



Machine Learning for Socially Responsible Portfolio Optimisation

Taisha Nundlall
taishanundlall@gmail.com
University of Johannesburg
Johannesburg, South Africa

Terence L. van Zyl
University of Johannesburg
Johannesburg, South Africa
tvanzyl@gmail.com

ABSTRACT

Socially responsible investors build investment portfolios intending to incite social and environmental advancement alongside a financial return. Although Mean-Variance (MV) models successfully generate the highest possible return based on an investor's risk tolerance, MV models do not make provisions for additional constraints relevant to socially responsible (SR) investors. In response to this problem, the MV model must consider Environmental, Social, and Governance (ESG) scores in optimisation. Based on the prominent MV model, this study implements portfolio optimisation for socially responsible investors. The amended MV model allows SR investors to enter markets with competitive SR portfolios despite facing a trade-off between their investment Sharpe Ratio and the average ESG score of the portfolio.

CCS CONCEPTS

• **Computing methodologies** → **Unsupervised learning**; • **Applied computing** → **Forecasting**.

KEYWORDS

portfolio optimisation, ESG, social responsibility, machine learning

ACM Reference Format:

Taisha Nundlall and Terence L. van Zyl. 2023. Machine Learning for Socially Responsible Portfolio Optimisation. In *2023 7th International Conference on Intelligent Systems, Metaheuristics & Swarm Intelligence (ISMSI 2023)*, April 23–24, 2023, Virtual Event, Malaysia. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3596947.3596966>

1 INTRODUCTION

Given the climate crisis and the ongoing socioeconomic issues, more investors are looking to participate in the markets through socially responsible investment (SRI) [2] decisions. Socially responsible (SR) investors build investment portfolios to initiate social and environmental advancement alongside a financial return [1]. The availability of Environmental, Social and Governance (ESG) ratings¹ provides valuable information that can be incorporated into SRI decisions, where both moral value and economic value

¹Also known as ESG issues, are non-financial components that are investigated to rank and compare the sustainability and ethical impact of various securities

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ISMSI 2023, April 23–24, 2023, Virtual Event, Malaysia

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9992-0/23/04...\$15.00
<https://doi.org/10.1145/3596947.3596966>

are of concern. During SRI asset allocation, an additional trade-off [3] exists between portfolio returns, risk and ESG-ratings (i.e. Sharpe Ratio and ESG). Although Mean-Variance (MV) [7] models successfully generate the highest possible return based on the investor's risk tolerance, MV models do not make provisions for additional constraints relevant to SR investors [11, 13]. In response to this problem, we propose to extend the MV Model to include ESG-rating optimisation. A tri-criterion (i.e. mean, variance, ESG-rating) portfolio optimisation model allows investors to better align investment decisions with moral values, benefiting from such decisions' economic value.

Vo et al. [15] provides a Deep Responsible Investment Portfolio (DRIP), a portfolio optimization model that uses deep learning and incorporates ESG-ratings. Stock prices and ESG-ratings are inputs in a Multivariate Bi-directional Long Short-Term Memory (LSTM) neural network to predict stock returns. Deep reinforcement learning is used to re-train neural networks and re-balance the portfolio. The DRIP framework achieved competitive financial performance and better social impact than traditional portfolio models, sustainable indexes and funds. This literature is relevant to our research and provides useful guidance on incorporating ESG-ratings into deep learning models.

This study uses machine learning [6, 8, 14] to calculate optimal SRI portfolios to evaluate the effect [9] of considering ESG-ratings in portfolio optimisation. Exchange-traded funds (ETF) closing prices are forecasted using a random forest (RF) regression algorithm to produce out-of-sample data for use during asset allocation. This study uses an amended MV model [15] to predict portfolio weightings to optimise the portfolio Sharpe Ratio [12] and ESG-ratings. The research value for this study is a look at the effect of ESG-ratings as a third criterion in portfolio optimisation when using machine learning to forecast ETF returns.

This paper is structured as follows. In Section 1, a brief introduction to the study is presented. In Section 2, the research methodology is discussed in detail. In Section 3, results from the study's machine learning and portfolio optimisation components are discussed, respectively. Section 4 presents the conclusions of this study.

2 METHOD

This study consists of two main components: the use of machine learning to forecast the closing prices of the ETF and the optimisation of a portfolio using traditional mean-variance and the ESG-MV methods proposed in Vo et al. [15]. Details on the research methodology are discussed in this section, and the general structure of the methodology is depicted in Figure 1.

2.1 Datasets

VettaFi owns databases containing various financial datasets related to Exchange-Traded Funds (ETF). ETFs are a relatively new pooled

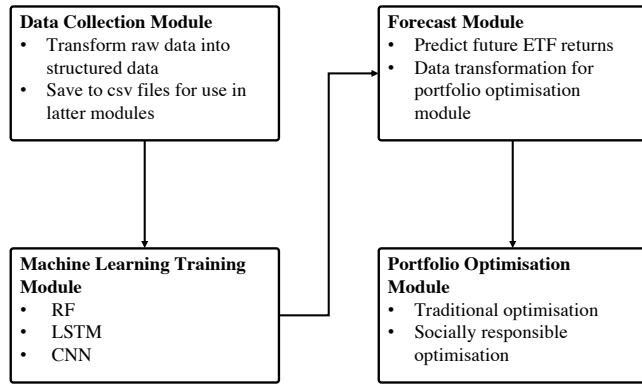


Figure 1: A diagram expressing the main idea of the study

investment security, with SPDR S&P 500 ETF Trust being the first ETF available on US exchanges, debuting as recently as 1993. VettaFi has gone one step further than their industry competitors to include ESG data for all ETFs captured on their system and has been made available for download to the public. We have downloaded the ESG data of U.S. ETFs traded on U.S. exchanges. Other ETF data are included in the VettaFi database, such as returns, leverage status, and year-to-date ETF prices. The dataset structure used can be seen in Figure 2. This study uses the ESG scores from this dataset, ranging from 0 to 10. The other features of the dataset are redundant, as they are merely sub-components of the ESG score.

Symbol	ETF Name	ESG Score	ESG Score Peer Percentile (%)	ESG Score Global Percentile (%)	Carbon Intensity (Tons of CO ₂ e / \$M Sales)	Sustainable Impact Solutions (%)
SPY	SPDR S&P 500 ETF Trust	8.06	68.18%	71.77%	150.02	5.65%
IIV	iShares Core S&P 500 ETF	8.06	67.75%	71.67%	150.12	5.64%
VTI	Vanguard Total Stock Market ETF	7.68	36.83%	62.54%	150.27	5.79%
VOO	Vanguard S&P 500 ETF	8.06	67.40%	71.62%	150.21	5.63%
QQQ	Invesco QQQ Trust	7.99	59.66%	69.84%	54.57	7.53%

Figure 2: ESG dataset released by VettaFi

Yahoo Finance has released datasets containing U.S ETFs and their historical prices from 1993-01-29 to 2021-11-30. The version used in this study is related to the financial values of November 2021. The structure of the dataset used can be seen in Figure 3. More than 80% of the historical data belong to trading days between 2011-11-30 and 2021-11-30. This is due to the increased number of ETFs available on the US exchange during this period. Our study limits the use of historical data to trading days between 2011-11-30 and 2021-11-30 to avoid excessive back-filling for the 80% of the ETFs that were non-existent before 2011-11-30. Before supplying the Yahoo Finance dataset to the portfolio optimisation component of the study, data related to ETFs not contained in the VettaFi dataset is removed to only consider ETFs with recorded ESG ratings. This reduces the number of ETFs in the dataset from 1644 to 1358.

2.2 Study Procedure

After performing the data preprocessing as detailed, the dataset is filtered to consider the adjusted closing values for one ETF at a time. The dataset is split 4:1 into training and testing datasets.

	fund_symbol	open	high	low	close	adj_close	volume
price_date							
2020-09-09	AAA	25.10	25.12	25.07	25.07	24.85	17300
2020-09-10	AAA	25.06	25.07	25.05	25.07	24.85	23500
2020-09-11	AAA	25.04	25.05	25.02	25.03	24.81	33400
2020-09-14	AAA	25.01	25.06	25.01	25.02	24.80	13100
2020-09-15	AAA	25.02	25.03	25.01	25.01	24.79	12100

Figure 3: ETF financial data released by Yahoo Finance

The historical data of one ETF is used to train and test each of the three machine learning models; random forest, LSTM, and Convolutional Neural Network (CNN). Using other ETFs, re-runs of the prediction models are performed to determine if the performance of the models is consistent. ETFs considered in the experiment are Betashares Australian High-Interest Cash ETF (AAA), SPDR S&P 500 ETF Trust (SPY), and Innovator US Equity Ultra Buffer ETF - October (UOCT).

Deep learning models have been constructed using the Adam optimiser, ReLu activation functions, and the mean absolute error (MAE) loss function. The CNN also made use of max pooling. The RF model is constructed with 10000 n_estimators.

All predictive models' mean absolute scaled error (MAE) MASE and root mean squared scaled error (RMSSE) are computed and evaluated to determine the best-performing model during training. The random forest model outperforms the other models and is further discussed in Section 3. The RF model forecasts daily closing prices for the two months following 2021-11-30. The historical and forecasted adjusted closing values, along with the ESG scores from the VettaFi dataset, were stored in a new dataset for the utility of the portfolio optimisation module.

For portfolio optimisation [10], we used historical data from 2011-12-01 to 2021-11-30, 10 years, and included our forecast data from the forecasting component of the study. As done in [15], portfolio optimisation, for both MV and ESG-MV portfolio optimization, is done using the Sequential Least Squares Programming (SLSQP) optimisation method.

No bounds have been placed on ESG scores in data pre-processing and portfolio optimisation modules. This is so that the entire market of 100 shares is considered in both the MV and the ESG-MV models. This will allow us to study the optimal weights allocated to assets with high and low ESG scores.

As an evaluation metric to measure the performance of a prediction model, a forecast error is computed from the difference between observed values and predicted values, that is, the ETF closing prices of the testing dataset vs. the predicted ETF closing prices. Prevalent forecast errors, such as MAE and RMSE, are scale dependent and cannot compare the time series of different units. As a way around scale dependency, scaled errors [5] compute the forecast errors using the MAE attained during training of the forecasting model. In this study, forecasting is performed on nonseasonal time series and scaled errors are computed using naïve forecasts, as provided in [4]. Scaled errors used to evaluate machine learning model performance during training are the mean absolute scaled error

(MASE) and the root mean squared scaled error (RMSSE). A scaled error is (favourably) lower than one if forecasts are more accurate than the average one-step naïve forecast computed on the training set. Scaled errors are computed as follows:

$$MASE = \text{mean}(|q_j|) \quad (1)$$

$$RMSSE = \sqrt{\text{mean}(q_j^2)} \quad (2)$$

where q_j is:

$$q_j = \frac{e_j}{\frac{1}{T-1} \sum_{t=2}^T |y_t - y_{t-1}|} \quad (3)$$

and q_j^2 is the square of the aforementioned q_j . e_j is the forecast error, $\{y_1, \dots, y_t\}$ is the training data, and T is the size of the training set.

To determine if the findings from the portfolio optimisation module are consistent, the optimisation algorithm is run multiple times using different (randomised) combinations of 100 ETFs. In every run, the expected annualized portfolio return, annualized risk, Sharpe ratio [12], and average portfolio ESG-rating are computed. Sharpe ratios are computed to relate the return of an investment to the risk of such an investment, where a favourable Sharpe ratio is greater than 1.

The optimisation function for the traditional MV model is calculated as a minimising function of the negative Sharpe ratio (S_p);

$$\min(-S_p) = \min\left(-\frac{r_p - r_f}{\sigma_p}\right) \quad (4)$$

where r_f is the risk-free return and portfolio return (r_p) is

$$r_p = \sum_{i=1}^N w_i r_i \quad (5)$$

and portfolio risk (σ_p) is,

$$\sigma_p = \sum_{i=1}^N \sum_{j=1}^N w_i \sigma_{ij} w_j \quad (6)$$

and w_i is the weight allocated to the ETF in the portfolio. Similarly, the ESG-MV model is calculated as a minimising function of the negative Sharpe ratio (\tilde{S}_p) but now includes the average portfolio ESG-rating, which is the mean of weighted ESG scores of individual ETFs;

$$\min(-\tilde{S}_p) = \min\left(-\frac{r_p - r_f}{\sigma_p} \text{ESG}_p\right) \quad (7)$$

Due to the available data on the ESG ratings, we cannot use historical ESG ratings to improve the precision of our portfolio optimisation, as was done in Vo et al. [15].

2.3 Tools and Libraries

All experiments were run on a 1.8GHz Dual-Core Intel Core i5 (5350U). All algorithms were implemented using Python 3.3 via the Jupyter Notebook open-source IDE. Python libraries used to implement the algorithms in this study include NumPy, SciPy, Statistics, Math, Matplotlib, Seaborn, Pandas, Tensorflow, and Scikit-Learn.

All related code, notebooks, datasets, and results can be viewed here.

3 RESULTS & DISCUSSION

3.1 Results of the prediction model

Adjusted closing prices for AAA ETF are predicted using machine learning algorithms. Figures 4 to 6 show the performance during training and testing.



Figure 4: Closing Prices using Random Forest Model




Figure 5: Closing Prices using LSTM Model




Figure 6: Closing Prices using CNN Model

It can be seen from Figures 4, 5, and 6, during testing and validation, the RF and CNN models show high variance (i.e. overfitting) whilst the LSTM shows underfitting. To specifically investigate the behaviour of the prediction models during training, Figures 7 to

9 plot the actual and predicted data, omitting the test data. When closely looking at Figure 7, the RF model overestimates the troughs and peaks. Figure 8 shows that the LSTM model is significantly under-fitting. Figure 9 shows that CNN does the opposite of RF and underestimates troughs and peaks.



Figure 7: Training and Validation using Random Forest Model




Figure 8: Training and Validation using LSTM Model




Figure 9: Training and Validation using CNN Model

Tables 1 to 3 provide the MASE and MAE of the models when predicting the adjusted closing prices of the AAA, SPY and UOCT ETF.

Table 1: MASE and RMSE when predicting AAA ETF closing prices

Model	MASE	RMSSE	Confidence Interval
RF	0.0014	0.0011	(25.0110, 25.0372)
LSTM	0.0126	0.0077	(25.0123, 25.0184)
CNN	0.3105	0.2500	(24.9973, 25.0468)

Table 2: MASE and RMSE when predicting SPY ETF closing prices

Model	MASE	RMSSE	Confidence Interval
RF	0.0000	0.0000	(148.93, 157.78)
LSTM	706.5726	479.5685	(283.95, 289.69)
CNN	0.0445	0.0337	(284.26, 291.01)

Table 3: MASE and RMSE when predicting UOCT ETF closing prices

Model	MASE	RMSSE	Confidence Interval
RF	0.001315	0.001055	(25.3791, 25.8541)
LSTM	0.337371	0.209528	(27.7632, 27.8420)
CNN	0.040433	0.030315	(27.6718, 27.9850)

From Table 2, the LSTM model performs poorly with wider confidence intervals. This could be attributed to the underfitting due to the model architecture.

To improve the performance of the three models, effective regularisation techniques need to be incorporated into the models to reduce generalisation errors. If implemented correctly, regularisation will improve the variance issues displayed in the RF and CNN models. The CNN model could be too complex for this time series model, leading to high variance.

Regardless of expected ETF prices, the RF outperforms the other models by having the lowest MASE and RMSE. Due to this, the RF model is used to predict the adjusted closing prices of ETFs for the 2021-11-30 to 2022-01-31 trading period. Figure 10 plots the forecasting of the AAA ETF for the 2021-11-30 to 2022-01-31 trading period using the RF model. This is a sample of the forecasts produced, stored and utilised in the portfolio optimisation component of the study.




Figure 10: Forecast Adjusted Closing Prices of AAA ETF

3.2 Portfolio optimisation model results

The results obtained from portfolio optimisation using the MV and ESG-MV methods are provided in Table 4. The table shows results from a market size of 100, where the market is randomly generated in every run (that is, every run contains a random selection of 100 ETFs to optimise over). Note that all values in the table are annualised.

Table 4: Tabulated Results from Optimisation Module

Run	Model	Sharpe Ratio	Risk	Return	Mean ESG Score
1	MV	2.174	11.486	31.218	5.204
	ESG MV	1.981	10.062	26.182	7.131
2	MV	2.031	6.265	18.974	5.725
	ESG MV	2.863	5.440	21.823	7.208
3	MV	2.475	9.441	29.617	5.899
	ESG MV	2.213	8.687	25.470	7.878
4	MV	1.933	8.751	23.166	1.715
	ESG MV	1.283	5.836	13.736	5.787
5	MV	1.831	9.447	23.545	5.903
	ESG MV	1.623	6.613	16.981	7.537
6	MV	3.277	9.430	37.156	6.375
	ESG MV	3.039	7.503	29.052	7.294
7	MV	1.679	12.03	26.445	3.584
	ESG MV	1.330	7.332	16.004	6.069
8	MV	1.589	12.58	26.242	3.044
	ESG MV	1.097	6.234	13.086	6.375
9	MV	2.120	5.550	18.017	6.179
	ESG MV	1.990	6.140	18.465	7.307
10	MV	3.172	6.482	26.813	5.555
	ESG MV	2.955	6.420	25.218	7.252
11	MV	1.920	4.841	15.546	4.300
	ESG MV	1.541	4.775	13.607	6.483
12	MV	2.294	7.709	23.933	2.662
	ESG MV	1.554	4.475	13.203	6.238

When comparing the results of the MV model to the ESG-MV models, the following are observed:

- There is an average decrease of 12.29% in the Sharpe ratio. SR portfolios perform slightly worse on average at the expense of achieving a drastic increase of 65.72% in mean portfolio ESG scores.
- There is an average decrease of 20% in annualised risk and an average reduction of 21.32% in annualised return. This implies that SR portfolios are less risky investments, but investors (as expected) suffer a lower return from said portfolios.

The ESG-MV model is still capable of achieving positive returns. This is achieved by maximizing the Sharpe Ratio in traditional MV models, which optimally selects ETFs with higher returns and avoids negative annualised returns. All Sharpe Ratios are favourably greater than 1, indicating that all optimal portfolio weights produced in the experiment have resulted in attractive investment portfolios.

In run two from Table 4, the Sharpe ratio using ESG-MV is higher than that of the MV model. This indicates that the ESG-MV model can achieve higher ESG ratings and Sharpe ratios than the MV model in particular markets. This shows that SR portfolios can have competitive financial performance compared to portfolios produced from the traditional MV model.

Table 5 provides the weights allocated to each of the 100 ETFs in the market generated during run 5 of the experiment. From Table

5, it is seen that ESG scores drive the optimal weights of ETFs with high ESG ratings (see WBIG ETF in Table 5).

In this particular run, both models consist of 33 non-zero weighted ETFs. However, zero-weighted ETFs in the MV model do not correlate with zero-weighted ETFs in ESG-MV. This is attributed to a higher ESG score being traded off by the risk and return of the ETF. This can be seen in the change in weightings of WBIG and SDP ETF. More specifically, WBIG (with a high ESG score of 9.38) is assigned 13.37% in the ESG-MV model but 0% in the MV model, while SDP (with a low ESG score of 0) is assigned 0% in the ESG-MV model but 6.62% in the MV model.

4 CONCLUSION

This study was designed to leverage machine learning in the computation of optimal SRI portfolios to evaluate the effect of considering ESG ratings in portfolio optimisation. Random forest, LSTM neural network, and CNN models were implemented to predict and forecast adjusted closing prices of selected ETFs. In evaluating model performance, the RF model outperformed the deep learning models during model testing. Consequently, the RF model was used to forecast the adjusted closing values of all ETFs.

Using historical and forecasted data, this study has shown that ESG-MV models produce competitive portfolios with a slight trade-off between a lower Sharpe ratio for higher mean ESG scores of the portfolio. The findings will interest SR investors seeking a portfolio with a balance of moral and economic value. The study is limited by the absence of historical data on ESG scores as utilised in Vo

Table 5: Asset allocation weights from MV optimisation vs MV-ESG optimisation, sorted by ascending ESG scores

ETF	ESG Score	MV Optimal weights (%)	ESG-MV Optimal weights (%)
SMMU	0	0	0
SDP	0	6.624	0
PAUG	0	0	0
IBMQ	0	0	0
FUMB	0	0	0
HIPS	0	2.42	0
DJUL	0	0	0
XBJL	0	0	0
BNKU	0	3.433	0
MLPA	3.73	7.052	2.324
OSCV	3.97	0	0
PSQ	5.81	0.158	0.012
SCHQ	5.81	0	0
SDOW	5.81	27.412	25.073
SPTS	5.81	0	0
FNDB	7.66	7.756	4.894
IVE	7.68	0.485	5.225
VOOV	7.68	0.057	0.003
FNCL	7.72	7.779	4.875
LRNZ	7.79	0	0
SPUU	7.88	4.102	3.259
EPS	7.89	0	0
USXF	9.22	0	2.305
WBIG	9.38	0	13.372
RBND	9.57	0	0

et al. [15]. The study should be repeated using regularisation in machine learning models to improve the overall performance of the models' forecasting accuracy.

ACKNOWLEDGMENTS

We thank the Nedbank Research Chair for his guidance, dedication, and invaluable insights while writing this paper. To my friends and family for their unwavering support and encouragement.

REFERENCES

- [1] Michael Branch, Lisa R Goldberg, and Pete Hand. 2019. A guide to ESG portfolio construction. *The Journal of Portfolio Management* 45, 4 (2019), 61–66.
- [2] Li Chen, Lipei Zhang, Jun Huang, Helu Xiao, and Zhongbao Zhou. 2021. Social responsibility portfolio optimization incorporating ESG criteria. *Journal of Management Science and Engineering* 6, 1 (2021), 75–85.
- [3] Jan De Spiegeleer, Stephan Höcht, Daniel Jakubowski, Sofie Reyners, and Wim Schoutens. 2021. ESG: A new dimension in portfolio allocation. *Journal of Sustainable Finance & Investment* (2021), 1–41.
- [4] R. J. Hyndman and G. Athanasopoulos. 2021. *Forecasting: principles and practice* (3 ed.). OTexts.
- [5] R. J. Hyndman and A.B. Koehler. 2006. Another look at measures of forecast accuracy. *International Journal of Forecasting* 22, 4 (2006), 679–688.
- [6] Siddeeq Laher, Andrew Paskaramoorthy, and Terence L Van Zyl. 2021. Deep learning for financial time series forecast fusion and optimal portfolio rebalancing. In *2021 IEEE 24th International Conference on Information Fusion (FUSION)*. IEEE, 1–8.
- [7] Harry Markowitz. 1952. Portfolio Selection. *Journal of Finance* 7, 1 (1952), 77–91.
- [8] Mufhumudzi Muthivhi and Terence L van Zyl. 2022. Fusion of Sentiment and Asset Price Predictions for Portfolio Optimization. In *2022 25th International Conference on Information Fusion (FUSION)*. IEEE, 1–8.
- [9] Abdessamad Ouchen. 2022. Is the ESG portfolio less turbulent than a market benchmark portfolio? *Risk Management* 24, 1 (2022), 1–33.
- [10] G. A. Vijayalakshmi Pai. 2018. *Metaheuristics for Portfolio Optimization- An Introduction using MATLAB*. Wiley-ISTE.
- [11] Andrew B Paskaramoorthy, Tim J Gebbie, and Terence L van Zyl. 2020. A framework for online investment decisions. *Investment Analysts Journal* 49, 3 (2020), 215–231.
- [12] F William Sharpe. 1966. Mutual Fund Performance. (1966), 119–138.
- [13] Tarrin Skeepers, Terence L van Zyl, and Andrew Paskaramoorthy. 2021. MA-FDRNN: Multi-asset fuzzy deep recurrent neural network reinforcement learning for portfolio management. In *2021 8th International Conference on Soft Computing & Machine Intelligence (ISCMI)*. IEEE, 32–37.
- [14] Terence L van Zyl, Matthew Woolway, and Andrew Paskaramoorthy. 2021. Par-den: Surrogate assisted hyper-parameter optimisation for portfolio selection. In *2021 8th international conference on soft computing & machine intelligence (ISCMI)*. IEEE, 101–107.
- [15] Nhi N Y Vo, Xuezhong He, Shaowu Liu, and Guandong Xu. 2019. Deep Learning for Decision Making and the Optimization of Socially Responsible Investments and Portfolio. *Decision Support Systems* (2019).