Model Risk Management for Deep Learning and Alpha Strategies

Ben Steiner
BNP Paribas Asset Management

The views expressed are those of the presenter and not his current employer

“... the engineers couldn't explain why the AI was executing the trades it was making. The creation was such a black box that even its creators didn't understand how it worked…”

“The Massive Hedge Fund Betting on AI”, Bloomberg, 9/27/2017

Agenda

1. **Introductory concepts**
   a. What is model risk management?
   b. Machine Learning for Alpha Strategies
2. **Challenges of Deep Learning**
   a. Non-stationarity
   b. Interpretation
   c. Learning what we already know
3. **Model Risk Management for Deep Learning and Alpha Strategies**
   a. Backtest evaluation
   b. Ongoing monitoring

Agenda

1. **Introductory concepts**
   a. What is model risk management?
   b. Machine Learning for Alpha Strategies
2. **Challenges of Deep Learning**
   a. Non-stationarity
   b. Interpretation
   c. Learning what we already know
3. **Model Risk Management for Deep Learning and Alpha Strategies**
   a. Backtest evaluation
   b. Ongoing monitoring

Definitions

A ***Model*** is a simplification of the real world into mathematical equations to forecast some future behavior.

Model Risk comes from either incorrect models (fundamental errors) or models being misapplied (incorrect or inappropriate usage).

Risk Management is the process of identifying, analysing and controlling uncertainty around objectives.

Model Risk Management is the understanding, analysing and controlling the risk inherent in using models.

1. Conceptual Soundness
2. Implementation Validation
3. Ongoing Monitoring

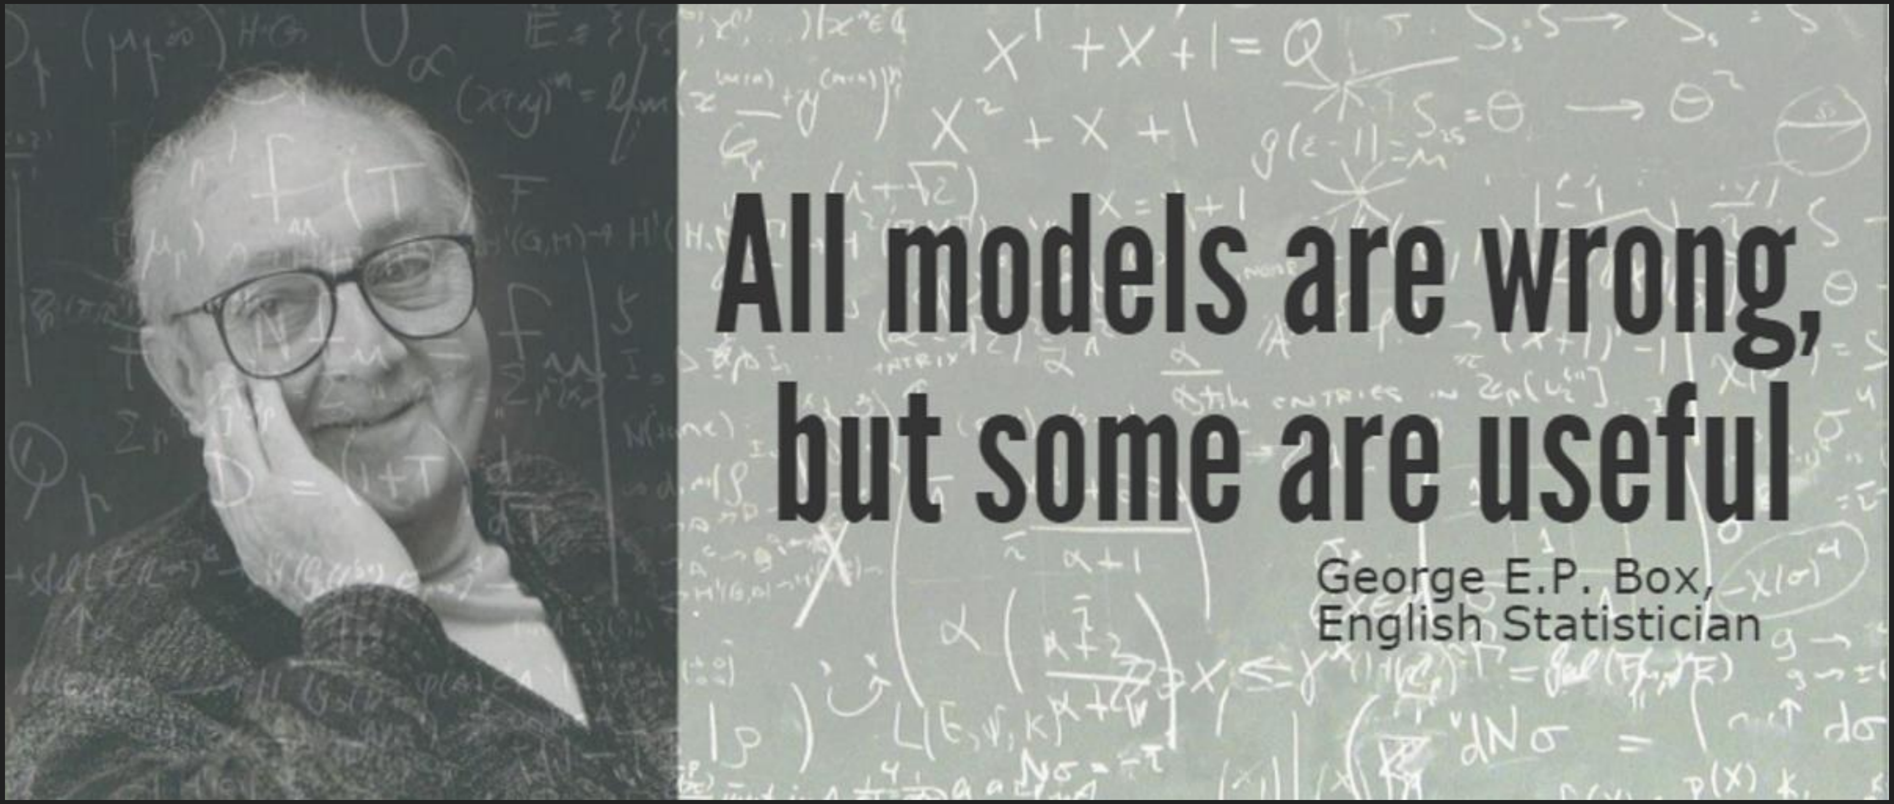
All models are wrong,
but some are useful
George E.P. Box,
English Statistician

Definitions

Machine Learning uses algorithms to learn from data without relying on rules-based programming

Deep Learning maps inputs to outputs using multiple layers of nonlinear processing units

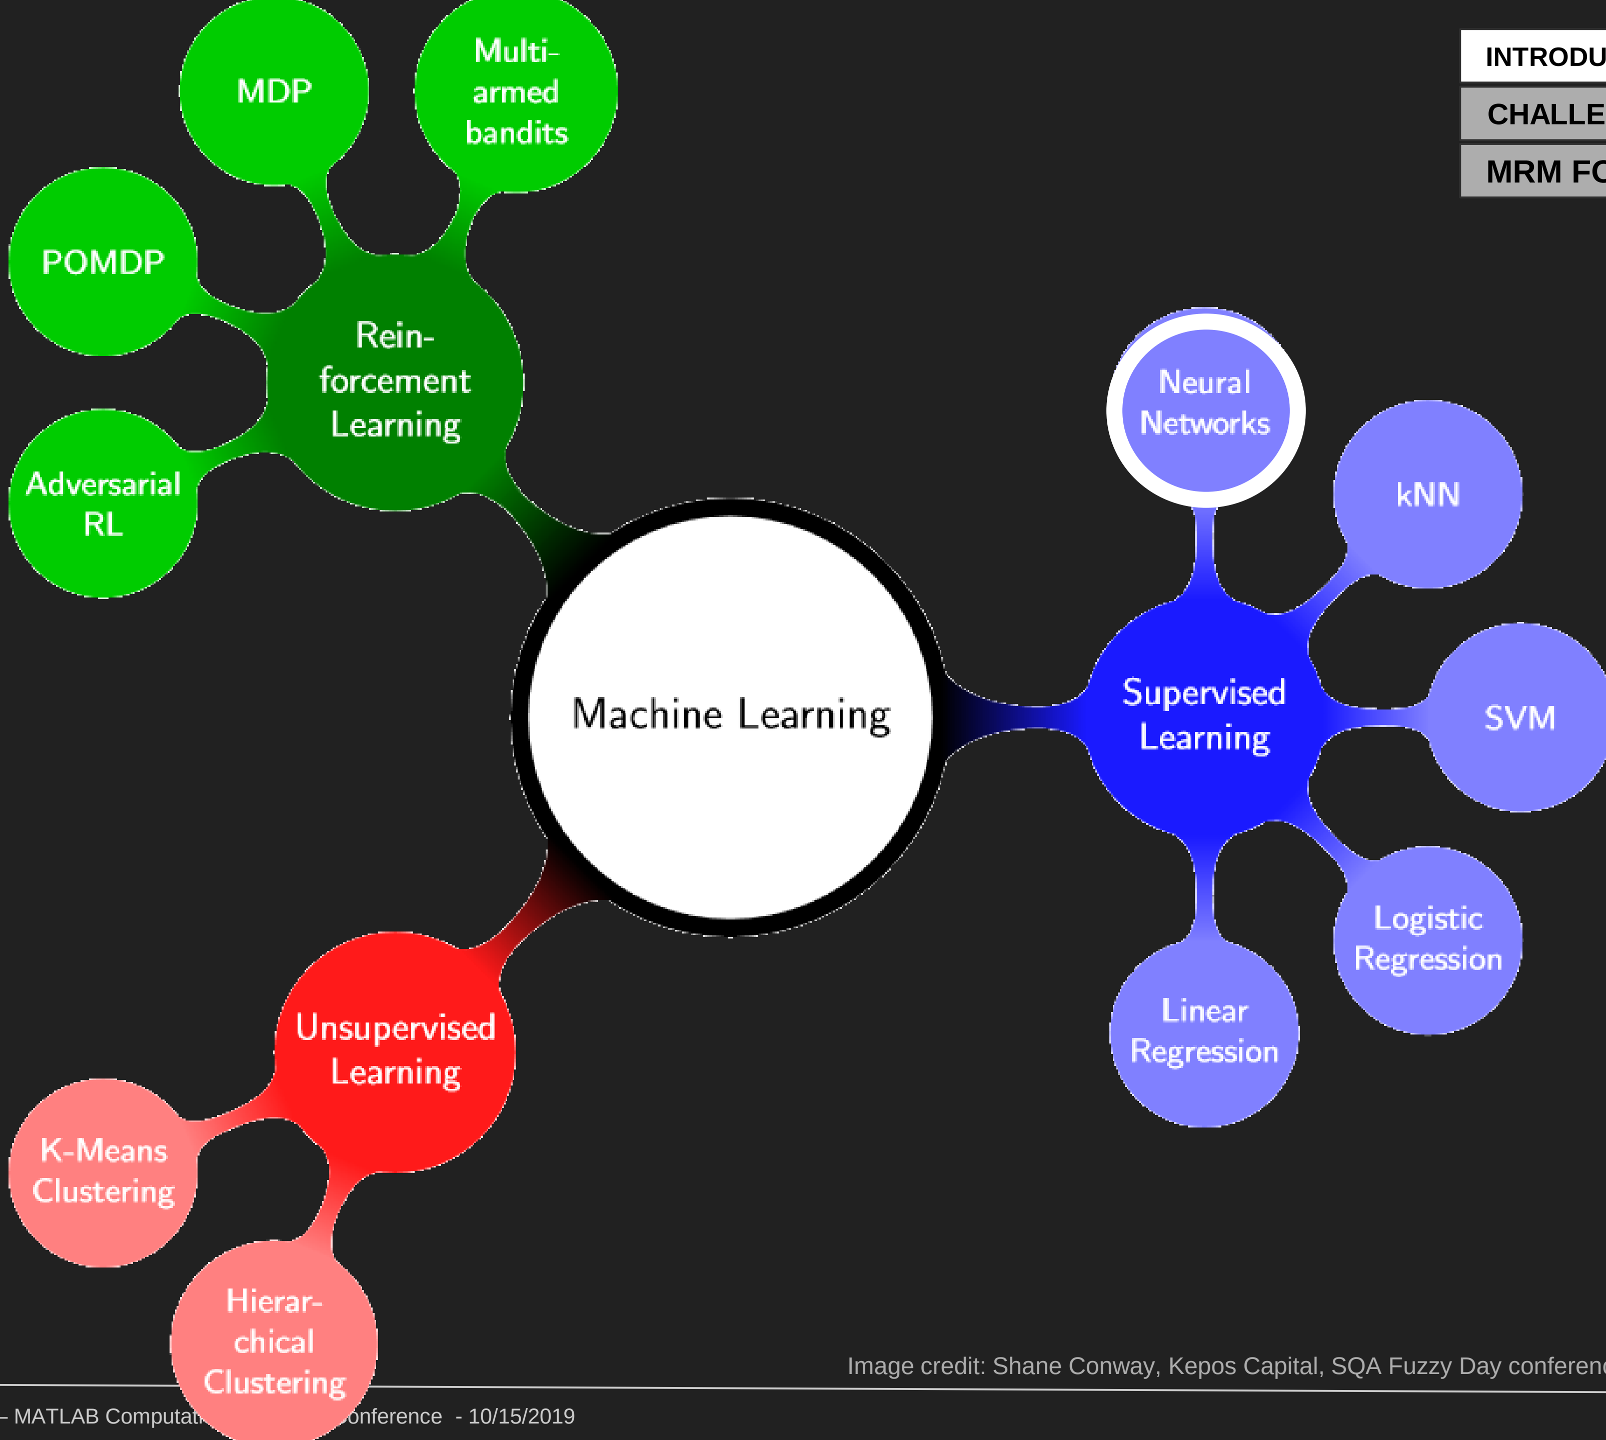

Image credit: Shane Conway, Kepos Capital, SQA Fuzzy Day conference 2017

Is Machine Learning still “a model”?

Statistical Science
2001, Vol. 16, No. 3, 199–231

Statistical Modeling: The Two Cultures

Leo Breiman

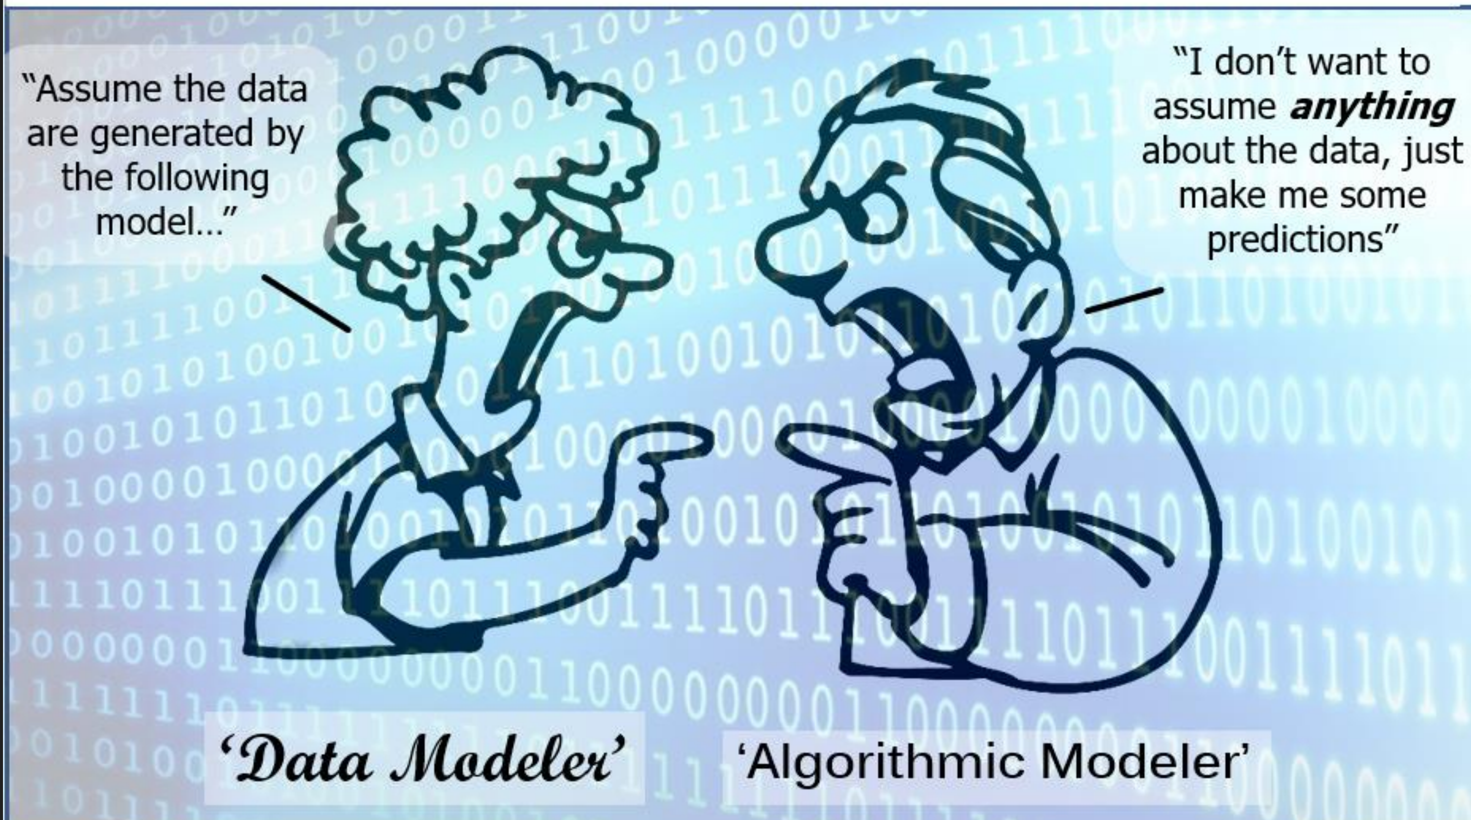

Deep Learning basics

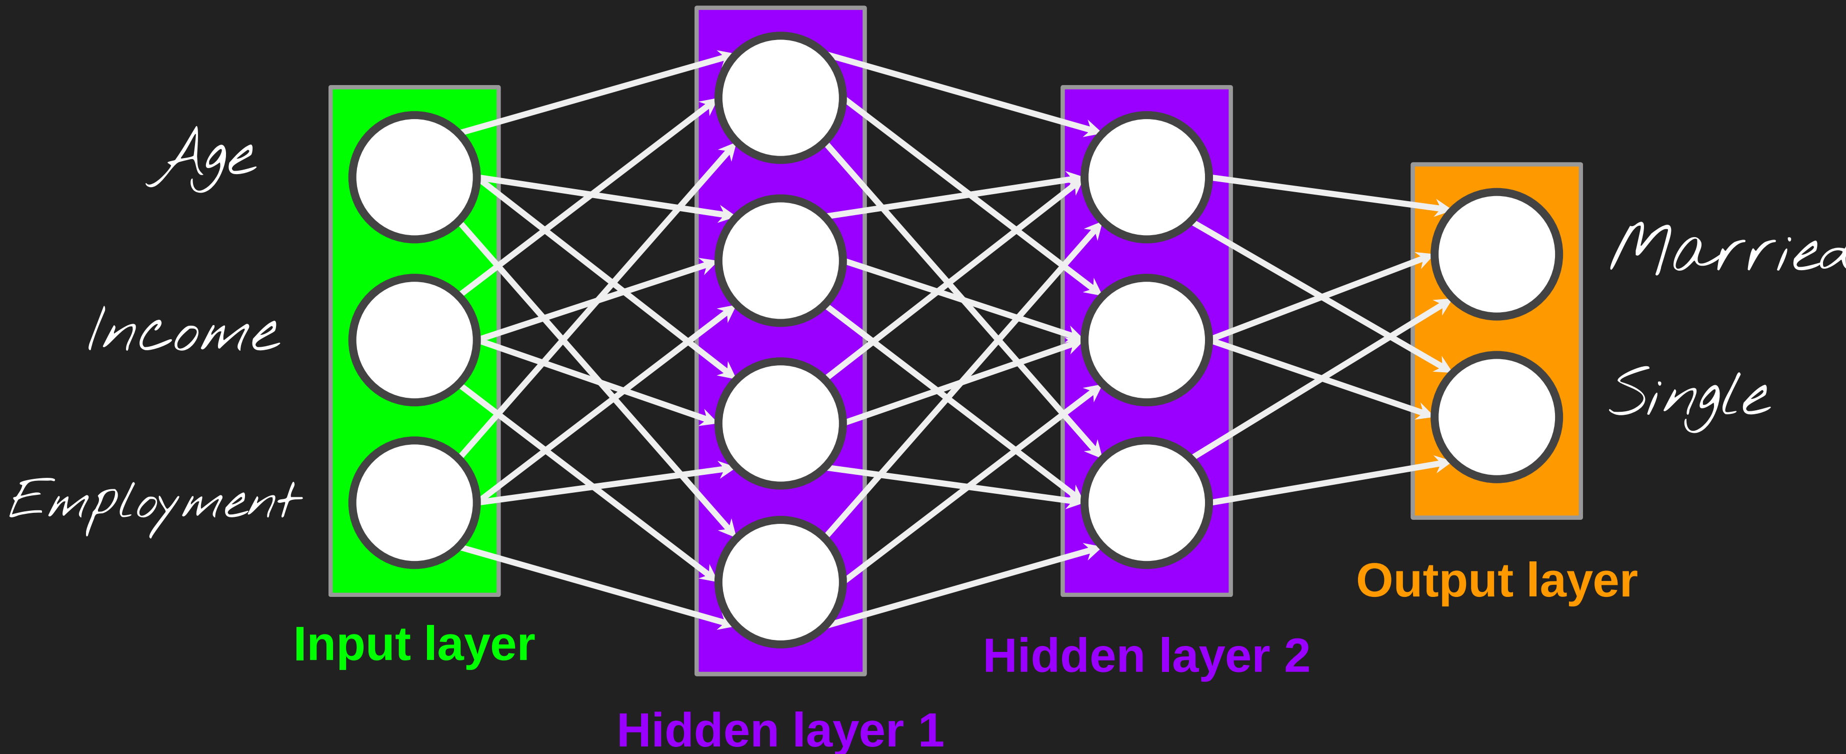

Input layer: features (or attributes)

Hidden layers: Bias and weights

Output layer: target variables (or responses)

Example: Arno Candel, H20.ai

Deep Learning basics

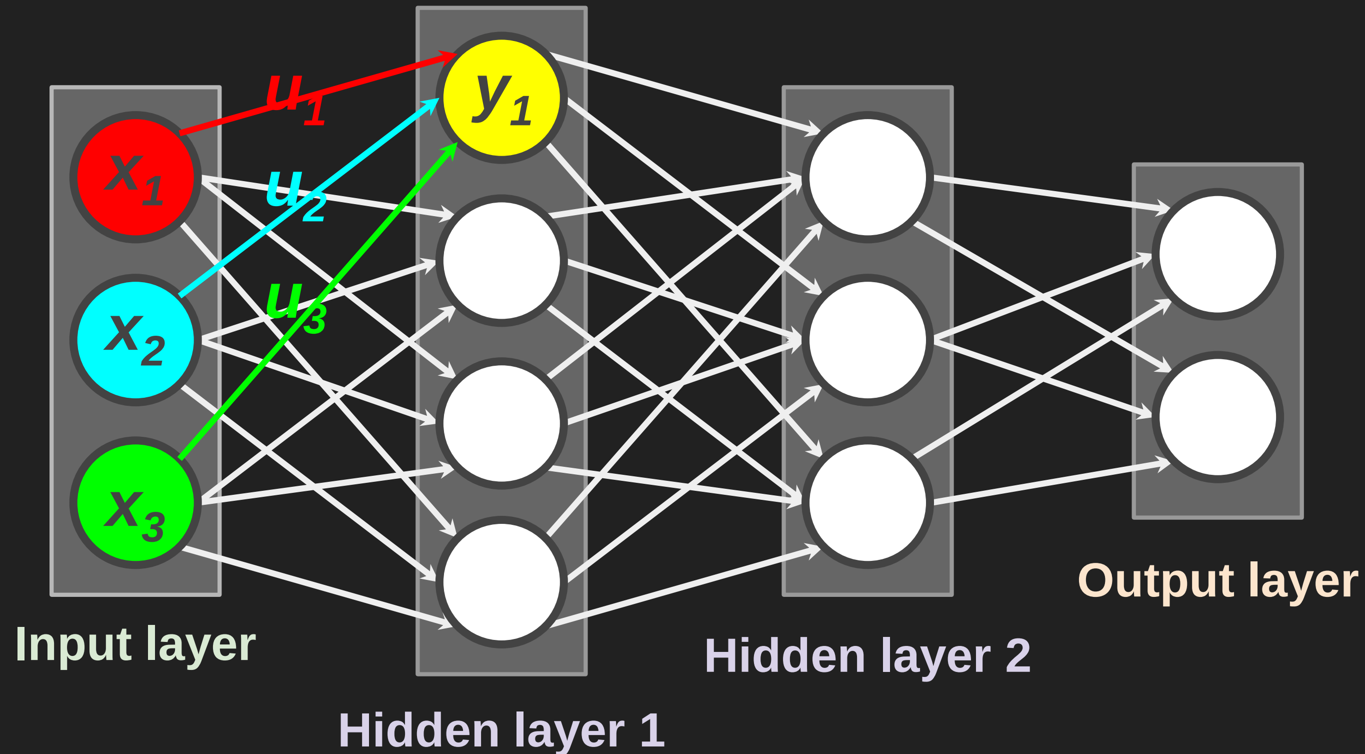

Neurons activate each other via **weighted sums**

$$y_1 = f((\, x_1u_1 + x_2u_2 + x_3u_3 \,) + b_1\,)$$

Non-linear activation function tanh: 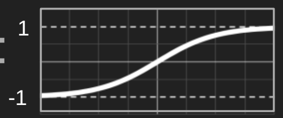

or rectilinear: 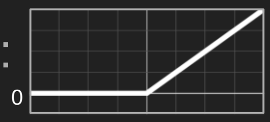

Deep Learning basics

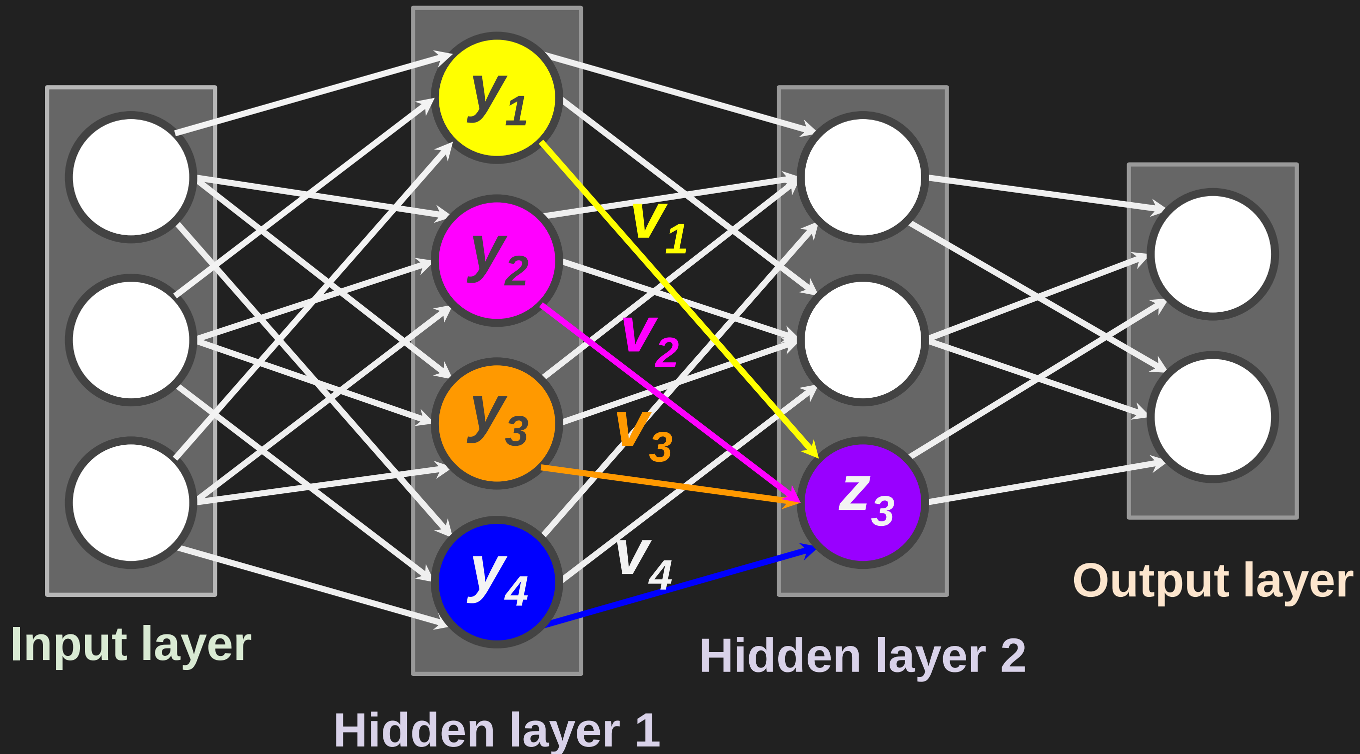

Neurons activate each other via **weighted sums**

$$z_3 = f((\, y_1 v_1 + y_2 v_2 + y_3 v_3 + y_4 v_4 \,) + c_3 \,)$$

Non-linear activation function tanh:

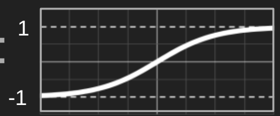

or rectilinear:

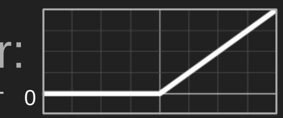

Deep Learning basics

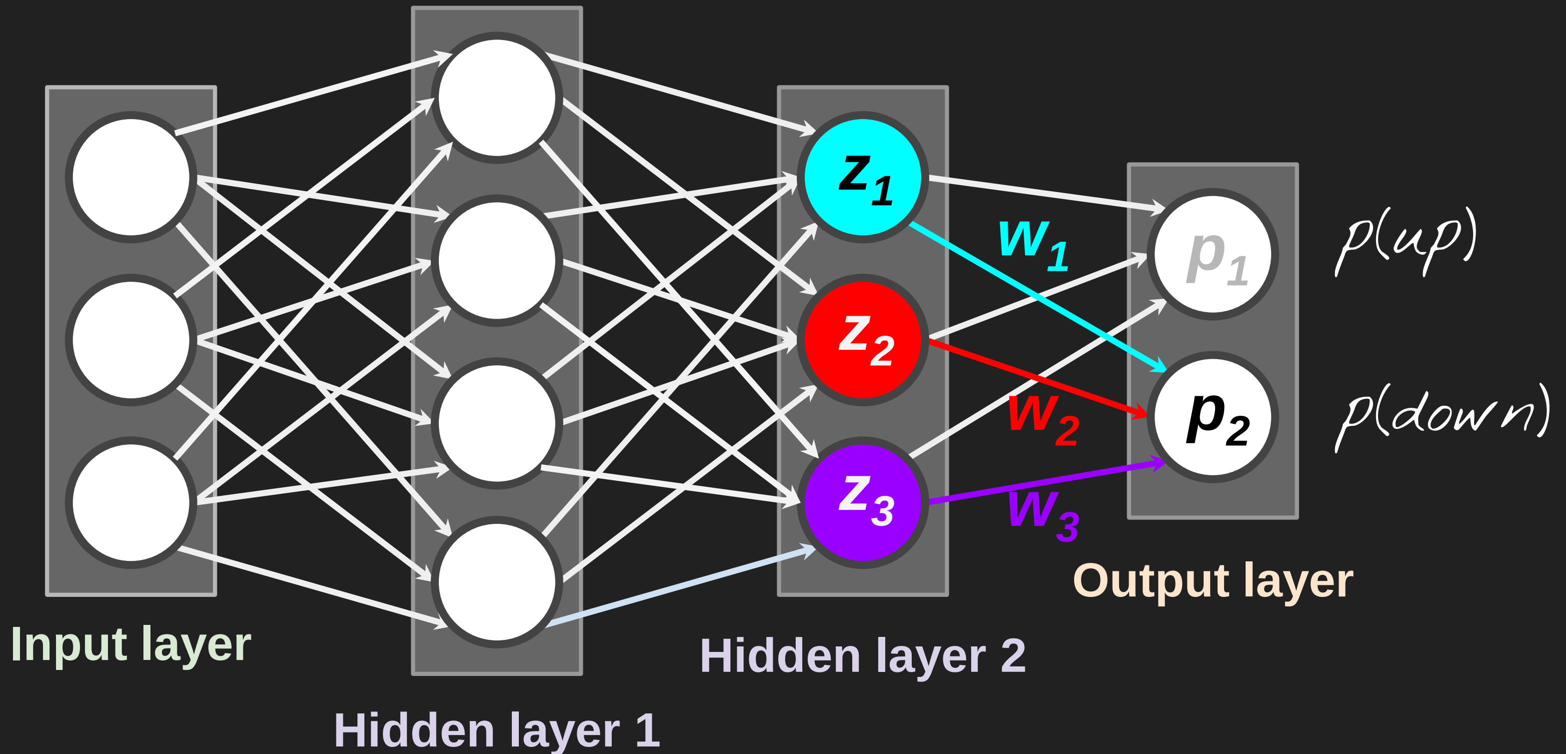

Neurons activate each other via **weighted sums**

$$p_2 = f((z_1 w_1 + z_2 w_2 + z_3 w_3) + d_2)$$

Non-linear activation function: softmax

Agenda

1. **Introductory concepts**
   a. What is model risk management?
   b. Machine Learning for Alpha Strategies
2. **Challenges of Deep Learning**
   a. Non-stationarity
   b. Interpretation
   c. Learning what we already know
3. **Model Risk Management for Deep Learning and Alpha Strategies**
   a. Backtest evaluation
   b. Ongoing monitoring

“A model may be reasonable, but the world itself may be unstable. What’s a good model today may be inappropriate tomorrow”

Emanuel Derman, 1996, GS research paper on model risk

Challenge 1: Non-Stationarity

Stationarity (or Nonstationarity) is a property of an underlying model and **not of observed data.**

A single realization from a stationary stochastic process can appear indistinguishable from a nonstationary deterministic process

'Change' is in the timeframe of the beholder

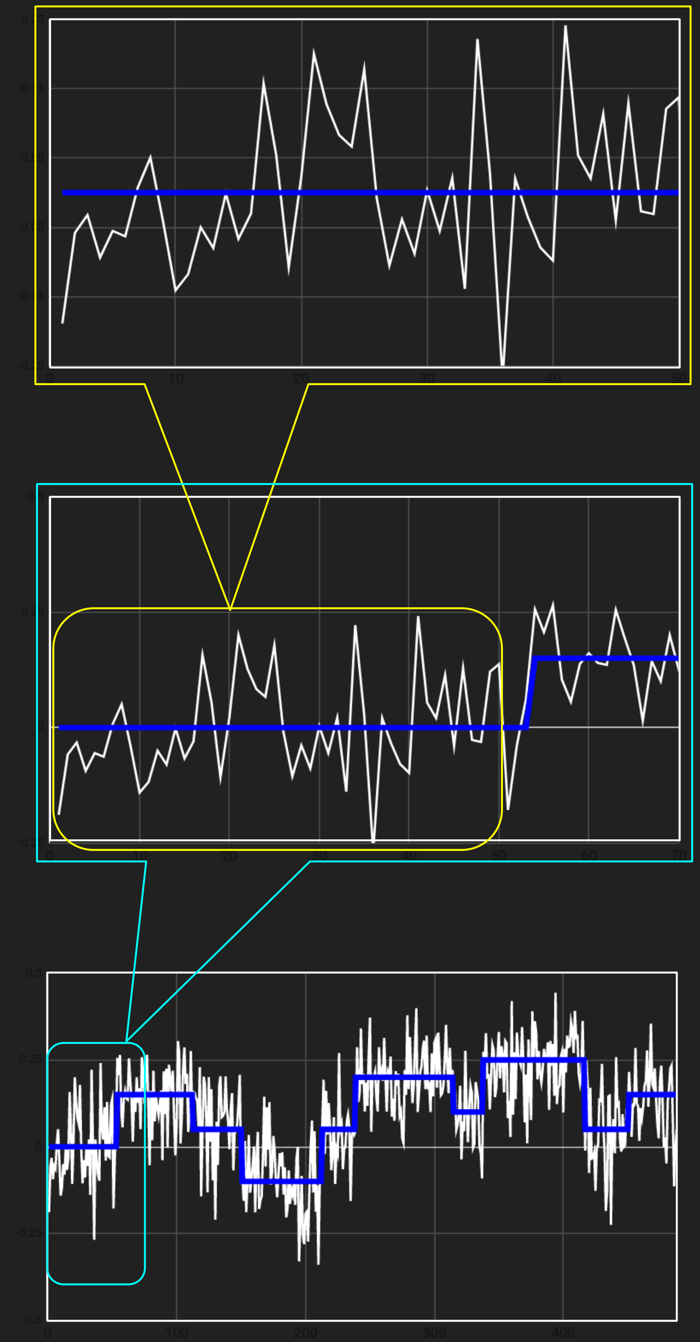

Challenge 1: Non-Stationarity

One problem, many names:

- Concept Drift
- Covariate Shift
- Dataset shift
- Nonstationarity
- Source component shift
- Temporal evolution

CONCEPT DRIFT

When the statistical properties of the target variable, which the model is trying to predict, change over time in **unforeseen** ways.

The **unforeseen** substitution of one data source $\boldsymbol{S_1}$ (with underlying probability distribution $\boldsymbol{\Pi_{S1}}$), with another source $\boldsymbol{S_2}$ (with distribution $\boldsymbol{\Pi_{S2}}$)

Not Concept Drift

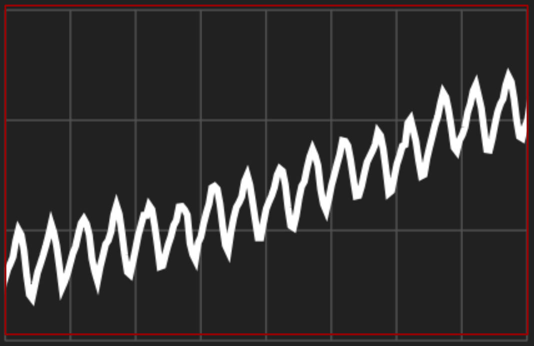

Different Types of Concept Drift

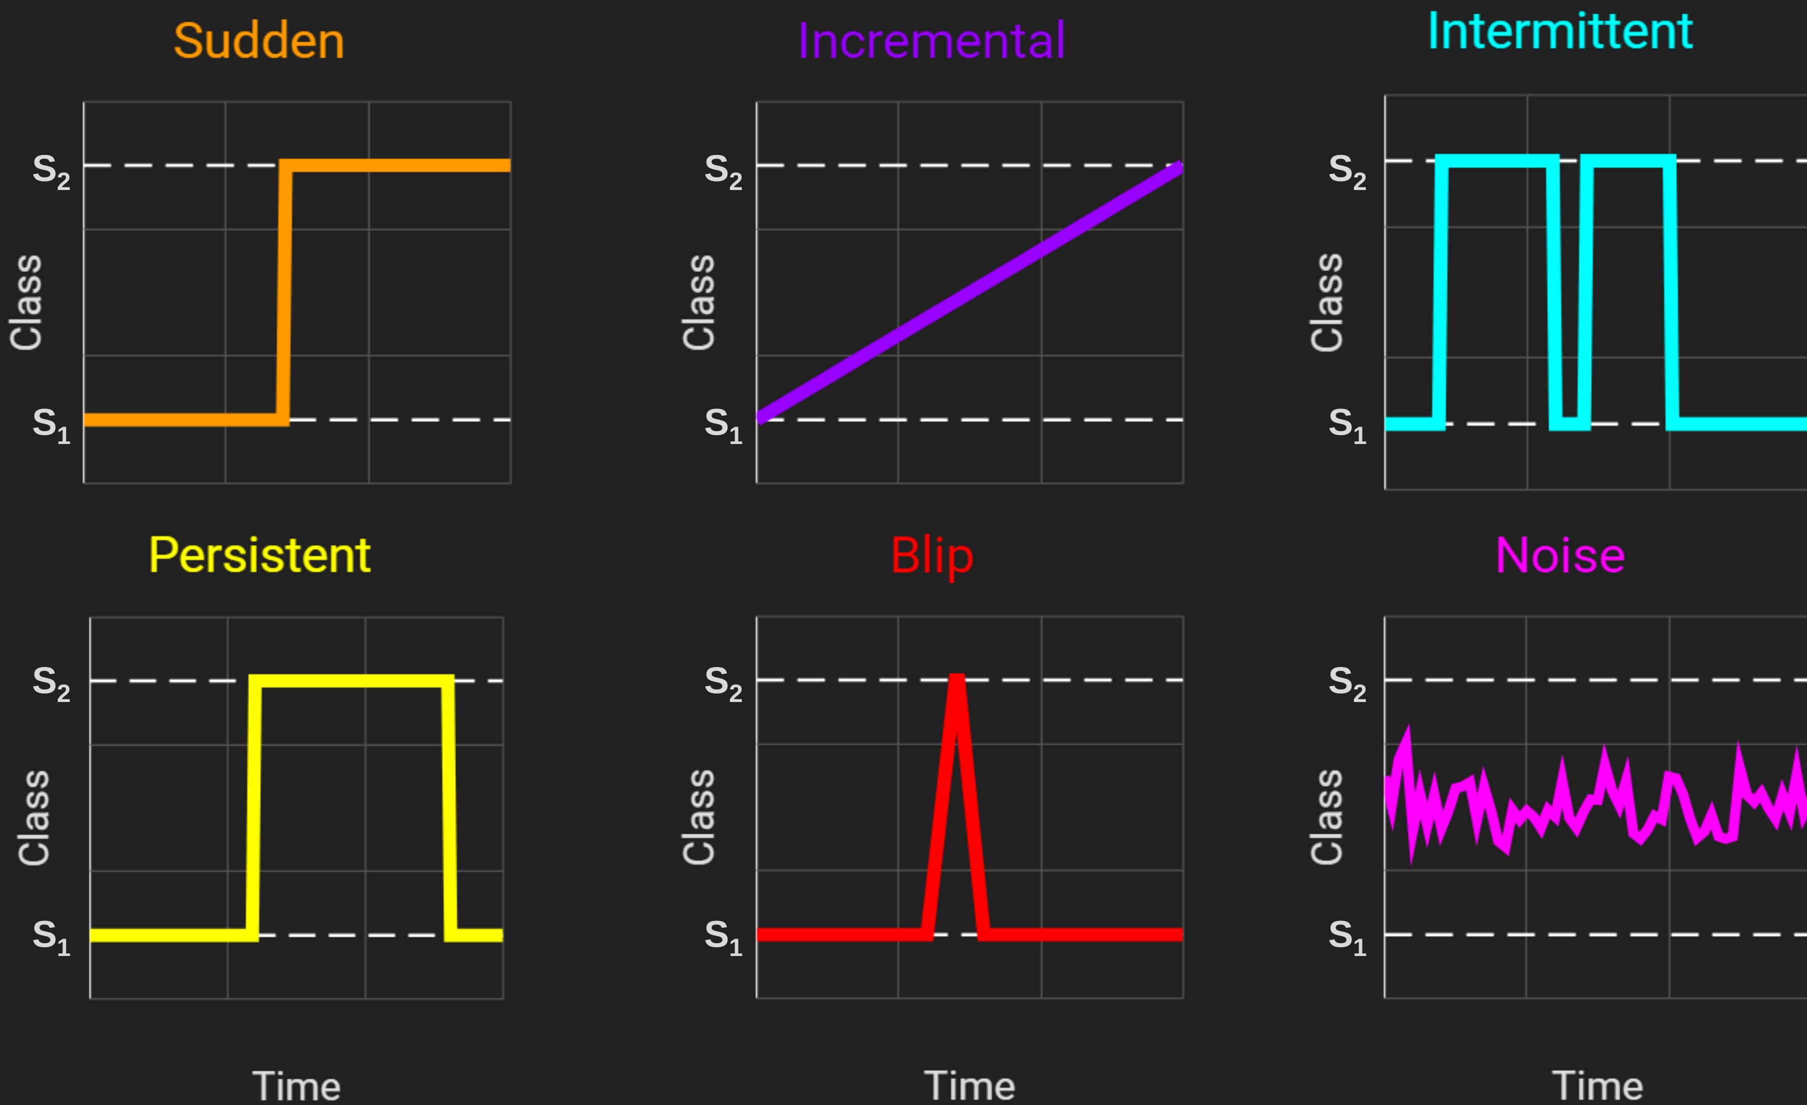

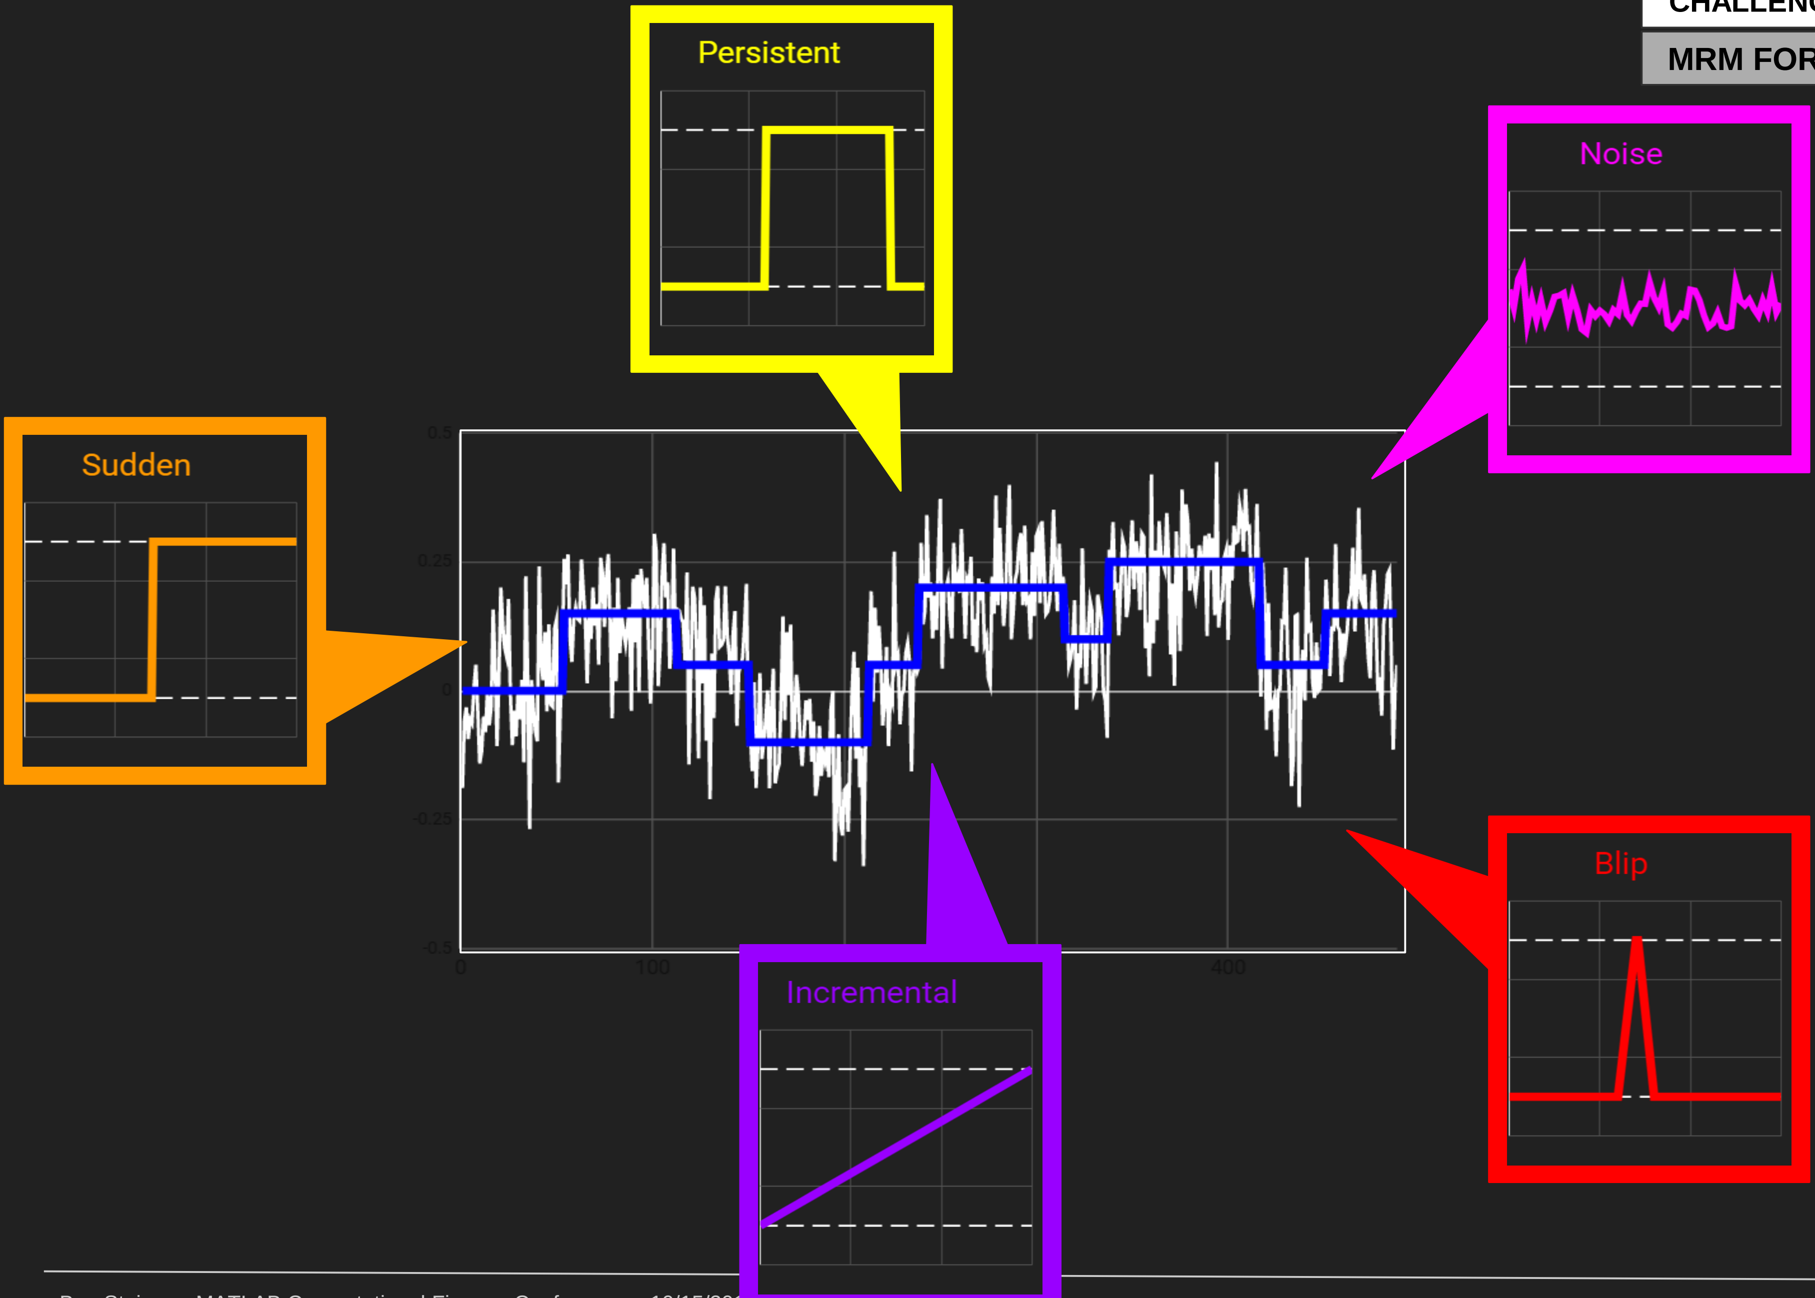
Persistent
Noise
Sudden
Blip
Incremental
0.5
0.25
0
-0.25
-0.5
0
100
400

Challenge 2: Interpretation

Why do we need interpretation?

What do we mean by interpretability?

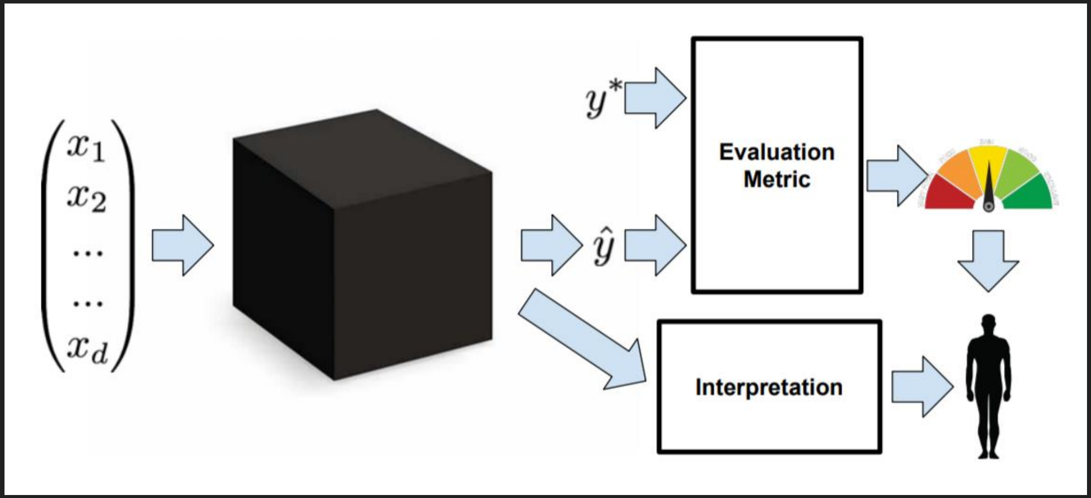

Zachary Lipton,UCSD, 2016 ICML Workshop on Human Interpretability in Machine Learning

- Causality
- Comprehension
- Decomposition
- Algorithmic transparency
- Post-hoc interpretation

“... the engineers couldn't explain why the AI was executing the trades it was making. The creation was such a black box that even its creators didn't understand how it worked…”

“The Massive Hedge Fund Betting on AI”, Bloomberg, 9/27/2017

Challenge 2: Look at the weights!

4 input features: 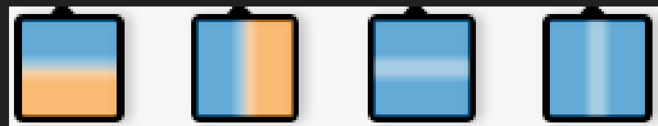

target classification:: 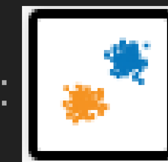

Network with 2 hidden layers of 3 & 2 neurons

Weight = thickness

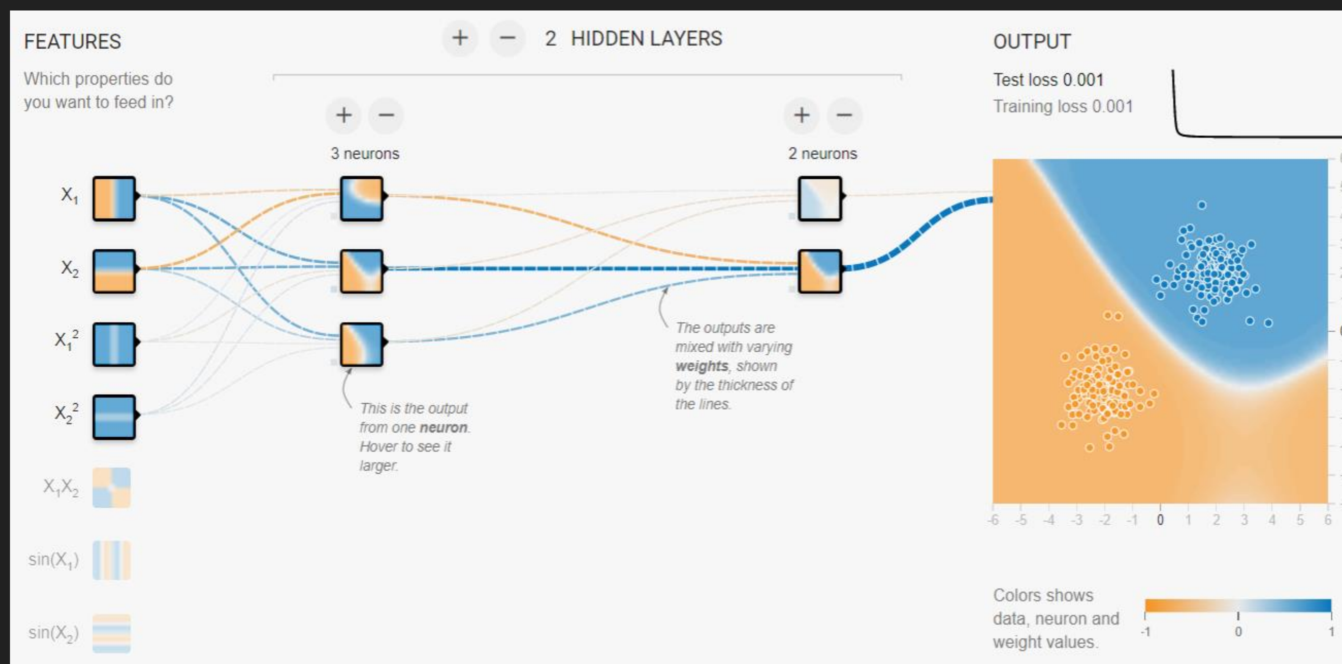

http://playground.tensorflow.org

Challenge 3: Learning what we already know

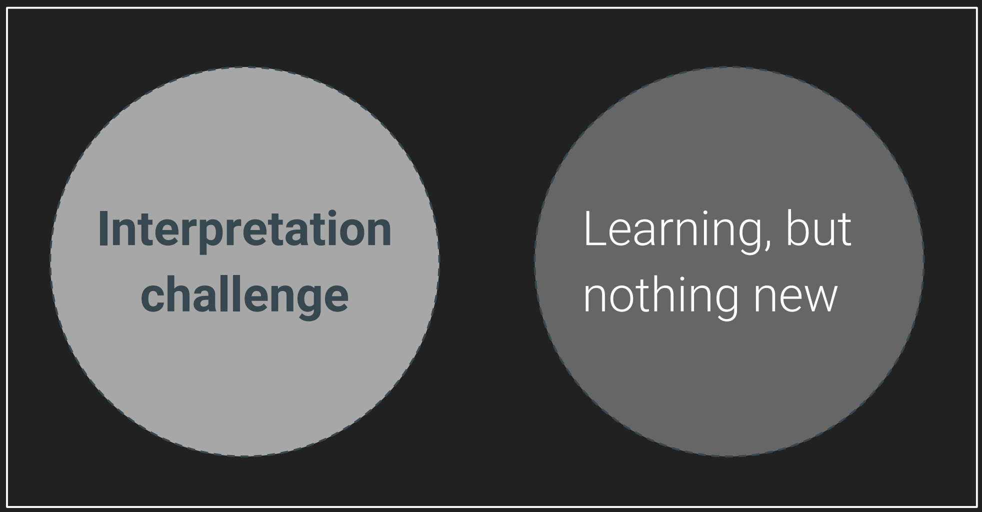

Good news: can't be both!

Solution:

Step 1: **Traditional multifactor model** (with known factors)

Step 2: **Deep Learning** (with residuals from step 1)

Challenge 3: Multifactor models

Multifactor models have the general form

$$R_{it} = \alpha_i + \boxed{\beta_{1i} f_{1t} + \beta_{2i} f_{2t} + \ldots + \beta_{Ki} f_{Kt}} + \varepsilon_{it} \quad (1)$$

$$= \alpha_i + \boxed{\beta'_i f_t} + \varepsilon_{it} \quad (2)$$

where

- R_{it} is the *simple return* (real or in excess of the risk free rate) on asset i $(i = 1, \ldots, N)$ in time period t $(t = 1, \ldots, T)$
- f_{kt} is the k^{th} *common factor* $(k = 1, \ldots, K)$
- β_{ki} is the *factor loading* or *factor beta* for asset i on the k^{th} factor
- α_i is the idiosyncratic return of asset i
- ε_{it} is the error term (or *'the return on asset i at time t from unknown factors'*)

Agenda

1. **Introductory concepts**
   a. What is model risk management?
   b. Machine Learning for Alpha Strategies
2. **Challenges of Deep Learning**
   a. Non-stationarity
   b. Interpretation
   c. Learning what we already know
3. **Model Risk Management for Deep Learning and Alpha Strategies**
   a. Backtest evaluation
   b. Ongoing monitoring

“The road to hedge fund failure is littered with good backtests”

Conceptual Soundness = Backtest Evaluation

- Return per unit risk per unit capital required
- Alpha Decay
- Temporal P&L
- Strategy Correlation
- Sensitivity Analysis
- Random Markets

Evaluation 1: Alpha Decay

Alpha term structure

Cost of implementation delay

Dictates execution speed

Too fast = alpha not capturable

Alpha not declining at all raises suspicion

Declining alpha indicates profit from trades at t=0.

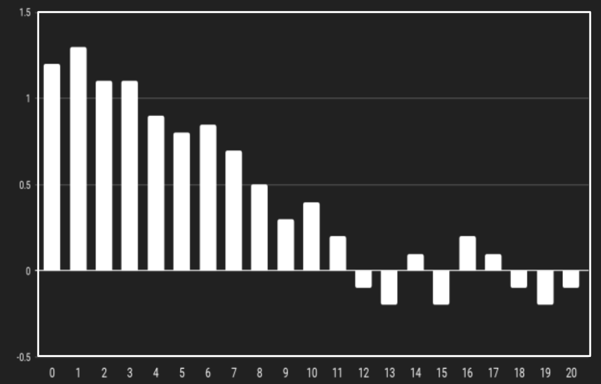

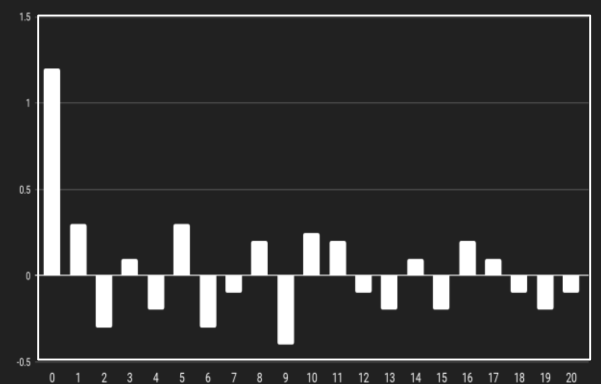

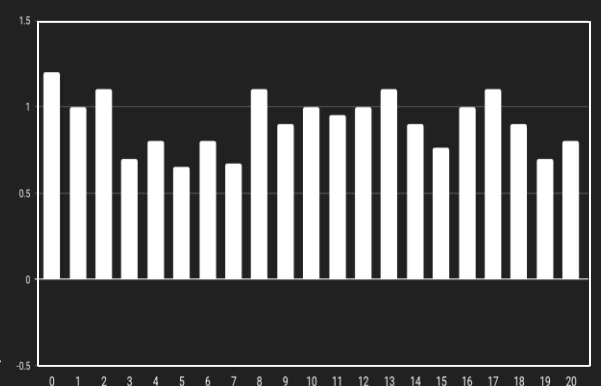

!?

Evaluation 2: Temporal P&L

Three strategies with same long run risk adjusted return but different temporal performance

Strategy decay: cyclical or secular?

1. Secular decay: avoid
2. Cyclical decay: Trend follow
3. Improving performance: yes please!

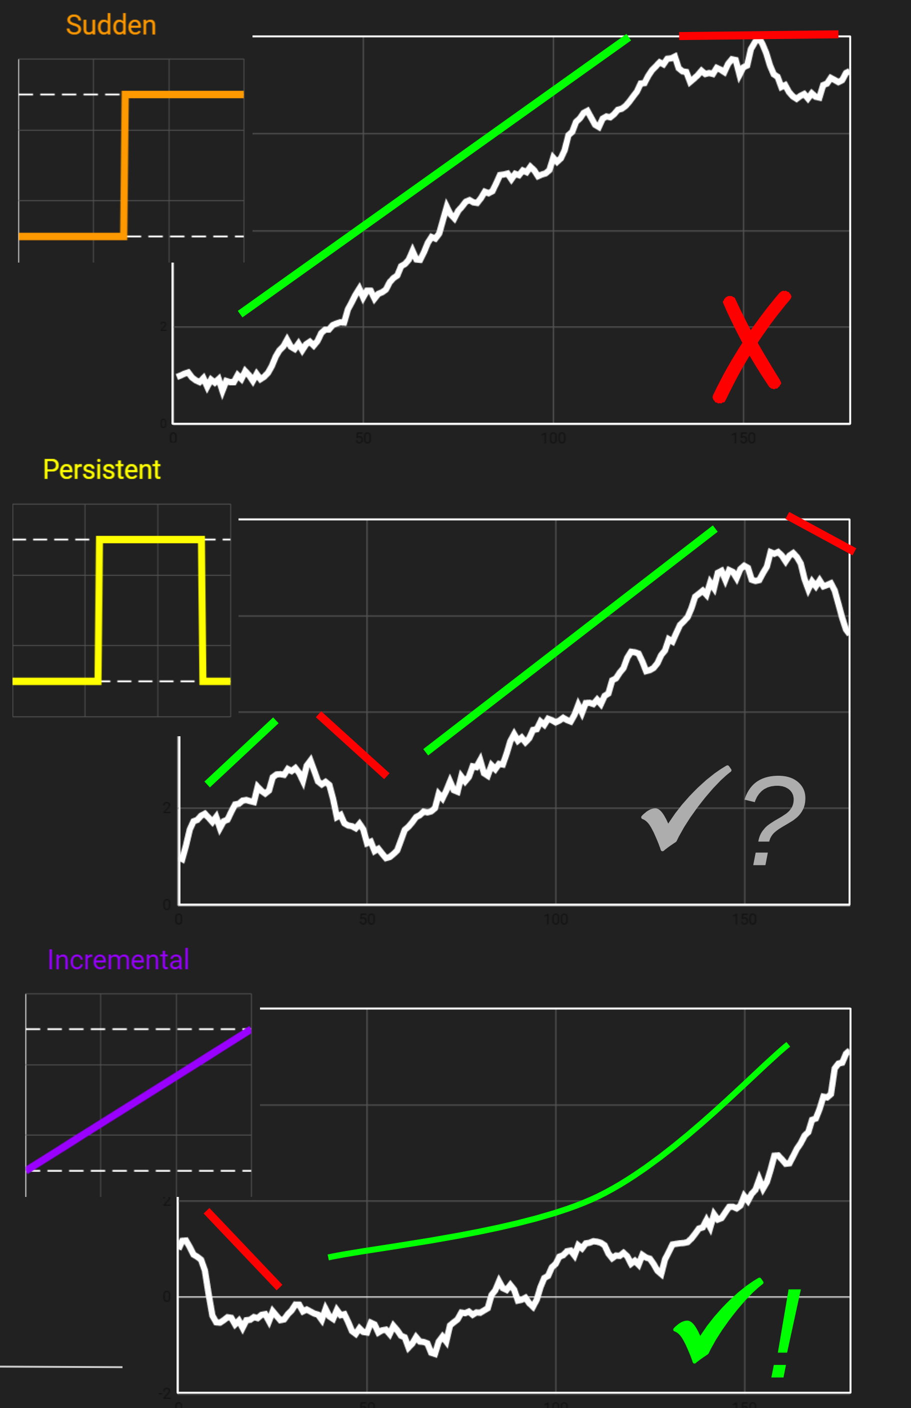

Evaluation 3: Strategy Sensitivity

Correlation with exogenous factors (eg: macro environment)

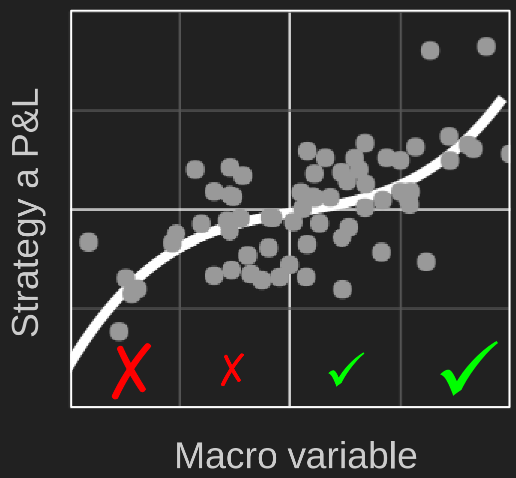

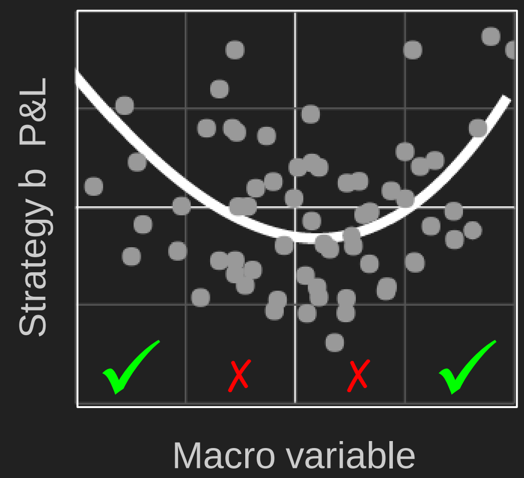

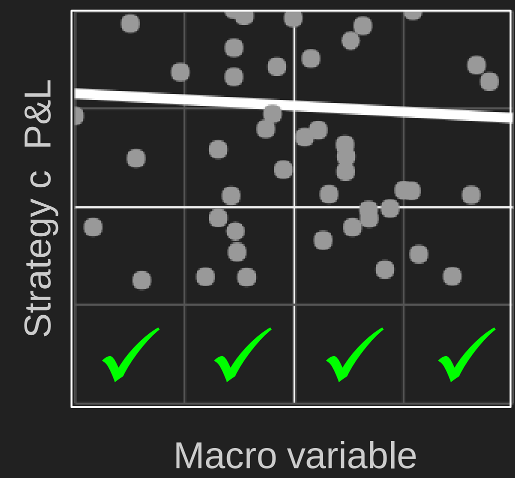

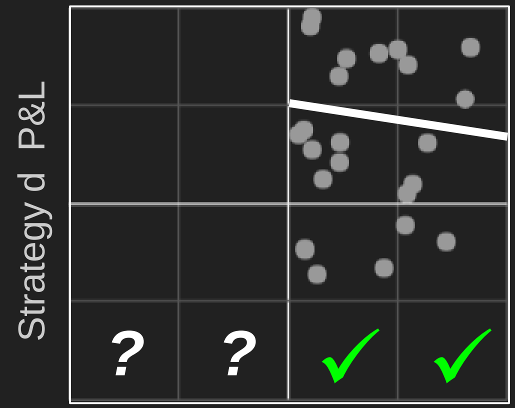

What don’t we know?

Evaluation 4: Random Portfolios

Random portfolio **weights**. No Deep Learning.

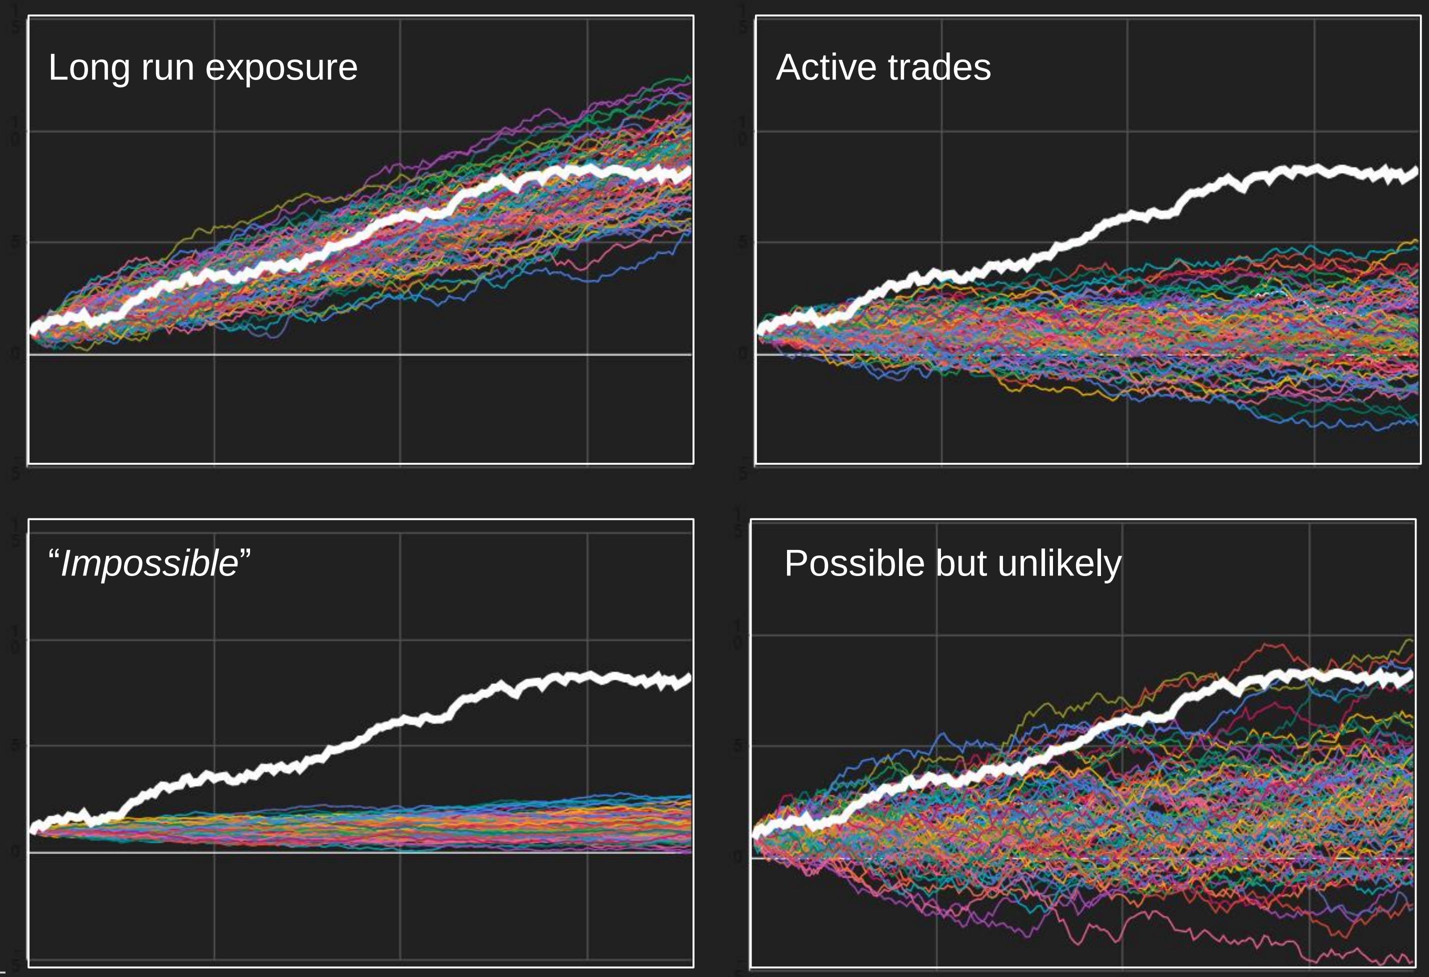

Evaluation 5: Random Returns

Randomize order of **returns**. Full retraining on noise…

- Break covariance between returns and features
- Break autocorrelation of returns
- Keep original features

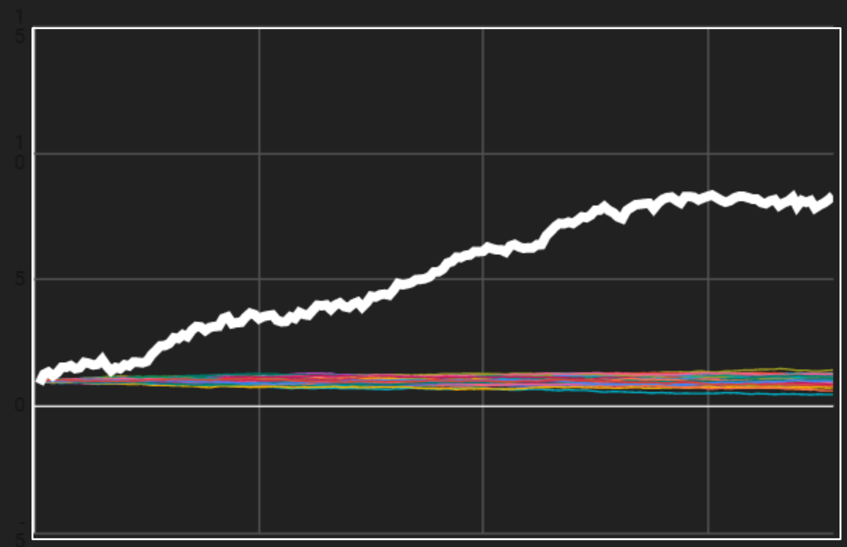

Strategy Monitoring

“You’ve never experienced your worst drawdown”

Is the strategy performing “*as intended*”

As intended = per backtest

How bad is bad?

Real world consequences

Out-of-sample:

- Distribution
- Small samples (SPC)

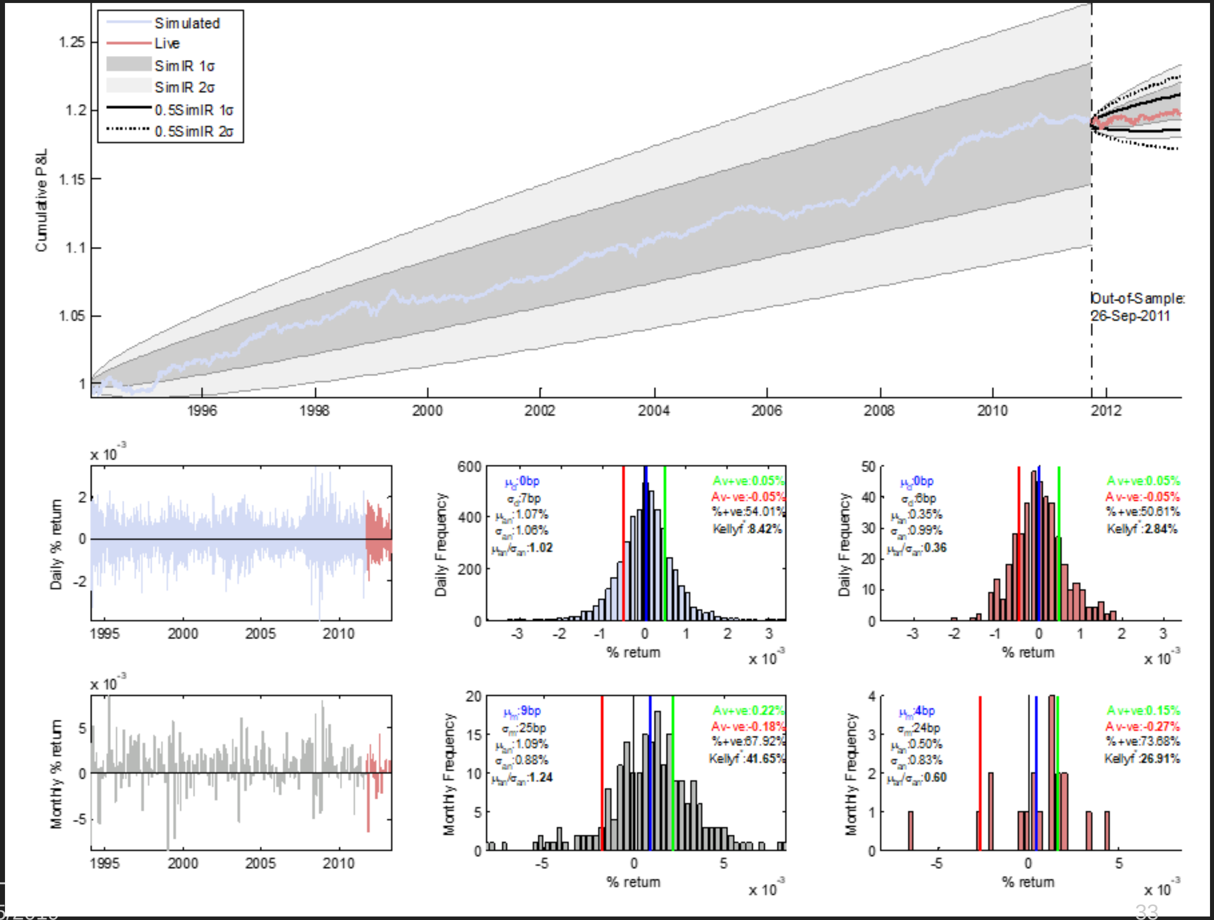

Conceptual Soundness

Evaluating Alpha Strategies with MATLAB

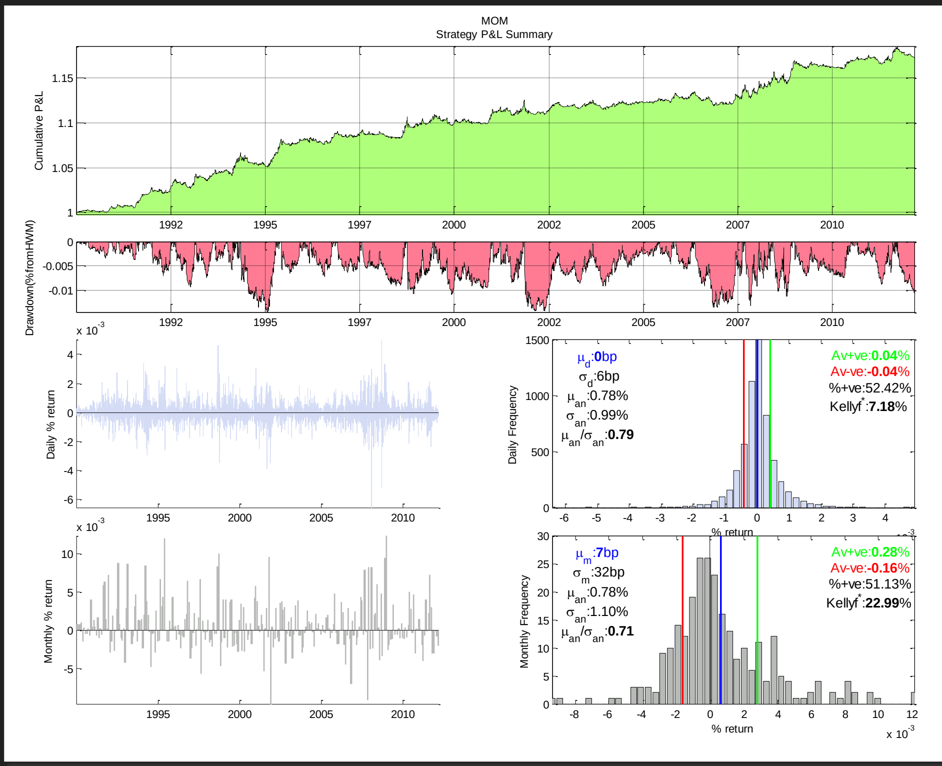

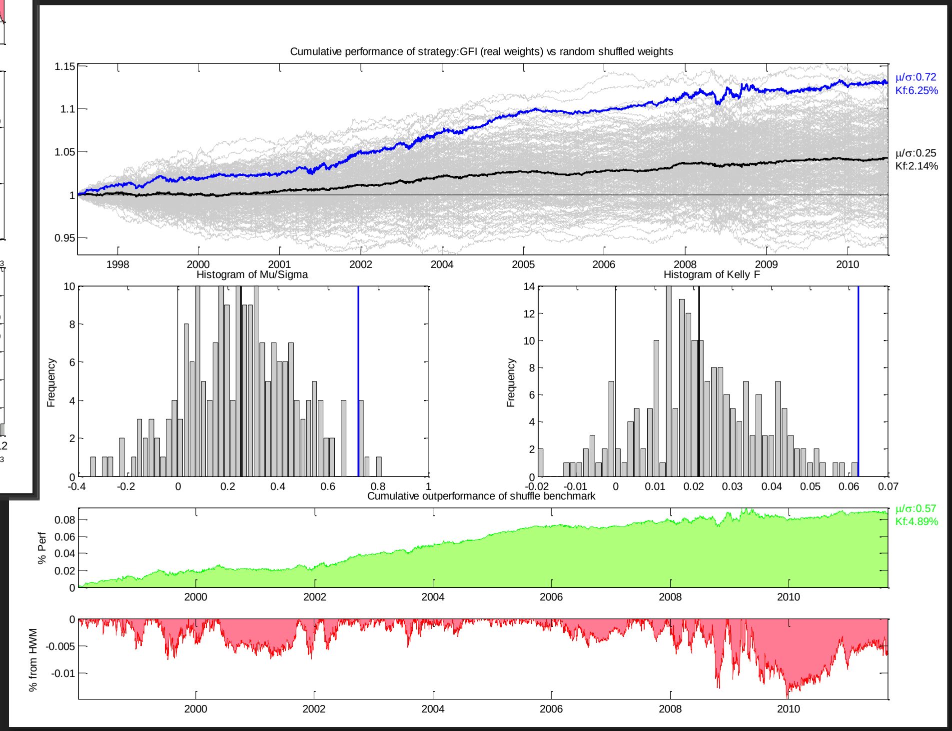

Summary

1. **Introductory concepts**
   a. Model Risk Management is controlling the risk a model is "wrong"
   b. Any Machine Learning is still a model
   c. Deep Learning maps inputs to outputs (eg: price direction) via non-linear functions
2. **Challenges of Deep Learning**
   a. Financial markets experience Concept Drift (video games do not!)
      - What type of Concept Drift is expected?
      - And how much data can be collected before the system changes again?
   b. Interpretation can be addressed (hint: look at weights!)
   c. Avoid learning what we already know (use residuals as the target)
3. **Model Risk Management for Deep Learning and Alpha Strategies**
   a. Conceptual Soundness via backtest evaluation
   b. Ongoing monitoring because "*you've never seen your largest drawdown*"

In conclusion, let’s make sure that...

The road to Alpha Strategy failure is <u>not</u> littered with Deep Learning backtests