

A Multi-Objective Evolutionary Algorithm Approach to Learn Disjunctive and Conjunctive Topical Queries

Rocío L. Cecchini[†] Carlos M. Lorenzetti[‡] Ana G. Maguitman[‡]

Departamento de Ciencias e Ingeniería de la Computación
Universidad Nacional del Sur, Av. Alem 1253,
(8000) Bahía Blanca, Argentina
phone: 54-291-4595135 fax: 54-291-4595136

[†] LIDeCC - Laboratorio de Investigación y Desarrollo en Computación Científica
[‡] LIDIA - Laboratorio de Investigación y Desarrollo en Inteligencia Artificial
e-mail: {rlc,cml,agm}@cs.uns.edu.ar

Abstract. Topical search refers to the process of formulating queries that reflect a thematic context. A combination of machine learning and information retrieval techniques can be applied to automate this process. In this work we propose the application of single- and multi-objective evolutionary algorithms to automatically evolve a population of topical queries. We report on the results of different strategies that attempt to evolve conjunctive and disjunctive queries. In our evaluations we observe that disjunctive queries have the potential to achieve better retrieval performance than conjunctive queries. In addition, our analysis reveals the limitations of the single-objective approach and highlights the advantages of applying multi-objective evolutionary algorithms for the problem at hand.

Keywords: topical search, conjunctive queries, disjunctive queries, multi-objective evolutionary algorithms.

1 Introduction

A major challenge for human information seekers is how to formulate queries that effectively reflect their information needs. Automatic topical search refers to automatically formulating queries with terms extracted from a thematic context. The resources collected by the formulation of topical queries can be used in different scenarios, such as responding to contextualized information needs [17, 3], fulfilling long term information needs [27], collecting resources for topical Web portals [6], or accessing the Deep Web [15], among others.

The effectiveness of a topical query depends on the task at hand. If the criteria for evaluating query performance can be quantitatively specified then the problem of topical search can be seen as an optimization problem where the objective function to be maximized quantifies the optimality of a query. In this optimization problem, therefore, the search space is defined as the set of possible queries that can be presented to a search interface. A particularity of this optimization problem is that the query space is a high-dimensional space, where each possible term accounts for a new dimension.

Usually, high-dimensional space problems are computationally very intensive and cannot be effectively solved using analytical methods. On the other hand, a query can be considered effective even if it is not an optimal one, at the same time as multiple queries can provide satisfactory results. Therefore, we may be interested in finding many near optimal queries rather than a single optimal one. Another aspect of this optimization problem is that several objective such as high precision and high recall can be used as criteria for evaluating query performance.

On the basis of the above discussion, Evolutionary Algorithms (EAs) [14, 11] are applicable to the problem of learning to automatically formulate high-quality topical queries. EAs are general-purpose search procedures based on the mechanisms of natural selection. An important component in EAs is the fitness function, which in combination with the selection mechanism determines which elements of the population are selected to be members of the next generation. Therefore, it is necessary to establish some criteria to determine if one solution is better than another. In the multi-objective case, there is not only one criterion to conclude whether one solution is better than another. The strategy adopted in this work applies the concept of Pareto optimality [22] as well as an aggregative technique based on the harmonic mean of the given objectives to rank the queries in a manner such that the most promising ones have a higher probability of being selected. The proposed framework starts by generating an initial population of queries using terms extracted from a topic description and incrementally evolves those queries based on their ability to retrieve results satisfying a number of objectives. This approach allowed us to address some interesting research question:

- **Potential for evolution of queries:** Is it possible to evolve queries in such a way that they significantly outperform those generated directly from the initial topic description?
- **Generalization of evolved queries:** Are the queries evolved from the training data useful on a new corpus?
- **Conjunctive vs. disjunctive queries:** For the objectives analyzed here, how good is the performance of conjunctive queries when compared to disjunctive queries?

2 An Overview of Evolutionary Algorithms

EAs [14, 11] are robust optimization techniques based on the principle of natural selection and survival of the fittest, which claims “in each generation the stronger individual survives and the weaker dies”. Therefore, each new generation would contain stronger (fitter) individuals in contrast to its ancestors.

To use EAs in optimization problems we need to define candidate solutions by chromosomes consisting of genes and a fitness function to be maximized. A population of candidate solutions (usually of a constant size) is maintained. The goal is to obtain better solutions after some generations. To produce a new generation EAs typically use selection together with the genetic operators of crossover and mutation. Parents are selected to produce offspring, favoring those parents with highest values of the fitness function. Crossover of population members takes place by exchanging subparts of the parent chromosomes (roughly mimicking a mating process), while mutation is the result of a

random perturbation of the chromosome (e.g., replacing a gene by another). Although selection, crossover and mutation can be implemented in many different ways, their fundamental purpose is to explore the search space of candidate solutions, improving the population at each generation by adding better offspring and removing inferior ones.

In Multi-Objective Optimization Problems (MOOPs) the quality of a solution is defined by its performance in relation to several, possibly conflicting, objectives. Traditional methods are very limited because, in general, they become too computationally intensive as the size of the problem grows [24, 18]. EAs are a suitable technique for dealing with MOOPs [8, 9, 11] and are called in this case Multi-Objective Evolutionary Algorithms (MOEAs). There are many approaches to multi-objective optimization using MOEAs, and in general, they can be classified in Pareto or non-Pareto EAs. In the first case, the evaluation is made following the Pareto dominance concept [22] discussed below. In the second case, the objectives are combined to obtain a single evaluation value.

There are some basic definitions based on the Pareto concept that must be considered:

Definition 1. Pareto Dominance [7]: A vector $\mathbf{u} = (u_1, u_2, \dots, u_k)$ is said to **dominate** another vector $\mathbf{v} = (v_1, v_2, \dots, v_k)$ (denoted by $\mathbf{u} \preceq \mathbf{v}$) if and only if \mathbf{u} is partially less than \mathbf{v} , i.e., $\forall i \in \{1, \dots, k\}, u_i \leq v_i \wedge \exists i \in \{1, \dots, k\} : u_i < v_i$.

Definition 2. Pareto Optimality [7]: A solution $\mathbf{x} \in \Omega$ is said to be **Pareto Optimal** with respect to a Ω , if and only if $\nexists \mathbf{x}^* \in \Omega$ for which $\mathbf{v} = F(\mathbf{x}^*) = (f_1(\mathbf{x}^*), \dots, f_k(\mathbf{x}^*))$ dominates $\mathbf{u} = F(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_k(\mathbf{x}))$. Where Ω is a feasible region for the MOOP.

Definition 3. Pareto Optimal Set [7]: For a given MOOP, $F(\mathbf{x})$, and a feasible region for that MOOP, Ω , the **Pareto Optimal Set**, \mathcal{P}^* , is defined as:

$$\mathcal{P}^* := \{\mathbf{x} \in \Omega, \nexists \mathbf{x}^* \in \Omega, F(\mathbf{x}^*) \preceq F(\mathbf{x})\}$$

Definition 4. Pareto Front [7]: For a given MOOP, $F(\mathbf{x})$, in a feasible region for that MOOP, Ω , and a **Pareto Optimal Set**, \mathcal{P}^* , the **Pareto Front** \mathcal{PF}^* is defined as:

$$\mathcal{PF}^* := \{\mathbf{u} = F(\mathbf{x}) \mid \mathbf{x} \in \mathcal{P}^*\}$$

Besides the Pareto or non-Pareto strategy, the EAs can be classified in elitist and non-elitist EAs. The difference resides in that the first uses a mechanism to retain the non-dominated individuals. In the last years, a great number of elitist Pareto-based EAs were developed. Several of them have shown very good performance in problems with objective space of size less or equal than four [9].

The Non-dominated Sorting Genetic Algorithm – II (NSGA-II) is one of the most studied and efficient EAs [10], consequently it was used in this work. The algorithm begins creating a random parent population P_0 of size n . The population is sorted based on the non-domination concept. Each solution is assigned a fitness (or rank) equal to its non-dominated level (1 if it belongs to the first front, 2 for the second front, and so on). In this order, minimization of fitness is assumed. After ranking the solutions, a population of n offspring, Q_0 , is created using binary tournament selection, recombination and mutation. The elitism is reached by comparing the current population with previously found best non-dominated solutions. The i th generation follows the next steps:

1. A combined population $R_i = P_i \cup Q_i$ of size $2n$ is formed.
2. R_i is ordered according to non-domination. Since all previously and current population members are included in R_i , elitism is ensured. Solutions belonging to the best front, \mathcal{F}_1 , are the best solution in the combined population R_i .
3. If the size of \mathcal{F}_1 is smaller than n , all members of the set \mathcal{F}_1 are chosen for the new population P_{i+1} . The remaining members of the population P_{i+1} are chosen from subsequent non-dominated fronts in the order of their ranking until no more sets can be accommodated. If \mathcal{F}_j is the last front from which individuals can be accommodated in the population, but not all the members can enter in the population, then a decision needs to be made to choose a subset of individuals from \mathcal{F}_j . In order to decide which members of this front will win a place in the new population, the NSGA-II uses a selection criterion based on a crowded-comparison operator that favors solutions located in lesser crowded regions.

In addition to the NSGA-II, an elitist non linear aggregation alternative was used. This scheme was implemented using an adaptation of the well known F_1 measure. This measure is reviewed in section 3.2. The PISA platform [2] was used to implement the strategies analyzed in this work.

3 Evolving Topical Queries with Multi-objective Evolutionary Algorithms

In order to evolve topical queries we start with a population of queries composed of terms extracted from an initial description of the given topic and rate the effectiveness of each query according to the quality of the search results. The best queries have higher chances of being selected for subsequent generations and therefore as generations pass, queries associated with improved search results will predominate. Furthermore, the mating process continually combines these queries in new ways, generating ever more sophisticated solutions.

3.1 Population and Representation of Chromosomes

Each chromosome corresponds to a query, which is represented as a list of terms. For our analysis of conjunctive queries, terms are assumed to be connected by the AND operator while for disjunctive queries terms are connected by the OR operator. Each term corresponds to a gene that can be manipulated by the genetic operators. The population is initialized with a fixed number of queries randomly generated with terms from the thematic description. Novel terms can be included in the queries after mutation takes place. These novel terms are obtained from a *mutation pool*, which is an ever increasing set of terms that may or may not be part of the initial context.

3.2 Fitness Function

The fitness function defines the criterion for assessing the quality of a query. One of the objective functions considered in our analysis is precision at rank 10 (*Precision@10*),

which is the fraction of the top 10 retrieved documents which are known to be relevant. To define this fitness function we associate with the search space Q and topics T a function $Precision@10 : Q \times T \rightarrow [0, 1]$. This objective function can numerically evaluate an individual query q in terms of precision at rank 10 for a given topic t as follows:

$$Precision@10(q, t) = \frac{|A_{q10} \cap R_t|}{|A_{q10}|},$$

where A_{q10} is the set of top-10 ranked documents returned by a search engine when q is used as a query, and R_t is the set containing all the documents associated with topic t , including those in its subtopics.

Another fitness function adopted in this work, $Recall : Q \times T \rightarrow [0, 1]$, is defined as the fraction of relevant documents R_t that are in the answer set A_q :

$$Recall(q, t) = \frac{|A_q \cap R_t|}{|R_t|}.$$

Finally we use a function $F^* : Q \times T \rightarrow [0, 1]$ that aggregates $Precision@10$ and $Recall$ as follows:

$$F^*(q, t) = \frac{2 \cdot Precision@10(q, t) \cdot Recall(q, t)}{Precision@10(q, t) + Recall(q, t)}.$$

The F^* is an adaptation of the F_1 measure, which is the weighted harmonic mean of precision and recall [26]. Maximization of fitness is assumed to prefer a query over another for the selection process. We have used a vector representation of the query together with the TFIDF weighting function [1] for assigning scores to the retrieved documents.

3.3 Genetic Operators

A new generation in our EAs is determined by a set of operators that select, recombine and mutate queries of the current population.

- **Selection:** A new population is generated by probabilistically selecting the highest-quality queries from the current set of queries. In the case when the query effectiveness can be codified as a scalar value (single-objective or aggregative methods) then two queries are chosen at random from the population and the one with highest effectiveness is selected for recombination and to populate the next generations. This method is known as 2-way tournament selection. In addition, elitism is applied to prevent losing the best queries. For the multi-objective case, selection is based on the elitist Pareto strategy described in section 2.
- **Crossover:** Some of the selected queries are carried out into the next generations as they are, while others are recombined to create new queries. The recombination of a pair of parent queries into a pair of offspring queries is carried out by copying selected terms from each parent into the descendants. The crossover operator used in our proposal is known as single-point. It results in new queries in which the first n terms are contributed by one parent and the remaining terms by the second parent, where the crossover point n is chosen at random.

- **Mutation:** Small random changes can be produced to the new population of queries. These changes consist in replacing a randomly selected query term t^q by another term t^p . The term t^p is obtained from the *mutation pool*, which is a set of terms that initially contains terms extracted from the thematic context and is incrementally updated with new terms from the relevant documents recovered by the system.

4 Evaluation

To run our evaluations we collected 448 topics from the Open Directory Project (ODP)¹. The topics were selected from the third level of the ODP hierarchy. A number of constraints were imposed on this selection with the purpose of ensuring the quality of our corpus. The minimum size for each selected topic was 100 URLs and the language was restricted to English. For each topic we collected all of its URLs as well as those in its subtopics. The total number of collected pages was more than 350K. The Terrier framework [21] was used to index these pages and to run our experiments. We used the stopword list provided by Terrier and Porter stemming was performed on all terms. We divided each topic in such a way that 2/3 of its pages were used for training and 1/3 for testing. The following scenarios were analyzed:

- A **Single-objective EA** with *Precision@10* as the objective function.
- A **Single-objective EA** with *Recall* as the objective function.
- **NSGA-II** with both *Precision@10* and *Recall* as objective functions.
- An **aggregative MOEA** with F^* as the objective function.

The EAs were run for 10 different topics. For each analyzed topic a population of 250 queries was randomly initialized using the topic ODP description. The size of each query was a random number between 1 and 32. The crossover probability was set to 0.7 and the mutation probability was 0.03.

4.1 Analyzing the Evolution of the Single-objective EAs

In our first experimental setting, we run a single-objective EA for 200 generations with the purpose of maximizing *Precision@10*. Both conjunctive (AND) and disjunctive (OR) queries were tested. In figure 1 we show the evolution of *Precision@10* (left) and *Recall* (right) for the ODP topic *BUSINESS/BUSINESS_SERVICES/CONSULTING* (CONSULTING). As can be observed in this figure, near-optimal queries were obtained for both AND- and OR-queries after a small number of generations. However, this was at the cost of very low *Recall* values.

Unsurprisingly, very low *Precision@10* values were achieved when the objective to be maximized was *Recall*, as shown in figure 2. Note, in addition that although high *Recall* values were achieved for OR-queries, this was not the case for AND-queries. Although these results are shown for a single topic, analysis of the rest of the topics yielded similar behavior.

While increasing the level of one performance measure at the cost of reducing the other is sometimes acceptable, we are typically interested in improving both measures.

¹ <http://dmoz.org>




Fig. 1. The evolution of *Precision@10* (left) and *Recall* (right) for the topic CONSULTING when the objective to be maximized is *Precision@10*.




Fig. 2. The evolution of *Precision@10* (left) and *Recall* (right) for the topic CONSULTING when the objective to be maximized is *Recall*.

4.2 Analyzing the Evolution of the NSGA-II

In order to evolve topical queries that simultaneously attempted to achieve high levels of *Precision@10* and *Recall* we run the NSGA-II algorithm for 300 generations. In figure 3 we plotted the performance achieved at each generation for the topic CONSULTING by looking at *Precision@10* (left), *Recall* (center) and *F** (right). It is interesting to note that when OR-queries were evolved, NSGA-II allowed to achieve very high levels of *Precision@10* without compromising *Recall*. In the case of AND-queries, although high *Precision@10* is eventually achieved, the values for *Recall* remain low.




Fig. 3. The evolution of *Precision@10* (left), *Recall* (center) and *F** (right) for the topic CONSULTING when applying the NSGA-II algorithm.

The trend observed in figure 3 for topic CONSULTING was also observed for the other topics in our corpus. The tables in figure 4 present the means over 10 topics for the first and last generations based on *Precision@10*, *Recall* and *F**. These comparison tables show that the NSGA-II algorithm achieved a substantial improvement throughout the successive generations.

NSGA-II: AND-queries			
TRAINING	mean <i>Precision@10</i>	mean <i>Recall</i>	mean <i>F*</i>
First Generation	0.038	0.059	0.020
Last Generation	0.689	0.196	0.176
NSGA-II: OR-queries			
TRAINING	mean <i>Precision@10</i>	mean <i>Recall</i>	mean <i>F*</i>
First Generation	0.055	0.049	0.022
Last Generation	0.953	0.653	0.766

Fig. 4. First generation vs. last generation of queries evolved with NSGA-II.

4.3 Analyzing the Evolution of the Aggregative MOEA

Finally, we monitored the evolution of the aggregative MOEA throughout 300 generations. The charts in figure 5 show these values for the topic CONSULTING while the tables in figure 6 summarize the mean performance achieved for the 10 topics considered in our evaluation.




Fig. 5. The evolution of *Precision@10* (left), *Recall* (center) and *F** (right) for the topic CONSULTING when applying the aggregative MOEA.

We observe that the performance of the aggregative MOEA is similar to that of NSGA-II. This allows us to conclude that for the objectives analyzed here the results of applying an aggregative approach to rank and evolve queries are comparable to those obtained by a non-aggregative, more computationally intensive approach. In addition, OR-queries have better potential for evolution than AND-queries.

Aggregative MOEA: AND-queries			
TRAINING	mean <i>Precision@10</i>	mean <i>Recall</i>	mean <i>F*</i>
First Generation	0.042	0.060	0.022
Last Generation	0.570	0.358	0.372
Aggregative MOEA: OR-queries			
TRAINING	mean <i>Precision@10</i>	mean <i>Recall</i>	mean <i>F*</i>
First Generation	0.054	0.049	0.022
Last Generation	0.948	0.635	0.749

Fig. 6. First generation vs. last generation of queries evolved with the aggregative MOEA.

4.4 Evaluating Query Performance on the Test Set

In order to determine if the evolved queries are effective when used on a new corpus we computed *Precision@10*, *Recall* and *F** for each of the 10 topics on the test set. The question addressed here is whether the evolved queries are superior to the baseline queries (i.e., queries generated directly from the initial topic description). The tables shown in figures 7 and 8 present this comparison for both NSGA-II and the aggregative MOEA. This comparison shows that the tested algorithms are able to evolve queries with quality considerably superior to that of the queries generated directly from the thematic context. In particular, OR-queries achieved much higher performance than AND-queries.

NSGA-II: AND-queries			
TESTING	mean <i>Precision@10</i>	mean <i>Recall</i>	mean <i>F*</i>
Baseline	0.016	0.075	0.011
Evolved Queries	0.409	0.193	0.156
NSGA-II: OR-queries			
TESTING	mean <i>Precision@10</i>	mean <i>Recall</i>	mean <i>F*</i>
Baseline	0.015	0.056	0.009
Evolved Queries	0.572	0.637	0.577

Fig. 7. Baseline vs. queries evolved with NSGA-II.

5 Conclusions and Future Work

This paper analyzes different strategies for evolving topical queries with single- and multi-objective EAs. We noted that the single-objective EAs present limitations that can be overcome by applying Pareto-based and aggregative techniques. Because the aggregative MOEA is less computationally intensive than NSGA-II, we conclude it is a better choice for the problem studied here. In addition, the aggregative techniques are

Aggregative MOEA: AND-queries			
TESTING	mean <i>Precision@10</i>	mean <i>Recall</i>	mean <i>F*</i>
Baseline	0.018	0.077	0.013
Evolved Queries	0.500	0.336	0.338
Aggregative MOEA: OR-queries			
TESTING	mean <i>Precision@10</i>	mean <i>Recall</i>	mean <i>F*</i>
Baseline	0.015	0.057	0.009
Evolved Queries	0.543	0.643	0.538

Fig. 8. Baseline vs. queries evolved with the aggregative MOEA.

more flexible, allowing to define aggregate fitness functions that favor one objective over the other.

We have also compared the potential for evolution that OR-queries have in comparison to AND-queries, and observed that the former are considerably superior to the latter. Interestingly, most popular search engines, such as Google or Yahoo, use conjunctive matching, which means that it is mandatory for all the query term to appear in a document (or to be associated with the document in some way) for it to be considered. A reason for using conjunctive matching is that user information needs are more intuitively defined using the AND semantics: as the number of terms increases the answer set is more selective and more precisely fits the information needs of the user. However, as observed here, automatically generated queries can achieve better performance for topical search when disjunctive matching is applied.

Another important result derived from our evaluations is that the evolved queries do not overfit the training data. Therefore, once a population of topical queries is available, it can be used to retrieve topical material from sources such as the Web, where no relevance assessments are typically available.

Other attempts to apply EAs in information retrieval include the design of techniques to evolve better document descriptions to aid indexing or clustering [13, 25], term-weight reinforcement in query optimization [12, 28, 23], and optimization of keywords and logical operators [20]. A related research area deals with the development of evolving agents that crawl the Web to search for topical material [19]. A comprehensive literature review of Web-based evolutionary algorithms can be found in [16].

Differently from most of the existing EA proposals to document retrieval, which attempt to tune the weights of the individual terms, our methods take each query as an individual. The proposed method is fully automatic as long as a training corpus is available and the objective functions have been defined. A powerful aspect of this method is the use of a mutation pool containing new candidate terms collected throughout the successive generations of queries. The use of this incrementally generated pool of terms has shown to be effective in aiding the exploration of query space [4].

The techniques presented in this article are applicable to any domain for which it is possible to generate term-based characterizations of a context. In [5] we proposed to apply single-objective genetic algorithms to evolve AND-queries. In that case we used the Web as a corpus for training the algorithm and the optimization criteria were based

on the similarity of the retrieved material to the topic of interest. In the present work, instead of using unlabeled material from the Web we take advantage of a taxonomy of topics from ODP and its associated webpages, which are labeled as relevant or irrelevant to the specific topics.

In the future we expect to run additional experiments applying other objective functions coming from the information retrieval and Web search communities as well as ad-hoc ones. Moreover, we plan to test different parameter settings for the EAs. In this work we look at queries with simple syntaxes. An interesting follow-up study concerns applying genetic programming to evolve queries with more complex syntaxes, including boolean operators and other special commands.

Acknowledgement

This research work is supported by CIC (TICs-Sinergia 2008), Agencia Nacional de Promoción Científica y Tecnológica (PICT 2005 Nro. 32373), Universidad Nacional del Sur (PGI 24/ZN13) and Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET).

References

1. Ricardo Baeza-Yates and Berthier Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley, 1999.
2. Stefan Bleuler, Marco Laumanns, Lothar Thiele, and Eckart Zitzler. PISA – A Platform and Programming Language Independent Interface for Search Algorithms. In *Evolutionary Multi-Criterion Optimization*. Springer-Verlag, 2003.
3. Jay Budzik, Kristian J. Hammond, and Larry Birnbaum. Information access in context. *Knowledge based systems*, 14(1–2):37–53, 2001.
4. Rocío L. Cecchini, Carlos M. Lorenzetti, Ana G. Maguitman, and Nélide B. Brignole. Genetic algorithms for topical web search: A study of different mutation rates. In *XIII Congreso Argentino de Ciencias de la Computación*, October 2007.
5. Rocío L. Cecchini, Carlos M. Lorenzetti, Ana G. Maguitman, and Nélide B. Brignole. Using genetic algorithms to evolve a population of topical queries. *Information Processing and Management*, 44(6):1863–1878, 2008.
6. Soumen Chakrabarti, Martin van den Berg, and Byron Dom. Focused crawling: a new approach to topic-specific Web resource discovery. *Computer Networks (Amsterdam, Netherlands: 1999)*, 31(11–16):1623–1640, 1999.
7. Carlos A. Coello Coello. Theoretical and numerical constraint handling techniques used with evolutionary algorithms: A survey of the state of the art. *Computer Methods in Applied Mechanics and Engineering*, 191(11-12):1245–1287, 2002.
8. Carlos A. Coello Coello, Gary B. Lamont, and David A. Van Veldhuizen. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Springer, New York, second edition, September 2007. ISBN 978-0-387-33254-3.
9. Kalyanmoy Deb. *Multi-Objective Optimization Using Evolutionary Algorithms*. John Wiley & Sons, 2004.
10. Kalyanmoy Deb, Samir Agrawal, Amrit Pratap, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions Evolutionary Computation*, 6(2):182–197, 2002.

11. Agoston E. Eiben and James E. Smith. *Introduction to Evolutionary Computing*. Springer, 2003.
12. Ophir Frieder and Hava T. Siegelmann. On the allocation of documents in multiprocessor information retrieval systems. In *Proceedings of the Fourteenth Annual International ACM/SIGIR Conference on Research and Development in Information Retrieval*, pages 230–239, 1991.
13. Michael Gordon. Probabilistic and genetic algorithms in document retrieval. *Communications of the ACM*, 31(10):1208–1218, 1988.
14. John H. Holland. *Adaptation in natural and artificial systems*. Ann Arbor: The University of Michigan Press, 1975.
15. Henry Kautz, Bart Selman, and Mehul Shah. The hidden Web. *AI Magazine*, 18(2):27–36, 1997.
16. Ibrahim Kushchu. Web-based evolutionary and adaptive information retrieval. *IEEE Transactions on Evolutionary Computation*, 9(2):117–125, April 2005.
17. David B. Leake, Travis Bauer, Ana Maguitman, and David C. Wilson. Capture, storage and reuse of lessons about information resources: Supporting task-based information search. In *Proceedings of the AAAI-00 Workshop on Intelligent Lessons Learned Systems*. Austin, Texas, pages 33–37. AAAI Press, 2000.
18. Fu Lin and Guiming He. An improved genetic algorithm for multi-objective optimization. In *PDCAT*, pages 938–940. IEEE Computer Society, 2005.
19. Filippo Menczer, Gautam Pant, and Padmini Srinivasan. Topical web crawlers: Evaluating adaptive algorithms. *ACM Transactions on Internet Technology (TOIT)*, 4(4):378–419, November 2004.
20. Zacharis Z. Nick and Panayiotopoulos Themis. Web search using a genetic algorithm. *IEEE Internet Computing*, 5(2):18–26, 2001.
21. Iadh Ounis, Christina Lioma, Craig Macdonald, and Vassilis Plachouras. Research directions in Terrier: a search engine for advanced retrieval on the web. *Novatica/UPGRADE Special Issue on Web Information Access, Ricardo Baeza-Yates et al. (Eds), Invited Paper*, VIII(1):49–56, February 2007.
22. Vilfredo Pareto. *Cours d'Economie Politique*. Droz, Genève, 1896.
23. Frederick E. Petry, Bill P. Buckles, Dev Prabhu, and Donald H. Kraft. Fuzzy information retrieval using genetic algorithms and relevance feedback. In *Proceedings of the ASIS Annual Meeting*, volume 30, pages 122–125, October 1993.
24. Gregorio Toscano Pulido. Optimización Multiobjetivo Usando un Micro Algoritmo Genético. Master's thesis, Maestría en Inteligencia Artificial, Universidad Veracruzana, Xalapa, Veracruz, México, September 2001.
25. Vijay Raghavan and Brijesh Agarwal. Optimal determination of user-oriented clusters: an application for the reproductive plan. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 241–246, Mahwah, NJ, USA, 1987. Lawrence Erlbaum Associates, Inc.
26. Cornelis J. Van Rijsbergen. *Information Retrieval*. Butterworth-Heinemann, 1979.
27. Gabriel Somlo and Adele E. Howe. Querytracker: An agent for tracking persistent information needs. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 488–495, Washington, DC, USA, 2004. IEEE Computer Society.
28. Jing-Jye Yang and Robert Korfhage. Query optimization in information retrieval using genetic algorithms. In *Proceedings of the 5th International Conference on Genetic Algorithms*, pages 603–613, San Francisco, CA, USA, 1993. Morgan Kaufmann Publishers Inc.