Stock Market Prediction Using Evolutionary Support Vector Machines: An Application to the ASE20 Index

Georgios Sermpinis
Andreas Karathanasopoulos
Sovan Mitra
Charalampos Stasinakis

Abstract. The main motivation for this paper is to introduce a novel hybrid method for the prediction of the directional movement of financial assets with an application to the ASE20 Greek stock index. Specifically, we use an alternative computational methodology named Evolutionary Support Vector Machine (ESVM) Stock Predictor for modeling and trading the ASE20 Greek stock index extending the universe of the examined inputs to include autoregressive inputs and moving averages of the ASE20 index and other four financial indices. The proposed hybrid method consists of a combination of genetic algorithms with support vector machines modified to uncover effective short term trading models and overcome the limitations of existing methods. For comparison purposes, the trading performance of the ESVM stock predictor is benchmarked with four traditional strategies (a Naïve strategy, a buy and hold strategy, a MACD and an ARMA models), and a MLP neural network model. As it turns out, the proposed methodology produces a higher trading performance, even during the financial crisis period, in terms of annualized return and information ratio, while providing information about the relationship between the ASE20 index and DAX30, NIKKEI225, FTSE100, S&P500 indices.

Keywords: ASE20 stock index, trading, genetic algorithms, support vector machines, leverage, transaction costs.

JEL Classification codes: C61 - Optimization Techniques; Programming Models; Dynamic Analysis, C63 - Computational Techniques; Simulation Modeling, E17 - Forecasting and Simulation: Models and Applications, E47 - Forecasting and Simulation: Models and Applications, F17 - Trade Forecasting and Simulation

1. Introduction

Stock market analysis is an area of growing quantitative financial applications. Modeling and trading financial indices remains nowadays a very challenging open problem for the scientific community. Forecasting financial time series is a difficult task because of their complexity and their nonlinear, dynamic and noisy behavior. Traditional methods such as ARMA models and moving average models fail to capture the complexity and the nonlinearities that exist in financial time series. Neural network approaches have given satisfactory results but there is clearly a need for more sophisticated techniques and approaches than the existing ones (Li, 2010).

The most important disadvantages of existing methods include the manual tuning of the algorithms parameters is the overfitting problem and the fact that modelling and trading are most of the times considered as different problems. Moreover, most existing approaches handle financial indexes as being independent and cut off from the global stock market. However, as expected several articles have been published lately reporting dependencies between various financial time series ((Doman and Doman, 2010), (Graham et al. 2013)). This phenomenon is expected to be more intense for financial indexes from smaller markets such as the ASE20 financial index. However, the dependencies of ASE20 with the rest of the financial market have not yet been studied.

The purpose of this paper is to present a novel robust method for the prediction of the directional movement of financial assets with an application to the ASE20 Greek stock index. In order to predict the movement direction of the index the present paper utilizes the Evolutionary Support Vector Machine (ESVM) Stock Predictor model which is a hybrid method of Genetic Algorithms and Support Vector Machines. Genetic algorithm (Box et al., 1994) is an evolutionary heuristic optimization algorithm which has been proved to perform extremely well in practical difficult problems where the search space is big and complicated. Support Vector Machine (SVM) is a supervised learning technique used for data analysis and pattern recognition mainly in classification problems. In our hybrid methodology, a genetic algorithm is used to optimize the SVM parameters and to find the optimal feature subset.

The proposed model is an evolutionary approach deploying genetic algorithms to tune the parameters of a Support Vector Machine Model, and select its inputs optimally in order to achieve the highest statistical and trading performance while keeping the final model simple to avoid overfitting. This model was applied to model and trade the ASE 20 index and it utilizes as candidate inputs autoregressive inputs and moving averages of the ASE20 index and of other financial indices. Therefore, our paper also tries to investigate the relationship between the ASE20, and some of the most important financial indexes worldwide. The selection of the proposed model for this problem is justified by its ability to work on high dimensional spaces, as many inputs are considered, its enhanced convergence properties and its ability to uncover high profitable trading strategies.

The ESVM Stock Predictor is benchmarked with five traditional methods (Naïve Strategy, Buy and Hold Strategy, MACD strategy, ARMA plus a Multilayer Perceptron neural network) and the results obtained seem very promising in terms of both statistical and trading metrics.

The rest of the paper is organized as follows: In Section 2, a brief literature review is presented while Section 3 describes the dataset used for our experiments. In Section 4 all benchmark models and the ESVM Stock Predictor are described in detail and in Section 5 comparative trading results for the ASE20 index are presented. Finally, Section 6 presents concluding remarks and some future directions for research.

2. Literature Review

Artificial intelligence algorithms have been extensively used for modeling and trading financial indexes in the latest decades. The initial approaches included the application of Artificial Neural Network variations such as Multilayer Perceptron Neural Networks, Higher Order Neural Networks, Recurrent Neural Networks and so on [Theofilatos et al, 2013]. The overfiting problem and the constraints of Artificial Neural Network approaches in handling problems with many inputs (curse of dimensionality) led to the introduction of Support Vector Machine (SVM) based methods.

SVM and their regression extension (Support Vector Regresion) have been applied extensively in modeling and trading various financial indexes. In 2003, two applications on SVM financial time series forecasting were developed. In (Cao and Tay, 2003), SVMs are applied to the problem of forecasting several futures contracts from the Chicago Mercantile Market showing the superiority of SVMs over Back Propagation and regularized Radial Basis Function Neural Networks. In (Kim, 2003), SVMs are used to predict the direction of change in the daily Korean composite stock index and they are benchmarked against Back Propagation Neural Networks and Case Base Reasoning. The experimental results show that SVMs outperform the other methods and that they should be considered as a promising methodology for financial time-series forecasting. In 2005, Huang *et al.* (2005) use SVM for predicting the directional movement of the NIKKEI225 index with very promising results. Furthermore in (Ince and Trafalis, 2008), Support Vector Regression is applied to the short-term forecasting of ten financial indices from the NASDAQ and outperforms all other traditional forecasting methods used. Lastly (Yuan et al. 2011) mentioned that the SVR models if can be modified in terms of changing the internal parameters such as kernel functions and lagranze multipliers can outperform and give profitable results forecasting the Taiwan Capitalization stock exchange index. Kara et al (2011) compared artificial neural networks with support vector machines when applied in modeling Instanbul Stocks and their experimental results indicated the superiority of SVMs over Neural Networks. Some recent research using SVM regression model from (Ding 2012) gave some interesting results pointing out the huge profitability and competitivity of this model. All these approaches provided encouraging results but suffered from the difficulties in tuning their parameters and selecting the most suitable inputs and from the reduced interpretability of the extracted predictors which cannot be further analyzed by experts.

Various extensions and hybrid methodologies have been proposed to overcome the limitations of simple artificial intelligence approaches. In particular many methods, such as (Yeh C. et al, 2011) and (Choudhury S. et al 2014), attempted to increase the performance of SVM based methods by either utilizing multiple Kernel Functions or deploying clustering methods alongside SVMs to increase their accuracy. Furthermore, many hybrid techniques that combine meta-heuristic algorithms with Artificial Neural Network and SVM approaches have been proposed. These approaches utilize the meta-heuristic algorithms in order tune the parameters of the deployed artificial intelligence predictor and/or select the optimal inputs for them. Some of the most representative approaches of these hybrid algorithms are the hybrid combination of Genetic Algorthms and Artificial Neural Networks

(Karathanasopoulos et al., 2010), of Genetic Algorithms and SVM (Dunis et al., 2013) and the combination of Genetic Algorithms and e-SVR (Yan, 2012). However, most of these approaches emphasize on the optimization of the statistical performance of the extracted predictions and do not optimize their trading performance. Moreover, they utilize as inputs only autoregressive inputs of a single financial index and the they utilize only simple meta-heuristic algorithms whose convergence properties could be improved.

Regarding modeling and trading ASE20 index, several papers focus on their one day ahead prediction and trading with simple trading strategies. In (Dunis et al., 2010) the problem of modeling and trading the ASE20 index is approached by using benchmark models such as a naïve strategy, a moving average model, an autoregressive moving average plus several neural network models like a Multilayer Perceptron (MLP), a Higher Order Neural Network (HONN) and a Recurrent Neural Network (RNN). The results show that the Multilayer Perceptron produces better results and outperforms all other neural networks and traditional statistical models in terms of annualized return. Finally, in (Karathanasopoulos et al., 2010) genetic algorithms are combined with a Multilayer Perceptron and this hybrid method is applied with encouraging results to the problem of predicting the one day ahead change of the ASE20 index.

The ESVM model combines genetic algorithms with Support Vector Machines. To the best of our knowledge, this is the first time that genetic algorithms have been combined with Support Vector Machines to optimize the parameters of the SVM models and on parallel select the optimal inputs among an extended list of candidate inputs for modeling and trading ASE20. Moreover, the proposed approach utilize a fitness function that enables the algorithm to extract highly profitable trading strategies.

Existing approaches in financial time-series modeling and prediction use information from previous values of the selected asset as inputs. In this research on the ASE20 index, we use previous values of that index plus ASE20, DAX30, Nikkei225, S&P500 and FTSE100 indices as inputs. Until now, the relationship between Greek stock market and foreign financial markets has not been studied thoroughly. The only important result was obtained in (Niarchos et al., 1999). In this paper, international information transmission between the U.S. and Greek stock markets was investigated using daily data from the ASE20 and the S&P500 index returns but the authors were not able to derive a clear conclusion about the influence of the U.S. stock market on the Greek one.

3. The ASE20 Index and Related Financial Data

After rapid economic growth during 1997-2000, Greece managed to join the EMU (European Monetary Union) and adopt the Euro currency on 1 January 2001. Prospects turned from moderate to very positive for the Greek economy and despite the crisis observed in 2001 and from 2009 until now, economists still believe that there still exists opportunities in investing in the Greek stock market. The Athens Stock Exchange (ASE) is the official market for share trading in Greece, for both private and institutional investors. The main indices calculated by the ASE are: the three FTSE/ASE indices (20 index, Mid Cap 40 index, Small Cap 80 Index), the ASE Composite Index of the Main Market, the ASE Composite index of the Parallel Market, the ASE sector indices, the All-share Index and the two (Main – Parallel) total return indices.

The blue chip index FTSE/ASE20 is the underlying asset for futures contracts on this index that are traded in derivatives markets. The ASE20 index is based on the 20 largest ASE stocks. It was developed in 1997 in partnership between ASE and FTSE International and is now the established benchmark for the Greek stock market. It represents over 50% of ASE's total capitalization and currently has a heavier weight on banking, telecommunication and energy stocks.

The ASE20 index is traded as a futures contract that is cash settled upon maturity of the contract with the value of the index fluctuating on a daily basis. The cash settlement of this index is simply determined by calculating the difference between the traded price and the closing price of the index on the expiration day of the contract. Whilst the futures contract is traded in index points the monetary value of the contract is calculated by multiplying the futures price by the multiple of 5 euro per point. For example, a contract trading at 1,400 points is valued at 7,000 EUR.

The ASE20 futures contract is therefore suited to institutional trading which justifies its choice for this empirical application.

Fig. 1: The ASE-20 Greek stock index

In order to train our models, we divided our dataset as follows:

Name of Period	Trading Days	Beginning	End
Total Dataset	3243	1 January 2001	31 December 2013
Training Dataset	1899	1 January 2001	11 August 2008
Validation Set	1344	12 August 2008	31 December 2013

Table 1: Total dataset

The observed ASE20 time series is non-normal (Jarque-Bera statistics confirms this at the 99% confidence interval) containing slight skewness and high kurtosis. It is also non-stationary and we decided to transform the ASE20 series into stationary series of rates of return[1].

Given the price level P1, P2,…,Pt, the rate of return at time t is formed by:

$$R_t = \left(\frac{P_t}{P_{t-1}}\right) - 1 \tag{1}$$

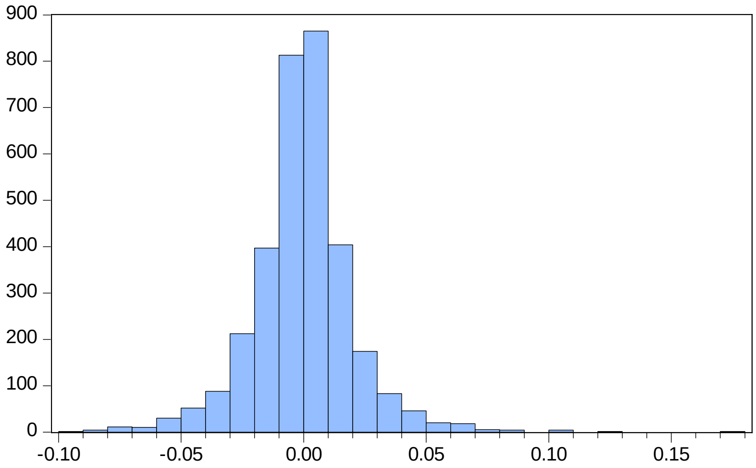

Series: RETURNS
Sample 1/03/2001 12/31/2013
Observations 3243

Mean	-0.000290
Median	0.000000
Maximum	0.177910
Minimum	-0.093318
Std. Dev.	0.020459
Skewness	0.342982
Kurtosis	7.887077
Jarque-Bera	3290.844
Probability	0.000000

Fig. 2: ASE20 returns summary statistics (total dataset)

As inputs to our algorithms and our networks, we selected 18 different inputs described in detail in Table 2 below.

Number	Variable	Lag
1	Athens Composite all share return	1
2	Athens Composite all share return	2
3	Athens Composite all share return	3
4	Athens Composite all share return	4
5	Athens Composite all share return	5
6	Athens Composite all share return	6
7	Athens Composite all share return	7
8	Athens Composite all share return	8
9	Athens Composite all share return	9
10	Athens Composite all share return	10
11	Dax30 index return	2
12	10 days moving average of Dax30 index return	2
13	Nikkei 225 index return	1
14	10 days moving average of Nikkei 225 index return	1
15	FTSE100 index return	2
16	10 days moving average of FTSE100 return	2
17	S&P 500 index return	2

[1] The percentage return is linearly additive but the log return is not linearly additive across portfolio components.

18	10 days moving average of S&P 500 return	2

Table 2: Explanatory variables

In order to train the ESVM we further divided our dataset as follows:

Name of Period	Trading Days	Beginning	End
Total Dataset	3243	1 January 2001	31 December 2013
Training Dataset	1439	1 January 2001	6 November 2006 2009
Test Dataset	460	7 November 2006	11 August 2008
Validation Set	1344	12 August 2008	31 December 2013

Table 3: ESVM Stock Predictor Datasets

In Figure 3 we show the Greek stock index for the out-of- sample period.

Fig. 3: The ASE-20 Greek stock index (out-of-sample validation period)

As inputs in our forecasting models except from previous values of the ASE20 stock index, we used previous values from other important financial indices. DAX30, S&P500 and FTSE100 indexes were included as being the most significant financial indexes in the European market where the ASE20 index is also located. NIKKEI225 was considered in our work as it is the most significant representative of the Asian Market, with great impact worldwide which additionally presented similar behavior with the ASE20 during the crisis period after 2009. The financial indices which were used are the following:

- DAX30 Index: The DAX30 (Deutscher Aktien IndeX, formerly Deutscher Aktien-Index) is the blue chip stock market index consisting of the 30 largest German companies in terms of order book volume and market capitalization trading on the Frankfurt Stock Exchange. Prices are taken from the electronic Xetra trading system.

- S&P500 Index: The Standard & Poor's 500 index is a basket of 500 stocks weighted by market value, and its performance is thought to be representative of the US stock market as a whole. Over 70% of all U.S. equity is tracked by the S&P500 which selects companies based upon market size, liquidity and sector.

- FTSE100 Index: The FTSE100 consists of the 100 largest companies by market capitalization on the London Stock Exchange. The composition of the FTSE 100 is reviewed quarterly.

- NIKKEI225 index: The Nikkei225 is the benchmark index for the Tokyo Stock Exchange (TSE). It is a price-weighted average (the unit is yen) consisting of 225 Japanese companies, and the components are reviewed once a year. Currently, the NIKKEI225 is the most widely quoted average of Japanese equities.

All examined financial indexes were downloaded from the Bloomberg

3.1 EXAMINING THE RELATION BETWEEN THE ASE20 INDEX AND OTHER FINANCIAL INDICES

In order to examine the relationship between the ASE20 index and the other financial indices we first compute the correlation between the ASE20 index values, the value of the previous day's closing price of the other financial index and the value of the moving average for the 10 previous closing prices. In Table 4 we present these results.

Financial Indices and Moving Averages	Pearson correlations (two tailed) with ASE-20 index	Significance level (2-tailed)
DAX30 index	-0.008	0.679
10-day DAX30 Moving Average	0.025	0.210
NIKKEI225 index	-0.035	0.080
10-day NIKKEI225 Moving Average	0.009	0.666
FTSE100 index	-0.008	0.702
10-day FTSE100 Moving Average	0.048	0.019
S&P500	-0.037	0.065
10-day S&P500 Moving Average	0.034	0.092

Table 4: Pearson correlations of foreign indices with the ASE20 index in the in sample dataset

Pearson's correlation measures the linear association between two variables. The values of the correlation coefficient range from -1 to 1. The sign of the correlation coefficient indicates the direction of the relationship (positive or negative). The absolute value of the correlation coefficient indicates the strength, with larger absolute values indicating stronger relationships. The significance level (or p-value) is the probability of obtaining results as extreme as the one observed. If the significance level is very small (for example 0.05 which is the default significance level used by most statistical packages like SPSS or R) then the correlation is significant and the two variables are linearly related. If the significance level is relatively large (for example, 0.50) then the correlation is not significant and the two variables are not linearly related. We can easily observe that NIKKEI225, S&P500 indices and the 10-day moving averages of FTSE100 and S&P100 indices are according to Pearson's correlation the most linearly correlated with the ASE20 Greek stock index. These conclusions encourage us to continue investigating the impact of foreign indices on the Greek stock market. In order to capture the complex, non-linear relations between foreign financial indices and ASE-20 index, the novel wrapper methodology presented in chapter 4.3 was applied.

4. Forecasting Models

4.1 Benchmark Models

In this paper, we benchmark our ESVM model with 4 traditional strategies, namely Buy and Hold strategy, an autoregressive moving average model (ARMA), a moving average convergence/divergence technical model (MACD) and a naïve strategy, plus a Multilayer (MLP) neural network.

Buy and Hold strategy is the simplest trading strategy derived when buying the index (asset) at the beginning of the review period and selling it back at the end.

The naïve strategy simply takes the most recent period change as the best prediction of the future change.

The MACD strategy used is quite simple. Two moving average series are created with different moving average lengths. The decision rule for taking positions in the market is straightforward. Positions are taken if the moving averages intersect. If the short-term moving average intersects the long-term moving average from below a 'long' position is taken. Conversely, if the long-term moving average is intersected from above a 'short' position is taken. The forecaster must use judgment when determining the number of periods n on which to base the moving averages. The combination that performed best over the in-sample sub-period was retained for out-of-sample evaluation. The model selected was a combination of the ASE20 and its 7-day moving average, namely n = 1 and 7 respectively or a (1, 7) combination. The performance of this strategy is evaluated solely in terms of trading performance.

Autoregressive moving average models (ARMA) assume that the value of a time series depends on its previous values (the autoregressive component) and on previous residual values (the moving average component) (Box et al., 1994). Using the correlogram in the training and the test sub periods as a guide we choose a restricted ARMA (21,21) model. All of its coefficients are significant at the 99% confidence interval. The null hypothesis that all coefficients (except the constant) are not significantly different from zero is rejected at the 99% confidence interval. The selected ARMA model takes the form:

$$Y_t = 0.0000213 - 0.045Y_{t-2} - 1.103Y_{t-7} + 0.036Y_{t-15} + 0.260Y_{t-21} - 0.062E_{t-2} - 1.234E_{t-7} + 0.034E_{t-15} + 0.253E_{t-21} \quad (5)$$

The model selected is retained for out-of-sample estimation. The performance of the strategy is evaluated in terms of traditional forecasting accuracy and in terms of trading performance.

Neural networks exist in several forms in the literature. Their most popular architecture, which is utilized in the present manuscript, is the Multilayer Perceptrons (MLP).

The network architecture of a 'standard' MLP looks as presented in figure 4:

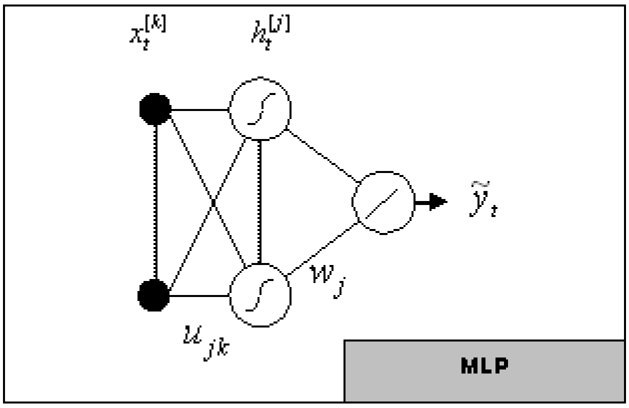

Fig. 4: A single output, fully connected MLP model

The training of the network (which is the adjustment of its weights in the way that the network maps the input value of the training data to the corresponding output value) starts with randomly chosen weights and proceeds by applying a learning algorithm called Error Back Propagation (EBP). The learning algorithm simply tries to find those weights which minimize an error function (normally the sum of all squared differences between target and actual values). The error function to be minimized is:

$$E(u_{jk}, w_j) = \frac{1}{T}\sum_{t=1}^{T}\left(y_t - \tilde{y}_t(u_{jk}, w_j)\right)^2, \quad \text{with } y_t \text{ being the target value} \tag{6}$$

4.2 Support Vector Machines

Support vector machines (SVM) are a group of supervised learning methods that can be applied to classification or regression. SVMs represent an extension to nonlinear models of the generalized algorithm developed by Vapnik (2000). They have been developed into a very active research area and have already been applied to many scientific problems. Specifically, SVM have already been applied in many prediction and classification problems in finance and economics ((Ince and Trafalis, 2008), (Cao and Tay, 2003), (Huang et al., 2005), (Kim, 2003)) although they are still far from mainstream and the few financial applications so far have only been published in statistical learning and artificial intelligence journals.

SVM models were originally defined for the classification of linearly separable classes of objects. For any particular linear separable set of two-class objects SVM are able to find the optimal hyperplanes that separates them providing the bigger margin area between the two hyperplanes. The mathematical explanation of this ability is described in Appendix A.1.

SVM can also be used to separate classes that cannot be separated with a linear classifier. In such cases, the coordinates of the objects are mapped into a feature space using nonlinear functions. The feature space in which every object is projected is a high-dimensional space in which the two classes can be separated with a linear classifier. This procedure is explained mathematically in Appendix B.1.

4.3 The Evolutionary SVM (ESVM) Stock Predictor

In this section, we describe the proposed methodology. The ESVM Stock Predictor is a hybrid method of GAs and SVMs specialized for trading financial assets.

When using SVMs, two major decisions must be made. The feature subset used as input to the classifier and the SVM parameters must be optimized. In order to optimize both, we used GA which is a heuristic evolutionary technique known for its potential in hard optimization problems.

GAs (Holland, 1995) are search algorithms inspired by the principle of natural selection. They are useful and efficient if the search space is big and complicated or there is not any available mathematical analysis of the problem. A population of candidate solutions, called *chromosomes,* is optimized via a number of evolutionary cycles and genetic operations, such as *crossovers* or *mutations*. Chromosomes consist of genes, which are the optimizing parameters. At each iteration (*generation*), a fitness function is used to evaluate each chromosome, measuring the quality of the corresponding solution, and the fittest chromosomes are selected to survive. This evolutionary process is continued until some termination criteria are met. It has been shown that GAs can deal with large search spaces and do not get trapped in local optimal solutions like other search algorithms (Holland, 1995).

In our approach, we use a simple GA where each chromosome comprises *feature genes* that encode the best feature subset and *parameter genes* that encode the best choice of parameters. Feature genes are represented with binary representation schema and parameter genes with continuous values representation. In particular, feature genes take values one or zero indicating if a specific gene should be used as input for the SVM or not. GA is also used to optimize the parameters C and γ (gamma) used by Support Vector Machines. As described in the previous section the parameter *C* is a kind of regularization parameter, that controls the tradeoff between learning capacity and training set errors and gamma is a parameter of the RBF Kernel function. The selection of these parameters is extremely critical for the performance of the trained SVM model. In our algorithm's setup, C and Gamma parameter genes take continuous values that range from 0 to 1000.

For the genetic algorithm used in our wrapper methodology, the simple one-point crossover and the mutation operators were used to avoid including even more critical parameters in the overall procedure. One-point crossover creates two offspring from every two parents. The parents are selected at random, a crossover point c_x, is selected at random, and two offspring are made by both concatenating the genes that precede c_x in the first parent with those that follow (and include) c_x in the second parent. The probability for selecting an individual as a parent for the crossover operator to be applied is named as crossover probability. The offspring produced by the crossover operator replace their parents in the population. The mutation operator places random values in randomly selected genes with a certain probability named as mutation probability. Mutation operator is very important for avoiding local optima and exploring a larger surface of the search space. Crossover and mutation probabilities for the GA where set to 0.9 and 0.1 respectively. Crossover is used in hope that new chromosomes will have good parts of old chromosomes and maybe the new chromosomes will be better. However it is good to leave some part of population survive to next generation. This is the reason a high (but not equal to one) crossover probability was used. As already mentioned, mutation is made to prevent falling GA into local extreme, but it should not occur very often, because then GA will in fact change to random search. That is the main reason why a small mutation probability was applied.

For the selection step of the GA, roulette selection (Holland, 1995) was used. In roulette selection chromosomes are selected according to their fitness. The better the chromosomes are, the more chances to be selected they have. In our approach, elitism was used to raise the evolutionary pressure in better solutions and to accelerate the evolution. By using elitism, we assured that the best solution is copied without changes to the new population, so the best solution found can survive at the end of every generation. The fitness function is defined as in equation (7)[2]:

$$fitness = accuracy + accumulated\_return \tag{7}$$

where *accuracy* is the SVM accuracy in the in sample test set and *accumulated_return* is the accumulated return of the SVM in the sample test set. In trading applications statistical accuracy is not always synonymous of financial profitability. The fitness function aims to balance the statistical accuracy and the financial effectiveness of the classifiers. Equation (7) allows our algorithm to extract prediction models that not only present high statistical performance (first factor), but also increased trading performance (second factor).

The size of the initial population was set to 30 chromosomes and the termination criterion is the maximum number of 50 generations to be reached combined with a termination method that stops the evolution when the population is deemed as converged. The population is deemed as converged when the average fitness across the current population is less than 5% away from the best fitness of the current population. Specifically, when the average fitness across the current population is less than 5% away from the best fitness of the population, the diversity of the population is very low and evolving it for more generations is unlikely to produce different and better individuals than the existing ones or the ones already examined by the algorithm in previous generations.

The flowchart of the proposed methodology is depicted in detail in Figure 5. Figure 5 starts with the description of the data splitting procedure. As described in section 3 the initial dataset is split in in-sample dataset and out-of-sample datasets. The in-sample dataset is further split in training and test datasets which are the datasets used to uncover the optimal predictors. This task is achieved through the application of the ESVM Stock Prediction methodology. Our algorithm creates a set of random solutions (every solution is described by a set of parameters and a set of selected inputs). These solutions are applied in the training dataset and evaluated using in the test dataset with the proposed fitness function. Then, the genetic operators of selection, crossover, mutation and evaluation are applied iteratively until one of the termination criteria is satisfied. Finally the best solution used to re-train an SVM model with the whole in sample dataset and the trained predictor's performance is measured on the out-of-sample dataset.

[2] We have also explored as fitness functions (equation (7)) the accumulated_return and the information ratio. In both cases, the trading performance of our models was slightly worse in the in-sample period.

The ESVM is applied to the problem of forecasting the one day ahead direction of ASE-20 Greek stock index and then trading it through future contracts. In Figure 6 we present the mean performance of the best member of the population in every generation in every algorithms execution. In the vertical axes the performance is measured using the fitness function (9). In Figure 7 we present the mean total performance of the population in every generation for every algorithms execution. In order to compute the total performance of every population, the performances of the individuals (using equation 9) are summed. Figures 6 and 7 indicate that after approximately 30 generations the proposed algorithm has covered the most promising areas of the search space and is close to convergence. This finding validated our selection of 50 generations as one of the termination criteria.

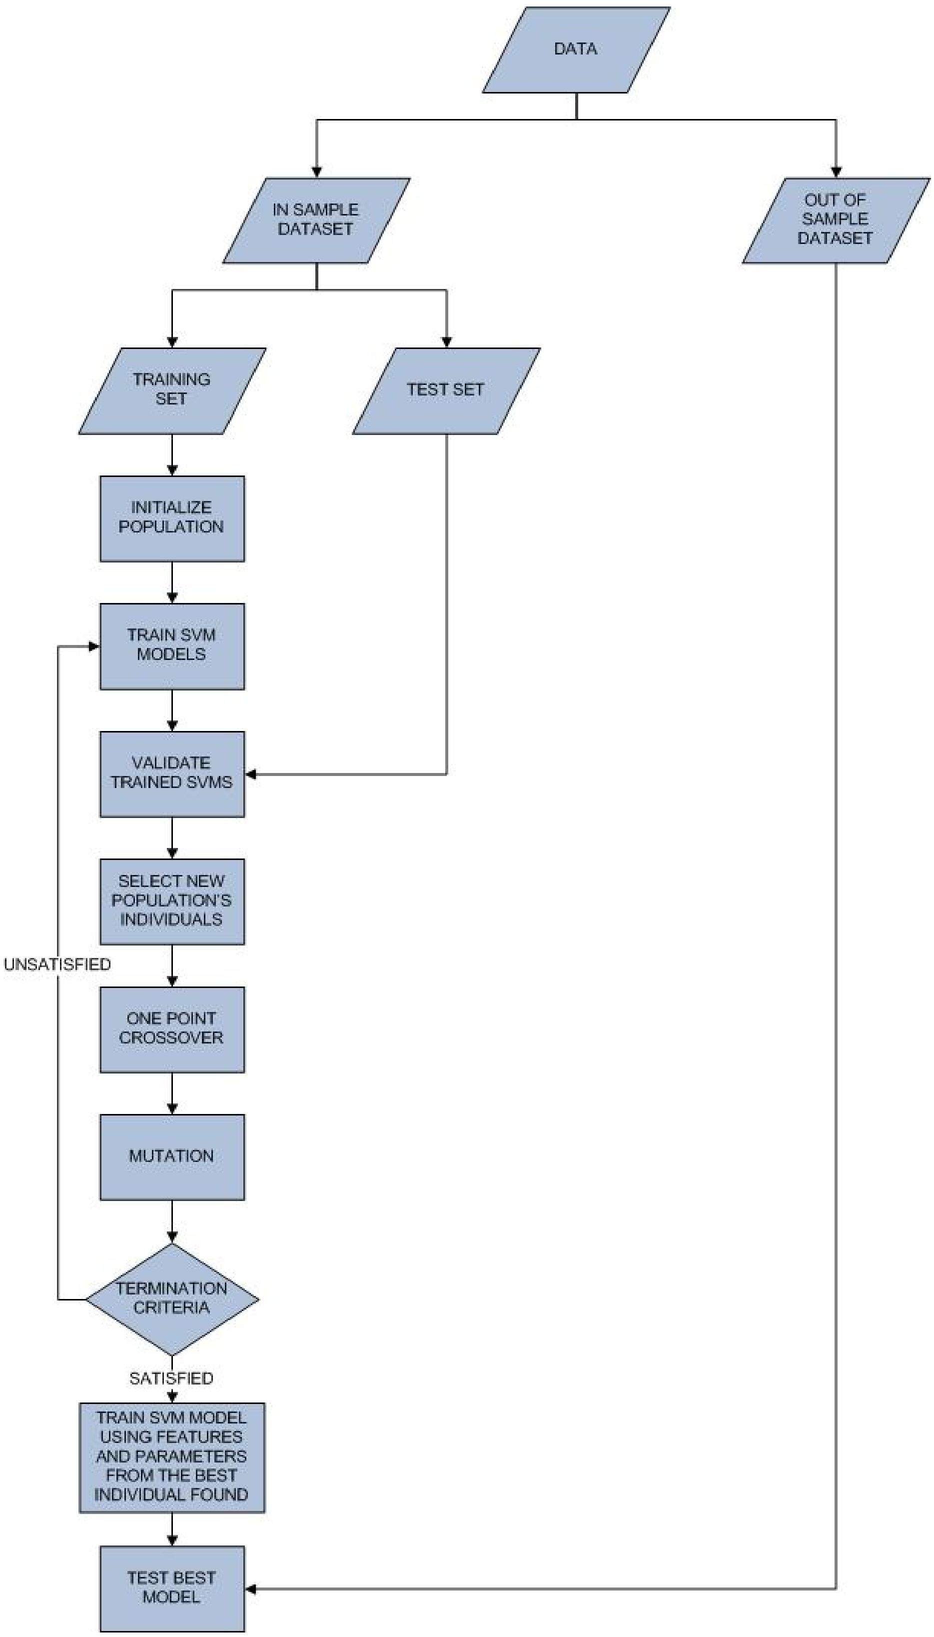

Fig. 5: Flowchart of ESVM Stock predictor's methodology

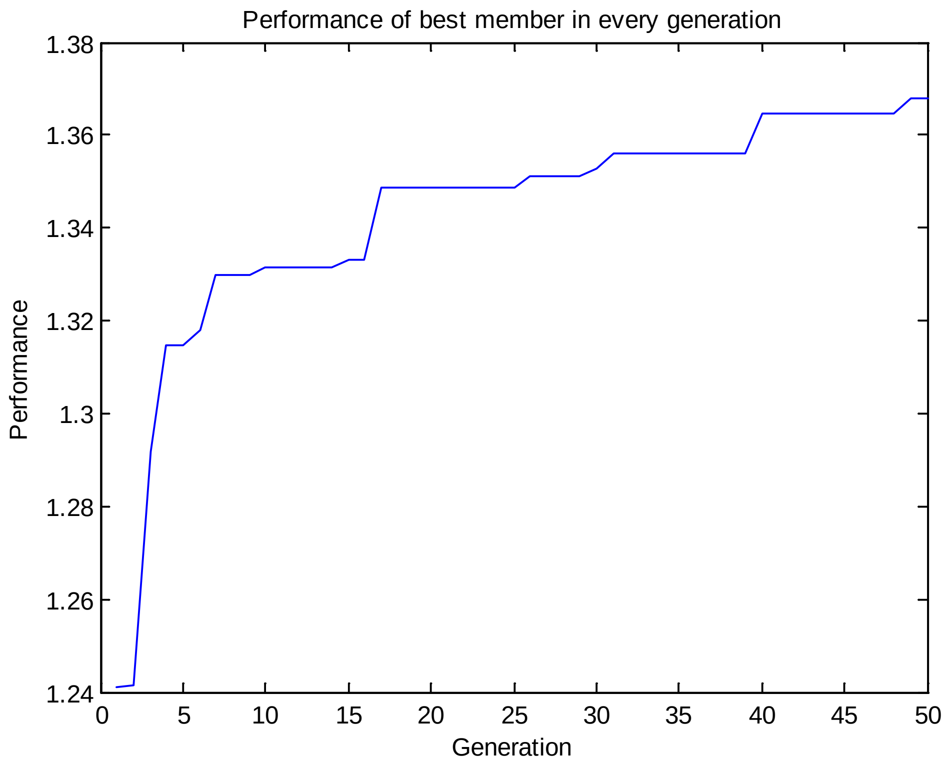

Fig. 6: Mean Performance of the best member of the population in every generation for 30 algorithms executions

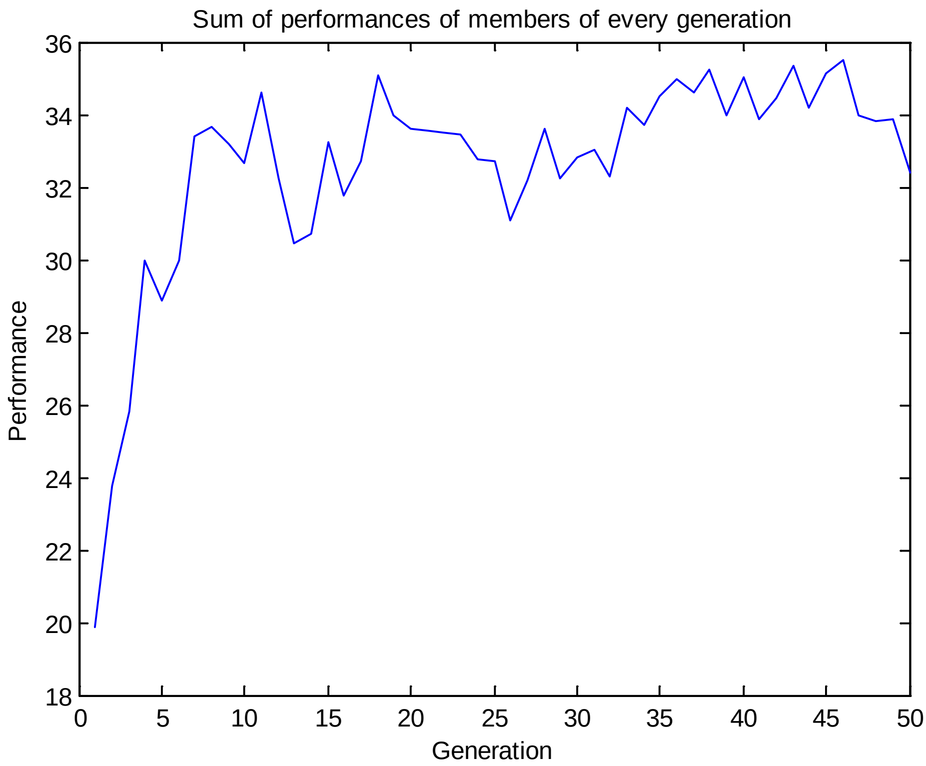

Fig. 7: Mean total performance of the population in every generation for 30 algorithms executions

The inputs that are selected in the best execution of the ESVM Stock Predictor are presented in table 5. The parameter C and gamma (see Appendix A.2) were set by the ESVM Stock Predictor to 2.82 and 274.374 respectively. This comparatively small value for the parameter C forces the ESVM Stock Predictor not to overfit the training data and thus to enhance its performance over the out-of- sample dataset.

Number	Variable	Lag
1	Athens Composite all share return N	1
2	Athens Composite all share return N	3
3	Athens Composite all share return N	5
4	Athens Composite all share return N	7
5	Athens Composite all share return N	8
6	Athens Composite all share return N	10
7	Dax30 index return N	2
8	10 days moving average of Dax30 index return N	2
9	10 days moving average of Nikkei 225 index return N	1
10	S&P 500 index return N	2
11	10 days moving average of S&P 500 return N	2

Table 5: Input variables selected by ESVM model

By further examining Table 5 we conclude that data from the FTSE100 index are not retained by the ESVM Stock Predictor even if it is strongly related to ASE20 index as shown in Table 4. The conclusion from this is that one cannot select inputs to be used in a machine learning algorithm like the ESVM Stock Predictor by resorting to a simple method like linear correlation. Simple linear methods cannot capture the complex multiple correlations that exist between the different inputs and only a more sophisticated and powerful technique like GAs can achieve this hard task.

5. Empirical Trading Simulation Results

In this section we present the results of the proposed methodology applied to trading future contracts on the ASE20 Greek stock index. These results are compared with the results of the retained benchmark models.

The trading performance of all the models considered in the out-of-sample subset is presented in the table below. Our trading strategy for the proposed methodology is simply the output of the best classifier found. Specifically, we go or stay long if the ESVM Stock Predictor's model forecasts a positive movement and go or stay short when a negative direction is forecast. The trading strategy applied in benchmark models is simple and identical for all of them: go or stay long when the forecast return is above zero and go or stay short when the forecast return is below zero. Because of the stochastic nature of the proposed methodology a simple run is not enough to measure its performance. This is the reason why 30 runs where executed and the mean results are presented in the next tables.

	Buy and Hold	**NAIVE**	**MACD**	**ARMA**	**MLP**	**ESVM**
Information Ratio (excluding costs)	-0,72	0,87	0,24	0,58	0,98	1,64
Sharpe Ratio (excluding costs)	-0,74	0,85	0,22	0,56	0,96	1,62
Correct Directional change	45,11%	49,02%	49,45%	48,96%	51,00%	54,12%
Annualised Volatility (excluding costs)	40,37%	41,25%	42,32%	40,30%	40,89%	40,05%
Annualised Return (excluding costs)	-29,05%	35,78%	10,97%	23,36%	40,44%	65,71%
Maximum Drawdown (excluding costs)	-67,24%	-31,27%	-42,26%	-40,15%	-30,75%	-39,12%
Positions Taken (annualised)	1	112	42	110	74	80

Table 6: Out of sample trading performance results

We can see that the ESVM Stock Predictor's model performs significantly better than the other benchmark methods in terms of information ratio and annualized return. In the peak of the crisis (2011-2012), ESVM generated an annualized return of 62.19% before transaction costs. In the previous years, 2009-2010 and 2010-2011, ESVM produced annualized returns of 70.27% and 67.9% respectively. A figure with the cumulative return and the drawdown of the EVSM in the out of sample is presented in Appendix A.1. The MLP outperforms other methods in terms of maximum and in terms of positions taken the MACD model, ARMA and Buy and Hold strategy trade less than all the other methods. It is noteworthy that even the simple traditional methods such as naïve trader and the MACD model performed provided profitable strategies. This is attributed to the nature of our out of sample dataset which included the financial crisis period from 2009 until 2011. During this period the positive and the negative returns are clustered. In order to verify this, we conduct the

Wald-Wolfowitz or runs test for randomness. The test confirms that the sequence of signs is not random at the 95% confidence interval.

This ability of ESVM stock predictor is mainly attributed to the external inputs it utilizes and to its increased generalization properties.

5.1 Trading costs and leverage

Up to now, we have presented the trading results of all our models without considering transaction costs. Since some of our models trade quite often, taking transaction costs into account might change the whole picture.

We therefore introduce transaction costs as well as leverage for each model. The aim is to devise a trading strategy that takes advantage of the relatively lower volatility of the return profile of some models compared to others.

5.1.1 Transaction costs

According to the Athens Stock Exchange, transaction costs for financial institutions and fund managers dealing a minimum of 143 contracts or 1 million Euros is 10 Euros per contract (round trip). Dividing this transaction cost of the 143 contracts by the average size deal (1 million Euros) gives us an average transaction cost for large players of 14 basis points or 0.14% per position.

	NAIVE	MACD	ARMA	MLP	ESVM
Annualised Return (excluding costs)	35,78%	10,97%	23,36%	40,44%	65,71%
Positions Taken (annualised)	112	42	110	74	80
Transaction costs	15,68%	5,88%	15,4%	10,36%	11,2%
Annualised Return (including costs)	20,1%	5,09%	7,82%	30,08%	54,51%

Table 7: Out-of-sample results with transaction costs

From Table 7 one can easily see that even when considering transaction costs, the ESVM Stock Predictor still significantly outperforms all other benchmark trading strategies in terms of annualized return.

5.1.2 Leverage to exploit high information ratios

In order to further improve the trading performance of our models we introduce a "level of confidence" to our forecasts, i.e. a leverage based on the test sub-period. The leverage factors applied are calculated in such a way that each model has a common volatility of 20%[3] on the test data set.

The transaction costs are calculated by taking 0.14% per position into account, while the cost of leverage (interest payments for the additional capital) is calculated

[3] Since most of the models have a volatility of about 20%, we have chosen this level as our basis. The leverage factors retained are given in table 8.

at 4% p.a. (that is 0.016% per trading day[4]). Our final results are presented in table 8 below.

Table 8 clearly shows that even when considering leverage, the ESVM Stock Predictor's model still significantly outperforms all other benchmark trading strategies in terms of annualized return.

	NAIVE	MACD	ARMA	MLP	ESVM
Information Ratio (excluding costs)	0,85	0,35	0,65	0,95	1,57
Sharpe Ratio (excluding costs)	0,83	0,33	0,63	0,93	1,55
Annualised Volatility (excluding costs)	44.44%	44.27%	43.38%	44.39%	44.50%
Annualized Return (excluding costs)	37.18%	15.77%	28.44%	42.19%	70.29%
Maximum Drawdown (excluding costs)	-33.42%	-44.26%	-44.75%	-33.85%	-45.13%
Leverage Factor	1.08	1.05	1.10	1.06	1.15
Positions Taken (annualized)	112	42	110	74	80
Transaction and leverage costs	15.40%	7.41%	19.45%	13.11%	19.72%
Annualized Return (including costs)	21,78%	8,36%	8,99%	29,08%	50,57%

Table 8: Out-of-sample trading performance - Final results

From Table 8, it is easily observed that even considering more advance trading techniques like leverage, ESVM Stock Predictor still outperforms significantly all other benchmark trading strategies in terms of annualized return and information ratio.

6. Concluding Remarks

In the present paper, we introduce a new hybrid methodology which combines genetic algorithms and support vector machines and applies it to the problem of forecasting the next day movement of the ASE20 Greek stock index extending the universe of the examined inputs to include autoregressive inputs and moving averages of the ASE20 index and of four other financial indices. For comparative purposes we also apply a naïve trading strategy, a MACD strategy, an ARMA

[4] The interest costs are calculated by considering a 4% interest rate p.a. divided by 252 trading days. In reality, leverage costs also apply during non-trading days so that we should calculate the interest costs using 360 days per year. But for the sake of simplicity, we use the approximation of 252 trading days to spread the leverage costs of non-trading days equally over the trading days. This approximation prevents us from keeping track of how many non-trading days we hold a position.

modeling approach and a MLP neural network. Previous values of the ASE20 index and of other important financial indices are used as inputs for our models.

The proposed ESVM Stock Predictor methodology produces the highest trading performance in terms of annualized return and information ratio before transaction costs. When leverage and transaction costs are considered, the ESVM Stock Predictor model continues to outperform all other benchmark models achieving higher values for the annualized return and information ratio.

It is noteworthy to mention that ESVM Stock Predictor achieves very high trading performance even if our out-of-sample period is placed in a period of crisis for the Greek Stock Market and the global economy. The proposed hybrid methodology inherits the high generalization ability of the SVM classifiers and this is the main reason for producing such a high trading performance in the out-of-sample dataset while maintaining high trading performance in the in-sample dataset. Furthermore, the genetic algorithm's global search enhances this SVM's trait and thus helps the proposed methodology to overwhelm the disparity between in-sample and out-of-sample trading performances which is one of the main problems in financial forecasting and trading (Rapacha and Wohar, 2006). Moreover, our proposed novel fitness function drives the candidate solutions to trading strategies of extremely high profit.

It is also important to note that the ESVM Stock Predictor methodology was able to uncover possible relations between the ASE20 stock index and the DAX, the NIKKEI225 and the S&P500 while showing that the FTSE100 index movements do not affect significantly the Greek stock market. These findings confirm our intuition that ASE20 stock index is highly related to other important international stock indexes. However, our results come in contrast with the results of the linear Pearson correlations presented in Table 4. One possible explanation for this contradiction is that the information given to our models by the FTSE100 index is probably the same as the information contained in the other selected inputs. Using highly correlated inputs that hold mutual information has been shown to deteriorate the performance of classifiers (Guyon and Elisseef, 2003). Furthermore, simple linear methods like Pearson correlation cannot capture the complex non-linear multiple correlations that exist between the different inputs and only a more sophisticated and powerful technique like GAs can achieve this hard task. The ESVM Stock Predictor, by using a wrapper methodology for selecting the optimal feature subset, manages to handle these correlations effectively and this is probably one of the reasons for achieving such promising results. Despite the promising results of the ESVM stock predictor and its ability to uncover the dependencies of ASE20 with several indexes, we are still far from uncovering most dependencies of ASE20 with foreign indexes and from extracting a near to optimal predictor. To achieve these goals, the universe of candidate inputs form our predictor should be enlarged to include more technical indicators such as different window moving averages, volume of transactions, volatility measures and inputs from even more foreign financial indexes such as unemployment rates, spread factors and related exchange rates. Moreover to further understand and interpret the extracted models, our approach should be completed with a mechanism to extract interpretable fuzzy rules from trained SVM models. Papadimitriou and Terzidis (2005) were initially proposed a robust methodology for this purpose and this method has already been applied for extracting fuzzy interpretable trading rules from trained svm predictors (Amorgianiotis et al, 2014).

The adaptive and evolutionary nature of our algorithm should make it successful in the less volatile stock market periods of the last years. However, its performance needs to be formally confirmed and its application for modeling and trading other financial assets is the next necessary step for our research. Furthermore, it will be interesting to use weights in the terms of the fitness function and search, how they affect the results. Another interesting future direction would be the study of the performance of our proposed methodology in periods of economic or political crises following the procedure proposed in (Kim et al., 2011).

REFERENCES

Amorgianiotis, T., Mitra, S., Theofilatos,K., Likothanassis, S. (2014) Uncovering Dependencies between European Markets using ESVM Fuzzy Inference Trader, In: Proceedings of Forecasting Financial Markets Conference, Marseilles, May 2014.

Boyaciglou, M. A. and Avci, D. (2010) An Adaptive Network-Based Fuzzy inference System (ANFIS) for Prediction of Stock Market return: The case of the Istanbul Stock Exchange . Expert Systems with Applications. 37. 7908-7912.

Box, G., Jenkins, G. and Gregory, G. (1994) *Time Series Analysis: Forecasting and Control*, Hoboken, NJ: Prentice-Hall.

Cao, L. and Tay, F. (2003) Support Vector Machine With Adaptive Parameters in Financial Time Series Forecasting. *IEEE Transactions on Neural Networks*. 14(6), 1506 – 1518.

Choudhury. S., Ghosh. S., Bhattacharya. A., Fremandes. K.R. and Tiwari. K. M., (2014). A real Time Clustering and SVM Based Price-Volatility Prediction for Optimal Trading Strategy. Neurocomputing 131. 419-426.

Cortes, C. and Vapnik, V. N. (1995) Support Vector Networks. *Machine Learning*. 20, 1-25.

Ding. Z., (2012) Application of Supporr Vector Machine Regression in Stock Price Forecasting. Business, Economics, Financial Sciences and Management Advances in Intelligent and Soft Computing. 143. 359-365.

Doman, M., and Doman, R. (2012) Dependencies between Stock Markets During the Period Including Late-2000s Financial Crisis. Procedia Economics and Finance, 1. 108-117.

Graham, M., Kiviaho, J., and Nikkinen, J. (2013) Short term and long-term dependencies of the S&P500 index and Commodity prices, 13(4), 583-592.

Dunis, C., Laws, J., Karathanasopoulos, A. (2010) Modeling and Trading the Greek Stock Market with Mixed Neural Network Models. Applied Financial Economic. 21(23), 1793-1808.

Dunis, C., Likothanassis, S., karathanasopoulos, A., Sermpinis, G. and Theofilatos, K. (2013) A hybrid genetic algorithm–support vector machine approach in the task of forecasting and trading. Journal of Asset Management , advance online publicaiton (doi:10.1057/jam.2013.2).

Guyon, I. and Elisseef, A. (2003) An Introduction to Variable and Feature Selection. Journal of Machine Learning Research. 3, 1157-1182.

Holland, J. (1995) *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*, Cambridge, Mass: MIT Press.

Huang, W., Nakamori, Y. and Wang, S. (2005) Forecasting Stock Market Movement Direction With Support Vector Machine. Computers & Operations Research. 32, 2513–2522.

Ince, H. and Trafalis, T. (2008) Short Term Forecasting with Support Vector Machines and Application to Stock Price Prediction. International Journal of General Systems. 37(6), 677–687.

Kara, W., Boyacioglou, M. E. and Baykan O. K., (2011) Predicting Direction of Stock Price Index Movement Using Artificial Neural Networks and Support vector machines: The Sample of the Istanbul Stock Exchange. Experts Systems with Applications. 38(5), 5311-5319.

Karathanasopoulos, A., Theofilatos K., Leloudas P. and Likothanassis S. (2010) Modeling the ASE20 Greek Index Using Artificial Neural Networks Combined with Genetic Algorithms, Artificial Neural Networks. Lecture Notes in Computer Science, 6352/2010, 428-435.

Kerthi, S. and Lin, C. J. (2003) Asymptotic Behaviors of Support Vector Machines with Gaussian Kernel. Neural Computation. 15, 1667-1689.

Kim, J., Shamsuddinb, A., Lim, K.P. (2011) Stock return predictability and the adaptive markets hypothesis: Evidence from century-long U.S. data. Journal of Empirical Finance, 18(5), 868-879.

Kim, K. (2003) Financial Time Series Forecasting Using Support Vector Machines, Neurocomputing. 55, 307-319.

Lee. S., Yeh. Y., and Huang C. (2014) A Multiple-Kernel Support Vector Regression Approach for Stock market Price Forecasting. Experts System with Applications. 131. 419-426.

Li, Y. (2010) Applications of Artificial Neural Networks in Financial Economics: A survey. Computational Intelligence and Design (ISCID, 2010 International Symposium on. 1, 211-214.

Niarchos, N., Tse, Y., Wu, C. and Young, A. (1999) International Transmission of Information: A Study of the Relationship Between the U.S.and Greek Stock Markets. Multinational Finance Journal.19-40.

Papadimitriou, S. and Terzidis, K. (2005) Efficient and Interpretable Fuzzy Classifiers from Data with Support Vector Learning. Intelligent Data Analysis. 9, 527-550.

Rapacha, D., Wohar, E. (2006) In-sample vs. out-of-sample tests of stock return predictability in the context of data mining. Journal of Empirical Finance. 13(2), 231-247.

Scholkopf, B., Mika, S., Burges, J. C., Knirsch, P., Muller, K.-R., Ratsch, G. and Smola, A. (1999) Input Space Versus Feature Space In Kernel-Based Methods, IEEE Transactions on Neural Networks, 10(5), 1000 - 1017.

Scholkopf, B. and Smola, A.J. (2002) *Learning with Kernels: Support Vector Machines, Regularization and Beyond*, Cambridge, Mass: MIT Press.

Theofilatos, K., Georgopoulos, E., Likothanassis, S., Mavroudi S. (2014) Computational Intelligence: Recent advances, perspectives and open problems, In: Computational Intelligence for Trading and Investment, Routledge. DOI: 978-0-415-63680-3

Vapnik, V. N. (2000) *The Nature of Statistical Learning Theory*, Springer, United States of America.

Yeh, C., Huang, C., Lee, S. (2011) A multiple-kernel support vector regression approach for stock market price forecasting. Expert Systems with Applications. 38(3), 2177–2186.

Yuan, F. (2012) Parameters Optimization Using Genetic Algorithms in Support Vector Regression for Sales Volume Forecasting, Applied Mathematics, 3(1), 1480 - 1486.

Wald, A. and Wolfowitz, J. (1940) On a test whether two samples are from the same population, Annals of Mathematical Statistics, 11, 147-162.

APPENDIX

A.1 EVSM Return-Drawdown Figure

In the figure below, the drawdown and the cumulative return of the EVSVM algorithm in the out-of-sample period is presented.

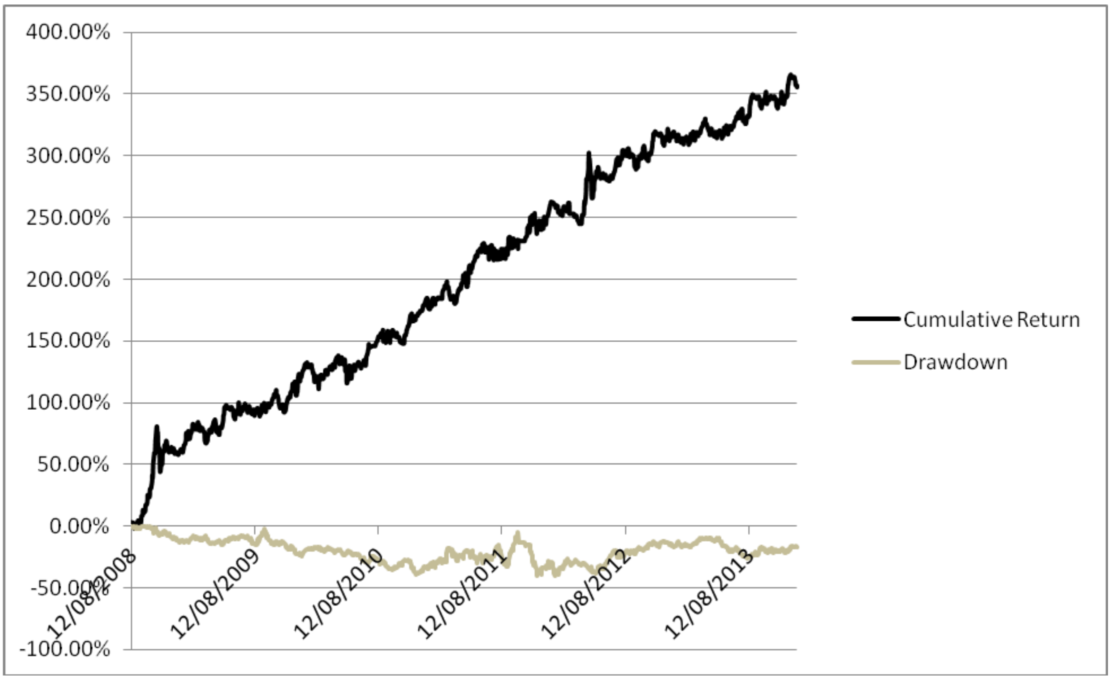

Fig. 8: Cumulative return and drawdown of the ESVM in the out-of-sample