Iosr Journal Of Economics And Finance (Iosr-Jef)
E-Issn: 2321-5933, P-Issn: 2321-5925.Volume 16, Issue 1 Ser. 3 (Jan. – Feb. 2025), Pp 59-70
www.Iosrjournals.Org

A Deep Learning Based Asset Allocation Methodology For Investment Portfolio Optimization Under Uncertainties

Marcus V. De Oliveira1 And Gilson A. O. P. Da Costa2
Instituto De Matemática E Estatística – Universidade Do Estado Do Rio De Janeiro (UERJ)
CEP 20550-900 - Maracanã, Rio De Janeiro – RJ – Brazil

Abstract
This paper presents a methodology for defining the best asset allocation for diversified investment portfolios based on a combination of the Modern Portfolio Theory and the use of a Deep Leaning model to forecast future volatility, in order to efficiently diversify the portfolio, considering both historical correlations between the assets and the forecast of individual volatility. Unlike traditional techniques, which use only the history, the proposed methodology takes into account the current trading conditions of assets, incorporating all aspects that are technically represented in the prices, specially the underlying current market conditions.

Date of Submission: 08-01-2025 Date of Acceptance: 18-01-2025

I. Asset Allocation And Risk Management

The task of defining asset allocation for an investment portfolio is highly complex, as it aims to generate higher returns while minimizing risk. Numerous studies have been conducted over the years to improve the quantification of risk and its impact on investment returns. However, it is important to recognize that the concept of risk is multifaceted and not easily defined. In Finance, risk can be defined in various ways, with the simplest definition being the probability of permanent capital losses, as discussed in [Mar18] and [Tal16].

Despite these challenges, it is necessary to quantify risks in order to improve our ability to minimize the impact of adverse events on the portfolio performance. However, it is important to recognize that not al risks can be quantified, and to remain vigilant in assessing potential risks, even those with a low probability of occurrence. Intuitively, it is easy to see that the greater the expected return on an investment, the greater its risk.

This happens because investors are risk averse, and if a low-risk investment had a more attractive return than a higher-risk one, investors would prefer the first over the second, which would lead to a reduction in the return of the first. Another relevant characteristic is that investors tend to underestimate risk during bullish market periods and overestimate it during bearish market periods, making the market moviments more abrupt than in a case where investor would have a better risk perception. Additionally, certain risks may have a low probability of occurring, but a high impact if they do occur. For example, consider the scenario of a plane crash, which has a low probability but a catastrophic impact if it occurs.

In this work, we will follow the approach adopted by most fund managers, which considers only the volatility of asset prices in the market. It is worth noting that volatility is not the sole risk associated with financial assets and is unlikely to be the most significant one, particularly if we define risk as the potential for permanent loss of capital. Nevertheless, volatility is an essential indicator of the relative risk of an asset compared to other assets and the market as a whole. Therefore, volatility is considered to be a proxy for other risks and can be easily quantified by variance of returns.

The next section will cover Modern Portfolio Theory, which uses variance as a risk metric and provides a methodology to determine the optimal asset allocation for a given portfolio, determining the weight of each asset in the portfolio composition. While there are limitations due to market inefficiencies in representing all factors that influence market performance at historical prices, Modern Portfolio Theory remains a well-structured and logical approach widely used by investment portfolio managers. Usually, portfolio management is based on the historical performance of the assets that make up the portfolio, with the strong assumption that future behavior will faithfully repeat the past. However, Data Mining techniques, as mentioned in [HW16], have great potential to improve predictive performance and enable more accurate incorporation of the current market scenario when defining asset allocation.

Investors face different types of risks when building their investment portfolios, including systematic and idiosyncratic risk. Systematic risk is also known as market risk, which refers to the overall risk that affects all assets in a particular market. On the other hand, idiosyncratic risk is the risk that is specific to a particular asset, group of assets, or industry. Creating a diversified portfolio can help mitigate unsystematic risk, but it cannot eliminate systematic risk, which affects the entire market. To mitigate idiosyncratic risk, investors can diversify their holdings within the same asset class or industry. For example, if an investor only holds stocks in the healthcare industry, they can diversify their holdings by investing in different healthcare companies with varying levels of idiosyncratic risk. This way, the risk associated with any single company's performance is reduced. However, the investor will still be subject to the systematic risk of healthcare sector.

Another way to protect against idiosyncratic risk is to conduct thorough research and analysis before investing in any particular asset. Investors can examine the financial statements, management team, competitive landscape, and industry trends of a particular company before investing in its stock. This approach can help investors identify potential red flags that may indicate higher idiosyncratic risk.

Finally, investors can also consider using financial instruments such as options, futures, or exchangetraded funds (ETFs) to hedge against both systematic and idiosyncratic risks. Options and futures contracts can provide protection against market volatility or sudden price changes in particular assets. ETFs, on the other hand, allow investors to diversify their holdings across various markets and asset classes, providing a cost-effective and convenient way to reduce portfolio risk. ETFs are usually have spread allocation in many assets, allowing the investor to be exposed to an investment thesis assuming low idiossyncratic risk.

In summary, investors can protect their investment portfolios against idiosyncratic risks by diversifying their holdings across different asset classes and markets, conducting thorough research and analysis, and using financial instruments to hedge against market volatility and sudden price changes. By combining these strategies, investors can effectively manage their portfolio risk and achieve their investment goals.

II. Assisted Asset Allocation Using Machine Learning

Portfolio managers generally apply techniques based on the history of expected return and volatility of the selected assets in order to define their allocation on the portfolio. However, these techniques usually ignore that market conditions change over time and these purely statistical models are not able to capture the particularities of the market at the time of evaluation. The objective of the proposed methodology is to train an Deep Learning model to learn, based on the past, how the market and the assets tend to behave.

The use of Deep Learning techniques for investment portfolio optimization is not an original idea itself, many researchers have already worked on that, applying very different approaches. As example is presented in [ZR20] where the authors aimed optimize an ETF portfolio Sharpe ratio. It is important to recognize that predicting future price of assets is a seemingly impossible task. Therefore, our objective is not to accurately predict the price of a stock or an ETF, commodity or any other type of asset, but to infer its probabilistic distribution, based on the model estimates and its prediction error on the test dataset. And this is relevant because, by performing an inference based on an intelligent model that considers the price dynamics, which was learned during the training process, we will be actively incorporating market conditions at the present time, not just the historical performance of the market and the assets.

In [FP18], the authors proposed a methodology that considers the dynamics of variables beyond the asset price, such as dividend yield ratio, book to market ratio, consumer price index (CPI) and "consumption wealth, income ratio". In the initial phases of this work, we strongly consider addressing variables linked to the fundamentals of assets, in addition to macroeconomic data, which certainly affect the behavior of asset prices in the financial market. However, there is an additional difficulty in using these data, its periodicity.

Overall, the use of Machine Learning in portfolio optimization is a rapidly growing field, with many researchers and practitioners exploring new techniques and approaches to improve investment decision-making. Machine Learning has the potential to transform the investment management industry by providing portfolio managers with powerful tools to make better investment decisions and improve their overall performance.

The graph in the Figure 5 shows the division by subject of the articles that make up the pre-selected list for the bibliographic review, not limited to those directly related to the topic of application of this work. It can be noted that the topic with the largest number of works is "Tranding", a category that consists of works that address applications whose objective is to deal with the modeling of high-frequency assets or derivatives, that is, those interested in the extra-short-term behavior assets, usually with the aim of profiting from arbitrages, quick trades or market anomalies.

There are at least two reasons for the great interest in high-frequency trading using machine learning methods. High-frequency trading involves using algorithms to make trades at high speeds, often measured in microseconds or milliseconds, to take advantage of small market inefficiencies. These trades can generate significant profits in a short amount of time, making it an attractive area for research and investment. But there is another technical reason for this: the availability of large amounts of data for high-frequency modeling. If the model is trained using only daily data, the availability of moving windows will be very limited (on average, there are 252 trading days per year). However, for high-frequency trading, there are several trading steps in a day, which allows overcoming one of the biggest challenges of Deep Learning: data scarcity. Essentially an Artificial Intelligence model can learn anything, as long as two conditions are met: (i) The data contains what must be learned (ie, there is something to learn) and (ii) There is a sufficient amount of data.

The categories ”Indirect Measures” and ”Asset Price Forecast” have the same frequency. ”Asset Price Forecast” includes all papers that aim to forecast the price of an asset, regardless of the methodology used. On the other hand, ”Indirect Measures” is the category for all papers that aim to forecast something related to the future behavior of the asset, but not its price directly. The most common indirect measure present in these papers is volatility. Many papers aim to forecast the volatility of an asset or the market as a whole, because of its relevance in taks like Risk Management, Hedging and Portfolio Management.

Although information about price would be much more useful than volatility, it is important to consider that the closer the forecast is to the future price, the lower the forecast reliability will be. The third and fourth most frequent subjects are ”Derivatives” and ”Portfolio”, respectively. The original objective of derivatives is to allow investors to protect their positions on the underlying asset. There are numerous papers addressing options pricing, aiming to identify discrepancies between the fair price of an option and its market price in order to trade them when such discrepancies are found, from a speculative perspective.

The ”Portfolio” group includes publications related to the optimization of investment portfolios, a subject much more related to the present work than the other topics. Despite not being the most frequently applied topic among those investigated, portfolio optimization has a large number of publications, with different methodologies that will be better described in the specific section on the topic. Observing the graph of figure 5 what can be noticed is that there is a group of subjects with a relatively similar amount of publications, ”Indirect Measures”, ”Asset Price Forecast”, ”Derivatives” and ”Portfolio”, all of them with much less publications than ”Trading”. We can understand that this behavior is greatly influenced by what has already been discussed regarding the aspects that make the use of Deep Learning attractive for algorithmic trading.

Figure 1: Number of papers per subject on the prior database selection for Literature Review

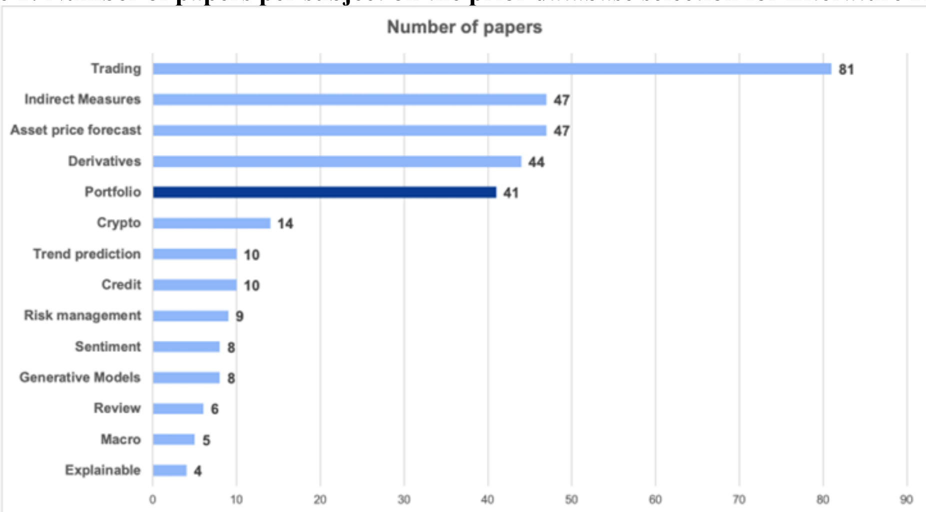

III. Fundamentals

Portfolio Theory: Classical Workflow

The Modern Portfolio Theory was developed by Harry Markowitzs and its fundamentals were introduced in [Mar52] which revolutionized the way investors think about constructing portfolios and managing risk. Although the Finance research has made significant improvements on portfolio and risk management, Markowitz's modern portfolio theory remains a key concept in portfolio management today. In general, the concepts related to the trade-off between risk and return are immune to the evolution of technology, as They are based on the logic that riskier assets have a greater expectation of return, simply by the laws of the market.

The Modern Portfolio Theory enables investors to construct a portfolio of assets that maximizes expected return for a given level of risk (measured by volatility, as quantified by the standard deviation of historical returns). This theory makes an important assumption: that investors are risk-averse, which means that for a given level of expected return, investors will always prefer the less risky portfolio. Although this may seem obvious, investors prefer higher return investments for a given level of volatility. Therefore, investors will only accept more risk if they are adequately rewarded.

An investment portfolio is comprised of various assets, each with a weight that corresponds to their portion of the total capital of the portfolio. Because the returns of assets are random variables, so are the returns of the portfolio. Let $r_k, k=1,2,...,N$ denote the return of the asset k and r_p the portfolio return. Therefore, the expected return of the portfolio is the weighted average of the expected returns of all individual assets, as shown in the following equation.

$$E[r_p] = \sum_{k=1}^{N} w_k E[r_k] \quad (1)$$

Similarly, it is possible to calculate the portfolio variance, σ_p^2, as a function of the variance of each individual asset, σ_k^2, the correlation coefficient between all combinations of two assets, ρ_{ij} (correlation coefficient between assets i and j) and the weights of each asset in the portfolio.

$$\sigma_p^2 = \sum_{k=1}^{N} w_k^2 \sigma_k^2 + \sum_{i=1}^{N} \sum_{j=1,\, j\neq i}^{N} w_i w_j \rho_{ij} \sigma_i \sigma_j \quad (2)$$

$$\rho_{ij} = \frac{cov(r_i, r_j)}{\sigma_i \sigma_j} = \frac{E\left[\left(r_i - E(r_i)\right)\left(r_j - E(r_j)\right)\right]}{\sigma_i \sigma_j} \quad (3)$$

Assets can be positively correlated (for example, Brent price and oil companies' stocks), negatively (for example, risky assets and gold) or even uncorrelated (very low correlation in Financial Market is not easy to find, because the global economy usually affects somehow the prices of all assets). The diversification is of paramount importance for a portfolio by two reasons:

The combination of assets allows the total variance of the portfolio to be smaller than the variance of individual assets and for many assets, the individual risk (idiosyncratic risk) is hugely minimized, making the correlation between assets being more relevant, which represents the second term of the equation (market risk).

Suppose an equally weighted portfolio, consisting of N assets. Therefore, the portfolio's variance can be written as

$$\sigma_p^2 = \frac{1}{N^2} \sum_{k=1}^{N} \sigma_k^2 + \sum_{i=1}^{N} \sum_{j=1,\, j\neq i}^{N} w_i w_j \rho_{ij} \sigma_i \sigma_j \quad (4)$$

Since $\frac{1}{N^2} \sum_{k=1}^{N} \sigma_k^2 \leq \frac{1}{N^2} N \times \max\{\sigma_k^2\} \to 0$ if $N \to \infty$, for many assets, the portfolio variance can be re-written as $\sigma_p^2 \approx \sum_{i=1}^{N} \sum_{j=1,\, j\neq i}^{N} w_i w_j \rho_{ij} \sigma_i \sigma_j$. Thus, diversification can almost eliminate the idiosyncratic risk, but cannot eliminate the systematic risk (market risk). Although we can use diversification to reduce the idiosyncratic risk, it will requires to include a really many assets to the portfolio and we will do that in this work by combining Exchange Traded Funds (ETFs), which are themselves a collection of assets from a given country or sector or having other specific characteristic.

When combining many assets, a very relevant question is: for a given desired expected return r_t, what is the allocation (weights of each asset) that leads to the smallest portfolio variance? This is very relevant because, as mentioned earlier, investors are risk averse, so, to obtain a certain desired expected return, they are always interested in the portfolio with the lowest variance, that is, the lowest volatility (or lowest risk).

Once we have, from historical data, the expected return on assets, their variances and correlation coefficients can be used to answer this question. Once we have this statistical measures it is enough to solve a simple minimization problem.

$$\min \sigma_p^2 = \sum_{k=1}^{N} w_k^2 \sigma_k^2 + \sum_{i=1}^{N} \sum_{j=1,\, j\neq i}^{N} w_i w_j \rho_{ij} \sigma_i \sigma_j \quad (5)$$

$$\text{Subject to} \begin{cases} \sum_{k=1}^{N} w_k E[r_k] = r_t \\ \sum_{k=1}^{N} w_k = 1 \end{cases}$$

Solving this minimization problem for all return values generates a curve called Minimal Variance Frontier, which represents the allocations that lead to the smallest variance for a given expected return. The Minimal Variance Frontier, shown in the Figure 2, is fundamental for defining the optimal portfolio according to the MPT. The upper portion of the curve is called "efficient frontier" – it is the combination of risky-assets that maximizes expected return for a given level of standard deviation. It doesn't make any sense to create a portfolio located at the lower part of the Minimal Variance Frontier, because it can be found other allocation with the same volatility level (same variance) but higher expected return. Therefore, any portfolio on this portion of the curve offers the best possible expected returns for a given level of risk.

Another fundamental concept for the Modern Portfolio Theory is the risk-free rate. The risk-free rate refers to the rate of return an investor expects to earn on an asset with zero risk. All assets carry some degree of risk, therefore, assets that generally have low default risks and fixed returns are considered risk-free. An example of a risk-free asset is a 3-month government US Treasury bill. For other countries, it could be used the government Treasury bills (Tesouro SELIC in Brazil) as risk-free rate at local currency for some analysis. However, it is paramount to understand that global diversification (meaning different currencies) can make assets that are really risk-free for one currency but not risk-free for other.

The Capital Allocation Line (CAL) is a line that depicts the risk-reward tradeoff of assets that carry idiosyncratic risk. The slope of the CAL is called the Sharpe ratio, which is the increase in expected return per additional unit of standard deviation (reward-to-risk ratio). According to the MPT, rational risk-averse investors should hold portfolios that fall on the efficient frontier (since they provide the highest possible expected returns for a given level of standard deviation). Thefore, the optimal portfolio (also called the "market portfolio") is the combination of assets which combines one risk-free asset with one risky asset.

Figure 2: Optimization problem to generate the efficient frontier and Optimal Allocation for the risk assets portfolio.

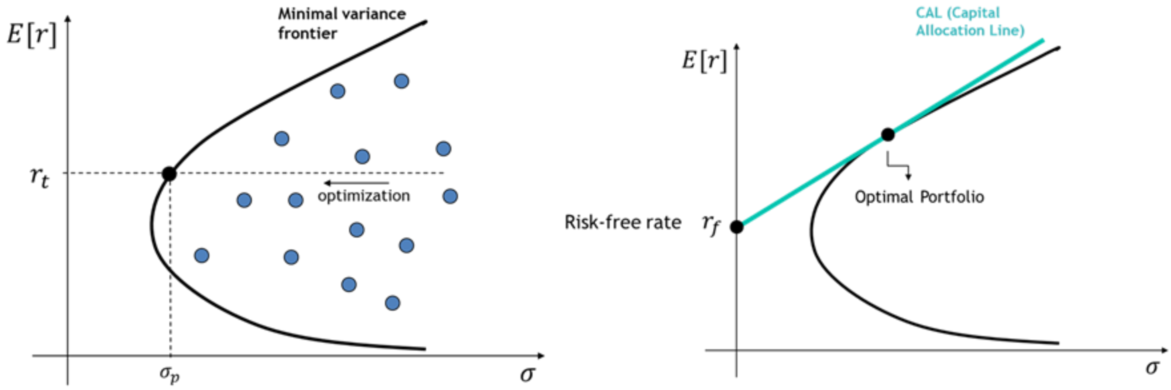

The slope of the Capital Allocation Line (CAL) is called the Sharpe Ratio (SR), which was introduced by William F. Sharpe in 1966 [Sha94]. The Sharpe Ratio is defined as the excess return of a portfolio over the risk-free rate per unit of its standard deviation. It measures the incremental expected return gained per unit of volatility and is used to evaluate the risk-adjusted performance of a portfolio. It can be calculated by the Equation 7, where E[rp] and σp are the expected return and standard deviation of the portfolio and rf is the risk-free rate.

$$SR = \frac{E[r_p] - r_f}{\sigma_p} \qquad (8)$$

MPT is not a one-size-fits-all solution and does not replace fundamental analysis or market insights. However, it provides a powerful framework for quantifying the trade-off between risk and reward and determining an optimal allocation for a given set of assets. In addition, historical data is commonly used to estimate expected returns and standard deviations, but this approach has limitations, as historical returns may not be representative of future returns. Hence, machine learning techniques can be applied to train models that can better capture the underlying relationships between asset prices and other market variables and can be used to extrapolate under uncertainty.

The Figure 3 presents the methodology we have just described in this section. The fundamental point is that the inputs for obtaining the minimun variance frontier (expected returns and variance) come from history. This is exactly what we intend to improve in this work. The methodology proposed actually uses historical data to train an intelligent model and extrapolate asset prices under uncertain conditions, in order to better represent the current market conditions.

Figure 3: Regular Asset Allocation Workflow.

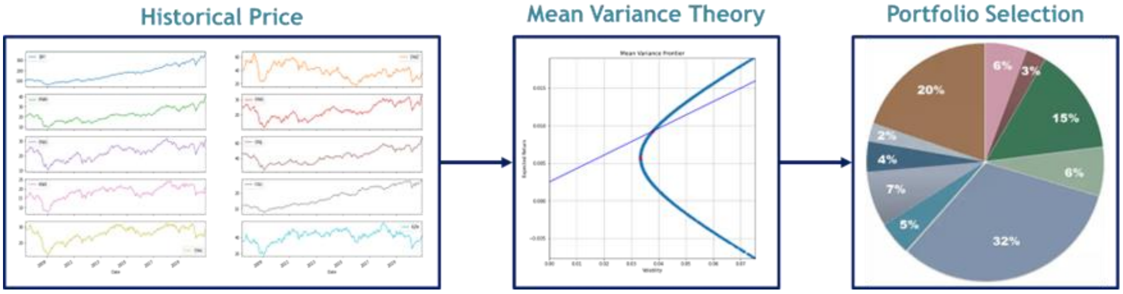

Long Short-Term Memory (LSTM)

Unlike standard feedforward neural networks, that a given layer feeds the following layer, Recurrent Neural Networks (RNN) have feedback connections. Long Short-Term Memory (LSTM) is a type of RNN that uses the concept of gates of reinforcement and forgetting to incorporate past information and use it in future predictions, allowing to process not only single data points, but also entire sequences. The network training will be carried out through sliding windows that will generate a dataset with a certain number of historical data (20) and only one forecast step.

Therefore, the LSTM can learn to predict based on a sequence of data instead of just one input, as the feedforward neural networks. It is very useful for applications related speech or video. For example, LSTM is applicable to tasks such as unsegmented, connected handwriting recognition, speech recognition and anomaly detection in network traffic. In the case of the problem we are addressing, the LSTM architecture is quite adequate, as we want to predict the future prices behavior based on the dynamics of asset prices in the past.

The Figure 4 shows LSTM architecture unit cell and its mathematical formulation. It consists of some gates containing a regular neural network operation (weighted sums and bias to be learned during the training) and activation functions. In all equations, σ represents the sigmoid function and the operator o denotes the Hadamard product (element-wise product).

Figure 4: LSTM Unit and its mathematical formulation

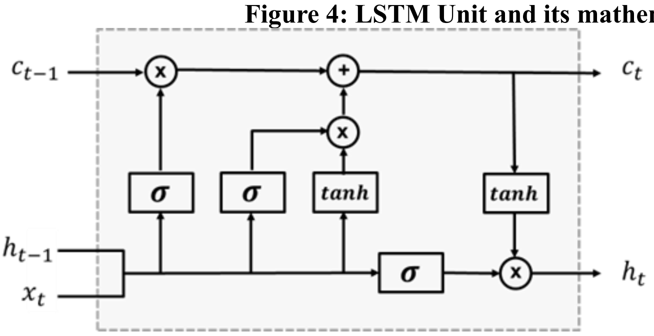

$$\begin{cases} f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f) \\ i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i) \\ o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o) \end{cases}$$

$$\begin{cases} \bar{c}_t = \tanh(W_c x_t + U_c h_{t-1} + b_c) \\ c_t = f_t \circ c_{t-1} + i_t \circ \bar{c}_t \end{cases}$$

$$h_t = o_t \circ \tanh(c_t)$$

The forget gate decides which information needs to be preserved and which can be ignored. The information from the current input x_t and hidden state h_(t-1) are combined like in a simple perceptron and passed through the sigmoid function. Sigmoid generates values between 0 and 1. It concludes whether the part of the old output is necessary, making its output closed to 1, otherwise, closed to 0. The weight matrix W_f and U_f and well as the bias vector b_f will be adjusted during the training process. The next figure shows the part of the architecture corresponding to the Forget Gate.

Input gate will learn relevant information to preserve. The architecture is exactly the same of the forget gate, the operation performed by a simple perceptron. The weight matrix W_i and U_i and well as the bias vector b_i will be adjusted during the training process. The next figure shows the part of the architecture corresponding to the Input Gate.

The objective of output gate is extracting useful information from the cell state at t-1 and from the input at t. The architecture is exactly the same of the Input and Forget Gates. The weight matrix W_o and U_o and well as the bias vector b_o will be adjusted during the training process. The next figure shows the part of the architecture corresponding to the Output Gate.

The purpose of memory cell is to consolidate what has been learned that must be forgotten and remembered. Therefore, for this step you need the results of the input and forget gates, as well as a simple perceptron type unit with tanh activation function. Usually the last trading days are more relevant for predicting the next behavior. However, it is not necessary to do any treatment in the data for the model to represent this fact.

This will be done automatically when training the gate weights and bias. Furthermore, depending on the trading volume of assets, older prices can influence more or less. However, we understand that 20 days is enough to incorporate short-term information into the price, and the long-term characteristics must be addressed a priori, when selecting assets based on fundamentals.

Hidden state vector calibrates the information that will be passed to the output. This gate has no parameters to be learned during the training process (the learning parameters were learned by the output gate).

IV. Methodology

Multi-Step Predictions under Uncertainties

Once we have estimates of future price dispersion, we can calculate the expected return and standard deviation. These values are used to apply classical portfolio theory. The methodology proposed in this paper consists of using the LSTM to predict, based on N previous time intervals, the next price of the asset. However, as our objective is to use MPT for long-term portfolio management, it is important to predict the behavior of assets some periods beyond the historical period (1 month, since we are using a montly based approach).

Thus, an important question arises: How to turn the single step estimate provided by the trained LSTM model into a multi-step estimate? First of all we have to recognize that, although LSTM has performed very well for this kind of task, the market behavior cannot be fully predicted and some significant errors are expected for some market conditions. Therefore, the proposed methodology consists of applying LSTM multiple times and adding two additional estimates based on the LSTM performance over the test dataset.

In summary, one point at t-1 generates three estimates at t: (i) LSTM estimate, (ii) LSTM estimate minus one LSTM error standard deviation and (iii) LSTM estimate plus LSTM error standard deviation. By adding these two new estimates we are creating a range of possible realizations that will lead at the end to a number of different prices for the asset under evaluation. Having these many estimates we can calculate both expected return and variance to be inputted into the MPT model to calculate the optimum allocation.

The Figure 5 schematically presents the proposed methodology. The starting point, which corresponds to the last point in the history, gives rise to 3 points, with one point (black) coming from the evaluation of the trained model and 2 points obtained by perturbation of this point (adding or subtracting a standard deviation). Thus a point generates 3, which generates 9, which generates 27, and so on.

Figure 5: Scheme of the multi-step prediction based on the proposed workflow.

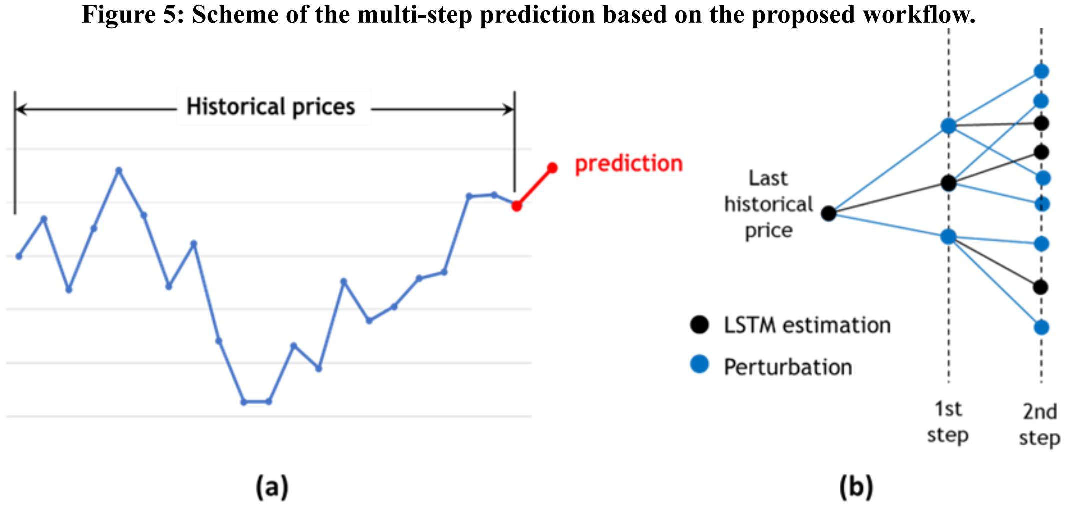

V. Application To A ETF Portfolio Optimization

Database

Since the main objective of portfolio management is to obtain good returns while reducing risk, by reducing volatility, diversification is the fundamental tool. Mutual funds usually have hundreds and even thousands of assets in their portfolios, the fundamentals of which are carefully analyzed, as well as the risks involved. In this work, we will use ETFs from 10 different countries. Each ETF contains a large amount of assets (SPY has, for example, shares from 500 companies).

Thus, we understand that sector diversification is covered because we are using ETFs and geographic diversification is covered because we are using ETFs from different countries, both developed and developing. Therefore, the dataset adopted mimic the allocation step, where the analyst will decide how much capital to invest in each asset of the selected list. Table 1 shows the ETFs selected to carry out this work.

Table 1 Selected ETFs

Ticker	Title
SPY	SPDR S&P 500 ETF Trust
EWZ	iShares MSCI Brazil ETF
EWN	iShares MSCI Netherlands ETF
EWQ	iShares MSCI France ETF
EWG	iShares MSCI Germany ETF
EWJ	iShares MSCI Japan ETF
EWA	iShares MSCI Australia ETF
FXU	First Trust Utilities AlphaDEX Fund
EWU	iShares MSCI United Kingdom ETF
EZA	iShares MSCI South Africa ETF

It can be noted in the Figure 6 that these assets have very different volatility levels, as well as their historical performance. There are assets that have appreciated strongly in the past, such as SPY and FXU, and others that had great volatility and did not provide good returns for investors in the considered time horizon. Thus, it is expected that, when carrying out an allocation analysis of a portfolio that allows for long and short positions, the best allocation will have a high weight of assets that had high return and low volatility and less weight for the more volatile assets, in addition to short position in assets that had very poor returns.

Figure 6: Time series corresponding to the 10 selected ETFs

The training, validation and testing datasets have 2678, 355 and 355 samples, respectively, considering 10% for training and 10% for validation.

Exploratory Data Analysis

The data was collected from the Yahoo Finance platform, which provides prices, trading volumes and other information at various frequencies. Since the objective of this work is not to deal with High Frequency Trading, we chose to use the closing price of related assets trades with daily frequency.

Since we are interested in diversification, it is very important to look carefully at the correlation matrix. If we look at European assets, there are very high correlations between them. One can imagine at least two reasons to explain this behavior (i) European assets have a portion of fundamentals associated to the EU and its macroeconomic dynamics is relevant to influence the behavior of the shares that compose the ETFs and (ii) These assets are denominated in Euros, and the ETF itself is denominated in USD, making that prices of these assets in USD to incorporate not only the performance of European assets, but also the exchange rate, which is the same for all ETFs.

A question that can arise when we notice that there are very high-correlated assets in the portfolio is whether they should even participate in it. Naturally, it is an allocation decision, but the fact that the assets have

DOI: 10.9790/5933-1601035970

a good correlation does not mean that the long-term performance will be similar. Just imagine that two assets go together perfectly (correlation coefficient = 1) but one varies twice the other, the more volatile asset will have an amplified long-term variation. Furthermore, according to the portfolio theory, diversification is very important to reduce the idiosyncratic risk of the assets. Note also that the decision to keep these assets in the portfolio is totally different from what it would be if we wanted to infer some property based on the combined prices of the ETFs, a situation in which it wouldn't be suitable to have so many variables with high correlation.

Another conclusion that we can take from the correlations is that assets from developing countries have a lower correlation with assets from developed countries. This tends to be positive for diversification, especially when short positioning is allowed, which is the case of the investments in which we are interested. However, assets from developing countries tend to be very volatile, both because they are less mature and riskier markets, and because these assets are originally denominated in emerging makers currencies, whose exchange rate against the dollar tends to vary widely.

Figure 7: Correlation between assets.

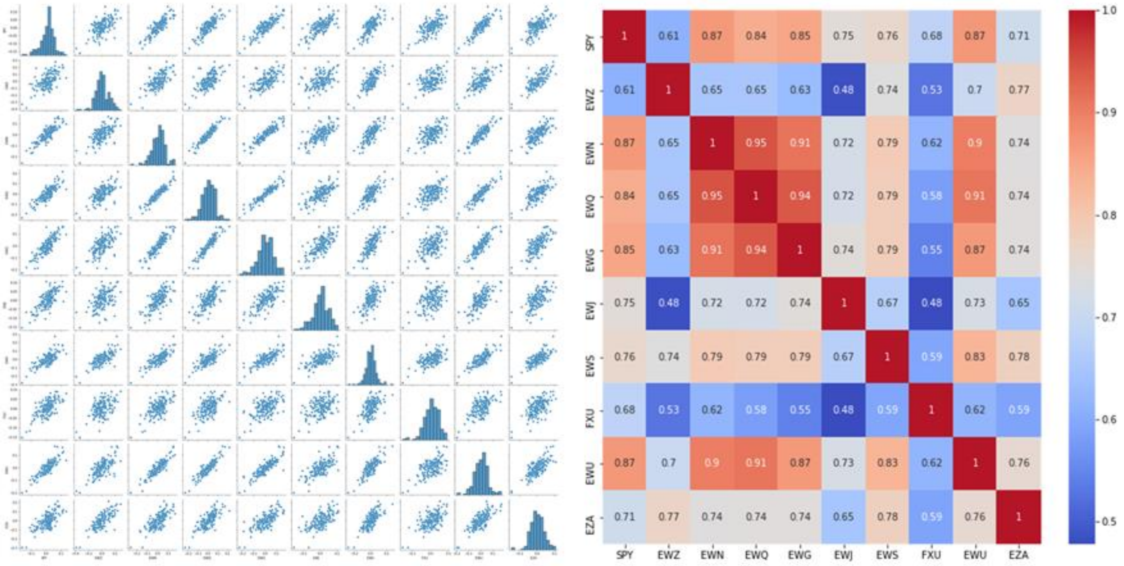

Pre-processing

Since the training will take place using sliding windows and prices vary in magnitude over time, it is important to make a normalization The proposed normalization consists of making the first price of the sliding window equal to 1

Figure 8: Proposed Normalization

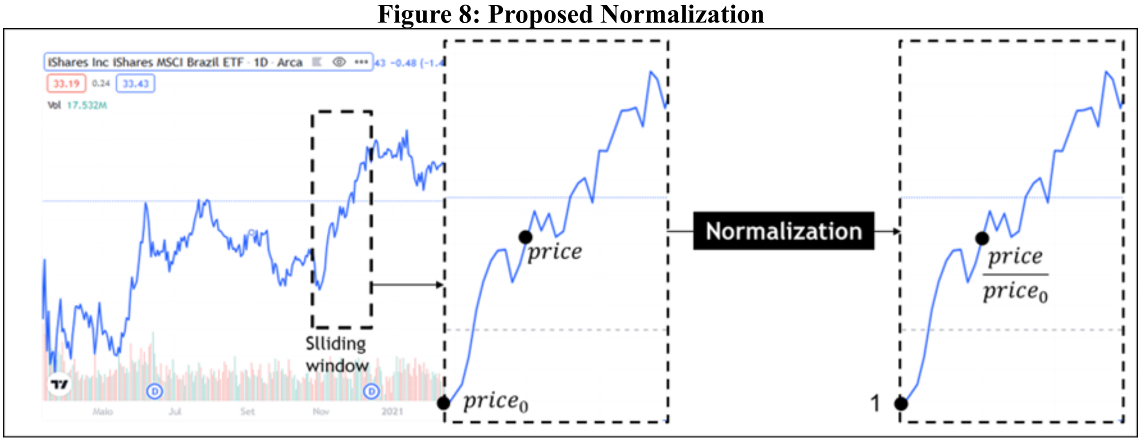

Since the assets in question are all traded in the United States, there was no need to deal with cases in which there is no data for a particular asset due to non-trading days. Furthermore, because the data came from a structured data platform and maintained by Yahoo, it was not necessary to do any missing data processing. All data on neociation days of interest were available to carry out the work. Naturally, days where no trading took place are disregarded and not reported by Yahoo Finance.

Model's Architecture

The proposed model is composed of layers with 50 LSTM units, with a probability dropout equal to 20%, in order to increase the generalizability of the trained model. At the end, a dense layer was included, which outputs the desired price in the next section of the market. The model has a total of 71051 parameters.

Figure 9: Model's arquitecture

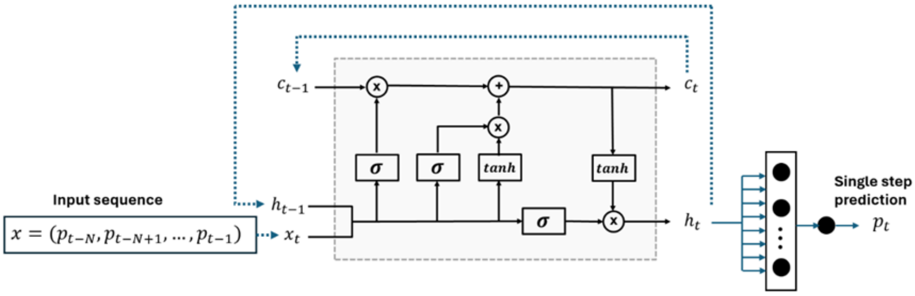

Results

Figure 10 shows the extrapolation step under uncertainty for the EWZ. Note that there are approximately 20 market sections per month and that the extrapolation was performed to 10 trading sessions (approximately two weeks). This option was motivated by the number of model evaluations that would be necessary for the extrapolation to occur over 20 trading sessions (3^20), which would make it computationally difficult. In this way, the expected value and standard deviation of the asset returns were estimated as proportional to the number of periods and proportional to the square root of the number of 10-day periods that exist in a month, as usual in Finance.

Figure 10: Extrapolation under uncertainties of the last window

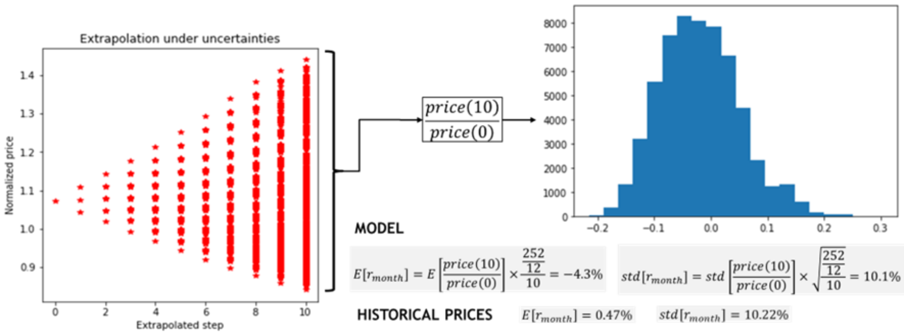

Figure 11 shows the historical performance of the optimal portfolio and its variance. Note that it is possible to obtain a portfolio that has a historical return very close to the return of the best asset (SPY), but with a lower variance than any individual asset, which is the ultimate goal of diversification. Naturally, this is only possible because we adopted long and short allocations, which allows us to leverage on the full long position. If it was not allowed to have short position (portfolio long only) this would not be possible.

Figure 11: Variance of each individual asset and the portfolio and historical return of the optimal portfolio

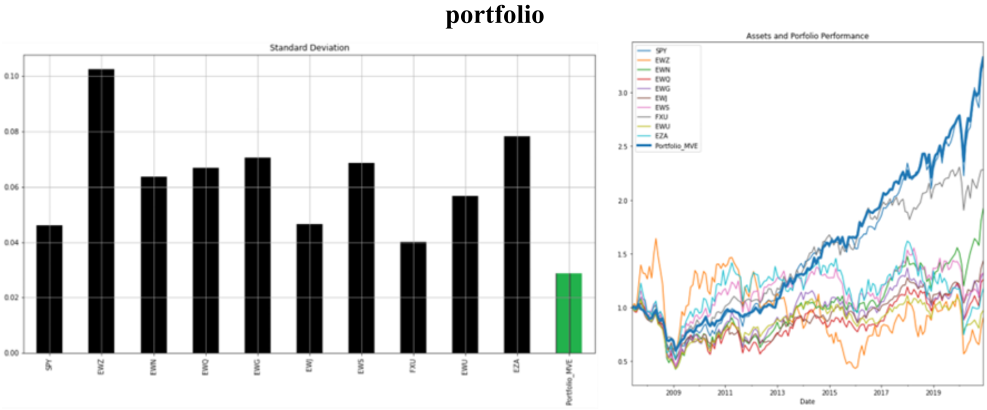

VI. Conclusions

This paper presented a methodology for using the LSTM to forecast time series under uncertainty and use these predictions to estimate expected return and variance for each asset in an investment portfolio, allowing you to use the Modern Portfolio Theory to define the best allocation for the portfolio in question.

Although the use of Machine Learning techniques, especially Deep Learning, in the Financial area has grown significantly in recent years, most of the work is focused on supporting trading strategies, especially in High Frequency Trading (HFT). Fewer works have been published on subjects related to the application of these techniques to longer-term portfolio management, typically used by mutual funds and hedge funds.

Although the performance of the models individually was not exceptional, it is important to consider that the extrapolation was performed under uncertainties, so that the error of the models led to an increase in uncertainty in the multi step procedure. Thus, it is important to consider that the eventual inaccuracy of the model reflects the unpredictability of price variation, which is greater in certain assets than in others.

Consequently, we consider that, on average, the presented methodology contributes to advancing towards using Deep Learning models to forecast volatility, rather than prices, which we believe is a more useful alternative for managing longer-term investment portfolios.

We have identified an opportunity to improve the technique by incorporating trading volume data from these assets, which we believe is very important to confirm trends (higher trading volumes) or make them less relevant (lower trading volumes) during training. A major advantage of using trading volume is the fact that the sampling is the same as for prices, unlike other variables, such as macroeconomic or corporate performance data, that are available only at some dates, depending on the asset. We hope that the incorporation of the volume will significantly contribute to the improvement of the method's performance.

References

[1] [Aiu12] Fernando Antônio Lucena Aiube. Modelos Quantitativos Em Finan.Cas: Com Enfoque Em Commodities. Bookman, 2012.

[2] [BSUM20] Eric Benhamou, David Saltiel, Sandrine Ungari, And Abhishek Mukhopadhyay. Bridging The Gap Between Markowitz Planning And Deep Reinforcement Learning. 2020.

[3] [CZHW21] Yijian Chuan, Chaoyi Zhao, Zhenrui He, And Lan Wu. The Success Of Adaboost And Its Application In Portfolio Management. 2021.

[4] [FF93] E. F. Fama And K. R. French. Common Risk Factors In The Returns On Stocks And Bonds. Journal Of Financial Economics, 33(1):3–56, 1993.

[5] [Fis18] Krauss C. Fisher, T. Deep Learning With Long Short-Term Memory Networks For Financial Market Predictions. European Journal Of Operational Research, 270:654–669, 2018.

[6] [FP18] He J. Feng, G. And J. Polson. Deep Learning For Predicting Asset Returns. Arxiv Mathematics, Computer Science, Economics. Arxiv:1804.09314v2 [Stat.ML], 2018.

[7] [HC20] J Huang And J. Chai. Deep Learning In Finance And Banking: A Literature Review And Classification. Frontiers Of Business Research In China, 2020.

[8] [HW16] N. Polson Heaton, J. And J. H. Witte. Deep Learning For Finance: Deep Portfolios. Applied Stochastic Models In Business And Industry, 33(1):3–12, 2016.

[9] [Ias20] Miquel Noguer I Alonso And Sonam Srivastava. Deep Reinforcement Learning For Asset Allocation In Us Equities. 2020.

[10] [IV22] Hugo Inzirillo And Ludovic De Villelongue. An Attention Free Long Short-Term Memory For Time Series Forecasting. 2022.

[11] [JXL17] Zhengyao Jiang, Dixing Xu, And Jinjun Liang. A Deep Reinforcement Learning Framework For The Financial Portfolio Management Problem. 2017.

[12] [KKL+23] Jiwon Kim, Moon-Ju Kang, Kanghun Lee, Hyungjun Moon, And Bo-Kwan Jeon. Deep Reinforcement Learning For Asset Allocation: Reward Clipping. 2023.

[13] [KSV16] Y. Kom Samo And A. Vervuurt. Stochastic Portfolio Theory: A Machine Learning Perspective. Proceedings Of The Thirty-Second Conference On Uncertainty In Artificial Intelligence, Page 657–665, 2016.

[14] [LPAK22] Jinho Lee, Sungwoo Park, Jungyu Ahn, And Jonghun Kwak. Etf Portfolio Construction Via Neural Network Trained On Financial Statement Data. 2022.

[15] [LYX+21] Yu Lin, Yan Yan, Jiali Xu, Ying Liao, And Feng Ma. Forecasting Stock Index Price Using The Ceemdan-Lstm Model. The North American Journal Of Economics And Finance, 57:101421, 2021.

[16] [Mar52] H. Markovitz. Portfolio Selection. The Journal Of Finance, 7(1):77–91, 1952.

[17] [Mar18] Howard Marks. The Most Important Thing Uncommon Sense For The Thoughtful Investor. Harper Business, 2018.

[18] [Men19] Xiang Meng. Dynamic Mean-Variance Portfolio Optimisation. 2019.

[19] [NE22] Hamid Nasiri And Mohammad Mehdi Ebadzadeh. Multi-Step-Ahead Stock Price Prediction Using Recurrent Fuzzy Neural Network And Variational Mode Decomposition. 2022.

[20] [OEC21] OECD. Artificial Intelligence, Machine Learning And Big Data In Finance: Opportunities, Challenges, And Implications For Policy Makers. 2021.

[21] [PCL21] Lifang Peng, Kefu Chen, And Ning Li. Predicting Stock Movements: Using Multiresolution Wavelet Reconstruction And Deep Learning In Neural Networks. Information, 12(10), 2021.

[22] [RPAGNV19] Eduardo Ramos-Perez, Pablo J. Alonso-Gonzalez, And Josâle Javier Nunez-Velazquez. Forecasting Volatility With A Stacked Model Based On A Hybridized Artificial Neural Network. Expert Systems With Applications, 129:1–9, 2019.

[23] [RWL20] Jinlong Ruan, Wei Wu, And Jiebo Luo. Stock Price Prediction Under Anomalous Circumstances. Dec 2020.

[24] [Sha94] W. Sharpe. The Sharpe Ratio. Journal Of Portfolio Management, 21(1):49—-58, 1994.

[25] [SZC+21] Srijan Sood, Zhen Zeng, Naftali Cohen, Tucker Balch, And Manuela Veloso. Visual Time Series Forecasting. In Proceedings Of The Second ACM International Conference On AI In Finance. ACM, Nov 2021.

[26] [Tal16] Nassim Nicholas Taleb. Antifragile: Things That Gain From Disorder. Random House, 2016.

[27] [War22] Xavier Warin. Deep Learning For Efficient Frontier Calculation In Finance. 2022.

[28] [ZR20] Zohren-S. Zhang, Z. And S. Roberts. Deep Learning For Portfolio Optimization. The Journal Of Financial Data Science Fall 2020, 2020.
[29] [ZSI23] Cheng Zhang, Nilam Nur Amir Sjarif, And Roslina Binti Ibrahim. Deep Learning Techniques For Financial Time Series Forecasting: A Review Of Recent Advancements: 2020-2022. 2023.