

## RESEARCH ARTICLE

# Fusing Sell-Side Analyst Bidirectional Forecasts Using Machine Learning

THENDO SIDOGI<sup>1,2</sup>, WILSON TSAKANE MONGWE<sup>1,2</sup>, RENDANI MBUVHA<sup>3,4</sup>,  
AND TSHILIDZI MARWALA<sup>1,2</sup>, (Senior Member, IEEE)

<sup>1</sup>School of Electrical Engineering, University of Johannesburg, Auckland Park 2000, South Africa

<sup>2</sup>Johannesburg Machine Intelligence Laboratory, Auckland Park 2006, South Africa

<sup>3</sup>School of Statistics and Actuarial Science, University of Witwatersrand, Johannesburg 2006, South Africa

<sup>4</sup>School of Electronic Engineering and Computer Science, Queen Mary University of London, London E1 4NS, U.K.

Corresponding author: Thendo Sidogi (thendos@gmail.com)

The work of Thendo Sidogi was supported by Tshilidzi Marwala's research account. The work of Wilson Tsakane Mongwe and Rendani Mbuva was supported by the Google Ph.D. Fellowships in Machine Learning. The work of Tshilidzi Marwala was supported by the National Research Foundation of South Africa.

**ABSTRACT** Sell-side analysts' recommendations are primarily targeted at institutional investors mandated to invest across many companies within client-mandated equity benchmarks, such as the FTSE/JSE All-Share index. Given the numerous sell-side recommendations for a single stock, making unbiased investment decisions is not often straightforward for portfolio managers. This study explores the use of historical sell-side recommendations to create an unbiased fusion of analyst forecasts such that bidirectional accuracy is optimised using random forest, extreme gradient boosting, deep neural networks, and logistic regression. We introduced 12-month rolling features generated from standard sell-side recommendations, such as analyst coverage, point and directional accuracy, while avoiding forward-looking biases. We introduce a novel "AI analyst" by fusing forecast features from numerous analysts using machine learning algorithms. We observed the added benefits of using these features from more than one analyst by systematically generating unbiased and incrementally better prediction accuracy from publicly available sell-side recommendations, with the Random forest algorithm showing the highest relative performance. In highly volatile sectors, like resources, the machine learning algorithms perform better than in low volatility sectors, suggesting the importance of rolling features in bi-directional prediction in the presence of high volatility. Using feature importance, we observe the incremental contribution of rolling features, showing the relationships between analyst coverage, volatility, and bidirectional forecast accuracy. Furthermore, parameters from logistic regression identify volatility features and initial and target price as some of the essential features when modelling analysts' directional predictions.

**INDEX TERMS** Accuracy, AI-analyst, analyst scores, DNN, feature importance, logistic regression, machine learning.

## I. INTRODUCTION

Institutional investors are primarily mandated to invest in an equity benchmark representing the most prominent companies listed on an exchange [1]. Among these companies within the equity benchmark, brokerage firms hire analysts to produce investment reports called sell-side equity analyst reports. Sell-side analyst reports are highly valued in the investment industry, and generally include three standard

quantitative outputs: earnings forecasts, target prices, and buy/sell recommendations [2].

The importance of sell-side reports has been demonstrated in several ways. Sell-side analysts play an essential role in interrogating company results and executive management of their strategies, while ensuring the reduction of any mismanagement or unethical behaviour [3]. Another important aspect of sell-side analyst reports is reducing the cost of capital whenever companies issue new equity to raise funds. This is done mainly by reducing information asymmetry between potential investors and their respective companies,

The associate editor coordinating the review of this manuscript and approving it for publication was Chao Tong<sup>1</sup>.

which improves stock liquidity. However, this introduces a potential conflict of interest because analysts can generate forecasts or recommendations for other economic incentives, such as securing underwriting business and boosting trading volume [4]. Other issues with sell-side analysts' reports were found in Chiang *et al.* [5], who indicate that analysts tend to herd towards a consensus when issuing recommendations, and this tendency increases with market sentiment. Market- or stock-specific sentiment has been shown to impact stock price movements [6].

Investors typically use sell-side analyst information to determine earnings growth and directional movements of stock prices, and hence, their investment decisions [7]. In summary, this approach is best described as obtaining market benchmark expectations from sell-side analyst reports. Therefore, assuming that equity markets are semi-strong, investors can generate alpha returns using these reports [8]. The question of whether sell-side reports can generate alpha performance has been of interest in previous research. Some studies have shown the existence of potential investment strategies in sell-side reports. Barber *et al.* [9] found that by using analyst ratings, the resultant investment strategies generate excess returns above 4%. Womack *et al.* [10] also found that brokerage sell-side analysts tend to produce good directional accuracy for stock price predictions.

Most prior studies on the value and potential alpha of these reports has focused on using analyst ratings. In contrast, researchers have only recently begun to focus on earnings and target price analyses because of lack of data. Major financial data vendors have only recently begun to capture target price data continuously [11]. There has been an increase in sell-side analysts in the South African brokerage market, resulting in more companies and sell-side reports being issued by various equity analysts. Inevitably, this has presented the need for investors to process sell-side data rapidly and in an unbiased manner while assisting with their investment processes [9]. Big investment houses have the resources to deal with more sell-side data, but smaller investment firms may need to think outside the box in order to remain competitive with the rest of the market.

In this study, we explore the idea of generating superior directional predictions by fusing sell-side data reports within companies belonging to the FTSE/JSE All-Share index using machine learning algorithms. We propose an artificial intelligence analyst or "AI analyst" to assist in using sell-side analyst data and generate investment recommendations, including directional predictions with little human intervention. We do this by considering investment signals from sell-side reports as quantitative input features to various machine-learning algorithms to predict the directional movement of individual stocks.

Machine learning algorithms have limited applications to sell-side earnings forecast data. To the best of our knowledge, at the time of writing, we do not know any studies investigating the application of these algorithms to sell-side earnings forecast data. A better approximation of market

consensus estimates using an iterative filtering algorithm compared to simply using the mean average of consensus estimates was introduced in [12]. The closest research on the topic was in [13], where analyst features, including analyst ratings, were used for stock price prediction. Analyst ratings and profitability of recommendations have been the focus of burgeoning literature [14], [15]. However, most prior studies have examined these outputs in isolation rather than studying the possible use of the combined outputs from the reports. Loh *et al.* [16] provide the value and use cases of sell-side earnings forecast data for investment recommendations.

The large number of resources that generate sell-side earnings forecasts suggests that they may be helpful in potential investment strategies. It is evident in [17] that in addition to earnings forecasts, several pieces of information on sell-side research data may provide additional information. Stock price forecasts by analysts typically have a short forecast horizon and allow investors to easily observe the expected direction and magnitude of equity price movements. Analysts may release revisions to their forecasts when new information becomes available. We use the methodology and findings of previous studies on sell side earnings forecast accuracy to conduct our analysis.

Accuracy in target prices or earnings forecasting is helpful for users of sell-side reports. Analysts have different abilities and skills within their respective coverage to forecast earnings and target prices for different time horizons [14]. The accuracy of near-term forecasts may depend on many factors, including analysts' ability to interpret market sentiment and the current economic and industrial themes [15]. Mikhail *et al.* [18] found that stock coverage, analyst following, and the ability to extract specific earnings guidance from insiders have a significant influence on analysts' forecast accuracy and in the long term, analysts' ability to predict these economic, industry, and company trends determines their forecast accuracy.

Prior evidence on the profitability of using sell-side data is mixed [11], [19]. If analysts with the most profitable recommendations, such as those described in [14], also exhibit superior forecasting skills, this would suggest that their stock-picking ability presents various phenomena occurring in the stock market [20]. Differentiation between sell side analyst may present an opportunity to introduce different investment strategies based on sell side data [11]. This study introduces 12-month rolling forecast accuracy scores to differentiate between sell-side analyst reports.

**Contributions:** Our contributions in this work can be summarised as follows:

- The first is generation of unbiased and improved directional predictions. Groyberg *et al.* [21] found that analysts may be inclined to generate sell-side reports to increase their compensation, particularly because factors such as forecast accuracy are not considered when making variable remuneration. This is mainly because of the compensation structure of the investment brokerage firms. The coverage of specific companies

can increase brokerage revenue by increasing trading activity. An unbiased AI analyst can provide a solution to investors and analysts facing conflicts of interest.

- Second, there has been an increase in sell-side analyst reports that are available to investors. Human consumption of these data can be challenging compared to using scalable technologies, such as artificial intelligence or machine learning algorithms. The direct valuation implications of target price forecasts make them potentially useful investment signals regarding directional stock price movements using the AI analyst.

The remainder of this paper is organised as follows: In section II, we discuss the collection of South African sell-side data describing various forecast accuracy and valuation recommendation metrics used. We also examine how the prior literature considers these metrics. We also describe how we introduced 12-month rolling features generated from standard outputs from sell-side reports from a research design perspective. A description of the machine learning algorithms used for the prediction is provided in section III. The results are presented in section IV, starting with the prediction accuracy for all the machine learning models and followed by a discussion on the feature importance of some machine learning algorithms in V. Section VI concludes the paper and outlines ideas for future research.

## II. DATA AND RESEARCH DESIGN

Similar to [11], the data used in this manuscript were sourced from Bloomberg. It is a collection of standard outputs from the historical sell-side reports of companies within the FTSE/JSE all-share index from January 2004 to June 2018. The main features generated from the reports include analysts' firm of employment at the date of the forecast, stock industry and sector, rank within the FTSE/JSE all-share index in terms of market capitalisation, forecast stock price including the date, and forecast stock rating which is split into five categories. We use some of these features to generate additional input features, which are discussed in the next section. To avoid any look-ahead bias, the first year of the dataset was removed to calculate 12 month rolling features. The dataset used for these experiments, including rolling features, ranged from July 2005 to June 2018.

### A. 12 MONTH ROLLING FORECAST ACCURACY METRICS

Before introducing rolling features, we first discuss generic definitions from previous research. Our research focus and experimental design have drawn considerable attention to directional accuracy measures of sell-side analysts. Directional forecast accuracy refers to the analysts' accuracy in predicting the future direction of stock prices. This measure is significant, as many investors choose to trade based on the expected directional movement (e.g. long/short or option and tactical strategies) rather than the expected price [22].

The standard methodology for the directional accuracy can be measured using a confusion matrix [23]. A confusion matrix contains the number of Type I and Type II errors

together with the number of correct forecasts. When there is a decrease in stock price, while the analyst reports forecasts an upward movement, it is defined as Type I (i.e. a false positive). A Type II error is defined as an increase in stock price when the analyst forecast moves downward (i.e. a false negative). Correct forecasts are then split into true-positives and true-negatives.

Directional accuracy (*DA*) can be considered as a measurement of classification accuracy. Our research focused on a binary classification problem, that is, upward or downward classes. Let true positives (*TP*) refer to the number of positive classes correctly predicted by the classifier. True negatives (*TN*) refers to the number of negative classes correctly predicted by the classifier. Type I error or false-positive (*FP*) and type II error or false negative (*FN*) refer to the number of classes incorrectly classified as either positive class as a negative or negative class as positive. *DA* can be defined as

$$DA = \frac{TN + TP}{TP + TN + FP + FN} \quad (1)$$

We introduced the directional accuracy score (*DAS*) based on the generic or classic definition of directional accuracy. The measurement is a cumulative score based on the number of times the analyst correctly (*TP* or *TN*) and incorrectly (*FP* or *FN*) predicted directional movement over the previous 12-month period. For a given period, the directional accuracy score can be measured as

$$DAS = \begin{cases} \sum_i x_i \times 0.75, & \text{if } x = 1 \\ \sum_i x_i \times 0.01, & \text{if } x = 0 \end{cases} \quad (2)$$

where  $x$  is the directional prediction and the scaling factor is chosen to differentiate between the analyst scores.

We also consider point accuracy to be the percentage difference between the target and realised prices. It measures the accuracy of an analyst's price prediction. In this study, we only focused on the unassigned point accuracy (*UPA*) [24], which is the absolute value of the point accuracy, defined as

$$UPA = \left| \frac{P_{tg} - P_{12}}{P_{12}} \right| \quad (3)$$

where  $P_{tg}$  and  $P_{12}$  is the price target and realised price after 12 months.

We introduce the point accuracy score (*PAS*) which is the average absolute point accuracy for an analyst over a 12 month period.

$$PAS = \frac{\sum_i |UPA_i|}{\sum n} \quad (4)$$

where  $n$  is the number of reports issued of the period.

We consider two metrics used in prior research to measure the term accuracy [11], [25], [26]. The first metric, MetEnd, measures whether the target prices are achieved at the end of the period by comparing the target price with the actual closing price at the end of the horizon. The second metric, (MetAny), measures whether target prices are achieved at any point during the forecast period. These are then binary

TABLE 1. Sample of input features. \* indicating 12 month rolling features.

Firm	Target Price	Rating	$P_0$	$\frac{P_0}{P_1}$	20dvol	Top40CSV	RelVol	DAS*	MES*	MAS*	PAS*	PTM*	AC*
61	210	1	247.8	0.847	0.014	0.051	0.264	0	0.01	0.01	-0.09	-0.02	15
61	5.22	1	5.18	1.008	0.008	0.074	0.103	2.75	2.25	3.75	-0.109	0.283	15
3	200	3	209.96	0.953	0.01	0.067	0.318	24	16.5	19.5	-0.109	0.292	165
38	405.70	5	311.5	1.302	0.017	0.057	0.302	13.5	9.0	21.75	-0.134	0.217	155
77	122.58	3	117.9	1.04	0.011	0.067	0.157	6.0	4.5	6.75	-0.079	0.123	51

variables where 1 indicates that the target price was achieved and 0 indicates that it was not achieved during the forecast period.

We introduce a MetEndScore (MES) and MetAnyScore (MAS) which is the cumulative number of times an analyst prediction meets the target price at the end of the forecast period or at anytime during the forecast horizon. Similar to DAS, these rolling scores are defined as

$$MES = \begin{cases} \sum_i x_i \times 0.90, & \text{if TP} \\ \sum_i x_i \times -0.05, & \text{if TN} \end{cases} \quad (5)$$

$$MAS = \begin{cases} \sum_i x_i \times 0.95, & \text{if TP} \\ \sum_i x_i \times -0.15, & \text{if TN} \end{cases} \quad (6)$$

where  $x$  is analysts' MetEnd in equation (5) and MetAny in equation (6) between periods  $i$  and  $n$  with the scaling factor chosen to differentiate between analyst scores.

For features relating to stock coverage and earnings momentum, we introduce analyst coverage AC, which shows the total number of reports for the past 12 months on a specific stock, defined as

$$AC = \sum_i^n x_i \quad (7)$$

where  $n$  is the previous 12 months from the sell-side analyst's issue date for stock  $x$ .

For earnings momentum, we use the price target momentum (PTM), defined as the percentage change in the current issued target price from the previously published target price.

$$PTM = \frac{CurrentTP - PreviousTP}{PreviousTP} \% \quad (8)$$

Similar to [11], we include analyst firm and forecast stock ratings as analyst report features. Stock volatility related features include rolling 20-day volatility (20dvol), monthly Top40 Index cross-sectional volatility (Top40CSV), and the ratio between the two, defined as:

$$RelVol = \frac{20dvol}{Top40CSV} \quad (9)$$

Table 1 shows a sample of the input features calculated from their definitions.

In this study, supervised machine learning techniques were used. The target to be predicted over the forecast period is calculated using the realised closing price  $P_{12}$  of the respective stock using the following formula:

$$Y_i = Sign(P_{12} - P_0) \quad (10)$$

where  $P_0$  is the price at issue date for analyst report  $i$ .

The target is labelled as either 1 or 0; a zero value of the target indicates that the price of the stock has fallen, whereas a

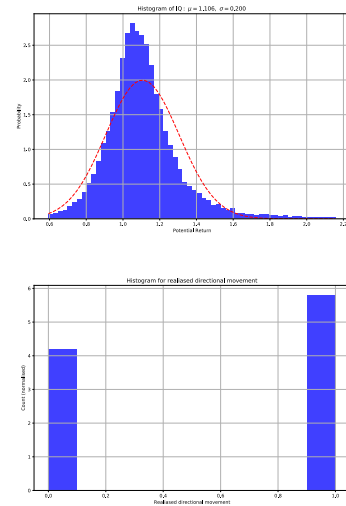


FIGURE 1. Class distribution of analyst directional prediction and realised stock directional movement.

positive value of the target indicates that the price of the stock has risen during the sell-side analysts' forecast period. This directional change framework is used to transform a financial time series into an event-based series [22]. Figure 1 shows the frequency distribution of potential return target and realised directional movement. We see that more analyst have positive return potentials. This is maybe due to the fact that markets have been trending upwards over the long term.

### B. RESEARCH DESIGN

Notably the starting period of the newly generated features was 12 months after the starting period of the sample data. This is because, to generate these scores for analysts, one would need previous forecast accuracy data from previous reports to prevent any forward look-head bias. From the definitions mentioned in Section 3.1, we generated 14 input features based on the FTSE/JSE all-share sell-side data as inputs to the machine learning models, which we discuss in the next section.

In summary, our proposed workflow is depicted in figure 2 shown below

### C. PERFORMANCE METRICS

Specific performance metrics were used to measure the performance of the classification algorithm. We used classification accuracy and precision scores to evaluate the performance of respective machine learning models in predicting stock price directional movement. The classification accuracy of a model is the ratio of the number of correct predictions made to the total number of predictions made, similar to equation (1).

The precision score is given by the ratio between the true positive predictions and the total positive predictions given by

$$PrecisionScore = \frac{TP}{TP + FP} \quad (11)$$

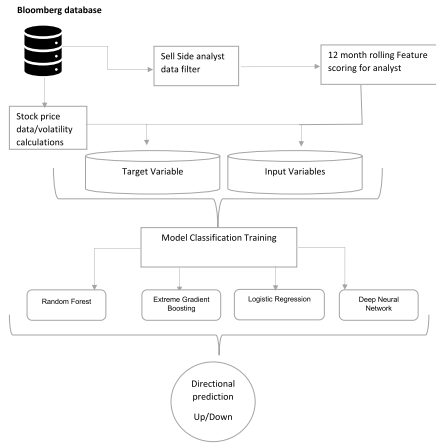


FIGURE 2. Proposed framework for stock price directional prediction.

### III. MACHINE LEARNING ALGORITHMS

To predict the directional movement of stock prices based on the features discussed, we used machine learning algorithms. These algorithms have become popular with the increase in computing power and data availability. We propose and design a feature-generating process using sell-side analyst data to predict the directional movement of stocks.

#### A. RANDOM FOREST (RF)

Random forest is a machine learning algorithm that uses an ensemble model that combines several base decision trees to produce an optimal predictive model and make point or classification predictions [27]. Decision trees are models with a high variance and low bias. When increasing the number of decision trees, the random forest reduces the high variance by having uncorrelated individual decision trees and taking the average class prediction from each decision tree while the bias remains constant. In each training step, samples from the training sets are selected randomly with replacements and a subset of features is selected randomly and whichever feature provides the best split or decision is used. The best split or branching of the tree is given by a feature where the root node gives the lowest Gini index in equation (12).

$$GiniIndex = 1 - \sum_{i=1}^C (f_i)^2 \tag{12}$$

where  $C$  is the number of classes and  $f_i$  is the frequency of label  $i$  on a node.

In our study, the stock price movement random forest is a classifier consisting of an ensemble of decision tree-structured classifiers  $\{h(x, \theta_k), k = 1, \dots\}$  where  $\{\theta_k\}$  are independent identically distributed random vectors and each decision tree predicts a unit class of either upward or downward movement to obtain an average class prediction of sell side analyst report features as inputs. The number of decision trees to be boosted was 100 with no maximum depth specified.

#### B. EXTREME GRADIENT BOOSTING (XGB)

State-of-art machine learning techniques for classification or regression tasks include extreme gradient boosting, first introduced by Friedman (2001) [28]. Gradient boosting is an optimisation problem, where the goal is to locate the minimum or minimise the model’s loss function by adding weak learners using gradient descent. During the learning process, decision trees are added at each stage, and the existing trees in the ensemble are not replaced. The contribution of the weak learner to the ensemble is based on the gradient descent optimisation process performed on each weak learner. The calculated contribution of each tree is then based on minimising the overall error of the composite learner.

In our experiment, the inputs  $x_i^p$  and  $y_i$  target set can be defined as  $\{(x_i^p, y_i)\}_{i=1}^n$  with a differentiable function  $L(y, F(x))$ , where  $F(x)$  is the function mapping  $x^p$  to  $y$  and is the weaker model within the ensemble tree with number of iterations  $m$ . The proposed optimisation problem is defined as follows

$$\gamma_{jm} = \arg \min_{\gamma} \sum_{x_i \in R_{jm}} L(y_i, F_{m-1}(x_i) + \gamma) \tag{13}$$

$$F_m(x) = F_{m-1} + \sum_{j=1}^{J_m} \gamma_{jm} 1_{R_{jm}}(x) \tag{14}$$

where chooses a separate optimal value  $\gamma_{jm}$  for each of the tree’s regions, instead of a single  $\gamma_m$ . The XGB model has the benefit of better interpretability in which we can estimate the relative importance of input features in predicting stock price directional movement. The number of decision trees to be boosted was 950 with a maximum depth of 4 for each tree.

#### C. DEEP NEURAL NETWORKS (DNN)

We also model stock price direction movement using deep neural networks. These are artificial neural networks (ANN) with several layers between the chosen architectures. ANNs are machine learning algorithms inspired by the human brain that help generate human intelligence without explicit programming.

We also employ a deep neural network to classify or predict the directional movement of stock prices. It has four layers of interconnected neuron units, through which the data are transformed. The input layer with 14 neurones represents the generated input features and market volatility features described earlier. The hidden layers have a ReLU activation function, whereas the output layer uses a sigmoid function. The first hidden layer has a dimension of 7, with a total of 105 parameters. The second hidden layer has a dimension of 3 with a total of 24 parameters. The last output layer with the sigmoid function has a total of 4 parameters. The adam optimiser was used to update all the 133 trainable parameters. The objective function is a binary cross-entropy loss function, that reduces the uncertainty of the distribution of upward or downward movements in stock prices. Using different input data such as fundamental company data, researchers have

TABLE 2. Experiment definitions.

Experiment number	Filter details	Sample size
1	No filter i.e., complete universe of analyst	29935
2	Filter on Resource stocks	10522
3	Filter on Financial stocks	7193
4	Filter on Industrial stocks	12220

TABLE 3. Accuracy score results.

Experiment	Accuracy				
	XGB	RF	LR	DNN	Consensus
1	75.50%	<b>79.02%</b>	60.06%	63.85%	54.20%
2	79.00%	<b>81.66%</b>	59.33%	63.32%	48.84%
3	77.00%	<b>78.04%</b>	66.01%	68.58%	57.33%
4	76.92%	<b>79.29%</b>	64.27%	66.32%	54.83%

TABLE 4. Precision score results.

Experiment	Precision Score				
	XGB	RF	LR	DNN	Consensus
1	75.24%	<b>79.26%</b>	58.64%	64.59%	50.59%
2	79.24%	<b>81.78%</b>	59.56%	61.64%	48.09%
3	75.36%	<b>77.34%</b>	62.91%	69.73%	52.69%
4	76.64%	<b>79.39%</b>	63.07%	69.14%	50.03%

shown that ANNs can successfully classify the directional movement of stock prices with 72% accuracy [29].

D. LOGISTIC REGRESSION (LR)

We modelled the bidirectional accuracy of the AI analyst using a logistic regression (LR). It is a parametric classification model based on a sigmoid function and its properties [30]–[32]. The input features described in section II-A are calculated as the weighted sum of the input features to generate probabilities between 0 and 1 using equation (15).

$$p(x) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_m x_m)}} \quad (15)$$

where  $p$  is the predicted probability of an observation belonging to a class,  $\beta_i$  is the parameter for  $i$ -th feature  $x_i$  and  $m$  is the number of input features.


To calibrate the model, the loss function  $l$  is the cross-entropy loss given by

$$l(D_x) = - \sum_i^{N_{obs}} Y_i \ln(p_i) + (1 - Y_i) \ln(1 - p_i) \quad (16)$$


where  $D_x$  represents the data and  $N_{obs}$  represents the number of observations,  $p_i$  represents the predicted probability of  $i$ -th observation. In this study, since there are 14 sell side report features,  $m = 14$ . These are the coefficients or parameters chosen such that they minimise the cross-entropy loss.

IV. EXPERIMENT RESULTS


Data normalisation was applied to obtain the input dataset at the same scale. The data is transformed to a range between  $-1$  and  $1$ . We begin with the entire universe after generating



(a) RF feature importance for experiment 1



(b) XGB feature importance for experiment 1



(c) LR feature importance for experiment 1


FIGURE 3. Feature importance for experiment 1.

the input score features. These start from 11 November 2005 to 29 June 2018 to give a total of 29935 analyst reports with their respective forecast accuracy scores. We run four experiments selecting between the complete analyst dataset and selecting a subset thereof. Table 2 outlines the experiments conducted as a subset of the complete universe.


To evaluate the performance of the model, the data were split at 80:20 for the training and validation sets. We used accuracy and precision scores to evaluate the performance of respective machine learning model in predicting directional movement of stock price. Table 3 The presents the results.

Next, we examined the performance of the model using the precision scores listed in Table 4.


Table 3 shows that machine learning algorithms have performed better than individual analyst predictions for directional movement. RF and XGB model also outperformed




(a) RF feature importance for experiment 2




(a) RF feature importance for experiment 3




(b) XGB feature importance for experiment 2



(b) XGB feature importance for experiment 3



(c) LR feature importance for experiment 2



(c) LR feature importance for experiment 3

FIGURE 4. Feature importance for experiment 2.


the deep neural networks. We believe that the reason for the outperformance of tree-based models is the nature of some of the categorical data as inputs to the feature inputs. We can also see from Table 3 that there is evidence that the accuracy of directional prediction increases with analyst specialization within sectors. RF showed the highest prediction accuracy of 81.78% in experiment 2, whereas the lowest accuracy score was in experiment 1. The DNN showed a similar differential performance between experiments. The LR showed the lowest performance among the machine learning models, with the highest accuracy score in experiment 3. Analyst consensus performance showed the highest accuracy in experiment 3, whereas the lowest performance was in experiment 2. We believe that this is because resource stocks are more difficult to forecast than financials because of the relatively higher volatility of resource stocks. However,

FIGURE 5. Feature importance for experiment 3.


it is important to note that the RF model performs best with resource stocks, indicating the benefits of using fused sell-side analysts' features. Table 4 shows precision score results which were similar to the respective accuracy scores across all the experiments.

### V. FEATURE IMPORTANCE


We explore the feature importance for XGB, RF and logistic regression when generating bidirectional stock price movement predictions. In figure 3, the feature importance highlights *Top40CSV*, *AC* and *P<sub>0</sub>* as strong features for directional prediction in experiment 1 under both XGB and RF while LR shows *MES* and *MAS* as important parameters for positive and negative class predictions respectively. Features in LR shows only firm, rating,  $\frac{TP}{P_0}$  and *AC* were statistically significant ( $\alpha = 5\%$ ). Ahmed et al. [11] also find that volatility



(a) RF feature importance for experiment 4



(b) XGB feature importance for experiment 4



(c) LR feature importance for experiment 4

FIGURE 6. Feature importance for experiment 4.

is a pivotal driver of directional accuracy. This is as expected because, given a specific mean return, volatility will drive the dynamics of whether the expected return is either negative or positive. We also found that analyst coverage is a vital feature in most experiments under both the RF and XGB models, but not under LR. The different results between experiments also show that the increase in industry coverage had spillover effects on analysts’ forecast accuracy [33], resulting in analyst coverage having a higher importance score, and, hence, higher prediction accuracy scores. In figure 4 and 5, the rolling score features are more important in experiments 2 and 3 under both XGB and RF. The rolling accuracy scores under XGB and RF show MAS and PAS with the highest feature importance in most experiments, whereas analysts’ stock ratings are the lowest in most experiments. This indicates that the framework to differentiate between analyst reports using

rolling features has added benefits. We expected the PTM feature to have a higher importance score in most experiments because the earnings momentum factor influences changes in analysts’ recommendations or target prices [34].

VI. CONCLUSION

There is increasing evidence that analysts’ reports are used in different investment processes because of their market influence. We conducted a study to build a generic fused AI analyst to generate directional predictions of stock price movements based on sell-side report outputs. The results show that a generic AI analyst performs better than an individual analyst directional prediction, with accuracy for generic fused AI analyst between 75%-79% and individual analyst accuracy between 54%-57%, similar to the findings of Ahmed et al. [11]. The precision score showed a similar performance.

Using machine learning algorithms, we have demonstrated that sell side report outputs including P<sub>0</sub>, AC and TP together with the rolling scores PAS, MAS and DAS are strong features for directional prediction of stock prices. It appears that the machine learning algorithms use P<sub>0</sub> to establish a reference point while using the other features to predict directional movement.

Possible investment strategies can be explored using AI analyst predictions to determine whether there are other factors or alpha contributions beyond standard factor returns within the FTSE/JSE all-share universe. We also find that features such as volatility and analyst coverage show heightened importance in all the experiments. It will be worthwhile to explore this relationship because one would suspect that analyst coverage increases accuracy in earnings or price predictions, and hence, lowers the volatility of the stock.

REFERENCES

- [1] S. Ramnath, S. Rock, and P. Shane, “The financial analyst forecasting literature: A taxonomy with suggestions for further research,” *Int. J. Forecasting*, vol. 24, no. 1, pp. 34–75, Jan. 2008.
- [2] K. K. Li and H. You, “What is the value of sell-side analysts? Evidence from coverage initiations and terminations,” *J. Accounting Econ.*, vol. 60, nos. 2–3, pp. 141–160, Nov. 2015.
- [3] T. Chen, J. Harford, and C. Lin, “Do analysts matter for governance? Evidence from natural experiments,” *J. Financial Econ.*, vol. 115, no. 2, pp. 383–410, Feb. 2015.
- [4] G. Niehaus and D. Zhang, “The impact of sell-side analyst research coverage on an affiliated broker’s market share of trading volume,” *J. Banking Finance*, vol. 34, no. 4, pp. 776–787, 2010.
- [5] M.-T. Chiang and M.-C. Lin, “Market sentiment and herding in analysts’ stock recommendations,” *North Amer. J. Econ. Finance*, vol. 48, pp. 48–64, Apr. 2019.
- [6] T. Sidogi, R. Mbuva, and T. Marwala, “Stock price prediction using sentiment analysis,” in *Proc. IEEE Int. Conf. Syst., Man, Cybern. (SMC)*, Oct. 2021, pp. 46–51.
- [7] C. Spence, M. Aleksanyan, Y. Millo, S. Imam, and S. Abhayawansa, “Earning the ‘write to speak’: Sell-side analysts and their struggle to be heard,” *Contemp. Accounting Res.*, vol. 36, no. 4, pp. 2635–2662, Dec. 2019.
- [8] E. F. Fama, “Efficient capital markets: A review of theory and empirical work,” *J. Finance*, vol. 25, no. 2, pp. 383–417, May 1970.
- [9] B. Barber, R. Lehavy, M. McNichols, and B. Trueman, “Can investors profit from the prophets? Security analyst recommendations and stock returns,” *J. Finance*, vol. 56, no. 2, pp. 531–563, Apr. 2001.



- [10] K. L. Womack, "Do brokerage analysts' recommendations have investment value?" *J. Finance*, vol. 51, no. 1, pp. 137–167, Mar. 1996.
- [11] S. Ahmed and E. Flint, "Evaluating equity analyst forecasts in South Africa," *Investment Anal. J.*, vol. 49, no. 3, pp. 199–214, Jul. 2020.
- [12] K. Kua and A. Ignjatovic, "Iterative filtering algorithms for computing consensus analyst estimates," in *Proc. IEEE Symp. Comput. Intell. Financial Eng. Econ. (CIFER)*, May 2022, pp. 1–8.
- [13] J. Singh and M. Khushi, "Feature learning for stock price prediction shows a significant role of analyst rating," *Appl. Syst. Innov.*, vol. 4, no. 1, p. 17, Mar. 2021.
- [14] A. Zaremba and P. Konieczka, "The profitability of following analyst recommendations on the Polish stock market," *Financial Internet Quart., e-Finanse*, vol. 11, no. 1, pp. 22–31, Mar. 2015.
- [15] S. J. Park and K. Y. Park, "Can investors profit from security analyst recommendations? New evidence on the value of consensus recommendations," *Finance Res. Lett.*, vol. 30, pp. 403–413, Sep. 2019.
- [16] R. K. Loh and G. M. Mian, "Do accurate earnings forecasts facilitate superior investment recommendations?" *J. Financial Econ.*, vol. 80, no. 2, pp. 455–483, May 2006.
- [17] K. Schipper, "Analysts' forecasts," *Accounting Horizons*, vol. 5, no. 4, p. 105, 1991.
- [18] M. B. Mikhail, B. R. Walther, X. Wang, and R. H. Willis, "Determinants of superior stock picking ability," 2004, pp. 1–33. [Online]. Available: [https://www4.gsb.columbia.edu/null?exclusive=filemgr.download&file\\_id=644981](https://www4.gsb.columbia.edu/null?exclusive=filemgr.download&file_id=644981)
- [19] M. B. Mikhail, B. R. Walther, and R. H. Willis, "Do security analysts exhibit persistent differences in stock picking ability?" *J. Financial Econ.*, vol. 74, no. 1, pp. 67–91, Oct. 2004.
- [20] J. L. Hall and P. B. Tacon, "Forecast accuracy and stock recommendations," *J. Contemp. Accounting Econ.*, vol. 6, no. 1, pp. 18–33, Jun. 2010.
- [21] B. Groysberg, P. M. Healy, and D. A. Maber, "What drives sell-side analyst compensation at high-status investment banks?" *J. Accounting Res.*, vol. 49, no. 4, pp. 969–1000, Sep. 2011.
- [22] A. Bakhach, E. P. K. Tsang, and H. Jalilian, "Forecasting directional changes in the FX markets," in *Proc. IEEE Symp. Ser. Comput. Intell. (SSCI)*, Dec. 2016, pp. 1–8.
- [23] J. Shen and M. O. Shafiq, "Short-term stock market price trend prediction using a comprehensive deep learning system," *J. Big Data*, vol. 7, no. 1, pp. 1–33, Dec. 2020.
- [24] S. Hutira, "Determinants of analyst forecasting accuracy," Joseph Wharton Scholars, 2016. [Online]. Available: [https://repository.upenn.edu/joseph\\_wharton\\_scholars/15](https://repository.upenn.edu/joseph_wharton_scholars/15)
- [25] S. Bonini, L. Zanetti, R. Bianchini, and A. Salvi, "Target price accuracy in equity research," *J. Bus. Finance Accounting*, vol. 37, nos. 9–10, pp. 1177–1217, Nov. 2010.
- [26] A. G. Kerl, "Target price accuracy," *Bus. Res.*, vol. 4, no. 1, pp. 74–96, Mar. 2011.
- [27] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
- [28] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," *Ann. Statist.*, vol. 29, no. 5, pp. 1189–1232, Oct. 2001.
- [29] L. Kryzanowski, M. Galler, and D. W. Wright, "Using artificial neural networks to pick stocks," *Financial Anal. J.*, vol. 49, no. 4, pp. 21–27, Jul. 1993.
- [30] A. F. Cabrera, "Logistic regression analysis in higher education: An applied perspective," in *Higher Education: Handbook for the Study of Higher Education*, vol. 10, J. C. Smart, Ed. New York, NY, USA: Agathon Press, 1994, pp. 225–256.
- [31] J. Aldrich and F. Nelson, "Logit and probit models for multivariate analysis with qualitative dependent variables," in *New Tools for Social Scientists*, W. D. Berry and M. S. Lewis-Beck, Eds. Newbury Park, CA, USA: Sage, 1986.
- [32] C.-Y. J. Peng, K. L. Lee, and G. M. Ingersoll, "An introduction to logistic regression analysis and reporting," *J. Educ. Res.*, vol. 96, no. 1, pp. 3–14, 2002.
- [33] K. Merkle, R. Michaely, and J. Pacelli, "Does the scope of the sell-side analyst industry matter? An examination of bias, accuracy, and information content of analyst reports," *J. Finance*, vol. 72, no. 3, pp. 1285–1334, Jun. 2017.
- [34] S. K. Mackinnon and R. Kruger, "Factors influencing changes in analyst consensus recommendations: Evidence from the Johannesburg stock exchange," *J. Appl. Bus. Res.*, vol. 30, no. 3, pp. 959–970, 2014.



**THENDO SIDOGI** was born in Venda, Limpopo, South Africa. He received the B.Sc. degree (Hons.) in mathematical techniques and the master's degree in mathematics from the University of Witwatersrand. He is currently pursuing the Ph.D. degree in machine learning in finance with the University of Johannesburg. He is also a Chartered Financial Analyst (CFA) Charterholder. His research interest includes the use of machine learning methods to improve quantitative investment methods.



**WILSON TSAKANE MONGWE** was born in Tembisa, Gauteng, South Africa. He received the B.Sc. degree in computing from the University of South Africa and the B.Bus.Sci. degree in actuarial science and the master's degree in mathematical finance from the University of Cape Town. His research interests include Bayesian machine learning and Markov chain Monte Carlo methods. He was a recipient of the Google Ph.D. fellowship in machine learning which supports his current Ph.D. research at the University of Johannesburg.



**RENDANI MBUVHA** was born in Venda, Limpopo, South Africa. He received the B.Sc. degree (Hons.) in actuarial science and statistics from the University of Cape Town and the M.Sc. degree in machine learning from the KTH Royal Institute of Technology, Sweden. He is currently a Senior Lecturer in statistics and actuarial science at the University of Witwatersrand. He is also a qualified Actuary and a Holder of the Chartered Enterprise Risk Actuary Designation. He was a recipient of the Google Ph.D. fellowship in machine learning which supported his Ph.D. research at the University of Johannesburg.



**TSHILIDZI MARWALA** (Senior Member, IEEE) was born in Duthuni, Venda, Limpopo, South Africa, in July 1971. He received the bachelor's degree in mechanical engineering from Case Western Reserve University, Cleveland, OH, USA, in 1995, the master's degree in mechanical engineering from the University of Pretoria, Pretoria, South Africa, in 1997, and the Ph.D. degree in engineering from the University of Cambridge (St. Johns College), Cambridge, U.K.,

in 2000. He has also completed other leadership courses from the Columbia Business School, National University of Singapore, GIBS University of Pretoria, Harvard Business School, and the University of South Africa. Some of his accomplishments include being a fellow of Cambridge Commonwealth Trust, in 1997, CSIR, in 2005, South African Academy of Engineering, in 2007, TWAS, The World Academy of Science, in 2010, African Academy of Science, in 2013, South African Institute of Electrical Engineers, in 2016. He is currently a registered Professional Engineer (Pr.Eng.) with the Engineering Council of South Africa. He is the also a Vice-Chancellor and the Principal of the University of Johannesburg. He has won many awards, including notably, the Bronze Order of Mapungubwe awarded by the President of the Republic of South Africa, in 2004, and the TWAS-AAS-Microsoft Award for Young Scientists, in 2009. His contributions in artificial intelligence field come in forms of over 15 books, he has authored, over 50 peer-reviewed chapters, over 50 journal publications, and over 150 conference publications. He also holds the Deputy Chairperson of the Presidential Fourth Industrial Revolution Commission of South Africa position and has contributed to over 50 articles to local press and journals on the subject of Fourth Industrial Revolution and Economics.

• • •