

Combining news sentiment and technical analysis to predict stock trend reversal

Original

Combining news sentiment and technical analysis to predict stock trend reversal / Cagliero, Luca; Attanasio, Giuseppe; Garza, Paolo; Baralis, ELENA MARIA. - ELETTRONICO. - (2019), pp. 514-521. (Intervento presentato al convegno 9th ICDM Workshop on Sentiment Elicitation from Natural Text for Information Retrieval and Extraction tenutosi a Beijing, Cina nel November 8, 2019) [10.1109/ICDMW.2019.00079].

Availability:

This version is available at: 11583/2753472 since: 2020-01-22T10:42:46Z

Publisher:

IEEE Computer Society Press

Published

DOI:10.1109/ICDMW.2019.00079

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Combining news sentiment and technical analysis to predict stock trend reversal

Luca Cagliero, Giuseppe Attanasio, Paolo Garza, Elena Baralis
Dipartimento di Automatica e Informatica
Politecnico di Torino, Turin, Italy
Email: name.surname@polito.it

Abstract—The use of machine learning techniques to predict the next-day stock direction is established. To make prediction models more robust, a common approach is to combine historical time series and news sentiment analysis. Most of the trading simulations performed in this field rely on trend following strategies, which are aimed at identifying and following an ongoing price trend that is likely to persist in the next days. Conversely, a more limited effort has been devoted to applying machine learning techniques to predict trend reversal, i.e., changes in price directions. This paper investigates the relevance of news information and time series descriptors derived from technical analysis to predict trend reversal in the next days. It compares the performance of various classification models trained on (i) technical indicators, which indicate short-term overbought or oversold conditions, (ii) news sentiment descriptors, which express the opinion of the financial community, (iii) the historical time series, to highlight recurrences in price trends, and (iv) a combination of the above. The results achieved on an 11-year dataset related to the stocks of the U.S. S&P 500 index show that the strategies combining the historical values of news sentiment and stock price indicators averagely perform better than all the other tested combinations. Hence, news information is worth considering by trend reversal strategies.

Index Terms—news sentiment analysis, quantitative trading, classification, trend reversal prediction

I. INTRODUCTION

Stock price forecasting is among the most popular financial problems. In the last decade, the diffusion of online trading systems together with the increasing availability of stock-related data have prompted the application of machine learning techniques to forecast the stock markets [1]. Quantitative trading consists in applying a data-driven approach to stock trading. Machine learning algorithms have been tightly integrated into quantitative trading systems. Supervised models are first trained on historical data and then applied in order to forecast the future stock prices [2].

Machine learning models can be trained on stock-related data with different characteristics. The analyzed features include (i) The historical time series of stock prices and exchanged volumes as well as the indicators/oscillators derived from technical analysis (e.g., moving averages, MACD, RSI [3]). (ii) The main accounting measures used in fundamental analysis to estimate the value of financial assets (e.g., EBITDA, gross profit, operational margin [4]). (iii) The opinion of analysts, investors, and media experts given through news, social networks, and online communities [5], [6]. This paper addresses stock price forecasting based on

machine learning models trained on features (i) and (iii), i.e., fundamental analysis is out of the scope of this work.

Classification is one of the most commonly used machine learning technique to forecast the stock markets [7]. It entails training predictive models to forecast the future stock direction (e.g., up, down, stable) in the near future (e.g., in the next day, one week ahead). Based on the outcomes of per-stock prediction models, the trading systems can automate the process of trading signal generation (i.e., open a short-selling position on a given stock at the opening of the trading day because its price is likely to decrease).

Trading systems may rely on different strategies. The most established approaches are (i) *trend following*, which entails identifying significant uptrends/downtrends and betting on their continuation, or (ii) *trend reversal*, which entails identifying the end of a trend and betting on a trend reversal (e.g., if a price uptrend is ending, the system opens a short-selling position betting on a forthcoming price decrease). Trend reversal approaches are commonly based on the analysis of historical time series combined with the use of technical oscillators (e.g., RSI, momentum) (e.g., [8], [9]). Conversely, the exploitation of news information to drive trend reversal strategies is still an open research issue.

This paper proposes a trend reversal approach to quantitative stock trading that combines news information with historical prices and technical indicators. The proposed approach relies on per-stock classification models, whose goal is to predict the stock direction in the very next days (i.e., from one to five days ahead). Classification models are trained on per-stock training datasets collecting stock-related data aggregated at a daily granularity. The trading system monitors the daily stock price variation to identify ongoing trends in stock price variations (i.e., sequences of W consecutive relevant increases/decreases of the daily prices). Whenever a significant trend is detected, the classification models are used to predict whether a reversal is likely to happen in the next days (e.g., a significant decrease in the stock price is going to occur in the next day). Based on the features considered during classifier training, the decisions of the trading system can be influenced by (i) the historical stock prices, (ii) the technical oscillators computed on the stock prices and exchanged volumes, (iii) the news sentiment indicators, or (iv) a combination of the above. The aim of this paper is to understand whether news sentiment indicators are worth considering during trend reversal strategies, eventually

in combination with other features. To this aim, it explores the performance of various systems relying on different combinations of features (including news sentiment and not) on 11-year stock data of the Standard & Poor index (i.e., the main American Stock Exchange index).

The research questions addressed in the paper can be summarized as follows:

- (RQ1) Is news information worth considering in classification-based trend reversal approaches?
- (RQ2) What is the impact of including news information on classifier performance?
- (RQ3) Which features should be combined with news sentiment indicators to maximize the average return of trend reversal strategies?

The results achieved on S&P data show that the strategies combining news sentiment and stock prices analysis averagely perform best. Hence, news sentiment analysis is worth considering in order to design trend reversal strategies. The most promising results were achieved by focusing on trends consisting of sequences of four significant and concordant price variations (decrease/increase). The classification models were able to produce reliable trend reversal predictions.

The rest of the paper is organized as follows. Section II overviews the related literature. Section II formulates the problem addressed in the paper. Section IV thoroughly describes the characteristics of the proposed trading system. Section V presents the results of the empirical evaluations, while Section VI draws conclusions and discusses the future extensions of this work.

II. LITERATURE REVIEW

Exploiting news information to improve the performance of stock market prediction methods is an established approach to quantitative stock trading [1]. The preliminary studies on the effects of news reports on stock prices have been carried on in [10] using linear models. The application of more advanced text mining algorithms and analytics have also been tested. For example, to investigate the impact of news content on stock market prediction in [11] the authors first generated a training news dataset by labeling news content according to the stock price variations from 5 to 20 minutes after news release. Next, they trained classification models to predict the future stock direction. Similar to [11], in [12] the authors explored the risk implications of information being newly available to market participants. To this end, they applied different classification techniques in order to detect patterns hidden in textual data that could explain increased risk exposure. Similar analyses have been carried out on microblogs [13]–[15] and financial reports [16]. Recently, in [5] the authors studied how to incorporate social and news opinion and sentiment to forecast the daily stock prices in the next 10 days. To tackle the above issue, they trained a Neural Network regression model on a dataset including historical closing prices and a couple of news sentiment features extracted using the NLTK toolkit (<https://www.nltk.org/>). Similar to [5], this paper combines news sentiment and time series descriptors

derived from technical analysis with the goal of predicting the stock prices in the next days. Unlike [5], it specifically investigates the usefulness of incorporating news information in a trend reversal trading strategy. To tackle the aforesaid issue, it explores various combinations of stock-related features including technical oscillators and volatility indices. The study recently presented in [17] performed a correlation analysis between the news content and the stock prices in a time span ranging from few days before the current time to 30 days ahead. It applies a Long Short-Term Memory (LSTM) recurrent neural network to predict the future direction of the stock prices. The results show that the information contained in the news articles is strongly correlated with the past stock price movements, whereas the strength of the correlation rapidly decreases while moving forward in time. The contribution of this paper can be deemed as complementary to [17], because it deepens the analyses performed by [17] in a trend reversal scenario using various classification methods and prediction horizons.

Quantitative trading systems based on trend reversal strategies have already been proposed in literature. For example, in [9] the authors presented a framework that combines Support Vector Machines with K-Nearest Neighbor classifiers. It has been applied to forecast the values of the Indian stock market indices. The proposed data analytics framework relies on various technical features, but it did not consider news sentiment. Alternative strategies to detect reversal patterns are based on (i) candlestick pattern recognition [18], [19] (ii) the analysis of technical oscillators such as momentum and RSI [20], and (iii) the volumes of targeted queries executed on the Google search engine [21]. To the best of our knowledge, this is the first attempt to investigate the applicability of news information in trend reversal approaches.

III. PROBLEM STATEMENT

Let \mathbf{S} be a subset of stocks under consideration. Hereafter, we will consider the stocks of the U.S. Standard & Poor 500 index. Each stock in \mathbf{S} is described by a set of discrete time series, where each time series consists of the values of a stock-related descriptor sampled on a daily basis.

Let $TS^D(s_j)$ be the time series corresponding to descriptor D and stock $s_j \in \mathbf{S}$. The stock descriptors belong to three main categories: (i) historical stock prices (*PRICE*, in short), (ii) technical oscillators and volatility indices (*OSC*), (iii) news sentiment (*NEWS*). For example, descriptor *cp* belongs to category *PRICE* and represents the daily closing price of a stock. The full list of the descriptors used in the experiments is reported in Section IV-A.

This paper addresses the prediction of the future *direction* of a given stock to support trend reversal strategies. The direction of a stock is expressed in terms of the percentage variation of the daily closing prices. More specifically, the direction $dir(s_j, d_t)$ of a stock $s_j \in \mathbf{S}$ on an arbitrary day d_t is defined as follows.

$$dir(s_j, d_t) = \begin{cases} up & \text{if } \frac{cp_t - cp_{t-1}}{cp_{t-1}} \cdot 100 \geq 1, \\ down & \text{if } \frac{cp_t - cp_{t-1}}{cp_{t-1}} \cdot 100 \leq -1, \\ stable & \text{otherwise} \end{cases}$$

A positive variation above 1% is denoted as *up* and indicates a significant price increase from day d_{t-1} to d_t . Oppositely, a negative variation below -1% indicates a significant price decrease (*down*). Finally, *stable* denotes a stationary situation.

Predicting on an arbitrary day d_t the future direction of stock s_j on day d_{t+1} entails inferring the function f

$$dir(s_j, d_{t+1}) = \mathbf{f}(dir(s_j, d_t), dir(s_j, d_{t-1}), \dots, dir(s_j, d_{t-x+1}))$$

where $dir(s_j, d_{t+1})$ is the value of the target variable, $dir(s_j, d_t), \dots, dir(s_j, d_{t-x+1})$ are the current and past values of stock direction in a predefined time window $[d_{t-x+1}, d_t]$, and $f(\cdot)$ is the prediction function that needs to be inferred. Notice that in the simplest case of *univariate analysis*, which is considered in the above formula, the future direction of a stock is assumed to exclusively depend on the current and past stock directions.

A more realistic model incorporates the cross-dependencies among the time series associated with multiple descriptors, i.e., the *multivariate analysis*. Let $descr_{t-x+1}^*, \dots, descr_t^*$ be the series of past and current values of an arbitrary stock descriptor. The problem of predicting the future direction of stock s_j on an arbitrary day d_{t+1} in the multivariate scenario entails inferring the function \mathbf{f}^*

$$dir(s_j, d_{t+1}) = \mathbf{f}^*(dir(s_j, d_t), \dots, dir(s_j, d_{t-x+1}), \\ descr^*(s_j, d_t), \dots, descr^*(s_j, d_{t-x+1}), \\ \dots, \\ descr^{**}(s_j, d_t), \dots, descr^{**}(s_j, d_{t-x+1}))$$

To design a trend reversal strategy, the concepts of *uptrend* and *downtrend* for a given stock are introduced. Specifically, an *uptrend* is a sequence of W consecutive *up* values in the series of stock directions (where W is a user-specified parameter). It indicates that the closing prices have significantly increased for W consecutive trading days. Analogously, a *downtrend* is a sequence of W consecutive *down* values.

Trend reversal trading systems aim at predicting when a trend is close to its end. Detecting excesses in price trends allows traders to bet on a forthcoming change of the stock price direction [22]. In our context, a trend reversal occurs when any of the following conditions are satisfied: (i) A *down* occurs in the series of stock directions immediately after or close to an *uptrend*. (ii) An *up* occurs immediately after or close to a *downtrend*.

Given an arbitrary stock s_j , a number N of consecutive days in the near future, the purpose of the prediction task addressed in this paper is to predict whether a trend reversal will occur in any of the N days following the detection of

an *uptrend/downtrend*. Notice that predicting a reversal after an *uptrend* entails predicting direction *down* in any of the days between d_{t+1} and d_{t+N} , while predicting a reversal after a *downtrend* entails predicting direction *up* on any of the aforesaid days.

IV. CLASSIFICATION-BASED TREND REVERSAL STRATEGY

This section presents the classification-based approach to quantitative stock trading. The main steps can be summarized as follows:

- 1) *Data retrieval and modelling*. Stock-related data are collected and organized into per-stock training datasets consisting of the historical values of all the considered descriptors aggregated on a daily basis (see Section IV-A).
- 2) *Trend detection*. The daily closing price variations of each stock are monitored to detect uptrends/downtrends of a user-specified duration (see Section IV-B).
- 3) *Classification model learning*. For each trend detected at the previous step, ad hoc classifiers are trained on the past and current per-stock data in order to predict the closing price variations in the following five days. The classifier predictions are used in the current day to predict whether a trend reversal will occur (see Section IV-C).
- 4) *Trade management*. Classifier outcomes are used to generate per-stock trading signals. Specifically, for each trend associated with a predicted reversal, an appropriate trading position is opened. The money management is based on a weekly strategy (see Section IV-D).

A. Data retrieval and modelling

This step entails collecting and preparing stock-related data for the trend detection and classification processes. For each stock it collects the historical values of all the considered descriptors. The three main categories of descriptors considered in this study are enumerated below.

Historical prices. This category, hereafter denoted as *PRICE*, includes all the series of historical prices acquired at a daily frequency (e.g., opening price, closing price, maximum price, minimum price). These values are commonly used in technical analysis to detect reversal patterns from candlestick charts [3]. Since we are interested in predicting the future stock directions, for each descriptor included in this category we consider the daily percentage variation with respect to the previous day instead of its absolute value.

Technical oscillators and volatility indices. This category, hereafter denoted as *OSC*, includes the series of historical values of all the main technical oscillators and volatility indices, i.e., the Moving Average Convergence/Divergence, the Aroon Oscillator (computed over 14 periods), the Average Directional Index (14 periods), the Difference between Positive Directional Index (DI+) and Negative Directional Index (DI-) (computed over 14 periods), the Percentage Price Oscillator (12 and 26 periods), the Relative Strength Index (14 periods), the Money Flow Index (14 periods), the True Strength Index,

the Stochastic Oscillator (14 periods), the Chande Momentum Oscillator (14 periods), the Average True Range Percent (14 periods), the Percentage Volume Oscillator (14 periods) [3].

Oscillators are largely used technical indicators whose values vary over time within a specified range (typically between 0 and 100). They are defined on either stock prices or volumes, or a combination of the above. They indicate the presence of overbought or oversold conditions. They can be used to recognize market excesses thus planning trend reversal strategies accordingly. Volatility indices are established technical indicators that indicate how quickly stock prices change. They are deemed as useful for detecting trend reversal because they estimate the level of bullish or bearish sentiment (i.e., the confidence of keeping a long-/short-selling positions open).

News sentiment. This feature category, hereafter denoted as NEWS, groups all the stock descriptors related to the sentiment or the opinion of the financial community. In this study, a list of sentiment descriptors is extracted from the news articles ranging over the considered stocks. Descriptor values are based on the occurrences of the words in a domain-specific sentiment dictionary (<https://sraf.nd.edu/textual-analysis/>). The full list of considered descriptors is given below.

- Total number of pertinent news articles released on the considered day.
- Total number of pertinent news articles released on the considered day and containing any word with a sentiment (either *positive* or *negative*).
- Total number of pertinent news articles released on the considered day that contain at least a word with *positive* sentiment (the complementary version of the descriptor for *negative* sentiment is available as well).
- Total number of pertinent news articles released on the considered day that contain any words with *positive* sentiment and do not contain *negative* words (the complementary version is also available).
- Total length (expressed as the number of contained words) of the pertinent news articles released on the considered day.
- Total number of words with *positive* sentiment in any news article released on the considered day (the complementary version is also available).
- Total number of words with *positive* sentiment in any news article released on the considered day and with strong correlation with the next-day stock direction in the past data (the complementary version is also available).

Notice that the descriptors' list and the corresponding categorization can be modified or extended according to domain experts' preferences.

The historical values of the selected descriptors are used to populate the per-stock training datasets. A per-stock training dataset is a relational table [23] consisting of a distinct attribute for each selected descriptor and a record per day in the recorded history. For each attribute a record takes the value of the corresponding descriptor computed on the corresponding day. For example, in the record corresponding to day d_t of the training dataset related to stock s_j feature *closing price*




Fig. 1. Trend detection and reversal prediction.

takes the value $cp(s_j, d_t)$. To identify the temporal patterns hidden in the analyzed data, each record contains not only the current descriptor values on the corresponding day, but also the past values within a sliding window of size S . To this end, for each descriptor the past values taken in the sliding window are stored in separate attributes. For example, the record corresponding to day d_t contains not only the closing price of the stock on day d_t , but also the history of the closing prices in the sliding window (stored in separate attributes), i.e., $cp(s_j, d_{t-1}), cp(s_j, d_{t-2}), \dots, cp(s_j, d_{t-S+1})$.

Each record of the training dataset has a label. The value of the label changes according to the target of the classification task. Specifically, to predict the next-day direction of stock s_j , an arbitrary record of the training dataset corresponding to day d_t is labeled as $dir(s_j, d_{t+1})$, i.e., the direction of stock s_j the days after. Analogously, to predict the direction of stock s_j N day after, the same record is labeled as $dir(s_j, d_{t+N})$.

B. Trend detection

The aim of this step is to automatically detect uptrends/downtrends in the series of the closing stock prices (see Section III). For each stock in $s_j \in \mathbf{S}$, the series of the daily price directions $dir(s_j, d_{t-x}), dir(s_j, d_{t-x+1}), \dots, dir(s_j, d_t)$ is monitored. Whenever an uptrend/downtrend is detected for stock s_j , the classification process, tailored to stock s_j , is triggered.

The sketch depicted in Figure 1 exemplifies how trend detection triggers the next classification step. Specifically, the dotted curve in the upper part of the figure shows the series of the past/future daily closing prices of the considered stock. The corresponding directions are indicated above the curve (u=up, d=down, s=stable). A dashed circle surrounding the direction indicates that it is a future stock directions (i.e., a target of the predictions), while non-dashed circles indicate known values (i.e., relative to the past or current days). At the closure of day d_t an uptrend (consisting of $W=3$ up values) is detected. This event triggers the generation of five classification models, each one predicting the future direction of the stock on a distinct day in the future, between d_{t+1} and d_{t+5} .

C. Classification model learning

This step focuses on predicting whether a trend reversal occurs in the N days following an uptrend/downtrend. To this end, per-stock classification models are separately trained to predict the stock directions in different days in the future.

Recalling the example in Figure 1, the classifier predictions generated at the end of day d_t (up, down, or stable) are depicted as squares in the bottom part of the figure. At the end of day d_t , the classifier predicts *down* for the next-day closing price (d_{t+1}), *stable* for d_{t+2} and d_{t+3} , *up* for d_{t+4} , and *down* for d_{t+5} .

The directions predicted in the next days are exploited to forecast a forthcoming trend reversal. Specifically, on an arbitrary future day d_{t+y} [$1 \leq y \leq 5$] if the classifier prediction is *down* whereas the ongoing trend is an *uptrend* then a trend reversal is detected. Vice versa, if the classifier prediction on the same day is either *up* or *stable* then an uptrend continuation is likely. Formally speaking, the inference of trend reversal on the future days will be tackled by defining the Boolean function $\text{reversal}(\cdot)$ as follows.

$$\text{reversal}(s_j, d_{t+y}) = \begin{cases} \text{true} & \text{if } (\text{uptrend} \wedge \mathbf{f}^*(s_j, d_{t+y}) = \text{down}) \\ & \vee (\text{downtrend} \wedge \mathbf{f}^*(s_j, d_{t+y}) = \text{up}) \\ \text{false} & \text{otherwise} \end{cases}$$

For example, in Figure 1 $\text{reversal}(s_j, d_{t+1})$ takes value *true* on day d_{t+1} because the ongoing trend is an uptrend whereas the classifier prediction is *down*.

D. Trade management

The trading system exploits the trend reversal forecasts to drive investments in the stock market. More specifically, the detection of a trend reversal in a future day (i.e., $\text{reversal}(\cdot)=\text{true}$) causes the opening of a multi-day trading position on the corresponding stock. A long selling is opened if the classification process was triggered by a downtrend (betting on a short-term increase of the stock prices). Conversely, a short-selling position is opened if the trigger was an uptrend (betting on a short-term decrease of the stock prices). In the example in Figure 1, the detection of a reversal on day d_{t+1} causes the opening of a short-selling position.

All the trading positions are opened at market opening (i.e., at the stock opening price of the considered day) and closed at market closure. Each position is automatically managed by the quantitative trading system by setting (i) a percentage stop loss value (i.e., the maximum admissible loss per trade), (ii) a take profit (at least twice the stop loss value to guarantee a favourable risk/reward ratio), and (iii) a maximum trade duration (expressed in days) to minimize market exposure of short-term trades.

V. EXPERIMENTS

We performed an extensive empirical evaluation of the performance of the proposed approach on real stock-related data to answer to the research questions posed in Section I.

Data. We crawled 11-year stock market data, i.e., from the beginning of year of 2007 to the end of year 2017. The acquired data are related to all stocks of the American Standard & Poor 500 index using the APIs of Yahoo Finance (<https://finance.yahoo.com/>). To study the influence of financial news on stock prices, we crawled also the stock-related news released in the same period by Reuters (<https://www.reuters.com/>). To prepare the training datasets we considered as descriptors the stock closing prices, the technical oscillators and volatility indices, and the news sentiment indicators described in Section IV-A. Since the market conditions change over time, we separately analyzed the stock and news data related to different years.

Algorithms. We analyzed the performance of various classification models, including a Support Vector Machines, a Bayesian Classifier, a feedforward Artificial Neural Network, an ensemble method, and a distance-based classifier. Specifically, we considered the algorithm implementations available in the Scikit-learn library (<https://scikit-learn.org/>). To configure the classification models we run a grid search separately for each algorithm. The main configuration settings are: (i) Support Vector Classifier (SVC): RBF Kernel, $C=1$, $\text{Gamma}=\frac{1}{|D|\cdot\sigma_X^2}$ (where $|D|$ is the number of data features, while σ_X^2 is the feature variance), (ii) Multinomial Naive Bayes (MNB): $\alpha=1.0$, (iii) Random Forest (RFC): $\text{Criterion}=\text{Gini}$, $\text{n_estimators}=100$, (iv) MultiLayer Perceptron (MLP): $\text{hidden_layers}=1$, $\text{hidden_layer_size}=50$, $\text{solver}=\text{lfbfgs}$, and (v) K-Nearest Neighbor (KNN): $\text{k}=5$, $\text{weights}=\text{uniform}$.

Trading simulation. We backtested the proposed trading system separately on each dataset by making the following assumptions: (i) an initial portfolio of 100,000 USD, (ii) a fixed transaction cost equal to 5 USD, (iii) a variable amount per trade equal to 5% of the current equity, (iv) an adaptive stop loss computed by considering as maximum absolute loss the 25% of the total price variation during the ongoing trend (i.e., the variation of the closing price from the beginning of the ongoing trend to the current day), and (v) a maximum trade duration equal to 5 days.

Classifier training. To train the classification models we applied an expanding window approach. More specifically, during the simulation we trained the classifiers at the times at which an uptrend/ongoing trend is detected. Let d_t be the day at which the ongoing trend is detected. The training dataset includes all the records corresponding to the current and past days from the beginning of the year to d_t . As described in Section IV-A, training records are labeled according to the target of the classifier as $\text{dir}(s_j, d_{t+y})$ [$1 \leq y \leq 5$]. Due to the lack of training data, we did not consider the trends detected in the first two months (i.e., for each dataset the backtesting phase was run from the beginning of March to the end of December).

Quality measures. We evaluated the performance of the trading system in terms of percentage total return. Furthermore, we evaluated the ability of classifiers to predict trend

TABLE I

AVERAGE TOTAL RETURN (%) PER ADOPTED CLASSIFIER AND FEATURES.

Features	KNN	MLP	MNB	RFC	SVC	Avg.
NEWS+PRICE	66%	78%	18%	73%	32%	53%
OSC+PRICE	43%	77%	26%	76%	27%	50%
OSC+PRICE+NEWS	45%	75%	23%	67%	23%	47%
PRICE	64%	54%	2%	54%	23%	39%

reversal under various aspects [23]. Specifically, we computed (i) the precision of each classifier while predicting a reversal (i.e., $\text{reversal}(s_j, d_{t+y}) = \text{true}$), as the percentage of correctly detected trend reversals over the total number of trend reversals predicted by the classifier. (ii) The recall of each classifier while predicting a reversal, as the percentage of correctly detected trend reversals over the total number of actual trend reversal in the test data. (iii) The Root Mean Square Error (RMSE), expressed in days, between the actual and predicted closest day of reversal with respect to the current day. It indicates to what extent classifiers are precise while predicting the closest day of reversal. To compute the RMSE, we combined classifier predictions generated by setting different values of prediction horizon N and we supposed the error to be bounded between zero (correct prediction) and five (missing prediction within the maximum prediction horizon).

The rest of this section is organized as follows. Sections V-A and V-B compare the returns achieved by the trading strategies based on different categories of descriptors (or combinations of categories) and on different classification models. The aim is to understand whether news information is worth considering in the trend reversal strategy (Research Question RQ1) and, if so, which categories of descriptors should be considered in combination with the news sentiment (research question RQ3). Section V-C analyzes the impact of the (user-specified) duration (W) of the ongoing trend on the returns of the proposed trading system. The goal of this analysis is to identify a proper configuration setting for the trend detection phase. Section V-D shows some examples of equity lines achieved by the proposed strategies. Section V-E evaluates the ability of the classification models, in terms of precision, recall, and RMSE, to correctly predict trend reversal in the future days. The aim is to measure the impact of including news information on classifier performance (research question RQ2). Finally, section V-F briefly discusses the time complexity of the classification-based approach.

All the experiments were run on a machine equipped with Intel® Xeon® X5650, 32 GB of RAM and running Ubuntu 18.04.1 LTS.

A. Comparison between different trading strategies

Table I reports, for each combination of stock descriptor category and classifier, the average total return (expressed in percentage) computed by averaging the total returns achieved over all the considered datasets (11 years of data). For each classifier the best performing strategy is written in boldface. The combination of NEWS with PRICE (NEWS+PRICE) yielded the highest returns for 3 out of 5 classifiers, while

TABLE II

AVERAGE RANKINGS AND FRIEDMAN TEST (CD=0.06). RANKING BASED ON TOTAL RETURN.

Features	KNN	MLP	MNB	RFC	SVC	AVG
NEWS+PRICE	1.55	2.09	2.73	1.64	2.09	2.02
OSC+PRICE	3.27	1.73	1.45	1.82	2.18	2.09
OSC+PRICE+NEWS	3.00	2.45	1.82	3.00	3.45	2.74
PRICE	2.18	3.73	4.00	3.55	2.27	3.15

TABLE III

IMPACT OF THE TREND DURATION W ON THE AVERAGE TOTAL RETURN (%). MLP CLASSIFIER.

Features	W=3	W=4	W=5
NEWS+PRICE	68%	105%	61%
OSC+PRICE	71%	101%	58%
OSC+PRICE+NEWS	69%	101%	57%
PRICE	50%	70%	42%

the combination of OSC and PRICE ranked second (best in 2 cases out of 5). Hence, news information appeared to be worth considering while planning trend reversal strategies. The MLP classifier performed significantly better than the other classification methods. The strategy based on the NEWS+PRICE combination achieved significantly higher returns than those relying only on the PRICE descriptors (+24% for MLP, +14% by averaging all the classifiers). Combining all the descriptor categories together appeared to be not convenient in the performed simulations, probably due to the higher complexity of the training phase while coping with high-dimensional data.

B. Statistical significance tests

We applied the Friedman test to evaluate the statistical significance of the achieved performance gaps between the strategies relying on different combinations of descriptors. The test was performed separately for each classification algorithm. For each classifier we applied the following procedure: (1) Separately for each dataset, the feature combinations were sorted by decreasing total return. (2) For each combination of descriptors, we computed its average rank over all the considered datasets (from 2007-data to 2017-data). (3) The observed average ranking gaps between each pair of descriptor combinations were compared with the critical difference threshold CD (0.06). The gap was statistically significant, at the 95% significance level, if its value was above CD.

Table II summarizes the achieved results. The NEWS+PRICE combination achieved the best average mean rank value for 3 classifiers out of 5 and the performance gaps are all statistically significant (including the best performing MLP classifier). By averaging the ranks over all the classifiers NEWS+PRICE turned out to be the best combinations of stock descriptors. Therefore, news sentiment is worth considering while predicting a trend reversal will occur in the near future.

C. Impact of the ongoing trend duration

Since the prediction phase is triggered by the detection of an uptrend/downtrend, we empirically analyzed the impact of the




Fig. 2. Equity line. MLP classifier. Trend duration $W=4$.

user-specified trend duration (W) on the trading system performance. Table III reports the average total returns achieved by different strategies while varying the ongoing trend duration W between 3 and 5 (i.e., the predictions are triggered by trends lasting 3, 4, or 5 days, respectively).

The results show that setting the ongoing trend duration to 4 yielded the highest returns. On the one hand, increasing the minimum duration of a trend decreases the number of opened trades (because it is less likely to predict a trend reversal). On the other hand, decreasing the duration resulted in generating more false signals (because trends become less reliable). Setting W to 4 yielded a good trade-off between speculative approaches and conservative ones.

D. Equity lines

Quantitative trading systems are requested to be both profitable and robust against unexpected draw-downs in the long run. Equity lines are established representations of the temporal variations of the equity value. Figures 2(a)-2(b) show the equity lines corresponding to two representative years (2011 and 2013) using different combinations of stock descriptors and setting the best system configuration, i.e., MLP algorithm with trend duration W set to 4. Year 2011 is an example of bearish market, while in 2013 the market was bullish.

The equity lines show the robustness of the proposed strategy. In 2011 NEWS+PRICE achieved a 130% final return with minimal draw-down and its performance is as good as that achieved by OSC+PRICE and significantly better than PRICE. In 2013 NEWS+PRICE performed better than OSC+PRICE and PRICE, especially in the second half of the year.

E. Prediction quality

To evaluate the robustness of the classification models, we measured the precision, recall, and RMSE achieved by the classifiers during the trading simulations. Figures 3 and 4 respectively plot the precision and recall values achieved by the MLP classifier in years 2011 and 2013 using different combinations of descriptor categories. Since separate classifiers are applied to predict reversal at different prediction horizons (N), the reported values of precision and recall vary according to the value of N . The precision values are fairly high, i.e., they are all above 79% in year 2011, while they are all above 69%

TABLE IV
RMSE OF MLP TREND REVERSAL PREDICTIONS IN 2011 AND 2013.

Features	2011			2013		
	W=3	W=4	W=5	W=3	W=4	W=5
NEWS+PRICE	1.67	1.55	1.58	1.88	1.64	1.09
OSC+PRICE	1.87	1.81	1.61	2.03	1.96	2.16

on 2013-data. The recall values for NEWS+PRICE are also promising (above 82% in 2011, above 62% in 2013). Hence, prediction models appeared to be reliable.

Table IV reports the RMSE (expressed in days) achieved by the MLP classifier in years 2011 and 2013 for the NEWS+PRICE and OSC+PRICE combinations. The NEWS+PRICE RMSE is always lower than those of OSC+PRICE and it is close to 1.5 by setting W to 4. Hence, trend reversal forecasts were timely generated to effectively support the quantitative trading system.

F. Execution time


The time complexity of the proposed approach mainly depends on the data preparation and on the classifier training phases. Data preparation took, on average, 130s. Classifier training took between 15s ($W=3$) and 400s ($W=5$) depending on the classification algorithm. The MLP classifier turned out to be the most computationally intensive approach (one order of magnitude slower than the others). Applying the trained classifiers to predict the future directions took negligible time.

VI. CONCLUSIONS AND FUTURE WORKS


This paper explores the usefulness of integrating news information for improving the effectiveness of classification-based trend reversal trading strategies. It compares the performance of different stock descriptor combinations (including news sentiment and not) and classification algorithms on an 11-year U.S. stock dataset enriched with news information. The results show that:

(RQ1) News information is worth considering to perform trend reversal detection.


(RQ2) Classifiers trained on news information are robust enough to be effectively integrated into trend reversal trading systems.




(a) Year 2011



(b) Year 2013

Fig. 3. Precision by varying the prediction horizon. MLP classifier, Trend duration $W=4$.

(a) Year 2011



(b) Year 2013

Fig. 4. Recall by varying the prediction horizon. MLP classifier. Trend duration $W=4$.

(RQ3) Combining news information with historical price series allowed us to achieve performance superior to (or, in the worst case, comparable to) that achieved by the most established oscillators used in technical analysis to detect trend reversal.

As future work, we will (i) incorporate additional news sentiment features (trained on social media as well), (ii) integrate deep learning approaches (e.g., LSTM), and (iii) study the portability of the applicability of the system towards diverse contexts (e.g., cryptocurrency trading).

REFERENCES

- [1] R. C. Cavalcante, R. C. Brasileiro, V. L. Souza, J. P. Nobrega, and A. L. Oliveira, "Computational intelligence and financial markets: A survey and future directions," *Expert Systems with Applications*, vol. 55, pp. 194 – 211, 2016.
- [2] E. Chan, *Algorithmic Trading: Winning Strategies and Their Rationale*, 1st ed. Wiley Publishing, 2013.
- [3] J. Murphy, *Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications*, ser. New York Institute of Finance Series. New York Institute of Finance, 1999.
- [4] A. O'Sullivan and S. Sheffrin, *Economics: Principles in Action*, ser. Prentice Hall Science/Social Studies. Prentice Hall, 2006.
- [5] Z. Wang, S.-B. Ho, and Z. Lin, "Stock market prediction analysis by incorporating social and news opinion and sentiment," in *IEEE International Conference on Data Mining Workshops*, 2018, pp. 1375–1380.
- [6] J. Bollen, H. Mao, and X. Zeng, "Twitter mood predicts the stock market," *Journal of computational science*, vol. 2, no. 1, pp. 1–8, 2011.
- [7] D. Enke and S. Thawornwong, "The use of data mining and neural networks for forecasting stock market returns," *Expert Systems with Applications*, vol. 29, no. 4, pp. 927 – 940, 2005.
- [8] L. A. Teixeira and A. L. I. de Oliveira, "A method for automatic stock trading combining technical analysis and nearest neighbor classification," *Expert Systems with Applications*, vol. 37, no. 10, pp. 6885 – 6890, 2010.
- [9] R. K. Nayak, D. Mishra, and A. K. Rath, "A naive SVM-KNN based stock market trend reversal analysis for indian benchmark indices," *Applied Soft Computing*, vol. 35, pp. 670 – 680, 2015.
- [10] P. C. Tetlock, "Does Public Financial News Resolve Asymmetric Information?" *The Review of Financial Studies*, vol. 23, no. 9, pp. 3520–3557, 2010.
- [11] A. Nikfarjam, E. Emadzadeh, and S. Muthaiyah, "Text mining approaches for stock market prediction," in *The 2nd International Conference on Computer and Automation Engineering*, vol. 4, 2010, pp. 256–260.
- [12] S. S. Groth and J. Muntermann, "An intraday market risk management approach based on textual analysis," *Decision Support Systems*, vol. 50, no. 4, pp. 680 – 691, 2011, enterprise Risk and Security Management: Data, Text and Web Mining.
- [13] E. J. Ruiz, V. Hristidis, C. Castillo, A. Gionis, and A. Jaimes, "Correlating financial time series with micro-blogging activity," in *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining (WSDM'12)*, 2012, pp. 513–522.
- [14] B. Li, K. C. Chan, C. Ou, and S. Ruifeng, "Discovering public sentiment in social media for predicting stock movement of publicly listed companies," *Information Systems*, vol. 69, pp. 81 – 92, 2017.
- [15] T. M. Nisar and M. Yeung, "Twitter as a tool for forecasting stock market movements: A short-window event study," *The Journal of Finance and Data Science*, vol. 4, no. 2, pp. 101 – 119, 2018.
- [16] R. P. Schumaker, Y. Zhang, C.-N. Huang, and H. Chen, "Evaluating sentiment in financial news articles," *Decision Support Systems*, vol. 53, no. 3, pp. 458 – 464, 2012.
- [17] A. Picasso, S. Merello, Y. Ma, L. Oneto, and E. Cambria, "Technical analysis and sentiment embeddings for market trend prediction," *Expert Systems with Applications*, vol. 135, pp. 60 – 70, 2019.
- [18] N. T. Son, L. Van Thanh, T. Q. Ban, D. X. Hoa, and B. N. Anh, "An analyze on effectiveness of candlestick reversal patterns for vietnamese stock market," in *Proceedings of the 2018 International Conference on Information Management & Management Science*. ACM, 2018, pp. 89–93.
- [19] K. Lee and G. Jo, "Expert system for predicting stock market timing using a candlestick chart," *Expert Systems with Applications*, vol. 16, no. 4, pp. 357 – 364, 1999.
- [20] X. Chu, Z. Gu, and H. Zhou, "Intraday momentum and reversal in chinese stock market," *Finance Research Letters*, vol. 30, pp. 83 – 88, 2019.
- [21] D. Heyman, M. Lescauwat, and H. Stieperaeer, "Investor attention and short-term return reversals," *Finance Research Letters*, vol. 29, pp. 1 – 6, 2019.
- [22] T. Williams and V. Turton, *Trading Economics: A Guide to Economic Statistics for Practitioners and Students*. Wiley, 2014.
- [23] P.-N. Tan, M. Steinbach, and V. Kumar, *Introduction to Data Mining*. Addison-Wesley, 2006.