


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

Stephen Dixon and Qi Zeng

University of Melbourne

June 2026

Jiang, Kelly and Xiu (2023) form high-performing portfolios with CNN predicted probability of positive future returns, using open, high, low, close charts with volume and moving average line.



However, **what does the machine see in a stock chart?** We

1. propose and apply a method to *visualize learnt features*; and
2. show that CNN learns a *novel form of reversal*; and
3. *quantify* such pattern's impact on CNN predictions.

These features can explain a large proportion of variation in both CNN predictions, as well as CNN portfolio returns.

Our paper contributes to three branches of literature:

1. Rapidly growing application of machine learning models to asset pricing (Chen et al., 2024; Gu et al., 2020; Jiang et al., 2023; Murray et al., 2024).
2. Traditional reversal and momentum portfolios (Cooper, 1999; Lehmann, 1990; Moskowitz et al., 2012).
3. The interpretation of such machine learning models (Bell et al., 2026).

Using AMEX, NASDAQ, and NYSE stocks:

1. Replication: Trained CNN over 1993 to 2000, then applied in 2001-2019 for weekly, monthly, and quarterly portfolios.
2. For long/short end of portfolio, **sum up matrix representation of images** without the moving average line. In some sense, we are calculating the distribution of pixels:

$$\begin{array}{|c|} \hline 0 \\ \hline 0 \\ \hline 0 \\ \hline 1 \\ \hline \end{array} + \begin{array}{|c|} \hline 0 \\ \hline 0 \\ \hline 0 \\ \hline 1 \\ \hline \end{array} + \begin{array}{|c|} \hline 0 \\ \hline 0 \\ \hline 1 \\ \hline 0 \\ \hline \end{array} + \begin{array}{|c|} \hline 0 \\ \hline 0 \\ \hline 0 \\ \hline 1 \\ \hline \end{array} = \begin{array}{|c|} \hline 0 \\ \hline 0 \\ \hline 1 \\ \hline 3 \\ \hline \end{array} \quad (1)$$

3. Regress CNN predictions against observed trends
 4. Compare CNN portfolio returns to that of the fitted models.
- Some discussion on economic linkage

JKX Replication - EW Decile Portfolios

$P(> 0)$ - CNN probability of positive next week/month return
 $I5$ - Return over the 5-day formation window


	5-day CNN				Weekly Return			
	Ret	SR	$P(> 0)$	$I5$	Ret	SR	$P(> 0)$	$I5$
Weekly								
Short	-0.32	-2.16	0.29	0.03	-0.15	-0.63	0.43	0.13
Long	0.48	2.47	0.64	-0.05	0.44	1.47	0.53	-0.11
L - S	0.80	7.18			0.59	3.15		
Monthly								
Short	-0.00	-0.03	0.39	0.02	0.02	0.09	0.48	0.13
Long	0.21	0.98	0.60	-0.04	0.21	0.65	0.53	-0.11
L-S	0.21	2.30			0.19	1.10		

bold indicates significance at 1% level.

What does the machine see?

CNN predicts that a stock has the **highest** (**lowest**) prob. of realising positive future returns when:

- ▶ Last day close is extremely **low** (**high**)
- ▶ Volume **rises** (**falls**)
- ▶ Difference between close and high-low average is **negative** (**positive**) compared to previous days



Price variables are the pixel positions! Given a vector $price^{raw}$ containing all open, high, low, and close prices:

$$price^{scaled} = \frac{price^{raw} - \min(price^{raw})}{\max(price^{raw}) - \min(price^{raw})} \quad (2)$$

For example, a close price at the lowest (highest) position has the value of 0 (1). Using the scaled prices, we measure the following:

$$RelativeClose_{it} = (C_{it5} - \frac{H_5 + L_5}{2}) - \frac{1}{4} \sum_{d=1}^4 (C_d - \frac{H_d + L_d}{2}) \quad (3)$$

Approximating 5-Day CNN Predictions

Relative close absorbs a large proportion of last close's coefficient, and produces a nice Sharpe ratio.

	$1\{Pr(R5 > 0) > 0.5\}$					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Open</i> _{5th day}	-0.14***				-0.13***	-0.13***
<i>Close</i> _{5th day}		-0.23***			-0.10***	-0.06***
<i>Volume</i> _{5th day}			0.09***		0.08***	0.08***
Relative Close				-0.15***	-0.15***	-0.16***
Formation Returns						-0.06***
No. Obs (Million)	29	30	30	30	29	29
Pseudo R-squared	0.09	0.29	0.03	0.09	0.38	0.39
OLS Adj. R-squared	0.10	0.37	0.05	0.19	0.54	0.54
	Applying the fitted logit					
Ret				0.47	0.59	0.71
SR				4.71	4.89	4.52
R-squared				0.37	0.80	0.80

Top panel reports the standard deviation times average marginal effect from a logistic regression. Bottom panel reports portfolio performance using the fitted model to make predictions, as well as the R-squared from regressing CNN returns against the model returns.

Volume and relative close matter more in longer horizon predictions.

	$1\{Pr(R20 > 0) > 0.5\}$					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Open</i> _{5th day}	-0.02***				-0.09***	-0.09***
<i>Close</i> _{5th day}		-0.15***			-0.02***	-0.01***
<i>Volume</i> _{5th day}			0.15***		0.15***	0.15***
Relative Close				-0.20***	-0.23***	-0.23***
Formation Returns						-0.03***
No. Obs (Million)	29	30	30	30	29	29
Pseudo R-squared	0.00	0.07	0.08	0.13	0.26	0.26
OLS Adj. R-squared	0.01	0.18	0.13	0.26	0.48	0.48
	Applying the fitted logit					
Ret				0.13	0.16	0.18
SR				1.70	1.88	1.92
R-squared				0.28	0.66	0.70

Top panel reports the standard deviation times average marginal effect from a logistic regression. Bottom panel reports portfolio performance using the fitted model to make predictions, as well as the R-squared from regressing CNN returns against the model returns.

Are there any economic or technical linkages?

Cooper (1999) attributes weekly short term reversal to investor overreaction, showing that the effect is strongest with reduction in volume. We document a **scaled price trend** reversal with important differences:

1. It is **pattern based**, scaled prices means the exact return cannot be inferred.
2. Rising volume trend facilitates upward reversal.
3. The position of close relative to midpoint of high-low is important. Weekly return does not consider even the unscaled high-low prices.

Our results suggest that overreaction, as learnt by CNN, is more likely when the price trend involves a sudden drop in the last day closing price.

So far, I have presented results for portfolios that rebalance weekly and monthly.

1. Similar short-horizon results when the model is trained on 20 or 60 day stock charts.
2. When the prediction horizon increases to 60 days, returns become driven by a time-series momentum pattern instead of the reversal pattern at shorter horizons.

What if CNN learns multiple patterns? Vectorize images and apply k-clustering to identify groups of patterns. Our main reversal results remain consistent.

Machine learning performs better, but **how**? We answer this question in the specific context of predicting equity returns with CNN by:

1. proposing an aggregate heat map methodology to visualize features learnt by CNN.
2. showing that a large proportion of variation in CNN prediction and returns are driven by novel reversal patterns.

Thanks for listening!