

Financial Computing and Analytics Group



<http://fincomp.cs.ucl.ac.uk/>

The Financial Computing and Analytics research group investigates socio-economic systems using methods from computer science, applied mathematics, computational statistics and network theory.



UCL
Centre for Blockchain Technologies



Systemic Risk Centre



Silvia Bartolucci



Philip Treleven



Guido Germano



Fabio Caccioli



Christopher Clack



Giacomo Livan



Paolo Barucca



Carolyn Phelan



Daniel Hulme



Robert E Smith



Jiahua Xu



Tomaso Aste



Simone Righi



Paolo Tasca



Geoff Goodell



Nikhil Vadgama



Jessica James



Nick Firoozye



Ariane Chapelle



Denise Gorse

Research

- computational finance
- data-driven modelling
- artificial intelligence
- financial risk management
- blockchain technology
- digital economy
- systemic risk
- numerical pricing of derivatives
- agent-based simulation
- empirical finance
- market microstructure
- algorithmic trading
- high-frequency trading
- data science
- big data analytics
- network analysis
- machine learning
- price formation
- portfolio optimization

EPSRC

Pioneering research and skills



Education:

- PhD Doctoral Training Centre in Financial Computing
- MSc Computational Finance
- MSc Financial Risk Management
- MSc Financial Technology, forthcoming 2021-22
- MSc Emerging Digital Technologies, forthcoming 2021-22

What can machines learn?



NEW NAVY DEVICE LEARNS BY DOING
Psychologist Shows Embryo of Computer Designed to Read and Grow Wiser

WASHINGTON, July 1 (UPI)—The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The embryo, called "Perceptron," is a machine that learns right and left after fifty attempts at the Navy demonstration on women.

The service said it would use the prototype to build the first of its Perceptron, costing \$250,000, and to read and write. It is expected to be finished in about a year at a cost of \$100,000.

Dr. Frank Rosenblatt, designer of the Perceptron, said that the machine would be the first one to think as the human brain. As its brain be-

New York Times 1958

lags, Perceptrons will make mistakes at first, but will grow wiser as it gains experience, he said.

Dr. Rosenblatt, a research psychologist at the Cornell Aeronautical Laboratory, Buffalo, said, "Perceptrons might be fine in the phone or mechanical dial industries."

Without Human Controls

The Navy said the perceptron would be the first man-made mechanism "capable of recognizing, recognizing and identifying its surroundings without any human help or control."

The "Perceptron" is designed to remember images and information it has perceived from ordinary computers remember only what is fed into them on punch cards or magnetic tape.

Later Perceptrons will be able to recognize people and call out their names and normally learn to read and write in English by inspecting or writing in another language if you prefer.

Mr. Rosenblatt said in principle it would be possible to build brains that could reproduce themselves on an assembly line and which would be conscious of their existence.

In the first fifty trials, the machine made 100 errors, one with the picture of the left side and the other with squares on the right side.

Learns by Doing

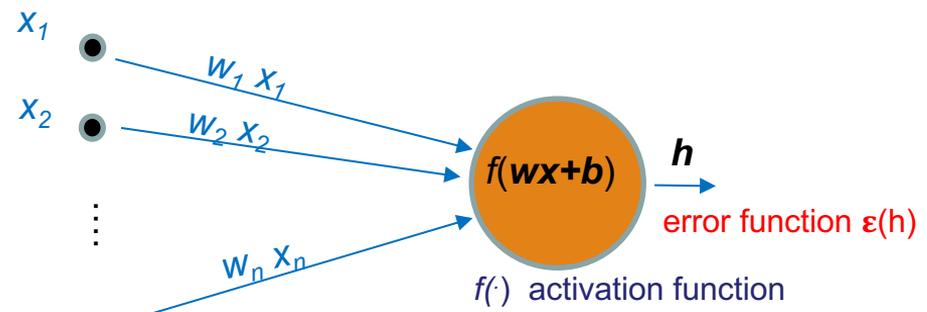
In the first fifty trials, the machine made 100 errors, one registering a "0" for the left picture and "0" for the right picture.

Dr. Rosenblatt said he could explain why the machine learned only in highly restricted terms. But he said the computer has now been a "self-organizer."

The first Perceptron will have about 1,000 electronic "associative cells" receiving electrical impulses from an eye-like sensing device with 400 photo-cells. The machine has 10,000,000,000 responsive cells, including 100,000,000 connections with the eyes.

The perceptron

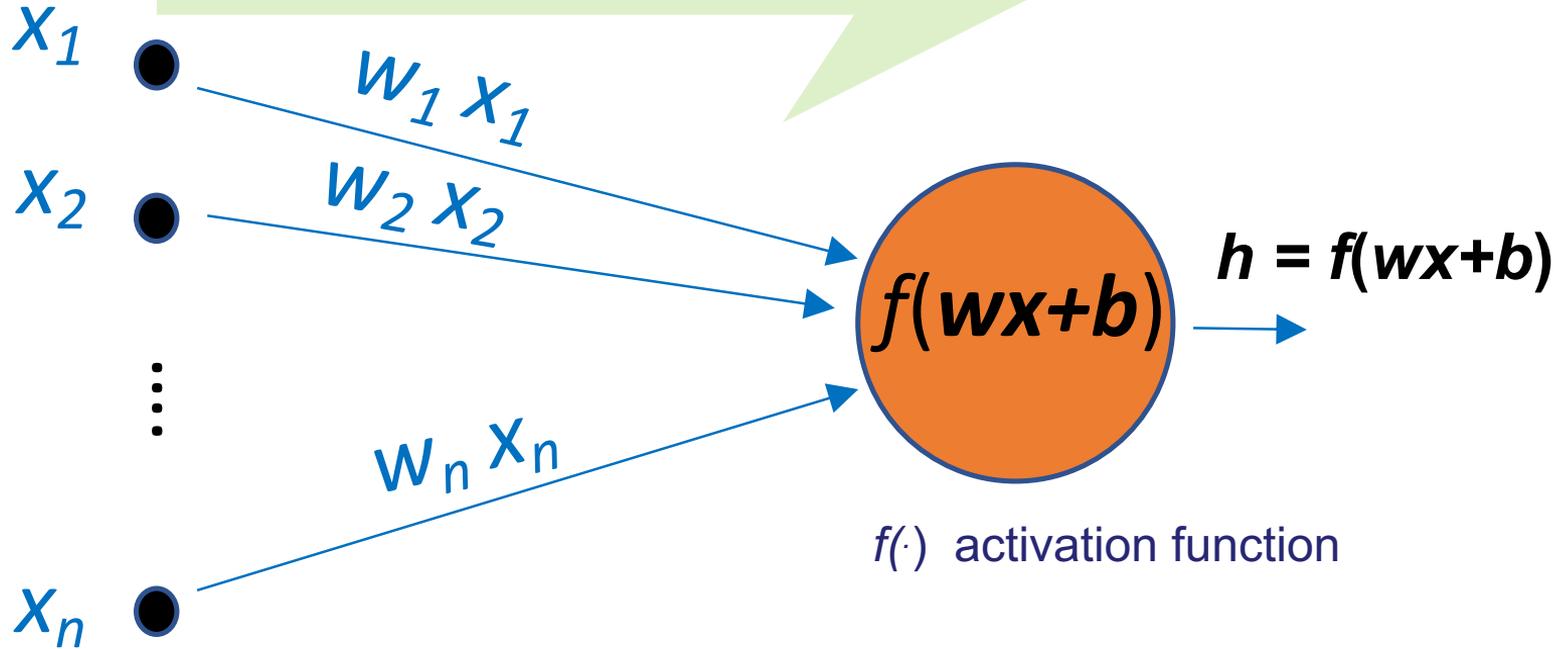
(Rosemblatt '57, Minsky Paper '69)



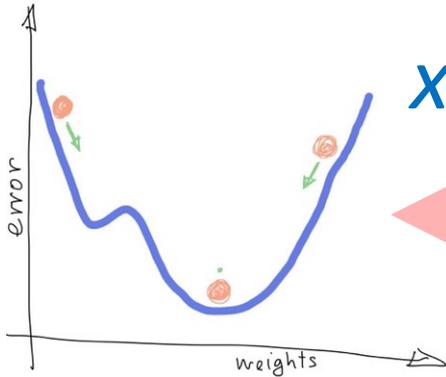
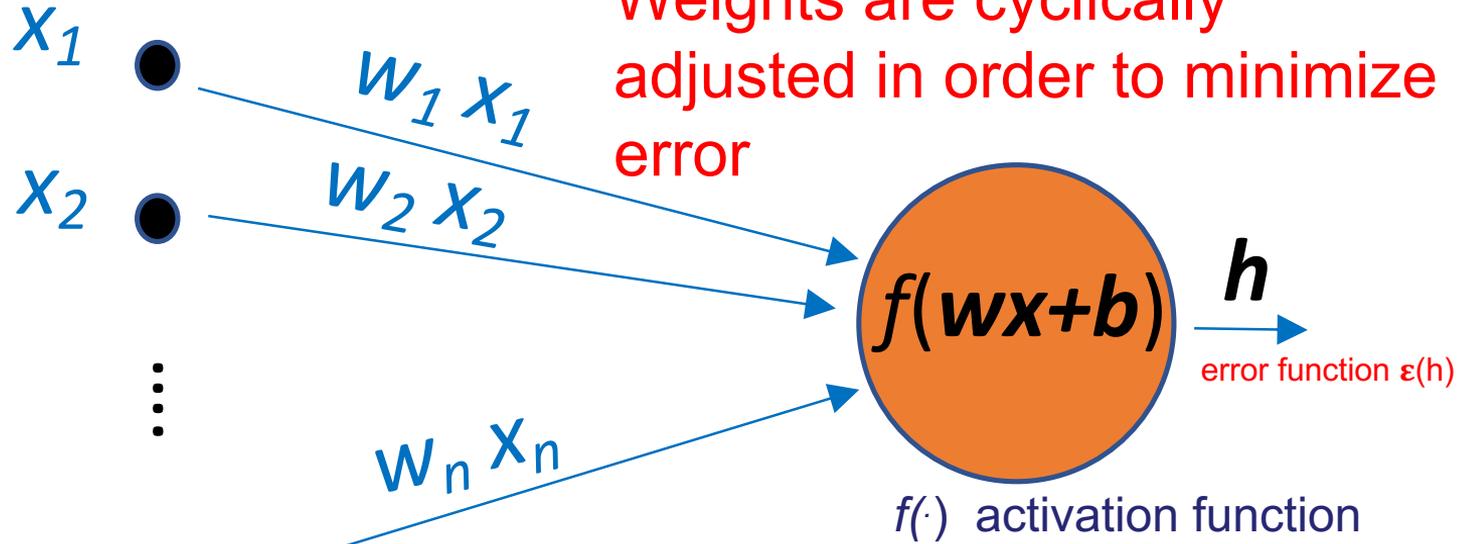
It learns from examples

Frank Rosenblatt with a Mark I Perceptron computer in 1960

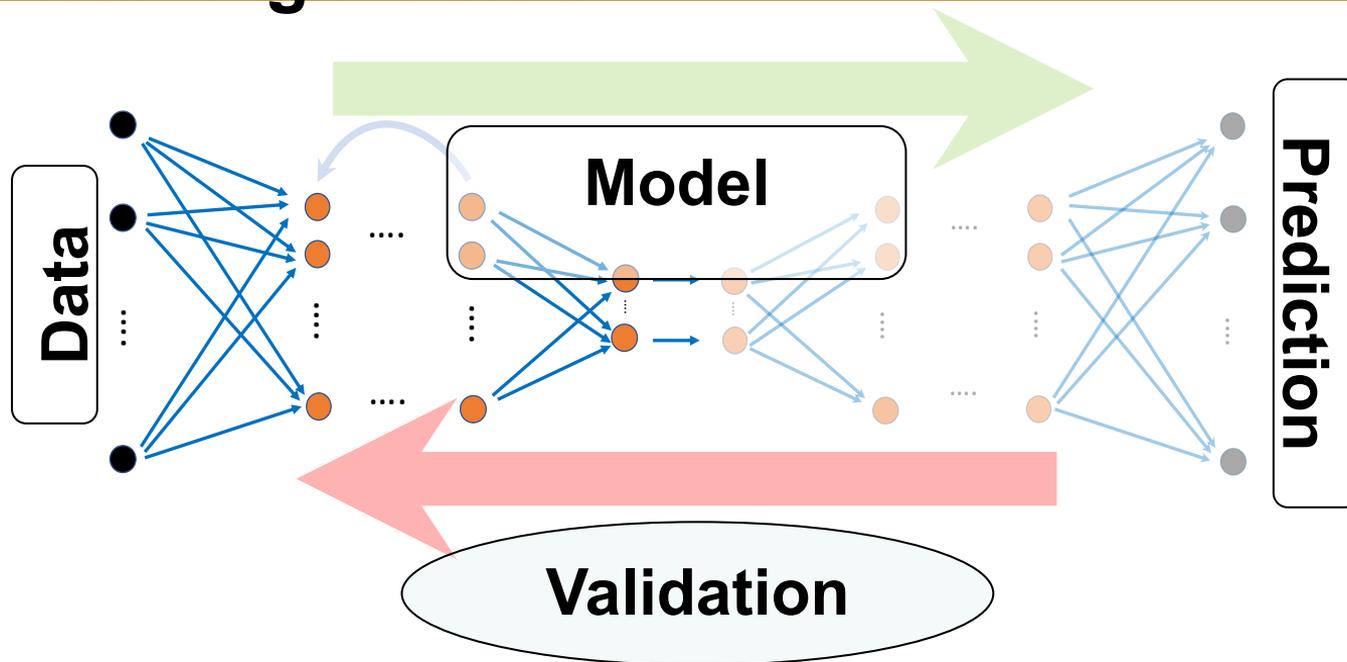
forward propagation of signal



Weights are cyclically adjusted in order to minimize error



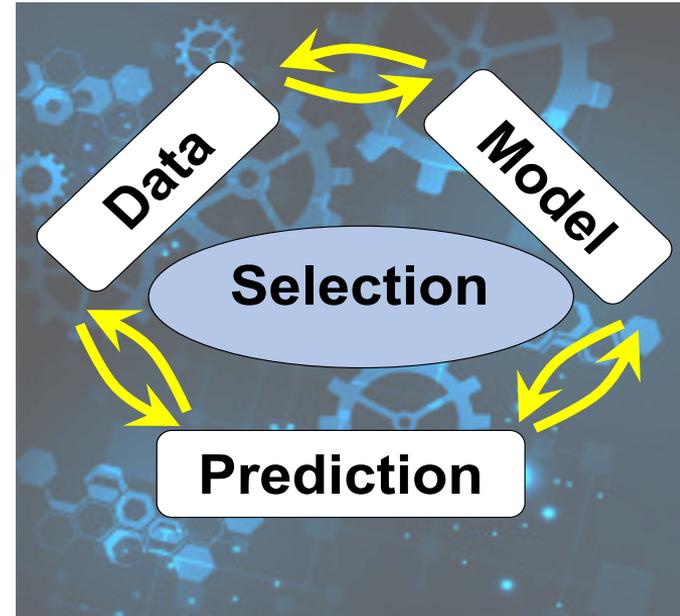
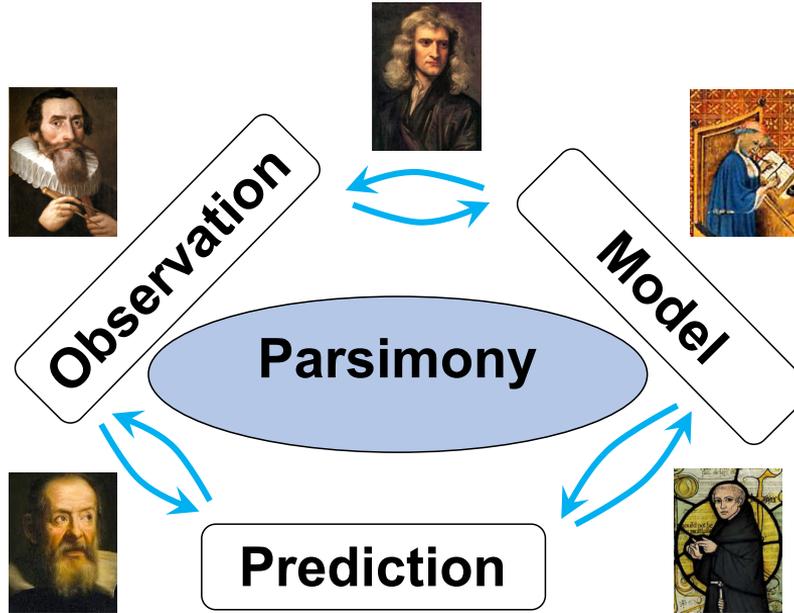
backward propagation of error



The circular learning approach where a model is automatically learned from data through validation and optimization is not new



Scientific method vs. machine learning method

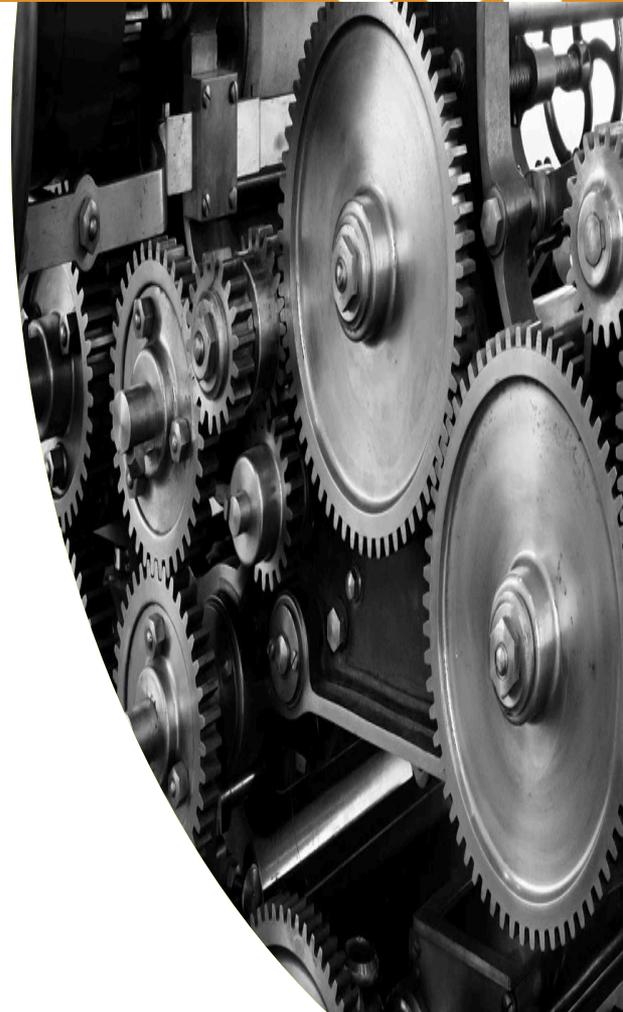


- What is novel is the 'bottom up' approach that does not need human intuition for the formulation of the model
- Parsimony is no longer central
- A very large number of models are automatically generated and the model selection part becomes central

What can machines learn?

Machine learning, artificial intelligence and deep learning had an enormous success in recent years

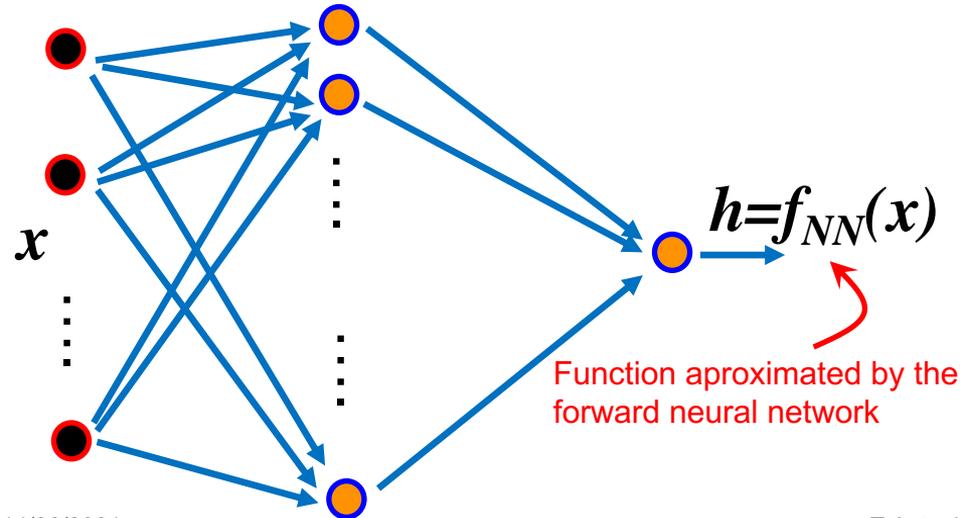
Their successful application has been mainly in the domain of image recognition/manipulation and games



Universal approximation theorem

Cybenko 1989, Hornik 1991

A forward neural network with more than one layer can approximate any function as far as the network has a large enough number of neurons



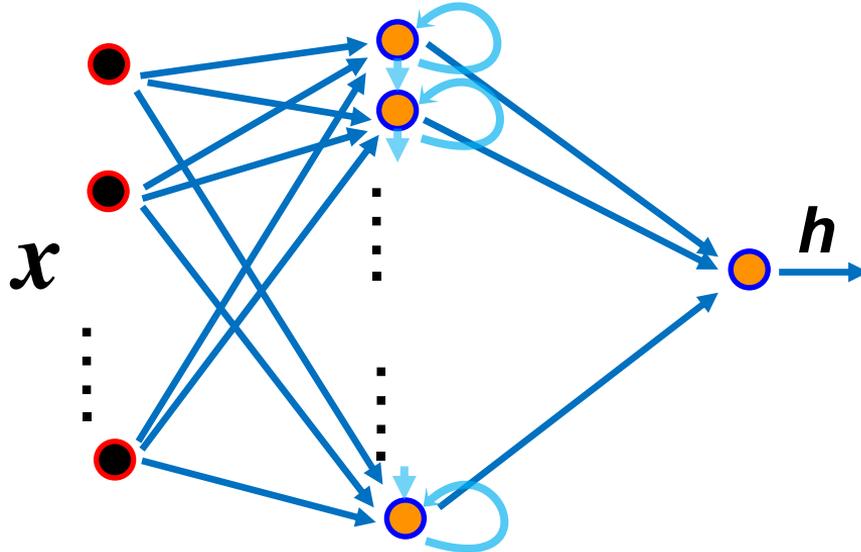
Reference function

$$|f_{NN}(x) - f(x)| < \epsilon$$



Recurrent neural networks are Turing complete

By adding loops (within a layer or backwards) a recurrent neural network is Turing complete and therefore it can perform any computation



Can machines learn markets?

Markets

are gigantic
computation
environments
where machines
and humans
interact to calculate
the price of things



Markets are complex systems

- Markets are complex systems where a large and heterogeneous number of variables interact within an intricate system of relations
- In markets trades are executed at speeds ranging from nanoseconds to years, this are 10^{15} order of magnitude (comparable to our distance to Proxima centaury in meters)
- Market variables have statistical properties that are non-normal with fat-tails power law distributions
- Markets adapt and change, they are not stationary

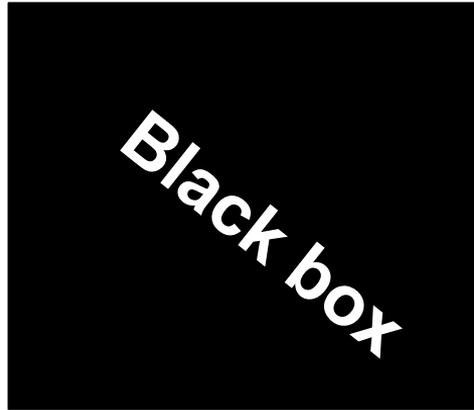
Can machines learn markets?

Has a machine that can trade successfully learned something about the market?

Can we define what does it mean learn in the case of markets?

A pragmatic perspective:
can machines learn to automate tasks so far performed by humans?

What can we learn from machines?



Deep learning models return equations or algorithms; these are the same kind of outputs of human-made models

They are however extremely complicated being

- **high-dimensional and**
- **non-linear**

High dimensional spaces are very different from the low-dimensional space we live in.

In particular, the subdivision of high-dimensional spaces into regions (the basins of optimal solutions) have non-intuitive properties

- Most of the volume of the region is near the surface
- The number of neighboring regions grows at least exponentially with dimension
- The number of interfaces between regions grow combinatorically with dimension
- Any 'gradient descent optimizer' will be always and unavoidably stuck in some saddle-point frontier

Tomaso Aste and , Denis Weaire, *The pursuit of perfect packing*. CRC Press, 2008.

T. Aste, and N. Rivier. "Random cellular froths in spaces of any dimension and curvature." *Journal of Physics A: Mathematical and General* 28, no. 5 (1995): 1381.

T Aste, "Dynamical partitions of space in any dimension." *Journal of Physics A: Mathematical and General* 31, no. 43 (1998): 8577.

Linear solutions of linear problems are neat and unique, they have convenient properties:

1. Small changes in the input produce proportionally small changes in the output
2. Small changes in the solution structure or parameters also produce small changes in the output
3. Approximate solutions with similar errors are similar

Non-linear solutions are very different

Non-linear solutions solutions are very different

1. Small changes in the input can produce very large changes in the output
2. Small changes in the solution structure or parameters can produce very large changes in the output
3. Approximate solutions with similar errors can be completely different and there is a combinatorial large number of them

Intriguingly, the way deep learning systems are trained (specifically methods such as: data sampling, drop off, knockout, data augmentations, loss functions, weight regularization...) seem to overcome several of these issues, at least in some cases. We have a lot of to learn about how these systems discover approximate solutions

What machines can learn about our complex world - and what can we learn from them?

They learn approximate solutions that are high-dimensional and non-linear

The structure of the solution tells us very little about the model

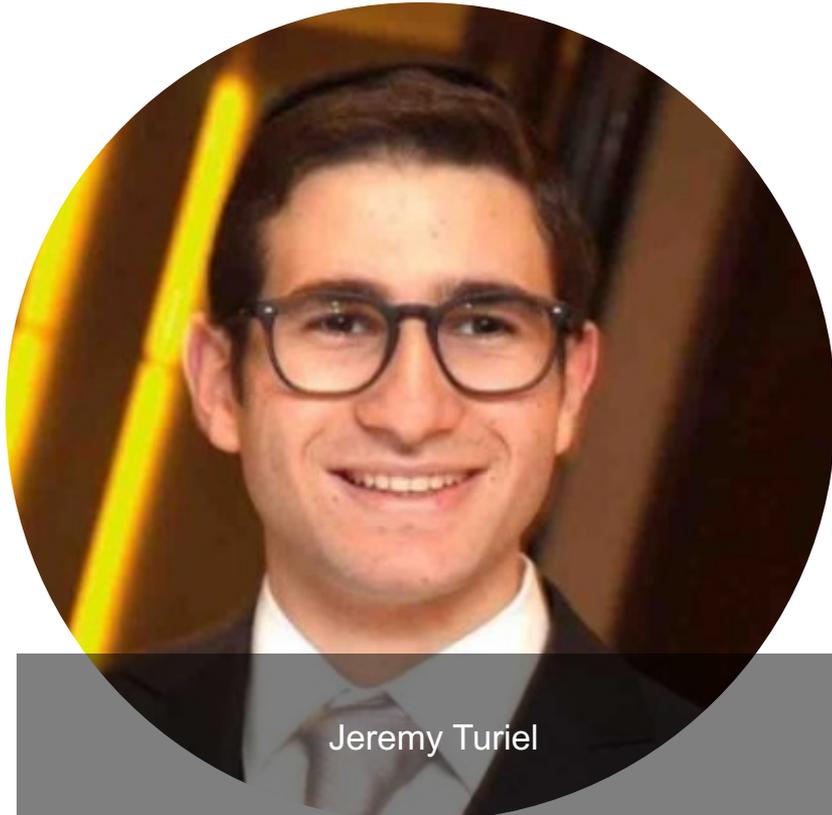
However, if the learning process is done properly, these solutions are quite robust and can work also in circumstances different from the training examples

Deep learning the limit Order Book

Experiment 1

A comparative perspective

<https://arxiv.org/pdf/2007.07319.pdf>



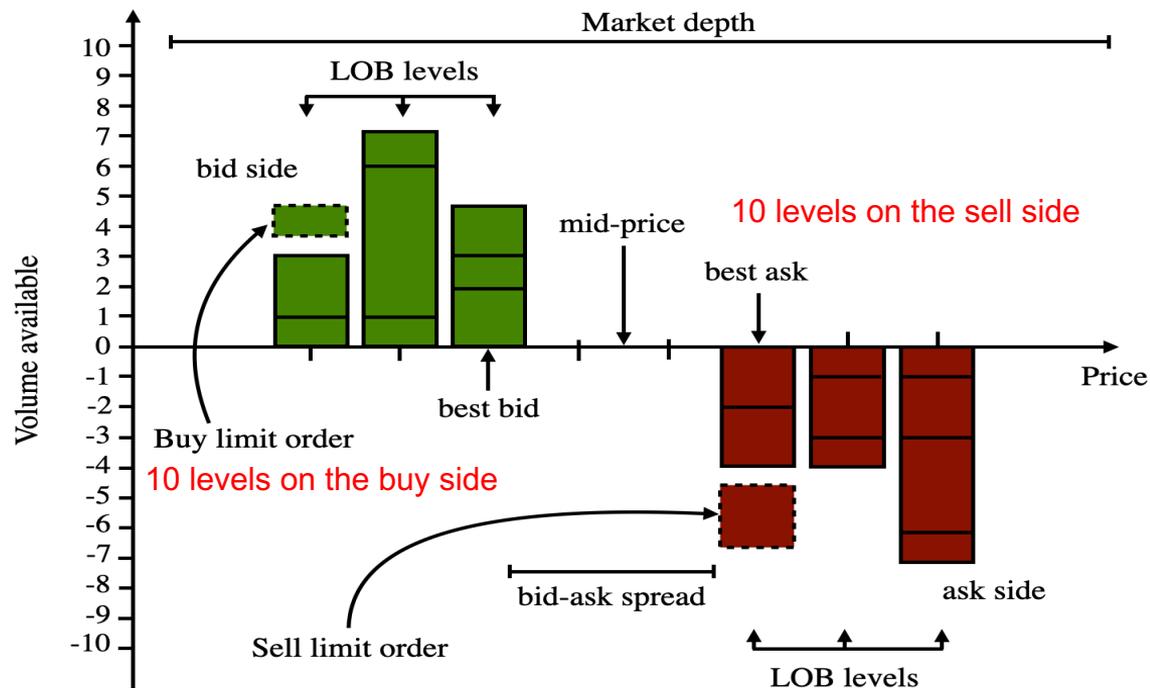
Jeremy Turiel



Antonio Briola

The limit order book (LOB)

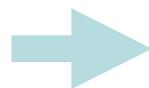
The Limit Order Book (LOB) is a self-organizing system where a large number of players interact with offers and bids and eventually agree on a transaction price



About 10 transactions per second for liquid assets on NASDAQ

INPUT

LOB prices and
volumes for 10
previous ticks
400-dimensions



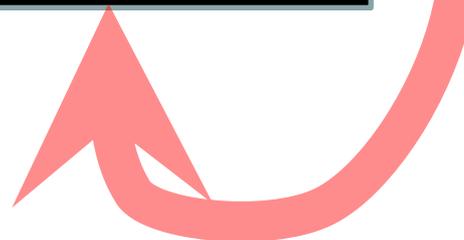
**Machine
learning
system
(7 models)**

Testing with 6 million
datapoints



OUTPUT

Transaction
price range at
a given
horizon
3-dimensions



Training with 18 million examples

$$[(p, v)_0^a, (p, v)_0^b, (p, v)_1^a, (p, v)_1^b, \dots, (p, v)_{10}^a, (p, v)_{10}^b]$$

20 levels of LOB

10 previous ticks (x_{t-9}, \dots, x_t)

$x_t = 40$ -dimensions

Overall a vector of 400-dimensions

Data:

NASDAQ, Intel Corp. (INTC) LOB data

Period:

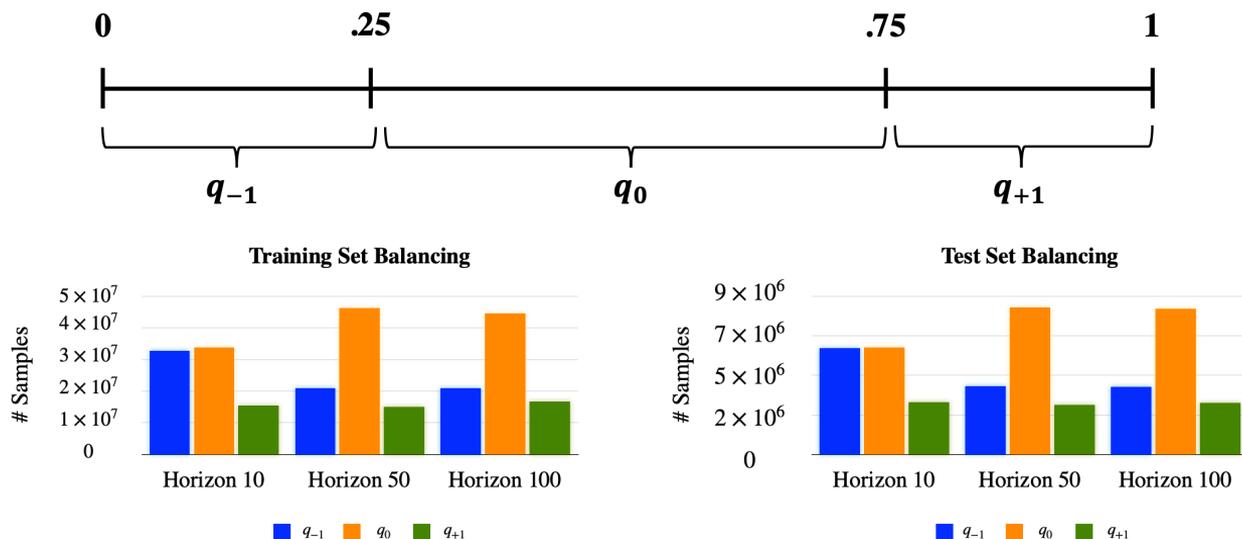
training 3 months:

04-02-2019 to 31-05-2019 -> ~ 18 million transactions

testing 1 month:

03-06-2019 to 28-06-2019 -> ~ 6 million transactions

Probability of transaction after 10, 50 or 100 ticks at a price within a given quantile range established from the training set

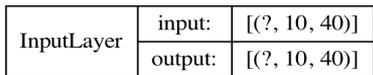


Horizons
10, 50, 100 ticks

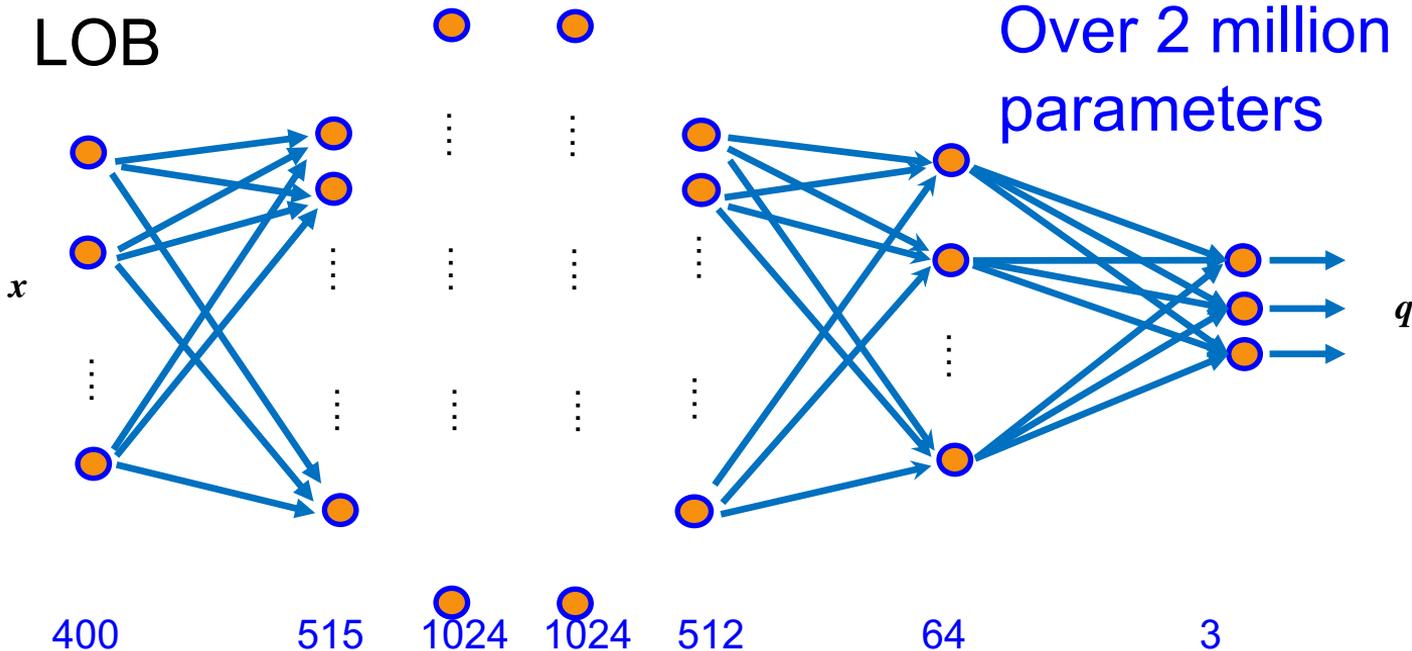
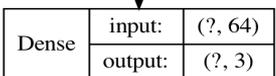
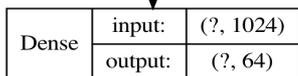
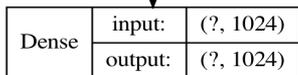
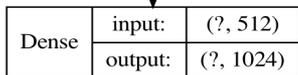
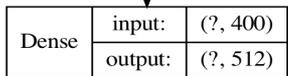
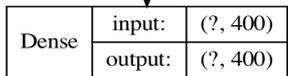
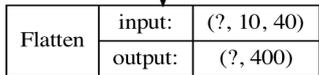
Can a Multilayer Perceptron learn the LOB efficiently?
And how will it do with respect to other simpler and more complicated models?

We investigate and compare 7 different models with increasing levels of complexity

1. **Random Model** – uniform probability outcome prediction
2. **Naive Model** – output most represented in training
3. **Logistic Regression** – simple perceptron
4. **Multilayer Perceptron** – deep learning model
5. **Shallow LSTM** – deep learning with memory
6. **Self-Attention LSTM** – deep learning with memory & loop
7. **CNN-LSTM** – state of the art deep learning

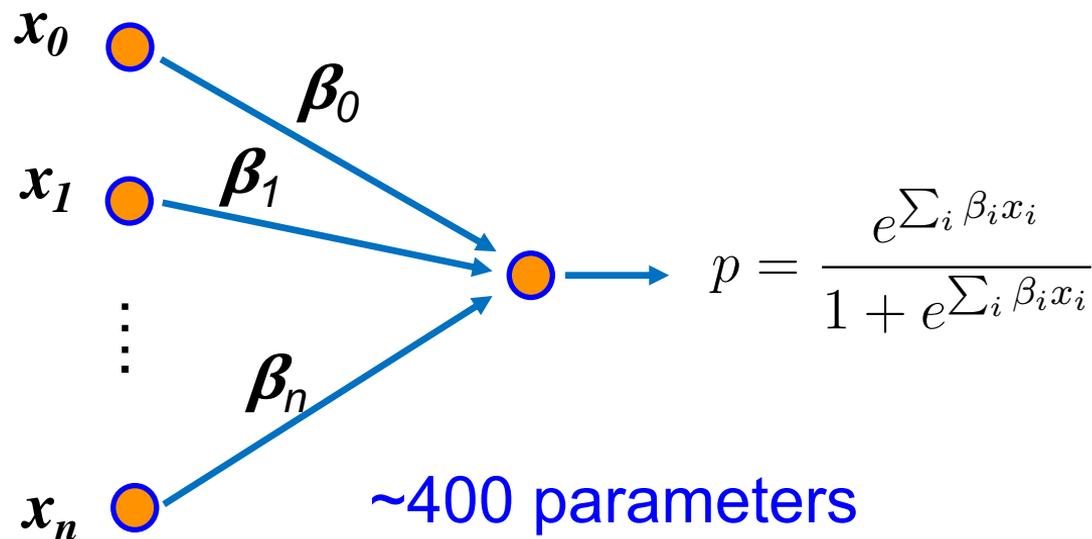


← LOB

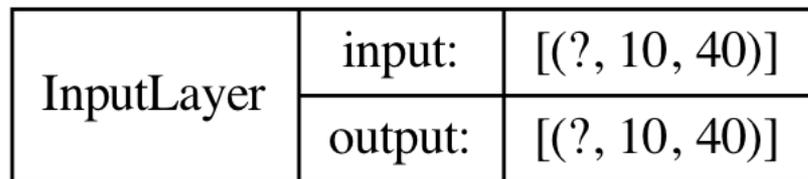


→ Quantile probability

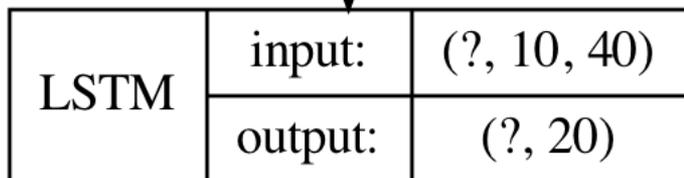
Layer (type)	Output Shape	Param #
input_1 (InputLayer)	[(None, 10, 40)]	0
flatten (Flatten)	(None, 400)	0
dense (Dense)	(None, 400)	160400
dense_1 (Dense)	(None, 512)	205312
dense_2 (Dense)	(None, 1024)	525312
dense_3 (Dense)	(None, 1024)	1049600
dense_4 (Dense)	(None, 64)	65600
dense_5 (Dense)	(None, 3)	195
Total params: 2,006,419. Trainable params: 2,006,419, Non-trainable params: 0		



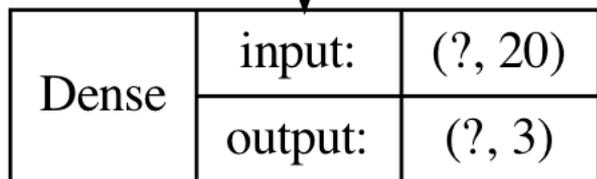
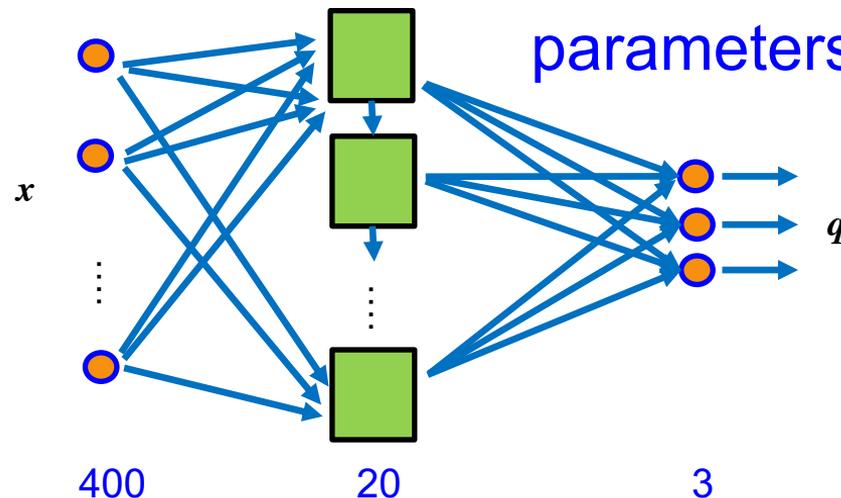
One-layer perceptron ($x_0=1$)



← LOB



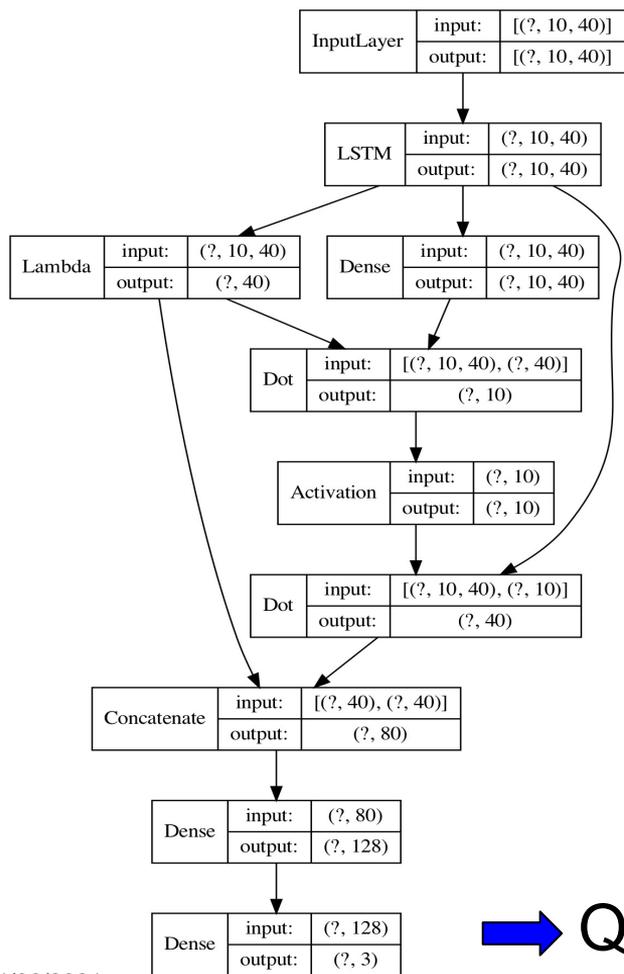
About 5 thousands parameters



→ Quantile probability

```

Model: "model"
-----
Layer (type)      Output Shape      Param #
-----
input_1 (InputLayer)  [(None, 10, 40)]      0
lstm (LSTM)         (None, 20)         4880
dense (Dense)       (None, 3)           63
-----
Total params: 4,943
Trainable params: 4,943
Non-trainable params: 0
    
```



← LOB

→ Quantile probability

About 25 thousands parameters

Layer (type)	Output Shape	Param #	Connected to
input_1 (InputLayer)	[(None, 10, 40)]	0	
lstm (LSTM)	(None, 10, 40)	12960	input_1[0][0]
attention_score_vec (Dense)	(None, 10, 40)	1600	lstm[0][0]
last_hidden_state (Lambda)	(None, 40)	0	lstm[0][0]
attention_score (Dot)	(None, 10)	0	attention_score_vec[0][0] last_hidden_state[0][0]
attention_weight (Activation)	(None, 10)	0	attention_score[0][0]
context_vector (Dot)	(None, 40)	0	lstm[0][0] attention_weight[0][0]
attention_output (Concatenate)	(None, 80)	0	context_vector[0][0] last_hidden_state[0][0]
attention_vector (Dense)	(None, 128)	10240	attention_output[0][0]
dense (Dense)	(None, 3)	387	attention_vector[0][0]

Total params: 25,187
 Trainable params: 25,187
 Non-trainable params: 0

The models range a parameter space dimension from zero to one million

Model	Input	Number of layers (*)	Number of parameters(**)	Learning rate	Training epochs
Random	-	0	0	—	-
Naive	LOB	1	1	—	-
Logistic Reg.	LOB	2	4×10^2	10^{-3}	30
MLP	LOB	7	2.0×10^6	10^{-3}	30
Shallow LSTM	LOB	3	4.9×10^3	10^{-3}	30
Self-Attention LSTM	LOB	4	2.5×10^4	10^{-3}	30
CNN-LSTM	LOB	28	6.0×10^4	10^{-3}	30

Table 1: Summary of the inputs and hyperparameters used in the models in this article. (*)The number of layers includes the input and the output layer. (**)The number of parameters is approximated to the nearest order of magnitude and truncated for readability.

We tested performances of the prediction for each of the quantile regions by computing:

Precision, Recall and F-measure

In order to correct for class imbalance we also tested:

balanced Accuracy, weighted Precision, weighted Recall and weighted F-score

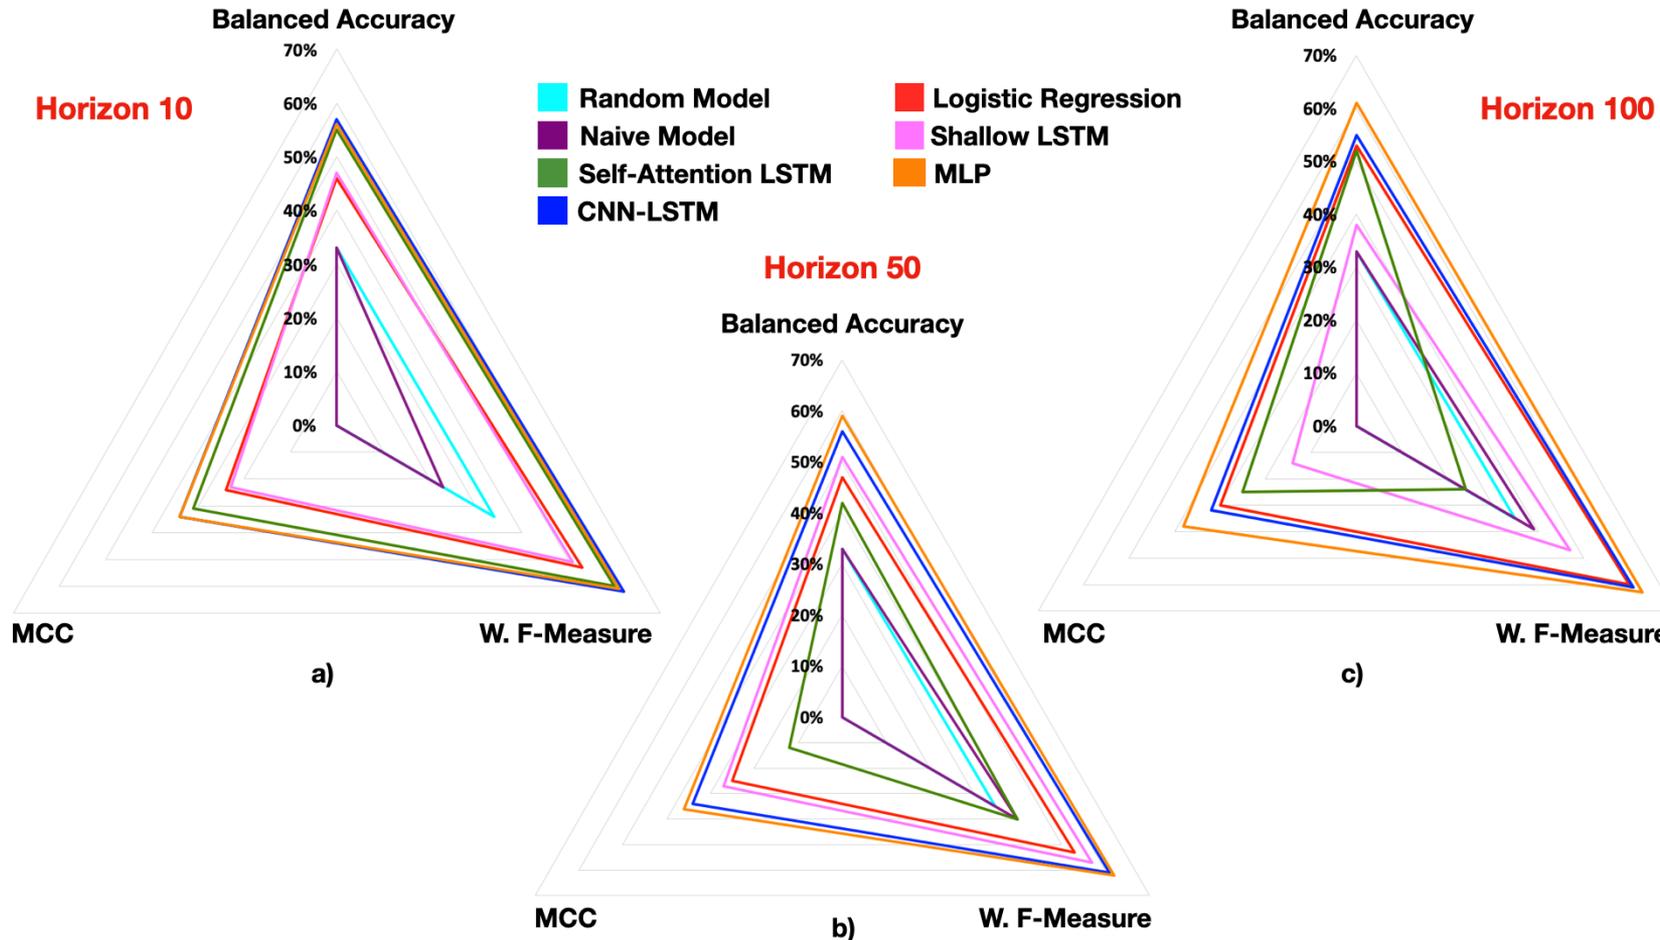
Furthermore, two multi-class correlation metrics between forecasted and real labels computed:

Matthews Correlation Coefficient (MCC) and Cohen's Kappa

Performances at three horizons: 10, 50, 100 ticks

	Random Model			Naive Model			Logistic Regression			Shallow LSTM			Self-Attention LSTM			CNN-LSTM			Multilayer Perceptron		
	H10	H50	H100	H10	H50	H100	H10	H50	H100	H10	H50	H100	H10	H50	H100	H10	H50	H100	H10	H50	H100
Balanced Accuracy	0.33	0.33	0.33	0.33	0.33	0.33	0.46	0.47	0.53	0.47	0.51	0.38	0.55	0.42	0.52	0.57	0.56	0.55	0.56	0.59	0.61
Weighted Precision	0.41	0.41	0.41	0.16	0.30	0.30	0.54	0.56	0.62	0.58	0.57	0.65	0.61	0.50	0.47	0.62	0.61	0.61	0.62	0.62	0.63
Weighted Recall	0.33	0.33	0.33	0.40	0.55	0.54	0.59	0.59	0.61	0.58	0.58	0.57	0.61	0.45	0.34	0.62	0.62	0.62	0.62	0.63	0.63
Weighted F-Measure	0.34	0.35	0.35	0.23	0.40	0.39	0.53	0.53	0.60	0.51	0.57	0.47	0.60	0.40	0.24	0.62	0.61	0.61	0.61	0.62	0.63
Precision quantile [0, 0.25]	0.26	0.26	0.26	0	0	0	0.57	0.57	0.57	0.57	0.56	0.58	0.60	0.37	0.55	0.59	0.59	0.59	0.59	0.59	0.59
Precision quantile [0.25, 0.75]	0.55	0.55	0.55	0.40	0.55	0.54	0.59	0.60	0.63	0.59	0.62	0.56	0.62	0.55	0.52	0.65	0.64	0.63	0.64	0.66	0.67
Precision quantile [0.75, 1]	0.20	0.20	0.20	0	0	0	0.31	0.38	0.58	0.57	0.43	0.97	0.57	0.52	0.26	0.57	0.57	0.57	0.59	0.58	0.57
Recall quantile [0, 0.25]	0.33	0.33	0.33	0	0	0	0.55	0.59	0.59	0.06	0.62	0.22	0.35	0.81	0.62	0.54	0.55	0.53	0.60	0.59	0.60
Recall quantile [0.25, 0.75]	0.33	0.33	0.33	1	1	1	0.82	0.81	0.76	0.85	0.69	0.93	0.76	0.44	0.01	0.71	0.72	0.74	0.74	0.70	0.68
Recall quantile [0.75, 1]	0.33	0.33	0.33	0	0	0	0	0	0.23	0.51	0.23	0	0.53	0.004	0.92	0.46	0.42	0.39	0.34	0.46	0.54
F-Measure quantile [0, 0.25]	0.29	0.29	0.29	0	0	0	0.56	0.57	0.58	0.11	0.59	0.31	0.44	0.503	0.58	0.57	0.57	0.56	0.59	0.59	0.59
F-Measure quantile [0.25, 0.75]	0.42	0.42	0.42	0.57	0.71	0.70	0.70	0.70	0.68	0.69	0.65	0.70	0.68	0.489	0.02	0.68	0.68	0.68	0.69	0.68	0.67
F-Measure quantile [0.75, 1]	0.25	0.25	0.25	0	0	0	0	0	0.33	0.54	0.30	0	0.55	0.009	0.40	0.51	0.48	0.46	0.43	0.51	0.56
MCC	0	0	0	0	0	0	0.24	0.25	0.30	0.23	0.27	0.14	0.31	0.120	0.25	0.34	0.34	0.32	0.34	0.36	0.38
Cohen's Kappa	0	0	0	0	0	0	0.21	0.22	0.30	0.20	0.27	0.10	0.30	0.105	0.16	0.34	0.33	0.32	0.33	0.36	0.38

Best



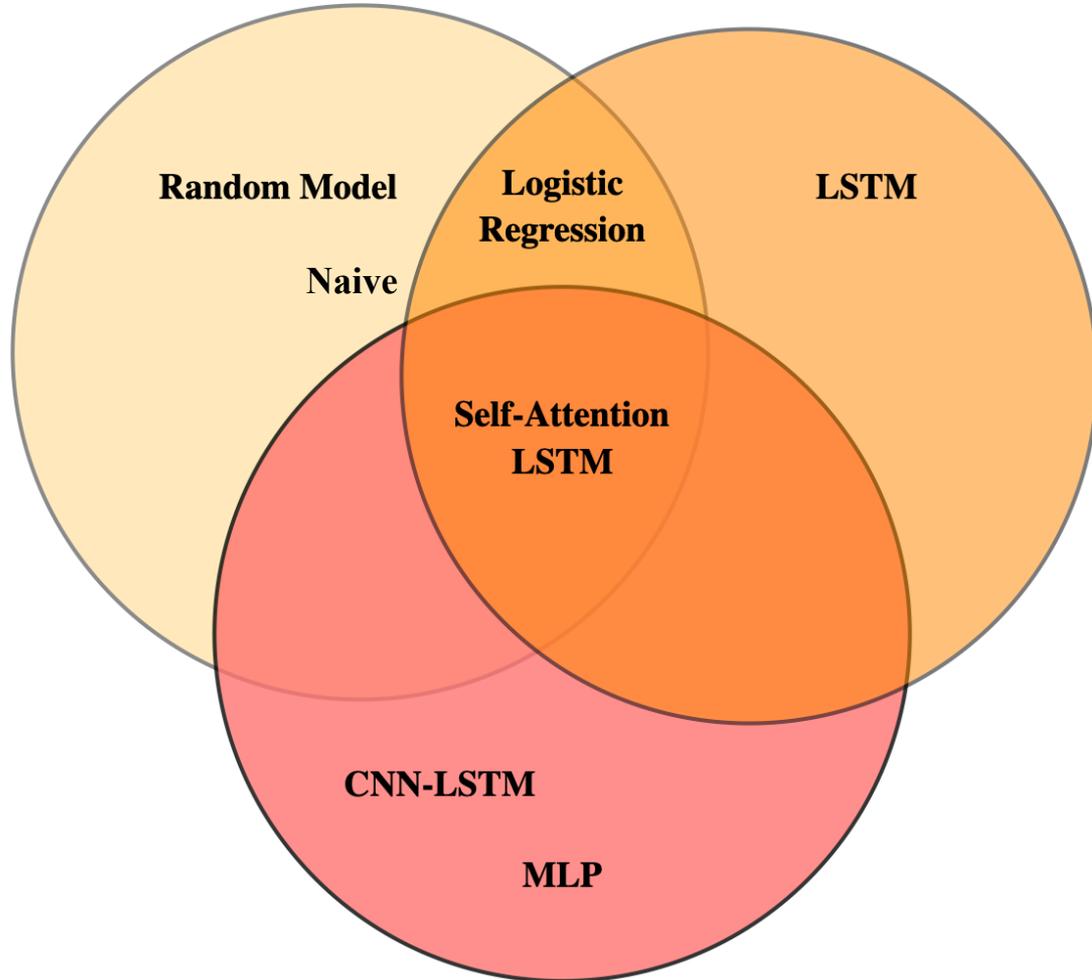
Bayesian correlated t-test from MCC measure

$H_{\Delta\tau} \Delta\tau = 10$	$H_{\Delta\tau} \Delta\tau = 50$	$H_{\Delta\tau} \Delta\tau = 100$
Multilayer Perceptron CNN - LSTM	Multilayer Perceptron CNN - LSTM	Multilayer Perceptron CNN - LSTM
Self-Attention LSTM	Shallow LSTM Multinomial Logistic Regression	Multinomial Logistic Regression
Shallow LSTM Multinomial Logistic Regression	Self-Attention LSTM	Self-Attention LSTM
Naive Model Random Model	Naive Model Random Model	Shallow LSTM
		Random Model

Bayesian correlated t-test from F measure

$H_{\Delta\tau} \Delta\tau = 10$	$H_{\Delta\tau} \Delta\tau = 50$	$H_{\Delta\tau} \Delta\tau = 100$
Multilayer Perceptron CNN - LSTM Self-Attention LSTM	Multilayer Perceptron CNN - LSTM	Multilayer Perceptron
Shallow LSTM Multinomial Logistic Regression	Shallow LSTM	CNN - LSTM Multinomial Logistic Regression
Naive Model	Multinomial Logistic Regression	Shallow LSTM
Random Model	Self-Attention LSTM	Self-Attention LSTM
	Naive Model	Naive Model
	Random Model	Random Model

- Multilayer perceptron is the best performing
- CNN-LSTM is second-best with comparable performances
- Logistic regression has good performances comparable with LSTM and self-attention LSTM
- Naïve and Random are worst



One can attempt to cluster the models together based on their performances

One can note that memory/recurrent loops (LSTM) play little role in performances. Most of the information is processed forwards from the LOB input

Deep learning the limit Order Book

Experiment 2

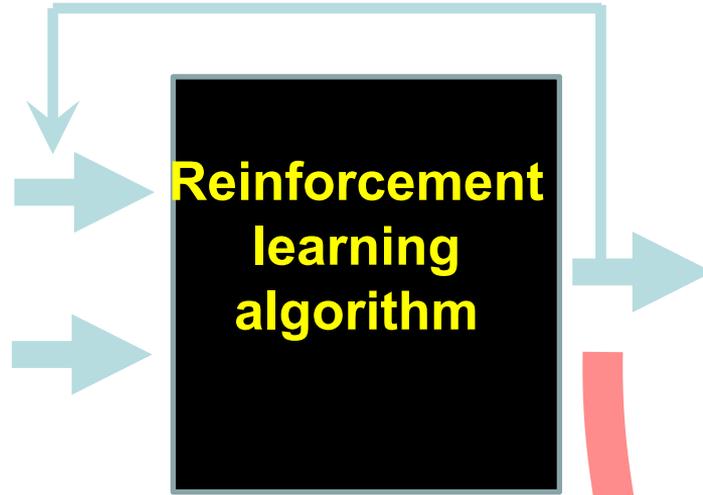
Reinforcement Learning

<https://arxiv.org/abs/2101.07107>

Reward

INPUT

LOB prices and volumes for 10 previous ticks
400-dimensions



Testing with 6 million examples

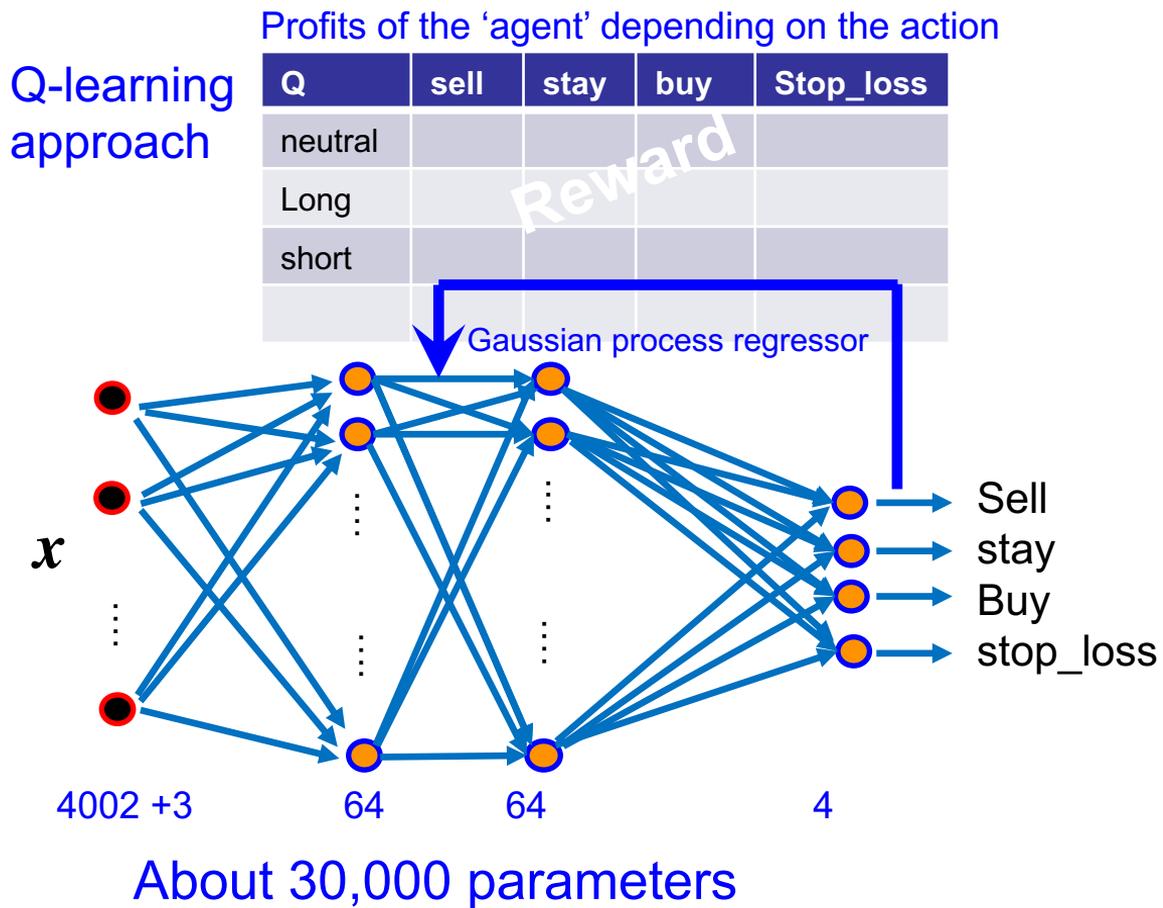
OUTPUT

Action:

1. Sell,
2. stay,
3. buy,
4. stop_loss

Training with ~6 millions selected examples

The algorithm buys or sells or holds one unit of Intel Corporation stock (INTC) on NASDAQ during the month of June 2019. It is trained during the previous 3 months.

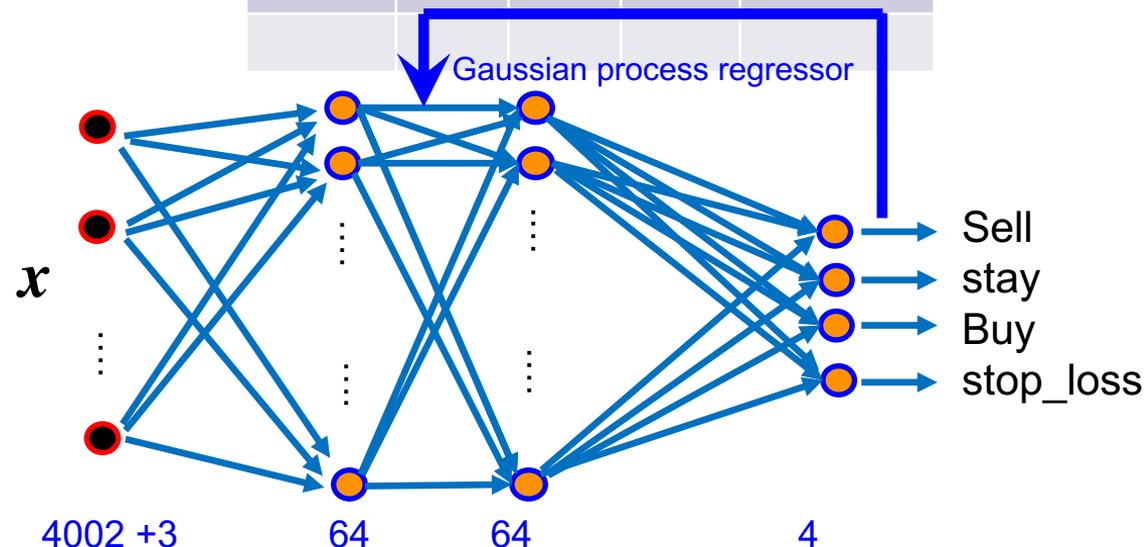


We train on 25 selected significant events containing about **250,000** ticks

- Training on 60 days
04/02/2019-30/04/2019
- Validation on 22 days
01/05/2019-31/05/2019
- Testing on 20 days
03/06/2019-28/06/2019

Profits of the 'agent' depending on the action

Q	sell	stay	buy	Stop_loss
neutral				
Long				
short				



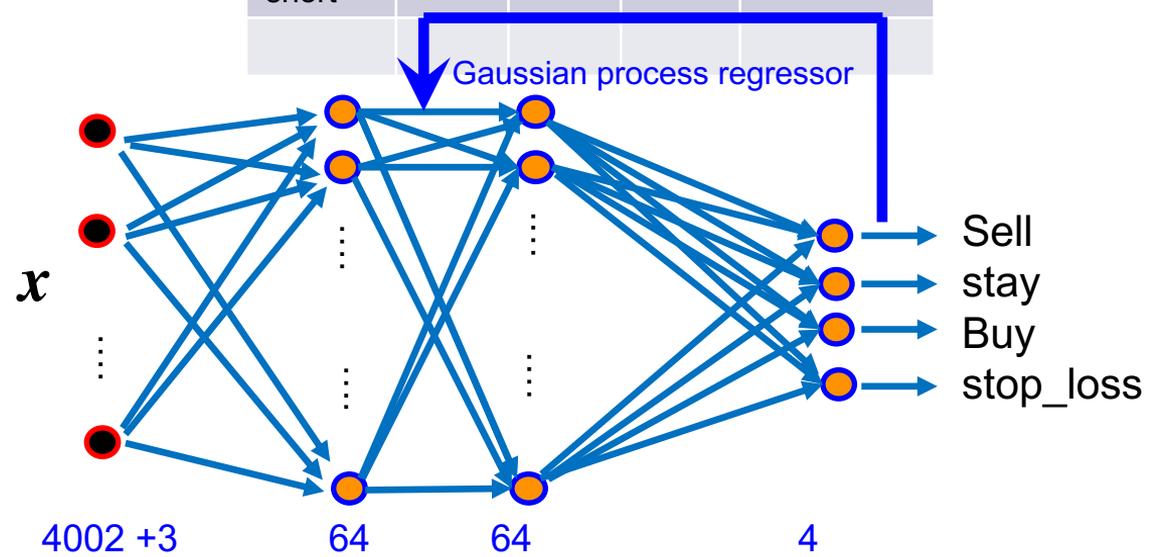
About 30,000 parameters

We test for three sets of input information; all have the full LOB (400 dimension) plus:

1. State of the agent
2. State of the agent & market price minus price paid for the unit (mark to market profit)
3. State of the agent & price paid for the unit (mark to market) & bid-ask spread

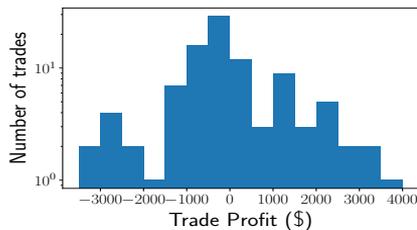
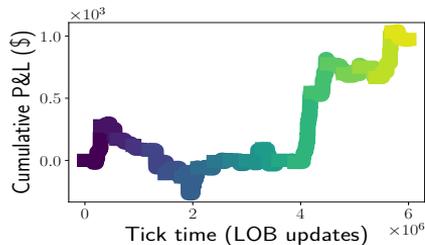
Profits of the 'agent' depending on the action

Q	sell	stay	buy	Stop_loss
neutral				
Long				
short				

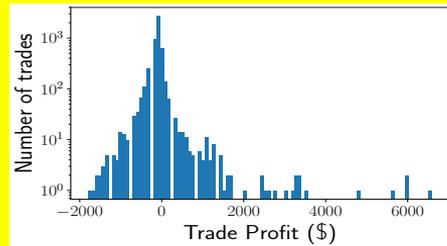
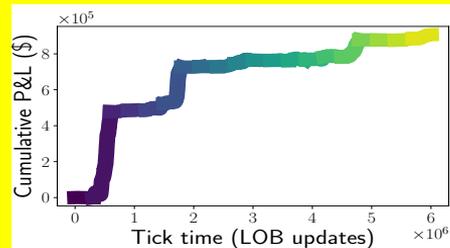


About 30,000 parameters

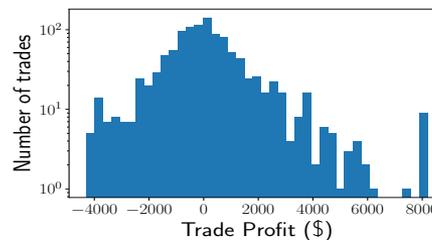
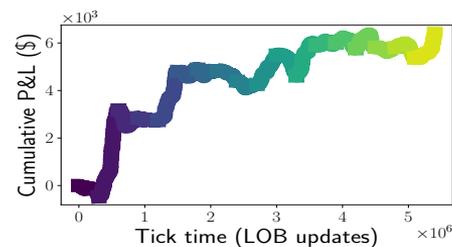
1. LOB + State of the agent



2. LOB + State of the agent & price paid for the unit (mark to market)



3. LOB + State of the agent & price paid for the unit (mark to market) & bid-ask spread



What the agent learned?

- The agent increases profits by about 100 folds by trading 10 to 100 times more often using information about the reference unit price (mark to market). The agent learns to increase profits while increasing trading frequency despite the bid-ask spread cost
- Risk is reduced considerably
- The extra information on the bid-ask spread does not increase performances

Conclusions

- Artificial intelligence is providing increasingly powerful new instruments
- It is almost a surprise that a complicated self-trained machine, such as the Multi Layer Perception, can learn something about the price formation mechanism on the LOB
- It is almost a surprise that a self-learning agent can discover trading strategies
- Results are very good but not ground-breaking. Are we at the beginning or at the end of this journey?

Links and references

LINKS

FCA Group Page:

<http://fincomp.cs.ucl.ac.uk/introduction/>

My articles:

<https://scholar.google.co.uk/citations?user=27pUbTUAAA&hl=en>

Software:

TMFG & Clique Forests

<https://github.com/cvborkulo/NetworkComparisonTest/pull/5>

<https://uk.mathworks.com/matlabcentral/fileexchange/56444-tmfg>

RELEVANT PAPERS

Aste, Tomaso. "What machines can learn about our complex world-and what can we learn from them?." Available at SSRN 3797711 (2021).

Briola, Antonio, Jeremy Turiel, Riccardo Marcaccioli, and Tomaso Aste. "Deep Reinforcement Learning for Active High Frequency Trading." arXiv preprint arXiv:2101.07107 (2021).

Briola, Antonio, Jeremy Turiel, and Tomaso Aste. "Deep Learning Modeling of the Limit Order Book: A Comparative Perspective." Available at SSRN 3714230 (2020).