

A Time Series Approach to Explainability for Neural Nets with Applications to Risk-Management and Fraud Detection

Marc Wildi: ZHAW, marc.wildi@zhaw.ch

Branka Hadji Misheva: BFH, branka.hadjimisheva@bfh.ch

September 28, 2022

Table of contents

- 1 Forecasting with NN
- 2 Black-Box: Why we Need Explainability
- 3 Classic XAI Approaches
- 4 Regression and Explainability
- 5 NN and Explainability: LPD
- 6 LPD Applied to Bitcoin
- 7 Application to S&P500
- 8 Summary

Neural Nets: Forecasting

- Review of **international forecast competitions**
- Before 2015: **classic linear** approaches **win**
 - M1 (1982), M2 (1993), M3 (2000), NN3 (2007) and NN5 (2009) competitions
- Recently: **hybrid approaches** based on a mix of ARIMA and neural nets **outperform**
 - M4 (2020) and M5 (2021)
- Accruing interest in **neural nets for forecasting** (in particular economic time series)
- BUT... **Black Box**

Black-Box: Why We Need Explainability

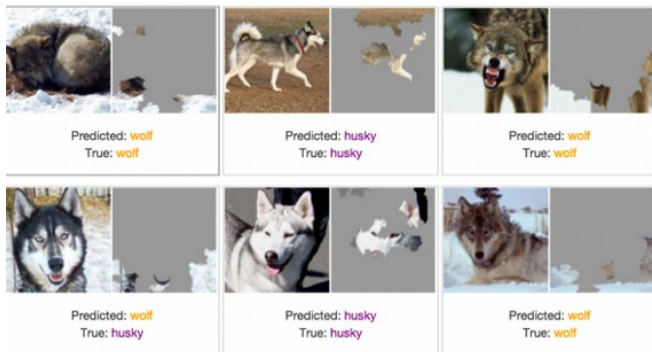
Let's consider the 'Husky vs Wolf' experiment results.



- The classifier makes one mistake!

Why Do We Need Explainability?

Next, we investigate which features drive the classification.



- The decision is based on white patches


Why Do We Need Explainability?

- Verify that **accuracy** is the result of **proper problem representation**
 - The model is capturing **relevant dependencies** between features.
 - This ensures **trust** in the system.
- **Communication**: convince layperson
- **Regulation** demands it.

No black box excuses – explainability/traceability of models is necessary and can improve the analysis process | It is the responsibility of supervised firms to ensure that BDAI-based decisions can be explained and are understood by third-party experts. Supervisory authorities take a critical view of models that are categorised purely as black boxes. New approaches allow firms using such models to at least gain some insight into how these models work and identify the reasons behind decisions. In addition, a better understanding of models provides an opportunity to improve the analysis process – allowing, for instance, the responsible units in the supervised firm to identify statistical problems.

Figure: Extract: Bafin AI and Big Data Report 2020

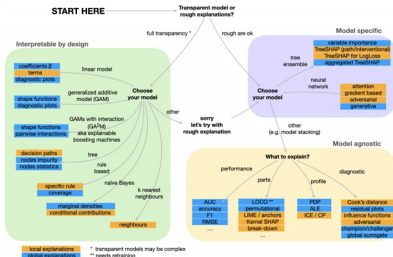
Deploying Explainability: Zoo of XAI models



Arya et al. (2019) proposed taxonomy based on questions about what is explained, how it is explained and at what level



Linaratos et al. (2021) taxonomy mind-map of Machine Learning Interpretability Techniques.



Maksymiuk et al. (2021) model-oriented taxonomy for XAI method

Figure: Machine Learning Interpretability Techniques

Deploying Explainability: Utility of Classical XAI Methods for Finance

- Classic approaches are **data-intensive** (sometimes odd/difficult to explain...)
- Address 'data-intensiveness'
 - Create 'fake' data (simulation) or
 - **Shuffle** available data
- Creating 'fake' data: **contradiction** (don't know true model)
- Shuffling: kills **dependencies** (trends, vola-cluster, draw-downs, extreme events, ...): a 09-27-2022 observation could sit next to a 01-03-1999 observation
- **Key limitation**: many classical methods **ignore feature dependence**

Explainability Example: Classic Regression

- Let

$$y_t = 1 + 0.5x_{1,t} + 1.4x_{2,t} + \epsilon_t$$

- **Interpretation:** what is the meaning of the parameters?
Partial Derivative!!!!
- **Communication:** regression is 'well-known'
- **Validation:** confront model to common sense/expert knowledge/**experience**

New XAI-Tool

- 1 Preserve **dependence** (no shuffling)
- 2 Avoid '**fake**' (no simulation)
- 3 Link to **regression** (no exoticism)

NN in a 'Regression' Perspective

- Consider **exact 'regression'** replication of net: **partial derivatives**

$$\begin{aligned}y_1 &= NN\left(\text{Input}(1)_1, \text{Input}(2)_1, \text{Input}(3)_1\right) \\ &= 0.1 + 0.2\text{Input}(1)_1 + 0.3\text{Input}(2)_1 - 0.2\text{Input}(3)_1\end{aligned}$$

$$\begin{aligned}y_2 &= NN\left(\text{Input}(1)_2, \text{Input}(2)_2, \text{Input}(3)_2\right) \\ &= 0.15 + 0.23\text{Input}(1)_2 + 0.36\text{Input}(2)_2 - 0.18\text{Input}(3)_2\end{aligned}$$

... = ...

$$\begin{aligned}y_t &= NN\left(\text{Input}(1)_t, \text{Input}(2)_t, \text{Input}(3)_t\right) \\ &= b_t + w_{1t}\text{Input}(1)_t + w_{2t}\text{Input}(2)_t + w_{3t}\text{Input}(3)_t\end{aligned}$$

XAI-Tool LPD (Linear Parameter Data)

- **LPD: Data Transformation**

$\text{Input}(1)_t, \text{Input}(2)_t, \text{Input}(3)_t \rightarrow \text{LPD}_t := b_t, w_{1t}, w_{2t}, w_{3t}$

- Dimension $T \times n$ -dim (n =number of inputs) irrespective of NN-architecture
- **New:** time-dependent **intercept** b_t , see [paper](#)
 - Differs from classic gradient approaches (saliency-maps,...)
 - Non-linear signal extraction (estimate of local time-dependent drift)
- **Derivation** of LPD: see [paper](#)
- **Generalization** of LPD: **X-functions**, see [paper](#)
 - Can address **departure from linearity, overfitting, trading performances**

Application of LPD to BTC Crypto-Currency

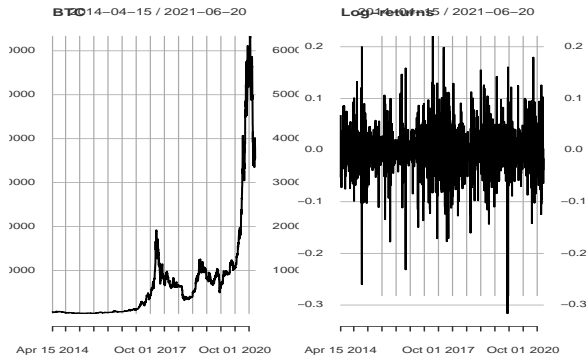


Figure: SP500: log-prices and log-returns

- Advise against an application of NN to **trending** data, see [paper](#)

Neural Net

- Forecast tomorrows log-return of BTC based on last week's data


$$\hat{BTC}_{T+1} = NN(BTC_T, BTC_{T-1}, \dots, BTC_{T-5})$$

- NN is a neural net with a single hidden layer with 100 neurons
- LPD: **drift** and **weights** assigned to $BTC_T, BTC_{T-1}, \dots, BTC_{T-5}$

$$LPD_t = (b_t, w_{1t}, w_{2t}, w_{3t}, w_{4t}, w_{5t}, w_{6t})$$

Mean LPD of All Input Variables

Mean out-of-sample LPD of 100 random nets



XAI-Outcome: Explain the Net!

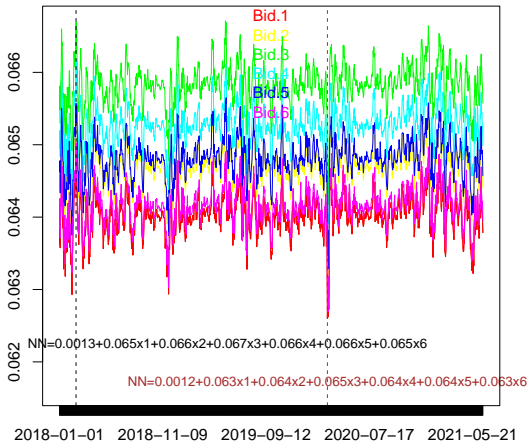
- **XAI: first order linear approximation**
 - Neural net is close to an unassuming **equally-weighted MA(6)**-filter

$$\bar{o}_t = \frac{1}{100} \sum_{i=1}^{100} o_{t,i} = \overline{\text{LPD}}_t \left(\begin{array}{c} 1 \\ \mathbf{x}_t \end{array} \right) \approx 0.0015 + 0.065 \sum_{j=1}^6 x_{jt}$$

- Can gain **'trust'** in the forecast (**simple forecast heuristic**)
- Unfortunately **not very useful**
- But... Let's look at the **second order departures from linearity**

Mean LPD of All Input Variables

Mean out-of-sample LPD of 100 random nets




Risk-Management (RM)

- Idea
 - Identify '**uncertain times**' or '**unusual states**' of the market in real-time
 - **Downsize** market exposure during **uncertain times**/unusual states
- Identification of **uncertain times** and **unusual states**
 - Unusually **weak LPD** (weak dependence structure): see [Paper](#)
 - Unusually **strong QPD** (strong non-linearity: see [paper](#))
 - Uncertainty: **spreading** LPDs (see [paper](#))
- **RM: does not rely on forecasting at all!!!**

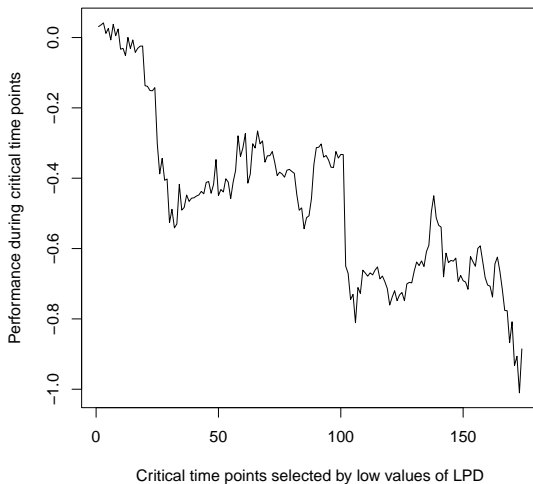
Market-Exit when $|LPD|$ Weak (Critical Time)

- **Out-of-sample** performances: net is trained up to 01-01-2018

y-and-hold (black) vs. LPD-down (blue): $< 1/7$ -quantile, length 9



Next-Day's Cumulated Performance at Critical Time-Points



Analysis: Critical, Neutral and Auspicious Time

	Next day's mean return
Today's LPD: Weak dependence	-0.509%
Today's LPD: Normal dependence	0.244%
Today's LPD: Strong dependence	0.305%
All time points	0.142%

- Novel RM-tool: does **not** rely on '**direction**' (**NN-forecasts are ignored!!!**)
- Does **not** rely on **volatility**
- Does rely on **market dependence structure**: weak dependence might be due to panic/herding/overreaction
- **Orthogonal** to classic RM-approaches

Application to S&P500

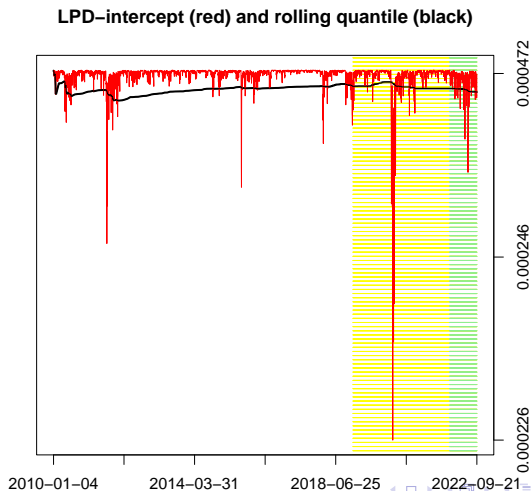
- **Equity-index** is subject to protracted down-turns during **economic recessions**
 - Frequency: once per decade
 - Duration: half a year (at climax)
 - Unconditional probability: $0.5/10=1/20$ (quantile)
- **Risk-management**: exit-market in case of 'severe' down-turns
 - LPD $<$ own historical 1/20 quantile (rolling-window)

Novelty: LPD-Intercept

- XAI-tool relies on **LPD-intercept**, see [paper](#)
 - Differs from classic gradient approaches to XAI (saliency-maps,...)
- Intercept is estimate of **local drift** in returns
 - Close to **non-linear signal extraction**: extraction of trend/drift for non-stationary DGP

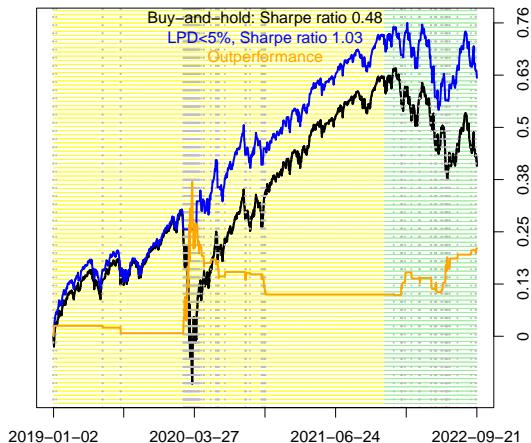
LPD Market-Exit Signals: 'Red' Below 'Black'

- Yellow area: out-of-sample (Pandemy); green area: 'truly' out-of-sample (inflation, interest rate hikes, Russian invasion)




S&P500: LPD vs. Buy-and-Hold

Buy-and-hold (black) vs. LPD (blue) and ouperformance (orange)



Draw-downs: LPD vs. Buy-and-Hold


Drawdowns: buy-and-hold (black) vs. LPD (blue)



Market-Exit if NN-Forecast Small/Negative (Bad Outlook)

- Conventional directional RM-tool: exit market if NN-forecast small/negative ($< 1/20$ quantile)

Buy-and-hold (black) vs. LPD (blue) and outperformance (orange)



LPD vs. Forecasts

- Forecasts are less informative than LPD (non-linear signal extraction)
- LPD much less noisy! (smooth drift estimate)
- LPD faster than linear MA-filters (NN reacts immediately if data-point 'suspect')

Summary

- XAI tool: exact, fast, clear/interpretable, preserve data integrity (no 'fake', do not alter dependence)
- LPD:
 - **XAI**: Explain and validate NN (first order linear approximation)
 - **Extension**: risk-management
 - **Story**: Chaotic markets → weak dependence → downsize exposure
 - **LPD faster** than classic MA-filters (relative anticipation)
- Fraud detection: see [paper](#)