Enhancing Profitability through Interpretability in Algorithmic Trading with a Multiobjective Evolutionary Fuzzy System

Adam Ghandar[1], Zbigniew Michalewicz[1,2], and Ralf Zurbruegg[3]

[1] School of Computer Science, University of Adelaide, Adelaide, SA 5005, Australia
[2] Institute of Computer Science, Polish Academy of Sciences, ul. Ordona 21, 01-237 Warsaw, Poland, and Polish-Japanese Institute of Information Technology, ul. Koszykowa 86, 02-008 Warsaw, Poland
[3] Business School, University of Adelaide, Adelaide, SA 5005, Australia

Abstract. This paper examines the interaction of decision model complexity and utility in a computational intelligence system for algorithmic trading. An empirical analysis is undertaken which makes use of recent developments in multiobjective evolutionary fuzzy systems (MOEFS) to produce and evaluate a Pareto set of rulebases that balance conflicting criteria. This results in strong evidence that controlling portfolio risk and return in this and other similar methodologies by selecting for interpretability is feasible. Furthermore, while investigating these properties we contribute to a growing body of evidence that stochastic systems based on natural computing techniques can deliver results that outperform the market.

1 Introduction

Algorithmic trading is an important part of the global financial services industry. In 2008 over 40% of executed market orders were attributed to algorithmic trading methods in major developed stock markets. Growth in the volume of trades generated by automatic signals has risen 30-40% per annum since, with the financial crisis over the period having little effect on the uptake of technology [1]. Financial portfolio management is a complex task that takes place in a highly dynamic and competitive environment with arguably immeasurable uncertainty. These factors make the problem quite different from other applications of computational intelligence in control, pattern recognition, etc, even though the tools and methods used are the same. It is therefore of value to examine the relationships between system design and parameters and specific and extensive performance markers and tools in financial applications distinctly.

This paper applies an evolving fuzzy system based on the representation described in [2] and from a technical viewpoint significantly extends that earlier research to make the models which are learned even closer to those used by financial practitioners. This is done by adding fundamental data (accounting and macro economic information) and by making use of a multiobjective EA to implement criteria for model parsimony. Here we make a contribution to answering

C.A. Coello Coello et al. (Eds.): PPSN 2012, Part II, LNCS 7492, pp. 42–51, 2012.

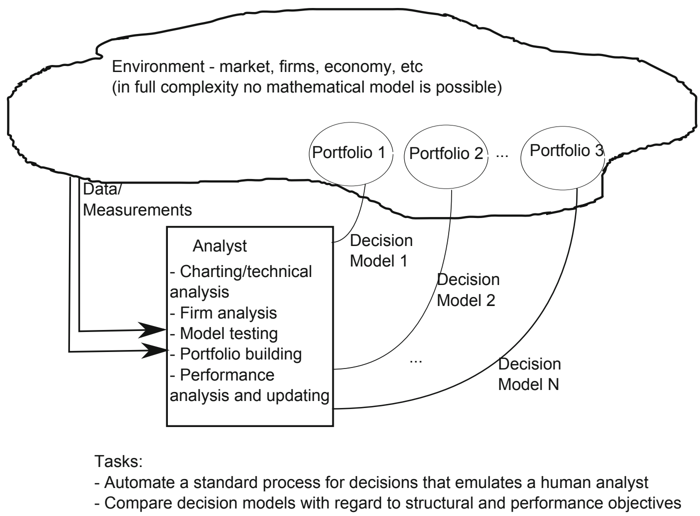

Fig. 1. Practical problem: design a control system to consolidate the tasks of the human portfolio analyst

questions about whether subjective criteria (such as human interpretability) can produce value in the application of heuristic, and specifically computational intelligence, approaches to problem solving in financial trading - a complex and dynamic activity performed on the basis of incomplete information. We find that by controlling rule intelligibility we are able to strongly influence the risk and return profile of portfolios managed by the algorithm. This angle is simply not able to be considered in classical modeling as it is currently espoused in finance because a trading model is viewed in simple terms essentially as a formula rather than an intelligent agent.

The particular methodology used to produce interpretable models is as justified as other techniques capable of representing expressive models computationally because it too has been found to perform at a standard equivalent to other state of the art machine learning approaches on standard benchmarks [3]. Many other multiobjective approaches are described in the literature [4]. The conclusions made here are also to some extent applicable to other methods. In addition, certain aspects of the particular approach make it particularly suitable for this analysis, notably the the structure that is imposed on the decision model representation to cause the optimzation process to in some respects imitate human reasoning in the application domain.

Rule based approximate reasoning systems can be applied to approximate human thinking when combined with a self-learning methodology. Efforts to automate human thinking in this way has been found to lead to improved performance in real world applications: for example in emulating a skilled human operator who controls complex machinery "without a formal quantitative model in mind" [5]. Figure 1 illustrates the rationale of the system presented here in adapting approximate reasoning principles to financial modeling. A financial analyst generally performs a sequence of specific activities: generating or selecting/tuning a range of indicators (these are often called rules in financial parlance and should not be confused with the term rulebase as it is used here) from raw data measurements or feeds; testing resultant models using historic data; building a portfolio using the model and lastly there is constant process of updating the approach based on portfolio management performance.

In the remainder of the paper we describe the design of the system in performing the tasks illustrated in Figure 1, provide empirical results on performance of different portfolios while varying model complexity, and make conclusions applicable more generally regarding the complexity of decision models in the financial problem domain.

2 Methodology

The approach mimics a real financial analyst (see Figure 1). The key tasks are: data measurement; transformation to generate unit strategies based on historic price and volume data (technical analysis) and information about the firms underlying the stocks (fundamental analysis); selecting and combining the numerous models in the previous step into a genotype comprising a fuzzy rulebase and a vector of parameters for the strategies; and lastly the implementation of decisions on portfolio contents. A multiobjective EA facilitates the study of model complexity.

2.1 Measurement, Information Set, and Initial Models

The information set defines the universe of discourse that can be represented from the environment. Technical (price/volume) and fundamental information on underlying firms is incorporated.

Fundamental Strategies. A set of fundamental variables found to be useful in other emprical studies was selected. The relationship between the variables considered and price changes has been found to be farily transient further accentuating the need for an adaptive methodology able to be achieved with a heuristic approach. Dividend yield measures the cash flow benefit with regard to the share investment. The power of dividend yield to forecast stock return has been noted in [6] as being “temporary” component of prices. Price to Book Value has been used in fundamental factor models for some time. The price Earnings ratio (PE) divides the share price, over the total earnings per share, in some periods and markets the PE ratio is a predictor of higher return with reduced risk, see [7]. Earnings per share (EPS) is calculated by taking the ratio of the profit of the firm over its market share. Stocks that have a higher earnings per share generate more income relative to the stock price and thus places upward pressure on the share price [8]. The debt to equity ratio looks at the liabilities per share; it has been shown to be positively correlated with stock price [9].

In a falling or volatile market liabilities may be of more importance. The last three fundamental variables (earnings before interest and tax, return on assets, and return on equity) divide firm income into classifications provide more fine grained picture.

Each variable is processed to obtain a rate of change measurement in the form of an oscillator (O). The oscillator measures the change relative to an earlier point $O = (v_t - v_{t-m})/v_{t-m}$. The parameter m measures this period and belongs to the set $10, 20, 30, \ldots, 260$.

Technical Strategies. Technical strategies use price and volume data. They are widely used in industry and theoretical justifications postulate the importance of behavioural factors and their detection using technical rules [10]. Technical trading rules may also be able to pick up institutional trading activity [11].

The technical inputs is given in table 1. The abbreviations have meaning: SMA, single moving average; DMA, double moving average; PPO, price oscillator; OBV, on balance volume indicator; RSI, relative strength index; MFI, money flow index; Vol. DMA, volume double moving average; PVO, percentage volume oscillator; DMI, directional movement index; %R, percent R. For the OBV indicator the value obv_t for each day t is calculated by initially at $t = 0$ $obv_0 = v_0$, then for each subsequent day t: if $p_t > p_{t-20}$ then $obv_t = obv_{t-1} + v_t$; else if $p_t < p_{t-20}$ then $obv_t = obv_{t-1} - v_t$, else if if $p_t > p_{t-20}$ then $obv_t = obv_{t-1}$.

2.2 Decision Model Representation

A solution aggregates the inputs from the processing described in the previous subsections and is represented using a set of fuzzy rules and an integer vector of parameters (time horizons - see restrictions in Table 1).

A rulebase is a mapping

$$D : \Re^n \rightarrow \Omega,$$

from vector of observations $\mathbf{x} = \{x_1, \ldots, x_n\} \in \Re^n$ to a signals $\{\omega_1, \ldots, \omega_c\} \in \Omega$ to buy or sell. The form of the rules is based on [13], the output is interpreted as a degree of certainty the buy or sell signal is correct, given antecedent and training data. A rule r_k in a ruleset M has the format: R_k : if x_1 is $A_1 \wedge \ldots \wedge$ x_n is A_n; then $(z_{k,1}, \ldots, z_{k,c})$ where $x_1 \ldots x_n$ are feature observations that are described by linguistic labels $A_1 \ldots A_n$, these are common in the different rules and precalculated as in [2].

$$z_{k,i} = \frac{\text{Sum of matching degrees of rule with } \omega_i}{\text{Sum of total matching degrees for all rules}}.$$

The mapping D uses the max operation to aggregate rules and the product t-norm to aggregate the antecedent conjunctions. The degree of certainty for i-th signal is $D : eval_i(\mathbf{x}) = \max_{k=1}^{M} \left\{ z_{k,i} \prod_{j=1}^{n} \{\mu_j(x_j)\} \right\}$, In this way, we specify a search space of possible rules to correspond to trading rules that a human expert trader could construct using the same information. A rule is

Table 1. Technical indicators and restrictions on parameters

Name	Formula	Restrictions
Price Change 1		$\delta = 20$
Price Change 2	$\ln\left(\frac{p_t}{p_{t-\delta}}\right)$	$\delta = 50$
Price Change 3		$\delta = 100$
SMA Buy	$\frac{p_t}{ma_t}$	$len_{ma} \in \{10, 20, 30\}$
SMA Sell	$\frac{ma_t}{p_t}$	$len_{ma} \in \{10, 20, 30\}$
DMA Buy 1		$len_{ma2} \in \{10, 20, 30\}$ $len_{ma2} \in \{40, 50, 60\}$
DMA Buy 2	$\frac{ma1_t}{ma2_t}$	$len_{ma1} \in \{60, 70, \ldots, 120\}$ $len_{ma2} \in \{130, 140, \ldots, 240\}$
DMA Buy 3		$len_{ma1} \in \{60, 70, \ldots, 120\}$ $len_{ma2} \in \{130, 140, \ldots, 240\}$
DMA Sell 1		$len_{ma2} \in \{10, 20, 30\}$ $len_{ma2} \in \{40, 50, 60\}$
DMA Sell 2	$\frac{ma2_t}{ma1_t}$	$len_{ma1} \in \{60, 70, \ldots, 120\}$ $len_{ma2} \in \{130, 140, \ldots, 240\}$
DMA Sell 3		$len_{ma1} \in \{60, 70, \ldots, 120\}$ $len_{ma2} \in \{130, 140, \ldots, 240\}$
PPO 1		$len_{ma2} \in \{10, 20, 30\}$ $len_{ma2} \in \{40, 50, 60\}$
PPO 2	$\frac{ma1_t - ma2_t}{ma1_t} \times 100$	$len_{ma1} \in \{60, 70, \ldots, 120\}$
PPO 3		$len_{ma2} \in \{130, 140, \ldots, 240\}$ $len_{ma1} \in \{60, 70, \ldots, 120\}$ $len_{ma2} \in \{130, 140, \ldots, 240\}$
DMI	see [12]	
%R	$\%R = \frac{p_t - min[p_{t-1}, \ldots, pt-10]}{max[p_{t-1}, \ldots, pt-10] - min[p_{t-1}, \ldots, pt-10]}$	
RSI	$RSI = 100 - \frac{100}{1+RS}$	$RS = \frac{totalgains \div n}{totallosses \div n}$
MFI	$MFI = 100 - \frac{100}{1+MR}$	$MR = \sum \frac{MF^+}{MF^-}$ $MF^+ = p_i \times v_t, where p_i > p_{i-1}$, and $MF^- = p_i \times v_t, where p_i < p_{i-1}$
Vol. DMA Buy 1	$\frac{vma1_t}{vma2_t}$	$len_{vma1} = 5, len_{vma2} = 20$
Vol. DMA Buy 2		$len_{vma1} = 20, len_{vma2} = 100$
Vol. DMA Sell 1	$\frac{vma2_t}{vma1_t}$	$len_{vma1} = 5, len_{vma2} = 20$
Vol. DMA Sell 2		$len_{vma1} = 20, len_{vma1} = 100$
OBV Buy	$\frac{\left(p_t - \max\left[p_{t-1}, \ldots p_{t-n}\right]\right)}{p_t} + \frac{\left(\max\left[obv_{t-1}, \ldots obv_{t-n}\right] - obv_t\right)}{obv_t}$	
OBV Sell	$\frac{\left(\min\left[p_{t-1}, \ldots, p_{t-n}\right] - p_t\right)}{p_t} + \frac{\left(obv_t - \min\left[obv_{t-1}, \ldots, obv_{t-n}\right]\right)}{obv_t}$	
PVO 1	$\frac{ma1_t - ma2_t}{SMt} \times 100$	$len_{ma1} = 5, len_{ma1} = 20$
PVO 2		$len_{ma1} = 20, len_{ma1} = 100$
Bol 1	$Bol = \frac{p_t - ma_t}{2 \times sd(P_t, \ldots, P_{t-\delta})}$	$\delta = 20$
Bol 2		$\delta = 50$

considered in this paper to be a statement in structured language that specifies that if some condition(s) hold, then a particular action ought to be taken. IF [Conditions], THEN do [buy/sell a particular stock] Each rulebase comprises several such statements. The genotype representation and operators are provided in [2].

2.3 Learning Process

The learning model uses a pareto based algorithm (SPEA2, [14]) to obtain a set of solutions balancing objectives of solution simplicity and in sample prediction accuracy. We can express the task to learn rules as an optimization problem in which there are two main criteria. These are to minimize the number of false signals produced by the strategies (accuracy), and reduce the model complexity:

$$\text{minimize } \mathbf{z} = f_{error}(\mathbf{x}), f_{complexity}(\mathbf{x}).$$

The accuracy objective performance is the error in determining buy and sell signals from a set of examples in training data T:

$$f_{error}(\mathbf{x}) = 1 - \frac{\#\text{ correct signals}}{\#\text{ false signals} + \#\text{ correct signals}}.$$

The number of correct signals is the count of the number of times the rulebase correctly anticipated a rise or fall in the share price, a false signal is the number of times the rulebase falsly predicted a rise or fall in the share price. The complexity of a strategy specified by a rulebase is defined by the number of rules and, within each rule, by the number of clauses.

A rulebase has two main sources of complexity which are the number of rules, and inputs quantified per rule. Therefore, $f_{complexity}$ may be decomposed into these components (which are modeled as separate objectives using the multiobjective evolutionary algorithm). In this paper we consider two complexity objectives, the *number of rules*, and the *average number of inputs used per rule (#inputs/#rules)* in the rulebase. Given two fuzzy rulebase solutions, $\mathbf{x_1}$ and $\mathbf{x_2}$, we can say that $\mathbf{x_1}$ dominates $\mathbf{x_2}$ if it is less than or equal to the other in all objectives being minimized or otherwise that it dominates in particular one of the objectives. The final source of complexity is from the definition of the linguistic variables - we set this deterministically prior to running the optimizer so it is not an objective here. The approach is verified using classical benchmarks such as Iris in [3].

3 Financial Portfolio Management

These experiments used historic data from the Standard and Poor ASX 200 oil and gas stocks between April 2000 and November 2009. All data was sourced from Data International. Instead of raw price data we used a total return index adjusted for stock splits, mergers and dividend payments. The oil and gas sector is very volatile in this period (global financial crisis and reversal as a result of the resources trade with China). There was a period of growth followed by fall and a subsequent recovery, this allowed for an evaluation of the system in three different epochs. Transaction costs were 0.25% to buy or sell, population size for SPEA was 100 and the archive size was 10. In the business model, transactions inspired by the decision model were applied to re-balance portfolios to sell and buy recommended stocks at fixed intervals of 20 trading days, for adaptation the optimization process was run before each re-balance using the previous year of data for "training".

Two benchmarks are used for comparison the first is a a buy and hold approach (BH) and the second is a standard active alpha strategy [1]. The alpha is based on the single-factor pricing model which relates a stocks excess return, $r_{i,t} - r_{f,t}$ to market return as follows

$$r_{i,t} - r_{f,t} = \alpha_i + \beta_i \left[r_{m,t} - r_{f,t} \right] + e_{i,t},$$

Table 2. Metrics of performance for the portfolios managed by solutions with varying structural complexity (SC) and linguistic complexity (LC) along with the 2 benchmark portfolios. Confidence bounds are at the 90% level based on 30 test runs.

	HP Ret	An. Ret	σ	IR	Sel.	Net Sel.
Benchmarks						
BH	6.1325	0.2899	0.3654	0.5836	0.2956	0.1703
Alpha	12.7924	0.3873	0.3988	0.7931	0.4095	0.2697
Low LC						
Low SC	16.0 [± 4.74]	0.366 [± 0.0375]	0.501 [± 0.0279]	0.699 [± 0.0694]	0.438 [± 0.0405]	0.262 [± 0.0369
Medium SC	10.4 [± 2.94]	0.314 [± 0.0351]	0.418 [± 0.0218]	0.621 [± 0.0773]	0.345 [± 0.0362]	0.199 [± 0.0351]
High SC	12.7 [± 4.55]	0.318 [± 0.0421]	0.453 [± 0.0297]	0.617 [± 0.0811]	0.368 [± 0.0450]	0.208 [± 0.0407]
Medium LC						
Low SC	16.5 [± 5.99]	0.368 [± 0.0372]	0.490 [± 0.0587]	0.667 [± 0.0621]	0.419 [± 0.0522]	0.246 [± 0.0366]
Medium SC	13.3 [± 3.04]	0.355 [± 0.0330]	0.526 [± 0.0581]	0.647 [± 0.0617]	0.421 [± 0.0408]	0.235 [± 0.0311]
High SC	10.8 [± 2.32]	0.337 [± 0.0287]	0.491 [± 0.0396]	0.615 [± 0.0563]	0.388 [± 0.0359]	0.215 [± 0.0292]
High LC						
Low SC	12.7 [± 1.80]	0.375 [± 0.0179]	0.394 [± 0.0152]	0.764 [± 0.0377]	0.394 [± 0.0204]	0.256 [± 0.0178]
Medium SC	31.6 [± 7.65]	0.482 [± 0.0279]	0.490 [± 0.0179]	0.912 [± 0.0456]	0.543 [± 0.0332]	0.370 [± 0.0290]
High SC	30.8 [± 7.67]	0.477 [± 0.0292]	0.491 [± 0.0214]	0.903 [± 0.0504]	0.537 [± 0.0341]	0.364 [± 0.0302]

index i indicates the stock and t refers to a day, e is an error term and $r_{i,t}$ is the stocks return on day t, $r_{f,t}$ is the risk free rate, $r_{m,t}$ is the market return. In an efficient market it would be possible to price stocks based solely on their risk components, here the excess return above the market. In the ideal theoretic case returns of any stock above the risk-free rate can be fully explained by the risk component, meaning that β_i would be one and α_i zero. If the alpha is actually positive the stock is outperforming relative to its level of risk and should be bought (conversely if it is negative the stock should be sold). In testing, the alpha portfolio produced an annual return of 38% compared with the buy and hold strategy which resulted in 29.56%. The information ratio of the alpha portfolio was 0.78 and for the buy and hold it was 0.58. Other metrics are given in table 2.

With some levels of complexity, the system was able to out perform both the active and passive benchmarks. Table 2 shows performance metrics achieved while varying linguistic and structural complexity parameters. Linguistic complexity refers to the granularity of the membership functions, and structural complexity refers to the relative number of rules and inputs per rule (this was controlled by selecting solutions from the Pareto front with different trade-offs between performance and complexity objectives).

As well as volatility, we consider risk using two more refined metrics. The first of these, the information ratio (IR), is superior to the commonly used Sharpe ratio (return/volatility) as it considers excess return. It is calculated as the portfolios average return in excess of the benchmark portfolio over the standard deviation of this excess return. It is used to evaluate the active stock-picking ability of the rulebase. The IR is calculated as: IR $= \frac{\sqrt{260}\alpha}{\sigma_e}$, where σ_e is the standard error of alpha in the capital asset pricing model expression of portfolio return given in the previous section and used to manage the alpha benchmark. Selectivity and Net Selectivity [15] provide further refinement of overall performance adjusted for risk. Any returns that a portfolio earns above the risk free rate are adjusted for both the returns that a market benchmark portfolio would earn if it had the same level of systematic risk and the same level of total risk.

Figure 3 shows differences in portfolio return and volatility due to complexity. Linguistic complexity is indicated by the number of membership functions

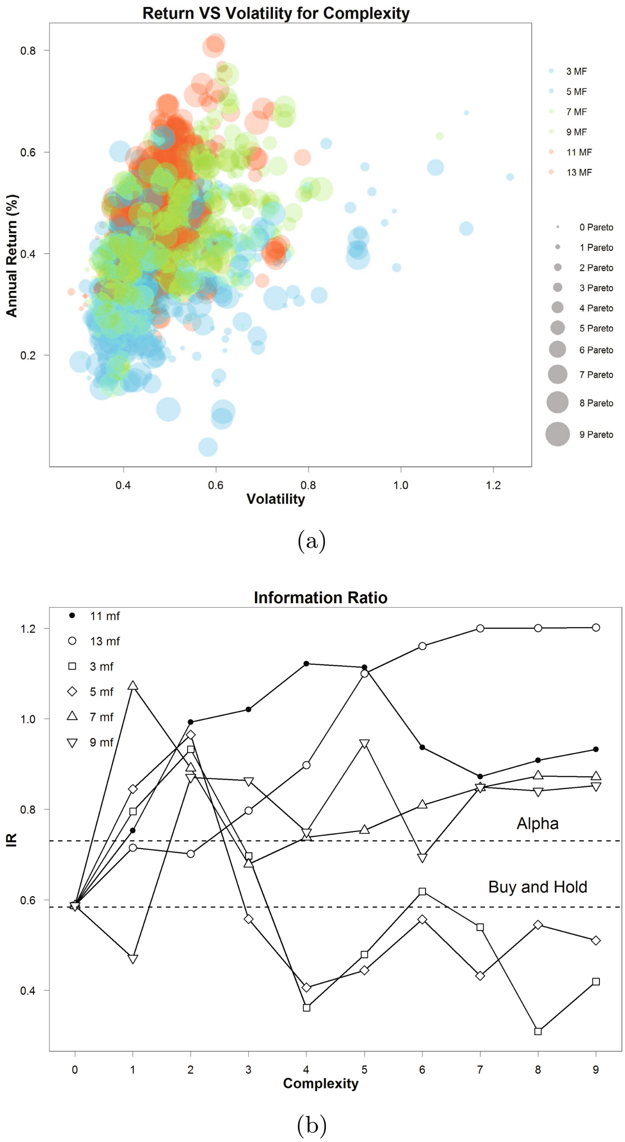

Fig. 2. The risk profile of portfolios managed with solutions of differing complexity. (a) shows Return vs. Volatility by complexity. (b) shows the Information Ratio of the portfolios and the benchmarks.

(i.e. 3 to 13 MF) and structural complexity by numbers (0 to 9 Pareto). Simpler models resulted in of lower return and (often) higher volatility, there was also lower stability between the different runs. The stability between different runs is an important risk as the method is stochastic. For more complex models higher return was observed at the cost of higher volatility (though here the information ratio shows the increase in volatility is justified by the return). In all cases there is a crux where increasing complexity above a certain level does not lead to gain and performance deteriorates. But only the average to higher linguistic complexity models provide the potential to out perform the benchmarks with appropriate levels of structural complexity.

4 Conclusion

In this paper we have described an evolutionary fuzzy system for portfolio management that first of all makes novel contributions by significantly building on earlier work in [2]. The experiments show the system performs well. There was considerable variation however for different levels of complexity and this can be used to hone performance. Almost all related research in finance is focused on relatively simple rules, we find such rules did not result in excess return above the benchmarks while the more complex models could. Therefore, it seems likely that controversy in some circles regarding the possibility of finding profitable rules somewhat misses the point, recent systems based on machine learning methods probably are able to do well by harnessing complexity. Another observation is that as the system approximates human reasoning, additional complexity may indicate to some extent why it is a fact that industry practitioners commonly attempt to generate profits by trading based on past data, despite there being in academia almost a consensus that there is limited possibility to do so.

When computational intelligence is used in algorithmic trading, it can lead to novel ways of controlling performance. For instance, as solution complexity is found to be a strong driver of risk and return, performance can be reliably shaped through identifying the locus where additional return starts to generate higher risk. It is also possible to use model complexity parameters improve stability and thus limit problems associated with the stochastic nature of the learning process which are often viewed as a drawback compared to other static modeling approaches in algorithmic trading.

References

1. Narang, K.: Inside the Black Box The Simple Truth About Algorithmic Trading. John Wiley & Sons, Inc., New York (2009)
2. Ghandar, A., Michalewicz, Z., Schmidt, M., To, T.-D., Zurbruegg, R.: Computational intelligence for evolving trading rules. IEEE Trans. Evolutionary Computation 13(1), 71–86 (2009)
3. Ghandar, A., Michalewicz, Z., Zurbruegg, R.: A case for learning simpler rule sets with multiobjective evolutionary algorithms. In: RuleML Europe, pp. 297–304 (2011)

4. Jin, Y., Sendhoff, B.: Pareto-Based Multiobjective Machine Learning: An Overview and Case Studies. IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews 38(3), 397–415 (2008)
5. Lee, C.-C.: A self-learning rule-based controller employing approximate reasoning and neural net concepts. International Journal of Intelligent Systems 6(1), 71–93 (1991)
6. Ray, Ball: Anomalies in relationships between securities' yields and yield-surrogates. Journal of Financial Economics 6(2-3), 103–126 (1978)
7. Basu, S.: Investment performance of common stocks in relation to their price-earnings ratios: A test of the efficient market hypothesis. The Journal of Finance 32(3), 663–682 (1977)
8. Beaver, W., Lambert, R., Morse, D.: The information content of security prices. Journal of Accounting and Economics 2(1), 3–28 (1980)
9. Bhandari, L.C.: Debt/equity ratio and expected common stock returns: Empirical evidence. The Journal of Finance 43(2), 507–528 (1988)
10. Barberis, N., Shleifer, A., Vishny, R.: A model of investor sentiment. Journal of Financial Economics 49(3), 307–343 (1998)
11. Kavajecz, K., Odders-White, E.: Technical Analysis and Liquidity Provision. Rev. Financ. Stud. 17(4), 1043–1071 (2004)
12. Wilder, J.: New Concepts in Technical Trading Systems. Trend Research (1978)
13. Cordón, O., Gomide, F.A.C., Herrera, F., Hoffmann, F., Magdalena, L.: Genetic fuzzy systems. new developments. Fuzzy Sets and Systems 141(1), 1–3 (2004)
14. Zitzler, E., Laumanns, M., Thiele, L.: Spea2: Improving the strength pareto evolutionary algorithm. Tech. Rep. (2001)
15. Fama, E.F.: Components of investment performance. Journal of Finance 27(3), 551–567 (1972)