variety of market situations in order to avoid overfitting and make the agent more robust.

2.2 BACKTEST USING CROSS-VALIDATION

Conventional methods (Jiang & Liang, 2017) used k -fold cross-validation (KCV) to backtest agents' trading performance. The KCV method partitions the dataset into k subsets, generating k folds. Then, for each trial, select one subset as the testing set and the rest $k - 1$ subsets as the training set. However, there are still risks of overfitting. First, a k -fold cross-validation method splits data by drawing from an IID process, which is a bold assumption for financial markets (Robinson & Sims, 1994). Second, the testing set generated by a cross-validation method could have a substantial bias (De Prado, 2018). Finally, informational leakage is possible because the training and testing sets are correlated (Farokhi & Kaafar, 2020). The researcher should exercise extreme caution to avoid leaking testing knowledge into the training set.

We will employ an improved approach. Existing research results ignore the backtest overfitting issue and give the false impression that the DRL-based trading strategy may be ready to deploy on markets (Varoquaux & Cheplygina, 2021; Bouthillier et al., 2021; Dodge et al., 2020). WF only tests a single market situation with high statistical uncertainty, and KCV takes a false IID assumption without control for leakage. Therefore, we suggest a combinatorial cross-validation method that tracks the *degree of overfitting* during the training process. The combinatorial CV method simulates a higher variety of market situations, reducing overfitting by averaging results. Also, our combinatorial CV method controls for leakage and allows the researchers to observe the degree of overfitting during the training process.

2.3 BACKTEST WITH HYPERPARAMETER TUNING

DRL agents are highly sensitive to hyperparameters, as can be seen from the implementations of Stable Baselines3 (Raffin et al., 2021), RLlib (Liang et al., 2018), Raylib (Liaw et al., 2018), UnityML (Juliani et al., 2018) and TensorForce (Kuhnle et al., 2017). The selection of hyperparameters takes a lot of time and strongly influences the learning result. Several cloud platforms provide hyperparameter tuning services.

In finance, FinRL-Podracr (Li et al., 2021) employed an evolutionary strategy for training a trading agent on a cloud platform, which ranks agents with different hyperparameters and then selects the one with the highest return as the result agent. Essentially, this method carried out a hyperparameter tuning process.

3 CRYPTOCURRENCY TRADING USING DEEP REINFORCEMENT LEARNING

First, we model a cryptocurrency trading task as a Markov Decision Process (MDP). Then, we build a market environment using historical market data and describe the general setting for training a trading agent. Finally, we discuss the backtest overfitting issue.

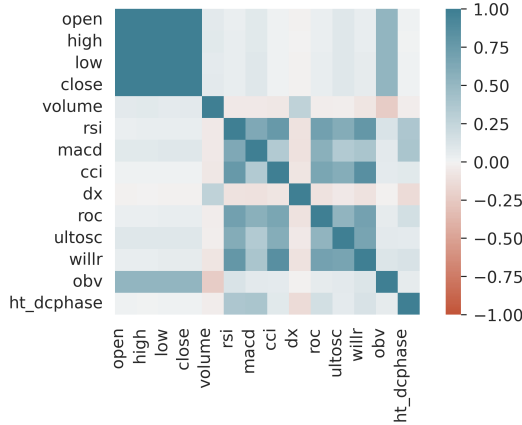


Figure 1: The correlation matrix of features.

3.1 MODELING CRYPTOCURRENCY TRADING

Assuming that there are D cryptocurrencies and T time stamps with $t = 0, 1, \dots, T - 1$. We use a deep reinforcement learning agent to make trading actions, which can be either *buy*, *sell*, or *hold*. An agent observes the market situation, e.g., prices and technical indicators, and takes actions to maximize the cumulative return. We model a trading task as a Markov Decision Process (MDP) (White, 1991) as follows

- State $\mathbf{s}_t = [b_t, \mathbf{h}_t, \mathbf{p}_t, \mathbf{f}_t] \in \mathbb{R}^{1+(I+2)D}$, where $b_t \in \mathbb{R}_+$ is the cash amount in the account, $\mathbf{h}_t \in \mathbb{R}_+^D$ denotes the share holdings, $\mathbf{p}_t \in \mathbb{R}_+^D$ denotes the prices at time t , $\mathbf{f}_t \in \mathbb{R}^{ID}$ is a feature vector for D cryptocurrencies and each has I technical indicators, and \mathbb{R}_+ denotes non-negative real numbers.
- Action $\mathbf{a}_t \in \mathbb{R}^D$: an action changes the non-negative share holdings $\mathbf{h}_t \in \mathbb{R}_+^D$, i.e., $\mathbf{h}_{t+1} = \mathbf{h}_t + \mathbf{a}_t$, where positive actions increase \mathbf{h}_t , negative actions decrease \mathbf{h}_t , and zero actions keep \mathbf{h}_t unchanged.
- Reward $r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}) \in \mathbb{R}$ is defined as a return when taking action \mathbf{a}_t at state \mathbf{s}_t and arriving at a new state \mathbf{s}_{t+1} . Here, we set it as the change of portfolio value, i.e., $r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}) = v_{t+1} - v_t$, where $v_t = \mathbf{p}_t^T \mathbf{h}_t + b_t \in \mathbb{R}_+$.
- Policy $\pi(\mathbf{a}_t | \mathbf{s}_t)$ is the trading strategy, which is a probability distribution over actions at state \mathbf{s}_t .

We consider 15 features that are used by existing papers (Xiong et al., 2018; Zhang et al., 2020; Yang et al., 2020; Liu et al., 2021), e.g., open-high-low-close-volume (OHLCV) and 9 technical indicators. Over the training period (from 02/02/2022 to 04/30/2022, as shown in Fig. 3), we compute Pearson correlations of the features and obtain a correlation matrix in Fig. 1. We list the 9 technical indicators as follows:

- Relative Strength Index (RSI) measures price fluctuation.

- Moving Average Convergence Divergence (MACD) is a momentum indicator for moving averages.
- Commodity Channel Index (CCI) compares the current price to the average price over a period.
- Directional Index (DX) measures the trend strength by quantifying the amount of price movement.
- The rate of change (ROC) is the speed at which variable changes over a period (Gerritsen et al., 2020).
- Ultimate Oscillator (ULTSOC) measures the price momentum of an asset across multiple timeframes (Gudelek et al., 2017).
- Williams %R (WILLR) measures overbought and oversold levels (Ni & Yin, 2009).
- On Balance Volume (OBV) measures buying and selling pressure as a cumulative indicator that adds volume on up-days and subtracts volume on down-days (Gerritsen et al., 2020).
- The Hilbert Transform Dominant (HT) is used to generate in-phase and quadrature components of a detrended real-valued signal to analyze variations of the instantaneous phase and amplitude (Nava et al., 2016).

As shown in Fig. 1, if two features have a correlation exceeding $\pm 60\%$, we drop either one of the two. Finally, $I = 6$ uncorrelated features are kept in the feature vector $\mathbf{f}_t \in \mathbb{R}^{6D}$, which are trading volume, RSI, DX, ULTSOC, OBV, and HT. Since the OHLC prices are highly correlated, \mathbf{p}_t in \mathbf{s}_t is chosen to be the close price over the period $[t, t + 1]$. Note that the close price of period $[t, t + 1]$ equals to the open price of period $[t + 1, t + 2]$. The feature vector \mathbf{f}_t characterizes the market situation. For the case $D = 10$, \mathbf{s}_t has size 81.

3.2 BUILDING MARKET ENVIRONMENT

We build a market environment by replaying historical data, following the style of OpenAI Gym (Brockman et al., 2016). A trading agent interacts with the market environment in multiple episodes, where an episode replays the market data (time series) in a time-driven manner from $t = 0$ to $t = T - 1$. At the beginning ($t = 0$), the environment sends an initial state \mathbf{s}_0 to the agent that returns an action \mathbf{a}_0 . Then, the environment executes the action \mathbf{a}_t and sends a reward value r_t and a new state \mathbf{s}_{t+1} to the agent, for $t = 0, \dots, T - 1$. Finally, \mathbf{s}_{T-1} is set to be the terminal state.

The market environment has the following three functions:

- `reset` function resets the environment to $\mathbf{s}_0 = [b_0, \mathbf{h}_0, \mathbf{p}_0, \mathbf{f}_0]$ where b_0 is the investment capital and $\mathbf{h}_0 = \mathbf{0}$ (zero vector) since there are no share holdings yet.
- `step` function takes an action \mathbf{a}_t and updates state \mathbf{s}_t to \mathbf{s}_{t+1} . For \mathbf{s}_{t+1} at time $t + 1$, \mathbf{p}_{t+1} and \mathbf{f}_{t+1} are accessible by looking up the time series of market data, and update b_{t+1} and \mathbf{h}_{t+1} as follows:

$$\begin{aligned} b_{t+1} &= b_t - \mathbf{p}_t^T \mathbf{a}_t, \\ \mathbf{h}_{t+1} &= \mathbf{h}_t + \mathbf{a}_t. \end{aligned} \tag{1}$$

- `reward` function computes $r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}) = v_{t+1} - v_t$ as follows:

$$r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}) = (b_{t+1} + \mathbf{p}_{t+1}^T \mathbf{h}_{t+1}) - (b_t + \mathbf{p}_t^T \mathbf{h}_t). \tag{2}$$

Trading constraints

1). **Transaction fees.** Each trade has transaction fees, and different brokers charge varying commissions. For cryptocurrency trading, we assume that the transaction cost is 0.3% of the value of each trade. Therefore, (2) becomes

$$r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}) = (b_{t+1} + \mathbf{p}_{t+1}^T \mathbf{h}_{t+1}) - (b_t + \mathbf{p}_t^T \mathbf{h}_t) - c_t, \tag{3}$$

and the transactions fee c_t is

$$c_t = \mathbf{p}_t^T |\mathbf{a}_t| \times 0.3\%, \tag{4}$$

where $|\mathbf{a}_t|$ means taking entry-wise absolute value of \mathbf{a}_t .

2). **Non-negative balance.** We do not allow short, thus we make sure that $b_{t+1} \in \mathbb{R}_+$ is non-negative,

$$b_{t+1} = b_t + \mathbf{p}_t^T \mathbf{a}_t^S + \mathbf{p}_t^T \mathbf{a}_t^B \geq 0, \quad \text{for } t = 0, \dots, T - 1, \quad (5)$$

where $\mathbf{a}_t^S \in \mathbb{R}_-^D$ and $\mathbf{a}_t^B \in \mathbb{R}_+^D$ denote the selling orders and buying orders, respectively, such that $\mathbf{a}_t = \mathbf{a}_t^S + \mathbf{a}_t^B$. Therefore, action \mathbf{a}_t is executed as follows: first execute the selling orders $\mathbf{a}_t^S \in \mathbb{R}_-^D$ and then the buying orders $\mathbf{a}_t^B \in \mathbb{R}_+^D$; and if there is not enough cash, a buying order will not be executed.

3). **Risk control.** The cryptocurrency market regularly drops in terms of market capitalization, sometimes even $\geq 70\%$. To control the risk for these market situations, we employ the Cryptocurrency Volatility Index, CVIX (Bonaparte, 2021a). Extreme market situations increase the value of CVIX. Once the CVIX exceeds a certain threshold, we stop buying and sell all our cryptocurrency holdings. We resume trading once the CVIX returns under the threshold.

3.3 TRAINING A TRADING AGENT

A trading agent learns a policy $\pi(\mathbf{a}_t | \mathbf{s}_t)$ that maximizes the discounted cumulative return $R = \sum_{t=0}^{\infty} \gamma^t r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1})$, where $\gamma \in (0, 1]$ is a discount factor and $r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1})$ is given in (3). The Bellman equation gives the optimality condition for an MDP problem, which takes a recursive form as follows:

$$Q^\pi(\mathbf{s}_t, \mathbf{a}_t) = \mathbb{E}_{\mathbf{s}_{t+1}} [r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1})] + \gamma \mathbb{E}_{\mathbf{a}_{t+1}} [Q^\pi(\mathbf{s}_{t+1}, \mathbf{a}_{t+1})]. \quad (6)$$

There are tens of DRL algorithms that can be adapted to crypto trading. Popular ones are TD3 (Fujimoto et al., 2018), SAC (Haarnoja et al., 2018), and PPO (Schulman et al., 2017).

Next, we describe a general flow of agent trading. At the beginning of training, we set hyperparameters such as the learning rate, batch size, etc. DRL algorithms are highly sensitive to hyperparameters, meaning that an agent’s trading performance may vary significantly. We have multiple sets of hyperparameters in the training stage for different trials. Each trial trains with one set of hyperparameters and obtains a trained agent. Then, we pick the DRL agent with the best performing hyperparameters and re-train the agent on the whole training data.

3.4 BACKTEST OVERFITTING ISSUE

Backtest (De Prado, 2018) uses historical data to simulate the market and evaluates the performance of an agent, namely, how would an agent have performed should it have been run over a past time period. Researchers often perform backtests by splitting the data into two chronological sets: one training set and one validation set. However, a DRL agent usually overfits an individual validation set that represents one market situation, thus, the actual trading performance is in question.

Backtest overfitting occurs when a DRL agent fits the historical training data to a harmful extent. The DRL agent adapts to random fluctuations in the training data, learning these fluctuations as concepts. However, these concepts do not exist, damaging the performance of the DRL agent on unseen states.

4 PRACTICAL APPROACH TO ADDRESS BACKTEST OVERFITTING

We propose a practical approach to address the backtest overfitting issue (Bailey et al., 2016). First, we formulate the problem as a hypothesis test and reject agents that do not pass the test. Then, we describe the detailed steps to estimate the probability of overfitting, $p \in [0, 1]$.

4.1 HYPOTHESIS TEST TO REJECT OVERFITTED AGENTS

We formulate a hypothesis test to reject overfitted agents. Mathematically, it is expressed as follows

$$\begin{cases} \mathcal{H}_0 : & p < \alpha, & \text{NOT overfitted,} \\ \mathcal{H}_1 : & p \geq \alpha, & \text{overfitted.} \end{cases} \quad (7)$$

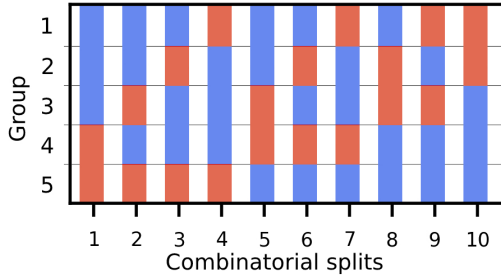


Figure 2: Illustration of combinatorial splits. Split the data into $N = 5$ groups, $k = 2$ groups for the train set (red) and the rest $N - k = 3$ groups for the second set (blue).

where $\alpha > 0$ is the level of significance.

The hypothesis test (7) is expected to reject two types of *false-positive* DRL agents. 1) Existing methods may have reported overoptimistic results since many authors were tempted to go back and forth between training and testing periods. This type of information leakage is a common reason for backtest overfitting. 2). Agent training is likely to overfit since DRL algorithms are highly sensitive to hyperparameters (Clary et al., 2019; Henderson et al., 2018; Mania et al., 2018). For example, one can train agents with PPO, TD3, or SAC algorithm and then reject agents that do not pass the test. We set the level of significance α according to the Neyman-Pearson framework (Neyman & Pearson, 1933).

4.2 ESTIMATING PROBABILITY OF OVERFITTING

We estimate the probability of overfitting using the cumulative return vector. The cumulative return vector is defined as $\mathbf{R} = (v_t - v_0)/v_0$, where v_t is the portfolio value at time step t and v_0 is the original capital. We estimate the probability of overfitting via three general steps.

- **Step 1:** For each hyperparameter trial, average the cumulative returns on the validation sets (of length T') and obtain $\mathbf{R}_{\text{avg}} \in \mathbb{R}^{T'}$.
- **Step 2:** For H trials, stack \mathbf{R}_{avg} into a matrix $\mathbf{M} \in \mathbb{R}^{T' \times H}$.
- **Step 3:** Based on \mathbf{M} , we compute the probability of overfitting p .

Cross-validation (Browne, 2000) allows training-and-validating on different market situations. Given a training time period, we perform the following steps:

- **Step 1** (Training-validation data splits): as shown in Fig. 2, divide the training period with T data points into N groups of equal size, k out of N groups construct an validation set and the rest $N - k$ groups as an training set, resulting in $J = \binom{N}{N - k}$ combinatorial splits. The training and validation sets have $(N - k)(T/N)$ and $T' = k(T/N)$ data points, respectively.
- **Step 2** (One trial of hyperparameters): set a new set of parameters for hyperparameter tuning.
- **Step 3:** In each training-validation data split, we train an agent using the training set and then evaluate the agent’s performance metric for each validation set, $i = 1, \dots, H$. After training on all splits, we take the mean performance metric over all validation sets.

Step 2) and Step 3) constitute *one hyperparameter trial*. Loop for H trials and select the set of hyperparameters (or DRL agent) that performs the best in terms of mean performance metric over all splits. This procedure considers various market situations, resulting in the best performing DRL agent over different market situations. However, a training process involving multiple trials will result in overfitting (Bailey et al., 2014). We want to measure the probability of overfitting.

Consider a probability space $(\mathcal{T}, \mathcal{F}, \mathbb{P})$, where \mathcal{T} represents the sample space, \mathcal{F} the event space and \mathbb{P} the probability space. A sample $c \in \mathcal{T}$ is a **split of matrix M across rows**. For instance, we split M into four subsets:

$$M^1, M^2, M^3, M^4 \in \mathbb{R}^{T'/4 \times H}. \quad (8)$$

A sample c could be any split of M , such as:

$$c = \begin{bmatrix} M^1 \\ M^2 \end{bmatrix}, \text{ IS set} \quad \bar{c} = \begin{bmatrix} M^3 \\ M^4 \end{bmatrix}, \text{ OOS set} \quad (9)$$

An agent is overfitted if its best performance in the IS set has an expected ranking that is lower than the median ranking in the OOS set (Bailey et al., 2016). For a sample $c \in \mathcal{T}$, let $\mathbf{R}^c \in \mathbb{R}^H$ and $\bar{\mathbf{R}}^c \in \mathbb{R}^H$ denote the IS and OOS performance of the columns of a sample c , respectively. We rank \mathbf{R}^c and $\bar{\mathbf{R}}^c$ from low values to high values, resulting in r^c and \bar{r}^c (for example, [1, 3, 2], where 3 corresponds to the best performance metric with index 2). Define ϵ as the index of the best performing IS strategy. Then, we check the corresponding OOS rank $\bar{r}^c[\epsilon]$ and define a relative rank ω^c as follows:

$$\omega^c = \frac{\bar{r}^c[\epsilon]}{H+1}, \text{ for } c \in \mathcal{T}. \quad (10)$$

Define a logit function (Barker, 2005) as follows:

$$\lambda^c = \ln \frac{\omega^c}{1-\omega^c}, \text{ for } c \in \mathcal{T}. \quad (11)$$

If $\omega^c < 0.5$, we have $\lambda^c < 0$, meaning that the best strategy IS has an expected ranking lower than the OOS set, which is overfitting. High logit values indicate coherence between IS and OOS performance, indicating a low level of overfitting. Finally, the probability of overfitting is computed as follows (Bailey et al., 2016):

$$p = \int_{-\infty}^0 f(\lambda) d\lambda, \quad (12)$$

where $f(\lambda)$ denotes the distribution function of λ .

Finally, we discuss how to set the level of significance α . In our problem, the percentage of selected IS strategies with an expected ranking lower than the median ranking of OOS determines the probability of overfitting. The relative OOS rank is a uniform distribution. According to (Heyde, 1963), such a uniform's logit is normally distributed. In agreement with a standard application of the Neyman-Pearson framework (Neyman & Pearson, 1933), we can set the level of significance α , which is the probability of a type I error. Because α is a probability, it ranges between 0 and 1. Thus, if an investigator selects $\alpha = 10\%$, they are allowing a 10% probability of incorrectly rejecting the null hypothesis in favor of the alternative when the null hypothesis is true.

5 PERFORMANCE EVALUATIONS

First, we describe the experimental settings and the compared methods and metrics. Then, we show that the proposed method can help reject two types of overfitted agents. Finally, we present the backtest performance to verify that our hypothesis test helps increase the chance of good trading performance.

5.1 EXPERIMENTAL SETTINGS

We select 10 cryptocurrencies with high trading volumes: AAVE, AVAX, BTC, NEAR, LINK, ETH, LTC, MATIC, UNI, and SOL. We assume a trade can be executed at the market price and ignore the slippage issue because a higher trading volume indicates higher market liquidity.

Data split: We use five-minute-level data from 02/02/2022 to 06/27/2022. We split it into a training period (from 02/02/2022 to 04/30/2022) and a testing period (from 05/01/2022 to 06/27/2022, during which the crypto market **crashed two times**), as shown in Fig. 3. The training period splits further into IS-OOS sets for estimating p as in Section 4.2.




Figure 3: Data split: train and validate an agent in the blue period, and test its performance in the green period.

Table 1: The hyperparameters and their values.

Hyperparameter	Description	Range
Learning rate	Step size during the training process	$[3e^{-2}, 2.3e^{-2}, 1.5e^{-2}, 7.5e^{-3}, 5e^{-6}]$
Batch size	Number of training samples in one iteration	$[512, 1280, 2048, 3080]$
Gamma γ	Discount factor	$[0.95, 0.96, 0.97, 0.98, 0.99]$
Net dimension	Width of hidden layers of the actor network	$[2^9, 2^{10}, 2^{11}]$
Target step	Explored target step number in the environment	$[2.5e^3, 3.75e^3, 5e^3]$
Break step	Total timesteps performed during training	$[3e^4, 4.5e^4, 6e^4]$

Training with combinatorial cross-validation: there are in total $T = 25055 = 87$ (days) \times 24 (hours) \times 60/5 (minutes) $- 1$ datapoints in the training time period, and $T' = 16704 = 58$ (days) \times 24 (hours) \times 60/5 (minutes) in the testing time period. We divide the training data set into $N = 5$ equal-sized subsets, each subset has 5011 datapoints. We perform the combinatorial cross-validation method with $N = 5$ and $k = 2$, where each data split has $k = 2$ validation sets and $N - k = 3$ training sets. Therefore, the total number of training-validation splits is $J = 10$.

Hyperparameters: we list six tunable hyperparameters in Table 1. Their values are based on the implementations of Stable Baselines3 (Raffin et al., 2021), RLlib (Liang et al., 2018), Ray Tune (Liaw et al., 2018), UnityML (Juliani et al., 2018) and TensorForce (Kuhnle et al., 2017). There are in total $2700 = 5 \times 4 \times 5 \times 3 \times 3 \times 3$ combinations.

Trials H : for any distribution over a sample space with a finite maximum, the maximum of 50 random observations lies within the top 5% of the actual maximum, with 90% probability. Specifically, $1 - (1 - 0.05)^H > 0.9$ requires $H > 50$ assuming the optimal region of hyperparameters occupies at least 5% of the grid space.

Volatility index CVIX (Crypto VIX) (Bonaparte, 2021b): the total market capitalization of the crypto market **crashed two times** in our testing period, namely, 05/06/2022 - 05/12/2022 and 06/09/2022 - 06/15/2022. We take the average CVIX value over those time frames as our threshold, $CVIX_t = 90.1$.

Level of significance α : we allow a 10% probability of incorrectly rejecting the null hypothesis in favour of the alternative when the null hypothesis is true (type I error), $\alpha = 10\%$.

5.2 COMPARED METHODS AND METRICS

We consider two cases where the proposed method helps reject overfitted agents.

5.2.1 CONVENTIONAL DEEP REINFORCEMENT LEARNING AGENTS

First, the walk-forward method (Park et al., 2020; Moody & Saffell, 2001; Deng et al., 2016; Yang et al., 2020; Li et al., 2019) applies a training-validation-testing data split. On the same training-validation set, we train with $H = 50$ different sets of hyperparameters, all with PPO algorithm (Schulman et al., 2017) and then calculate p (Bailey et al., 2014). Second, we train another conventional agent using the PPO algorithm and the k -fold cross-validation (KCV) method with $k = 5$.

5.2.2 DEEP REINFORCEMENT ALGORITHMS WITH DIFFERENT HYPERPARAMETERS

DRL algorithms are highly sensitive to hyperparameters, resulting in high variability of trading performance (Clary et al., 2019; Henderson et al., 2018; Mania et al., 2018). We use the probability of overfitting to measure the likelihood that an agent is overfitted. We tune the hyperparameters in

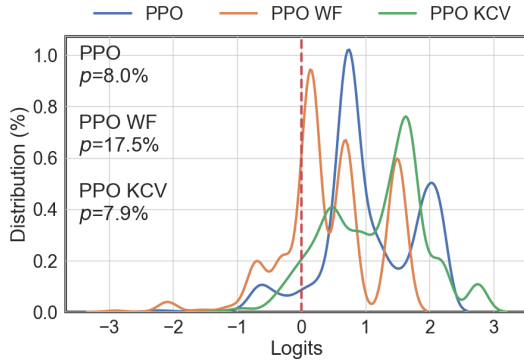


Figure 4: Logit distribution $f(\lambda)$ of three conventional agents.

Table 1 for three agents, TD3 (Fujimoto et al., 2018), SAC (Haarnoja et al., 2018) and PPO, and calculate p for each agent with each set of hyperparameters for $H = 50$ trials.

5.2.3 PERFORMANCE METRICS

We use three metrics to measure an agent’s performance:

- **Cumulative return** $R = \frac{v-v_0}{v_0}$, where v is the final portfolio value, and v_0 is the original capital.
- **Volatility** $V = \text{std}(R_t)$, where $R_t = \frac{v_t-v_{t-1}}{v_{t-1}}$, and $t = 0, \dots, T - 1$.
- **Probability of overfitting** p . We split M into 14 submatrices, analyze all possible combinations, and set the threshold $\alpha = 10\%$ (Neyman & Pearson, 1933).

The cumulative return measures the profits. The widely used volatility measures the degree of variation of a trading price series over time; the amount of risk.

5.2.4 BENCHMARKS

We compare with two benchmark methods: an equal-weight portfolio and the S&P Cryptocurrency Broad Digital Market Index (S&P BDM Index).

- Equal-weight strategy: at time t_0 , distribute the available cash b_0 equally over all D available cryptocurrencies.
- S&P BDM index (Indices & Methodology, 2022): the S&P broad digital market index S&P tracks the performance of cryptocurrencies with a market capitalization greater than \$10 million.

5.3 REJECT CONVENTIONAL AGENTS

Fig. 4 shows the logit distribution function $f(\lambda)$ for conventional agents described in Section 5.2.1. The area under $f(\lambda)$ for the domain $[-\infty, 0]$ is the probability of overfitting p . The peaks represent the outperforming DRL agents, as an agent’s success is sensitive to the hyperparameter set. These outperforming trials deliver the same relative rank more often, and the logits are a function of the relative rank. We compare our PPO approach to conventional PPO WF and PPO KCV. The WF and KCV methods have $p_{WF} = 17.5\% > \alpha$ and $p_{KCV} = 7.9\% < \alpha$, respectively. For the WF method, we accept the alternative hypothesis \mathcal{H}_1 and conclude that it is overfitting.

5.4 REJECT OVERFITTED AGENTS

Table 2 presents the hyperparameters selected for each agent. Fig. 5 shows the logit density function for the DRL agents in Section 5.2.2. The probabilities of overfitting are: $p_{PPO} = 8.0\% < \alpha$,


Figure 5: Logit distribution $f(\lambda)$ of three DRL agents.

Table 2: Selected hyperparameters for each DRL algorithm.

Hyperparameter	PPO	TD3	SAC
Learning rate	7.5e-3	3e-2	1.5e-2
Batch size	512	3080	3080
Gamma	0.95	0.95	0.97
Net dimension	1024	2048	1024
Target step	5e4	2.5e3	2.5e3
Break step	4.5e4	6e4	4.5e4

$p_{TD3} = 9.6\% < \alpha$ and $p_{SAC} = 21.3\% > \alpha$. We accept alternative hypothesis \mathcal{H}_1 and conclude that SAC is overfitted. Finally, TD3 has a dominant part of the logit distribution at a high logit domain (≈ 4). High logit values indicate coherence between IS and OOS performance.

5.5 BACKTEST PERFORMANCE

Fig. 6 and Fig. 7 show the backtest results. When the CVIX indicator surpasses 90, the agent stops buying and sells all cryptocurrency holdings. Fig. 6 and Table 3 compare conventional agents, market benchmarks and our approach. Compared to PPO WF and PPO KCV, our method outperforms the other two agents with at least a 15% increase in the cumulative return. The lower volatility of PPO indicates that our method is more robust to risk. Fig. 7 and Table 3 show the backtest results of the DRL agents. The cumulative return of the PPO agent is significantly better ($\approx 24\%$) than those of agents TD3 and SAC. Also, in terms of volatility, the PPO agent is superior.

Finally, compared to the benchmarks, the performance of our approach is more excellent in terms of cumulative return and volatility. From our method and experiments, we conclude that the superior agent is PPO.

6 CONCLUSION

In this paper, we have shown the importance of addressing the backtesting overfitting issue in cryptocurrency trading with deep reinforcement learning. Results show that the least overfitting agent PPO (with combinatorial CV method) outperforms the conventional agents (WF and KCV methods), two other DRL agents (TD3 and SAC), and the S&P DBM Index in cumulative return and volatility, showing good robustness.

Table 3: Performance comparison.

Metrics \times Method	S&P BDM Index	Equal-weight	PPO WF	PPO KCV	PPO	TD3	SAC
Cumulative return	-50.78%	-47.78%	-49.39%	-55.54%	-34.96%	-59.08%	-59.48%
Volatility	5.81e-2	4.19e-3	2.79e-3	3.49e-3	2.01e-3	3.51e-3	3.78e-3
Prob. of overfitting p	-	-	17.5%	7.9%	8.0%	9.6%	21.3%

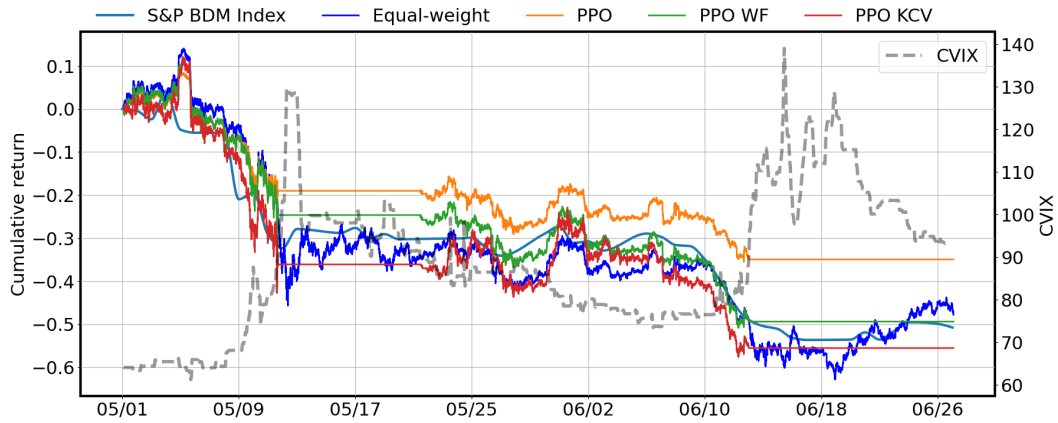


Figure 6: The average trading cumulative return curves for conventional agents. From 05/01/2022 to 06/27/2022, the initial capital is \$1,000,000.




Figure 7: The average trading cumulative return curves for DRL algorithms. From 05/01/2022 to 06/27/2022, the initial capital is \$1,000,000.

Future work will be interesting to 1). explore the evolution of the probability of overfitting during training and for different agents; 2). test limit order setting and trade closure; 3). explore large-scale data, such as all currencies corresponding to the S&P BDM index; and 4). consider more features for the state space, including fundamentals and sentiment features.

REFERENCES

- Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Deep reinforcement learning at the edge of the statistical precipice. *Advances in Neural Information Processing Systems*, 34:29304–29320, 2021.
- Andrew Ang, Tom Morris, et al. Asset allocation with crypto: Application of preferences for positive skewness. 2022.
- David H Bailey, Jonathan M Borwein, Marcos López de Prado, and Qiji Jim Zhu. Pseudomathematics and financial charlatanism: The effects of backtest over fitting on out-of-sample performance. *Notices of the AMS*, 61(5):458–471, 2014.
- David H Bailey, Jonathan Borwein, Marcos Lopez de Prado, and Qiji Jim Zhu. The probability of backtest overfitting. *Journal of Computational Finance*, forthcoming, 2016.
- Lawrence E Barker. Logit models from economics and other fields, 2005.
- Yosef Bonaparte. Introducing the Cryptocurrency VIX: CVIX. *SSRN Electronic Journal*, pp. 1–9, 2021a. doi: 10.2139/ssrn.3948898.
- Yosef Bonaparte. Introducing the cryptocurrency vix: Cvix. Available at SSRN 3948898, 2021b.
- Xavier Bouthillier, Pierre Delaunay, Mirko Bronzi, Assya Trofimov, Brennan Nichyporuk, Justin Szeto, Nazanin Mohammadi Sepahvand, Edward Raff, Kanika Madan, Vikram Voleti, et al. Accounting for variance in machine learning benchmarks. *Proceedings of Machine Learning and Systems*, 3:747–769, 2021.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. pp. 1–4, 2016.
- Michael W Browne. Cross-validation methods. *Journal of Mathematical Psychology*. *Journal of Mathematical Psychology*, 44:108–132, 2000.
- Stephanie CY Chan, Samuel Fishman, John Canny, Anoop Korattikara, and Sergio Guadarrama. Measuring the reliability of reinforcement learning algorithms. *arXiv preprint arXiv:1912.05663*, 2019.
- Kaleigh Clary, Emma Tosch, John Foley, and David Jensen. Let’s play again: variability of deep reinforcement learning agents in atari environments. *arXiv preprint arXiv:1904.06312*, 2019.
- Christian Conrad, Anessa Custovic, and Eric Ghysels. Long- and short-term Cryptocurrency volatility components: A GARCH-MIDAS Analysis. *Journal of Risk and Financial Management*, 11(2):23, 2018.
- Marcos Lopez De Prado. *Advances in financial machine learning*. John Wiley & Sons, 2018.
- Yue Deng, Feng Bao, Youyong Kong, Zhiqian Ren, and Qionghai Dai. Deep direct reinforcement learning for financial signal representation and trading. *IEEE Transactions on Neural Networks and Learning Systems*, 28(3):653–664, 2016.
- Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah Smith. Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping. *arXiv preprint arXiv:2002.06305*, 2020.
- Fan Fang, Carmine Ventre, Michail Basios, Leslie Kanthan, David Martinez-Rego, Fan Wu, and Lingbo Li. Cryptocurrency trading: a comprehensive survey. *Financial Innovation*, 8(1):1–59, 2022.

- Farhad Farokhi and Mohamed Ali Kaafar. Modelling and quantifying membership information leakage in machine learning. *arXiv preprint arXiv:2001.10648*, 2020.
- Scott Fujimoto, Herke Van Hoof, and David Meger. Addressing Function Approximation Error in Actor-Critic Methods. *35th International Conference on Machine Learning, ICML 2018*, 4: 2587–2601, 2018.
- Dirk F Gerritsen, Elie Bouri, Ehsan Ramezanifar, and David Roubaud. The profitability of technical trading rules in the bitcoin market. *Finance Research Letters*, 34:101263, 2020.
- M Ugur Gudelek, S Arda Boluk, and A Murat Ozbayoglu. A deep learning based stock trading model with 2-d cnn trend detection. In *IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1–8. IEEE, 2017.
- Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- Ben Hambly, Renyuan Xu, and Huining Yang. Recent advances in reinforcement learning in finance. *arXiv preprint arXiv:2112.04553*, 2021.
- Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pp. 3207–3214, 2018.
- Chris C Heyde. On a property of the lognormal distribution. *Journal of the Royal Statistical Society: Series B (Methodological)*, 25(2):392–393, 1963.
- Dow Jones Indices and Index Methodology. S&P Digital Market Indices Methodology. (January), 2022.
- Zhengyao Jiang and Jinjun Liang. Cryptocurrency portfolio management with deep reinforcement learning. In *Intelligent Systems Conference (IntelliSys)*, pp. 905–913. IEEE, 2017.
- Arthur Juliani, Vincent-Pierre Berges, Ervin Teng, Andrew Cohen, Jonathan Harper, Chris Elion, Chris Goy, Yuan Gao, Hunter Henry, Marwan Mattar, et al. Unity: A general platform for intelligent agents. *arXiv preprint arXiv:1809.02627*, 2018.
- Alexander Kuhnle, Michael Schaarschmidt, and Kai Fricke. Tensorforce: a tensorflow library for applied reinforcement learning. Web page, 2017. URL <https://github.com/tensorforce/tensorforce>.
- Xinyi Li, Yinchuan Li, Yuancheng Zhan, and Xiao-Yang Liu. Optimistic bull or pessimistic bear: Adaptive deep reinforcement learning for stock portfolio allocation. *arXiv preprint arXiv:1907.01503*, 2019.
- Zechu Li, Xiao-Yang Liu, Jiahao Zheng, Zhaoran Wang, Anwar Walid, and Jian Guo. Finrl-podracar: High performance and scalable deep reinforcement learning for quantitative finance. In *Proceedings of the Second ACM International Conference on AI in Finance*, pp. 1–9, 2021.
- Eric Liang, Richard Liaw, Robert Nishihara, Philipp Moritz, Roy Fox, Ken Goldberg, Joseph Gonzalez, Michael Jordan, and Ion Stoica. RLlib: Abstractions for distributed reinforcement learning. In *International Conference on Machine Learning*, pp. 3053–3062. PMLR, 2018.
- Richard Liaw, Eric Liang, Robert Nishihara, Philipp Moritz, Joseph E Gonzalez, and Ion Stoica. Tune: A research platform for distributed model selection and training. *arXiv preprint arXiv:1807.05118*, 2018.
- Jimmy Lin, Daniel Campos, Nick Craswell, Bhaskar Mitra, and Emine Yilmaz. Significant improvements over the state of the art? a case study of the ms marco document ranking leaderboard. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 2283–2287, 2021.

- Xiao-Yang Liu, Hongyang Yang, Jiechao Gao, and Christina Dan Wang. Finrl: Deep reinforcement learning framework to automate trading in quantitative finance. In *Proceedings of the Second ACM International Conference on AI in Finance*, pp. 1–9, 2021.
- Yukun Liu and Aleh Tsyvinski. Risks and returns of cryptocurrency. *The Review of Financial Studies*, 34(6):2689–2727, 2021.
- Horia Mania, Aurelia Guy, and Benjamin Recht. Simple random search provides a competitive approach to reinforcement learning. *arXiv preprint arXiv:1803.07055*, 2018.
- John Moody and Matthew Saffell. Learning to trade via direct reinforcement. *IEEE Transactions on Neural Networks*, 12(4):875–889, 2001.
- Alan Moreira and Tyler Muir. Volatility-managed portfolios. *The Journal of Finance*, 72(4):1611–1644, 2017.
- Noemi Nava, Tiziana Di Matteo, and Tomaso Aste. Time-dependent scaling patterns in high frequency financial data. *The European Physical Journal Special Topics*, 225(10):1997–2016, 2016.
- Jerzy Neyman and Egon S Pearson. The testing of statistical hypotheses in relation to probabilities a priori. In *Mathematical proceedings of the Cambridge philosophical society*, volume 29, pp. 492–510. Cambridge University Press, 1933.
- He Ni and Hujun Yin. Exchange rate prediction using hybrid neural networks and trading indicators. *Neurocomputing*, 72(13-15):2815–2823, 2009.
- Hyungjun Park, Min Kyu Sim, and Dong Gu Choi. An intelligent financial portfolio trading strategy using deep q-learning. *Expert Systems with Applications*, 158:113573, 2020.
- Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dornmann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 2021.
- PM Robinson and CA Sims. Time series with strong dependence. In *Advances in Econometrics, sixth World Congress*, volume 1, pp. 47–95, 1994.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Gaël Varoquaux and Veronika Cheplygina. How i failed machine learning in medical imaging—shortcomings and recommendations. *arXiv preprint arXiv:2103.10292*, 2021.
- Chelsea C White. A survey of solution techniques for the partially observed markov decision process. *Annals of Operations Research*, 32(1):215–230, 1991.
- Zhuoran Xiong, Xiao-Yang Liu, Shan Zhong, Hongyang Yang, and Anwar Walid. Practical deep reinforcement learning approach for stock trading. *NeurIPS Workshop on Challenges and Opportunities for AI in Financial Services*, 2018.
- Hongyang Yang, Xiao-Yang Liu, Shan Zhong, and Anwar Walid. Deep reinforcement learning for automated stock trading: An ensemble strategy. In *Proceedings of the First ACM International Conference on AI in Finance*, pp. 1–8, 2020.
- Zihao Zhang, Stefan Zohren, and Stephen Roberts. Deep reinforcement learning for trading. *The Journal of Financial Data Science*, 2(2):25–40, 2020.

A APPENDIX

You may include other additional sections here.