
Inferring agent objectives at different scales of a complex adaptive system

D. Hendricks* A. Cobb R. Everett J. Downing S.J. Roberts

Machine Learning Research Group
Oxford-Man Institute of Quantitative Finance
University of Oxford

*Corresponding author: dieter.hendricks@eng.ox.ac.uk

Abstract

We introduce a framework to study the effective objectives at different time scales of financial market microstructure. The financial market can be regarded as a complex adaptive system, where purposeful agents collectively and simultaneously create and perceive their environment as they interact with it. It has been suggested that multiple agent classes operate in this system, with a non-trivial hierarchy of top-down and bottom-up causation classes with different effective models governing each level. We conjecture that agent classes may in fact operate at different *time scales* and thus act differently in response to the same perceived market state. Given scale-specific temporal state trajectories and action sequences estimated from aggregate market behaviour, we use Inverse Reinforcement Learning to compute the effective reward function for the aggregate agent class at each scale, allowing us to assess the relative *attractiveness* of feature vectors across different scales. Differences in reward functions for feature vectors may indicate different objectives of market participants, which could assist in finding the scale boundary for agent classes. This has implications for learning algorithms operating in this domain.

1 Introduction

Equity financial markets consist of multiple competing agents operating through a centralised electronic exchange, giving rise to non-linear interactions both through time and among agent classes. *Investors* use their understanding of asset dynamics to time buying and selling decisions for financial gain, *traders* use their understanding of market dynamics to plan trades and minimise the cost of realising investment decisions and *market makers* use their understanding of investor demand to profit from liquidity provision. The field of market microstructure [1] studies the dynamics of price formation in this system at intraday time scales, considering how mechanistic rules, regulatory oversight and social behaviours of participants interact to manifest the observed time series'. The *Complexity Economics* [2] paradigm seeks to explain observed behaviours through the lens of complex adaptive systems, where competing agents continually adapt their actions and strategies based on the observed state they mutually create. Wilcox and Gebbie [3] take this further, proposing a mechanism for bottom-up and top-down causation, with level-specific effective models governing actors and inter-level interaction via noise terms. Actors at each level perceive the system in a different way, which invalidates the use of hierarchies of the same effective model to capture system complexity.

We are interested in developing learning algorithms in this domain, where an understanding of scale-specific state representation, in the context of competing agents, is key to ensuring relevant features can be exploited and useful learning can take place faster than the natural time-scale of the system. Galla, Farmer and Sanders [4, 5] investigate the nature of agent learning in complicated games, using Experience Weighted Attraction (EWA) to evaluate the propensity for asymptotic learning. They demonstrate the importance of understanding agent payoff (reward) correlations when

setting learning rates, to ensure that learning is feasible and chaotic regimes are avoided. This study seeks to understand the landscape of competing agents at different time scales, to interrogate the importance of scale for learning policies in this domain.

We use Inverse Reinforcement Learning (IRL) to compute the effective reward function at different scales of equity market microstructure, using scale-specific temporal state trajectories and action sequences estimated from aggregate market behaviour. This allows us to identify *attractive* states and assess the attractiveness of the associated feature vector across different scales. This is a first step towards understanding relative objectives of agent classes at different scales of this system.

2 Method

2.1 Inverse Reinforcement Learning for state attractiveness

IRL aims to infer the reward function in a Markov decision process (MDP) defined by the tuple $m = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \gamma, \mathcal{R} \rangle$, describing the state space \mathcal{S} , the action space \mathcal{A} , the transition function \mathcal{P} , the discount factor γ and the reward \mathcal{R} . Given access to sample state-action space trajectories collected from observing an agent’s behaviour, the objective is to find a reward function that induces agents to follow trajectories matching the expert trajectories.

While many IRL algorithms have been proposed ([6, 7, 8, 9]), we will use Maximum entropy (MaxEnt) IRL [10], which represents a move towards looking at probability distributions over paths. Ziebart et al. employ the principle of maximum entropy to pick the least informative set of parameters for a linear reward function $r = \theta^\top \mathbf{f}_\zeta$, a linear combination of state-action trajectories \mathbf{f}_ζ that match the feature expectations between an agent’s observed trajectories and the learner’s behaviour.

This matching equation,

$$\sum_{\text{Path } \zeta_i} P(\zeta_i) \mathbf{f}_{\zeta_i} = \tilde{\mathbf{f}} \quad (1)$$

provides the constraints for maximum entropy, where $\tilde{\mathbf{f}} = \frac{1}{m} \sum_i \mathbf{f}_{\zeta_i}$ is the average empirical feature count from m observed trajectories. Therefore summing the feature mappings along a trajectory results in $\mathbf{f}_{\zeta_i} = \sum_{s_j \in \zeta} \mathbf{f}_{s_j}$.

Ziebart et al. evaluate a probability distribution over trajectories through space $P(\zeta_i | \theta, T)$ that give trajectories with equal rewards, equal probabilities and assign an exponentially higher preference to higher rewards. Note that T is the transition model of the MDP. The optimal set of parameters θ^* is then obtained by maximising the likelihood, $L(\theta)$, i.e.

$$\theta^* = \underset{\theta}{\operatorname{argmax}} L(\theta) = \underset{\theta}{\operatorname{argmax}} \sum_{\text{examples}} \log P(\tilde{\zeta} | \theta, T) \quad (2)$$

We will use the MaxEnt IRL implementation provided by Matthew Alger [11].

2.2 Determining temporal state trajectories

Hendricks et al. [12] propose an approach for the detection and online estimation of intraday temporal states from equity market microstructure features. They found an interesting hierarchy of system behaviour at different calendar time scales, with results suggesting there may be different universality classes characterising behaviour at each scale. The approach uses an analogy to the q-state Potts model to develop an unsupervised clustering technique consistent with finding meta-stable object configurations in complex systems. Time periods are clustered into states based on observed *trade price, spread, volume* and *volume imbalance* features for major stocks on an exchange. Significant states are identified, with the associated State Signature Vectors (SSV) used for online state detection and assignment.

We will construct temporal state trajectories at different calendar scales based on the SSV state assignment using the method in [12].

Given the objective of categorising agent *classes* at different scales, we will use the sequence of *average price returns* at each scale to determine associated actions. The motivation is that if the behaviour at a certain scale resulted in a *positive (negative) price return*, this could indicate *net buying (selling)* decisions for that class. We will thus use the sign of average price return to assign the action to one of three states, $\{BUY, SELL, NEUTRAL\}$.

3 Experiments

Data State and action trajectories were computed from tick-level trades and top-of-book quotes for 42 stocks on the Johannesburg Stock Exchange (JSE) from 1 November 2012 to 30 November 2012, sourced from the Thomson Reuters Tick History (TRTH) database. This data was resampled based on the scale considered. We use four features for each stock: Change in i) trade *price*, ii) quote *spread* (Ask price - Bid price), iii) trade *volume* and iv) quote *volume imbalance* ($\frac{\text{Ask volume}}{\text{Bid volume}}$).

Results Figure 1 shows the feature vectors and estimated reward functions for each state at the four candidate time scales: 5, 15, 30 and 60 minutes. It is interesting that low *spread* states are unfavourable, as there are fewer opportunities when markets are highly liquid, whereas high *spread* states provide opportunity for profit via market dislocations. Upon examination of the relative rewards and feature vectors, it appears there are differences in attractiveness across scales.

Figure 2 shows the relative reward across scales for similar feature vectors. We used a simple *K*-means clustering algorithm to group *all* states into 6 clusters on the basis of feature vector similarity. We then sum the (normalised) rewards for all states at the same scale in each cluster, so that we have a net reward for each scale given the same feature vector. Clusters 1-5 all show both positive and negative rewards, providing some evidence that agents react differently to observed features at different scales. Interpreting the rewards using relative cluster centroid *spread* and *volume imbalance* provides some intuition.

For Cluster 1, negative *volume imbalance* and negative *spread* could indicate a larger quantity of aggressive buy orders, which could have a short-term negative effect of buying at a higher price, but if the stock is held at longer time scales, there could be a net gain in asset value.

For Cluster 5, positive *volume imbalance* and negative *spread* indicate a larger quantity of aggressive sell orders. This could be positive at the 5-min scale, where a short-term sell is more easily matched, but the increase in aggressive ask quotes could translate into downward price pressure, leading to negative rewards at higher scales.

For Cluster 6, negative *volume imbalance* and positive *spread* could be enable agents to buy larger quantities of stocks at lower prices, with the imbalance leading to further price appreciation across all scales.

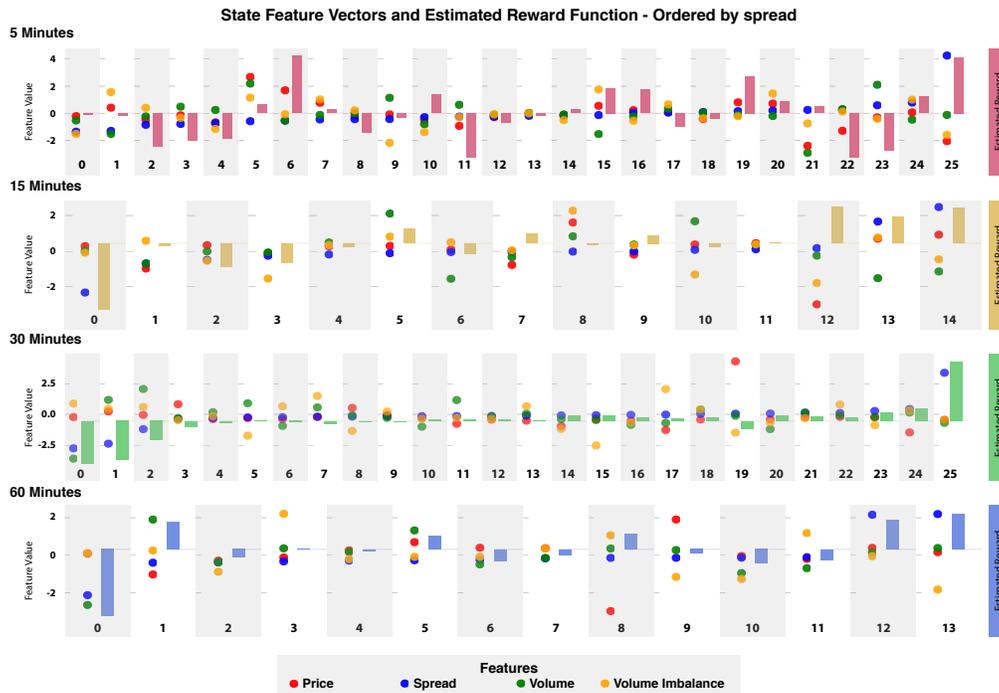


Figure 1: This figure shows the feature vectors (dots) and estimated reward function (bars) for each state, at 4 candidate time scales (5, 15, 30 and 60 minutes)

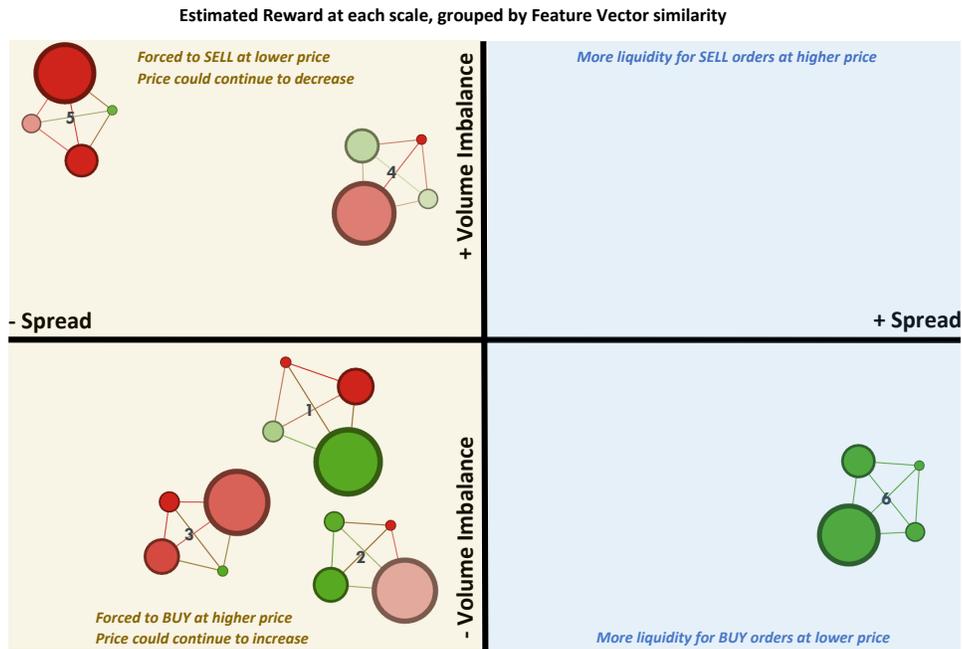


Figure 2: This figure illustrates the relative reward across scales for similar feature vectors. Node size indicates scale (smallest=5min, largest=60min), node colour indicates reward (red=low negative, green=high positive), nodes are clustered by feature vector similarity, and clusters are positioned by the relative *spread* and *volume imbalance* feature values of the cluster centroids.

4 Discussion and future work

We have provided a framework which allows us to interrogate the relative attractiveness of feature vectors at different scales of a complex adaptive system. This will help to assess whether a scale boundary exists for agent classes, to better inform learning algorithm specification. Further work will assess a complete spectrum of event time scales, consistent with how machine algorithms process information, to provide clarity on boundaries between high-frequency (algorithmic) market making, machine trading, human trading and investment decisions.

References

- [1] M. O'Hara. *Market Microstructure Theory*. Blackwell publishing, 1998.
- [2] W.B. Arthur. *Complexity and the economy*. Oxford University Press, Oxford, UK, 2014.
- [3] D. Wilcox and T. Gebbie. Hierarchical causality in financial economics. *Working paper*, available at SSRN: <https://ssrn.com/abstract=2544327>, 2014.
- [4] T. Galla and J.D. Farmer. Complex dynamics in learning complicated games. *Proceedings of the National Academy of Sciences*, 110(4):1232–1236, 2013.
- [5] J.B.T. Sanders, J.D. Farmer, and T. Galla. The prevalence of chaotic dynamics in games with many players. *Working paper*, *arXiv:1612.08111 [q-fin.EC]*, 2016.
- [6] S. Russell. Learning agents for uncertain environments. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 101–103. ACM, 1998.
- [7] A.Y. Ng and S.J. Russell. Algorithms for inverse reinforcement learning. In *ICML*, pages 663–670, 2000.
- [8] S. Levine, Z. Popovic, and V. Koltun. Nonlinear Inverse Reinforcement Learning with Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 19–27, 2011.
- [9] M. Wulfmeier, P. Ondruska, and I. Posner. Maximum Entropy Deep Inverse Reinforcement Learning. *Working paper*, *arXiv:1507.04888 [cs.LG]*, 2016.
- [10] B.D. Ziebart, A.L. Maas, J.A. Bagnell, and A.K. Dey. Maximum Entropy Inverse Reinforcement Learning. In *AAAI*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.
- [11] M. Alger. *Inverse Reinforcement Learning*, 2016.
- [12] D. Hendricks, T. Gebbie, and D. Wilcox. Detecting intraday financial market states using temporal clustering. *Quantitative Finance*, 16(11):1657–1678, 2016.