
Adversarial recovery of agent rewards from latent spaces of the limit order book

Jacobo Roa-Vicens
JP Morgan Chase & Co.*
jacobo.roavicens@jpmorgan.com

Yuanbo Wang
Twitter, Inc. †
yuanbow@twitter.com

Virgile Mison
JP Morgan Chase & Co.
virgile.mison@jpmchase.com

Yarin Gal
Department of Computer Science
University of Oxford, United Kingdom
yarin@cs.ox.ac.uk

Ricardo Silva
Department of Statistical Science
University College London, United Kingdom
ricardo.silva@ucl.ac.uk

Abstract

Inverse reinforcement learning has proved its ability to explain state-action trajectories of expert agents by recovering their underlying reward functions in increasingly challenging environments. Recent advances in adversarial learning have allowed extending inverse RL to applications with non-stationary environment dynamics unknown to the agents, arbitrary structures of reward functions and improved handling of the ambiguities inherent to the ill-posed nature of inverse RL. This is particularly relevant in real time applications on stochastic environments involving risk, like volatile financial markets. Moreover, recent work on simulation of complex environments enable learning algorithms to engage with real market data through simulations of its latent space representations, avoiding a costly exploration of the original environment. In this paper, we explore whether adversarial inverse RL algorithms can be adapted and trained within such latent space simulations from real market data, while maintaining their ability to recover agent rewards robust to variations in the underlying dynamics, and transfer them to new regimes of the original environment.

1 Introduction

Reinforcement learning (RL) achieves robust performance in a wide number of fields, with particularly relevant success in model-free applications [1, 2] where agents explore an environment with no prior knowledge about its underlying dynamics, and learn a policy that maximizes certain cumulative reward function. Such learning process typically requires recurrent access of the agent to the environment on a trial-and-error based exploration; however, reinforcement learning in risk-critical tasks such as automatic navigation or financial risk control would not allow such an exploration, since decisions have to be made in real time in a non-stationary environment where the risks and

*Opinions expressed in this paper are those of the authors, and do not necessarily reflect the view of JP Morgan.

†Research completed at JP Morgan Chase & Co.

costs inherent to a trial-and-error approach can be unaffordable. In addition, such RL performance generally requires that the designer manually specifies a reward function that represents the task adequately and can be optimized efficiently [3].

In the context of learning from expert demonstrations, inverse reinforcement learning has proved capable of recovering through inference the reward function of expert agents through observations of their state-action trajectories [4, 5] with decreasing dependence on pre-defined assumptions about linearity or the general structure of the underlying reward function, generally under a maximum entropy framework [6].

Recent advances in inverse RL have extended its application to high-dimensional state spaces with unknown dynamics and arbitrary non-linear reward structures, thanks to the use of neural networks to represent the reward function [7]. Adversarial inverse reinforcement learning (AIRL) [8] extends inverse RL further, achieving the recovery of rewards robust to variations in the dynamics of the environment, while learning at the same time a policy to perform the task; AIRL builds on the equivalences found by [9] between inverse RL under maximum entropy and the training process of generative adversarial networks (GANs) [10].

1.1 Contributions

Financial markets are a particularly challenging case for inverse RL: high-dimensional, stochastic environments with non-stationary transition dynamics unknown to the observer, and with variable reactions to actions from agents. This makes AIRL particularly interesting to test on real financial data, aiming at learning from experts robust reward functions that can then be transferred to new regimes of the original environment.

In this paper we explore the adaptation of AIRL to a volatile financial environment based on real tick data from a limit order book (LOB) in the stock market, attempting to recover the rewards from three expert market agents through an observer with no prior knowledge of the underlying dynamics, where such dynamics can also change with time following real market data, and where the environment reacts to the agent’s actions. This is a challenging environment where we would expect AIRL to remain better suited to learn from the expert agent trajectories than methods aimed directly at recovering the policy (generative adversarial imitation learning [11], or GAIL), given the expected added value of AIRL regarding robustness of the recovered rewards with respect to varying dynamics of the environment. GAIL provides means analogous to generative adversarial networks that allow extraction of policies directly from data through a model-free approach for complex behaviours in high-dimensional environments. The performance of each method is then evaluated against the proportion of the expert agent’s total cumulative reward that can be obtained by policies recovered through adversarial learning.

1.2 Model environment based on real data

The adversarial learning algorithms used in the experiment will require a model of the environment where the observed agent trajectories took place, in order to evaluate the iterative estimations of rewards and policies most likely to have generated the observations. In practice, we would observe expert trajectories from agents as training data for adversarial learning, and then transfer the learnt policies to new test market data from the real environment.

The latent space model of the equity LOB described in [12] shows how the Mixture-Density Recurrent Network (RNN-MDN) architecture presented in World Models [13] to learn the transition dynamics of the environment can also be applied to financial time series, learning a causal representation of LOB market data where RL experts can be trained to learn policies then transferable back to the original environment. Building on this work, we train three expert traders in the latent space market model through advantage actor critic (A2C) [14], double DQN [2], and Policy Gradient [15] respectively, whose learnt policies remain profitable when tested on subsequent time series out of sample. The collections of expert state-action trajectories generated by each trained agent serve as input for the AIRL and GAIL algorithms in our experiments, following the implementation by [8]. Our conclusions will then examine the proportion of the experts’ cumulative rewards produced by the policies learnt through either AIRL or GAIL from each expert agent.

2 Background

2.1 Related work

A number of previous works have applied inverse RL to financial data, focusing on evaluations of feature vectors for state representations at different scales to explore a market of competing agents [16], and assuming linear structures for the reward functions. Other works have focused on assessing the comparative performance of probabilistic inverse RL based on Bayesian neural networks [17] as an alternative to inverse RL based on Gaussian processes [5], applied to a simulated financial market under a maximum entropy framework and allowing non-linear structures in the reward functions of the agents.

On the other hand, model-based approaches have attempted to recover specific parameters such as risk aversion implied by data [18], where the observer assumes a certain structure for the underlying utility of the agent. Research with simulations of real environments through neural networks [19] allows to extend the original action and reward spaces to produce observations in the same spaces. The representation of an environment through generative models has also been previously described by World Models [13] and its adaptation to limit order books [12], where the authors obtain latent representations of the environment enabling agents to learn a policy efficiently, and to transfer it back to the original environment. Other authors have explored applications of Gaussian inverse RL to learning investor sentiment models from historic data [20], to then extract tradable signals from the model.

2.2 Adversarial Learning under Maximum Causal Entropy

The adversarial IRL experiments presented in this paper follow the AIRL implementation described in [8], framed through a standard Markov decision process defined by a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, r, \gamma \rangle$ consisting of a state space \mathcal{S} ; a set \mathcal{A} of eligible actions; a model \mathcal{T} of transition dynamics $\mathcal{T}(s', a, s)$, where each $p(s'|s, a)$ represents the transition probability from state s through action a to state s' ; the unknown reward function $r(s, a)$ that we aim at recovering; and the discount factor γ taking values between $[0, 1]$.

In general terms, forward RL seeks an optimal policy π^* that maximizes an expected cumulative reward $\mathbb{E}[\sum_{t=0}^T \gamma^t r(s_t) | \pi^*]$ under the transition dynamics reflected in $\mathcal{T}(s', a, s)$ and the policy π^* . We denote the state-action trajectories derived from such expert policy as $\mathbf{x} = \{(s_t, a_t)\}_{t=0}^T$.

Under this framework, inverse RL attempts to recover the reward function r that would most likely have generated a given set of expert trajectories $\mathcal{D} = \{\mathbf{x}_n\}_{n=1}^N$ under the MDP $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \gamma \rangle$ of unknown reward. The adversarial learning methods considered will observe the trajectories in \mathcal{D} to infer a reward \hat{r} (that yields a policy $\hat{\pi}$) to explain \mathcal{D} . The performance of each model is then evaluated through the total cumulative reward $\mathbb{E}[\sum_{t=0}^T \gamma^t r(s_t) | \hat{\pi}]$ that $\hat{\pi}$ can recover against the total reward $\mathbb{E}[\sum_{t=0}^T \gamma^t r(s_t) | \pi^*]$ obtained by the expert agent under policy π^* .

One of the main challenges of inverse RL resides in its ill-posed nature: firstly, a given set of state-action trajectories \mathcal{D} may be explained by several different optimal policies [3]; the maximum entropy principle as proposed in [4] addresses this problem, assuming that the optimal policy $\pi^*(a|s)$ that would have generated the trajectories in \mathcal{D} follows $\exp\{Q'(s, a)\}$ as developed in [4, 21]. The soft Q' function refers to the learning process under the standard RL formulation, where the objective is regularized against a metric of differential entropy.

Secondly, the optimal policy determined may also be explained by a number of different reward functions, where an external observer lacks the means to distinguish true rewards from those product of non-stationary environment dynamics. The latter, when transferred to an environment with new dynamics, may fail to produce an optimal outcome. Adversarial inverse RL [8] introduces the concept of *disentangled rewards*, aiming at learning reward functions robust to variations in the environmental dynamics. This makes AIRL particularly attractive to study agents in financial markets, given their inherent need of continuous adaptation to changing dynamics.

The connection between inverse RL under maximum causal entropy and GANs as described by [9] compares the iterative cycles between generator and discriminator in the GAN with cases of inverse RL that employ neural nets to learn generic reward functions under unknown environment dynamics

[7, 22]. Moreover, a probabilistic approach to the problem allows to capture the stochastic nature of the relationship between the reward objective and the policy actually executed.

Following the cited previous works, we use as benchmark a GAIL implementation adapted to our market model: while GAIL is able to learn a policy from expert demonstrations, we expect AIRL to outperform in an environment with volatile, non-stationary dynamics, besides recovering a reward function in addition to the policy.

3 Experimental setup and results

The first requirement of our experiments is a model environment based on real financial data, that allows training of RL agents and is also compatible with the AIRL and GAIL learning algorithms. Learning a rich representation of the environment adds the general advantage of allowing RL models that are simpler, smaller and less expensive to train than model-free counterparts for a certain target performance of the learnt policy, as they search in a smaller space. Our experimental setup is based on an OpenAI Gym structure [23] used to adapt the AIRL and GAIL implementations by [8] to a financial version of a World Model environment as follows:

3.1 World models learnt from observed environments

World Models [13] provides an unsupervised probabilistic generative method to learn world representations in space and time through neural networks, including the transition dynamics as our experiment requires. RL agents operating in the real world would adapt their policies to their expectations of such transition dynamics, i.e. probabilities $p(s'|s, a)$ of certain future state s' given the current state s and choice of action a contained in $\mathcal{T}(s', a, s)$. RL agents can then be trained within the world model, with their learnt policies remaining optimal once transferred back to the original environment (or even outperforming the agents trained in the original environment in some instances).

These models capture three elements of the original environment: firstly, an auto-encoder learns a latent representation of each state vector from its original high-dimensional space, compressing all the observed information of that state into a vector z within a latent space of lower dimension (Fig. 1).

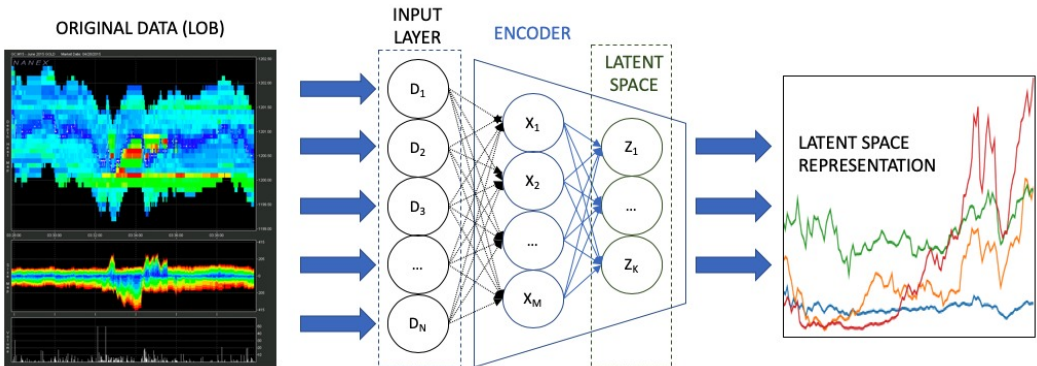


Figure 1: An auto-encoder extracts a latent-space representation from the original LOB data [24]

Secondly, after learning the latent vector for each state, the model learns the transition dynamics connecting consecutive states through certain action choices a_t . This model of transition probabilities $p(s'|s, a)$ of $\mathcal{T}(s', a, s)$ in the MDP is learnt through an RNN [25, 26, 27], able to learn sequential information from time series along various time horizons by learning $P(z_{t+1}|a_t, h_t, z_t)$ where h_t represents a hidden state of the RNN.

The choice of RNN architectures is based on their suitability to capture the transition model of sequential data [28], thus learning some predictive modelling of future states. In order to better reflect the stochastic nature of the target environment, the deterministic output of the RNN is then fed into a

Mixture Density Network (MDN) [29] to produce a probabilistic prediction of next latent state. This combined RNN-MDN architecture is detailed further in [13, 30].

Finally, an RL module decides which action to take given all the above information, including the transition predicted in the latent space. As noted in [13], this approach has the additional advantage of providing RL agents with access to the latent space learnt by the model, which includes representations of both the current state and an indication of what to expect in upcoming states through the transitions modelled, sampling from the probability distribution of future states provided by the RNN-MDN.

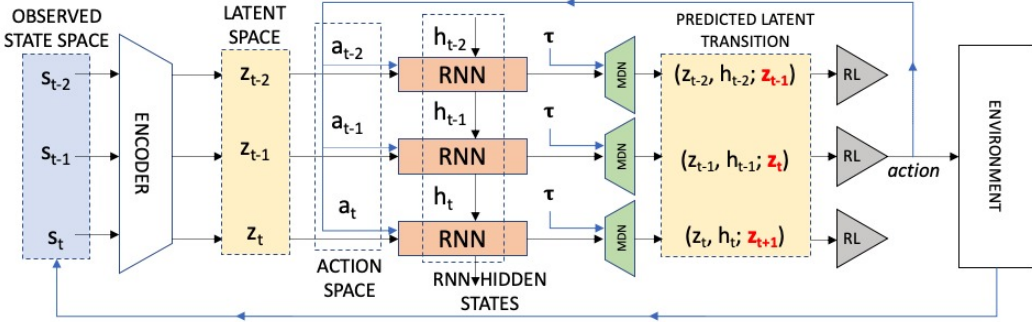


Figure 2: Model for prediction of latent space transitions. Our RL processes are run in the latent state space.

3.2 LOB version of world models learnt from real market data

The original version of World Models learns from environments based on computer simulations, such as `CarRacing-v0` [31] and `VizDoom` [32]. We base our experiments on the extension of this architecture to financial environments presented in [12], where the auto-encoder is trained to learn a latent space from sequences of snapshots of actual limit order book data, and the RNN-MDN models the transitions between consecutive time frames as follows:

The data used to train the model consists of LOB time series from shares of Tencent Holdings Ltd. (HK-700) traded in the Hong Kong stock exchange along sixty trading days between January and March, 2018. Data from the following twenty trading days in April is then used as testing reference for the adversarial learning algorithms. Each state contains the sequence of the last 10 data ticks for 3 LOB levels, so that the sequential information necessary to learn the transition dynamics is captured in the data for each state.

LOB data is of high dimension, generally denominated Level II market data, and features several relevant data points for each timestamp. Firstly, it contains various levels of bid (buy) and ask (sell) prices for each timestamp, typically ranked in decreasing order of competitiveness. The *mid price* is generally derived as an average between the highest bid and lowest ask in the first LOB level. Secondly, it includes the trading volume associated with each of such prices offered. The combined progression in time of this data structure is often represented as a tensor of three dimensions [33]. Finally, trade stamp series contain the price and size of the last transactions executed out of previous LOB states, used in this model as RL exploration.

In order to allow the RNN-MDN to model the transition dynamics, each state s_t consists of 10 consecutive LOB data ticks for each of the above features. The auto-encoder learns a compressed version of the data in lower dimensionality, through a latent space representation of short sequences of original market data (Fig. 1). The model used in [12] features a CNN-based auto-encoder, looking back 10 updates in the limit order book that contains 3 levels of orders, and condenses the $(10 * 3 * 4)$ matrix into a z vector of dimension 12.

The RNN-MDN used to model the transition dynamics assumes that $p(s'|s, a)$ follows a Gaussian mixture distribution [12, 13], where the model learns the parameters of the distribution $p(s'|s, a) = \sum_{k=1}^K w_k(s, a) \mathcal{D}(s' | \mu_k(s, a), \sigma_k^2(s, a))$ with $\mathcal{D}(\cdot)$ assumed to be Gaussian, $k = 5$ and 128 neurons

per layer in the RNN. Once the prediction of the next state is sampled in the latent space, the reward with respect to the present state and the action chosen is produced through a regression model.

3.3 Training experts in the LOB latent space

Once the full world model is trained and integrated into an OpenAI Gym structure, we train in it the expert agents whose trajectories will serve as input to the adversarial learning experiments. We follow the selection of RL algorithms in [12] proved to learn policies in the world model that remain profitable when transferred back to the real environment: double DQN [2], advantage actor critic (A2C) [14] and policy gradient [15].

We now use each of the three trained expert agents to generate collections of 100 state-action trajectories for each agent, with a length of 1000 (s_t, a_t) pairs each. Among the three experts, A2C seems to be the best performing agent in terms of high mean and low variance of cumulative rewards, followed by DQN (Fig. 3):

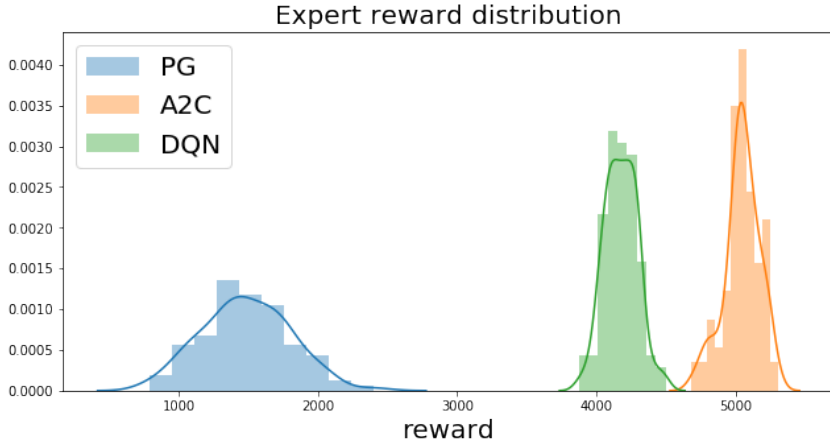


Figure 3: Distribution of rewards obtained by each expert RL agent.

3.4 Adversarial learning from Reinforcement Learning trading experts

As in general GAN structures, adversarial inverse reinforcement learning is implemented through a generator and a discriminator: two neural networks that contest with each other. The generator learns to produce an output mapping from a latent space to a target data distribution, while the discriminator tries to classify which samples come from the true distribution and which were produced by the generator.

In the case of AIRL, the generator learns to produce state-action pairs $(s_t, a_t)_i^L$ based on those observed in the trajectories from the expert demonstrations $(s_t, a_t)_i^E$, while the discriminator D_θ tries to classify which state-action pairs were produced by the generator against those actually produced by an expert agent, training through binary logistic regression. The reward estimate \hat{r}_θ is then updated from the GAN minimax loss: $\log D_\theta(s, a, s') - \log(1 - D_\theta(s, a, s'))$. An estimate $\hat{\pi}$ of the policy is then updated from \hat{r}_θ iteratively.

Our experiments are initialized with the expert trajectories $(s_t, a_t)_i^E$ produced by executing each expert policy π^E within the world model learnt from data gathered between January and March. We then initialize a policy $\hat{\pi}_0$ and the discriminator D_θ . The adversarial observer then updates iteratively D_θ and the reward \hat{r}_θ that produces a policy $\hat{\pi}$, which generates samples of $(s_t, a_t)_i^L$ increasingly similar to $(s_t, a_t)_i^E$. The evaluations of cumulative rewards produced by each $\hat{\pi}$ take place with a model environment run with test data from April.

We present in Fig. 4 and Table 1 the results of running AIRL and GAIL on 200 training iterations: AIRL is found to outperform GAIL in all the three cases considered (learning from demonstrations of each expert trained), consistently with the initial motivation of the experiment.

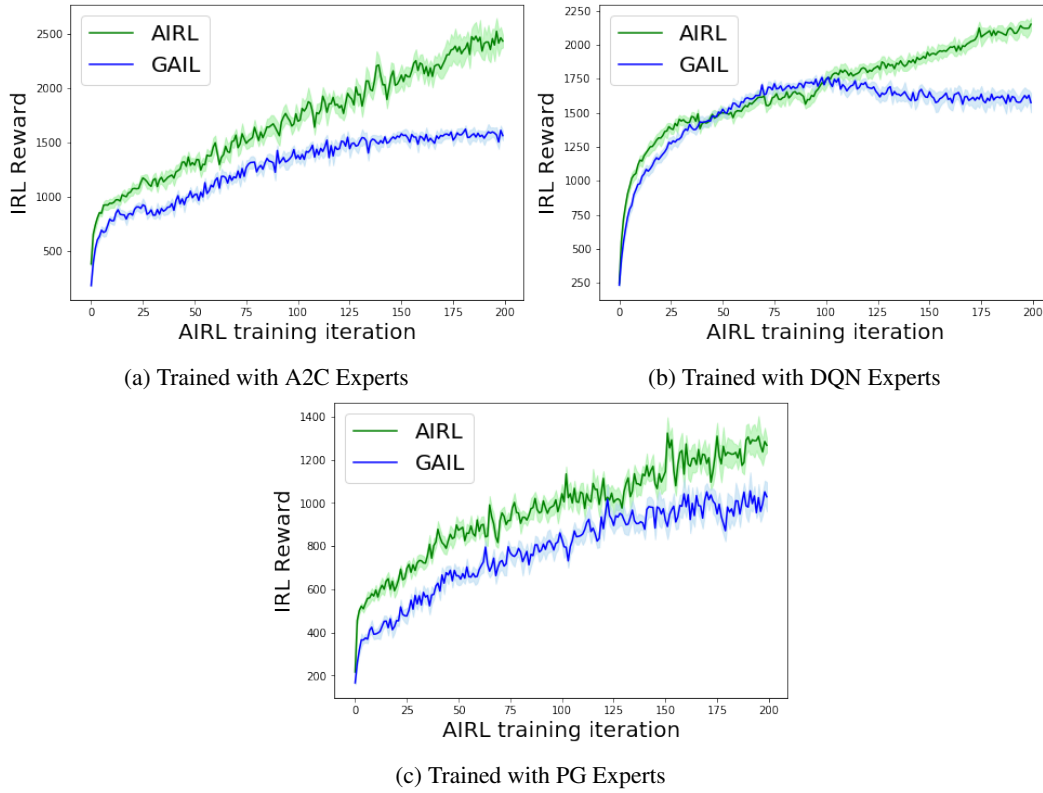


Figure 4: Comparative of rewards obtained through Airl and Gail by learning from each expert agent.

Agent	Average Return	Std Return
Expert A2C	5047.3	125.9
Airl	2433.0	59.7
Gail	1563.1	43.4

(a) Adversarial learning from A2C Experts

Agent	Average Return	Std Return
Expert DQN	4177.4	116.6
Airl	2151.1	45.6
Gail	1574.5	70.8

(b) Adversarial learning from DQN Experts

Agent	Average Return	Std Return
Expert PG	1483.3	313.2
Airl	1267.7	33.8
Gail	1029.1	67.9

(c) Adversarial learning from PG Experts

Table 1: Summary of cumulative rewards obtained by Airl, Gail and each expert agent in the order book.

4 Conclusions

We have presented an experimental setup to adapt adversarial inverse reinforcement learning and generative adversarial imitation learning to a volatile financial environment, in order to evaluate their ability to learn from expert trajectories in latent space simulations: the latent model employed allows the training of the adversarial algorithms based on real market data. The results obtained show that both methods can be trained against a latent space model of the market, while the specific advantage of AIRL for learning rewards robust to variable environment dynamics allows it to outperform GAIL in the cumulative rewards recovered from each of the three agents as the number of iterations increases.

A review of the original data series (Fig. 5) confirms the variability of mid price series and of their volatility levels (measured across various horizons) when comparing the data corresponding to the training period against that of the testing period. The rewards learnt by AIRL from the training model are robust to such variable dynamics, hence its outperformance over GAIL when evaluated on the test period.

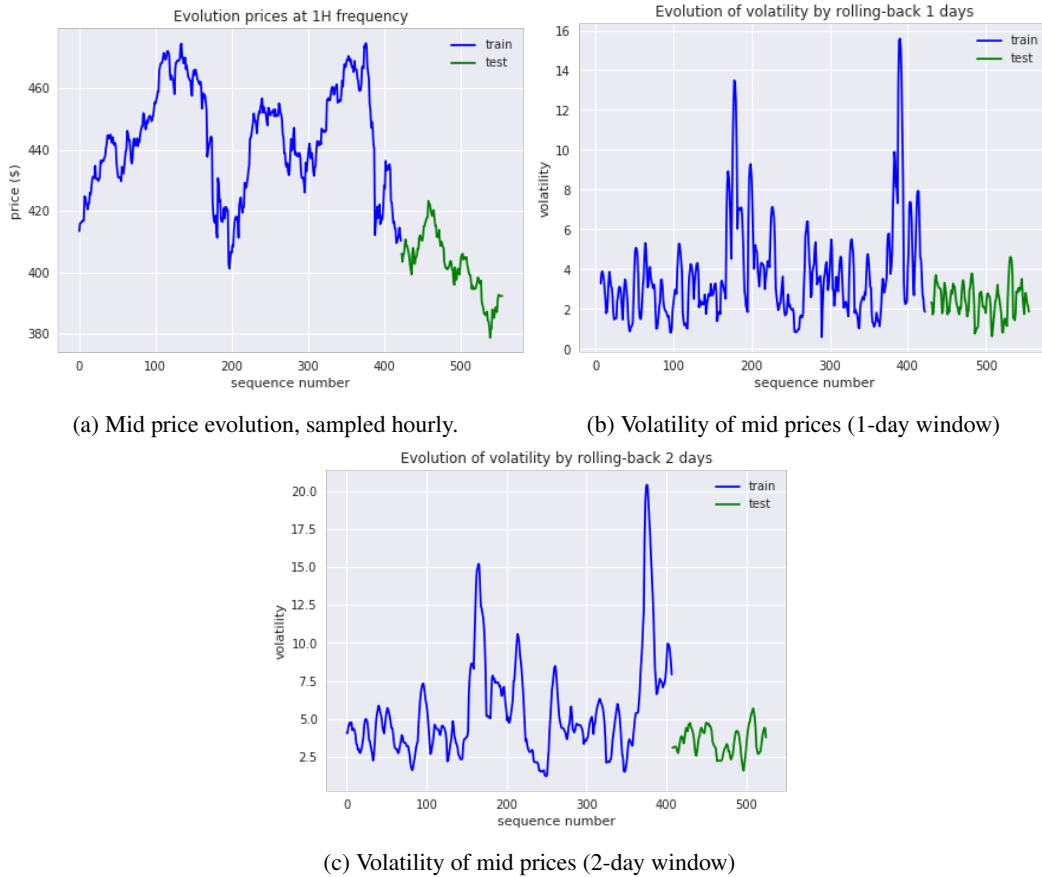


Figure 5: The original data series used to train the latent space model and the experts (blue) and to test the adversarial learning methods (green) contain significant variability of price and volatility levels.

References

- [1] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [2] Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. *CoRR*, abs/1509.06461, 2015.
- [3] Andrew Y. Ng, Daishi Harada, and Stuart J. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Proceedings of the Sixteenth International Conference on Machine Learning, ICML '99*, pages 278–287, 1999.
- [4] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.
- [5] Sergey Levine, Zoran Popovic, and Vladlen Koltun. Nonlinear inverse reinforcement learning with gaussian processes. In *Advances in Neural Information Processing Systems*, pages 19–27, 2011.
- [6] Brian D. Ziebart. *Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy*. PhD thesis, Pittsburgh, PA, USA, 2010. AAI3438449.
- [7] Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *International Conference on Machine Learning*, pages 49–58, 2016.
- [8] Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2017.
- [9] Chelsea Finn, Paul F. Christiano, Pieter Abbeel, and Sergey Levine. A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. *CoRR*, abs/1611.03852, 2016.
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.
- [11] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 4565–4573. Curran Associates, Inc., 2016.
- [12] Haoran Wei, Yuanbo Wang, Lidia Mangu, and Keith Decker. Model-based reinforcement learning for predictions and control for limit order books, 2019. Cite Arxiv 1910.03743.
- [13] David Ha and Jürgen Schmidhuber. World models. *CoRR*, abs/1803.10122, 2018.
- [14] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1928–1937, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [15] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Machine Learning*, pages 229–256, 1992.
- [16] Dieter Hendricks, Adam Cobb, Richard Everett, Jonathan Downing, and Stephen J. Roberts. Inferring agent objectives at different scales of a complex adaptive system. Papers 1712.01137, arXiv.org, December 2017.
- [17] Jacobo Roa Vicens, Cyrine Chtourou, Angelos Filos, Francisco Rullan, Yarin Gal, and Ricardo Silva. Towards inverse reinforcement learning for limit order book dynamics. In *ICML Workshop 'AI in Finance: Applications and Infrastructure for Multi-Agent Learning' at the 36th International Conference on Machine Learning*, 2019.
- [18] Igor Halperin and Ilya Feldshteyn. Market self-learning of signals, impact and optimal trading: Invisible hand inference with free energy. *arXiv preprint arXiv:1805.06126*, 2018.

- [19] Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H. Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, Ryan Sepassi, George Tucker, and Henryk Michalewski. Model-based reinforcement learning for atari. *CoRR*, abs/1903.00374, 2019.
- [20] Steve Y. Yang, Yangyang Yu, and Saud Almahdi. An investor sentiment reward-based trading system using gaussian inverse reinforcement learning algorithm. *Expert Syst. Appl.*, 114:388–401, 2018.
- [21] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1352–1361. JMLR. org, 2017.
- [22] Abdeslam Boularias, Jens Kober, and Jan Peters. Relative entropy inverse reinforcement learning. In Geoffrey Gordon, David Dunson, and Miroslav Dudík, editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 182–189, Fort Lauderdale, FL, USA, 11–13 Apr 2011. PMLR.
- [23] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016. Cite arxiv:1606.01540.
- [24] LOB representation from NxCore API. 2019. Nanex Corp.
- [25] Jürgen Schmidhuber. Making the world differentiable: On using self-supervised fully recurrent neural networks for dynamic reinforcement learning and planning in non-stationary environments. Technical report, 1990.
- [26] Jürgen Schmidhuber. An on-line algorithm for dynamic reinforcement learning and planning in reactive environments. In *In Proc. IEEE/INNS International Joint Conference on Neural Networks*, pages 253–258. IEEE Press, 1990.
- [27] Jürgen Schmidhuber. Reinforcement learning in markovian and non-markovian environments. In R. P. Lippmann, J. E. Moody, and D. S. Touretzky, editors, *Advances in Neural Information Processing Systems 3*, pages 500–506. Morgan-Kaufmann, 1991.
- [28] David Ha and Douglas Eck. A neural representation of sketch drawings. *CoRR*, abs/1704.03477, 2017.
- [29] Christopher M. Bishop. Mixture density networks. Technical report, 1994.
- [30] Alex Graves. Generating sequences with recurrent neural networks. *CoRR*, abs/1308.0850, 2013.
- [31] Oleg Klimov. Carracing-v0. 2016. <https://gym.openai.com/envs/CarRacing-v0/>.
- [32] Michal Kempka, Marek Wydmuch, Grzegorz Runc, Jakub Toczek, and Wojciech Jaskowski. Vizdoom: A doom-based AI research platform for visual reinforcement learning. *CoRR*, abs/1605.02097, 2016.
- [33] Dat Thanh Tran, Martin Magris, Juho Kannianen, Moncef Gabbouj, and Alexandros Iosifidis. Tensor representation in high-frequency financial data for price change prediction. *CoRR*, abs/1709.01268, 2017.

Disclaimer

Opinions and estimates constitute our judgement as of the date of this Material, are for informational purposes only and are subject to change without notice. This Material is not the product of J.P. Morgan’s Research Department and therefore, has not been prepared in accordance with legal requirements to promote the independence of research, including but not limited to, the prohibition on the dealing ahead of the dissemination of investment research. This Material is not intended as research, a recommendation, advice, offer or solicitation for the purchase or sale of any financial product or service, or to be used in any way for evaluating the merits of participating in any transaction. It is not a research report and is not intended as such. Past performance is not indicative of future results. Please consult your own advisors regarding legal, tax, accounting or any other aspects including suitability implications for your particular circumstances. J.P. Morgan disclaims any responsibility or liability whatsoever for the quality, accuracy or completeness of the information herein, and for any reliance on, or use of this material in any way.

Important disclosures at: www.jpmorgan.com/disclosures