

How does Convolutional Neural Network (CNN) predict stock returns?

Stephen Dixon* and Qi Zeng[†]

This version: July 2024

Abstract

We interpret the image-based asset pricing factors produced with convolutional neural network (CNN) by proposing a methodology through which the aggregate pattern learned by vision models can be visualized in a heat map. We show and quantify the CNN captured price and volume paths, including the level, the slope, the curvature, and, most importantly, the last day change in the difference between close price and high-low average. Four simple geometric measures can explain 31% (24%) of variation in predicted probability of positive weekly (monthly) returns, and 73% (48%) of variation in weekly (monthly) CNN returns. We further hypothesize and confirm that the machine can identify time series persistence and reversal patterns. Separating them enables the formation of long-short portfolios with better performances than does the existing CNN factor.

Keywords: machine learning, convolutional neural network (CNN), technical analysis, stock returns predictability

*University of Melbourne. Email address: dixosd@student.unimelb.edu.au.

[†]University of Melbourne. Email address: qzeng@unimelb.edu.au.

Recent advancements in Machine Learning (ML) have revived research interest in technical analysis, and suggest that charting could predict future returns (Murray et al., 2024). A prominent example is Jiang et al. (2023, JKX therein), who uses convolutional neural networks (CNN) to construct an imaged-based factor that is distinct from known ones. Although they show some economic linkage to the model’s predictive abilities, the actual patterns learnt by CNN remain elusive. As such, our paper addresses one central question: What does CNN *see* in stock charts when making predictions?

Answering the above question is no trivial task; these networks not only intake high dimensional inputs, but also fit large number of parameters in a recursive manner. JKX attempt to address this issue by regressing the predicted probability of a stock realising positive returns against stock characteristics, and the data points used to construct the image. They show the exceptional predictive abilities of CNN stems from scaling the price variables, but they acknowledge interpretation is still lacking given the black box nature of ML. In this paper, We provide a new methodology to visualise CNN captured commonalities in open, high, low, close, and volume (OHLCV) charts. In comparison to the gradient-weighted class activation mapping (grad-CAM) method provided by Selvaraju et al. (2019), our methodology shows the ‘aggregate’ patterns captured by the portfolios resulting from predictions from CNN instead of the patterns learned by CNN itself.

We begin by visualizing the top (long) and bottom (short) decile CNN portfolios, which seem to capture a non-traditional reversal characterized by declines (rises) in price components and, strikingly, an extremely low (high) closing price on the last day of the formation period. While there are distinct trends for volume and high-low spread, the last day closing price is both optically and numerically more important. Building on this, we then construct a set of quantitative measures for which the CNN predictions are regressed against. These measures capture the level, slope and the curvature of the trends. The most significant pattern, however, is the last day change. For the 27 Million out-of-sample weekly and monthly prediction, 15% of variation is explained by a measure of the extreme change in last day

closing price relative to previous days. 30% (24%) of variation in weekly (monthly) prediction is explained by further including last day volume and close price relative to that of previous days, and formation returns. We then directly tie these trends to CNN’s portfolio performance; regressing CNN weekly (monthly) returns against returns of the four individually sorted long-short portfolios results in 73% (48%) r-squared. Thus, we provide both graphical and quantitative evidences on features driving the machine’s ability to forecast returns.

So is there any economic reasoning behind these machine learnt trends? The overall patterns are consistent with Weekly Short Term Reversal (WSTR), which was first documented by Lehmann (1990), and more recently scrutinized by Cooper (2015) to proxy for investor overreaction. CNN predictions are similar to that of Cooper (2015), as both predicts next week or month’s return is more likely to be positive when volume increase or price decreases. We further document that CNN outperforms WSTR by focusing on a steeper price and volume growth rate towards the end of formation window. Broadly speaking, the price and volume path through the formation window matters, in addition to some simple average or slope of stock characteristics. Since overreactions are quickly corrected to some extent (i.e., short sale constraints), it is expected that the aforementioned trend will become less important for longer holding periods. Indeed, as the prediction horizon increases, the importance placed on the last day closing price and volume is reduced, and the portfolio performance also plummets. More support is provided by recent work of Chu (2024), who illustrate that CNN is profitable only when it disagrees with the herding of technical retail traders, but not when it agrees with such investors. It is therefore completely possible that CNN learnt, on a large part, patterns that represent market overreaction.

The discussion so far have centered around CNN’s ability to capture a novel form of short term reversal, which can explain a large proportion of, but not all, variation in CNN predictions and the subsequent portfolio returns. While WSTR’s long (short) end are, by construct, all stocks with negative (positive) returns, CNN’s long (short) end has 24% of

stocks with positive (negative) returns, on average. The previous two facts combined calls for further investigation into vision model’s predictive abilities. Specifically, we conjecture that trend reversal does not comprehensively describe the CNN’s learning abilities, as the model could potentially learn about persistency in trends (i.e., positive (negative) future returns, given positive (negative) formation returns). From an asset pricing perspective, persistent trend is similar to time series momentum as described by Murray et al. (2024). To test this hypothesis, we decompose the JKK factor by applying the model in sub-sample of equities with positive return over the formation period, and the sub-sample with negative returns over the same period.

Consistent with our hypothesis, in the sample of stocks with positive (negative) returns, the top (bottom) decile portfolio display an obvious upward (downward) trend with a slow-down/small reversal in the last day, while realising annualized 23% (-19%) returns¹ over the next 5 days ². In addition, the reversal signals are much cleaner, compared to the noise in JKK patterns. We observe in the negative (positive) sub-sample, the high (low) decile portfolio shows a clear decreasing (increasing) trend in the formation with the large drop (rise) in closing price, while realising annualized 62% (-43%) returns.

It is important to consider the implication of the previous sub-sample decomposition on portfolio formation. We can, of course, long stocks with high probability of realising positive return, and short those with low probability in negative or positive sub-samples. However, the economic intuition is unclear: the portfolio shorts trend persistency (reversal) and longs trend reversal (persistency) in negative (positive) sample. Each of the two high-low portfolios incorporates two distinctly different signals. We propose forming two positions, each of which earns a spread on the same signal; Long (Short) trend persistency in the positive (negative) sample and Long (Short) trend reversal in the negative (positive) sample.

The resulting reversal portfolio generates 105%, 27% and 10% annualized return over

¹As in JKK, returns of 5-day portfolios may not be practically achievable, but it is still a good measure of performance.

²We also conduct sub-sample training and application of the same model, which provides visually and numerically the same findings.

weekly, monthly, and quarterly holding period, respectively. On the other hand, the persistency portfolio realizes 42%, 5%, and (insignificant) 3% returns over the three holding periods. One reason for the poor performance of persistency in monthly and quarterly holding period is that while the long end realises a significant 9% and 11%, respectively, the short end realises volatile and insignificant positive returns.

Our results indicate that the image-based factors have some unique features that are not captured by the simple summary characteristics over the price history. It is the evolution of that history (the image) that matters most. JKX observes the same point when they shows that normalization of the time series data affects the performance the most. In other words, it is the “shape” that matters. Compared to JKX, our approach shows the ”shape”(s) that drive pattern portfolios, provides basic intuition for the shapes, and establishes a link to literature on traditional asset pricing and technical analysis. Further, we show that such intuition derived from opening the black box of machine learning could also have an affect on the final performance, which, in our case, is portfolio formation.

Our work contributes to two bodies of literature, the three-decade theoretical and empirical justification of why technical analyses does or does not work, and the recent influx of research on the application of machine learning in asset pricing. Lo et al. (2000) define and test ten commonly used patterns by chartists, and find some support for the informativeness of charting in NASDAQ. In a follow up discussion, Jegadeesh (2000) expresses concern for the lack of theoretical foundation, subjective choices in the empirical design, and the lack of significance in NYSE/AMEX. While papers as early as Brown and Jennings (1989) and Barberis et al. (1998) provide theoretical motivation for technical predictability, such recent papers as Zhu and Zhou (2009) and Han et al. (2016) provide both empirical and theoretical support.

Bajgrowicz and Scaillet (2012) and Ready (2002) show that even in-sample returns to technical analysis is offset by transaction costs, and the former underlines and control for data snooping. Indeed, as long as the research pre-define enough patterns, surely one will be

significant. The recent advances in ML addresses such problem; rather than testing patterns specified by the experimenter, cleverly designed algorithms can learn patterns in historical information and make some 'educated' prediction. Jiang et al. (2023) introduces vision models to predict if an equity will be positive in the next period using OHLCV images, and Murray et al. (2024) combines four different neural network models including CNN and CNN-based model (CNN-LTSM) and finds a different image-based predictors. Although most of the work pay special attention to interpreting the model, it is notoriously hard to observe the inner workings of these algorithms due to the large number of parameters. Our work tackles this issue and show the specific stock chart features learnt by convolutional neural networks. We further reinforce our findings by quantifying these patterns which explain a considerable variation a non-linear output in standard regression settings. To the best of our knowledge, we are the first to show, on an aggregate level, the patterns learned by a vision model.

The remainder of the paper is organised as follows. In the next section, we will briefly review the CNN models including the image formation, and refer any curious reader to JKK for more in depth detail and justification. In Section 3 we present the data, our heat-map construction and the main results. We provide sub-sample analysis in Section 4, and conclude in Section 5.

1 Stock Charts and CNN

This section provides a brief overview of representing market data with images, and the CNN model(s) that uses such inputs. We provide some intuition and justification for the use of CNN, and refer to the original discussion in JKK for most technical details.

Classic technical analysis uses charts, or more precisely, the movement and trend in prices, as inputs to predict stock returns. JKK uses the so-called OHLC bar, which represents the daily open price, high-low bar, and close price of a stock. An example of 20-day OHLC

chart is shown in Figure 1. Specifically, each day of historical data is represented with three columns of pixels, open (one pixel), high-low bar, and close (one pixel). In addition, each image also includes the trading volume (mid-column bottom panel) and the moving average with the period the same as the image window. The open, high-low bar, close, volume, and moving average is referred to as graph components here forth. In the image, each component is represented by a white pixel, with black background. JKKX considers only three windows, 5 days, 20 days and 60 days, which correspond to weekly, monthly and quarterly horizon. Within each n -day window, price components minus the minimum value are scaled³ with the maximal value of that window to ensure the n -day historical data fits in the fixed size image.

The most popular method of learning image data is through CNN or its variations. In the usual feed-forward neural network, the nodes are fully connected between layers, and between inputs and layers, This requires the fitting of a huge amount of parameters, which is often not sustainable with the available data. CNN instead uses convolution layers which include a number of small-sized “blocks”, or filters. Each filter moves through the input image and produces a smaller-than-input-size output through a convolution operation. One can regard the convolution operation as a variation of kernel smoothing. Each filter essentially captures some unique features in the original image that is independent of the location and orientation of the feature. Figure 2 demonstrates the key steps (model structure) that transform a 5-day image input into predicted probabilities of being in some defined group. The key innovation is that features are created from row and colum-wise convolution of kernels, making it possible to capture nonlinear spatial associations among various price curves (Jiang et al., 2023, p.3202). In the context of technical analysis, one, if not the most, important reason to use CNN is that the supposed return predicting pattern(s) are no longer specified by the practitioner or researcher, but instead learned by the machine. This mitigates, to some extent, data snooping and subjective choices made, in exchange for

³We still adjust price variables by the CRSP adjust returns to reduce some noise.

interpretability.

We focus on understanding the features learned by CNN, and extend the analysis to different sub-sample applications. The model is trained to predict whether the next 5, 20, 60 day returns will be positive by minimizing the standard cross-entropy loss function

$$L(y, \hat{y}) = y \cdot \log(\hat{y}) + (1 - y) \cdot \log(1 - \hat{y}) \quad (1)$$

where \hat{y} is the softmax output (prediction), and y is the actual realization. We do explore k -class predictions by changing the final fully connected layer to output k features, with and without a weighted cross entropy loss function⁴. However, such changes were implemented either as a robustness check or as a test unrelated to our main empirical analyses, so we report the findings in the appendix.

2 Data and Methodology

2.1 Data and Simple Returns Reversal

Our data construction closely follows JKX, which mostly follows Gu et al. (2020). We use monthly total individual equity returns from CRSP for all firms listed in the NYSE, AMEX, and NASDAQ. The sample starts on January 8th, 1993, and ends on December 31, 2019. This is because the daily open, high, low and close only became available from Jan 8th, 1993 on. Table 1 shows the summary statistics of the sample. Like in JKX, we do not winsorize or remove data, and use the full sample instead. In the table, we report the statistics for daily prices, open, high, low, and close. We use ex-dividend returns as in JKX and the dollar

⁴Consider a binary prediction where class one is 99% of the sample. Assume perfect prediction is not possible, and improving the prediction accuracy of one class reduces that of the other equally. Then accuracy can simply be maximized by making the prediction of class one for any input. The weighted cross entropy loss function

$$L(y, \hat{y}) = -w_n[y \cdot \log(\hat{y}) + (1 - y) \cdot \log(1 - \hat{y})]$$

address such issue by assigning some weight w_n to the n^{th} class's loss. If the weight is equal across classes, then we return to the standard form. Our applications in the appendix use the inverse of class frequencies so that each class is equally important.

volume. These form the base for the images. In addition to these variables, we also report the statistics for the shares outstanding and the total market value. One can see that the samples include some extreme values in almost every variable used. It will be problematic to use the raw values in the standard econometric analysis, but for image-based analysis, this is not a severe issue. Firstly, by scaling the price variable, extreme values are still contained in an image of fixed size⁵. Secondly, CNN objectively looks for distinct patterns in the prices and volumes, and, therefore, reduces the effects from noises and enhances the stability of the parameters with respect to the sample.

Using the full universe of CRSP equities, we generate around 53 million images for 5-, 20-, and 60-day horizons from 1993 to 2019. The images serve as input for training and making predictions with a CNN model. Models are trained with images from 1993 to 2000 to predict if an observation is in a predefined class. Specifically, as in JKX, we predict whether the next 5-, 20-, and 60-day returns are positive or not. The final models are applied to images from 2001 to 2019 to obtain the probabilities of being in each class. We will discuss the 5-day window in the main analysis, and leave the results for 20- and 60-days later.

Finally, We form decile portfolios using the out-of-sample probabilities and report the weekly performance and average characteristics in the top panel of Table 2. Here we sort stocks according to the CNN predicted probability of realizing positive 5-day returns (R_5), which are shown in the column denoted as $P(R_5 > 0)$. The predicted probability increases from 0.29 to 0.64 across the deciles. Look across the equal-weighted portfolio, the returns of the portfolios monotonically increase from negative 32 percent to positive 48 percent per annum. The long-short portfolio generates a stunning 80% annualized return and a Sharpe ratio of 7.18 for an equal-weighted portfolio. Even for a value-weighted portfolio, it can generate an annualized return of 23 % and a Sharpe Ratio of 1.57. These are consistent with those reported in JKX, validating our algorithm. JKX admits that this result might not be achievable in practice. Our focus here is not so much on the actual performances

⁵The image size is dependent on the days of historical information included. For the 5-day case, the height and width of all images is 32 and 15 pixels, respectively.

achieved. Instead, we want to understand what CNN learns from the images to generate this performance, even if it is only theoretically achievable. The monthly and quarterly performance is reported in 8 for comparison.

For this purpose, we also show some characteristics across the portfolios. How exactly does CNN determine the assigned probabilities and how do they relate to the portfolio characteristics? The first thought is the return relation between the past five days (denoted as $I5$) and the next five days. The last column shows this result. One can see a monotonic decrease in $I5$, from 2% in decile 1 to negative 4% in decile 10. This indicates some reversal in returns, but the slope of $I5$ is much smaller than that of $R5$, indicating that there are additional patterns in the image than just past returns. To further reinforce this observation and show it is not driven by outliers, we show in columns 5 and 6, the proportion of holdings with positive $I5$ and $R5$ returns. Indeed, in the high decile that generates 47% annualized returns, only 24% of holdings had positive formation period returns. This proportion then monotonically increases to 66% in the low decile, which realizes annualized returns of -32%. This says that the monotonic patterns in $I5$ and $R5$ are not just from some extreme values, but also the number of the stocks. Furthermore, the performances do not seem related to the size effect. Looking at the column on the market cap (MC), there is no monotone relation between market capitalization and probability. If anything, the small cap stocks are concentrated in decile 1 and decile 10, and are much smaller than those in the middle deciles, indicating the extreme returns are at least partially driven by smaller stocks.

Is this purely a mean-reverting effect? In the bottom panel, we show the results from portfolios sorted by Weekly Short-Term Reversal (WSTR), defined as the past weekly returns ($I5$). As expected, the change in both the magnitude of formation returns, and the proportion of stocks that positive formation returns across the decile is much larger than observed in CNN sorted portfolios. However, the long and short leg of CNN portfolios not only generates wider spreads than WSTR, but also achieves a 4% (2%) higher (lower) proportion of holdings with positive realization in the long (short) end. WSTR only achieves 59%

annualized return and 3.15 Sharpe ratio for the equal-weighted portfolio, and 25% and 1.10 for value-weighted portfolios. In summary, there are some trend-reversal, but not in the traditional asset pricing sense. We next move beyond the simple past returns and focus on characterizing the performance driving *return and volume path patterns*.

2.2 Visualising features Across the Portfolios formed from CNN

It is notoriously difficult to understand what ML models learn, and, in the case of CNN, to identify the patterns useful for prediction. We provide a simple methodology to visualize the captured commonalities, under some restrictions on the image. Consider a set of images associated with positive or negative holding period returns $I = \{A^{(1)}, A^{(2)}, \dots, A^{(k)}\}$ each with height h and width w ,

$$A^{(k)} = \begin{bmatrix} a_{11}^{(k)} & a_{12}^{(k)} & \cdots & a_{1w}^{(k)} \\ a_{21}^{(k)} & a_{22}^{(k)} & \cdots & a_{2w}^{(k)} \\ \vdots & \vdots & \ddots & \vdots \\ a_{h1}^{(k)} & a_{h2}^{(k)} & \cdots & a_{hw}^{(k)} \end{bmatrix} \quad (2)$$

where $a_{ij}^{(k)}$ is one if the pixel is white, and zero otherwise⁶. Assume all images have a common pattern, characterized by the same pixel placements with some unknown noise distribution, with the rest of the pixels following a random uniform distribution (or a distribution distinct from the noise process). Then we can observe the pattern through a 'heat map' calculated

⁶The binary representation here is also important because mixing colored pixels likely interferes with interpreting the output of our methodology.

as the weighted ($w^{(k)}$) average of all images

$$H = \begin{bmatrix} \sum_{k=1}^n w^{(k)} a_{11}^{(k)} & \sum_{k=1}^n w^{(k)} a_{12}^{(k)} & \cdots & \sum_{k=1}^n w^{(k)} a_{1w}^{(k)} \\ \sum_{k=1}^n w^{(k)} a_{21}^{(k)} & \sum_{k=1}^n w^{(k)} a_{22}^{(k)} & \cdots & \sum_{k=1}^n w^{(k)} a_{2w}^{(k)} \\ \vdots & \vdots & \ddots & \vdots \\ \sum_{k=1}^n w^{(k)} a_{h1}^{(k)} & \sum_{k=1}^n w^{(k)} a_{h2}^{(k)} & \cdots & \sum_{k=1}^n w^{(k)} a_{hw}^{(k)} \end{bmatrix} \quad (3)$$

then normalizing each element by either the maximum value of their respective graph component in the same column, to observe the within-column distribution, or in the whole image, to observe the dispersion across all days. We plot the heat-map across the three types of components in an image: OHLC prices, moving average, and the volume.

Similar to our work in Section 2.1, research on the cross-section of stock returns generally forms some portfolio, then checks whether stock characteristics change with portfolio bins. Such an approach ignores the within-formation period movements of characteristics for a good reason: there is no guarantee a measure accounting for such changes is not a result of data mining, and there is no economic intuition for why a particular trend can persistently predict equity returns. Machine learning addresses the former, as CNN objectively learns some patterns, and provides 20 years of robust out-of-sample portfolio performance with no retraining. To address the latter concern, we study the geometric pattern of returns and volumes picked up by CNN models. We plot the heat-maps of images *across the portfolios*, and show the long-short strategy’s underlying pattern(s), if any. In particular, we also quantify any predominant trends captured by CNN, and estimate its effect on the predicted probability of realising a positive return.

Before we continue with the application, we wish to compare our method with the popular Gradient-Weighted Class Activation Mapping (Grad-CAM) approach proposed by Selvaraju et al. (2019). Grad-CAM essentially identifies the regions important for final prediction and it provides an intuitive way to understand the decision-making process of CNNs by visualizing the importance or relevance of these regions. Our method is different. We want

to search for the salient patterns exhibited in the resulting portfolios formed from the CNN predictions. These are the patterns that an investors can see, while Grad-CAM is the patterns that are important for the CNN itself. Therefore, it is important for understanding the investors' behavior, which enable researchers to explore potential economical sources of return predictability of the CNN findings.

Our application revolves around equal weighted images to reduce the number of figures, and to match the later regressions ⁷. Figure 3 and 7 show the equal-weighted aggregate of all 5 and 20-day OHLCV images from 2001 to 2019 in the full, positive return, and negative return samples. Each graph component is normalized by the maximum count of the specific component across the whole graph, and its average daily level is shown in green to reflect some column-wise properties. To summarise, the daily averages give some indications of the average trend, while the heat map shows the importance of a component on a day relative to that on other days.

A sensible concern can be raised with regard to normalizing by different numbers across the image, rather than by the maximum count number of the image. However, dividing by the maximum count does not take into account the underlying distribution of each component. For example, the high-low bar covers all the areas between high and low, the counts for this component by construction is higher that that for open and close pixels, which are just one in each day. So if the aggregate image is normalized by the maximum value in the OHLC panel, then the open and close pixel columns will all be close to 0.

Back to the 'Null hypothesis' images, there seems to be a combination of downward and upward trend patterns over the prices but overall there are not many distinguishing trends. For the sub-sample aggregations, the daily average shows an obvious upward trend for the positive sample and a downward one for the negative sample. From the heat map, we see that the first open price and the last close price matter the most. This conforms to our intuition as those two determine the 5-day returns. For now, we will focus on the full sample

⁷Most trends still hold for value-weighted aggregates, and even less noisy for some.

images as it is null for JKX portfolio patterns, and it also motivates our more thorough analyses and discussion in the sub-sample later on.

3 Results

3.1 Heat maps of CNN Portfolios

We now apply our main methodological contribution, and examine it's usefulness. Only images from out-of-sample (2001-2019) is used to remove trends applicable only to the training period. The following discussion revolves around equal weighted heat maps to maintain conciseness. The conclusions mostly hold for value weighted versions and are provided in Appendix A. Figure 3 shows the heat map for all images, and displays noisy or flat trend for all components.

Figure 4 and 8 visualises the 5- and 20-day image portfolios with weekly, monthly, and quarterly holding period in columns one, two and three, respectively. The top row represents the tenth decile (long), and the bottom row represents the first decile (short)⁸. For each specification, we first plot the image on the left, and separately plot the open, high/low, and close on the right. The average value in each column, for each component is shown in green outline. We focus on the 5-day window for now. Consistent with the univariate analysis in the data summary statistics section, the first thing we notice is that the long-end (decile 10) exhibits a trend reversal. Strong downward sloping prices and upward sloping volume over the formation window is associated with highest probability of positive future returns. Although the opposite holds for the short-end (decile 1), the price and volume slopes are less steep⁹.

The relation between prices and dates is not quite linear. For the long-end, the rate of

⁸Images for the other deciles can be made available upon request. We do not report them as the fifth decile, for example, is quite similar to the null image, and because we wish to only report the most useful figures.

⁹For the value weighted, decile 1 reversal is much steeper and is nearly the inverse of decile 10.

decrease in open price, high-low bar, and volume increase towards the last formation day. Again, this is much less obvious for the short-end. **The single, most prominent feature** is the extremely low (high) last day close price for long (short) end. In fact, the last day closing price is so important that previous days' weight is nearly empty, whilst the lowest (highest) pixel of the last day for long (short) end holds almost all weight. Recall the heat map is the weighted average of all images in that portfolio. The darker the image, the higher the density of the images in that pixel. This is the only consistent and obvious pattern regardless of portfolio or image weighting method. This suggest sorting on the average trend (return) over the 5-day is not enough, but rather, the change in price closer to formation day may matter more.

The observed pattern changes as the prediction period increases from weekly to quarterly: The downward (upward) sloping first four-day prices and upward (downward) sloping volume inverses. Nevertheless, the large change in last day close persists. We also notice that as horizon increases, the weight of first four close prices also increases.

Overall, we observe some clear features of the price and volume picked up by CNN models through learning about patterns associated with positive next period return. The most important component seems to be the fifth-day close price; the lowest (highest) decile CNN portfolio are almost certain to have the highest (lowest)¹⁰ last day closing price. This is unobserved in the aggregate for the full sample. The other feature are the nonlinear monotone pattern over the five-days, which include the prices and volumes. We now turn these observation into quantitative measures.

3.2 Measures of Image Patterns and Regression Analysis

Visual patterns are neither theoretical nor empirical evidence. However, it does provide reasonable justification for constructing within formation period trend measures to verify the graphical intuition. To capture the above observations, we first normalize all five day

¹⁰To be more precise, highest (lowest) position in the image.

prices by the maximum of all open, high, low, close and moving average value in the window. This essentially results in a continuous measure of a data point's position in the image. Similarly for volume, we normalise by the maximum of the 5 day volume. Define the normalized measure as $x_{t,f} \in [0, 1]$, where $t \in [1, 5]$ is the day in the image, and $f \in [open, high, low, close, volume]$ is the type of data. We propose three general measure of trend, the level, slope, and curvature

$$\text{Level}(f) = \frac{1}{5} \sum_{i=1}^5 x_{i,f} \quad (4)$$

$$\text{Slope}(f) = x_{5,f} - x_{1,f}; \quad (5)$$

$$\text{Curvature}(f) = (x_{5,f} - x_{3,f}) - (x_{3,f} - x_{1,f}). \quad (6)$$

which are applied to open, high minus low, and volume. We further propose three specific measures to capture the focus on last day features

$$\text{Weekly Relative Last Close (WRLC)} = x_{5,close} - \frac{\sum_{i=1}^4 x_{i,close}}{4} \quad (7)$$

$$\text{Weekly Relative Last Volume (WRLV)} = x_{5,volume} - \frac{\sum_{i=1}^4 x_{i,volume}}{4} \quad (8)$$

$$\text{Weekly Relative Last Close Change (WRLCC)} = (x_{5,close} - x_{5,low}) - \frac{\sum_{i=1}^4 (x_{i,close} - x_{i,low})}{4} \quad (9)$$

where we apply relative differences to volume, close price and difference between close price and high-low average. Relative differences are crucial here to capture if previous days, any other component is also at the lowest pixel position¹¹. In addition, we use the average formation period return as a measure of the general trend and as a control for WSTR as documented by Cooper (2015) and Lehmann (1990).

We test the independent variables under two specifications. First, a simple regression is

¹¹Notice in the heat maps that in the top and bottom deciles, the close price on last day is the only component concentrated at the lowest and highest position.

ran

$$Pr(R5 > 0)_{i,t} = \alpha + \beta_1 I5_{i,t} + \beta_2 WRLC_{i,t} + \beta_3 WRLV_{i,t} + \beta_4 WRLCC_{i,t} + X_{i,t} + \epsilon_{i,t} \quad (10)$$

to determine how much variation in $Pr(R5 > 0)_{i,t}$, the CNN predicted probability of realising a positive return for stock i at time t , can be explained linearly by the proposed measures of trend. $X_{i,t}$ includes the three level, three slope, and three curvature measure. Secondly, we further estimate the overall marginal effect of pattern variables by employing the logistic regression

$$Logit[Pr(R5 > 0)_{i,t} > 0.5] = \alpha + \beta_1 I5_{i,t} + \beta_2 WRLC_{i,t} + \beta_3 WRLV_{i,t} + \beta_4 WRLCC_{i,t} + X_{i,t} + \epsilon_{i,t} \quad (11)$$

where the dependent is encoded into one if the predicted probability is above 50%. The cutoff may seem arbitrary, but when the model is trained, 50% is also the cutoff applied to the binary prediction used to calculate cross-entropy loss.

Table 3 and 4 reports the average marginal effects estimated by the logistic regression times standard deviation of each independent variable ¹² using 5-day and 20-day formation dates, respectively. For both tables, the top, middle, and bottom panel reports the results for weekly, monthly, and quarterly holding period. For each panel, we report the pseudo-r-squared, as well as the adjusted r-squared from OLS with unreported coefficient.

Focusing on the 5-day formation for prediction, the Logit results are consistent with visually observed trends underlying the R5, R20 and R60 predictions. A standard deviation increase in adjusted return and WRLC are associated with a significant 21% decrease in the predicted probability, and each can individually explain 8% and 16% of variation in CNN weekly predictions. Across all specifications, WRLCC reduces the predicted probability by 15% with a standard deviation increase, and can individually explain 15% of variation.

¹²Since our independent variables are not standard in the literature, we make this choice to improve the interpretation of the results.

Lastly, a standard deviation increase in WRLV increases the CNN output by 9%, and explains only 6% of variation. These four variables together can linearly explain 31% of variation in the 27 million CNN prediction, and the nine trends barely adds any further explanatory power.

As the prediction horizon increases, the explanatory power and economic magnitude shifts from the relative last day variables and formation returns to the level variable, with the sole exception of relative last close change (RLCC). This is not to say they have lost all explanatory power as even at quarterly horizon, the four main variables are still significant and explains 11% of variation (although RLCC by itself can explain 10% of variation). The shift is not symmetric; the explanatory power of spread and volume level increases from 0% to 4%, and a standard deviation is associated with -4% and 7% change in the quarterly predictions, respectively. Further, a more convex open price (high-low spread) path, is related to lower (higher) probability, underlining the CNN's ability to capture some non-linearity.

The above observations generally hold for the monthly formation window regressions with monthly trend variables, but OLS R-squared is low by construct as the input information quadruples compared to five day formation windows¹³. The economic magnitude is also quite low for relative last day measures for holding period longer than one week.

The robustness and the strong effects of the last day changes in terms of CNN outputs does not directly tie them to the actual portfolio performances. To do so, we form portfolios using the four last day change measures based on past week. Specifically, we form decile portfolios using the weekly relative last close change (RLCC), relative last close (RLC), and relative last volume (RLV) and compare them with the WSTR. Table 5 shows the results. One can see that the RLCC sorted long-short portfolio and RLC sorted long-short portfolio outperform WSTR. R5 long-short portfolio sorted by RLCC exhibits 49 percent annualized return with the Sharpe Ratio of 3.82. Even for R20, it has 14 percent annualized returns with a Sharpe Ratio of 1.38. The Sharpe Ratios for WSTR are 3.01 and 1.10 respectively for R5

¹³The CNN network architecture is also more complex.

and R20. While RLC and RLV sorted long-short portfolio have relative similar performances as that of WSTR, these are completely new factors results from our portfolio based heat-map approach.

Since CNN captures all of these important “factors”, we further illustrate their significance by regressing the JKK factor returns (I5R5, I5R20 and I5R60) on the returns earned by these factors and report the results in Table 6. RLCC based returns can explain 37% of variation in weekly CNN returns, and 73% when RLC, RLV and WSTR is included. At the monthly horizon, the former can still explain 22% of CNN returns, and 48% with the inclusion of the latter portfolios. Altogether, the four last day change variables can explain a large proportion of variation in not only the CNN predictions, but also the large returns generate by these predictions.

4 Subsample Analysis

Although we have explained a large proportion of CNN’s abilities with simple factors, 27% and 52% of variation CNN weekly and monthly returns, respectively, remains unexplained. Therefore, we further investigate the abilities of CNN by working off three points of related but unaddressed inconsistency. Firstly, although there are upward trends in quarterly portfolio’s closing price, the regression results still supports the existence of downward trends. Secondly, a visual inspection of quarterly closing prices in panel C and F of Figure 4 shows a bi-modal closing price. Thirdly, even in the weekly portfolio, 24% of stocks have formations returns opposing reversal in both ends. It is indicative of CNN’s ability to capture at least two patterns with high and low closing price for each of end of the portfolio. Collectively, it is reasonable to hypothesize that CNN also capture the persistency in trends, in addition to reversal. More precisely, it could reliably predict positive (negative) holding period return, conditioned on positive (negative) formation period return. Such hypothesis is in part motivated by time series momentum, where stocks exhibiting high (low) historical returns

continue to do so (Moskowitz et al., 2012).

We sort decile portfolios in the sub-sample of stocks with positive ($I > 0$) and negative ($I \leq 0$) formation period return, and report the performance and characteristics in the top and bottom panel of Table 7, respectively. Notice that both sorts perform similarly to the full sample sort, and the difference in CNN probability is still around 30% between the top and bottom bin. Under traditional reversal, a positive (negative) I implies a negative (positive) R . Consistent with intuition, the WSTR portfolio is only significant in the reversal decile in the two sub-samples, that is, the decile with lowest return in the negative sample, or the decile with highest return in the positive sample. On the other hand, CNN generate returns in both extreme deciles in both sub-samples in the equally weighted case. In the positive sample, the top decile with a formation return of 3% and 59% probability of realising positive returns earn an annualized return (Sharpe) of 24% (1.44). In the negative sample, the bottom decile with 31% chance of realising positive return earn an annualized return (Sharpe) of -22% (-1.14). Although not significant in the value weighted case ¹⁴, it still indicates that CNN also learns persistency in trends.

As such, we currently have four positions of interest, **1**) Positive sample, decile 1: $R < 0|I > 0$, **2**) Positive sample, decile 10: $R > 0|I > 0$, **3**) Negative sample, decile 1: $R < 0|I < 0$, and **4**) Negative sample, decile 10: $R > 0|I < 0$. Positions 2 and 3 exhibit return persistence, and positions 1 and 4 exhibit return reversal. Figure 6's first and second row shows the heat-maps of positions 2 and 3 for 5-, 20-, 60-day holding periods, respectively. Across all aggregate images, we can observe an obvious upward trend for position 2, leading to a positive return, and an obvious downward trend for position 3, leading to a negative return. Both the univariate statistics and pattern is extremely different from that observed for the reversal positions. The landmark pattern of reversal, extremely high/low last day close, is also no longer observed in the persistence positions.

With this in mind, we should expect to see a more extreme trend in the two reversal

¹⁴Arguably, value weighted is less relevant for the learning abilities of CNN since all observations are weighted equally in the training process

portfolios. Figure 5's first and second row show the corresponding aggregate images for the two reversal portfolios. Consistent with our expectation, the increasing (decreasing) trend is much steeper and focused for position 1 (4), leading to a negative (positive) returns. With the observations also holding for 20-day persistent and reversal portfolios in Figures 10 and 9, the motivating issue have mostly been addressed. In some sense, we have decomposed the JKX factor to show two unique patterns captured by the CNN.

Nevertheless, can the decomposition provide further value than simply some intuition about return driving patterns? We argue that it is counter intuitive to long and short two different signals in a sub-sample. Instead, the portfolio should consist of opposite ends of the same signal. Consequently, we form a trend reversal portfolio by buying position 4 and selling position 1, and a trend persistency portfolio by buying position 2, and selling position 3. The performances across horizons and weighting schemes are documented in Table ???. The long-short persistence portfolio produce an annualized return of 42% and Sharpe ratio 4.41. This is quite different from the short-term reversal pattern we observed. Even at monthly horizon, the persistence portfolio can produce 5% annualized abnormal returns with a Sharpe ratio 52%. It is at the 60-day horizon that the persistence portfolio produce insignificant returns. Under value-weighted portfolio, the return differences disappear, or even become negative under 20-day window, implying that the persistence effect is likely to be from small firms. Nevertheless, this does not discredit the finding that CNN can capture some trends indicative of persistency.

The reversal long-short portfolio produces significant, and better equal weighted results for all three windows. Under 5-day horizon, it produces a stunning annualized return of 105% and a Sharpe ratio 6.74. Under 20-day and 60-day horizon, it produces annualized returns of 27% and 10 % respectively, and Sharpe ratios of 2.45 and 1.21 respectively. They are far larger than those in JKX. For value-weighted portfolios, the reversal strategy at weekly horizon generates 43% annualized returns, nearly double that of JKX, with a 2.11 Sharpe ratio. At the monthly level, JKX value weighted portfolio is not significant at all, while

the reversal strategy earns a significant 8% return with 58% sharpe ratio. The correlation between the JKK, reversal, and persistency portfolio is reported in Table 9 to 11. While reversal is highly correlated to JKK portfolio, it only has a 20% correlation with persistency.

5 Conclusion

Image-based machine learning models have been shown to produce unique and powerful predictors for the cross-sectional returns. Still, one main challenge remains for the vision models applied and machine learning in general: How can we interpret such models, and what do they truly learn? We provide an intuitive way of visualising stock chart patterns learned by Convolutional Neural Networks. Namely, we use heat maps to visualise the machine learnt trends and importance of each input component.

Using this approach, we quantify the learnt patterns and find that characteristics related to the path of graph components play an important role in the CNN models' predictions. These characteristics include the level, slope, curvature, formation period return, and most importantly, the relative change in last day closing price and volume in the learning window. The last feature has significant explanatory power for the variation in the complex predictions made by CNN; Four simple variables can, in fact, explain 31% and 29% of weekly and monthly variation, and 71% and 48% of variation in weekly and monthly CNN returns.

Further analysis of the sorted portfolio heat-maps reveals that there are two predominant type of patterns captured by CNN: persistency and reversal. Although reversal is still the dominant force, it is interesting to observe persistency predicted using only 5-days of historical data. We then form long-short portfolios on the two signal and achieves even better performances than the impressive performances shown in JKK.

Our results naturally lead to one question: How can the characteristics of the return and trading volume path predict future returns? This presents a new challenge for any theoretical asset pricing models.

References

- Bajgrowicz, P., & Scaillet, O. (2012). Technical trading revisited: False discoveries, persistence tests, and transaction costs. *Journal of Financial Economics*, *106*(3), 473–491. <https://doi.org/10.1016/j.jfineco.2012.06.001>
- Barberis, N., Shleifer, A., & Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, *49*(3), 307–343. [https://doi.org/10.1016/S0304-405X\(98\)00027-0](https://doi.org/10.1016/S0304-405X(98)00027-0)
- Brown, D. P., & Jennings, R. H. (1989). On Technical Analysis. *The Review of Financial Studies*, *2*(4), 527–551.
- Chu, B. (2024). Technical analysis with machine learning classification algorithms: Can it still ‘beat’ the buy-and-hold strategy? [Available at SSRN: <https://ssrn.com/abstract=4765615> or <http://dx.doi.org/10.2139/ssrn.4765615>].
- Cooper, M. (2015). Filter Rules Based on Price and Volume in Individual Security Overreaction. *The Review of Financial Studies*, *12*(4), 901–935. <https://doi.org/10.1093/rfs/12.4.901>
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, *33*, 2223–2273.
- Han, Y., Zhou, G., & Zhu, Y. (2016). A trend factor: Any economic gains from using information over investment horizons? *Journal of Financial Economics*, *122*(2), 352–375. <https://doi.org/10.1016/j.jfineco.2016.01.029>
- Jegadeesh, N. (2000). Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation: Discussion. *The Journal of Finance*, *55*(4), 1765–1770.
- Jiang, J., Kelly, B., & Xiu, D. (2023). (re-)imag(in)ing price trends. *The Journal of Finance*, *78*(6), 3193–3249. <https://doi.org/https://doi.org/10.1111/jofi.13268>
- Lehmann, B. N. (1990). Fads, martingales, and market efficiency. *The Quarterly Journal of Economics*, *105*(1), 1–28. Retrieved July 28, 2024, from <http://www.jstor.org/stable/2937816>
- Lo, A. W., Mamaysky, H., & Wang, J. (2000). Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The Journal of Finance*, *55*(4), 1705–1765. <https://doi.org/https://doi.org/10.1111/0022-1082.00265>
- Moskowitz, T. J., Ooi, Y. H., & Pedersen, L. H. (2012). Time series momentum [Special Issue on Investor Sentiment]. *Journal of Financial Economics*, *104*(2), 228–250. <https://doi.org/https://doi.org/10.1016/j.jfineco.2011.11.003>
- Murray, S., Xia, Y., & Xiao, H. (2024). Charting by machines. *Journal of Financial Economics*, *153*, 103791. <https://doi.org/10.1016/j.jfineco.2024.103791>
- Ready, M. J. (2002). Profits from Technical Trading Rules. *Financial Management*, *31*(3), 43–61. <https://doi.org/10.2307/3666314>
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2019). Grad-cam: Visual explanations from deep networks via gradient-based localization. *International Journal of Computer Vision*, *128*(2), 336–359. <https://doi.org/10.1007/s11263-019-01228-7>

Zhu, Y., & Zhou, G. (2009). Technical analysis: An asset allocation perspective on the use of moving averages. *Journal of Financial Economics*, 92(3), 519–544. <https://doi.org/10.1016/j.jfineco.2008.07.002>

Tables

Table 1: Summary Statistics

This table reports the mean, standard deviation, minimum, and percentiles in steps of 25% for non-winsorized variables for generating images from 1990 to 2019. Open, ask high, bid low, and close price refers to the first, highest, lowest, and last trading price in a day. Returns are adjusted for dividends, and market value is computed as the product between number of shares and closing price.

	N.Obs	Mean	S.D.	Min	25%	50%	75%	Max
Open Price	44.24M	40.64	1776.33	0.00	6.12	14.19	27.50	340807.25
Ask High	46.53M	39.82	1742.70	0.00	6.06	14.12	27.33	342250.00
Bid Low	46.53M	38.95	1721.13	0.00	5.75	13.75	26.53	339312.59
Close Price	46.53M	39.40	1732.21	0.00	5.94	13.94	26.97	340380.00
Return	46.53M	0.00	0.05	-0.97	-0.01	0.00	0.01	19.00
Volume	46.53M	0.61M	4.17M	0	9118	56900	291100	1897.90M
No. Shares	46.53M	74521.46	0.31m	1.00	7389.00	19214.00	50072.00	29.21M
Market Value	46.53M	2.60M	14.68M	1.75	0.05M	0.21M	0.94M	1288.15M

Table 2: Weekly Portfolio Return and Characteristics

This table reports the mean characteristics and annualized return (Ret) and sharpe ratio (SR) of decile portfolio formed over 2001 to 2019 with one week holding period. The top panel portfolio sorts on $P(R5 > 0)$, the CNN predicted probability of positive next week returns, with the bottom panel sorting on the one week return ($I5$). $\frac{N(I5>0)}{N}$ and $\frac{N(R5>0)}{N}$ is the proportion of stocks with positive formation and realised returns, respectively. Market capitalization (MC) is reported in billions. $I5$ is the one week return prior to portfolio formation, and $R5$ is the realised one week return. **Bold** and *italic* indicates significance at the 1% and 5% level, respectively.

	EW		VW		Mean Characteristics				
	Ret	SR	Ret	SR	$\frac{N(I5>0)}{N}$	$\frac{N(R5>0)}{N}$	MC	$P(R5 > 0)$	$I5$
<hr/> $P(R5 > 0)$ <hr/>									
Low	-0.32	-2.16	<i>-0.09</i>	-0.52	0.66	0.44	1.65	0.29	0.02
2	<i>-0.09</i>	-0.52	-0.05	-0.27	0.64	0.48	2.97	0.38	0.02
3	-0.02	-0.08	0.00	0.00	0.63	0.49	3.58	0.42	0.02
4	0.04	0.21	-0.00	-0.02	0.60	0.50	3.95	0.44	0.02
5	0.07	0.36	0.01	0.04	0.56	0.51	4.18	0.46	0.01
6	<i>0.10</i>	0.50	0.03	0.16	0.52	0.51	4.27	0.48	0.00
7	0.14	0.67	0.05	0.29	0.47	0.52	4.26	0.50	0.00
8	0.19	0.91	0.08	0.43	0.41	0.52	4.09	0.53	-0.01
9	0.26	1.23	0.12	0.62	0.34	0.53	3.70	0.56	-0.02
High	0.48	2.47	0.14	0.71	0.24	0.56	2.41	0.64	-0.04
H - L	0.80	7.18	0.23	1.57					
<hr/> Weekly Short Term Reversal (WSTR) <hr/>									
Low	0.45	1.47	<i>0.17</i>	0.56	0.00	0.52	1.22	0.53	-0.11
2	0.17	0.74	0.15	0.62	0.01	0.52	2.79	0.51	-0.05
3	<i>0.11</i>	0.57	0.12	0.61	0.08	0.52	3.74	0.49	-0.03
4	0.08	0.43	0.06	0.33	0.23	0.52	4.17	0.48	-0.01
5	0.06	0.36	0.04	0.26	0.44	0.52	4.36	0.47	0.00
6	0.05	0.28	0.02	0.15	0.64	0.52	4.53	0.46	0.00
7	0.04	0.26	-0.01	-0.04	0.79	0.51	4.56	0.45	0.01
8	0.03	0.17	0.01	0.04	0.91	0.50	4.36	0.44	0.03
9	0.02	0.10	-0.04	-0.19	0.98	0.49	3.63	0.44	0.05
High	-0.15	-0.63	-0.08	-0.35	1.00	0.46	1.72	0.43	0.13
H - L	-0.59	-3.15	-0.25	-1.10					

Table 3: Explaining CNN binary forecast with observable 5-day characteristics.

Panel A, B and C reports the logistic regression results of 5-day image based weekly, monthly, and quarterly CNN binary prediction $1(Pr(R > 0) > 0.5)$ against general characteristics observed in the 5-day heat map visualisation. Adjusted return is the 5-day cumulative return adjusted for dividends. The remaining geometrical variables are calculated with the scaled price variables. Given a price, the scaled price is $\frac{price - \min(allprices)}{\max(allprices)}$ where *allprices* is a vector of all 5-day price variables. Defining close change as the difference between close price and average of highest and lowest price, relative last close, close change, and volume are the last day value minus the average of the previous four days'. Level is the five-day average, slope is last day value minus the first, and curve is as the difference between day 3-5 slope and day 1-3 slope. The overall marginal effect times the standard deviation of the independent variable is reported. Open Level is excluded due to co-linearity. We also show the adjusted R-squared from OLS, in addition to pseudo R-squared. *** indicates significance at the 1% level.

Panel A: Weekly Prediction (I5R5)											
Adjusted Return	-0.20***										
Relative Last											
Close		-0.21***									
Close Change			-0.15***								
Volume				0.10***	0.09***						
Spread Level						0.01***					
Volume Level						-0.01***					
Open Slope							-0.11***				
Spread Slope							0.03***				
Volume Slope							0.08***				
Open Curve								-0.02***	-0.06***	-0.04***	-0.04***
Spread Curve								-0.00***	-0.00***	-0.00***	0.00***
Volume Curve								0.04***	0.01***	0.02***	0.01***
nobs	30M	27M	27M	30M	27M	27M	27M	27M	27M	27M	27M
pseudo.rsq	0.09	0.13	0.07	0.05	0.24	0.00	0.08	0.01	0.27	0.25	0.25
adjusted.rsq	0.08	0.16	0.15	0.06	0.31	0.00	0.09	0.01	0.34	0.31	0.31

Panel B: Monthly Prediction (I5R20)												
Adjusted Return	-0.05***									0.01***	-0.04***	-0.04***
Relative Last												
Close		-0.11***								0.02***	-0.02***	-0.02***
Close Change			-0.17***							-0.22***	-0.18***	-0.18***
Volume				0.10***	0.11***					0.02***	0.02***	0.02***
Spread Level						-0.02***				-0.00***	-0.00***	
Volume Level						0.07***				0.07***	0.07***	0.07***
Open Slope							-0.00***			-0.10***		
Spread Slope							0.00***			0.03***	0.03***	0.03***
Volume Slope							0.11***			0.10***	0.10***	0.10***
Open Curve								-0.03***		-0.05***	-0.04***	-0.04***
Spread Curve								-0.01***		-0.01***	-0.01***	-0.01***
Volume Curve								0.01***		0.01***	0.01***	0.01***
nobs	30M	27M	27M	30M	27M	27M	27M	27M	27M	27M	27M	27M
pseudo.rsq	0.01	0.03	0.07	0.03	0.13	0.02	0.04	0.00	0.17	0.16	0.16	
adjusted.rsq	0.02	0.07	0.15	0.07	0.24	0.03	0.08	0.01	0.30	0.29	0.29	
Panel C: Quarterly Prediction (I5R60)												
Adjusted Return	-0.01***									0.01***	-0.02***	-0.03***
Relative Last												
Close		-0.05***								0.06***	0.02***	0.03***
Close Change			-0.14***							-0.19***	-0.16***	-0.16***
Volume				0.04***	0.04***					0.00***	0.00***	0.01***
Spread Level						-0.04***				-0.05***	-0.05***	
Volume Level						0.07***				0.07***	0.07***	0.08***
Open Slope							0.03***			-0.06***		
Spread Slope							-0.05***			-0.06***	-0.06***	-0.05***
Volume Slope							0.06***			0.06***	0.06***	0.06***
Open Curve								-0.03***		-0.05***	-0.04***	-0.05***
Spread Curve								0.01***		0.00***	0.00***	0.01***
Volume Curve								0.01***		0.01***	0.01***	0.01***
nobs	30M	27M	27M	30M	27M	27M	27M	27M	27M	27M	27M	27M
pseudo.rsq	0.00	0.01	0.05	0.00	0.05	0.02	0.01	0.00	0.09	0.09	0.08	
adjusted.rsq	0.00	0.02	0.10	0.01	0.11	0.04	0.02	0.00	0.17	0.17	0.16	

Table 4: Explaining CNN binary forecast with observable 20-day characteristics.

Panel A, B and C reports the logistic regression results of 20-day image based weekly, monthly, and quarterly CNN binary prediction $1(Pr(R > 0) > 0.5)$ against general characteristics observed in the 20-day heat map visualisation. Adjusted return is the 20-day cumulative return adjusted for dividends. The remaining geometrical variables are calculated with the scaled price variables. Given a price, the scaled price is $\frac{price - \min(allprices)}{\max(allprices)}$ where *allprices* is a vector of all 20-day price variables. Defining close change as the difference between close price and average of highest and lowest price, relative last close, close change, and volume are the last day value minus the average of the previous 19 days'. Level is the 20-day average, slope is last day value minus the first, and curve is as the difference between day 10-20 slope and day 1-10 slope. The overall marginal effect times the standard deviation of the independent variable is reported. Open Level is excluded due to co-linearity. We also show the adjusted R-squared from OLS, in addition to pseudo R-squared. *** indicates significance at the 1% level.

Panel A: Weekly Prediction (I20R5)											
Adjusted Returns	-0.01***				0.01***				0.02***	0.02***	0.01***
Relative Last											
Close		-0.15***			-0.12***				-0.08***	-0.11***	-0.11***
Close Change			-0.14***		-0.11***				-0.13***	-0.12***	-0.11***
Volume				0.04***	0.04***				0.02***	0.02***	0.02***
Spread Level						-0.04***			-0.06***	-0.06***	
Volume Level						0.05***			0.04***	0.04***	0.06***
Open Slope							-0.05***		-0.03***		
Spread Slope							0.01***		0.00***	0.00***	0.01***
Volume Slope							0.04***		0.03***	0.02***	0.02***
Open Curve								-0.05***	-0.03***	-0.01***	-0.01***
Spread Curve								0.00***	0.00***	0.00***	0.01***
Volume Curve								0.01***	0.00***	0.00***	0.00***
nobs	30M	26M	26M	29M	26M	26M	26M	26M	26M	26M	26M
pseudo.rsq	0.00	0.05	0.05	0.01	0.09	0.02	0.01	0.01	0.11	0.11	0.10
adjusted.rsq	0.00	0.10	0.12	0.01	0.17	0.04	0.03	0.02	0.21	0.21	0.20

Panel B: Monthly Prediction (I20R20)											
Adjusted Return	0.02***				0.03***				0.03***	0.03***	0.03***
Relative Last											
Close		-0.05***			-0.04***				-0.04***	-0.03***	-0.03***
Close Change			-0.06***		-0.05***				-0.05***	-0.05***	-0.05***
Volume				0.02***	0.02***				0.01***	0.01***	0.01***
Spread Level						-0.05***			-0.05***	-0.05***	
Volume Level						0.05***			0.05***	0.05***	0.06***
Open Slope							-0.00***		0.01***		
Spread Slope							0.00***		-0.00***	-0.00***	-0.00***
Volume Slope							0.02***		0.02***	0.02***	0.02***
Open Curve								-0.03***	-0.00***	-0.01***	-0.01***
Spread Curve								0.00***	0.00***	0.00***	0.00***
Volume Curve								-0.00***	-0.00***	-0.00***	-0.00***
nobs	30M	26M	26M	29M	26M	26M	26M	26M	26M	26M	26M
pseudo.rsq	0.00	0.01	0.01	0.00	0.02	0.02	0.00	0.00	0.03	0.03	0.03
adjused.rsq	0.00	0.01	0.02	0.00	0.03	0.04	0.00	0.00	0.07	0.07	0.06
Panel C: Quarterly Prediction (I20R60)											
Adjusted Return	0.02***				0.03***				0.03***	0.03***	0.03***
Relative Last											
Close		-0.03***			-0.02***				-0.03***	-0.01***	-0.01***
Close Change			-0.03***		-0.03***				-0.03***	-0.03***	-0.03***
Volume				0.01***	0.01***				-0.01***	-0.01***	-0.01***
Spread Level						-0.07***			-0.07***	-0.07***	
Volume Level						0.07***			0.07***	0.07***	0.09***
Open Slope							0.00***		0.01***		
Spread Slope							0.00***		-0.01***	-0.01***	-0.01***
Volume Slope							0.02***		0.03***	0.03***	0.03***
Open Curve								-0.02***	-0.00***	-0.01***	-0.01***
Spread Curve								-0.00***	-0.00***	-0.00***	-0.00***
Volume Curve								-0.01***	-0.00***	-0.00***	-0.00***
nobs	30M	26M	26M	29M	26M	26M	26M	26M	26M	26M	26M
pseudo.rsq	0.00	0.00	0.00	0.00	0.01	0.03	0.00	0.00	0.04	0.04	0.03
adjused.rsq	0.00	0.00	0.01	0.00	0.01	0.07	0.00	0.00	0.09	0.09	0.06

Table 5: Weekly Factor Portfolios

This table reports the annualized return and Sharpe ratio of weekly, monthly, and quarterly portfolios formed with 5-day geometric factors, and weekly short term reversal over 2001-2019. Significance at the 5 and 10 percent level are indicated by **bold** and *italic*, respectively.

		R5		I5 R20		R60	
		RET	SR	RET	SR	RET	SR
Relative Last Close Change (RLCC)	Low	0.35	1.33	0.16	0.58	0.12	0.37
	High	-0.14	-0.54	0.03	0.11	0.07	0.22
	L-H	0.49	3.82	0.14	1.38	0.05	0.53
Relative Last Close (RLC)	Low	0.42	1.40	0.19	0.61	0.12	0.35
	High	-0.12	-0.53	0.02	0.07	0.06	0.22
	L-H	0.54	3.06	0.17	1.12	0.05	0.40
Weekly Adjusted Return (WSTR)	Low	0.44	1.45	0.20	0.63	0.12	0.36
	High	-0.13	-0.57	0.02	0.07	0.06	0.22
	L-H	0.57	3.01	0.18	1.10	0.06	0.40
Relative Last Volume (RLC)	High	0.16	0.93	0.13	0.65	0.11	0.49
	Low	-0.01	-0.06	0.07	0.34	0.08	0.38
	H-L	0.17	2.67	0.06	1.04	0.03	0.47

Table 6: Explaining 5-day Image Based Returns

OLS regression of the weekly, monthly, and quarterly returns based on 5-day images against returns of geometric factors. WRLCC, WRLC, WSTR, and WRLV stands for the weekly relative last close change, weekly relative last close, weekly short term momentum, and weekly relative last volume. Specifically, the I5R5/20/60 timeseries returns are regressed against the returns achieved by sorting on observed patterns. Newey West standard errors are reported in brackets. ***, **, and * indicate significance at the 1%, 5%, and 10% level, respectively.

I5R5 Weekly Return				
const	-0.52*** (0.03)	-0.65*** (0.02)	-0.73*** (0.01)	-1.03*** (0.04)
WRLCC	0.53*** (0.03)	0.28*** (0.02)	0.37*** (0.01)	0.36*** (0.01)
WRLC		0.38*** (0.02)	0.08*** (0.02)	0.08*** (0.02)
WSTR			0.29*** (0.02)	0.27*** (0.02)
WRLV				0.33*** (0.04)
R-squared	0.37	0.64	0.69	0.73
R-squared Adj.	0.37	0.64	0.69	0.73
N	4779	4779	4779	4779
I5R20 Monthly Return				
const	-0.37*** (0.02)	-0.47*** (0.02)	-0.52*** (0.02)	-0.85*** (0.04)
WRLCC	0.38*** (0.02)	0.26*** (0.02)	0.31*** (0.02)	0.30*** (0.01)
WRLC		0.22*** (0.02)	0.08*** (0.03)	0.08*** (0.02)
WSTR			0.13*** (0.03)	0.13*** (0.02)
WRLV				0.35*** (0.02)
R-squared	0.22	0.39	0.41	0.48
R-squared Adj.	0.22	0.39	0.41	0.48
N	4779	4779	4779	4779
I5R60 Quarterly Return				
const	-0.18*** (0.02)	-0.12*** (0.02)	-0.10*** (0.02)	-0.19*** (0.05)
WRLCC	0.20*** (0.02)	0.25*** (0.02)	0.23*** (0.02)	0.23*** (0.02)
WRLC		-0.12*** (0.02)	-0.06*** (0.03)	-0.06*** (0.03)
WSTR			-0.06*** (0.02)	-0.06*** (0.02)
WRLV				0.09*** (0.04)
R-squared	0.05	0.10	0.10	0.10
R-squared Adj.	0.05	0.10	0.10	0.10
N	4779	4779	4779	4779

Table 7: Sub-sample Weekly Portfolio Characteristics

This table reports the mean characteristics, and annualized return and sharpe ratio (SR) of decile portfolio formed over 2001 to 2019 with one week holding period, using the sample of equities with positive return in formation period. I is the one week return prior to portfolio formation, and R is the realised one week return. The top panel portfolio is formed using CNN predicted probabilities of $R > 0$, with the bottom panel using I . **Bold** and *italic* indicates significance at the 1% and 5% level, respectively.

		EW		VW		Mean Characteristics				
		Ret	SR	Ret	SR	$\frac{N(I>0)}{N}$	$\frac{N(R>0)}{N}$	Cap	$Pr(> 0)$	I
$I5 > 0$	$Pr(R5 > 0)$									
	Low	-0.45	-2.88	-0.22	-1.14	1.00	0.41	1.51	0.26	0.05
	4	-0.05	-0.28	-0.06	-0.28	1.00	0.49	3.90	0.42	0.05
	7	0.05	0.25	-0.04	-0.24	1.00	0.50	4.49	0.47	0.04
	High	0.24	1.44	-0.04	-0.17	1.00	0.53	2.89	0.59	0.03
	H - L	0.69	6.45	0.18	0.92					
	WSTR									
	Low	-0.06	-0.42	<i>-0.09</i>	-0.54	1.00	0.52	3.58	0.45	0.00
	4	0.02	0.11	-0.02	-0.09	1.00	0.50	4.53	0.44	0.02
	7	0.01	0.03	-0.01	-0.04	1.00	0.48	3.99	0.44	0.04
	High	-0.30	-1.16	-0.13	-0.43	1.00	0.44	1.06	0.43	0.19
	H - L	-0.25	-1.41	-0.03	-0.15					
$I5 \leq 0$	$Pr(R5 > 0)$									
	Low	-0.20	-1.25	-0.01	-0.05	0.00	0.46	1.54	0.31	-0.02
	4	<i>0.13</i>	0.59	0.05	0.27	0.00	0.51	3.52	0.47	-0.04
	7	0.25	1.10	0.12	0.60	0.00	0.53	3.81	0.53	-0.05
	High	0.64	3.01	0.22	1.09	0.00	0.59	1.97	0.67	-0.06
	H - L	0.84	6.42	0.23	1.49					
	WSTR									
	Low	0.68	1.85	0.23	0.66	0.00	0.51	0.76	0.53	-0.15
	4	0.21	0.95	0.16	0.74	0.00	0.52	3.05	0.50	-0.04
	7	<i>0.10</i>	0.58	0.08	0.42	0.00	0.52	4.12	0.49	-0.02
	High	0.03	0.20	0.02	0.13	0.00	0.52	3.56	0.46	0.00
	H - L	-0.65	-2.44	-0.21	-0.78					

Table 8: Portfolio Returns Across Different Horizons

This table reports the annualised return and Sharpe ratio of all out-of-sample (2001-2019) portfolios. Specifically, the weekly (R5), monthly (R20), and quarterly (R60) re-balancing equal/value weighted portfolios formed with 5 and 20 day image factors, which are JKX, Reversal, and Persistency. Significance at 0.05, and 0.01 is reported in **bold**, and *italic* font, respectively.

		R5				R20				R60			
		I5		I20		I5		I20		I5		I20	
		RET	SR	RET	SR	RET	SR	RET	SR	RET	SR	RET	SR
EW													
JKX	Low	-0.31	-2.08	-0.34	-2.01	-0.01	-0.05	0.01	0.05	0.06	0.24	0.06	0.22
	High	0.48	2.53	0.40	2.15	0.18	0.91	0.14	0.72	0.12	0.62	0.10	0.52
	H-L	0.80	7.08	0.74	7.10	0.19	2.39	0.13	1.84	0.06	0.87	0.05	0.50
Reversal	Low	-0.44	-2.78	-0.44	-2.72	-0.05	-0.26	-0.01	-0.08	0.04	0.18	0.04	0.18
	High	0.63	3.14	0.59	2.77	0.22	1.08	0.19	0.86	0.14	0.65	0.11	0.50
	H-L	1.07	6.84	1.02	6.27	0.27	2.40	0.20	1.86	0.10	1.09	0.06	0.66
Persistency	Low	-0.20	-1.21	-0.28	-1.48	0.02	0.12	0.02	0.10	0.07	0.27	0.06	0.23
	High	0.24	1.46	0.20	1.32	0.10	0.61	0.08	0.49	0.10	0.55	0.09	0.51
	H-L	0.44	4.57	0.48	4.03	0.08	0.93	0.06	0.52	0.03	0.31	0.03	0.21
VW													
JKX	Low	-0.10	-0.56	-0.08	-0.41	-0.01	-0.03	0.02	0.10	0.03	0.14	<i>0.02</i>	0.12
	High	0.15	0.77	0.09	0.49	<i>0.04</i>	0.20	<i>0.03</i>	0.19	0.03	0.17	0.03	0.19
	H-L	0.25	1.76	0.17	1.44	0.04	0.34	<i>0.01</i>	0.16	0.00	0.03	0.01	0.11
Reversal	Low	-0.17	-0.86	-0.14	-0.77	<i>-0.03</i>	-0.18	-0.00	-0.02	0.01	0.06	0.00	0.01
	High	0.19	0.89	0.16	0.76	0.06	0.29	0.07	0.40	0.03	0.18	0.05	0.25
	H-L	0.35	1.82	0.31	1.72	0.09	0.53	0.08	0.59	0.02	0.14	0.04	0.31
Persistency	Low	-0.03	-0.14	-0.05	-0.26	0.01	0.05	0.02	0.11	0.03	0.15	0.03	0.14
	High	-0.01	-0.05	-0.02	-0.09	<i>-0.04</i>	-0.20	<i>-0.03</i>	-0.19	-0.00	-0.01	-0.00	-0.02
	H-L	0.02	0.10	0.04	0.20	-0.04	-0.28	-0.06	-0.37	-0.03	-0.22	-0.03	-0.23

Table 9: Weekly Returns Time Series Correlation

This table reports the time series correlation between the JKX, reversal, and persistency portfolios over 2001 to 2019, with weekly holding period.

		EW			VW		
		JKX	Persistency	Reversal	JKX	Persistency	Reversal
EW	JKX	1.000000	0.391931	0.899473	0.475947	0.049218	0.497973
	Persistency	0.391931	1.000000	0.207607	0.141980	0.254371	0.069875
	Reversal	0.899473	0.207607	1.000000	0.374053	-0.050168	0.487939
VW	JKX	0.475947	0.141980	0.374053	1.000000	0.293464	0.598614
	Persistency	0.049218	0.254371	-0.050168	0.293464	1.000000	-0.010979
	Reversal	0.497973	0.069875	0.487939	0.598614	-0.010979	1.000000

Table 10: Monthly Returns Time Series Correlation

This table reports the time series correlation between the JKX, reversal, and persistency portfolios over 2001 to 2019, with monthly holding period.

		EW			VW		
		JKX	Persistency	Reversal	JKX	Persistency	Reversal
EW	JKX	1.000000	0.418583	0.891834	0.540917	0.159124	0.424493
	Persistency	0.418583	1.000000	0.176512	0.244811	0.439558	0.018730
	Reversal	0.891834	0.176512	1.000000	0.476663	0.078156	0.439232
VW	JKX	0.540917	0.244811	0.476663	1.000000	0.414779	0.679822
	Persistency	0.159124	0.439558	0.078156	0.414779	1.000000	0.060179
	Reversal	0.424493	0.018730	0.439232	0.679822	0.060179	1.000000

Table 11: Quarterly Returns Time Series Correlation

This table reports the time series correlation between the JKX, reversal, and persistency portfolios over 2001 to 2019, with quarterly holding period.

		EW			VW		
		JKX	Persistency	Reversal	JKX	Persistency	Reversal
EW	JKX	1.000000	0.756213	0.733984	0.468076	0.367740	0.243299
	Persistency	0.756213	1.000000	0.209047	0.439192	0.513012	0.193362
	Reversal	0.733984	0.209047	1.000000	0.310423	0.091714	0.209783
VW	JKX	0.468076	0.439192	0.310423	1.000000	0.724288	0.554279
	Persistency	0.367740	0.513012	0.091714	0.724288	1.000000	0.288049
	Reversal	0.243299	0.193362	0.209783	0.554279	0.288049	1.000000

Figures

Figure 1: Example of 20-day OHLC Image with volume and moving average.

This is an example 20-day OHLC image. Each day of historical data uses three columns, the first of which represents the daily open price, second of which represents the high-low bar, and last of which represents the close price. The bottom proportion shows volume bars, and a 20-day moving average line is overlaid onto the OHLC proportion.




Figure 2: 5-day Image CNN Structure.

This figure shows the network structure of models that use 5-day images as inputs. The left most black and white image is an example input. $D \times H \times W$ stands for D channels, each with height H and width W . n is the number of classes the model predicts. C stands for a convolution layer, and M stands for a max pool layer, followed by their respective kernel sizes. After each convolution layer, the output is passed through a batch normalization layer, then a leaky ReLU layer. The flattened features are passed through a fully connected layer with n outputs, and the final softmax layer yields n probabilities (FCS).




Figure 3: 'Null' 5-day Images

This figure shows the aggregate 5-day pattern of the full, positive return, and negative return sample from 2001 to 2019 in panel A, B, and C, respectively. Aggregation is equal weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level of a component within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all 5-days of open together. Moving average line is excluded during aggregation.




Figure 4: Visualized Full Sample Weekly, Monthly, and Quarterly 5-day Image Portfolios

This figure, in columns one, two, and three, shows the aggregate 5-day pattern of weekly, monthly, and quarterly long-short portfolio formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return. Aggregation is equal weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.

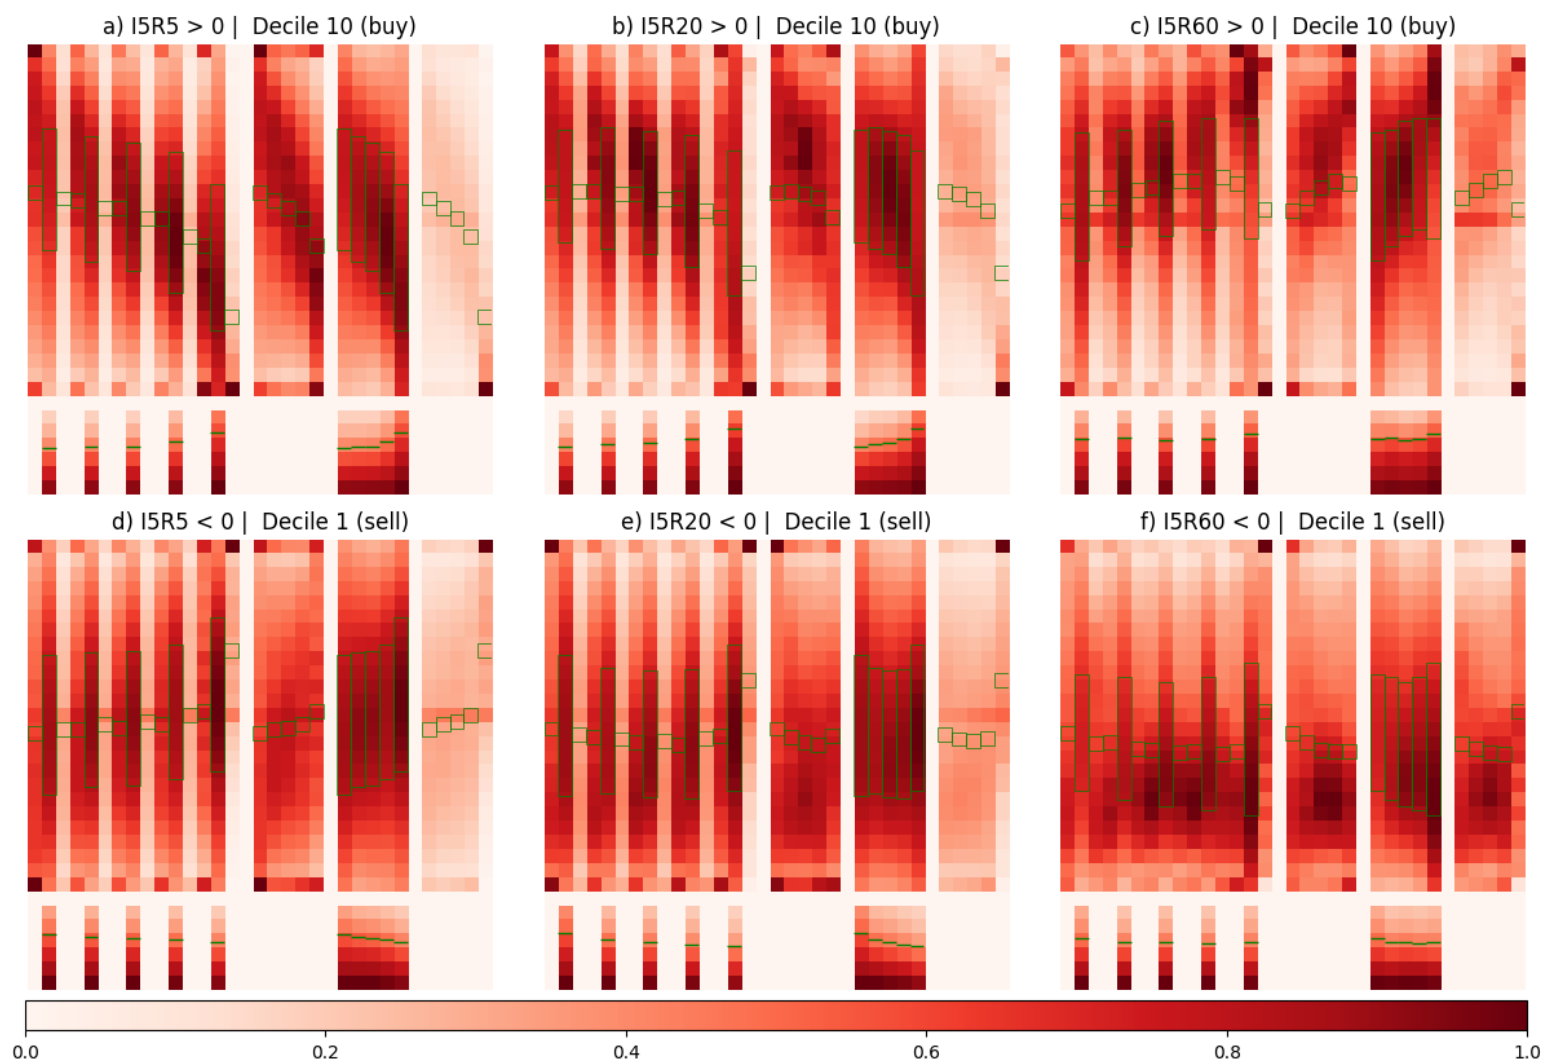


Figure 5: Visualized Weekly, Monthly, and Quarterly 5-day Trend Reversal Portfolios

This figure, in columns one, two, and three, shows the aggregate 5-day pattern of weekly, monthly, and quarterly long-short *reversal* portfolio formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return, in the sample with negative (positive) returns. Aggregation is equal weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.

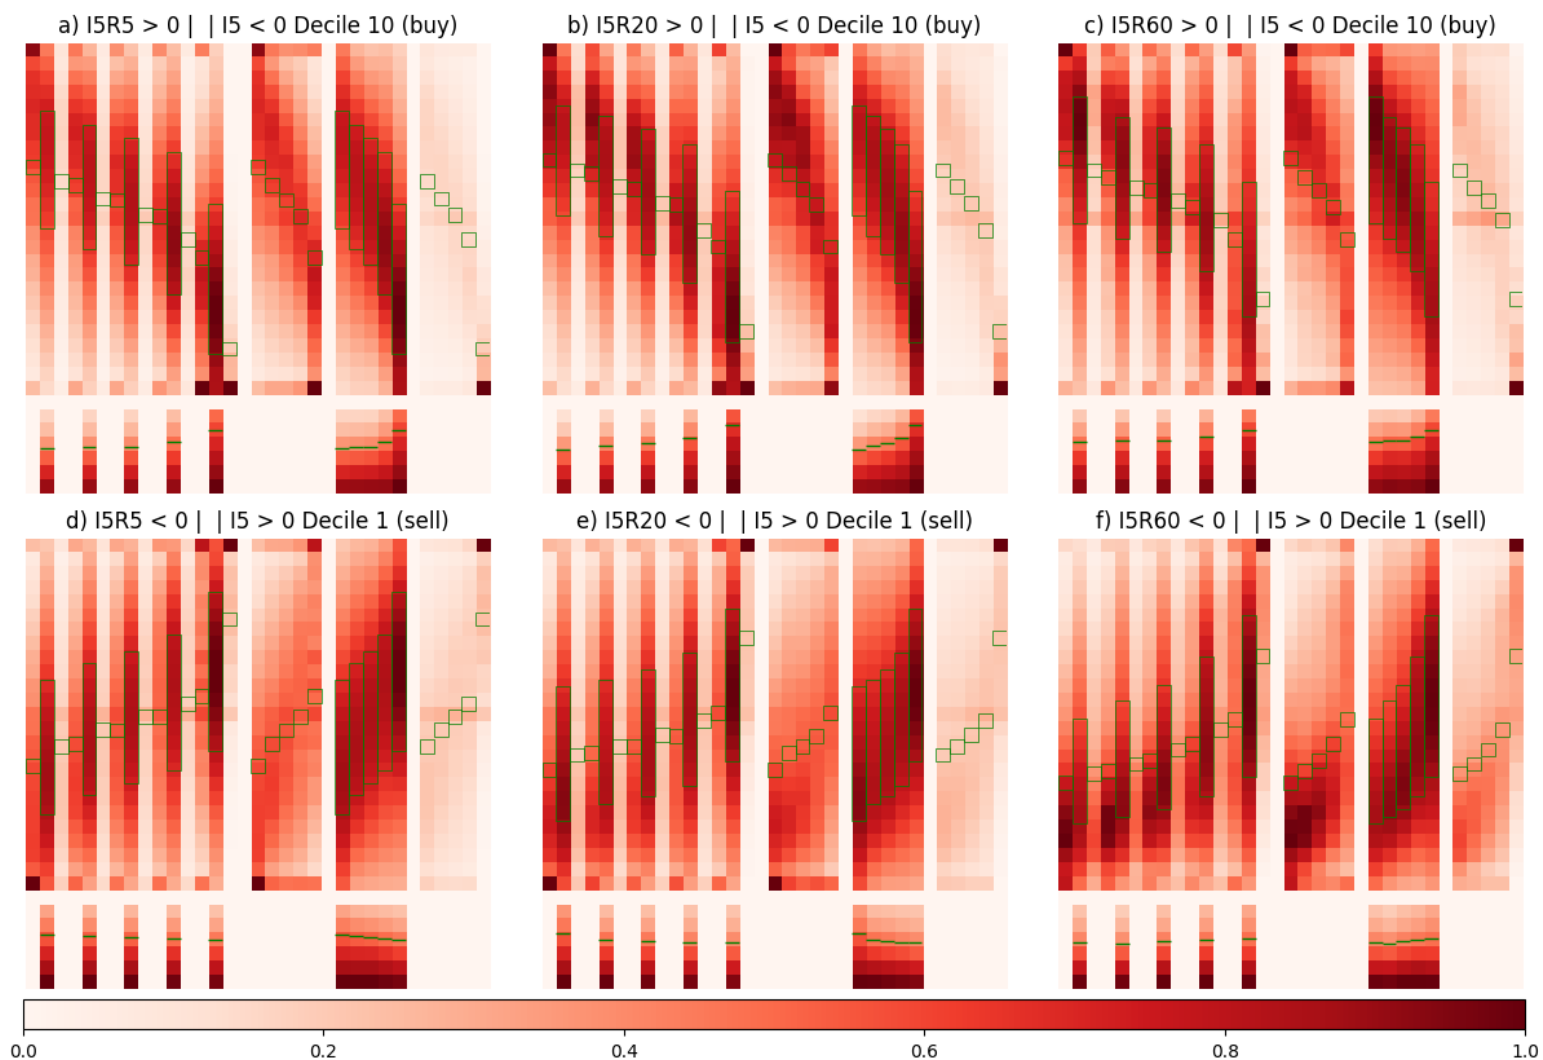
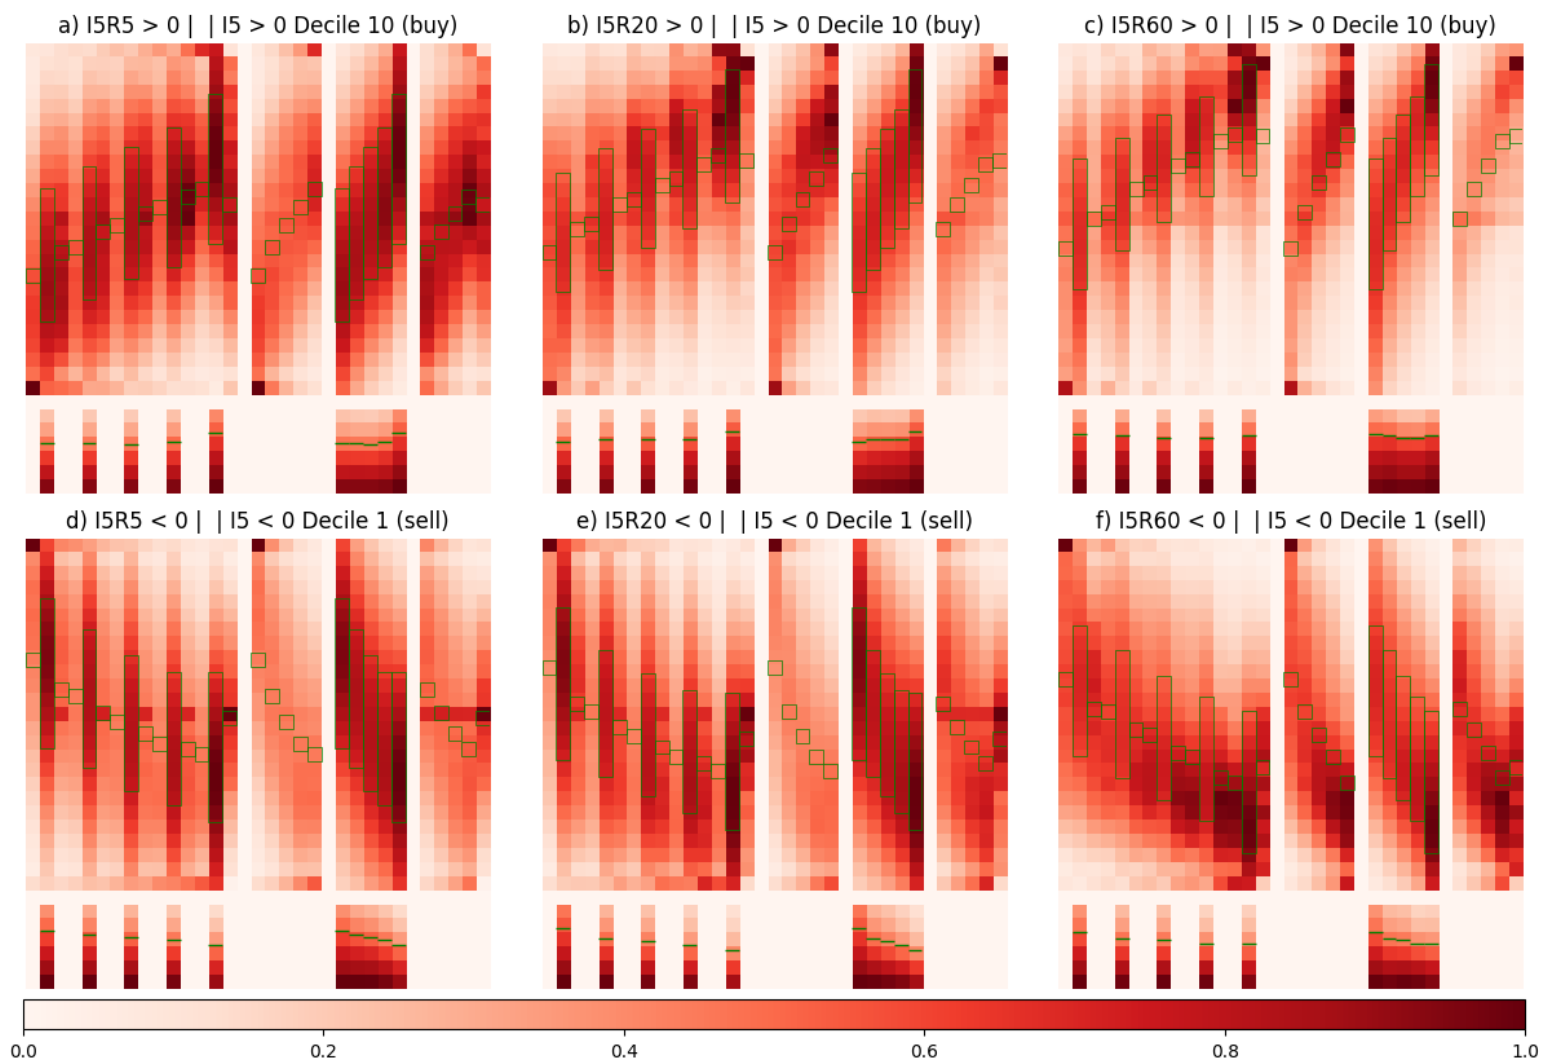


Figure 6: Visualized Weekly, Monthly, and Quarterly 5-day Trend Persistency Portfolios

This figure, in columns one, two, and three, shows the aggregate 5-day pattern of weekly, monthly, and quarterly long-short **persistence** portfolio formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return, in the sample with positive (negative) returns. Aggregation is equal weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.



Appendix A. Value Weighted Heatmaps

We show the value weighted heat-maps here for reference.

Figure 7: 'Null' 20-day Images

This figure shows the aggregate 20-day pattern of the full, positive return, and negative return sample from 2001 to 2019 in panel A, B, and C, respectively. Aggregation is equal weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all 20-days of open together.




Figure 8: Visualized Full Sample Weekly, Monthly, and Quarterly 20-day Image Portfolios

This figure, in columns one, two, and three, shows the aggregate 20-day pattern of weekly, monthly, and quarterly long-short portfolio formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return. Aggregation is equal weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.

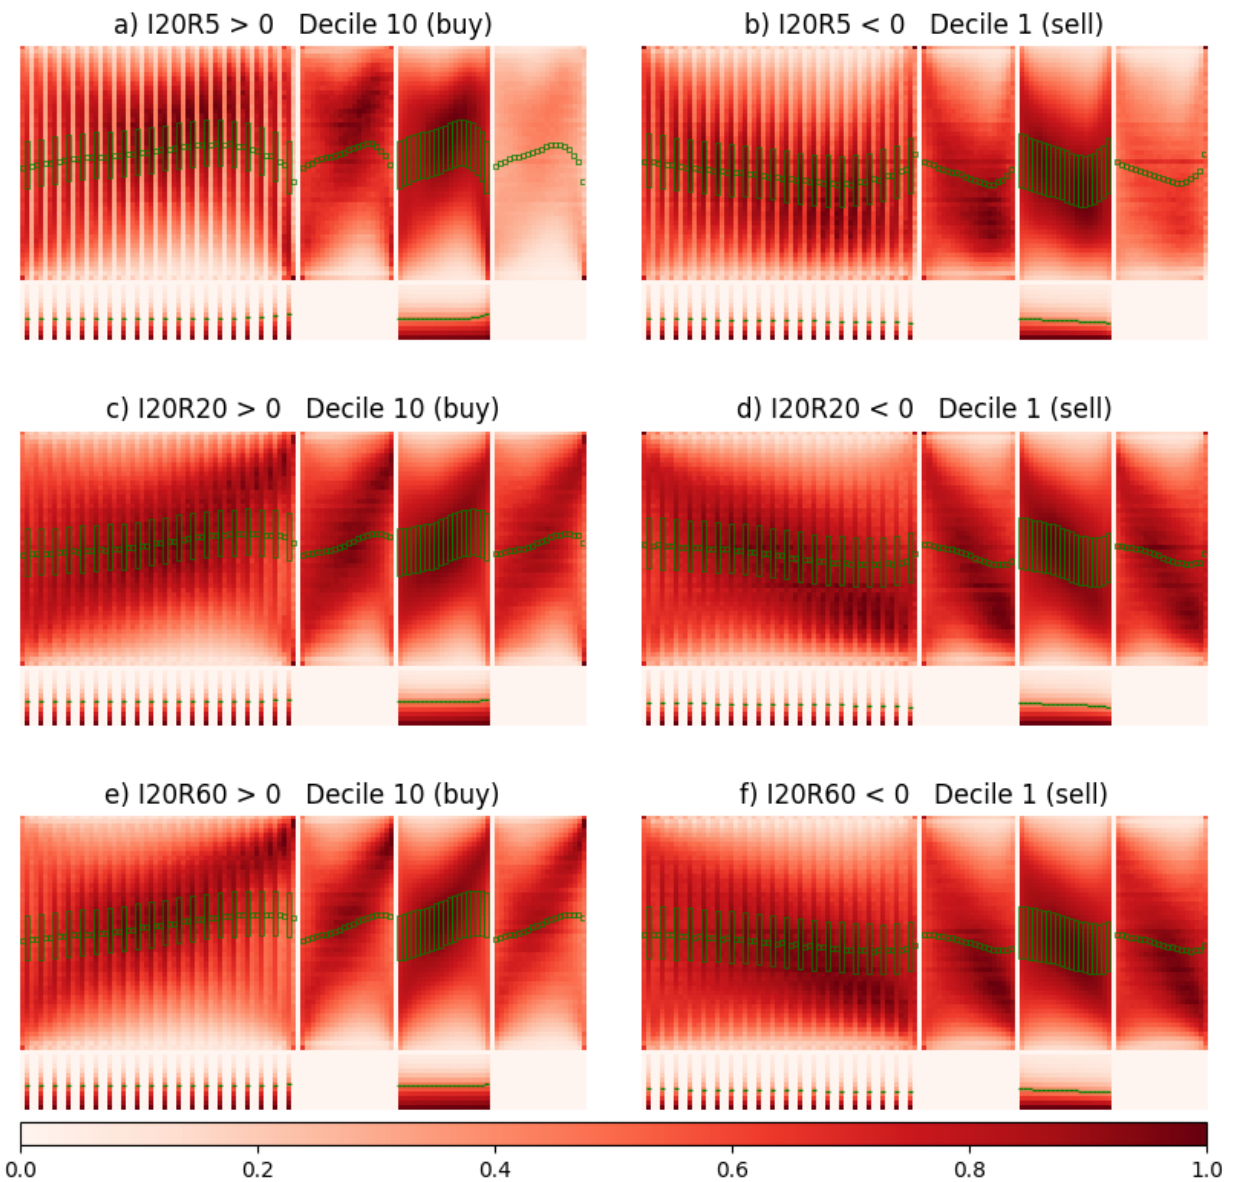


Figure 9: Visualized Weekly, Monthly, and Quarterly 20-day Trend Reversal Portfolios

This figure, in columns one, two, and three, shows the aggregate 20-day pattern of weekly, monthly, and quarterly long-short *reversal* portfolio formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return, in the sample with negative (positive) returns. Aggregation is equal weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.

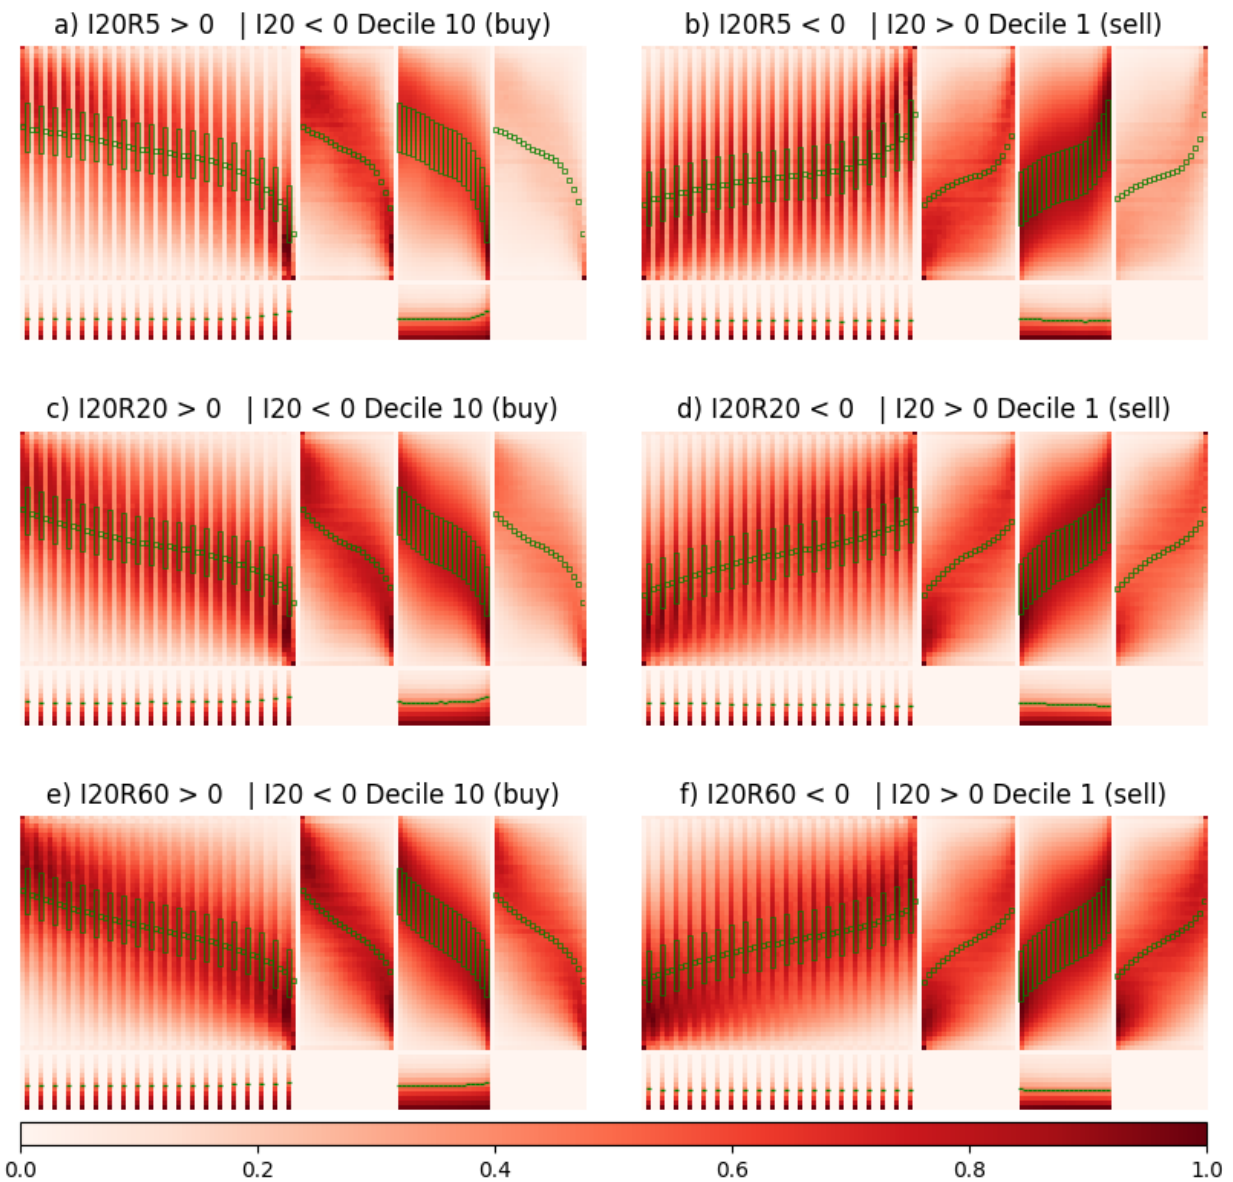


Figure 10: Visualized Weekly, Monthly, and Quarterly 20-day Trend Persistence Portfolios

This figure, in columns one, two, and three, shows the aggregate 20-day pattern of weekly, monthly, and quarterly long-short **persistence** portfolio formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return, in the sample with positive (negative) returns. Aggregation is equal weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.




Figure 11: Visualized Full Sample Weekly, Monthly, and Quarterly 5-day Image Portfolios (Value Weighted)

This figure, in columns one, two, and three, shows the aggregate 5-day pattern of weekly, monthly, and quarterly long-short portfolios formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return. Aggregation is value weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.

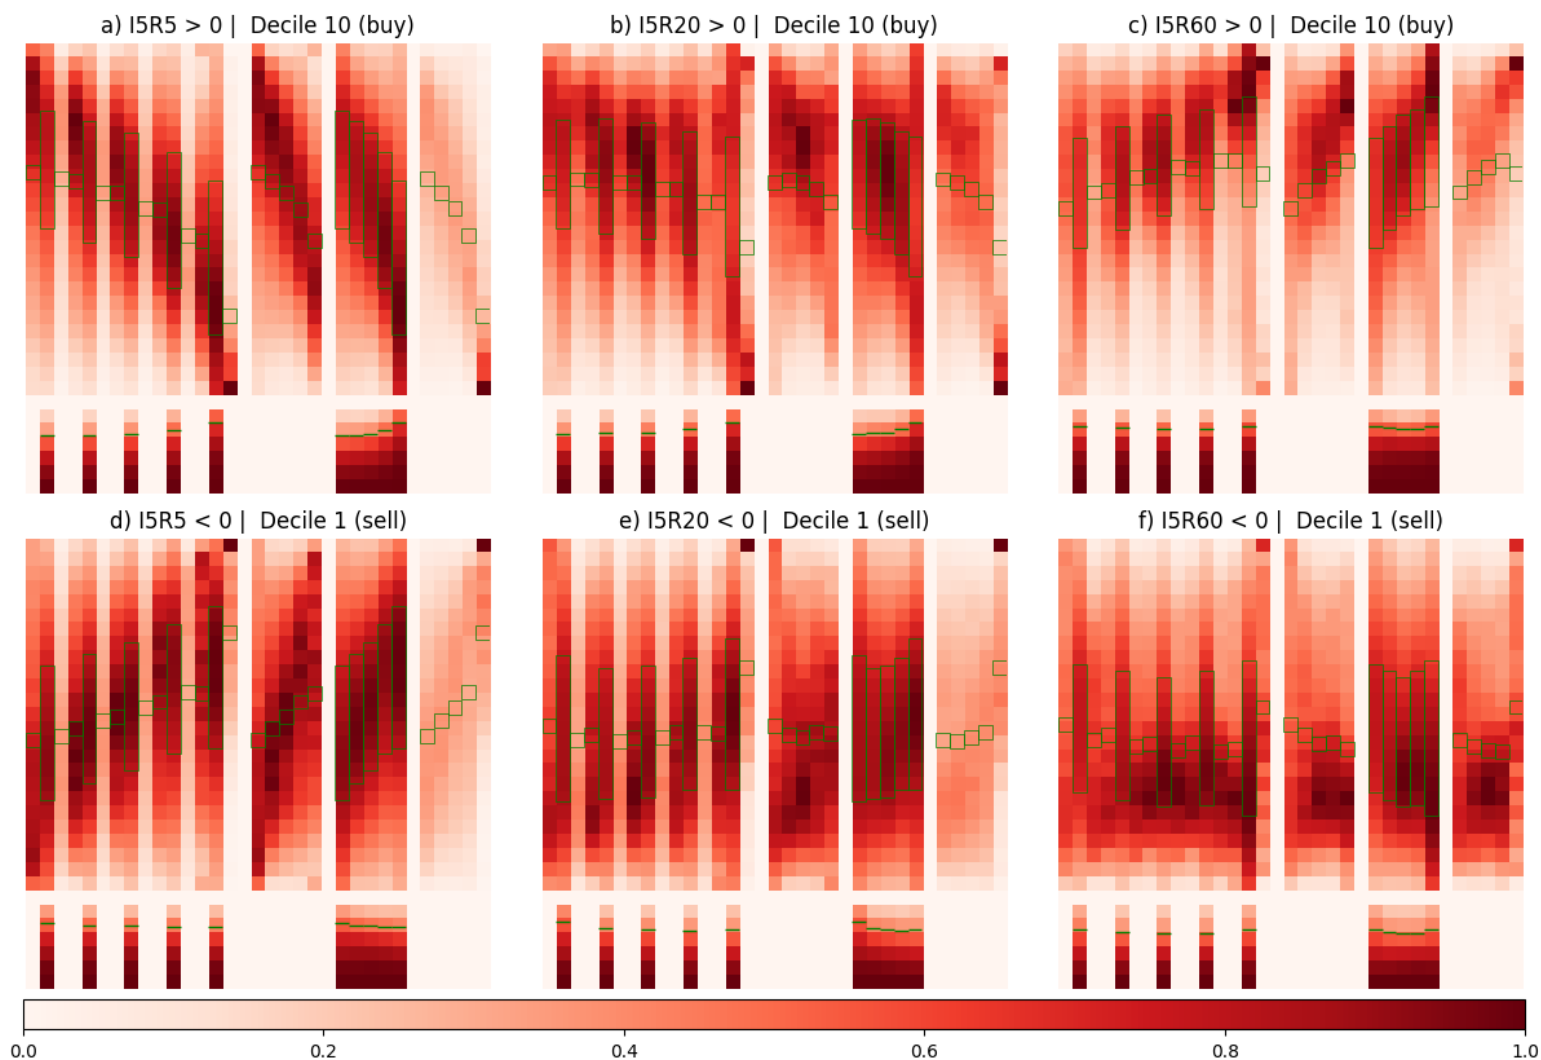


Figure 12: Visualized Weekly, Monthly, and Quarterly 5-day Trend Reversal Portfolios (Value Weighted)

This figure, in columns one, two, and three, shows the aggregate 5-day pattern of weekly, monthly, and quarterly long-short *reversal* portfolio formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return, in the sample with negative (positive) returns. Aggregation is value weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.

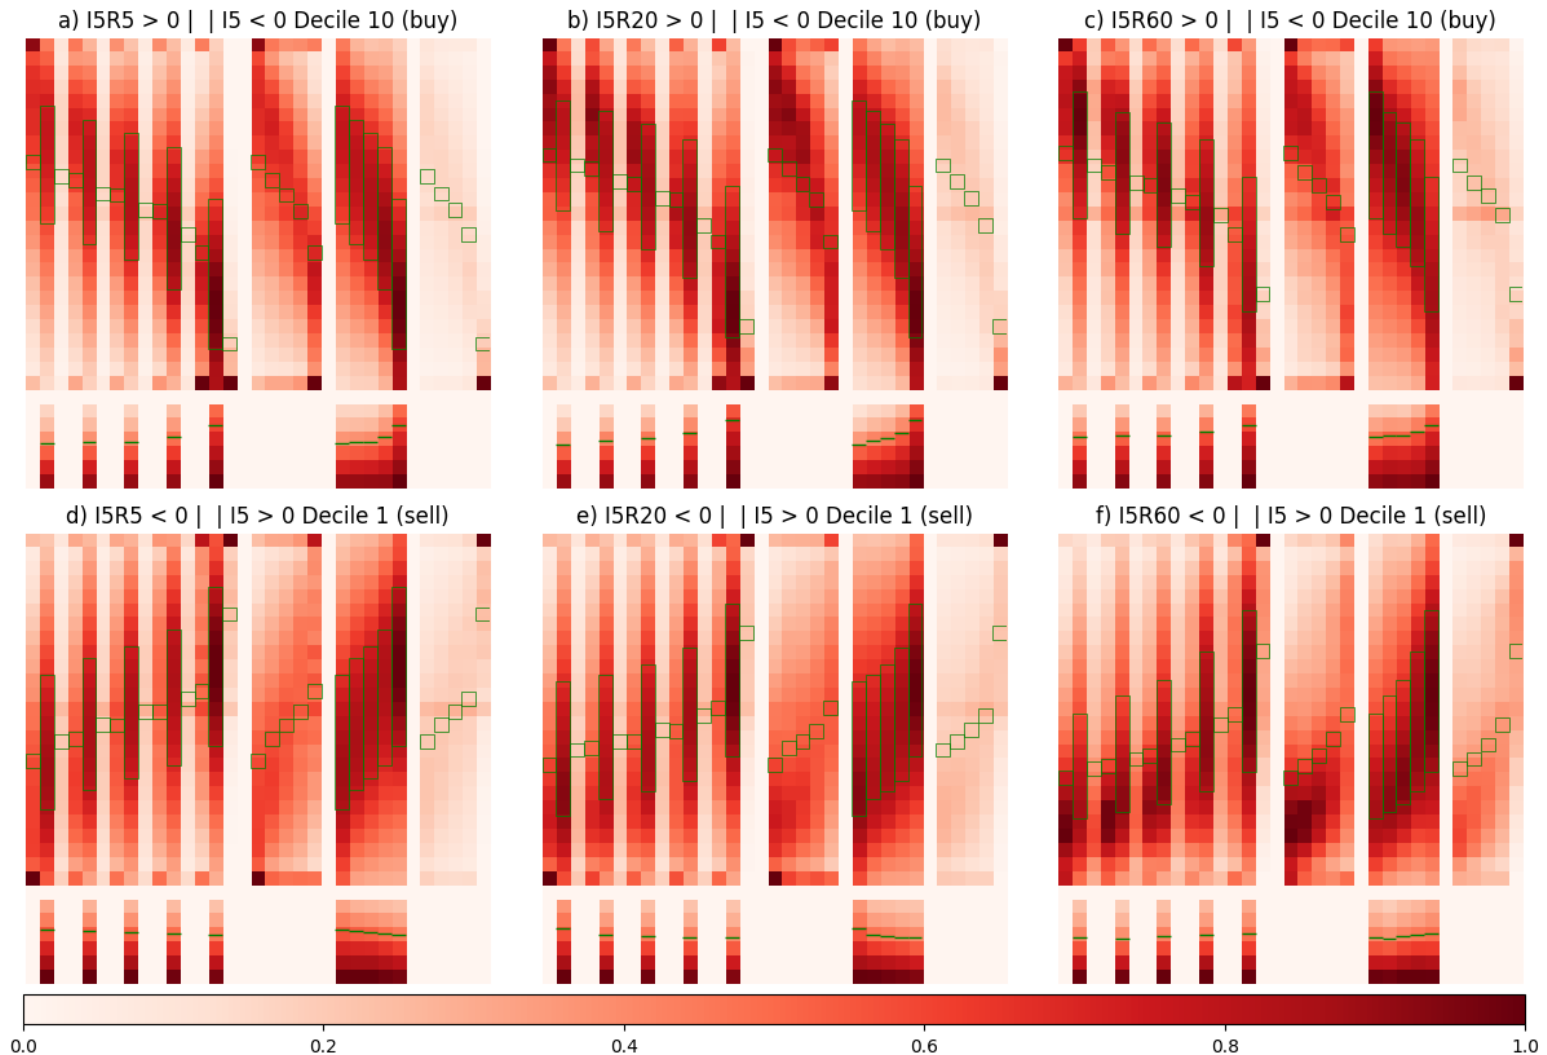


Figure 13: Visualized Weekly, Monthly, and Quarterly 5-day Trend Persistence Portfolios (Value Weighted)

This figure, in columns one, two, and three, shows the aggregate 5-day pattern of weekly, monthly, and quarterly long-short **persistence** portfolio formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return, in the sample with positive (negative) returns. Aggregation is value weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.

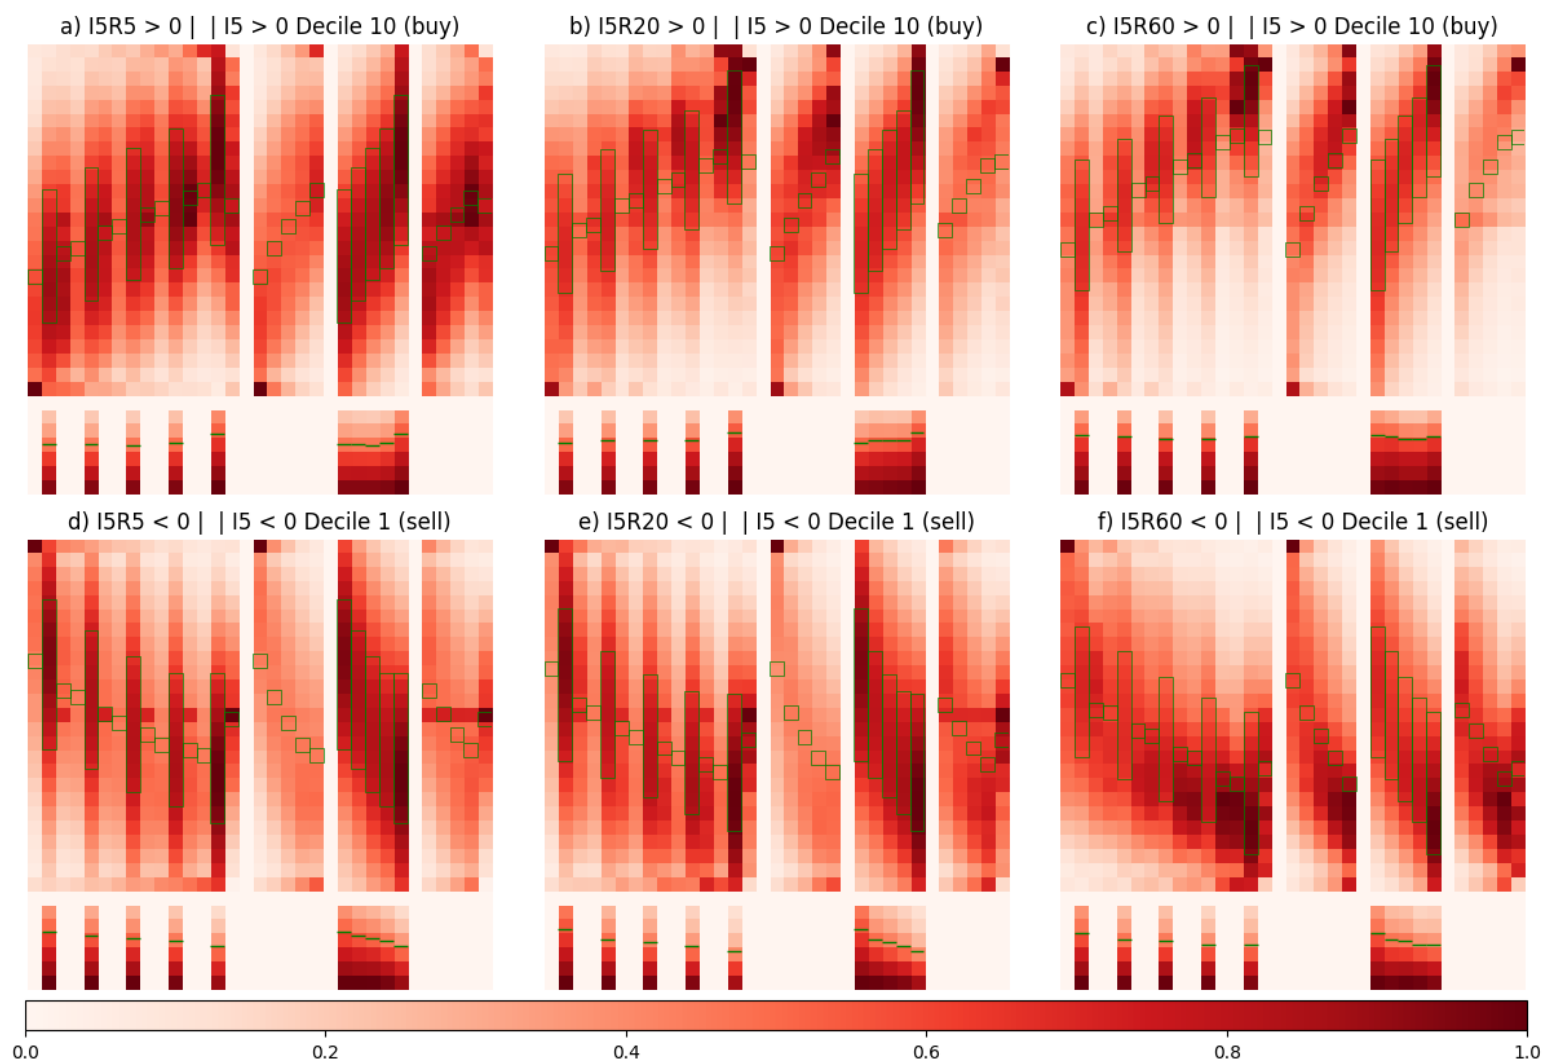


Figure 14: Visualized Full Sample Weekly, Monthly, and Quarterly 20-day Image Portfolios (Value Weighted)

This figure, in columns one, two, and three, shows the aggregate 20-day pattern of weekly, monthly, and quarterly long-short portfolio formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return. Aggregation is value weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.

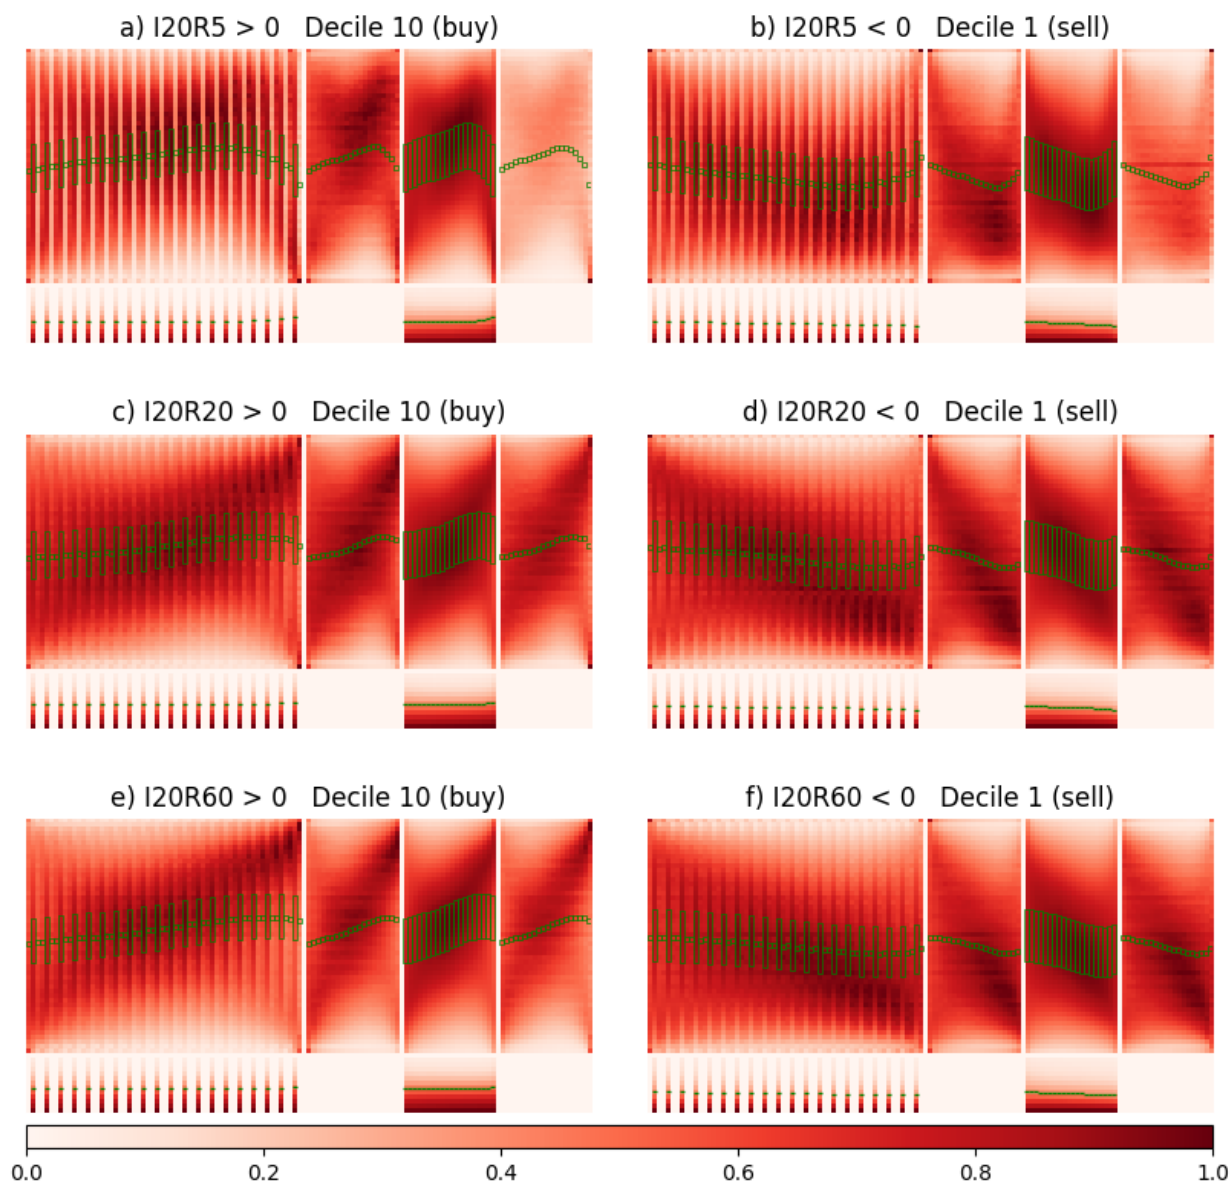


Figure 15: Visualized Weekly, Monthly, and Quarterly 20-day Trend Reversal Portfolios (Value Weighted)

This figure, in columns one, two, and three, shows the aggregate 20-day pattern of weekly, monthly, and quarterly long-short *reversal* portfolio formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return, in the sample with negative (positive) returns. Aggregation is value weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.




Figure 16: Visualized Weekly, Monthly, and Quarterly 20-day Trend Persistency Portfolios (Value Weighted)

This figure, in columns one, two, and three, shows the aggregate 20-day pattern of weekly, monthly, and quarterly long-short **persistence** portfolio formed, respectively, in the full sample from 2001 to 2019. The top (bottom) row shows the long (short) end with the highest (lowest) CNN predicted probability of positive return, in the sample with positive (negative) returns. Aggregation is value weighted and normalization takes place component wise (i.e., all close price pixel counts are divided by the largest close price pixel count in the image). The green outline shows the average level within a column. In each panel, we show the raw aggregate plot, then the segmented plots in order of open, high-low, and close components. A segmented plot of open price, for example, is all open columns together.

