CLAIMS

WHAT IS CLAIMED IS:

1. A method of portfolio management, the method comprising:

retrieving most recent transaction data of one or more financial markets;

forecasting, using an ensemble of prediction models , one or more prediction features based on the most recent transaction data;

outputting, by a reinforcement learning based intelligent agent, an action that specifies a portfolio weight distribution over a plurality of financial investment products based on the forecasted one or more prediction features and the most recent transaction data, the plurality of financial investment products being traded on the one or more financial markets; and

initiating, completely autonomously, a set of transactions on the plurality of financial investment products to achieve the portfolio weight distribution in response to the outputting of the action that specifies the portfolio weight distribution by the reinforcement learning based intelligent agent.

2. The method of claim 1, wherein at least one prediction model from the ensemble of prediction models comprises a temporal convolutional neural network.

3. The method of claim 1, wherein at least one prediction model from the ensemble of prediction models is trained using historical transaction data of the one or more financial markets.

4. The method of claim 1, wherein the one or more prediction features comprise at least one of one or more price related features or one or more sentiment related features.

5. The method of claim 1, wherein the reinforcement learning based intelligent agent comprises a set of feature extraction networks, a set of feature analysis networks, and a set of network improvement networks.

6. The method of claim 5, wherein outputs from at least one prediction model from the ensemble of prediction models are fed directly to at least one of the set of feature extraction networks or the set of feature analysis networks.

7. The method of claim 1, wherein the reinforcement learning based intelligent agent is trained using historical transaction data of the one or more financial markets and historical outputs from at least one prediction model from the ensemble of prediction models.

8. The method of claim 1, wherein the reinforcement learning based intelligent agent comprises a replay buffer for training of the reinforcement learning based intelligent agent, wherein data stored in the replay buffer is sampled via prioritized replay based on temporal difference errors.

9. An apparatus for portfolio management, comprising:

a memory; and

at least one processor coupled to the memory and configured to:

retrieve most recent transaction data of one or more financial markets;

forecast, using an ensemble of prediction models , one or more prediction features based on the most recent transaction data;

output, by a reinforcement learning based intelligent agent, an action that specifies a portfolio weight distribution over a plurality of financial investment products based on the forecasted one or more prediction features and the most recent transaction data, the plurality of financial investment products being traded on the one or more financial markets; and

initiate, completely autonomously, a set of transactions on the plurality of financial investment products to achieve the portfolio weight distribution in response to the outputting of the action that specifies the portfolio weight distribution by the reinforcement learning based intelligent agent.

10. The apparatus of claim 9, wherein at least one prediction model from the ensemble of prediction models comprises a temporal convolutional neural network.

11. The apparatus of claim 9, wherein the one or more prediction features comprise at least one of one or more price related features or one or more sentiment related features.

12. The apparatus of claim 9, wherein the reinforcement learning based intelligent agent comprises a set of feature extraction networks, a set of feature analysis networks, and a set of network improvement networks.

13. The apparatus of claim 12, wherein outputs from at least one prediction model from the ensemble of prediction models are fed directly to at least one of the set of feature extraction networks or the set of feature analysis networks.

14. The apparatus of claim 9, wherein the reinforcement learning based intelligent agent comprises a replay buffer for training of the reinforcement learning based intelligent agent, wherein data stored in the replay buffer is sampled via prioritized replay based on temporal difference errors.

15. An apparatus for portfolio management, comprising:

means for retrieving most recent transaction data of one or more financial markets;

means for forecasting one or more prediction features based on the most recent transaction data;

means for outputting, based on the forecasted one or more prediction features and the most recent transaction data, an action that specifies a portfolio weight distribution over a plurality of financial investment products that are traded on the one or more financial markets; and

means for initiating, completely autonomously, a set of transactions on the plurality of financial investment products to achieve the portfolio weight distribution in response to the outputting of the action that specifies the portfolio weight distribution.

16. The apparatus of claim 15, wherein the means for forecasting comprises a temporal convolutional neural network or an ensemble of such networks.

17. The apparatus of claim 15, wherein the one or more prediction features comprise at least one of one or more price related features or one or more sentiment related features.

18. The apparatus of claim 15, wherein the means for outputting the action that specifies the portfolio weight distribution comprises a set of feature extraction networks, a set of feature analysis networks, and a set of network improvement networks.

19. The apparatus of claim 18, wherein outputs from the means for forecasting are fed directly to at least one of the set of feature extraction networks or the set of feature analysis networks.

20. The apparatus of claim 15, wherein the means for outputting the action that specifies the portfolio weight distribution comprises a replay buffer for training of the means for outputting the action that specifies the portfolio weight distribution, wherein data stored in the replay buffer is sampled via prioritized replay based on temporal difference errors.