

Multi-Factor Inception: What to Do with All of These Features?

Tom Liu
University of Oxford
Oxford, UK

Stefan Zohren
Oxford-Man Institute of Quantitative Finance, University
of Oxford
Oxford, UK

ABSTRACT

Cryptocurrency trading represents a nascent field of research, with growing adoption in industry. Aided by its decentralised nature, many metrics describing cryptocurrencies are accessible with a simple Google search and update frequently, usually at least on a daily basis. This presents a promising opportunity for data-driven systematic trading research, where limited historical data can be augmented with additional features, such as hashrate or Google Trends. However, one question naturally arises: how to effectively select and process these features? In this paper, we introduce Multi-Factor Inception Networks (MFIN), an end-to-end framework for systematic trading with multiple assets and factors. MFINs extend Deep Inception Networks (DIN) [21] to operate in a multi-factor context. Similar to DINs, MFIN models automatically learn features from returns data and output position sizes that optimise portfolio Sharpe ratio. Compared to a range of rule-based momentum and reversion strategies, MFINs learn an uncorrelated, higher-*Sharpe* strategy that is not captured by traditional, hand-crafted factors. In particular, MFIN models continue to achieve consistent returns over the most recent years (2022-2023), where traditional strategies and the wider cryptocurrency market have underperformed.

KEYWORDS

Machine Learning, Cryptocurrency trading, Alternative data.

DISCLAIMER

This work reflects the analysis and personal view of the author Tom Liu. No reader should interpret this work to represent the views of any third party. Assumptions, opinions, views and estimates constitute the author's judgement as of the date given and are subject to change without notice and without duty to update.

1 INTRODUCTION

Cryptocurrencies have emerged as a novel and increasingly popular asset class in recent years. Since the inception of Bitcoin [26] in 2009, the global cryptocurrency market has grown rapidly, reaching a market capitalisation of \$1.2 trillion in 2023 [9]. This is fast approaching established equity indices, such as FTSE 100 and EURO STOXX 50. In a similar trend, researchers are actively exploring strategies to capitalise on cryptocurrency trading opportunities.

Cryptocurrency markets exhibit parallels with equities, including high volatility, positive asset correlations, and idiosyncratic risk. Consequently, one approach involves adapting rule-based strategies initially developed for equities into the cryptocurrency domain. Studies have demonstrated the presence of analogous effects, such as momentum [8, 18, 23], reversion [17], value and carry [13], in cryptocurrency data. Notably, [22] reveals a strong Time Series Momentum (TSMOM) effect in both price data and alternative data

sources like Google searches and tweets. This indicates the wider applicability of technical indicator-based strategies beyond conventional price and volume data for cryptocurrencies.

In addition to traditional strategies, researchers increasingly apply Machine Learning (ML) techniques in cryptocurrency analysis. Although cryptocurrencies have a relatively short history, their decentralised nature facilitates access to higher frequency data, including minute-by-minute prices from platforms like Binance [15], as well as alternative data sources [5, 6], which can help train more complex ML models. Transfer learning from related assets, such as Foreign Exchange (FX), has also been used to train Fused Encoder Networks in data-scarce scenarios [28].

Prior ML research often relies on hand-crafted input features. For example, [28] extracts raw and normalised price returns over different time frames. While recent models can learn features directly from data, they are typically limited to a single input factor [14, 21, 24, 34]. For ML models, we define factors as a type of input data, including price, volume or alternative data, such as hashrate. Moving to a multi-factor context with alternative data, [29] provides indicators for change, momentum, and decay in each factor as the inputs for their AutoEncoder model. This scenario presents a common challenge: when developing new multi-factor ML strategies, a crucial question is how to select and process features. This can significantly increase the dimensionality of the input for time series models that leverage prior historical data. Simultaneously, it restricts model inputs to a predefined set of hand-crafted features.

Our approach addresses these challenges by providing a single time series of returns for each asset and factor. We then develop a model that can learn useful features automatically from the data. To achieve this, we introduce Multi-Factor Inception Networks (MFIN) as a novel way to combine price and alternative data across multiple assets. This draws on prior work in Deep Inception Networks (DIN) [21], and extends the concept to multiple factors, including price, volume, hashrate and tweets. Our experiments demonstrate that MFIN strategies remain profitable after transaction costs and learn uncorrelated behaviours to traditional strategies, such as momentum and reversion. MFIN can also be combined with an existing portfolio of traditional strategies, contributing positively to Sharpe ratio and breakeven transaction costs, whilst limiting drawdowns.

2 DATA

We evaluate strategy performance on 7 cryptocurrencies, according to price and volume data from CoinMarketCap (CMC) [9] over January 2018 to March 2023. These are spot rates against USD for BCH, BTC, DASH, DOGE, ETH, LTC, ZEC. Given that these coins are highly liquid, prior work has shown that CMC prices are reliable and accurately reflect the underlying market processes [31].

To avoid survivorship bias, we use Wayback Machine [1] to select the top 25 coins by market capitalisation on CMC as of January

2019: the start of the first test set. Stablecoins, such as Tether and USD coin, are excluded. We then filter to the final 7 coins according to the availability of alternative data on BitInfoCharts (BIC) [5] in January 2019. We supplement and cross-check alternative data with Blockchain (BC) [6], where data for 6 coins is available from March 2019. The final coin (ZEC) is added to BC in 2020.

Google Trends data G is missing for 3 of 7 coins in BIC, so we obtain this from the Google Trends website [10]. Due to website limits, we download data in 90-day segments to allow for daily resolution. We link segments with backwards proportional adjustment [30] from Eq. 1. This preserves percentage changes, since Google Trends are scaled in proportion to the largest signal in each segment.

$$G_{adjusted} = G_{raw} * \frac{g_2}{g_1}, \quad (1)$$

g_1 = last datum of preceding segment on roll date,
 g_2 = first datum of new segment on roll date.

To avoid lookahead bias, all strategies trade at market open on date t , using price and alternative data from prior dates up to $t - 1$. Although new open prices are available at date t , we do not include this to improve tradeability. The features from each data source are listed in Table 1.

Table 1: Features from CoinMarketCap (CMC), BitInfoCharts (BIC) and BlockChain (BC) data sources. All monetary values are in USD.

CMC Price	BIC Alternative	BC Alternative
open	transactions	fee-reward ratio
high	block size	chain size increase
low	sent addresses	coin days destroyed
close	sent USD	cost per transaction
volume	difficulty	
market cap.	mining profitability	
	hashrate	
	av. transaction size	
	av. transaction value	
	confirmation time	
	tweets	
	google trends	

3 TRADITIONAL STRATEGIES

We first apply rule-based strategies that capture momentum and reversion behaviour separately.

3.1 Momentum Strategies

For each alternative feature j from the BIC and BC data sources, we consider a Time Series Momentum (TSMOM) strategy that buys/shorts assets with positive/negative signal strength,

$$w_{i,t} = \text{sign}(s_{i,t}). \quad (2)$$

3.1.1 Moskowitz, Ooi & Pedersen (MOP) [25]. MOP is the seminal TSMOM strategy. The authors propose a score based on past 12-month returns for a strategy with monthly investment horizon. In this paper, we modify the signal $s_{i,t} = r_{i,j,t}^{(k)}$ to be returns over past $k \in \{5, 21, 63, 126, 252\}$ days.

3.1.2 Baz et al. (BAZ) [3]. BAZ proposes using MACD indicators to capture the cross-over of exponentially-weighted moving averages (EWMA). When the EWMA with short timescale $S_k \in \{4, 8, 16, 32\}$ is greater than the EWMA with long timescale $L_k \in \{12, 24, 48, 96\}$, then we expect the asset to have upwards momentum and vice versa. The resulting BAZ signal is $s_{i,t} = \text{MACD}_{i,j,t}^{(k)}$.

3.2 Reversion Strategies

3.2.1 Bollinger Bands (REV) [7]. REV is a reversion strategy, based on the spread $\delta_{i,j,t} = r_{i,open,t}^{(k)} - r_{i,j,t}^{(k)}$ between the open price and an alternative feature j on the same date t . The return length $k \in \{1, 5, 10, 21\}$ days is a tuneable parameter. We compute a z-score signal $s_{i,t}$ from $\delta_{i,j,t}$ in a rolling 63-day exponentially weighted window. The portfolio weights short each asset according to the entry $z_u \in \{1.5, 1.75, 2.0\}$ and exit $z_l \in \{0.5, 0.75, 1.0\}$ z-score thresholds.

$$w_{i,t} = \begin{cases} -\text{sign}(s_{i,t}), & \text{if } |s_{i,t}| \geq z_u \text{ and } w_{i,t-1} = 0 \\ w_{i,t-1}, & \text{if } |s_{i,t}| \geq z_l \text{ and } w_{i,t-1} \neq 0 \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

We verify that the spread $\delta_{i,j,t}$ is stationary by applying an Augmented Dicky Fuller (ADF) test, requiring a p-value $\leq 1\%$.

3.3 Additional Benchmarks

3.3.1 Combined (CMB). CMB is a volatility-scaled portfolio of MOP, BAZ and REV strategies and we compare this with the inherent ability of MFIN models to combine information from different input factors.

3.3.2 Long-only. This benchmark represents market returns and assigns equal weight $w_{i,t} = 1$ to all assets i .

3.4 Portfolio Construction

Each strategy outputs portfolio weights $w_{i,t}$ for all assets i at time t . We perform volatility scaling across $N_A = 7$ assets to $\sigma_{tgt} = 15\%$, to improve Sharpe ratio and decrease the likelihood of extreme returns [11]. The portfolio returns $R_{p,t}$ are defined as,

$$R_{p,t+1} = \frac{\sigma_{tgt}}{N_A} \sum_{i=1}^{N_A} \frac{w_{i,t}}{\sigma_{i,t}} \cdot r_{i,t+1} - C \left| \frac{w_{i,t}}{\sigma_{i,t}} - \frac{w_{i,t-1}}{\sigma_{i,t-1}} \right|, \quad (4)$$

given asset returns $r_{i,t}$, annualised ex-ante volatility $\sigma_{i,t}$ and transaction cost coefficient C . This is consistent with [20, 21, 32, 33].

When comparing different strategies, we apply a second layer of volatility scaling to portfolio returns, so that each strategy takes a similar overall risk. During live trading, it is not possible to scale portfolio volatility to exactly $\sigma_{tgt} = 15\%$, since future volatility can only be estimated. For more realistic results, we perform the second stage of volatility scaling with a rolling 21-day exponentially-weighted standard deviation. This is applied to all strategies, including Long-only.

As shown in [20, 21], the performance of Machine Learning models, such as LSTM and DIN, varies across random initialisations. Therefore, to ensure robust results, we ensemble MFINs over 10 random seeds. Each model is trained with the same hyperparameters and we take a simple average of their outputs.

3.5 Back-testing Details

We optimise MOP, BAZ, and REV strategies across features j and parameters k, z_u, z_l . For data exploration purposes, we produce an overestimate of strategy performance by selecting the two best feature-parameter combinations, according to Sharpe ratio. This is an ex-post analysis, using all available test data from April 2019.

However, realistic strategies must avoid lookahead bias. We simulate this by using expanding-window train-test splits. This is the same procedure for MFIN models, and is shown in Figure 1. We rank all feature-parameter combinations based on Sharpe ratio in each training set, and evaluate an equally weighted portfolio of the top two combinations in the corresponding test set. For diversification, we ensure that the chosen combinations do not share the same feature. This prevents selecting similar strategies with slightly different signal parameters.




Figure 1: Expanding-window train-test splits.

4 MULTI-FACTOR INCEPTION NETWORKS

Deep Inception Networks (DINs) [21] are an end-to-end framework for systematic investing, which learn features and position sizing directly from daily returns. Although originally developed for single-factor scenarios, we can extend DINs to multiple input features with Multi-Factor Inception Networks (MFIN). As shown in Figure 2, we process a $(T \times N_I)$ matrix of standardised daily returns for each asset with a DIN and add a dense output layer to optimise the multi-asset portfolio. T is the sequence length and N_I is the number of input features. The same DIN is used for all assets, which reduces the model complexity in proportion to the total available data. For example, increasing the number of assets N_A from 1 to 50 would not increase the number of model parameters by 50x.

Following the underlying DINs, the MFIN model outputs portfolio weights that directly optimise Sharpe ratio, adjusted for transaction costs and correlation to Long-only.

$$\text{Loss} = -\sqrt{252} \cdot \frac{\text{Mean}(R_{p,t})}{\text{Std}(R_{p,t})} + K \cdot |\rho|, \quad (5)$$

where $R_{p,t}$ are volatility-scaled portfolio returns adjusted for transaction costs. $R_{b,t}$ are portfolio returns of the Long-only benchmark and ρ is Pearson correlation between $R_{p,t}$ and $R_{b,t}$.

As defined in Eq. 4, $R_{p,t}$ depends on transaction cost coefficient C . A large value of C encourages lower portfolio turnover. The correlation regularisation coefficient K helps prevent MFIN models from adopting a static behaviour that may fail to generalise out-of-sample. Hyperparameters C and K are automatically tuned based on losses in the validation sets. In this project, we set $K_{\text{valid}} = C_{\text{valid}} = 0$ to avoid manually scaling the contribution of each term in the validation loss.

4.1 Model Inputs and Outputs

MFINs take 2 input tensors, of size $(T \times N_A \times N_I)$ and $(T \times N_A \times 2)$, respectively. Here, we use sequence length $T = 100$, number of assets $N_A = 7$ in the portfolio, number of input features $N_I = 22$, from both price and alternative data.

- (1) **Past returns** $X_t^{(j)}$: daily returns for each asset i and input feature j , standardised by 63-day exponentially-weighted standard deviation $\frac{r_{i,j,t}}{\sigma_{i,j,t}}$. The last row of $X_t^{(j)}$ is the most recent returns data, computed from $t-2$ to $t-1$ to prevent lookahead bias.
- (2) **Future returns** $Y_t^{(1)}$: $X_t^{(\text{open})}$ shifted forwards by 2 days and scaled to σ_{tgt} . The last row is “one-step-ahead” returns in open price, from t to $t+1$.
- (3) **Volatility scaling** $Y_t^{(2)}$: volatility scaling factors $\frac{\sigma_{tgt}}{\sigma_{i,1,t}}$, corresponding to $Y_t^{(1)}$.

At each timestep t , DIN models predict a vector of position sizes $w_{i,t} \in [-1, 1]$ for portfolio construction. The X_t tensor is used for predictions and Y_t is used to evaluate the loss function. In this paper, we use a batch size equivalent to sequence length T .

4.2 Deep Inception with OrigCIM-LSTM

Due to the limited amount of cryptocurrency data, we favour simpler Deep Inception Network (DIN) models. For a fixed $N_I = 22$, one potential selection is OrigCIM-LSTM. The architecture of this DIN variant is illustrated in Figure 3. As shown in [21], OrigCIM has fewer trainable parameters than more scalable Feature Extractors, such as FlexCIM, when the input width is relatively small. However, in order to capture cross-factor interactions between input features using the Feature Extractor, we do not use simpler Time Series models, such as DeepLOB [35].

4.2.1 Original Custom Inception Module (OrigCIM). OrigCIM is a Feature Extractor that uses a single layer of convolutional filters. We extract features in each time series (TS), across input features (CS), and across both time and input features (combined). These correspond to filters of size $(ts_filter_length \times 1)$, $(1 \times N_I)$, and $(ts_filter_length \times N_I)$, respectively. (1×1) filters propagate the original returns data to the next stage, with a learnt scaling factor. In total, there are $N_T = 4$ feature types, which are learnt in parallel with an Inception Module structure. We use $n_filters$ filters of each type to learn a variety of TS, CS and combined features. This intermediate feature tensor has size $(T \times N_A \times N_F)$, where the number of learnt features $N_F = N_T \times n_filters$. Finally, we perform Dimensionality Reduction by convolving across learnt features.

4.2.2 Long Short-Term Memory (LSTM) [20, 12]. To constrain model complexity, the largest choice of hyperparameter ts_filter_length




Figure 2: High-level model architecture for Multi-Factor Inception Networks (MFINs).




Figure 3: OrigCIM-LSTM variant of Deep Inception Network (DIN).

is 20. Therefore, we use an LSTM in the Position Sizer to capture longer-term interactions in the data. For $N_H = hidden_layer_size$ LSTM cells, the output size is $(T \times N_H)$. Different to the original DINs proposed in [21], a final dense layer in the OrigCIM-LSTM variant produces an output width of 1. Portfolio optimisation across assets occurs within the encapsulating MFIN model instead.

4.3 Model Training

Following [20, 27, 32, 33], we train our model with an expanding window approach with a 90%/10% train/validation split for hyperparameter optimisation. The cryptocurrency dataset uses 1-year increments, as shown in Figure 1.

All MFIN models are implemented in *Tensorflow*. We customise the model training pipeline so that training losses are calculated from “one-step-ahead” predictions of the entire batch, and disable data shuffling during training. This ensures that the transaction costs in the loss function of Eq. 5 are correct. Consistent with [20, 21, 27, 32, 33], we use the Adam optimiser [16].

Given prior work which shows HB to be effective for tuning the underlying DINs [21], we use 30 iterations of Hyperband (HB) for hyperparameter optimisation. MFIN models train for a maximum of 250 epochs. Early stopping after 25 epochs of non-decreasing validation loss is used to avoid overfitting to training data. A full list of fixed and tuneable parameters is provided in Table 2.

5 RESULTS AND DISCUSSION

5.1 Metrics

We evaluate strategies using similar metrics to [21], with modifications due to dynamic volatility scaling at the portfolio level.

- (1) **Profitability:** mean annual returns (*MAR*).
- (2) **Risk:** annualised volatility (*VOL*) and maximum drawdown (*MDD*). Since strategies can have differences in *VOL*, it is most meaningful to compare *MDD* as a number of standard deviations σ (i.e. a multiplier on *VOL*).
- (3) **Risk-adjusted profitability:** annualised *Sharpe*, *Sortino*, and *Calmar* ratios.
- (4) **Correlation (*CORR*):** low (absolute) Pearson rank correlation coefficient to Long-only helps produce consistent returns and avoid drawdowns during market crashes.
- (5) **Transaction costs:** breakeven transaction cost (*BRK*) is expressed in basis points (bps), after volatility scaling at the asset level and portfolio level with $\sigma_{i,t}$ and σ_t , respectively.

$$BRK = \frac{\sum_{t=1}^T \sum_{i=1}^{N_A} \frac{w_{i,t}}{\sigma_{i,t} \cdot \sigma_t} \cdot r_{i,t+1}}{\sum_{t=1}^T \sum_{i=1}^{N_A} \left| \frac{w_{i,t}}{\sigma_{i,t} \cdot \sigma_t} - \frac{w_{i,t-1}}{\sigma_{i,t-1} \cdot \sigma_{t-1}} \right|} \quad (6)$$

- (6) **Statistical significance:** we interpret the Probabilistic Sharpe Ratio (*PSR*) [2] as significant when above a 99% confidence level against a benchmark *Sharpe* of 0. Minimum Track

Table 2: Hyperparameters for MFIN models.

Parameter	Category	Value
Objective	Fixed	Eq. 5
Max epochs	Fixed	250
Early stopping	Fixed	25
Train/valid ratio	Fixed	90%/10%
Batch size	Fixed	100
Sequence length T	Fixed	100
Activation	Fixed	ELU
C_{valid}	Fixed	0
K_{valid}	Fixed	0
C	Tuned	[0, 0.5, 1, 2, 5]
K	Tuned	[0, 1, 2, 5]
$dropout_rate$	Tuned	[0.1, 0.2, 0.3]
$learning_rate$	Tuned	[1e-3, 1e-4, 1e-5]
hidden_layer_size, N_H	Tuned	[32, 64, 96, 128]
n_filters	Tuned	[16, 32, 48, 64]
ts_filter_length	Tuned	[3, 5, 10, 15, 20]
max_epochs	Hyperband	10
hyperband_iterations	Hyperband	1
factor	Hyperband	3

Record (MTR) is the minimum number of observations, in days, needed for a significant result.

5.2 Performance of Traditional Strategies

5.2.1 Data Exploration. Data exploration deliberately overestimates strategy performance, but is useful to check for momentum and reversion behaviour in the dataset.

In Table 3, we list the top 2 feature-parameter combinations for each strategy type. We find that the number of unique sender addresses (*sent addresses*), is a strong momentum indicator and provides the highest *Sharpe* for both trend and crossover indicators.

Both fast and slow momentum are present. This can be seen from the best BAZ strategies using the shortest and longest MACD indicators, respectively, without much difference in performance before transaction costs. Notably, BAZ with *block size* has low correlation to Long-only despite a slow signal. This is because the strategy is highly correlated during the bull market of 2021, and anti-correlated through the drawdown thereafter.

Although the 6 selected strategies have similar Sharpe ratios between 1.4 to 1.7, combining information from different features is still beneficial. A simple volatility-scaled CMB portfolio increases *Sharpe* to 3.07, with a high breakeven transaction cost of 54.0 bps (basis points). This shows that alternative data can add value to cryptocurrency trading strategies. However, we expect performance of realistic strategies, even simple ones, to decay out-of-sample.

5.2.2 Realistic Strategies. An implementable strategy must instead make investment decisions based wholly on past data, without

lookahead bias from knowing which feature-parameter combinations perform well on the entire test set. This is achieved through the train-test splits described in Section 3.5.

We tabulate results for realistic MOP, BAZ, and REV strategies in Table 4. Compared to the Data Exploration strategies, performance for MOP and REV decays by approximately 50%. Only BAZ retains a similar *Sharpe*: the *sent addresses* MACD signal is consistently profitable throughout the test set and can be readily identified based on only ex-ante data. The full list of feature-parameter combinations is provided in Appendix A.

The volatility-scaled CMB portfolio still improves Sharpe, but this is less noticeable: a 13% increase from the best underlying strategy, compared to the 83% increase achieved during data exploration. As discussed in Sections 4 and 5.3, a fully data-driven approach to combining factors may be useful for extracting further value from alternative data.




Figure 4: Cumulative returns for MFIN, traditional strategies without lookahead bias and Long-only at 0 transaction costs.

5.3 Performance of MFIN Models

We apply MFIN to standardised daily returns in each factor. This reduces the dimensionality of the input and features are instead learnt automatically from the data.

As shown in Table 4, MFIN outperforms traditional strategies and benchmarks prior to transaction costs. Inspecting cumulative returns in Figure 4, MFIN performs consistently across time and is able to adapt to varying market conditions. For example, MFIN is the only strategy to avoid a large drawdown following November 2021, as investors sell risk assets including equities and cryptocurrencies. Accounting for returns, $Calmar = 2.05$ is a threefold increase compared to Long-only, and double that of CMB.


Considering fees, MFIN is also able to deliver positive cost-adjusted Sharpe ratios up to $BRK = 17.5$ bps, above the standard Binance trading fee of 10.0 bps [4]. At lower transaction costs up to 5.0 bps, MFIN maintains a higher *Sharpe* than traditional strategies and benchmarks. Beyond 7.5 bps, BAZ, CMB and Long-only outperform due to less turnover, as shown in Table 5. In particular,

Table 3: Data exploration for traditional strategies. These strategies cannot be implemented with only ex-ante information. Best values in comparable metrics are underlined.

Strategy	Feature	Signal Parameters	Sharpe	Sortino	Calmar	VOL %	MDD σ	CORR %	BRK bps
MOP	sent addresses	$k = 21$	1.67	2.62	1.82	14.0	0.98	17.9	28.3
MOP	hashrate	$k = 21$	1.49	2.27	1.39	14.0	1.14	48.6	29.9
BAZ	sent addresses	$(S_k, L_k) = (4, 12)$	1.52	2.37	1.08	14.0	1.49	-15.4	47.7
BAZ	block size	$(S_k, L_k) = (32, 96)$	1.47	2.24	1.31	14.4	1.18	<u>-0.2</u>	<u>193.7</u>
REV	coin days destroyed	$(k, z_u, z_l) = (5, 1.75, 0.75)$	1.64	2.65	1.43	13.7	1.23	83.4	35.7
REV	google trends	$(k, z_u, z_l) = (5, 1.75, 0.75)$	1.52	2.80	1.56	11.5	1.01	68.5	63.5
CMB	-	-	<u>3.07</u>	<u>5.40</u>	<u>4.50</u>	14.0	<u>0.83</u>	37.7	54.0
Long-only	-	-	1.03	1.47	0.66	14.4	1.56	-	-

Table 4: Comparison of MFIN performance against realistic traditional strategies. Best values in comparable metrics are underlined. Significant PSR, compared to benchmark Sharpe ratio of 0, are highlighted in bold.

Strategy	MAR %	HR %	PNL	Sharpe	Sortino	Calmar	VOL %	MDD σ	CORR %	BRK bps	PSR %	MTR days
MFIN	24.2	<u>54.5</u>	1.08	1.64	2.55	2.05	13.8	0.86	-8.2	17.5	99.9	727
MOP	9.7	51.1	1.07	0.74	1.11	0.75	13.7	0.95	<u>0.3</u>	9.1	93.0	3539
BAZ	17.9	52.1	1.11	1.25	1.89	0.91	13.9	1.41	-7.2	<u>36.8</u>	99.3	1259
REV	11.1	52.7	1.06	0.85	1.28	0.48	13.4	1.72	81.6	27.6	95.4	2724
CMB	21.0	52.0	<u>1.15</u>	1.42	2.21	1.01	14.1	1.48	25.0	20.4	99.8	963
CMB + MFIN	<u>29.1</u>	54.1	1.14	<u>1.89</u>	<u>3.05</u>	<u>3.21</u>	14.0	<u>0.64</u>	15.7	20.9	100.0	<u>535</u>
Long-only	14.8	53.7	99.6	1.03	1.47	0.66	14.4	1.56	-	-	97.8	1898

**Figure 5: Correlation between MFIN, traditional and benchmark strategies. Values above and below the diagonal are Pearson and Spearman correlation coefficients, respectively.****Table 5: Strategy performance at different transaction cost coefficients C . The best strategy at each C is underlined. Strategies that outperform a volatility-scaled Long-only benchmark are highlighted in bold.**

Strategy	Sharpe after C (bps) of costs					
	0.0	2.5	5.0	7.5	10.0	12.5
MFIN	1.64	1.41	1.18	0.94	0.71	0.48
MOP	0.74	0.54	0.34	0.14	-0.06	-0.26
BAZ	1.25	1.17	1.09	1.01	0.93	0.84
REV	0.85	0.77	0.70	0.63	0.55	0.48
CMB	1.42	1.25	1.08	0.90	0.73	0.56
CMB + MFIN	<u>1.89</u>	<u>1.67</u>	<u>1.44</u>	<u>1.21</u>	0.97	0.74
Long-only	1.03	1.03	1.03	1.03	<u>1.03</u>	<u>1.03</u>

BAZ trades on slower MACD signals and REV holds no position until identifying a reversion opportunity.

In Figure 5, we plot a correlation matrix for traditional, MFIN and benchmark strategy returns. Pearson correlation coefficients are shown above the leading diagonal, and these agree with the Spearman coefficients below the diagonal. We find that MFIN is minimally correlated to all strategies. This shows that MFIN learns patterns that are fundamentally different from traditional rule-based strategies. We can also show that such ML strategies contribute additively to traditional strategies: a volatility-scaled combination of the CMB and MFIN portfolio weights (CMB + MFIN) has higher $Sharpe = 1.89$ and $BRK = 20.9$ bps, whilst experiencing lower $MDD = 0.64$ σ than either underlying strategy.

6 CONCLUSION

We analyse the predictive power of alternative data for cryptocurrency trading. Momentum and reversion patterns are observed in the data, which can be combined additively to enhance individual Sharpe ratios from 1.4–1.7 to a combined $Sharpe = 3.07$. Removing lookahead bias, BAZ strategies with MACD indicators continue to perform out-of-sample, with only an 18% decay in Sharpe ratio.

We introduce Multi-Factor Inception Networks (MFIN) as a fully data-driven approach to learn from multiple factors, on multiple assets. MFIN has higher Sharpe ratio than traditional strategies and benchmarks up to 5 bps, with a breakeven transaction cost of $BRK = 17.5$ bps. For risk-sensitive investors, MFIN performs consistently over time and avoids large drawdowns, with a three-fold increase in $Calmar = 2.05$ compared to the underlying cryptocurrency market. This data-driven strategy can also be combined with traditional strategies to further improve performance to $Sharpe = 1.89$, $Calmar = 3.21$ and $BRK = 20.9$ bps.

Future work can incorporate additional alternative data, such as Reddit subscribers and Github stars, which represent the popularity of a cryptocurrency on social media and within the technical community, respectively. As MFINs are a general framework for multi-factor models, another avenue of research could focus on enhancing the performance and interpretability of the underlying Deep Inception Network (DIN). One potential DIN variant is DeepLOB-TFT, which instead learns cross-factor interactions in the interpretable Temporal Fusion Transformer [19, 32] Position Sizer.

ACKNOWLEDGMENTS

We would like to thank the Oxford-Man Institute of Quantitative Finance for computing support.

REFERENCES

- [1] Internet Archive. 2023. Wayback machine. <https://archive.org/web/>. Accessed: 2023-04-05. (2023).
- [2] David H Bailey and Marcos Lopez de Prado. 2012. The Sharpe ratio efficient frontier. *Journal of Risk*, 15, 2, 13.
- [3] Jamil Baz, Nicolas Granger, Campbell R Harvey, Nicolas Le Roux, and Sandy Rattray. 2015. Dissecting investment strategies in the cross section and time series. *Available at SSRN 2695101*.
- [4] Binance. 2023. Trading fees. <https://binance.com/en/fee/trading>. Accessed: 2023-07-10. (2023).
- [5] BitInfoCharts. 2023. BitInfoCharts. <https://bitinfocharts.com/>. Accessed: 2023-04-05. (2023).
- [6] Blockchair. 2023. Blockchair. <https://blockchair.com/>. Accessed: 2023-07-18. (2023).
- [7] John Bollinger. 2002. *Bollinger on Bollinger bands*. McGraw-Hill New York.
- [8] Guglielmo Maria Caporale and Alex Plastun. 2020. Momentum effects in the cryptocurrency market after one-day abnormal returns. *Financial Markets and Portfolio Management*, 34, 3, 251–266.
- [9] CoinMarketCap. 2023. CoinMarketCap. <https://coinmarketcap.com/>. Accessed: 2023-04-13. (2023).
- [10] Google. 2023. Google Trends. <https://trends.google.com/home>. Accessed: 2023-04-05. (2023).
- [11] Campbell R Harvey, Edward Hoyle, Russell Korgaonkar, Sandy Rattray, Matthew Sargison, and Otto Van Hemert. 2018. The impact of volatility targeting. *The Journal of Portfolio Management*, 45.
- [12] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9, 8, 1735–1780.
- [13] Stefan Hubrich. 2017. ‘Know when to hodl’em, know when to fodd’em’: an investigation of factor based investing in the cryptocurrency space.
- [14] Zhengyao Jiang and Jinjun Liang. 2017. Cryptocurrency portfolio management with deep reinforcement learning. In *2017 Intelligent systems conference (IntelliSys)*. IEEE, 905–913.
- [15] Kaggle. 2022. Binance full history. <https://kaggle.com/datasets/jorijnsmit/binance-full-history>. Accessed: 2023-07-10. (2022).
- [16] Diederik P Kingma and Jimmy Ba. 2014. Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [17] Tim Leung and Hung Nguyen. 2019. Constructing cointegrated cryptocurrency portfolios for statistical arbitrage. *Studies in Economics and Finance*.
- [18] Yi Li, Andrew Urquhart, Pengfei Wang, and Wei Zhang. 2021. Max momentum in cryptocurrency markets. *International Review of Financial Analysis*.
- [19] Bryan Lim, Serkan Ö Arık, Nicolas Loeff, and Tomas Pfister. 2021. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37.
- [20] Bryan Lim, Stefan Zohren, and Stephen Roberts. 2019. Enhancing time-series momentum strategies using deep neural networks. *The Journal of Financial Data Science*, 1, 4, 19–38.
- [21] Tom Liu, Stephen Roberts, and Stefan Zohren. 2023. Deep inception networks: a general end-to-end framework for multi-asset quantitative strategies. *arXiv preprint arXiv:2307.05522*.
- [22] Yukun Liu and Aleh Tsyvinski. 2021. Risks and returns of cryptocurrency. *The Review of Financial Studies*, 34, 6, 2689–2727.
- [23] Yukun Liu, Aleh Tsyvinski, and Xi Wu. 2022. Common risk factors in cryptocurrency. *The Journal of Finance*, 77, 2, 1133–1177.
- [24] Giorgio Lucarelli and Matteo Borrotti. 2020. A deep Q-learning portfolio management framework for the cryptocurrency market. *Neural Computing and Applications*, 32, 17229–17244.
- [25] Tobias J Moskowitz, Yao Hua Ooi, and Lasse Heje Pedersen. 2012. Time series momentum. *Journal of financial economics*, 104, 2, 228–250.
- [26] Satoshi Nakamoto and A Bitcoin. 2008. A peer-to-peer electronic cash system. *Bitcoin*.—URL: <https://bitcoin.org/bitcoin.pdf>, 4, 2, 15.
- [27] Daniel Poh, Bryan Lim, Stefan Zohren, and Stephen Roberts. 2021. Building cross-sectional systematic strategies by learning to rank. *The Journal of Financial Data Science*, 3, 2, 70–86.
- [28] Daniel Poh, Stephen Roberts, and Stefan Zohren. 2022. Transfer ranking in finance: applications to cross-sectional momentum with data scarcity. *arXiv preprint arXiv:2208.09968*.
- [29] Yinghe Qing, Jifeng Sun, Ying Kong, and Jianwu Lin. 2022. Fundamental multi-factor deep-learning strategy for cryptocurrency trading. In *2022 IEEE 20th International Conference on Industrial Informatics (INDIN)*. IEEE, 674–680.
- [30] Thomas Stridsman. 1998. Data pros and cons. *Futures*, 27, 6, 54–56.
- [31] David Vidal-Tomás. 2022. Which cryptocurrency data sources should scholars use? *International Review of Financial Analysis*, 81, 102061.
- [32] Kieran Wood, Sven Giegerich, Stephen Roberts, and Stefan Zohren. 2021. Trading with the momentum transformer: an intelligent and interpretable architecture. *arXiv preprint arXiv:2112.08534*.
- [33] Kieran Wood, Stephen Roberts, and Stefan Zohren. 2022. Slow momentum with fast reversion: a trading strategy using deep learning and changepoint detection. *The Journal of Financial Data Science*, 4, 1, 111–129.
- [34] Chao Zhang, Zihao Zhang, Mihai Cucuringu, and Stefan Zohren. 2021. A universal end-to-end approach to portfolio optimization via deep learning. *arXiv preprint arXiv:2111.09170*.
- [35] Zihao Zhang, Stefan Zohren, and Stephen Roberts. 2019. Deeplob: deep convolutional neural networks for limit order books. *IEEE Transactions on Signal Processing*, 67, 11, 3001–3012.

Table 6: Feature-parameter selections for rule-based strategies.

Test set	MOP		BAZ		REV	
	Feature	Parameters	Feature	Parameters	Feature	Parameters
2019 - 2020	transactions	$k = 5$	sent addresses	$(S_k, L_k) = (4, 12)$	sent USD	$(k, z_u, z_l) = (21, 1.75, 0.75)$
	sent addresses	$k = 5$	confirmation time	$(S_k, L_k) = (16, 48)$	google trends	$(k, z_u, z_l) = (5, 1.5, 1.00)$
2020 - 2021	transactions	$k = 5$	confirmation time	$(S_k, L_k) = (32, 96)$	sent USD	$(k, z_u, z_l) = (21, 1.75, 0.75)$
	sent USD	$k = 5$	sent addresses	$(S_k, L_k) = (8, 24)$	fee-reward ratio	$(k, z_u, z_l) = (5, 1.5, 0.5)$
2021 - 2022	confirmation time	$k = 126$	confirmation time	$(S_k, L_k) = (16, 48)$	sent USD	$(k, z_u, z_l) = (21, 1.75, 0.75)$
	sent addresses	$k = 21$	sent addresses	$(S_k, L_k) = (4, 12)$	fee-reward ratio	$(k, z_u, z_l) = (5, 1.5, 0.5)$
2022 - 2023	sent addresses	$k = 21$	sent addresses	$(S_k, L_k) = (4, 12)$	sent USD	$(k, z_u, z_l) = (21, 1.75, 0.75)$
	coin days destroyed	$k = 126$	transactions	$(S_k, L_k) = (8, 24)$	coin days destroyed	$(k, z_u, z_l) = (5, 1.75, 0.75)$

Table 7: Number of trainable parameters N_p for MFIN.

		N_A			
		1	7	20	50
N_H	32	64K	170K	399K	929K
	64	80K	185K	415K	945K
$(n_filters = 40)$	96	103K	209K	438K	968K
	$(ts_filter_length = 10)$ 128	135K	241K	470K	1,000K
$n_filters$	16	49K	66K	103K	190K
	32	74K	142K	289K	629K
$(N_H = 80)$	48	110K	262K	591K	1,354K
	$(ts_filter_length = 10)$ 64	156K	426K	1,012K	2,366K
ts_filter_length	3	68K	107K	190K	384K
	5	74K	132K	257K	547K
	10	90K	196K	425K	955K
$(N_H = 80)$	15	106K	260K	593K	1,363K
	$(n_filters = 40)$ 20	122K	324K	761K	1,771K
Datapoints		42K	295K	843K	2,108K

A FEATURE-PARAMETER SELECTIONS

In Table 6, we record the feature-parameter combinations selected by realistic traditional strategies in Section 5.2.2. BAZ performs best and correctly selects *sent addresses* with short timescales in all test sets. In contrast, MOP only identifies the optimal timescale for *sent addresses* in the final two test sets. REV consistently chooses *sent USD*, which fails to generalise to unseen regimes: a positive *Sharpe* = **1.26** in the first 3 test sets is offset by *Sharpe* = **-2.31** in the final test set.

B MFIN COMPLEXITY

For a given number of assets N_A , the complexity of MFIN models depends on hyperparameters N_H , $n_filters$, and ts_filter_length . To inspect the impact of one hyperparameter, we keep the others fixed at a median value. In OrigCIM-LSTM, the Feature Extractor contributes the majority of parameters (sometimes more than 90%) via convolutional filters. With alternative DIN variants, such as DeepLOB-TFT, which only learn cross-factor interactions in the Position Sizer, there can be a more balanced contribution.