

A Multi-Objective Evolutionary Algorithm for Mining Quantitative Association Rules

Diana Martín and Alejandro Rosete

*Dept. Artificial Intelligence and Infrastructure of Informatic Systems
Higher Polytechnic Institute José Antonio Echeverría, Cujae
19390 La Habana, Cuba
Email: {dmartin, rosete}@ceis.cujae.edu.cu*

Jesús Alcalá-Fdez and Francisco Herrera

*Dept. Computer Science and Artificial Intelligence
CITIC-UGR, University of Granada
18071 Granada, Spain
Email: {jalcala, herrera}@decsai.ugr.es*

Abstract—Data mining is most commonly used in attempts to induce association rules from database. Recently, some researchers have suggested the extraction of association rules as a multi-objective problem, removing some of the limitations of current approaches. In this way, we can jointly optimize quality measures which can present different degrees of trade-off depending on the database used and the type of information can be extracted from it.

In this work, we extend the well-known multi-objective evolutionary algorithms NSGA-II to perform an evolutionary learning of the intervals of attributes and a condition selection in order to mine a set of quantitative association rules with a good trade-off between interpretability and accuracy. To do that, this method considers three objectives, maximize the interestingness, comprehensibility and performance. Moreover, this method follows a database-independent approach which does not rely upon minimum support and minimum confidence thresholds. The results obtained over two real-world databases demonstrate the effectiveness of the proposed approach.

Keywords—Data Mining; Quantitative Association Rules; Multi-Objective Evolutionary Algorithms; NSGA-II.

I. INTRODUCTION

Discovering association rules is one of several Data Mining (DM) techniques described in the literature [15]. Association rules are used to represent and identify dependencies between items in a database [25]. These are an expression of the type $X \rightarrow Y$, where X and Y are sets of items and $X \cap Y = \emptyset$. It means that if all the items in X exist in a transaction then all the items in Y are also in the transaction with a high probability, and X and Y should not have a common item [2].

Many previous studies for mining association rules focused on databases with binary or discrete values, however the data in real-world applications usually consists of quantitative values. Designing DM algorithms, able to deal with various types of data, presents a challenge to workers in this research field.

In the last years, many researchers have proposed Evolutionary Algorithms (EAs) [9] for mining quantitative association rules [19], [24] from databases with quantitative values. EAs, particularly Genetic Algorithms (GAs) [14], are considered to be one of the most successful search techniques

for complex problems and it has proved to be an important technique for learning and knowledge extraction. The main motivation for applying GAs to knowledge extraction tasks is that they are robust and adaptive search methods that perform a global search in place of candidate solutions (for instance, rules or other forms of knowledge representation).

Recently, some researchers have suggested the extraction of association rules as a multi-objective problem (instead of a single objective), removing some of the limitations of current approaches. Several objectives are considered in the process of extracting association rules, obtaining a set with more interesting rules and accurate [1], [13]. In this way, we can jointly optimize measures such as support, confidence, and so on, which can present different degrees of trade-off depending on the database used and the type of information can be extracted from it. Since this approach presents a multiobjective nature the use of Multi-Objective Evolutionary Algorithms (MOEAs) [4], [7] to obtain a set of solutions with different degrees of trade-off between the different measures could represent an interesting way to work (by considering these measures as objectives).

In this work, we propose an extension of the well-known MOEA Non-dominated Sorting Genetic Algorithm II (NSGA-II) [8] to mine a set of quantitative association rules (NSGA-II-QAR) with a good trade-off between interpretability and accuracy. To do that, this method performs an evolutionary learning of the intervals of the attributes and a condition selection for each rule considering three objectives, maximize the interestingness, comprehensibility and performance, understanding for performance the product between confidence and support. Moreover, this method follows a database-independent approach which does not rely upon the minimum support and the minimum confidence thresholds which are hard to determine for each database.

This paper is arranged as follows. Next section presents a brief study of the existing MOEAs for general purpose [29]. In Section III we present our proposal to learn the intervals of the attributes and to perform a condition selection in order to obtain a set of high quality association rules. Section IV shows the results of the proposed mining algorithm applied

over two real-world databases. Finally, Section V points out some concluding remarks.

II. MULTI-OBJECTIVE EVOLUTIONARY ALGORITHMS

EAs simultaneously deal with a set of possible solutions (the so-called population) which allows to find several members of the Pareto optimal set in a single run of the algorithm. Additionally, they are not too susceptible to the shape or continuity of the Pareto front (e.g., they can easily deal with discontinuous and concave Pareto fronts).

The first hint regarding the possibility of using EAs to solve a multi-objective problem appears in a Ph.D. thesis from 1967 [21] in which, however, no actual MOEA was developed (the multi-objective problem was restated as a single-objective problem and solved with a genetic algorithm). David Schaffer is normally considered to be the first to have designed a MOEA during the mid-1980s [22]. Schaffer's approach, called Vector Evaluated Genetic Algorithm (VEGA) consists of a simple genetic algorithm with a modified selection mechanism. However, VEGA had a number of problems from which the main one had to do with its inability to retain solutions with acceptable performance, perhaps above average, but not outstanding for any of the objective functions.

After VEGA, the researchers designed a first generation of MOEAs characterized by its simplicity where the main lesson learned was that successful MOEAs had to combine a good mechanism to select non-dominated individuals (perhaps, but not necessarily, based on the concept of Pareto optimality) combined with a good mechanism to maintain diversity (fitness sharing was a choice, but not the only one). The most representative MOEAs of this generation are the following: Nondominated Sorting Genetic Algorithm (NSGA) [23], Niched-Pareto Genetic Algorithm (NPGA) [16] and Multi-Objective Genetic Algorithm (MOGA) [12].

A second generation of MOEAs started when elitism became a standard mechanism. In fact, the use of elitism is a theoretical requirement in order to guarantee convergence of a MOEA. Many MOEAs have been proposed during the second generation (which we are still living today). However, most researchers will agree that few of these approaches have been adopted as a reference or have been used by others. In this way, the Strength Pareto Evolutionary Algorithm 2 (SPEA2) [27] and the NSGA-II [8] can be considered as the most representative MOEAs of the second generation, also being of interest some others as the Pareto Archived Evolution Strategy (PAES) [17] and the Multiobjective Evolutionary Algorithm Based on Decomposition (MOEA/D). Table I shows a resume of the most representative MOEAs of both generations.

Finally, we have to point out that nowadays NSGA-II is the paradigm within the MOEA research community since the powerful crowding operator that this algorithm uses

Table I
CLASSIFICATION OF MOEAs

Reference	MOEA	1 st Gen.	2 nd Gen.
[12]	MOGA	✓	
[16]	NPGA	✓	
[23]	NSGA	✓	
[3]	micro-GA		✓
[28], [18]	MOEA/D & MOEA/D-DE		✓
[10]	NPGA 2		✓
[8]	NSGA-II		✓
[17]	PAES		✓
[5], [6]	PESA & PESA-II		✓
[26], [27]	SPEA & SPEA2		✓

usually allows to obtain the widest Pareto sets in a great variety of problems, which is a very appreciated property in this framework.

III. A MOEA FOR MINING QUANTITATIVE ASSOCIATION RULES

The proposed algorithm extends the well-known MOEA NSGA-II [8] in order to mine a set of quantitative association with a good trade-off between interpretability and accuracy. To do that, we consider as objectives the comprehensibility, interestingness and performance of the rules.

In the following, the main characteristics of this approach are presented: coding scheme, initial gene pool, objectives, genetic operators and genetic model.

A. Coding scheme and initial gene pool

Each chromosome is a vector of genes that represent the attributes and intervals of the rule. We have used a positional encoding, where the i -th attribute is encoded in the i -th gene has been used. To combine the condition selection with the learning of the intervals, each gene consist of three parts:

- The first part (ac) represents when a gene is involved or not in the rule. When this part is '-1', this attribute is not involved in the rule, and when this part is '0' or '1' this attribute is part of the antecedent or consequent of the rule, respectively. All genes that have '0' on their first parts will form the antecedent of the rule while genes that have '1' will form the consequent of the rule.
- The second part represents the lower bound (lb) of the interval of the attribute.
- The third part represents the upper bound (ub) of the interval of the attribute.

Notice that, lb and ub will be equal in the intervals of nominal attributes. Finally, a chromosome C_T is coded in the following way, where m is the number of attributes in the database.

$$Gene_i = (ac_i, lb_i, ub_i), \quad i = 1, \dots, m, \\ C_T = Gene_1 Gene_2 \dots Gene_m$$

In order to avoid the intervals to grow up until spanning the total domain, we define *amplitude* as the maximum size the interval of a determined attribute can get. Thus, the *amplitude* of a attribute i is defined as:

$$Amplitude_i = (Max_i - Min_i) / \delta$$

where Max_i and Min_i are the maximum and minimum values of the domain of attribute i respectively, and δ is a value given by the system expert that determines the tradeoff between generalization and specificity of the rules.

The initial population will be consisted of a rule set (with only one attribute in the consequent) with a good coverage of the database. To do that, first we select at random the attributes that will be part of the antecedent and consequent of the rule. Then we select at random an example from database and generate the interval of each attribute with a size equal to 50% of the *amplitude* of each attribute and with the values of the example selected in the center of each of them. Finally, the examples covered for this rule are removed of the database. This process is repeated until initial population is completed. Notice that, if all examples are removed of the database all of them will be added again to the database.

B. Objectives

Three objectives are maximized for this problem: Interestingness, Comprehensibility and Performance. Performance is the result of the product between confidence and support which allow us to obtain accurate rules and a good trade-off between local and general rules. These measures (support and confidence) for a rule $X \rightarrow Y$ are defined as:

$$Support(X \rightarrow Y) = SUP(XY) / |D|$$

$$Confidence(X \rightarrow Y) = SUP(XY) / SUP(X)$$

where $SUP(XY)$ is the number of examples of the database covered by the antecedent and consequent of the rule, and $SUP(X)$ is the number of examples of the database covered by the antecedent of the rule.

Interestingness measures how much interesting the rule is which allow us to extract only those rules that may be more interesting to the users. In this case, we have used the interestingness measure lift [20] which represents the ratio between the confidence of the rule and the expected confidence of the rule. This is defined as:

$$Lift(X \rightarrow Y) = \frac{SUP(XY) / |D|}{SUP(X) / |D| * SUP(Y) / |D|}$$

where $SUP(Y)$ is the number of examples of the database covered by the consequent of the rule.

Finally, comprehensibility tries to quantify the understandability of the rule [11]. The generated rules may have a large number of attributes involved thereby making it difficult to understand. If the generated rules are not




Figure 1. A simple example of the crossover operator

understandable to the user, the user will never use them. Here, the comprehensibility of a rule $X \rightarrow Y$ is measured by the number of attributes involved in the rule and is defined as:

$$Comprehensibility(X \rightarrow Y) = 1 / Attr_{X \rightarrow Y}$$

where $Attr_{X \rightarrow Y}$ is the number of attributes involved in the rule.

C. Genetic Operators

The crossover operator generates two offspring interchanging randomly the genes of the parents (exploration). Figure 1 shows a simple example of the performance of this operator.

The mutation operator consists in modifying randomly the interval (lb and ub) and ac of a gene selected at random. This operator selects at random one of the bounds of the interval and increases or decreases its value randomly. We have to be specially careful in not overcoming the fixed value of amplitude. The value for ac is randomly selected within the set $\{-1, 0, 1\}$.

D. Repairing operator

After mutation operator, if any rule doesn't have antecedent or consequent or has more than one attribute in the consequent, a repairing operator is performed to modify these rules. If there are more than one attribute in the consequent, one attribute is randomly selected as consequent between them and the remaining of attributes are passed to the antecedent. If there is not any attribute in the antecedent and/or consequent these are randomly selected between the attributes not involved.

Finally, the size of the intervals are decreased until the number of examples covered is smaller than the number of examples covered by the original intervals in order to obtain more simple rules.

E. NSGA-II Genetic Model

As in other EAs, first NSGA-II generates an initial population. Then an offspring population is generated from the current population by selection, crossover and mutation. The next population is constructed from the current and offspring populations. The generation of an offspring population and the construction of the next population are iterated until a stopping condition is satisfied. The NSGA-II algorithm has two features, which make it a high-performance MOEA. One is the fitness evaluation of each solution based on Pareto

ranking and a crowding measure, and the other is an elitist generation update procedure.

Each solution in the current population is evaluated in the following manner. First, Rank 1 is assigned to all non-dominated solutions in the current population. All solutions with Rank 1 are tentatively removed from the current population. Next, Rank 2 is assigned to all non-dominated solutions in the reduced current population. All solutions with Rank 2 are tentatively removed from the reduced current population. This procedure is iterated until all solutions are tentatively removed from the current population (i.e., until ranks are assigned to all solutions). As a result, a different rank is assigned to each solution. Solutions with smaller ranks are viewed as being better than those with larger ranks. Among solutions with the same rank, an additional criterion called a crowding measure is taken into account.

The crowding measure for a solution calculates the distance between its adjacent solutions with the same rank in the objective space. Less crowded solutions with larger values of the crowding measure are viewed as being better than more crowded solutions with smaller values of the crowding measure.

A pair of parent solutions are selected from the current population by binary tournament selection based on the Pareto ranking and the crowding measure. When the next population is to be constructed, the current and offspring populations are combined into a merged population. Each solution in the merged population is evaluated in the same manner as in the selection phase of parent solutions using the Pareto ranking and the crowding measure. The next population is constructed by choosing a specified number (i.e., population size) of the best solutions from the merged population. Elitism is implemented in NSGA-II algorithm in this manner.

Considering the components previously defined and the descriptions of the authors in [8], NSGA-II consists of the next steps:

- 1) A combined population R_t is formed with the initial parent population P_t and offspring population Q_t (initially empty).
- 2) Generate all non-dominated fronts $F = (F_1, F_2, \dots)$ of R_t .
- 3) Initialize $P_{t+1} = 0$ and $i = 1$.
- 4) Repeat until the parent population is filled.
- 5) Calculate crowding-distance in F_i .
- 6) Include i -th non-dominated front in the parent population.
- 7) Check the next front for inclusion.
- 8) Sort in descending order using crowded-comparison operator.
- 9) Choose the first $(N - |P_{t+1}|)$ elements of F_i .
- 10) Use selection, crossover, mutation and repairing operator to create a new population Q_{t+1} .
- 11) Increment the generation counter.

Table II
PARAMETERS CONSIDERED FOR COMPARISON

Method	Parameters
GENAR	$PopSize = 100, N_{Eval} = 50,000, P_{sel} = 0.25, P_{cro} = 0.7, P_{mut} = 0.1, nRules = 30, FP = 0.7, AF = 2$
MODENAR	$PopSize = 100, N_{Eval} = 50,000, Threshold = 60, CR = 0.3, W_{sup} = 0.8, W_{conf} = 0.2, W_{comp} = 0.1, W_{ampl} = 0.4$
NSGA-II-QAR	$PopSize = 100, N_{Eval} = 50,000, P_{mut} = 0.1, \delta = 2$

IV. EXPERIMENTAL RESULTS

In order to analyze the performance of the proposed mining algorithm, we have considered two real-world databases:

- *Stulong*: It is a database concerning a study of the risk factors of atherosclerosis in a population of 1419 middle-aged men in the years 1976 - 1999¹. Here, we extract five quantitative attributes out of a total of 64 attributes. The selected attributes are height, weight, systolic blood pressure, diastolic blood pressure and cholesterol level.
- *House_16H*: It concerns a study to predict the median price of the houses in a region by considering both the demographic composition and the state of housing market. This data was collected as part of the 1990 US census. For the purpose of this database, only a level State-Place was used and data from all states was obtained. This database contains 22,784 examples and 17 quantitative attributes².

We compare the proposed algorithm with a mono-objective algorithm and a MOEA for mining quantitative association rules, GENAR [19] and MODENAR [1], respectively. The parameters of the analyzed methods are shown in Table II. With these values for our proposal we have tried to facilitate comparisons, selecting standard common parameters that work well in most cases instead of searching for very specific values. The parameters of the remaining methods were selected according to the recommendation of the corresponding authors within each proposal. Furthermore, for all the experiments conducted in this study, the results shown in the tables always refer to association rules having a minimum confidence greater than or equal to 0.8.

The results obtained for the database *Stulong* by the analyzed methods are shown in Table III, where $\#R$ is

¹The study (STULONG) was performed at the 2nd Department of Medicine, 1st Faculty of Medicine of Charles University and Charles University Hospital, under the supervision of Prof. F. Boudk with collaboration of M. Tomeckov and Ass. Prof. J. Bultas. The data were transferred to electronic form by the European Centre of Medical Informatics, Statistics and Epidemiology of Charles University and Academy of Sciences. The data resource is on the web page <http://euromise.vse.cz/challenge2004>. At present, the data analysis is supported by the grant of the Ministry of Education CR Nr LN 00B 107

²This database was designed on the basis of data provided by US Census Bureau [<http://www.census.gov>] (under Lookup Access [<http://www.census.gov/cdrom/lookup>]: Summary Tape File 1).

Table III
RESULTS FOR THE DATABASE STULONG

Algorithm	#R	Av_{Sup}	Av_{Conf}	Av_{Lift}	Av_{Amp}	%Tran
GENAR	30	0.88	0.98	1.00	5.0	95.27
MODENAR	85	0.56	0.97	1.39	2.6	99.57
NSGA-II-QAR	90	0.52	0.93	43.74	3.1	100.00

Table IV
RESULTS FOR THE DATABASE HOUSE_16H

Algorithm	#R	Av_{Sup}	Av_{Conf}	Av_{Lift}	Av_{Amp}	%Tran
GENAR	30	0.43	0.98	1.01	17.0	87.29
MODENAR	98	0.62	0.97	1.00	6.0	96.64
NSGA-II-QAR	96	0.46	0.96	443.91	3.6	99.99

the number of the generated association rules, Av_{Sup} and Av_{Conf} are, respectively, the average support and the average confidence of the rules, Av_{Lift} is the average value for the measure lift of the rules, Av_{Amp} is the average length of the rules in terms of attributes involved, and %Tran is the percentage of transactions covered by the rules on the total examples in the database. Analysing the results presented in this table, we can present the following conclusions:

- The MOEAs returned sets of rules with less number of attributes and better coverage of the database than the mono-objective algorithm, giving the advantage of easier understanding from a user's perspective.
- The method proposed allow us to obtain a set of association rules involving few attributes and with the best average lift and coverage of the database, providing the user with high quality rules.

The results obtained for the database House_H16 by the analyzed methods are shown in Table IV. Analysing the results presented in Table IV, we can stress the following facts:

- In this database with more attributes, the rule sets obtained by MOEAs involve again few attributes in the rules and present a good coverage of the database.
- Moreover, the method proposed mine again the rules set with best average lift and coverage of the database.

Fig. 2 shows the relationship between the confidence and support of the rules obtained by our proposal and the numbers of evaluations for the database Stulong. It can be easily seen from this figure that the average confidence of the rules increases with the increase of the number of evaluations and the distribution of the rules (different supports) is maintained in the population. It means that the proposed method allows us to obtain an association rule set with a high confidence and a good trade-off between specific and general rules.

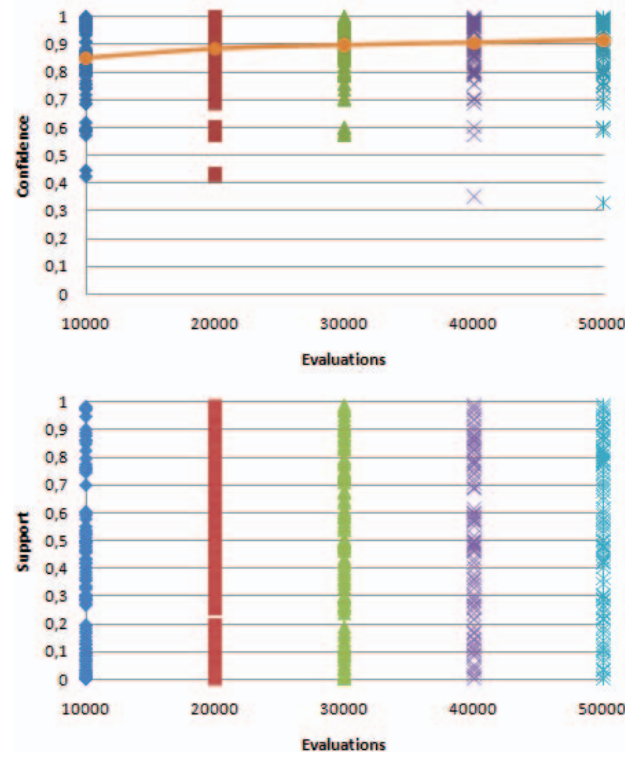


Figure 2. Relationship between confidence/support of the rules in the population and number of evaluations in the database Stulong. The line represents the average confidence of the rules in the population.

V. CONCLUSION

In this work, we have proposed an extension of the well-known MOEA NSGA-II to mine a set of quantitative association rules with a good trade-off between interpretability and accuracy. To do that, this method performs an evolutionary learning of the intervals of attributes and a condition selection for each rule considering three objectives, maximize the interestingness, comprehensibility and performance, understanding for performance the product between confidence and support. Moreover, this method follows a database-independent approach which does not rely upon the minimum support and the minimum confidence thresholds which are hard to determine for each database.

The results obtained over two real-world databases have shown how the method proposed let us to mine rule sets with a good trade-off between the different objectives, obtaining association rules with few attributes and with the best average lift and coverage of the dataset, providing the user with high quality rules.

ACKNOWLEDGMENT

This paper has been supported by the Spanish Ministry of Education and Science under grants TIN2008-06681-C06-01.

REFERENCES

- [1] B. Alatas, E. Akin, and A. Karci, *MODENAR: Multi-objective differential evolution algorithm for mining numeric association rules*. Applied Soft Computing 8, pp. 646-656, 2008.
- [2] R. Agrawal, T. Imielinski and A. Swami. *Mining association rules between sets of items in large databases*. SIGMOD, Washington D.C., May 1993, pp. 207-216.
- [3] C. Coello, and G. Toscano, *A Micro-Genetic Algorithm for multiobjective optimization*. First Int. Conf. on Evolutionary Multi-Criterion Optimization, LNCS 1993 (London, UK, 2001) pp. 126-140.
- [4] C. Coello, G. Lamont, and D. Van Veldhuizen, *Evolutionary Algorithms for solving multi-objective problems*, Kluwer Academic Publishers, 2002.
- [5] D. Corne, J. Knowles, and M. Oates, *The Pareto Envelope-based Selection Algorithm for multiobjective optimization*. Parallel Problem Solving from Nature VI Conf., LNCS 1917 (Paris, France, 2000) pp. 839-848.
- [6] D. Corne, N. Jerram, J. Knowles, and M. Oates, *PESA-II: Region based selection in evolutionary multiobjective optimization*. Genetic and Evolutionary Computation Conf. (San Francisco, CA, 2001) pp. 283-290.
- [7] K. Deb, *Multi-objective optimization using evolutionary algorithms*, Kluwer Academic, EE.UU., 2001.
- [8] K. Deb, S. Agrawal, A. Pratab, and T. Meyarivan, *A fast and elitist multiobjective genetic algorithm: NSGA-II*. IEEE Trans. Evolutionary Computation 6:2, pp. 182-197, 2002.
- [9] A.E. Eiben, and J.E. Smith, *Introduction to Evolutionary Computing*, SpringerVerlag, 2003.
- [10] M. Erickson, A. Mayer, and J. Horn, *The Niche Pareto Genetic Algorithm 2 applied to the design of groundwater remediation systems*. First Int. Conf. on Evolutionary Multi-Criterion Optimization, LNCS 1993 (London, UK, 2001) pp. 681-695.
- [11] M.V. Fidelis, H.S. Lopes, and A.A. Freitas, *Discovering comprehensible classification rules with a genetic algorithm*. Proceedings of the 2000 Congress on Evolutionary Computation, (CA, USA, 2000) pp. 805-810.
- [12] C.M. Fonseca, and P.J. Fleming, *Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization*. 5th Int. Conf. on Genetic Algorithms (San Mateo, CA, 1993) pp. 416-423.
- [13] A. Ghosh, and B. Nath, *Multi-objective rule mining using genetic algorithms*. Information Sciences 163, pp. 123-133, 2004.
- [14] D.E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc., 1989.
- [15] J. Han and M. Kamber, *Data Mining: Concepts and Techniques*. Second Edition, Morgan Kaufmann, 2006.
- [16] J. Horn, N. Nafpliotis, and D.E. Goldberg, *A niched pareto genetic algorithm for multiobjective optimization*. First IEEE Conf. on Evolutionary Computation, IEEE World Congress on Computational Intelligence 1 (Piscataway, NJ, 1994) pp. 82-87.
- [17] J.D. Knowles, and D.W. Corne, *Approximating the non dominated front using the Pareto archived evolution strategy*. Evolutionary Computation 8:2, pp. 149-172, 2000.
- [18] H. Li, and Q. Zhang, *Multiobjective Optimization Problems With Complicated Pareto Sets, MOEA/D and NSGA-II*. IEEE Transactions on Evolutionary Computation 13:2, pp. 284-302, 2009.
- [19] J. Mata, J. Alvarez, and J. Riquelme, *Mining Numeric Association Rules with Genetic Algorithms*. 5th International Conference on Artificial Neural Networks and Genetic Algorithms (Prague, 2001) pp. 264-267.
- [20] S. Ramaswamy, S. Mahajan, and A. Silberschatz, *On the Discovery of Interesting Patterns in Association Rules*. 24rd International Conference on Very Large Data Bases (San Francisco, CA, USA, 1998), pp. 368-379.
- [21] R.S. Rosenberg, *Simulation of genetic populations with biochemical properties*. MA Thesis, Univ. Michigan, Ann Harbor, Michigan, 1967.
- [22] J.D. Schaffer, *Multiple objective optimization with vector evaluated genetic algorithms*. First Int. Conf. on Genetic Algorithms (Pittsburgh, USA, 1985) pp. 93-100.
- [23] N. Srinivas, and K. Deb, *Multiobjective optimization using nondominated sorting in genetic algorithms*. Evolutionary Computation 2, pp. 221-248, 1994.
- [24] X. Yan, C. Zhang, and S. Zhang, *Genetic algorithm-based strategy for identifying association rules without specifying actual minimum support*. Expert Systems with Applications 36:2, pp. 3066-3076, 2009.
- [25] Ch. Zhang and S. Zhang, *Association Rule Mining: Models and Algorithms*, Lecture Notes in Computer Science, LNAI 2307, 2002.
- [26] E. Zitzler, and L. Thiele, *Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach*. IEEE Trans. Evolutionary Computation 3:4, pp. 257-271, 1999.
- [27] E. Zitzler, M. Laumanns, and L. Thiele, *SPEA2: Improving the strength Pareto evolutionary algorithm for multiobjective optimization*. Evolutionary Methods for Design, Optimization and Control with App. to Industrial Problems (Barcelona, Spain, 2001) pp. 95-100.
- [28] Q. Zhang, and H. Li, *MOEA/D: A Multiobjective Evolutionary Algorithm Based on Decomposition*. IEEE Transactions on Evolutionary Computation 11:6, pp. 712-731, 2007.
- [29] A. Zhou, B.-Y. Qu, H. Li, S.-Z. Zhao, P. Nagarathnam, Q. Zhang, *Multiobjective evolutionary algorithms: A survey of the state of the art*. Swarm and Evolutionary Computation 1, pp. 32-49, 2011.