

Spatio-Temporal Momentum: Jointly Learning Time-Series and Cross-Sectional Strategies

Wee Ling Tan, Stephen Roberts, and Stefan Zohren

Wee Ling Tan

is a doctoral candidate with the Machine Learning Research Group and the Oxford-Man Institute of Quantitative Finance at the University of Oxford in Oxford, UK.

weeling@robots.ox.ac.uk

Stephen Roberts

is the RAEng/Man Group Professor of Machine Learning at the University of Oxford and a faculty member of the Oxford-Man Institute of Quantitative Finance in Oxford, UK.

sjrob@robots.ox.ac.uk

Stefan Zohren

is a research fellow in financial machine learning at the Oxford-Man Institute of Quantitative Finance and a faculty member of the Department of Engineering Science, both at the University of Oxford in Oxford, UK, and is a fellow at the Turing Institute and a principal quant at Man Group plc in London, UK.

stefan.zohren@eng.ox.ac.uk

KEY FINDINGS

- Although both time-series and cross-sectional momentum strategies are designed to systematically capture momentum risk premiums, these strategies are regarded as distinct implementations and do not account for the concurrent predictability between time-series and cross-sectional momentum features of different assets.
- The work introduces spatio-temporal momentum, a class of data-driven strategies that simultaneously learns both time-series and cross-sectional momentum in a single framework.
- Backtesting on portfolios of equities and continuous futures contracts, spatio-temporal momentum was shown to outperform traditional strategies in the presence of transaction costs of up to 5–10 basis points, demonstrating superior performance over a simple blend of time-series and cross-sectional momentum.

ABSTRACT

The authors introduce spatio-temporal momentum strategies, a class of models that unify both time-series and cross-sectional momentum strategies by trading assets based on their cross-sectional momentum features over time. Although both time-series and cross-sectional momentum strategies are designed to systematically capture momentum risk premiums, these strategies are regarded as distinct implementations and do not consider the concurrent relationship and predictability between temporal and cross-sectional momentum features of different assets. They model spatio-temporal momentum with neural networks of varying complexities and demonstrate that a simple neural network with only a single fully connected layer learns to simultaneously generate trading signals for all assets in a portfolio by incorporating both their time-series and cross-sectional momentum features. Back testing on portfolios of 46 actively traded US equities and 12 equity index futures contracts, they demonstrate that the model is able to retain its performance over benchmarks in the presence of high transaction costs of up to 5–10 basis points. In particular, they find that the model when coupled with least absolute shrinkage and turnover regularization results in the best performance over various transaction cost scenarios.

Momentum strategies form a class of systematic strategies that rely on the premise of a persistence in the direction of asset returns over time (Jegadeesh and Titman 1993; Moskowitz, Ooi, and Pedersen 2012). These strategies are constructed to exploit the continuation of the underlying trend by increasing or assuming long positions during uptrends and decreasing or assuming short positions

during downtrends. Momentum strategies are often accompanied by volatility targeting, allowing momentum strategies to leverage positions taken during periods of low volatility, and reduce exposures during periods of high volatility. Volatility targeting at both the asset and portfolio level has been shown to boost Sharpe ratios, reducing the probability of extreme tail returns and minimizing maximum drawdowns for risk asset portfolios (Kim, Tse, and Wald 2016; Harvey et al. 2018).

Momentum strategies can be distinguished as belonging to a time-series or cross-sectional strategy. In terms of predictability, the former is primarily driven by the persistence in the trend of individual and market-level returns whereas the latter is usually deployed as a relative value strategy due to its perceived market neutrality (Baz et al. 2015). As such, current work has mainly regarded time-series and cross-sectional momentum as distinct implementations.

In time-series momentum strategies (Moskowitz, Ooi, and Pedersen 2012; Lim, Zohren, and Roberts 2019; Garg et al. 2021; Wood, Roberts, and Zohren 2022), trading signals for individual assets are constructed based on the asset's own historical returns. In addition, the signals for individual assets in the portfolio are typically constructed independently of other assets. In the presence of a defined universe or portfolio of assets, this strategy does not take into consideration any form of mutual interactions and predictability between assets. We believe that momentum features from other assets represent sources of information because these assets collectively represent the market state at any point in time and should be considered when constructing trading signals for a given asset.

On the other hand, cross-sectional momentum strategies (Jegadeesh and Titman 1993; Rouwenhorst 1998; Pirrong 2005; Poh et al. 2021, 2022) require a momentum score to first be quantified for each individual asset in the portfolio before computing a relative ranking of these scores in order to formulate positions for a select group of assets. In the first step of calculating a momentum score, cross-sectional momentum strategies ultimately consider only an asset's own historical returns, independent of returns from other assets. In principle, this is still identical to the way a trading signal is constructed in a time-series momentum strategy. Subsequently in the ranking step, a typical cross-sectional momentum strategy would allocate positions by assuming a maximum long position for assets ranked in the top decile, while taking a maximum short position for assets ranked in the bottom decile. We argue that this maximum long-short allocation does not accurately reflect the underlying trading signal for the selected assets. Moreover, this approach leaves a large majority of assets classified in intermediate deciles with no directional positions. As such, a cross-sectional strategy is unable to exploit any underlying trends for these intermediate assets.

In this work, we combine both types of strategies by considering a form of multi-asset momentum that simultaneously constructs trading signals based on momentum features from multiple assets over time. We model the spatio-temporal momentum (STMOM) strategy using a variety of neural networks trained in a data-driven manner, each representing different complexities. In predicting the trend and position size for any given asset, we consider the asset's own momentum features as well as the momentum features from other assets in the portfolio. In contrast to a time-series momentum strategy, our model learns to incorporate information collectively from the universe of assets into signal construction. Our proposed strategy also directly generates trading signals for every asset as a multitarget output of the neural network, effectively bypassing the need to manually rank assets relative to one another as in the case of cross-sectional momentum. We directly optimize the models with the Sharpe ratio (Sharpe 1994), allowing the STMOM networks to learn from risk-adjusted performance (Lim, Zohren, and Roberts 2019).

Our strategy closely resembles multitask learning (Caruana 1997; Crawshaw 2020), a training framework that aims to train a learner on multiple tasks simultaneously.

We observe that the prediction of momentum signals for different assets can be treated as multiple different but closely related prediction tasks. In the context of STMOM, the model leverages on the usefulness of a shared feature representation between different assets and is trained by optimizing over an aggregate Sharpe ratio computed from the multioutput model.

RELATED WORK

Momentum strategies are typically classified as either time-series or cross sectional. The work of Moskowitz, Ooi, and Pedersen (2012) incorporates volatility scaling (Harvey et al. 2018) into a time-series momentum strategy and documents significant excess returns arising from trading based on the sign of an asset's own returns over the past 12 months, backtesting across 58 instruments and more than 25 years of data. Since then, other works have introduced more complex trend estimation techniques, such as using volatility normalized moving average convergence divergence (MACD) signals (Baz et al. 2015) and blending slow and fast momentum strategies with various weights to reduce downside exposure (Garg et al. 2021). On the other hand, Jegadeesh and Titman (1993) show that a cross-sectional momentum strategy of buying winners and selling losers over a lookback horizon of 3 to 12 months led to significant excess returns for a portfolio of NYSE and AMEX stocks. The cross-sectional momentum effect has also been observed in international equity markets (Rouwenhorst 1998) and futures contracts (Pirrong 2005).

Machine learning algorithms have been increasingly developed to perform predictions by extracting and modeling features from data. Deep learning has eliminated the need for manual feature engineering, allowing for hierarchical feature representations to be learned directly from data (LeCun, Bengio, and Hinton 2015). With recent advances in deep learning methods and open source libraries (Abadi et al. 2016; Paszke et al. 2019), deep neural networks have been applied to model financial datasets such as high-frequency limit order book data (Zhang, Zohren, and Roberts 2019) and to construct portfolios (Zhang, Zohren, and Roberts 2020). Very often, neural networks that depend on recurrence and backpropagation through time, like recurrent neural networks (RNN) and long short-term memory networks (LSTM) (Hochreiter and Schmidhuber 1997), are used to model temporal relationships commonly present in financial data.

Deep learning algorithms have also been applied to momentum strategies. The work of Lim, Zohren, and Roberts (2019) introduces deep momentum networks (DMNs), a series of neural network architectures directly optimized for portfolio Sharpe ratio, in place of standard regression and classification methods for time-series momentum. The models directly output positions for individual contracts, combining trend estimation and position sizing into a single function. However, we note that the single-output DMN model is trained on batches of input features belonging to different asset classes. This requires the model to learn to construct and output trading signals individually for all types of instruments, regardless of the possibility of negative transfer between asset classes. Our approach adopts a multitask learning approach that incorporates a multioutput target, permitting specialization for individual assets while retaining the advantage of having only a single model that leverages on a common shared representation between different assets.

Multitask learning (Caruana 1997; Ruder 2017; Crawshaw 2020) aims to train a model on multiple tasks simultaneously to improve generalization performance across tasks. It has been successfully adopted in a range of applications, including computer vision (Zhang et al. 2014; Misra et al. 2016; Long et al. 2017) and natural language processing (Dong et al. 2015; Luong et al. 2015; Rei 2017). Training a deep

learning-based multitask model involves backpropagation with training signals from multiple related tasks. This can be seen as a form of model regularization via shared representations, specifically by coupling a main task with auxiliary tasks to introduce an inductive bias to the model (Ruder 2017). We take the approach in which all tasks (assets in the portfolio) are treated equally, unlike some multitask learning models that distinguish between a main task and one or more auxiliary tasks.

MOMENTUM STRATEGIES

Following the definition of Moskowitz, Ooi, and Pedersen (2012) and Baltas and Kosowski (2020), the overall returns of a time-series momentum (TSMOM) strategy that equally diversifies over N_t assets at time t is

$$r_{t,t+1}^{\text{TSMOM}} = \frac{1}{N_t} \sum_{i=1}^{N_t} X_t^{(i)} \frac{\sigma_{\text{tgt}}^{(i)}}{\sigma_t^{(i)}} r_{t,t+1}^{(i)} \quad (1)$$

where $X_t^{(i)} \in [-1, 1]$ denotes the trading signal or position for asset i at time t . Given the differences in volatility across individual assets, we scale the realized returns $r_{t,t+1}^{(i)}$ by their volatility to target equal risk assignments. We set the annualized volatility target σ_{tgt} to be 15% and estimate the ex ante volatility $\sigma_t^{(i)}$ with a 60-day exponentially weighted moving standard deviation of daily returns.

Construction of Trading Signals

Most momentum strategies are concerned with designing a proper trading signal $X_t^{(i)}$. We illustrate this with the following examples that we incorporate as benchmarks in our work.

Long only. In the simplest case, a long-only strategy takes a maximum long position $X_t^{(i)} = 1$ for each asset and performs a daily rebalancing according to changes in the ex ante volatility estimate.

TSMOM. In the TSMOM strategy of Moskowitz, Ooi, and Pedersen (2012), the position taken for an asset is based on the sign of the asset's returns over the past 12 months: $X_t^{(i)} = \text{sgn}(r_{t-252,t}^{(i)})$.

MACD. In Baz et al. (2015), volatility normalized MACD indicators are used in place of the sign of returns as signals for trend estimation:

$$\begin{aligned} \text{MACD}(i, t, S, L) &= m(i, t, S) - m(i, t, L) \\ \text{MACD}_{\text{norm}}(i, t, S, L) &= \frac{\text{MACD}(i, t, S, L)}{\text{std}(p_{t-63:t}^{(i)})} \\ Y_t^{(i)} &= \frac{\text{MACD}_{\text{norm}}(i, t, S, L)}{\text{std}(\text{MACD}_{\text{norm}}(i, t-252:t, S, L))} \\ X_t^{(i)} &:= \tilde{Y}_t^{(i)} = \frac{1}{3} \sum_{k=1}^3 \phi(Y_t^{(i)}(S_k, L_k)) \end{aligned} \quad (2)$$

where $\text{MACD}(i, t, S, L)$ is the MACD value of asset i at time t with a short time scale S and long time scale L . Further, $m(i, t, j)$ is defined as the exponentially weighted moving average of asset i prices at time t , with a time scale j that corresponds to a half-life of $HL = \log(0.5)/\log\left(1 - \frac{1}{j}\right)$. The MACD value is normalized by $\text{std}(p_{t-63:t}^{(i)})$, the 63-day rolling standard deviation of asset i prices. Multiple intermediate MACD

signals over different short and long time scales $S_k \in \{8, 16, 32\}$ and $L_k \in \{24, 48, 96\}$ are combined in an equally weighted sum to yield a position $X_t^{(i)}$ (or an aggregated

MACD signal $\tilde{Y}_t^{(i)}$), where $\phi(y) = \frac{y \exp(\frac{-y^2}{4})}{0.89}$ is a response function as defined in Baz et al. (2015).

DMNs. Lim, Zohren, and Roberts (2019) use deep neural networks to directly generate trading signals for individual assets, combining trend estimation and position sizing into a single function as approximated by the learned model f using time-series momentum features $\mathbf{u}_t^{(i)}$:

$$X_t^{(i)} = f(\mathbf{u}_t^{(i)}; \theta) \quad (3)$$

The model parameters θ are learned by optimizing over loss metrics that include volatility characteristics of returns. As a benchmark, we consider f to be the Sharpe-optimized LSTM.

CSMOM. Following Jegadeesh and Titman (1993), we consider a cross-sectional momentum (CSMOM) strategy that scores and ranks an asset based on its returns computed over the past 12 months. The strategy utilizes a decile portfolio, taking a maximum long and short position for the top and bottom 10% of ranked assets, respectively.

SPATIO-TEMPORAL MOMENTUM

Multitask Learning Premise

We observe that the prediction of individual momentum signals $X_t^{(i)}$ for different assets can be treated as multiple different but closely related prediction tasks. In the context of an equally weighted portfolio, we take the approach in which all assets (tasks) are treated equally in a multitask setting. In addition, we follow the definition of multitask learning in Caruana (1997), in which all tasks share the same inputs, unlike some definitions of multitask learning in which different inputs are utilized for different tasks.

Multitask Learning in STMOM

We model the simultaneous prediction of multiple trading signals $X_t^{(i)}$ with model f as a multitask learning problem, with each task i corresponding to predicting the trading signal for asset i . Given time t , our goal is to construct a multitask learning model f that is able to perform N^t tasks of predicting an aggregate trading signal $\mathbf{X}_t \in [-1, 1]^{N^t}$ over all assets.

In general, we have an input spatio-temporal tensor $\mathbf{u}_t \in \mathbb{R}^{N^t \times \tau \times d}$, where N^t is the number of assets, τ is the temporal history, d is the number of features, and $\mathbf{u}_t(i, j, k)$ represents the k -th feature of the i -th asset at time $t - j$. We want to learn a model f parameterized by θ :

$$\mathbf{X}_t = f(\mathbf{u}_t; \theta) \quad (4)$$

where

$$\mathbf{X}_t = \begin{bmatrix} X_t^{(1)} \\ X_t^{(2)} \\ \vdots \\ X_t^{(N^t)} \end{bmatrix} \quad (5)$$

In this framework, trading signals $X_t^{(i)}$ are directly computed using features from the universe of assets in the form of the spatio-temporal tensor \mathbf{u}_t . Specifically, this reduces to a time-series momentum strategy for $N^t = 1$. The realized returns of this framework still follows Equation 1.

Deep Learning Architectures

Given that the choice of model architecture is crucial in modeling the relationship between the trading signals and the STMOM features, we examine a range of architectures of different complexities that are able to serve as candidate end-to-end functions.

Single layer perceptron (SLP). We consider the simplest case of a fully connected neural network with a single hidden layer that computes a linear combination of the input features prior to activation:

$$\mathbf{X}_t = f(\mathbf{u}_t; \theta) = g(\mathbf{W}^\top \mathbf{u}_t + \mathbf{b}) \quad (6)$$

where $\mathbf{W} \in \mathbb{R}^{m \times N^t}$, $\mathbf{u}_t \in \mathbb{R}^m$ with $m = N^t \cdot \tau \cdot d$, $\mathbf{b} \in \mathbb{R}^{N^t}$ and $g = \tanh$ is the activation function.

Multilayer perceptron (MLP). We consider a fully connected neural network with two hidden layers, representing a step up in model complexity from the SLP model:

$$\mathbf{X}_t = f(\mathbf{u}_t; \theta) = g[\mathbf{W}^{[2]\top} \sigma(\mathbf{W}^{[1]\top} \mathbf{u}_t + \mathbf{b}^{[1]}) + \mathbf{b}^{[2]}] \quad (7)$$

where $\mathbf{W}^{[j]}$ and $\mathbf{b}^{[j]}$ are the weights and biases of the hidden layers and $g = \sigma = \tanh$ are the activation functions. Given the nonlinear activation function σ , the model incorporates nonlinearity with respect to the features of \mathbf{u}_t prior to mapping into trading signals.

Convolutional neural networks (CNNs). Apart from their applications in computer vision (Simonyan and Zisserman 2014; He et al. 2016), CNNs have been used to model multivariate time-series by incorporating convolutional filters that extract features from temporal data (Binkowski, Marti, and Donnat 2018). These models rely on autoregressive causal convolutions and optionally, dilated convolutions (Oord et al. 2016) that allow the encoding of features over various receptive fields. The use of causal convolutions preserves the autoregressive ordering of temporal features in that predictions made at time t cannot depend on any future time steps $t + 1, \dots, T$.

Long Short-Term Memory (LSTM). Given the issues with exploding and (primarily) vanishing gradients in learning neural networks especially with long sequential data (Bengio, Simard, and Frasconi 1994), RNNs that incorporate memory cell states and gating mechanisms such as LSTMs (Hochreiter and Schmidhuber 1997) have been used to address these limitations (Lim and Zohren 2021). The combination of memory cell states and gating functions helps in backpropagating gradients through time, allowing RNNs to better learn long-range dependencies in sequences.

We consider both the CNN and LSTM architectures for modeling STMOM. For full details of the implementations, we refer the reader to Appendix A.

Training Details

Loss function. Incorporating risk-adjusted metrics such as the Sharpe ratio during optimization have allowed models to better incorporate risk into learning and generation of trading signals (Lim, Zohren, and Roberts 2019). Given a set $\mathcal{D} = \{(\mathbf{u}_t, \mathbf{X}_t = f(\mathbf{u}_t; \theta)) \mid \mathbf{X}_t \in [-1, 1]^{N^t}\}_{t=1}^T$ of spatio-temporal tensors and their corresponding signals, we define the loss function $\mathcal{L}_{\text{Sharpe}}(\theta)$ over \mathcal{D} as the annualized Sharpe ratio:

$$\mathcal{L}_{\text{sharpe}}(\theta) = \sum_{i=1}^{N^t} \lambda_i \cdot \mathcal{L}_{\text{sharpe}}^{(i)}(\theta) \quad (8)$$

$$\mathcal{L}_{\text{sharpe}}^{(i)}(\theta) = - \frac{\sum_{t=1}^T R_i(t) \times \sqrt{252}}{\sum_{t=1}^T R_i(t)^2 - \left[\sum_{t=1}^T R_i(t) \right]^2} \quad (9)$$

$$R_i(t) = X_t^{(i)} \frac{\sigma_{\text{tgt}}^{(i)}}{\sigma_t^{(i)}} r_{t,t+1}^{(i)} \quad (10)$$

where $R_i(t)$ represents the volatility-scaled captured returns for asset i from time t to $t + 1$. The multitask loss weights λ_i balance task importance for the individual task-specific loss functions $\mathcal{L}_{\text{sharpe}}^{(i)}(\theta)$. In the case in which all tasks (assets) are treated equally, we have $\lambda_i = 1/N^t$.

Shrinkage penalty. Taking into account the high dimensionality of the weight matrix \mathbf{W} as per Equation 6, we incorporate L_1 regularization as an additional penalty term to the loss function of the SLP. This encourages feature selection and model sparsity as the weight coefficients of features that are of less relevance to prediction are shrunk toward zero. We incorporate L_1 regularization in the form of a penalty term on the sum of absolute weights of matrix \mathbf{W} in the overall loss function:

$$\mathcal{L}(\theta) = \mathcal{L}_{\text{sharpe}}(\theta) + \alpha \sum_{w_{ij} \in \mathbf{W}} |w_{ij}| \quad (11)$$

where α is a hyperparameter controlling the shrinkage penalty term.

Optimization. In processing our datasets, we conduct a train-validation split with the earlier 90% of data for training and the most recent 10% reserved for validation. To minimize the overall empirical loss, we perform backpropagation using minibatch stochastic gradient descent with the Adam optimizer (Kingma and Ba 2014) and initiate early stopping based on the validation set loss. For all machine learning methods, we conduct hyperparameter optimization with 100 iterations of random search for selecting optimal candidate models. We refer the reader to Appendix A for a detailed description of training parameters and specific details for the models.

PERFORMANCE EVALUATION

Overview of Datasets

US equities. Our first dataset consists of 46 actively traded US equities from the financials sector with data obtained from the Center for Research in Security Prices (CRSP) (2022). We work with daily returns data, rebalancing the portfolio daily and performing strategy backtesting from 1990 to 2022.

Equity index futures. Our second dataset comprises 12 ratio-adjusted continuous equity index futures contracts with data obtained from the Pinnacle Data Corp. (2021) CLC Database. We work with daily returns data, performing portfolio rebalancing daily and backtesting from 2003 to 2020. The full list of instruments is detailed in Appendix B.

Backtest Details

We utilize an expanding window approach, in which all models are trained with every iteration of five additional years as it becomes available. Taking US equities as an example, the first iteration would involve training and validating on the period from 1990 to 1995, then fixing the model weights and evaluating the model on out-of-sample data from 1995 to 2000. The second iteration would involve training and validating from 1990 to 2000 and evaluating from 2000 to 2005, and so on. We conduct each experiment over multiple random seeded runs and report the aggregate performance.

Momentum Features

In constructing a general input spatio-temporal tensor $\mathbf{u}_t \in \mathbb{R}^{N^t \times \tau \times d}$ as per Equation 4, we include the below d momentum features:

1. Volatility normalized returns—We use $r_{t-k,t}^{(i)} / (\sigma_t^{(i)} \sqrt{k})$, representing asset returns normalized by daily volatility estimates scaled to a time scale $k \in \{1, 20, 63, 126, 252\}$, corresponding to daily, monthly, quarterly, semiannual, and annual returns.
2. MACD—We take volatility normalized MACD signals $Y_t^{(i)}(S_k, L_k)$ as per Equation 2 as input features, using short and long time scales $S_k \in \{8, 16, 32\}$ and $L_k \in \{24, 48, 96\}$.

Results and Discussion

In evaluating the performance of all strategies, we use the following annualized metrics:

1. Profitability—expected returns ($\mathbb{E}[\text{Returns}]$), hit rate
2. Risk—volatility (vol.), downside deviation, maximum drawdown (MDD)
3. Performance ratios—Sharpe, Sortino, and Calmar ratios, average profit over loss $\left(\frac{\text{Ave. P}}{\text{Ave. L}} \right)$

For US equities (and, respectively, for equity index futures), we report the aggregated out-of-sample performance of all strategies from 1995 to 2022 (2008 to 2020) for overall returns computed in accordance with Equation 1. In this section, we compute the performance of the strategies in the absence of transaction costs to understand the raw predictive ability of the strategies. We provide an analysis of the impact of transaction costs in a later section. We first present the performance of all strategies from their raw signal outputs in Exhibit 1 (Exhibit 3). In order to facilitate the comparison between different strategies, we apply to all strategies an additional layer of volatility scaling at the portfolio level to target an annualized volatility of 15% and report the performance in Exhibit 2 (Exhibit 4) and cumulative returns in Exhibit 5 (Exhibit 6).

US Equities. From Exhibit 1, we first observe that the cross-sectional decile portfolio was an unprofitable strategy as seen from the CSMOM strategy delivering negative returns over the backtest period. In addition, classical time-series momentum portfolios like TSMOM and MACD generally underperformed compared to a simple long-only approach. Two machine learning methods, the DMN and SLP, were able to outperform the long-only approach as seen from their performance ratios.

EXHIBIT 1

Performance Metrics for Strategies—Raw Signal Outputs (US equities)

	E[Return]	Vol.	Downside Deviation	MDD	Sharpe	Sortino	Calmar	Hit Rate	Ave. P Ave. L
Benchmarks									
Long Only	0.068	0.102	0.072	0.235	0.667	0.944	0.290	0.541	0.951
TSMOM	0.012	0.067	0.050	0.287	0.177	0.240	0.042	0.526	0.932
MACD	0.001	0.045	0.033	0.237	0.021	0.029	0.004	0.519	0.931
CSMOM	−0.033	0.048	0.036	0.631	−0.702	−0.937	−0.053	0.494	0.911
Reference									
DMN	0.056	0.028	0.018	0.067	2.043	3.145	0.924	0.598	1.055
	(0.008)	(0.008)	(0.005)	(0.029)	(0.263)	(0.452)	(0.258)	(0.005)	(0.054)
STMOM									
SLP	0.032	0.029	0.020	0.075	1.114	1.609	0.435	0.581	0.955
	(0.006)	(0.005)	(0.004)	(0.015)	(0.182)	(0.288)	(0.092)	(0.007)	(0.041)
MLP	0.013	0.044	0.032	0.176	0.296	0.410	0.082	0.544	0.902
	(0.005)	(0.008)	(0.006)	(0.044)	(0.106)	(0.148)	(0.040)	(0.007)	(0.038)
CNN	0.010	0.051	0.037	0.167	0.195	0.273	0.066	0.516	0.980
	(0.006)	(0.006)	(0.005)	(0.037)	(0.106)	(0.150)	(0.043)	(0.007)	(0.029)
LSTM	0.014	0.044	0.031	0.159	0.320	0.458	0.103	0.546	0.903
	(0.004)	(0.008)	(0.006)	(0.053)	(0.108)	(0.158)	(0.061)	(0.009)	(0.036)

NOTE: Standard deviation shown in parentheses.

EXHIBIT 2

Performance Metrics for Strategies—Rescaled to Target Volatility (US equities)

	E[Return]	Vol.	Downside Deviation	MDD	Sharpe	Sortino	Calmar	Hit Rate	Ave. P Ave. L
Benchmarks									
Long Only	0.131	0.155	0.109	0.344	0.841	1.197	0.380	0.541	0.976
TSMOM	0.056	0.157	0.112	0.470	0.358	0.501	0.119	0.526	0.960
MACD	0.038	0.157	0.112	0.524	0.245	0.343	0.073	0.519	0.968
CSMOM	−0.101	0.154	0.115	0.964	−0.655	−0.880	−0.105	0.494	0.919
Reference									
DMN	0.487	0.167	0.105	0.260	2.920	4.647	1.887	0.598	1.179
	(0.019)	(0.001)	(0.002)	(0.029)	(0.119)	(0.229)	(0.181)	(0.005)	(0.020)
STMOM									
SLP	0.423	0.162	0.102	0.301	2.609	4.161	1.428	0.581	1.170
	(0.048)	(0.001)	(0.002)	(0.033)	(0.282)	(0.524)	(0.252)	(0.007)	(0.034)
MLP	0.175	0.169	0.110	0.410	1.040	1.590	0.439	0.544	1.032
	(0.030)	(0.013)	(0.002)	(0.070)	(0.178)	(0.270)	(0.109)	(0.007)	(0.020)
CNN	0.035	0.200	0.139	0.665	0.192	0.314	0.079	0.516	0.977
	(0.041)	(0.059)	(0.056)	(0.318)	(0.223)	(0.340)	(0.082)	(0.007)	(0.032)
LSTM	0.178	0.182	0.136	0.544	1.015	1.422	0.405	0.546	1.025
	(0.040)	(0.028)	(0.035)	(0.264)	(0.305)	(0.511)	(0.202)	(0.009)	(0.028)

With the introduction of volatility scaling at the portfolio level, we observe improvements in performance across all strategies as shown in Exhibit 2. In particular, we see that STMOM methods, with the exception of CNN, experienced the largest increase in performance ratios among machine learning methods, with the Sharpe ratio of the SLP increasing by about 134% (MLP: 251%; LSTM: 217%) as compared to an increase of 43% for the DMN. This increase led to the MLP and LSTM outperforming the long-only strategy. Both the DMN and SLP demonstrated the best performance and exhibited a large gap above all other methods as seen from Exhibit 2 and the cumulative returns plot in Exhibit 5. This can be attributed to the DMN and SLP both capturing higher expected returns while being subject to lower downside risks.

Within STMOM models, we observe a deterioration in performance of the STMOM strategy with an increased level of model complexity, as seen from the underperformance of the MLP, LSTM, and CNN relative to a simple SLP model trained with a shrinkage penalty. This is contrary to the expectation that a model of greater complexity would be better suited to model the dynamics of STMOM. It is also possible that the underperformance in complex architectures is associated with difficulties in training models with multiple model configurations, especially the CNN as seen from its poor performance, echoing a similar conclusion from Lim, Zohren, and Roberts (2019).

The underperformance of some STMOM models as compared to the DMN may be attributed to the relatively low signal-to-noise ratio of returns data, coupled with the limited amount of data available to STMOM strategies. During training, STMOM deep learners are exposed to only t samples, whereas deep TSMOM strategies like the DMN have access to a much larger pool of $t \times N^i$ samples. As a result, overly complex STMOM models also run the risks of overfitting to in-sample data.

Equity index futures. Similar to US equities, it is clear from Exhibit 3 that the cross-sectional decile portfolio was again unprofitable, with the CSMOM strategy delivering negative returns. Both classical time-series momentum portfolios like TSMOM and MACD and the reference DMN were profitable but underperformed a long-only approach. With the exception of CNN, all STMOM methods demonstrated the best performance, with the SLP and MLP models outperforming the DMN by more than four times, as seen from their performance ratios.

Incorporating volatility scaling at the portfolio level, we notice improvements in the performance for all strategies as shown in Exhibit 4. The increase in performance ratios for the benchmarks strategies and DMN were minimal, whereas STMOM methods experienced larger increases in their performance ratios with volatility scaling. This increase drove STMOM strategies to further outperform all other strategies, with the SLP model outperforming the DMN by more than six times in risk-adjusted returns. Echoing the observation for US equities, the simplest SLP model again demonstrated the best performance above all other STMOM methods as seen from Exhibit 4 and the cumulative returns plot in Exhibit 6, further lending support to the notion that a model of lower complexity is better suited to model STMOM.

For subsequent sections, we focus our analysis on the US equities dataset.

Signal Diversification

From Exhibit 7, the SLP exhibits mostly zero to moderately positive correlation with other momentum strategies, with the correlation between SLP and DMN at about 46%. The correlations of the SLP with other strategies are also fairly unstable, with the SLP periodically displaying a negative correlation with TSMOM, CSMOM, and DMN, translating into possible benefits from signal diversification.

EXHIBIT 3

Performance Metrics for Strategies—Raw Signal Outputs (equity index futures)

	E[Return]	Vol.	Downside Deviation	MDD	Sharpe	Sortino	Calmar	Hit Rate	Ave. P Ave. L
Benchmarks									
Long Only	0.054	0.125	0.093	0.275	0.427	0.574	0.195	0.549	0.883
TSMOM	0.020	0.107	0.078	0.251	0.191	0.262	0.081	0.523	0.944
MACD	0.006	0.068	0.050	0.175	0.081	0.111	0.032	0.529	0.905
CSMOM	−0.048	0.081	0.060	0.554	−0.584	−0.786	−0.086	0.489	0.946
Reference									
DMN	0.012 (0.017)	0.034 (0.024)	0.026 (0.018)	0.102 (0.066)	0.301 (0.220)	0.410 (0.307)	0.110 (0.088)	0.525 (0.009)	0.971 (0.042)
STMOM									
SLP	0.047 (0.017)	0.049 (0.019)	0.035 (0.015)	0.152 (0.081)	1.242 (0.834)	1.856 (1.433)	0.657 (0.781)	0.574 (0.018)	0.992 (0.142)
MLP	0.037 (0.014)	0.032 (0.013)	0.022 (0.010)	0.088 (0.063)	1.319 (0.593)	2.050 (1.065)	0.683 (0.478)	0.559 (0.014)	1.075 (0.127)
CNN	−0.016 (0.014)	0.065 (0.020)	0.049 (0.015)	0.291 (0.111)	−0.215 (0.170)	−0.283 (0.228)	−0.045 (0.033)	0.518 (0.014)	0.886 (0.048)
LSTM	0.039 (0.013)	0.057 (0.014)	0.040 (0.010)	0.172 (0.070)	0.746 (0.329)	1.079 (0.529)	0.289 (0.198)	0.550 (0.016)	0.973 (0.098)

EXHIBIT 4


Performance Metrics for Strategies—Rescaled to Target Volatility (equity index futures)

	E[Return]	Vol.	Downside Deviation	MDD	Sharpe	Sortino	Calmar	Hit Rate	Ave. P Ave. L
Benchmarks									
Long Only	0.070	0.154	0.114	0.319	0.456	0.616	0.221	0.549	0.887
TSMOM	0.033	0.154	0.112	0.355	0.212	0.291	0.092	0.523	0.946
MACD	0.023	0.156	0.114	0.390	0.145	0.198	0.058	0.529	0.915
CSMOM	−0.107	0.154	0.115	0.838	−0.697	−0.929	−0.128	0.489	0.929
Reference									
DMN	0.055 (0.026)	0.162 (0.006)	0.115 (0.002)	0.391 (0.086)	0.340 (0.165)	0.478 (0.233)	0.156 (0.093)	0.525 (0.009)	0.966 (0.028)
STMOM									
SLP	0.333 (0.084)	0.161 (0.003)	0.104 (0.004)	0.244 (0.076)	2.066 (0.498)	3.228 (0.910)	1.619 (0.957)	0.574 (0.018)	1.121 (0.071)
MLP	0.288 (0.058)	0.162 (0.004)	0.107 (0.005)	0.238 (0.059)	1.776 (0.333)	2.699 (0.560)	1.302 (0.452)	0.559 (0.014)	1.123 (0.048)
CNN	0.030 (0.046)	0.164 (0.015)	0.112 (0.002)	0.443 (0.121)	0.174 (0.279)	0.270 (0.412)	0.089 (0.101)	0.518 (0.014)	0.963 (0.023)
LSTM	0.251 (0.085)	0.195 (0.115)	0.103 (0.004)	0.298 (0.063)	1.389 (0.384)	2.451 (0.892)	0.890 (0.415)	0.550 (0.016)	1.118 (0.087)

We evaluate the combination of time-series momentum (TSMOM and DMN) with CSMOM and compare it to a STMOM strategy. Focusing on the Sharpe ratio, we see from Exhibit 8 that the combinations of both TSMOM+CSMOM (0.177) and DMN+CSMOM (2.115) portfolios do not outperform a single SLP portfolio (2.609),

EXHIBIT 5

Cumulative Returns—Rescaled to Target Volatility (US equities)




demonstrating that the benefit of using a STMOM strategy outweighs a portfolio that combines both time-series and cross-sectional momentum, thus indicating that the model is learning meaningful interactions between the time-series and cross-sectional domain which are not captured by separate models. In addition, we analyze the combination of a time-series momentum (DMN) with a STMOM strategy (SLP). The combination of DMN (2.920) and SLP (2.609) yielded an overall higher Sharpe ratio (3.304), highlighting substantial scope for strategy diversification by incorporating an STMOM strategy.

Transaction Costs and Turnover Regularization

In this section, we analyze the impact of transaction costs on the performance of all strategies. Following Baltas and Kosowski (2020), we define turnover $TO_t^{(i)}$ as the absolute daily change in the trading signal of an asset:

$$TO_t^{(i)} = \sigma_{\text{tgt}} \left| \frac{X_t^{(i)}}{\sigma_t^{(i)}} - \frac{X_{t-1}^{(i)}}{\sigma_{t-1}^{(i)}} \right| \quad (12)$$

which encompasses the amount of rebalancing as determined by shifts in both volatility estimates and underlying positions. In Exhibit 10, we plot the distributions of turnover averaged across all assets for individual strategies. We observe that machine learning models trade considerably more than the benchmarks, with the

EXHIBIT 6**Cumulative Returns—Rescaled to Target Volatility (equity index futures)**

SLP having a lower turnover than the DMN. To study the impact of transactions costs on performance, we compute ex cost Sharpe ratios using captured returns net of transaction costs $\tilde{r}_{t,t+1}^{\text{TSMOM}}$:

$$\tilde{r}_{t,t+1}^{\text{TSMOM}} = \frac{1}{N_t} \sum_{i=1}^{N_t} \left(X_t^{(i)} \frac{\sigma_{\text{tgt}}}{\sigma_t^{(i)}} r_{t,t+1}^{(i)} - c \cdot \text{TO}_t^{(i)} \right) \quad (13)$$

where c is a constant approximating the specific transaction cost scenario. Based on the ex cost Sharpe ratios in Exhibit 9, the nonregularized DMN and SLP models are able to retain their performance over benchmarks for transactions costs of up to $c = 5$ to 10 basis points, which are considered to be realistic transaction costs for equity markets. Following the approach of Lim, Zohren, and Roberts (2019), we investigate the effect of incorporating turnover regularization during training in the form of optimizing for ex cost Sharpe ratios computed using $\tilde{r}_{t,t+1}^{\text{TSMOM}}$. For nonrecurrent models that do not generate sequential signals, such as the SLP, we introduce a localized minibatch turnover regularization in the form


$$\widetilde{\text{TO}}_t^{(i)} = \sigma_{\text{tgt}} \left| \frac{X_t^{(i)}}{\sigma_t^{(i)}} - \frac{X_{t^*}^{(i)}}{\sigma_{t^*}^{(i)}} \right| \quad (14)$$

where $t \neq t^*$ with t and t^* representing consecutive samples within a given training minibatch. Compared to its nonregularized counterpart, incorporating turnover

EXHIBIT 7

Correlations (US equities)

Panel A: Correlation of Strategies



Panel B: One-Year Rolling Correlations of SLP




EXHIBIT 8

Performance Metrics for Strategy Combinations—Rescaled to Target Volatility (US equities)

	E[Return]	Vol.	Downside Deviation	MDD	Sharpe	Sortino	Calmar	Hit Rate	Ave. P Ave. L
Combination									
TSMOM+CSMOM	0.028	0.156	0.113	0.577	0.177	0.245	0.048	0.521	0.946
DMN+CSMOM	0.340	0.161	0.106	0.307	2.115	3.222	1.132	0.572	1.099
	(0.025)	(0.001)	(0.001)	(0.045)	(0.153)	(0.261)	(0.189)	(0.005)	(0.017)
DMN+SLP	0.551	0.167	0.101	0.230	3.304	5.441	2.402	0.603	1.229
	(0.027)	(0.001)	(0.001)	(0.013)	(0.151)	(0.308)	(0.147)	(0.005)	(0.026)

regularization for the DMN resulted in a lower turnover, but better performance is only observed at a transaction cost of $c = 10$ basis points, while performance for all other transaction cost scenarios deteriorated. On the other hand, turnover regularization consistently improved the SLP's performance across all transaction cost scenarios. Referring to the distribution of average turnover for the regularized SLP, we observe a lower mean turnover but an increase in the spread between the minimum, maximum, and interquartile range. This is likely a direct consequence of performing stochastic gradient descent with the localized minibatch turnover regularization as per Equation 14, requiring batching over a shuffled training set.

Interpretability

Apart from demonstrating the best performance, the SLP is the simplest model out of all STMOM architectures. We perform an analysis on the interpretability of the trained SLP using Shapley additive explanations (SHAP) (Lundberg and Lee 2017), a


EXHIBIT 9

Impact of Transactions Costs on Sharpe Ratio—Rescaled to Target Volatility (US equities)

Transaction Cost (basis points)	0.0	0.5	1.0	2.0	3.0	4.0	5.0	10.0
Benchmarks								
Long Only	0.841	0.839	0.838	0.835	0.832	0.829	0.826	0.812
TSMOM	0.358	0.347	0.336	0.315	0.293	0.271	0.249	0.140
MACD	0.245	0.238	0.232	0.219	0.207	0.194	0.182	0.119
CSMOM	−0.655	−0.683	−0.710	−0.765	−0.820	−0.875	−0.930	−1.204
Reference								
DMN	2.920	2.844	2.768	2.615	2.462	2.308	2.153	1.375
DMN+Reg	2.073	2.044	2.015	1.957	1.899	1.840	1.782	1.486
STMOM								
SLP	2.609	2.518	2.427	2.243	2.060	1.876	1.691	0.762
SLP+Reg	2.672	2.603	2.534	2.395	2.256	2.116	1.976	1.271

EXHIBIT 10

Average Turnover of Strategies (US equities)




model-agnostic interpretation method for computing feature importance and effects for machine learning models. SHAP values associated to each individual feature measures the change in the expected value of the model's prediction when conditioning on that feature. We approximate SHAP values with Deep SHAP (Lundberg and Lee 2017), a high-speed approximation algorithm that builds on DeepLIFT (Shrikumar, Greenside, and Kundaje 2017).

Prediction for a single asset. We analyze the importance and effects of the spatio-temporal features on the predicted signal of a single asset—Bank of America Corp (BAC), by plotting SHAP values for the top 20 features ordered according to their importance in Exhibit 11. It is evident that the top features are dominated by volatility normalized MACD, indicating that the MACD features of assets are, on average, contributing more to the predicted signal of BAC as compared to other types of features like volatility normalized returns. This could mean that the cross over of exponentially weighted moving averages are considered significant momentum features for the model in generating its trading signals. Interestingly, the top features do not contain much of the BAC's own features, and instead other assets' features in the cross section are considered important in its own predicted signal.

EXHIBIT 11

Impact on Predicted Signal Output for a Single Asset—BAC (US equities)




Moving from low to high feature values and vice versa, there are clear patterns in the resultant impact of spatio-temporal features on the predicted signal output as seen from the gradual transitions in color. Momentum and mean-reversion effects of individual features are also relatively distinct and can be inferred from the SHAP plot. Taking the first two features as examples, the first feature, `AJG_t-4_MACD_32_96`, exhibits a mean-reversion effect because higher feature values lead to a lower predicted signal for BAC, whereas the second feature, `JEF_t-3_MACD_32_96`, shows a momentum effect as higher feature values lead to a higher predicted signal for BAC.

Prediction for all assets. To examine global feature importance, we compute the mean absolute SHAP values for each spatio-temporal feature on the predicted signal for all assets in US equities and plot their cumulative impact in Exhibit 12. Ordered according to importance, the top 20 features are again dominated by volatility normalized MACD features. We note that the cumulative absolute impact of a feature does not distinguish between its momentum or mean-reversion effects and serves only as a proxy of its feature importance over all assets.

EXHIBIT 12

Cumulative Mean Absolute Impact on Predicted Signal Output for All Assets (US equities)



CONCLUSIONS

We introduce spatio-temporal momentum (STMOM) a class of machine learning models that combine time-series and cross-sectional momentum strategies. This class of momentum strategies simultaneously generates trading signals for a portfolio of assets in a multitask setting, using both time-series and cross-sectional features from all assets in the portfolio. We demonstrate that the SLP, a neural network with a single hidden layer trained with a shrinkage penalty, is able to outperform more complex models in modelling STMOM, demonstrating its effectiveness over benchmarks on two different datasets.

We show that the performance of a STMOM strategy can be superior to a simple blend of time-series and cross-sectional momentum strategies. Exhibiting unstable correlation with other momentum strategies, the STMOM strategy when combined with a time-series momentum strategy like the DMN yields higher risk-adjusted returns, making STMOM a valuable strategy for diversification.

Given various cost scenarios, we examine the impact of transaction costs on the performance of all strategies and show that both the DMN and SLP were able to retain their performance over benchmarks for costs up to 5 to 10 basis points when tested on the US equities data. We incorporate turnover regularization for the DMN and a different localized minibatch turnover regularization for nonrecurrent models like the SLP, resulting in the SLP achieving better performance ratios across all transaction cost scenarios.

With the simplicity of the SLP model, it is possible to directly visualize and interpret its weights corresponding to each STMOM feature. We analyze the importance of individual STMOM features using the model-agnostic SHAP method, revealing

clear patterns in momentum and mean-reversion effects of individual features for predicting the trading signal of an individual asset, as well as showing global feature importance for all assets.

In future works, we would like to study the performance of STMOM using alternative feature representations. Another extension of this work includes adapting attention-based deep learning architectures (Bahdanau, Cho, and Bengio 2014; Vaswani et al. 2017; Wood, Roberts, and Zohren 2023) and investigating their effectiveness in modeling STMOM.

APPENDIX A

DEEP LEARNING ARCHITECTURES AND OPTIMIZATION

Deep Learning Architectures

For the CNN and LSTM models, we consider the spatio-temporal tensor $\mathbf{u}_t \in \mathbb{R}^{\tau \times m'}$ with $m' = N^t \cdot d$, representing an input sequence of momentum features of all assets over a temporal history τ .

CNN. We consider a one-dimensional autoregressive CNN of the following:

$$\mathbf{h}_t = \mathcal{P}\sigma[\mathbf{W}_c^{[2]} * \sigma(\mathbf{W}_c^{[1]} * \mathbf{u}_t + \mathbf{b}_c^{[1]}) + \mathbf{b}_c^{[2]}] \quad (\text{A1})$$

$$\mathbf{X}_t = f(\mathbf{u}_t; \theta) = g[\mathbf{W}^{[2]\top} \sigma(\mathbf{W}^{[1]\top} \mathbf{h}_t + \mathbf{b}^{[1]}) + \mathbf{b}^{[2]}] \quad (\text{A2})$$

where $\mathbf{W}_c^{[j]}$ represent convolutional kernels with associated bias terms $\mathbf{b}_c^{[j]}$ and activations $\sigma = \tanh$, and $\mathbf{w} * \mathbf{u}$ representing the causal convolution operation between the input sequence \mathbf{u} and kernel \mathbf{w} . Additionally, we interface an average pooling layer \mathcal{P} that performs a downsampling step prior to propagating the activations \mathbf{h}_t into an MLP as per Equation 7.

LSTM. We consider a single-layer LSTM model:

$$\Gamma_t^i = \sigma_s(\mathbf{W}_i \mathbf{u}_t + \mathbf{V}_i \mathbf{h}_{t-1} + \mathbf{b}_i) \quad (\text{A3})$$

$$\Gamma_t^f = \sigma_s(\mathbf{W}_f \mathbf{u}_t + \mathbf{V}_f \mathbf{h}_{t-1} + \mathbf{b}_f) \quad (\text{A4})$$

$$\Gamma_t^o = \sigma_s(\mathbf{W}_o \mathbf{u}_t + \mathbf{V}_o \mathbf{h}_{t-1} + \mathbf{b}_o) \quad (\text{A5})$$

$$\tilde{\mathbf{c}}_t = g(\mathbf{W}_c \mathbf{u}_t + \mathbf{V}_c \mathbf{h}_{t-1} + \mathbf{b}_c) \quad (\text{A6})$$

$$\mathbf{c}_t = \Gamma_t^i \odot \tilde{\mathbf{c}}_t + \Gamma_t^f \odot \mathbf{c}_{t-1} \quad (\text{A7})$$

$$\mathbf{h}_t = \Gamma_t^o \odot g(\mathbf{c}_t) \quad (\text{A8})$$

$$\mathbf{X}_t = g(\mathbf{W}_{\text{dist}} \mathbf{h}_t + \mathbf{b}) \quad (\text{A9})$$

where \mathbf{W}_* , \mathbf{V}_* represent the weights and \mathbf{b}_* the biases of the input, forget, and output gates Γ_t^i , Γ_t^f , Γ_t^o respectively, with activation functions $\sigma_s = \text{sigmoid}$ and $g = \tanh$. Subsequently, the LSTM computes the memory cell state \mathbf{c}_t , hidden state activation \mathbf{h}_t with \odot representing the Hadamard product. Given the spatio-temporal input of the form $\mathbf{u}_t \in \mathbb{R}^{\tau \times m'}$, the network then maps the hidden state activation \mathbf{h}_t to a sequence of trading signals $\mathbf{X}_t \in [-1, 1]^{\tau \times N^t}$ by a fully connected layer with time-distributed weights \mathbf{W}_{dist} .

EXHIBIT A1

Fixed Parameters for Machine Learning Models

Parameters	DMN	SLP	MLP	CNN	LSTM
Epochs	100	500	500	500	500
Patience	25	25	25	25	25
Random Search Iterations	100	100	100	100	100
Temporal History	63	5	5	63	63

EXHIBIT A2

Hyperparameter Search Range for Machine Learning Models

Hyperparameters	Random Search Grid
Minibatch Size	32, 64, 128, 256
Dropout Rate	0.1, 0.2, 0.3, 0.4, 0.5
Hidden Layer Size	5, 10, 20, 40, 80, 160
Learning Rate	10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , 10^0
Max Gradient Norm	10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , 10^0 , 10^1
L1 Regularization Weight (α)	10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , 10^0

Fixed Parameters and Hyperparameter Optimization

We perform hyperparameter optimization with 100 iterations of random search for all machine learning models, using the search range as shown in Exhibit A2. We initiate early stopping with a patience of 25 epochs for the validation loss. We include regularization via dropout (Srivastava et al. 2014) for the reference DMN, as well as the MLP, CNN, and LSTM STMOM models.

APPENDIX B

DATASET DETAILS

All individual instruments in our datasets have less than 10% missing data. To reduce the effect of extreme outliers, we winsorize all data to limit values to be within five times its 252-day exponentially weighted moving standard deviation from its exponentially weighted moving average.

Equity index futures (Exhibit B1) consists of 12 ratio-adjusted continuous equity index futures contracts obtained from the Pinnacle Data Corp CLC Database. We perform backtesting from 2003 to 2020.

US equities (Exhibit B2) consists of actively traded US equities with data obtained from the CRSP. We screen for common stocks of domestic US companies listed on the NYSE, AMEX, and NASDAQ with market capitalization from small (300 million to 2 billion) to mega (>200 billion) using the Stock Screener provided by NASDAQ. We perform a random sample of 46 stocks from the financials sector. We perform backtesting from 1990 to 2022 (Exhibits B1 and B2).

EXHIBIT B1**Equity Index Futures Dataset**

Ticker Symbol	Ticker Description
SP	S&P 500,day session
YM	Mini Dow Jones (\$5.00)
EN	NASDAQ, MINI
ER	RUSSELL 2000, MINI
MD	S&P 400 (Mini electronic)
XU	DOW JONES EUROSTOXX50
XX	DOW JONES STOXX 50
CA	CAC40 INDEX
LX	FTSE 100 INDEX
AX	GERMAN DAX INDEX
HS	HANG SENG
NK	NIKKEI INDEX

EXHIBIT B2**US Equities Dataset**

Ticker Symbol	Ticker Description
AFG	AMERICAN FINANCIAL GROUP INC NEW
AFL	AFLAC INC
AJG	GALLAGHER ARTHUR J & CO
BAC	BANK OF AMERICA CORP
C	CITIGROUP INC
CADE	CADENCE BANK
CBSH	COMMERCE BANCSHARES INC
CFR	CULLEN FROST BANKERS INC
CHCO	CITY HOLDING CO
CPF	CENTRAL PACIFIC FINANCIAL CORP
CVBF	C V B FINANCIAL CORP
FITB	FIFTH THIRD BANCORP
GL	GLOBE LIFE INC
IBCP	INDEPENDENT BANK CORP MICH
INDB	INDEPENDENT BANK CORP MA
JEF	JEFFERIES FINANCIAL GROUP INC
JPM	JPMORGAN CHASE & CO
KEY	KEYCORP NEW
MCO	MOODYS CORP
MCY	MERCURY GENERAL CORP NEW
MKL	MARKEL CORP
NTRS	NORTHERN TRUST CORP
NWLI	NATIONAL WESTERN LIFE GROUP INC
PNC	P N C FINANCIAL SERVICES GRP INC
SBCF	SEACOAST BANKING CORP FLA
SCHW	SCHWAB CHARLES CORP NEW
SEIC	S E I INVESTMENTS COMPANY
SIGI	SELECTIVE INSURANCE GROUP INC
SIVB	S V B FINANCIAL GROUP
SNV	SYNOVUS FINANCIAL CORP

(continued)

EXHIBIT B2 *(continued)***US Equities Dataset**

Ticker Symbol	Ticker Description
SRCE	1ST SOURCE CORP
STT	STATE STREET CORP
TFC	TRUIST FINANCIAL CORP
TMP	TOMPKINS FINANCIAL CORP
TRC	TEJON RANCH CO
TROW	T ROWE PRICE GROUP INC
TRST	TRUSTCO BANK CORP NY
UBSI	UNITED BANKSHARES INC
UFCS	UNITED FIRE GROUP INC
USB	U S BANCORP DEL
VALU	VALUE LINE INC
VLY	VALLEY NATIONAL BANCORP
WABC	WESTAMERICA BANCORPORATION
WFC	WELLS FARGO & CO NEW
WRB	BERKLEY W R CORP
ZION	ZIONS BANCORPORATION N A

ACKNOWLEDGMENTS

We thank Daniel Poh for his valuable comments as well as the Oxford-Man Institute of Quantitative Finance for data access and compute resources.

REFERENCES

- Abadi, M., P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. “TensorFlow: A System for Large-Scale Machine Learning.” In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pp. 265–283. 2016. Savannah, GA, USA: USENIX Association.
- Bahdanau, D., K. Cho, and Y. Bengio. 2014. “Neural Machine Translation by Jointly Learning to Align and Translate.” *arXiv* 1409.0473.
- Baltas, N., and R. Kosowski. 2020. “Demystifying Time-Series Momentum Strategies: Volatility Estimators, Trading Rules and Pairwise Correlations.” In *Market Momentum: Theory and Practice*. John Wiley & Sons, Ltd, 3: 30–67.
- Baz, J., N. Granger, C. R. Harvey, N. Le Roux, and S. Rattray. 2015. “Dissecting Investment Strategies in the Cross Section and Time Series.” *SSRN* 2695101.
- Bengio, Y., P. Simard, and P. Frasconi. 1994. “Learning Long-Term Dependencies with Gradient Descent Is Difficult.” *IEEE Transactions on Neural Networks* 5 (2): 157–166.
- Binkowski, M., G. Marti, and P. Donnat. 2018. “Autoregressive Convolutional Neural Networks for Asynchronous Time Series.” In *International Conference on Machine Learning*, pp. 580–589. PMLR, Stockholm, Sweden.
- Caruana, R. 1997. “Multitask Learning.” *Machine Learning* 28 (1): 41–75.
- Center for Research in Security Prices (CRSP). 2022. *Calculated (or Derived) Based on Data from CRSP Daily Stock 2022*.
- Crawshaw, M. 2020. “Multi-Task Learning with Deep Neural Networks: A Survey.” *arXiv* 2009.09796.

- Dong, D., H. Wu, W. He, D. Yu, and H. Wang. 2015. "Multi-Task Learning for Multiple Language Translation." In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1723–1732. Beijing, China: Association for Computational Linguistics.
- Garg, A., C. L. Goulding, C. R. Harvey, and M. Mazzoleni. 2021. "Momentum Turning Points." SSRN 3489539.
- Harvey, C. R., E. Hoyle, R. Korgaonkar, S. Rattray, M. Sargaison, and O. Van Hemert. 2018. "The Impact of Volatility Targeting." *The Journal of Portfolio Management* 45 (1): 14–33.
- He, K., X. Zhang, S. Ren, and J. Sun. 2016. "Deep Residual Learning for Image Recognition." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778. Los Alamitos, CA, USA: IEEE Computer Society.
- Hochreiter, S., and J. Schmidhuber. 1997. "Long Short-Term Memory." *Neural Computation* 9 (8): 1735–1780.
- Jegadeesh, N., and S. Titman. 1993. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *The Journal of Finance* 48 (1): 65–91.
- Kim, A. Y., Y. Tse, and J. K. Wald. 2016. "Time Series Momentum and Volatility Scaling." *Journal of Financial Markets* 30: 103–124.
- Kingma, D. P., and J. Ba. 2014. "Adam: A Method for Stochastic Optimization." *arXiv* 1412.6980.
- LeCun, Y., Y. Bengio, and G. Hinton. 2015. "Deep Learning." *Nature* 521 (7553): 436–444.
- Lim, B., and S. Zohren. 2021. "Time Series Forecasting with Deep Learning: A Survey." *Philosophical Transactions of the Royal Society A* 379 (2194): 20200209.
- Lim, B., S. Zohren, and S. Roberts. 2019. "Enhancing Time-Series Momentum Strategies Using Deep Neural Networks." *The Journal of Financial Data Science* 1 (4): 19–38.
- Long, M., Z. Cao, J. Wang, and P. S. Yu. 2017. "Learning Multiple Tasks with Multilinear Relationship Networks." *Advances in Neural Information Processing Systems* 30: 1593–1602.
- Lundberg, S. M., and S.-I. Lee. 2017. "A Unified Approach to Interpreting Model Predictions." *Advances in Neural Information Processing Systems* 30: 4768–4777.
- Luong, M.-T., Q. V. Le, I. Sutskever, O. Vinyals, and L. Kaiser. 2015. "Multi-Task Sequence to Sequence Learning." *arXiv* 1511.06114.
- Misra, I., A. Shrivastava, A. Gupta, and M. Hebert. 2016. "Cross-Stitch Networks for Multi-Task Learning." In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3994–4003. Los Alamitos, CA, USA: IEEE Computer Society.
- Moskowitz, T. J., Y. H. Ooi, and L. H. Pedersen. 2012. "Time Series Momentum." *Journal of Financial Economics* 104 (2): 228–250.
- Oord, A. v. d., S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu. 2016. "Wavenet: A Generative Model for Raw Audio." *arXiv* 1609.03499.
- Paszke, A., S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. 2019. "PyTorch: An Imperative Style, High-Performance Deep Learning Library." *Advances in Neural Information Processing Systems* 32.
- Pinnacle Data Corp. 2021. "Backtest Futures With Our CLC Database." <https://pinnacledata2.com/clc.html>.
- Pirrong, C. 2005. "Momentum in Futures Markets." SSRN 671841.

- Poh, D., B. Lim, S. Zohren, and S. Roberts. 2021. "Building Cross-Sectional Systematic Strategies By Learning to Rank." *The Journal of Financial Data Science* 3 (2): 70–86.
- . 2022. "Enhancing Cross-Sectional Currency Strategies by Context-Aware Learning to Rank with Self-Attention." *The Journal of Financial Data Science* 4 (3): 89–107.
- Rei, M. 2017. "Semi-Supervised Multitask Learning for Sequence Labeling." *arXiv* 1704.07156.
- Rouwenhorst, K. G. 1998. "International Momentum Strategies." *The Journal of Finance* 53 (1): 267–284.
- Ruder, S. 2017. "An Overview of Multi-Task Learning in Deep Neural Networks." *arXiv* 1706.05098.
- Sharpe, W. F. 1994. "The Sharpe Ratio." *The Journal of Portfolio Management* 21 (1): 49–58.
- Shrikumar, A., P. Greenside, and A. Kundaje. 2017. "Learning Important Features through Propagating Activation Differences." In *International Conference on Machine Learning*, pp. 3145–3153. Sydney, NSW, Australia: Shrikumar.
- Simonyan, K., and A. Zisserman. 2014. "Very Deep Convolutional Networks for Large-Scale Image Recognition." *arXiv* 1409.1556.
- Srivastava, N., G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. 2014. "Dropout: A Simple Way to Prevent Neural Networks from Overfitting." *The Journal of Machine Learning Research* 15 (1): 1929–1958.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. 2017. "Attention Is All You Need." *Advances in Neural Information Processing Systems* 30: 6000–6010.
- Wood, K., S. Giegerich, S. Roberts, and S. Zohren. 2023. "Trading with the Momentum Transformer: An Intelligent and Interpretable Architecture." *arXiv* 2112.08534.
- Wood, K., S. Roberts, and S. Zohren. 2022. "Slow Momentum with Fast Reversion: A Trading Strategy Using Deep Learning and Change-point Detection." *The Journal of Financial Data Science* 4 (1): 111–129.
- Zhang, Z., P. Luo, C. C. Loy, and X. Tang. 2014. "Facial Landmark Detection by Deep Multi-Task Learning." In *European Conference on Computer Vision*, pp. 94–108. Cham: Springer.
- Zhang, Z., S. Zohren, and S. Roberts. 2019. "DeepLOB: Deep Convolutional Neural Networks for Limit Order Books." *IEEE Transactions on Signal Processing* 67 (11): 3001–3012.
- . 2020. "Deep Learning for Portfolio Optimization." *The Journal of Financial Data Science* 2 (4): 8–20.