

Elsevier required licence: © <2019>. This manuscript version is made available under the CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>. The definitive publisher version is available online at [insert DOI]

Deep Learning for Decision Making and the Optimization of Socially Responsible Investments and Portfolio

Abstract

A socially responsible investment portfolio takes into consideration the environmental, social and governance aspects of companies. It has become an emerging topic for both financial investors and researchers recently. Traditional investment and portfolio theories, which are used for the optimization of financial investment portfolios, are inadequate for decision-making and the construction of an optimized socially responsible investment portfolio. In response to this problem, we introduced a Deep Responsible Investment Portfolio model that contains a Multivariate Bidirectional Long Short Term Memory neural network, to predict stock returns for the construction of a socially responsible investment portfolio. The deep reinforcement learning technique was adapted to retrain neural networks and rebalance the portfolio periodically. Our empirical data revealed that the DRIP framework could achieve competitive financial performance and better social impact compared to traditional portfolio models, sustainable indexes and funds.

Keywords: Socially Responsible Investment, Portfolio Optimization, Time Series Analytics, Deep Reinforcement Learning, Decision Support System

1. Introduction

Traditionally, investors have focused on the investment returns by actively looking at the financial reports to find the best performing stocks. With the recent mindset change towards sensitive topics like global warming or refugees, investors are becoming concerned with other aspects of companies rather than just earnings. They are shifting their

investment towards companies which are actively doing good things for the environment, contributing to the society and operating with transparency. According to the 2018 Biennial Report On US Sustainable, Responsible and Impact Investing Trends [51], socially responsible investment (SRI) assets accounted for \$12 trillion out of \$47 trillion in total assets under professional management in the United States in 2018, representing a sharp increase of 38% since 2016.

Conventional investment and portfolio theory focuses on financial performance, i.e., the returns and risks of the portfolio [55]. Direct application of the theory might not be suitable for SRI because it focuses more on non-monetary objectives [6]. Therefore, socially responsible investors need a modified version of the modern portfolio theory that can serve their purpose better [38]. Besides, SRI investors currently have to read corporate social responsibility (CSR) reports to find good companies to invest in, which is time-consuming and difficult. The lack of effective quantitative approaches for SRI makes it more difficult for not only professional investors, but also the vast majority of lay investors. Therefore, this research will provide an easy and automated way of doing such investment in an ethical manner, which greatly benefits their decision-making and secures the optimal investment returns. This is one of the main motivational purposes for this research.

Recently, the Global Reporting Initiatives (GRI) and the United Nation Sustainable Development Goals (SDGs) has provided standardized metrics and frameworks for companies to disclose more information regarding their sustainability practices [14]. For example, environmental, social and governance (ESG) metrics of companies have been derived from reports and news articles (e.g. CSR reports, news articles, carbon disclosure project ratings), evaluating the company in different prospects (e.g. air emissions and waste management, employee health and safety control, board transparency and diversity) including their controversies (e.g. involvement in adult entertainment or gambling). These metrics have been consolidated into the combined ESG ratings (see Fig. 1). The availability of ESG ratings has led to an emerging research topic in SRI portfolio.

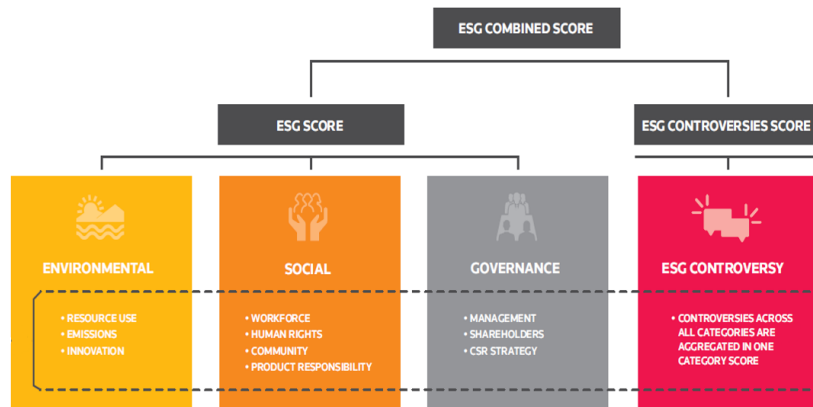


Figure 1: Combined ESG ratings [50]

With the availability of ESG metrics, quantitative methods can now be applied effectively to address the SRI portfolio construction problem. Current data mining approaches in this research field face a number of challenges. The first challenge is the accuracy of multivariate time series predictions. Stock return forecasts have been extensively studied with various quantitative finance and machine learning models [24, 44]. Most of these works have been focused on univariate time series predictions because it is expected that multivariate data would contain too much noise for the neural network to perform well [20, 33]. However, stock movements in the financial market are highly correlated; thus a multivariate model can learn these deep insights better than a combination of univariate networks. Following recent advances in neural networks research, especially Long Short Term Memory networks (LSTM) [26], the application of deep learning in the predictive investment field has become an alternative approach to the traditional financial model. In this paper, we propose a novel Multivariate Bidirectional LSTM neural network to predict multiple time series for stock returns.

The second challenge faced by current approaches in SRI is the application of multi-objective portfolio construction. Existing portfolio optimization methods are evolving around the standard Mean-Variance (MV) Portfolio [32], which focuses on maximizing returns and minimizing risks. To incorporate corporate responsibility performance

into our optimization problem, we introduced a modified MV model for SRI portfolio construction by integrating ESG ratings.

The third challenge for SRI portfolio is building a model that can adapt to market movements. As SRI in particular, and financial investment in general, are sensitive to market volatility, model parameters should be tuned up periodically to achieve both financial performance and ESG rating objectives. By adopting reinforcement learning techniques, we introduced a Deep Responsible Investment Portfolio (DRIP) model to retrain the prediction model and rebalance the portfolios effectively and autonomously.

An advantage of our proposed approach, which incorporates a multivariate BiLSTM neural network and MV-ESG, is that the framework can be generalized and extended to other scenarios with a similar multivariate prediction and multi-objective optimization problem. The developed deep reinforcement learning framework could also accommodate different neural networks and AI algorithms to tackle other types of complex and highly intercorrelated problems.

The main contributions of our research are:

- A novel, deep, responsible, investment portfolio framework to integrate deep neural networks, multi-objective optimization and reinforcement learning. The framework could be applied to other similar contexts of multivariate predictive analytics.
- A novel DRIP model that can forecast the returns quarterly and yearly on investment instead of just daily, which is a more realistic scenario for investors. The model has been fully tested and deployed on real-life datasets containing 100 stocks over a period of 30 years.
- The first report (to the best of our knowledge) leverages deep learning and incorporates ESG ratings into a portfolio optimization model.

This paper is organized as follows. In Section 1, we introduce the current background of socially responsible investment portfolios and their limitations, as well as

give our motivation for our proposed reinforcement learning solution. In Section 2, we review the literature on socially responsible investment research, and the application of deep learning by focusing on recent methodologies that are closely related to this paper. In Section 3, we introduce the technical details of our DRIP model and in Section 4, we present our empirical studies that applied our novel algorithms to real-life financial datasets to construct socially responsible investment portfolios, evaluate the performance against some baseline models and explore the potential for further and related research. We provide our conclusions in Section 5.

2. Literature Review

2.1. Socially Responsible Investment

The optimization of financial portfolios has been researched extensively. Many approaches have been developed to build decision support systems for stock trading. This includes standard mathematical finance modeling, e.g. Mean-Variance (MV) [32], AutoRegressive Moving Average (ARMA) and Generalized AutoRegressive Conditional Heteroskedasticity (GARCH) models [18], text mining of financial news [35] and social media [25].

However, limited research has been carried out on socially responsible investment. Although the socially responsible investment was proposed in the 1980s [22], it only became a topic of interest for academia and industry in the past decade [15]. During this time, research has correlated ESG ratings with the financial performance of companies [23, 16] or socially responsible funds [28, 34, 2]. The availability of environmental, social and governance (ESG) ratings has enabled more research and application in this area in academia [54] and industry [51].

Many sustainability funds have offered portfolios with certain values to attract investors to SRI. In management funds [46], there has been an increasing demand from

sustainably conscious investors to have more SRI options [37]. Multiple sustainable indexes and funds have been constructed based on areas of investor interests (e.g. water treatment, clean tech, renewable energy, gender equality and diversity). The literature has shown that companies or sustainable funds with higher ESG ratings can outperform the lower ones financially in long-term investments [19].

The literature of qualitative research in SRI has focused on reviewing the performance of companies [4] and socially responsible indexes or funds [49], and not on a data-driven approach to incorporate sustainability into an investment system. Some of the research has criticized the current stock screening process of SRI funds [53] and has proposed that the full integration of ESG ratings would be more beneficial [1]. These findings underpin the main motivation for our research to develop a framework with full integration of ESG ratings. Our paper contributes to the current knowledge of the application of deep learning for the prediction of stock returns and ESG-based SRI portfolio optimization.

2.2. Deep Learning for Stock Returns Forecasting

Researchers have undertaken extensive study to solve the time series forecasting problem of stock returns using deep learning [13, 33, 9]. Many have suggested that different types of Recurrent Neural Networks (RNN) outperform traditional financial time series models in different markets [7, 3, 44]. RNN contains feedback loops in its recurrent layer, which enables the storage of information in the “memory cell” over time. However, it does not perform well when the learning requires long-term temporal dependencies.

Long Short-Term Memory (LSTM) is a special type of RNN that has been proven to be effective in text mining to predict stock returns [30]. LSTM contains “memory cells” that are able to retain information for longer periods of time [26]. Consequently, LSTM often performs better in sequential data and financial time series predictions compared with RNN [36, 27], particularly in the SRI context where investors are concerned more

about long-term returns rather than the volatility of the short-term market.

Researchers have also compared the performance of different RNN architectures like LSTM and Gated Recurrent Unit (GRU) networks [40]. Others have suggested that Bi-directional LSTM (BiLSTM) might be a better option in a similar sequence prediction problem [8]. While the LSTM and GRU, with the unidirectional flow of information, might be adequate in most sequence prediction problems, the BiLSTM model reads the data one more time backward [42] which helps improve prediction accuracy, particularly in forecasting sequential data like financial time series.

Recently it was suggested that back-testing results could have given rise to false positives due to the normalization of testing data and prediction of the next time step only [43]. The next-time step prediction is only suitable for high-frequency trading strategies using intra-daily data, such as foreign exchange markets. In SRI, investors are more interested in long-term returns on investment. Conversely, research has been conducted on the long-term prediction for financial indexes with 1-year and 2-year time gaps, suggesting that long-term forecasting is possible for stock returns [17].

Our paper contributes to the current deep learning methodologies through the design of a novel BiLSTM neural network that predicts a long-term multivariate time series. To avoid false positive results, the financial returns data is not normalized and the model predicts multiple steps ahead. By constructing the baseline models using different types of LSTM networks as undertaken previously, we evaluate the prediction accuracy of the LSTM networks in the forecasting of SRI stock returns.

2.3. Portfolio Optimization

Few socially responsible investment models have been developed and proposed that utilize ESG ratings [52]. [21], for example, suggested a modification to the standard portfolio selection model with ESG scores. They utilized the Mean-Variance Stochastic Goal Programming (MV-SGP) model with a statistical approach for ESG screening on stocks based on scores and controversy risk. However, they did not consider predictive

analytics; they only used past returns and volatility to test their hypotheses. Furthermore, they did not validate their models with real financial data.

Multiple optimization functions are available, including the Expectation Maximization (EM) algorithm [12], quasi-Newton [5] or Powell methods [39]. However, most of them are not multi-objective or allow the special limit conditions that are required in a complex context like in socially responsible investment. The Sequential Least Squares Programming (SLSQP) method proposed by [29], for example, can be used to minimize a function of various variables with a different combination of bounds, equality and inequality constraints. However, its greedy behavior leads to a skewed distribution for the weights of stocks in the portfolio. This is not an optimum choice for investors who are worried about non-diversified portfolios with extreme exposure risk.

We have developed a financial model to construct a socially responsible investment portfolio that incorporates the Mean-Variance portfolio theory and ESG ratings (MV-ESG). Our model is not based on the ESG screening approach. Instead, it filters and leverages the ESG ratings in a multi-objective optimization function based on the SLSQP method. It also considers both past and predicted the future performance of stocks in a portfolio selection. This is one of the first mathematical models for constructing a socially responsible investment portfolio that achieves both better ESG ratings and competitive financial performance.

3. Methodology

Our DRIP framework consists of three main components: a multivariate BiLSTM neural network to predict stock returns quarterly and yearly; these predicted values are then combined with ESG ratings in our MVP- ESG model for portfolio construction; reinforcement learning techniques are then leveraged to automatically retrain the prediction models and re-balance our MVP-ESG portfolios after each period. The full rein-

forcement learning DRIP framework is as shown in Figure 4.

3.1. Multivariate BiLSTM neural networks

Standard feature engineering often includes normalization step, which transforms the data range to $[0,1]$. This common approach can help to improve the prediction accuracy of the neural networks. However, in the time series model, this approach implicitly tells the trained model the movement range of future stock prices, which makes out-of-bag testing results unrealistically accurate. We processed the input data for our neural networks in a different approach. In our DRIP model, we did not normalize data but instead fed the stock returns directly into the neural networks. We also trained the model to predict values with a longer time gap instead of a next period prediction, which is a more suitable scenario for stock investors in real-life trading.

Let $p_i(t)$ be the price at time t ($t = 1, \dots, T$) for stock i ($i = 1, \dots, N$). Δt was the time gap ($1 < \Delta t < T$). The return $r_i(t)$ for stock i at time t was $r_i(t) = p_i(t) - p_i(t - \Delta t)$. In the DRIP model, we used the sliding window technique to perform rolling forecast. Let δt be the sliding window size. The train features matrix $X_i(t)$ and return vector $Y_i(t)$ for stock i at time t were:

$$X_i(t) = \begin{bmatrix} r_i(t - T - \delta t) & r_i(t - T - \delta t + 1) & \cdots & r_i(t - T) \\ r_i(t - T + 1 - \delta t) & r_i(t - T + 1 - \delta t + 1) & \cdots & r_i(t - T + 1) \\ \cdots & \cdots & \cdots & \cdots \\ r_i(t - \delta t) & r_i(t - \delta t + 1) & \cdots & r_i(t) \end{bmatrix} \quad (1)$$

$$Y = \begin{bmatrix} r_i(t - T + \Delta t) \\ r_i(t - T + 1\Delta t) \\ \cdots \\ r_i(t + \Delta t) \end{bmatrix} \quad (2)$$

As suggested by [36], LSTM networks would outperform other neural networks in solving similar problems due to its information persistence characteristic. We considered three types of LSTM neural networks:

- **LSTM**, initially proposed by [26], is a special kind of RNN, which is capable of learning long-term dependencies. For each input vector x_t at time step t , LSTM networks uses multiple gating functions: the input gate i_t , forget gate f_t , and output gate o_t , together with a memory cell C_t to preserve long-term information and keeps track of its flow. The forget gate f_t and input gate i_t generated at each time step t are defined as follows:

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \quad (3)$$

$$i_t = \sigma(W_i[h_{t-1}, x_t] + b_i) \quad (4)$$

In the next step, a tanh layer generates a new memory cell \tilde{C}_t . LSTM then updates the old memory cell C_t and generates the output gate o_t and hidden state h_t :

$$i_t = \tanh(W_C[h_{t-1}, x_t] + b_C) \quad (5)$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \quad (6)$$

$$o_t = \sigma(W_o[h_{t-1}, x_t] + b_o) \quad (7)$$

$$h_t = o_t \odot \tanh(C_t) \quad (8)$$

where σ is the sigmoid function and \odot is the element-wise multiplication. W is the weight matrix and b is the bias vector to be learned by the LSTM at each specific gate.

- **BiLSTM** is a variation of the bidirectional RNN, firstly introduced by [42]. It concatenates a forward and backward unidirectional LSTM on the stock return

time series $Combined(h_t) = [\vec{h}_t, \overleftarrow{h}_t]$. Unidirectional LSTM only preserves long-term information of the past, while BiLSTM can preserve information from both past and future by using the combined two hidden states $Combined(h_t)$.

- **GRU** is a more recent alteration of LSTM, suggested by [10]. It concatenates both the forget gate f_t and input gate i_t into a single update gate z_t , and merges the cell state C_t and hidden state h_t . The architecture of GRU is simpler than the standard LSTM one. The hidden state h_t generated at each time step t is defined as follows:

$$z_t = \sigma(W_z \cdot [h_{t-1}, x_t]) \quad (9)$$

$$r_t = \sigma(W_r \cdot [h_{t-1}, x_t]) \quad (10)$$

$$\tilde{h}_t = \tanh(W \cdot [r_t \odot h_{t-1}, x_t]) \quad (11)$$

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t \quad (12)$$

By simplifying the architecture of the LSTM, GRU may learn the data at the combined gate. However, this single update gate might not learn some hidden information effectively. Hence, the performance of GRU networks may be less effective in forecasting long-term time series.

For our DRIP model, we designed a special type of BiLSTM to perform multivariate time series prediction. The data input shape for the multivariate BiLSTM neural network was in the form of a three-dimensional matrix with sizes $(T - \Delta t - \delta t, \delta t, N)$, where N was the number of stocks in total, δt was the sliding window size, and Δt was the prediction time gap (see Figure 4).

We also replicated neural network models with LSTM and GRU networks as in [40] to predict returns for every single stock in the portfolio. We constructed the neural networks with recurrent layers using the Adam optimizer from the ‘‘Keras’’ package [11]. This network also contained a dense layer and a final output layer with the ‘‘linear’’ activation function to predict the stock returns $r_i(t + \Delta t)$ in Δt periods of time.

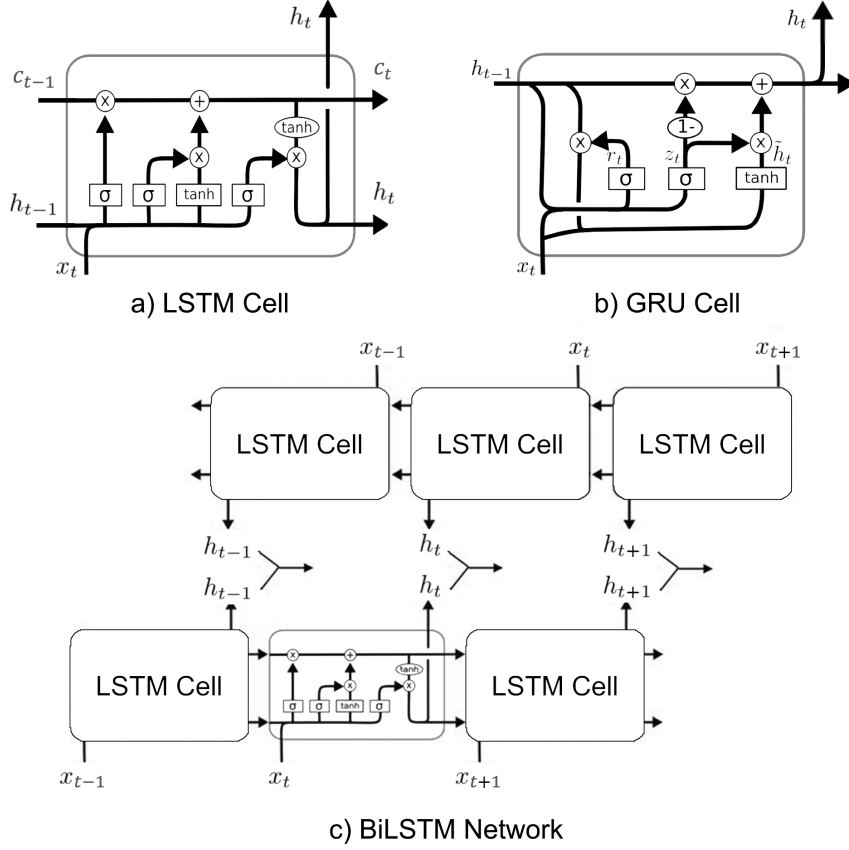


Figure 2: Graphical illustration of LSTM, GRU and BiLSTM

3.2. MV-ESG Model

The Mean Variance portfolio (MV) of [32] has always been the standard portfolio selection model. Its mathematical principle is constructed by two main components: maximizing the return r_p and minimizing the risk σ_p . The output of this optimization process is the efficient frontier, which is a set of investment portfolios with a greater return than any other with the same or less risk, and a lower risk than any other with the same or greater return. For illustration, the efficient frontier is plotted in Figure 3 with the risk on the horizontal axis and the return on the vertical axis.

The optimal portfolio based on the efficient frontier is commonly known as the max-

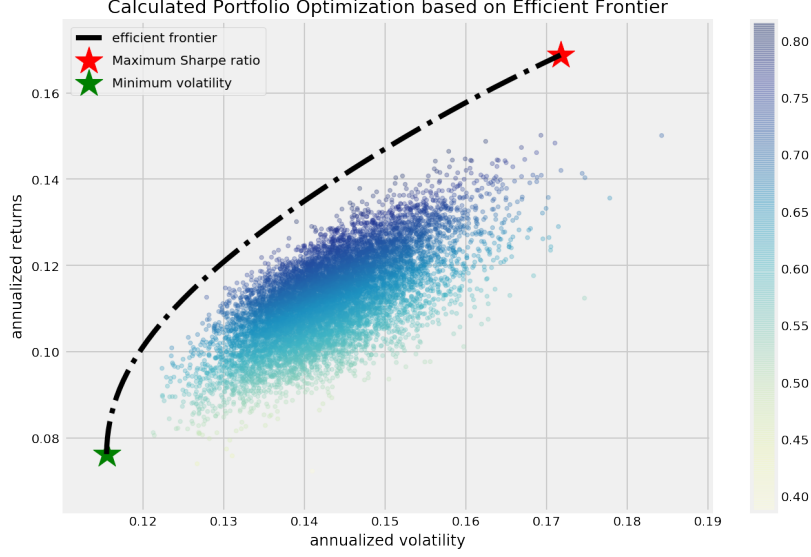


Figure 3: Standard MV Portfolio with Efficient Frontier

imum Sharpe portfolio (MAX-S), where the portfolio has a maximum Sharpe ratio calculated as $S_p = (r_p - r_f)/\sigma_p$. For the MAX-S portfolio, considering the risk free rate r_f (normally the return on bond investment or the bank interest rate), it minimizes the negative Sharpe Ratio [45]:

$$\min(-S_p) = \min\left(-\frac{r_p - r_f}{\sigma_p}\right) \quad (13)$$

$$r_p = \sum_{i=1}^N w_i r_i \quad (14)$$

$$\sigma_p = \sum_{i=1}^N \sum_{j=1}^N w_i \sigma_{ij} w_j \quad (15)$$

where w_i and w_j are the weights of stock i and j , with the boundary limit $w_i, w_j \in [0, 1]$, and σ_{ij} is the covariance matrix of the two stock i and j in the portfolio. The initial weight of each stock in the computation will be equally allocated according to the total number of stocks N in the portfolio, $w_i(0) = w_j(0) = 1/N$.

In our MV-ESG model, we built a multi-objective algorithm based on the SLSQP method [29] with three objectives: maximizing returns, minimizing volatility and maximizing ESG ratings. This algorithm minimized: $\min_{w_{esg}} \| -G \|$ with G being a three-dimensional matrix of constraints of the three objectives and w_{esg_i} being the ESG weights subject to boundary limits inferred from the companies' ESG ratings.

For comparison, we constructed a maximum ESG portfolio (MAX-ESG) for investors with low risk averse to compare with the standard MAX-S portfolio. In MAX-ESG, we minimized the negative Sharpe Ratio with the portfolio ESG ratings (ESG_p) as a new variable of the objective function.

$$\min(-\tilde{S}_p) = \min(-ESG_p \frac{r_p - r_f}{\sigma_p}) \quad (16)$$

$$ESG_p = \sum_{i=1}^N w_{esg_i} \frac{ESG_i + E\bar{S}G_i}{2} \quad (17)$$

where ESG_i was the combined ESG ratings of company i in the past year, $E\bar{S}G_i$ was the combined ESG ratings at the current prediction year, and w_{esg_i} was the ESG weight of stock i in the portfolio.

In the traditional MV model, r_p and σ_p are the past returns r_i and volatility σ_i , which is often called ex-post MV. In recent years, researchers and investors have been using the expected returns \bar{r}_i and volatility $\bar{\sigma}_i$. This approach called ex-ante MV is more suitable for predictive analytics in real-world financial trading. In our MV-ESG model, we combined both ex-post MV and ex-ante MV for portfolio selection and replaced the standard weight boundary with our ESG ones calculated based on the combined ESG ratings for each stock. Our MV-ESG model was computed using:

$$r_p = \sum_{i=1}^N w_{esg_i} \frac{r_i + \bar{r}_i}{2} \quad (18)$$

$$\sigma_p = \sum_{i=1}^N \sum_{j=1}^N w_{esg_i} \frac{\sigma_{ij} + \bar{\sigma}_{ij}}{2} w_{esg_j} \quad (19)$$

where r_i and \bar{r}_i were the ex-post and ex-ante returns, σ_{ij} and $\bar{\sigma}_{ij}$ were the ex-post and ex-ante covariance matrix of the two stock i and j in the portfolio. $wesg_i$ and $wesg_j$ were the ESG weight of stock i and j in the portfolio, with the boundary limit $wesg_1, \in [0, 1]$ for the company with the highest combined ESG score, then gradually decreasing to $wesg_N \in [0, 0]$ for the company with the lowest combined ESG score. This means the allocation of the company “ N ” in the portfolio was zero, indicating no investment in that company. The initial weight of each stock in the computation was not be equally allocated but assigned according to the ESG ratings.

3.3. DRIP Model with Reinforcement Learning

We combined the multivariate BiLSTM neural networks and the MV-ESG models into a single integrated reinforcement learning model named Deep Responsible Investment Portfolio (DRIP). Starting with a set of agent states S and a set of possible portfolio allocation sets A , we had the probability of the DRIP model select the specific portfolio allocation (the “action”) a when in state s at time step t as:

$$\pi : S \times A \rightarrow [0, 1] \quad (20)$$

$$\pi(a|s) : Pr(a_t = a|s_t = s) \quad (21)$$

We defined a simple state-value function V_π^s as the expected reward starting with the state $s_0 = s$ and Re_t denoting the reward function calculated as the sum of future discounted rewards:

$$V_\pi^s = E[Re] = E\left[\sum_{t=0}^{\infty} \gamma^t Re_t | s_0 = s\right] \quad (22)$$

$$Re = \sum_{t=0}^{\infty} \gamma^t \tilde{S}_t (1/MS E_t) \quad (23)$$

where $\gamma \in [0, 1]$ was the discount rate. \tilde{S}_t was the ESG-adjusted Sharpe Ratio [45], and $MS E$ was mean squared error of the prediction model. The DRIP model found a set of portfolio allocation to maximize the expected return.

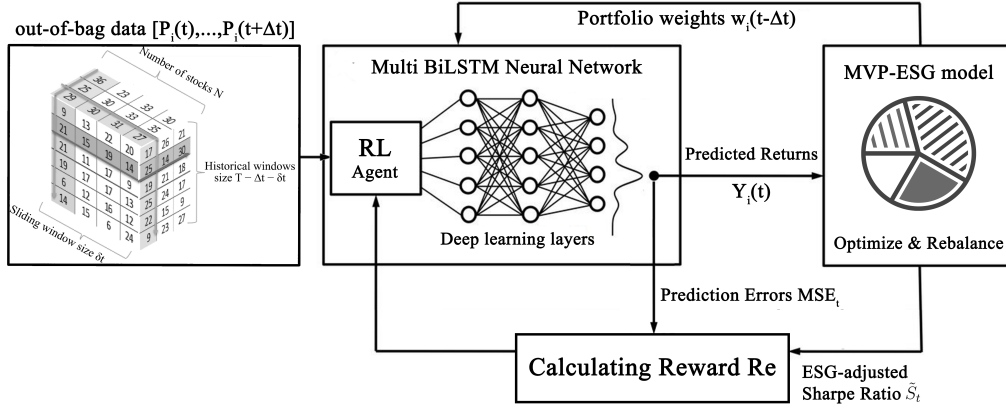


Figure 4: Reinforcement Learning DRIP Model

After each time gap Δt , DRIP retrained the prediction model with new stock prices data and then, together with the portfolio performance and stock weights from the previous period of time, constructed a new portfolio with updated allocation weights. The reinforcement learning was repeated on a predefined period basis, to improve the prediction model accuracy and the performance of the portfolio over time. The design of DRIP enabled its self-learning with the least human involvement as possible. The reinforcement learning model is shown in Figure 4.

4. Empirical Experiment

We designed experiments to test our proposed model in two ways: 1) DRIP model forecasts for quarterly and yearly returns of multivariate stock time series during the three year period from 2016 to 2018; and 2) Socially Responsible Portfolios optimization using the predicted returns and reinforcement learning DRIP model framework.

4.1. Datasets

Currently, there are various ESG rating services available [41], many of which offer a subscription fee for data access which limits its availability to the public. In 2018

however, Yahoo Finance made some of the ESG ratings obtained from Sustainalytics [48] available publicly. In this research, we utilized Yahoo Finance to obtain both financial stock prices and public ESG rating datasets in our reinforcement learning DRIP framework.

We downloaded the daily closing prices of all stocks in the Standard and Poor 500 list (S&P500) from the past 30 years from 31 December 1988 to 31 December 2018. In order to ensure a sufficient number of data points, we removed all the stocks which did not have a market price on 31 December 1988, which left us with 262 companies. SRI investors do not invest in companies with low ESG ratings; therefore we used a simple stock screening process to remove these unwanted stocks. From the shortlisted 262 stocks, we selected the top 100 companies with the highest combined ESG ratings according to Sustainalytics to construct the final dataset that contained a total of 756,000 data points.

We separated the train and test datasets using an out-of-bag approach, which excludes the testing period data from the past historical data at time t to avoid feeding the model any unknown future information. Our data splitting ratio is 9:1, which meant that the training data was from the year 1989 to 2015 for each stock, and the testing data was the three-year period from the year 2016 to 2018. We also adopt the rolling forecast approach to further split the data in the testing period into validation and test sets. For quarterly return prediction, we used “Q4/2015” and “Q1/2016” as the validation and test set for the first period. We then moved to the next quarter period until “Q3/2018” and “Q4/2018” as validation and test sets in the last period. We applied the same data splitting process to the yearly return prediction dataset.

4.2. Evaluation Metrics

To test our DRIP model, we used the Mean Absolute Error (MAE) and the Root Mean Squared Error (RMSE) as the evaluation metrics for the absolute value prediction:

$$\text{MAE} = 1/N \sum_{i=1}^N |\tilde{r}_i - r_i| \quad (24)$$

$$\text{RMSE} = \sqrt{1/N \sum_{i=1}^N (\tilde{r}_i - r_i)^2} \quad (25)$$

where $N = 100$ was the total number of stocks in the portfolio and \tilde{r}_i and r_i were the predicted and actual return of stock i for that period.

We also converted the predicted value to a binary label to evaluate the performance of uptrend or downtrend forecast using the prediction accuracy metric and the Area Under the Curve (AUC) scores with the Receiver Operating Characteristic (ROC) curve. The lower MAE and RSME together with the higher prediction accuracy and AUC scores indicate the better performance of the prediction model. Our baseline models for comparison are the LSTM and GRU neural networks as in [40] and a univariate standard BiLSTM model (Uni).

To evaluate the performance of our socially responsible portfolios using MV-ESG model, we compared its Sharpe Ratio against those of the standard MV portfolios and the reported financial performance from similar sustainable indexes and funds. The Sharpe Ratio was defined as $S = (r_p - r_f)/\sigma_p$ where r_p was the portfolio annualized return, σ_p was the portfolio annualized volatility, and $r_f = 2\%$ was the nominal risk-free rate. A better performing portfolio had a higher Sharpe Ratio, which yielded higher returns if the risks were similar or a lower risk if the returns were the same.

The sustainable indexes for comparison were: Dow Jones Sustainability World Index (DJSI World), Dow Jones Sustainability World Diversified Select Index (DJSI WD), and S&P500 ESG Factor Weighted Index (S&P500 ESG). All indexes data were obtained on the 31 December 2018 from S&P Dow Jones Indices, a division of S&P Global. The sustainable exchange traded funds (ETF) with their symbol codes in the brackets were: iShares Global Clean Energy ETF (ICLN), Invesco Solar ETF (TAN), iShares MSCI USA ESG Select ETF (SUSA) and Workplace Equality Portfolio (EQLT). All funds data were obtained on the 31 December 2018 from Morningstar.

4.3. Prediction Model Results

First of all, we tested the performance of DRIP model on the prediction of quarterly returns. The hyperparameters in our neural networks were set as: number of units in the deep learning layers equaled to 100, batch size equaled to 1, the loss was the mean squared error and random seed equaled 0. We used Adam optimizer with learning rate $lr = 0.001$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, fuzz factor $\epsilon = 1e - 7$, and decay equaled to 0. We also set the number of epochs equaled to 10 with a checkpoint after each epoch and only saved the best model for prediction.

Our experiment setup was as follows: $\Delta t = 63$, $\delta t = 63$ and the time gap was set to 63 representing the total number of trading days in a quarter. This meant that the model predicted the prices and returns three months ahead in time. After each period, the model was retrained and validated with the out-of-bag 3 month data and predicted the next return in 63 days. The testing data for each quarter of each year from 2016 to 2018 were referred to as “Q1”, “Q2”, “Q3” and “Q4” respectively.

We then tested the performance on the prediction of yearly returns with $\Delta t = 252$, $\delta t = 252$ representing the 252 trading days in a typical year. The other setup was the same as in the quarterly returns prediction model. The empirical results in Table 1 showed the performance evaluation for the quarterly and yearly returns prediction models using the multivariate financial time series as input. The reported RMSE and AUC Scores were averages for 100 stocks in each time period.

Our DRIP models, which used multivariate financial returns as the input significantly outperformed the other baseline models for most prediction periods in term of MAE and RMSE. We can conclude that the prediction model using multivariate financial returns and BiLSTM neural networks in our design was a better solution for this predictive analytic problem. Focusing on the trend prediction accuracy, except for the slightly worse results in “Q2/16” and “Q3/18”, our DRIP models that used BiLSTM achieved higher prediction accuracy and AUC scores regardless of the time periods or of the quarterly or yearly returns. These results demonstrated the effectiveness of our approach, that the

Table 1: DRIP Model Evaluation

| Mean Absolute Error (MAE) | | | | | | | | | | | | | | | |
|-----------------------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| | Q1/16 | Q2/16 | Q3/16 | Q4/16 | Q1/17 | Q2/17 | Q3/17 | Q4/17 | Q1/18 | Q2/18 | Q3/18 | Q4/18 | 2016 | 2017 | 2018 |
| DRIP | 0.0547 | 0.0880 | 0.0578 | 0.0814 | 0.0600 | 0.0551 | 0.0555 | 0.0601 | 0.0692 | 0.0667 | 0.0771 | 0.0948 | 0.0754 | 0.0830 | 0.1017 |
| Uni | 0.0656 | 0.0930 | 0.0678 | 0.0999 | 0.0633 | 0.0704 | 0.0679 | 0.0683 | 0.0714 | 0.0882 | 0.0854 | 0.1096 | 0.0821 | 0.0974 | 0.1157 |
| LSTM | 0.0771 | 0.0960 | 0.0734 | 0.1043 | 0.0667 | 0.0767 | 0.0707 | 0.0718 | 0.0650 | 0.0986 | 0.0788 | 0.1404 | 0.0889 | 0.1074 | 0.1122 |
| GRU | 0.0652 | 0.0950 | 0.0723 | 0.1139 | 0.0623 | 0.0793 | 0.0776 | 0.0729 | 0.0801 | 0.0993 | 0.1003 | 0.0936 | 0.0819 | 0.1018 | 0.1332 |
| Root Mean Squared Error (RSME) | | | | | | | | | | | | | | | |
| | Q1/16 | Q2/16 | Q3/16 | Q4/16 | Q1/17 | Q2/17 | Q3/17 | Q4/17 | Q1/18 | Q2/18 | Q3/18 | Q4/18 | 2016 | 2017 | 2018 |
| DRIP | 0.0672 | 0.1080 | 0.0754 | 0.1209 | 0.0768 | 0.0741 | 0.0750 | 0.0790 | 0.0858 | 0.0843 | 0.1202 | 0.1165 | 0.1014 | 0.1062 | 0.1273 |
| Uni | 0.0822 | 0.1189 | 0.0863 | 0.1534 | 0.0806 | 0.0933 | 0.0887 | 0.0939 | 0.0926 | 0.1120 | 0.1333 | 0.1402 | 0.1135 | 0.1241 | 0.1461 |
| LSTM | 0.0957 | 0.1225 | 0.0911 | 0.1637 | 0.0847 | 0.0991 | 0.0944 | 0.1029 | 0.0849 | 0.1233 | 0.1238 | 0.1739 | 0.1287 | 0.1355 | 0.1425 |
| GRU | 0.0812 | 0.1254 | 0.0914 | 0.1708 | 0.0801 | 0.1041 | 0.0952 | 0.0981 | 0.1057 | 0.1238 | 0.1536 | 0.1191 | 0.1086 | 0.1286 | 0.1659 |
| Area Under the Curve (AUC) Scores | | | | | | | | | | | | | | | |
| | Q1/16 | Q2/16 | Q3/16 | Q4/16 | Q1/17 | Q2/17 | Q3/17 | Q4/17 | Q1/18 | Q2/18 | Q3/18 | Q4/18 | 2016 | 2017 | 2018 |
| DRIP | 0.8392 | 0.5809 | 0.8387 | 0.8286 | 0.8495 | 0.7659 | 0.7521 | 0.8170 | 0.8165 | 0.8045 | 0.7024 | 0.9407 | 0.9525 | 0.9546 | 0.8989 |
| Uni | 0.7443 | 0.5803 | 0.7876 | 0.6603 | 0.8522 | 0.6052 | 0.6297 | 0.7486 | 0.7441 | 0.6477 | 0.6952 | 0.8502 | 0.9339 | 0.9397 | 0.8899 |
| LSTM | 0.7115 | 0.5946 | 0.7719 | 0.6320 | 0.8790 | 0.5069 | 0.5739 | 0.7001 | 0.7516 | 0.5692 | 0.7549 | 0.7278 | 0.9348 | 0.9106 | 0.8989 |
| GRU | 0.6821 | 0.5954 | 0.7522 | 0.5202 | 0.8281 | 0.5427 | 0.5632 | 0.7287 | 0.6642 | 0.5692 | 0.6884 | 0.8820 | 0.9142 | 0.9537 | 0.8720 |
| Trend Prediction Accuracy (%) | | | | | | | | | | | | | | | |
| | Q1/16 | Q2/16 | Q3/16 | Q4/16 | Q1/17 | Q2/17 | Q3/17 | Q4/17 | Q1/18 | Q2/18 | Q3/18 | Q4/18 | 2016 | 2017 | 2018 |
| DRIP | 77% | 50% | 74% | 76% | 80% | 66% | 70% | 77% | 72% | 76% | 73% | 85% | 92% | 92% | 80% |
| Uni | 69% | 50% | 73% | 62% | 77% | 58% | 65% | 75% | 65% | 64% | 70% | 76% | 91% | 91% | 78% |
| LSTM | 64% | 53% | 76% | 61% | 75% | 52% | 66% | 75% | 66% | 57% | 74% | 60% | 92% | 85% | 79% |
| GRU | 67% | 57% | 70% | 48% | 75% | 56% | 58% | 73% | 58% | 58% | 62% | 84% | 92% | 92% | 73% |

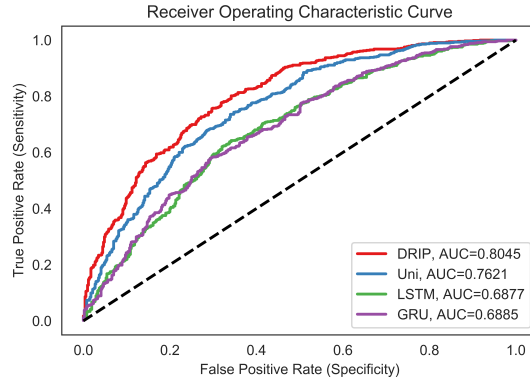


Figure 5: ROC curves

reinforcement learning had successfully captured the underlying hidden information in the inter-correlated multivariate series and improved itself over time.

The value predictions of quarterly returns generally had lower MAE and RSME than the yearly forecast. This result was expected as the time gap was smaller; hence, it was easier to forecast the absolute stock return values. The ROC curves in Figure 5 showed a performance lift in the quarterly returns prediction model compared to other baselines for the entire 3 year testing period. Conversely, the trend prediction was more accurate in yearly return models, which proved that the reinforcement learning model could filter out the market noise in short term price changes. Overall, our DRIP model effectively and accurately predicted the annual returns in all three years and the quarterly returns in 10 out of 12 testing periods. It showed that our prediction model was not over-fitted to a certain dataset period, and it could be generalized for similar applications.

4.4. Robustness Test

To test the robustness of our model, we first benchmarked the prediction model using different combinations of the neural network hyperparameters. We split the dataset into train, validation and test sets with the ratio 8:1:1. The hyperparameter sets were: number of units in the deep learning layers was in [100, 200, 300], batch size (BS) was

in [1, 10, 20] and learning rate (LR) was in [0.0001,0.001,0.01] accordingly. The MAE and prediction accuracy results for both validation and test set are presented in Table 2.

Table 2 showed that different hyperparameter sets could result in varied MAE and prediction accuracy. The gap between validation and test results are not significantly large, which indicates that our model was not overfitted. Our setting to generate the best results in the test set was: number of units equaled 100, batch size equaled 1 and learning rate equaled 0.01. In our rolling forecast and reinforcement learning model, the hyperparameters could be automatically tuned using grid search after each period.

We then used this set of hyperparameters to test the prediction model on three different datasets with 50, 100 and 200 randomly selected stocks (denoted as “Random50”, “Random100”, “Random200”). We also split these datasets into train, validation and test sets with the ratio 8:1:1. The results for this experiment are presented in Table 3.

Table 3 showed that our model still achieved a good prediction accuracy in randomly selected stock datasets. It is worth noticed that the MAE and the prediction accuracy are not worsen for the larger dataset but varied due to the randomness of stock selection. These results indicated that our prediction model is robust and generalizable with different data sizes.

4.5. Portfolio Optimization Model Results

We used the predicted returns from the DRIP model as input for our MV-ESG model to construct socially responsible investment portfolios. We constructed the MAX-ESG portfolios using predicted returns. The nominal risk free rate was set to 2%, $r_f = 0.02$. After obtaining the stock allocation in each portfolio, we calculated the actual annualized returns and volatility using real stock prices for that period. The annualized returns, volatility, Sharpe Ratio and ESG Score given in Table 4 were averaged for the entire year, for each year in the testing period.

The results showed that our MAX-ESG portfolios had consistently higher ESG ratings (3 to 5 points above). Even though the MAX-S portfolios had better financial returns

Table 2: Benchmarking prediction model with multiple hyperparameters

| Units | BS | LR | Validation set | | Test Set | |
|-------|----|--------|----------------|----------------|----------------|----------------|
| | | | MAE | Accuracy | MAE | Accuracy |
| 100 | 1 | 0.001 | 0.05468 | 0.89635 | 0.09335 | 0.86111 |
| 100 | 20 | 0.001 | 0.04906 | 0.93016 | 0.09682 | 0.83270 |
| 300 | 10 | 0.001 | 0.04515 | 0.94413 | 0.09705 | 0.83841 |
| 300 | 20 | 0.001 | 0.04514 | 0.93460 | 0.10181 | 0.81857 |
| 100 | 10 | 0.0001 | 0.08325 | 0.78762 | 0.09744 | 0.81921 |
| 200 | 1 | 0.0001 | 0.06186 | 0.86063 | 0.09823 | 0.80857 |
| 300 | 1 | 0.001 | 0.04553 | 0.94000 | 0.09875 | 0.83413 |
| 100 | 1 | 0.0001 | 0.06901 | 0.82190 | 0.09913 | 0.78571 |
| 200 | 10 | 0.0001 | 0.07367 | 0.82063 | 0.09957 | 0.77841 |
| 300 | 20 | 0.0001 | 0.07198 | 0.79937 | 0.09988 | 0.79444 |
| 300 | 20 | 0.01 | 0.05908 | 0.86794 | 0.10055 | 0.80841 |
| 100 | 20 | 0.01 | 0.05021 | 0.92175 | 0.10056 | 0.82317 |
| 200 | 20 | 0.001 | 0.04866 | 0.93952 | 0.10097 | 0.78095 |
| 300 | 10 | 0.0001 | 0.07087 | 0.83746 | 0.10174 | 0.75841 |
| 100 | 10 | 0.01 | 0.04902 | 0.91206 | 0.10176 | 0.78556 |
| 100 | 10 | 0.001 | 0.05525 | 0.89238 | 0.10188 | 0.78444 |
| 200 | 10 | 0.001 | 0.05029 | 0.91651 | 0.10234 | 0.81032 |
| 200 | 20 | 0.0001 | 0.08412 | 0.79683 | 0.10321 | 0.78000 |
| 300 | 1 | 0.0001 | 0.05370 | 0.90238 | 0.10436 | 0.78762 |
| 100 | 1 | 0.01 | 0.06871 | 0.82365 | 0.10709 | 0.75921 |
| 200 | 20 | 0.01 | 0.05432 | 0.91683 | 0.11083 | 0.78190 |
| 200 | 1 | 0.001 | 0.04542 | 0.93571 | 0.11203 | 0.75000 |
| 100 | 20 | 0.0001 | 0.08296 | 0.75698 | 0.11208 | 0.70444 |
| 200 | 10 | 0.01 | 0.08924 | 0.82016 | 0.12027 | 0.74190 |
| 300 | 10 | 0.01 | 0.08880 | 0.76238 | 0.12863 | 0.74048 |
| 300 | 1 | 0.01 | 0.09793 | 0.74810 | 0.13317 | 0.65349 |
| 200 | 1 | 0.01 | 0.11110 | 0.76635 | 0.13671 | 0.74825 |

Table 3: Benchmarking model with randomly selected datasets

| Data | Validation set | | Test Set | |
|-----------|----------------|----------|----------|----------|
| | MAE | Accuracy | MAE | Accuracy |
| Random50 | 0.064246 | 0.85205 | 0.058508 | 0.779762 |
| Random100 | 0.058567 | 0.809696 | 0.059457 | 0.755221 |
| Random200 | 0.067143 | 0.827499 | 0.056684 | 0.818358 |

Table 4: MV-ESG Model Evaluation

| | 2016 | | 2017 | | 2018 | |
|--------------|---------------|---------------|--------|---------------|---------------|---------------|
| | MAX-S | MAX-ESG | MAX-S | MAX-ESG | MAX-S | MAX-ESG |
| Return | 32.73% | 28.47% | 47.76% | 50.78% | 30.33% | 26.60% |
| Volatility | 17.22% | 14.89% | 19.37% | 19.18% | 16.84% | 14.31% |
| Sharpe Ratio | 1.7845 | 1.7777 | 2.3624 | 2.5431 | 1.6823 | 1.7191 |
| ESG Score | 70 | 74 | 70 | 75 | 68 | 71 |

in 2016 and 2018, they also showed a relatively higher volatility level. Conversely, our MAX-ESG portfolios still achieved great financial returns with lower risk. The Sharpe Ratios of the MAX-ESG portfolios were higher than those of the MAX-S ones for 2017 and 2018. In 2017, the MAX-ESG portfolio achieved a better financial return 50.78% at a lower risk level 19.19%, compared with 47.76% return at 19.37% volatility in the MAX-S portfolio. These findings showed that achieving a socially responsible investment portfolio, with higher ESG ratings, and without the sacrifice of a large financial return, was achievable with our MV-ESG model.

We compared the performance of our final MAX-ESG portfolio with reported financial returns in 2018 obtained from similar sustainable indexes and funds. The results in Table 5 show that our portfolio outperformed other indexes and funds in terms of financial performance and achieved the Sharpe Ratio of 2.0634. Our portfolio had the

Table 5: Benchmarking MAX-ESG portfolio with Sustainable Indexes and Funds in 2018

| | Period Returns | | | 3-year Annualized | | |
|----------------|----------------|--------|---------|-------------------|---------------|---------------|
| | 2016 | 2017 | 2018 | Return | Volatility | Sharpe Ratio |
| MAX-ESG | 28.47% | 50.78% | 26.60% | 35.28% | 16.13% | 2.0634 |
| S&P500 | 11.29% | 23.28% | -3.35% | 10.41% | 10.88% | 0.7727 |
| S&P500 ESG | 14.52% | 21.24% | -8.44% | 9.11% | 11.76% | 0.6043 |
| DJSI World | 8.23% | 27.98% | -8.03% | 9.39% | 11.52% | 0.6418 |
| DJSI WD | 10.71% | 24.00% | 9.54% | 14.75% | 10.48% | 1.2166 |
| ICLN | -16.91% | 21.48% | -9.02% | -1.48% | 17.11% | -0.2036 |
| TAN | -43.23% | 54.39% | -25.66% | -4.83% | 23.00% | -0.2971 |
| SUSA | 12.15% | 22.53% | -5.65% | 9.68% | 11.42% | 0.6722 |
| EQLT | 13.93% | 21.33% | 9.22% | 14.83% | 11.80% | 1.0870 |

best 3-year annualized return of 35.28%. Particularly in 2018, all indexes and funds had negative returns because many large stocks were in the downtrend. Our MAX-ESG portfolio was still able to achieve a positive return. This was mainly because the portfolio constructed was based on the maximization of Sharpe Ratio in the MV-ESG model, which optimally selects stocks with higher returns.

Our model's 3-year annualized volatility was in third place with 16.13%. This higher level of risk aligned with common investment knowledge on diversification [47]. Because the indexes often consist of a larger number of stocks, they generally had lower risks. However, the level diversification of our SRI portfolio was sufficient for individual investors. Our best MAX-ESG portfolio, for example, consisted of 7 stocks in 2016, 18 stocks in 2017 and 12 stocks in 2018 with the allocation as shown in Figure 6. The model could be enhanced to construct a more diversified portfolio for sustainable investment funds with further constraints on weights.

Since all these indexes and funds published different types of sustainability metrics,




Figure 6: MAX-ESG portfolio allocation (labels are trade symbols of companies)

we could not directly compared our portfolio ESG ratings to their benchmarks. We also could not report the net returns on investment due to the lacking of fund fees and tax calculation. In general, these results showed the effectiveness of our DRIP framework, not only for the optimization of socially responsible investment portfolios but also for financial stock investments in general.

4.6. Reinforcement Learning Test

Our reinforcement learning DRIP framework could be used to construct socially responsible portfolios with higher ESG ratings that still achieved competitive financial returns. Our DRIP model could predict multiple time steps ahead, which is an important feature for stock investors. Furthermore, the model significantly outperformed the univariate networks in both prediction accuracy and training speed with the same epoch size in terms of both prediction accuracy and training speed. In our experiments, it took one hour to perform reinforcement learning with the multivariate BiLSTM: a combination of

Table 6: Reinforcement Learning Test Results

| Prediction Model | | | | |
|--------------------------|---------------|---------------|---------------|------------|
| | MAE | RSME | AUC | Accuracy |
| DRIP | 0.0867 | 0.1117 | 0.9354 | 88% |
| Non-RL | 0.1098 | 0.1499 | 0.5515 | 55% |
| MAX-ESG Portfolio | | | | |
| | Return | Volatility | Sharpe Ratio | ESG Score |
| DRIP | 35.28% | 16.13% | 2.0634 | 73 |
| Non-RL | 5.30% | 14.00% | 0.2357 | 68 |

100 univariate neural networks with 10 epochs typically takes 100-times more training duration compared to our approach. This finding could lead to a better computationally efficient approach because the multivariate BiLSTM takes N times less in total training duration.

We also tested the performance of reinforcement learning by comparing the results to those without prediction model retraining and portfolio rebalancing (Non-RL). In the “Non-RL” framework, we still used the Multivariate BiLSTM networks for prediction model and the MV-ESG for SRI portfolio optimization. However, the models were retrained after each testing period (each quarter or each year) without any pre-trained model and parameter learning from previous periods. The results in Table 6 showed that reinforcement learning had significantly improved the model performance in terms of both the prediction of stock returns and the optimization of portfolios. Since we were working with multivariate time series, retraining models and rebalancing portfolios were proven to be essential. Therefore, our reinforcement learning approach within the DRIP system was suitable for this time-sensitive data analytics problem.

4.7. Discussion of Research

Overall, our research demonstrated a promising trend in applying deep learning techniques for the selection of socially responsible investment portfolios. With the current progress in artificial intelligence, we believe it will bring further breakthroughs in socially responsible investment research. Our research will not only contribute directly to current literature in various disciplines but also translate into benefits for responsible investors, funds or indexes in markets. In the AI research field, our prediction model with a BiLSTM network could serve as a baseline for further research of long-term stock return forecasting using neural networks. In this research, we only used a single type of neural networks and structured data (stock prices and ESG ratings) as input. Studying the different variations and combination of the deep neural networks, as well as incorporating unstructured data (news, company reports or social media content) with text mining approaches in SRI, was beyond the scope of this paper. However, our framework was designed with the flexibility to adopt different data mining approaches, neural networks or optimization algorithms in our future research.

Please be aware that by focusing on policies and rewards, our system might fail under the extreme situations, e.g. a financial crisis. Our model's performance and applicability are subject to the hypothesis of stable company performances and normal finance market scenario. To safeguard investments in such cases, we would need additional failsafe measures when applying our model in practice.

From the financial research aspect, the MV-ESG model was one of the first to combine ESG ratings with a math finance model. Further research on how to incorporate this with other quantitative finance models such as GARCH [18] would be relevant to both SRI scholars and investors. As many researchers are working on similar approaches for oil price forecasting [31], we believe a further investigation into this direction would be beneficial for SRI researchers. For simplicity purpose, we did not take into account income tax rates, inflation rates, trading fees and other financial fund management costs. Further calculation of these fees would help the model implementation in the real-world

investment scene. Moreover, ESG ratings are not the only metrics to measure corporate social responsibility. The integration of hundreds of ESG sub-categorical ratings (e.g. greenhouse gas emissions or community support) could improve the model significantly. In our future research, we could study the potential of using deep learning approaches for a personalized stock recommendation system in the SRI context.

5. Conclusions

Socially responsible investment is an emerging research topic with potential for long-term social impact. In this research, we proposed the DRIP model, which leveraged deep learning techniques to predict financial returns and construct a socially responsible investment portfolio. Validated with real-world data, our DRIP model, with a multivariate time series model, was able to accurately predict the stock returns three months ahead of time. It is possible that our framework could be generalized to build decision-support systems for similar multivariate prediction problems.

The socially responsible portfolios that we constructed using our novel MV-ESG model and reinforcement learning achieved much higher ESG ratings and a competitive financial performance overall compared with standard MV portfolio models and similar sustainable indexes and funds. With this rising trend in socially responsible investment, financial capital will diverge into good companies that contribute to a cleaner environment and a better society. This research also highlights a new direction for the use of more advanced deep learning approaches for quantitative finance research.

Acknowledgment

(removed for blind review)

References

- [1] Amel-Zadeh, A., & Serafeim, G. (2018). Why and how investors use ESG information: Evidence from a global survey. *Financial Analysts Journal*, 74, 87–103.
- [2] Auer, B. R., & Schuhmacher, F. (2016). Do socially (ir) responsible investments pay? New evidence from international ESG data. *The Quarterly Review of Economics and Finance*, 59, 51–62.
- [3] Bao, W., Yue, J., & Rao, Y. (2017). A deep learning framework for financial time series using stacked autoencoders and long-short term memory. *PLoS one*, 12, e0180944.
- [4] Bose, I., & Pal, R. (2012). Do green supply chain management initiatives impact stock prices of firms? *Decision support systems*, 52, 624–634.
- [5] Broyden, C. G., Dennis Jr, J., & Moré, J. J. (1973). On the local and superlinear convergence of quasi-Newton methods. *IMA Journal of Applied Mathematics*, 12, 223–245.
- [6] Calvo, C., Ivorra, C., & Liern, V. (2016). Fuzzy portfolio selection with non-financial goals: exploring the efficient frontier. *Annals of Operations Research*, 245, 31–46.
- [7] Chen, K., Zhou, Y., & Dai, F. (2015). A LSTM-based method for stock returns prediction: A case study of China stock market. In *Proceedings of the 2015 IEEE International Conference on Big Data* (pp. 2823–2824). IEEE.
- [8] Chen, T., Xu, R., He, Y., & Wang, X. (2017). Improving sentiment analysis via sentence type classification using BiLSTM-CRF and CNN. *Expert Systems with Applications*, 72, 221–230.
- [9] Chiang, W., Enke, D., Wu, T., & Wang, R. (2016). An adaptive stock index trading decision support system. *Expert Systems with Applications*, 59, 195–207.
- [10] Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, .
- [11] Chollet, F. et al. (2015). Keras. <https://keras.io>.
- [12] Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39, 1–22.
- [13] Di Persio, L., & Honchar, O. (2016). Artificial neural networks architectures for stock price prediction: comparisons and applications. *International Journal of Circuits, Systems and Signal Processing*, 10, 403–413.
- [14] Dumay, J., Guthrie, J., & Farneti, F. (2010). GRI sustainability reporting guidelines for public and third sector organizations: A critical review. *Public Management Review*, 12, 531–548.
- [15] Eccles, N., & Viviers, S. (2011). The origins and meanings of names describing investment practices that integrate a consideration of ESG issues in the academic literature. *Journal of Business Ethics*, 104, 389–402.

- [16] Fatemi, A., Glaum, M., & Kaiser, S. (2018). ESG performance and firm value: The moderating role of disclosure. *Global Finance Journal*, 38, 45 – 64. Special Issue on Corporate Social Responsibility and Ethics in Financial Markets.
- [17] Feuerriegel, S., & Gordon, J. (2018). Long-term stock index forecasting based on text mining of regulatory disclosures. *Decision Support Systems*, 112, 88–97.
- [18] Francq, C., & Zakoian, J.-M. (2019). *GARCH models: structure, statistical inference and financial applications*. Wiley.
- [19] Friede, G., Busch, T., & Bassen, A. (2015). ESG and financial performance: aggregated evidence from more than 2000 empirical studies. *Journal of Sustainable Finance and Investment*, 5, 210–233.
- [20] Gadre-Patwardhan, S., Katdare, V. V., & Joshi, M. R. (2016). A Review of Artificially Intelligent Applications in the Financial Domain. In *Artificial Intelligence in Financial Markets* (pp. 3–44). Springer.
- [21] Garcia-Bernabeu, A., Pla, D., Bravo, M., & Perez-Gladish, B. (2015). Mean-variance stochastic goal programming for sustainable mutual funds' portfolio selection. *Rect@*, 16, 135.
- [22] Gray, H. (1983). *New directions in the investment and control of pension funds*. Investor Responsibility Research Center.
- [23] Halbritter, G., & Dorfleitner, G. (2015). The wages of social responsibility—where are they? A critical review of ESG investing. *Review of Financial Economics*, 26, 25–35.
- [24] Henrique, B. M., Sobreiro, V. A., & Kimura, H. (2019). Literature review: Machine learning techniques applied to financial market prediction. *Expert Systems with Applications*, 124, 226 – 251.
- [25] Ho, C.-S., Damien, P., Gu, B., & Konana, P. (2017). The time-varying nature of social media sentiments in modeling stock returns. *Decision Support Systems*, 101, 69–81.
- [26] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9, 1735–1780.
- [27] Jiang, Q., Tang, C., Chen, C., Wang, X., & Huang, Q. (2019). Stock Price Forecast Based on LSTM Neural Network. In J. Xu, F. L. Cooke, M. Gen, & S. E. Ahmed (Eds.), *Proceedings of the Twelfth International Conference on Management Science and Engineering Management* (pp. 393–408). Cham: Springer International Publishing.
- [28] Kempf, A., & Osthoff, P. (2007). The effect of socially responsible investing on portfolio performance. *European Financial Management*, 13, 908–922.
- [29] Kraft, D. (1988). A software package for sequential quadratic programming. *Forschungsbericht-Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt*, .
- [30] Kraus, M., & Feuerriegel, S. (2017). Decision support from financial disclosures with deep neural networks and transfer learning. *Decision Support Systems*, 104, 38–48.

- [31] Kristjanpoller, W., & Minutolo, M. C. (2016). Forecasting volatility of oil price using an artificial neural network-GARCH model. *Expert Systems with Applications*, 65, 233–241.
- [32] Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, 7, 77–91.
- [33] Moghaddam, A. H., Moghaddam, M. H., & Esfandyari, M. (2016). Stock market index prediction using artificial neural network. *Journal of Economics, Finance and Administrative Science*, 21, 89 – 93.
- [34] Munoz, F., Vargas, M., & Marco, I. (2014). Environmental mutual funds: Financial performance and managerial abilities. *Journal of Business Ethics*, 124, 551–569.
- [35] Nam, K., & Seong, N. (2019). Financial news-based stock movement prediction using causality analysis of influence in the Korean stock market. *Decision Support Systems*, 117, 100 – 112.
- [36] Nelson, D. M., Pereira, A. C., & de Oliveira, R. A. (2017). Stock market's price movement prediction with LSTM neural networks. In *Proceedings of the 2017 International Joint Conference on Neural Networks (IJCNN)* (pp. 1419–1426). IEEE.
- [37] Nilsson, J. (2008). Investment with a conscience: Examining the impact of pro-social attitudes and perceived financial performance on socially responsible investment behavior. *Journal of Business Ethics*, 83, 307–325.
- [38] Peylo, B. T. (2012). A Synthesis of Modern Portfolio Theory and Sustainable Investment. *The Journal of Investing*, 21, 33–46.
- [39] Powell, M. J. (1964). An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The Computer Journal*, 7, 155–162.
- [40] Samarawickrama, A., & Fernando, T. (2017). A recurrent neural network approach in predicting daily stock prices an application to the Sri Lankan stock market. In *Proceedings of the 2017 IEEE International Conference on Industrial and Information Systems (ICIIS)* (pp. 1–6). IEEE.
- [41] Schäfer, H. (2016). Corporate Social Responsibility Rating. *A Handbook of Corporate Governance and Social Responsibility*, (p. 449).
- [42] Schuster, M., & Paliwal, K. K. (1997). Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45, 2673–2681.
- [43] Selvin, S., Vinayakumar, R., Gopalakrishnan, E., Menon, V. K., & Soman, K. (2017). Stock price prediction using LSTM, RNN and CNN-sliding window model. In *Proceedings of the 2017 International Conference on Advances in Computing, Communications and Informatics* (pp. 1643–1647). IEEE.
- [44] Sermpinis, G., Karathanasopoulos, A., Rosillo, R., & de la Fuente, D. (2019). Neural networks in financial trading. *Annals of Operations Research, Special Issue: Networks and Risk Management*, 11–16.

- [45] Sharpe, W. F. (1966). Mutual fund performance. *The Journal of Business*, 39, 119–138.
- [46] Siddiqui, A. I., Marinova, D., Hossain, A., & Todorov, V. (2011). Socially Responsible Investment in Australia. *Sustainability And Development In Asia And The Pacific: Emerging Policy Issues*, (p. 249).
- [47] Statman, M. (2004). The diversification puzzle. *Financial Analysts Journal*, 60, 44–53.
- [48] Stay, C. (2010). *Corporate Social Responsibility*. Technical Report Sustainalytics.
- [49] Stephen, S. (2018). Financial Performance of Environmentally Responsible Investment Funds: A Systematic Review. In *Academy of Management Proceedings* (p. 12451). Academy of Management Briarcliff Manor, NY 10510 volume 2018.
- [50] Thomson Reuters (2019). *Thomson Reuters ESG Scores*. Technical Report Thomson Reuters.
- [51] US SIF Foundation (2018). *2018 Biennial Report On US Sustainable, Responsible And Impact Investing Trends*. Technical Report US SIF Foundation.
- [52] Van Duuren, E., Plantinga, A., & Scholtens, B. (2016). ESG integration and the investment management process: Fundamental investing reinvented. *Journal of Business Ethics*, 138, 525–533.
- [53] Verheyden, T., Eccles, R. G., & Feiner, A. (2016). ESG for all? The impact of ESG screening on return, risk, and diversification. *Journal of Applied Corporate Finance*, 28, 47–55.
- [54] Von Wallis, M., & Klein, C. (2015). Ethical requirement and financial interest: a literature review on socially responsible investing. *Business Research*, 8, 61–98.
- [55] Zopounidis, C., Galariotis, E., Doumpos, M., Sarri, S., & Andriosopoulos, K. (2015). Multiple criteria decision aiding for finance: An updated bibliographic survey. *European Journal of Operational Research*, 247, 339 – 348.