

Research Issues in Multi Criteria Decision Making (MCDM): The Impact of Uncertainty in Solution Evaluation

Piero P. Bonissone

General Electric Global Research

Niskayuna, NY 12309, USA

bonissone@crd.ge.com

Abstract

We consider Multi Criteria Decision Making (MCDM) as the conjunction of three components: *search*, *preference tradeoffs*, and *interactive visualization*. The first MCDM component is the search process over the space of possible solutions to identify the non-dominated solutions that compose the Pareto set. The second component is the preference tradeoff process to select a single solution (or a small subset of solutions) from the Pareto set. The third component is the interactive visualization process to embed the decision-maker in the solution refinement and selection loop. We focus on the intersection of these three components and we highlight some research challenges, representing gaps in the intersection. We introduce a requirement framework to compare most MCDM problems, their solutions, and analyze their performances. We focus on the impact of uncertainty in each of these components and illustrate it with a real-world application.

Keywords: Multi Criteria Decision Making, Pareto search, preference aggregation, interactive visualization, uncertainty management, MCDM research challenges.

1. Multi Criteria Decision-Making

We view Multi Criteria Decision Making (MCDM) as the intersection of three fundamental areas that, traditionally, have been addressed and developed in isolation.

1) ***Solution generation via search.*** We need to search over the space of feasible solutions. We must perform efficient searches in multi- (or sometimes many-) dimensional spaces to identify the non-dominated solutions that

compose the Pareto set. There are usually significant challenges in the search process, such as non-convex solution spaces, complex coupling among objectives and constraints, and high-dimensional objective spaces. This search is driven by solution evaluations that might be nonlinear and probabilistic or imprecise, rather than linear and deterministic. The development of efficient search algorithms has been the goal of Multi-Objective Optimization (MOO), from classical mathematical programming to evolutionary approaches. However MOO's emphasis has been on generating densely sampled, well-distributed Pareto sets, without worrying about the solution selection phase.

2) ***Solution selection via preference aggregation and tradeoff.*** We need to elicit, represent, evaluate, and aggregate the decision-maker's preferences to select a single solution (or a small subset of them) from the Pareto set. These preferences may be ill defined, and state or time-dependent rather than constant values. The aggregation mechanism may be as simple as a linear combination or as complex as knowledge-driven models. The development of methods to capture and aggregate preferences has been the goal of Bayesian and Fuzzy decision-making techniques. However, their emphasis has been on the aggregation mechanisms to select a solution, rather than the solution generation phase.

3) ***Interactive visualization.*** We need a process to enhance our cognitive model of the problem and enable us to perform progressive decisions. We often need to embed the decision-maker in the solution refinement and selection loop. To this end, we need to understand and present the impacts that intermediate tradeoffs in one sub-space could have in the other ones, while allowing him/her to retract or modify any

intermediate decision steps to strike appropriate tradeoff balances [25].

These three research areas are the components of Horn definition of MCDM [11]: search for solutions, aggregation of preferences, and decision selection. Horn consider these cases:

1) Aggregation before search. We reduce the dimensionality of the problem by adding more ordering information, thus transforming a multi-criteria decision making problem into a single-criterion one. This is quite common in optimal control problems, when we select a control action that optimizes an aggregated performance function. However, this computational efficiency has some drawbacks:

- i) No Pareto set is generated, so the DM needs to describe all tradeoffs at once, without knowing the space of alternatives.
- ii) Only one global tradeoff is used to generate decision (usually a linear tradeoff based on a weight vector)
- iii) In such case, it is unsuitable when the performance set is not convex

2) Search before aggregation. We postpone tradeoffs until large numbers of inferior, dominated solutions are eliminated and the efficient Pareto set has been identified. Then we perform a **global** aggregation, considering all objectives at once and making a selection in a one-step aggregation. This approach suffers of drawbacks ii) and iii) above described, since we cannot understand the impacts of intermediate tradeoffs, retracting local decisions, or transforming some objectives into constraints once an acceptable level of performance is achieved. Furthermore, this approach lacks transparency in the decision-making process, since we cannot easily document the decision path that led to the selection.

3) Iterative integration of search aggregation. We start with a multi-criteria search to provide the DM with a preliminary idea of possible tradeoffs. The DM then makes multi-criteria decisions, reducing the search space dimensionality. If needed, a new search is performed in this region of the solution space.

The first two cases represent the independent development of the search and aggregation components. This third case, requiring the intersection of the three components, is needed in complex problems, where we need to justify

and explain our decision making process (for legal, compliance, business or ethical reasons).

2. MCDM Framework

With this perspective, we will focus on the intersection of *search, preference tradeoff, and interactive visualization* and highlight some research challenges, representing gaps in the intersection. We introduce a requirement framework to compare most MCDM problems, their solutions, and analyze their performances. Specifically, we consider the following criteria:

- 1) **Deployment requirements**, e.g., real-time vs. batch mode. Stringent requirement for on-board MCDM deployments, MCDM applications to real-time information streams, etc.
- 2) **Deployment architecture**, e.g., centralized or distributed architectures.
- 3) **Response evaluation**, e.g., deterministic, uncertain, vague, or imprecise. Many MCDM's application use unrealistic assumptions of perfect information in inputs and solution evaluations.
- 4) **Search complexity**, e.g., scalability in high-dimensional performance spaces. Many (evolutionary) search algorithms do not to maintain their performance in 10+ dimensional spaces (typical in design problems). In such cases, we need to use hybrid search methods, leverage our domain knowledge and problem structure, and interactively guide the search.
- 5) **Objectives and constraints complexity**, e.g., non-convex regions that prevents the use of fast searches or aggregations.
- 6) **Uncertainty management**, e.g., vagueness, uncertainty, or imprecision in solutions evaluation.
- 7) **Leveraging domain knowledge** in decision-making process, e.g., internal and external knowledge representations to improve search [4], content-dependent preference aggregation, and customized (case-based) visualization configuration management.
- 8) **Preferences representation/aggregation**, e.g., complete or partial ordering, linear or nonlinear aggregations, etc.
- 9) **Decision-making requirements and methods**, e.g., automated decisions for real-time applications, interactive and progressive decision-making for batch applications.

MCDM ATTRIBUTES			
1) Deployment Requirements	<i>Batch (Off-board)</i>	<i>Real-time (On-board)</i>	
2) Deployment Architecture	<i>Centralized</i>	<i>Distributed</i>	
3) Response Evaluation	<i>Deterministic Models</i>	<i>Stochastic Models Functional Approximation (data-driven models)</i>	<i>Fuzzy models - Qualitative evaluation</i>
4) Search Method	<i>Simple EMOO</i>	<i>Knowledge-enhanced, hybrid EMOO</i>	<i>Fusion of multiple search methods</i>
5) Objectives & Constraints Complexity	<i>Few objectives; Convex regions</i>	<i>High-dimensional objectives; Non-convex regions</i>	
6) Uncertainty Management	<i>Uncertainty measures implicitly captured in objectives</i>	<i>Explicit externally driven uncertainty management (e.g. fusion on multiple models)</i>	
7) Leveraging Domain Knowledge	<i>Customized data structures & operators in evolutionary search</i>	<i>Meta-heuristics to guide evolutionary search; fuzzy rule-based preference aggregation functions</i>	<i>Self-tuning fuzzy rule-based functions; context-dependent (Case-based) visualization config mgmt</i>
8) Preferences	<i>Complete Ordering (numerical, ordinal or cardinal)</i>	<i>Partial Ordering (Imprecise, fuzzy, linguistic, preferences assigned to objective subsets)</i>	
9) Decision-making Requirements & Methods	<i>Automated decision-making via constraints and weights</i>	<i>Interactive graphical with human in loop</i>	
10) Update Requirements for Solution fidelity	<i>Implicit via update of problem descriptors in database</i>	<i>Explicit via periodic retraining of data-driven models, or other adaptation mechanisms</i>	<i>Explicit via autonomous retraining of data-driven models, or other adaptation mechanisms</i>

Figure 1. Framework to describe Multi Criteria Decision Making Problems

10) *Update requirements* for solution fidelity, e.g., retraining/updating deployed data-driven or knowledge-driven evaluators.

An earlier version of this framework was first proposed in [3] and used in [23] to illustrate two MCDM applications.

3. MCDM Research Challenges

In this paper we revisit this framework to distill and describe six major research challenges in MCDM. Then, we will focus on the first one, related to the management of uncertainty.

- 1) Uncertainty, fuzziness, and imprecision in the inputs and in the solution evaluation methods
- 2) High-dimensional objective space and interactive decision-making process
- 3) Fuzzy preferences, context dependent aggregation
- 4) Leveraging domain knowledge in decision making process
- 5) Real-time, high throughput requirements and model learning update
- 6) Deployment and maintenance of adaptive, distributed decision-making systems

The first challenge, perhaps the most relevant to the topics of this conference, addresses questions such as the **impact** of uncertainty on the search algorithms - in terms of dominance, distances, filtering, etc. - and the **management** of uncertainty as it impacts solution sensitivity, extrapolation errors, tracking region of confidence competence etc.

The second challenge is caused by the **complexity** of searching for solutions in **high-dimensional spaces**. Many multi-objective evolutionary algorithms perform very well (i.e., they produce accurate, densely sampled Pareto sets) in low-dimensional spaces, but their time-complexity increases exponentially (e.g., they do not scale-up) in high-dimensional spaces [21], [18]. Researchers have proposed learning-follows decomposition methods of the space [14], hierarchical fair competition (HFC) [12], and dimensionality reduction via interactive human interaction – transforming objectives into constraints after partial search [9]. Beyond managing search complexity, we also need to properly visualize and interact with such high-dimensional objective space. This usually stresses our cognitive limitations, especially if the evaluations are uncertain.

The third challenge is related to the difficulties of **eliciting, representing, and aggregating preferences** to guide tradeoff and selection decisions. It has been noted in [6] that “numerical scales (either cardinal or ordinal) provide a complete order over the preferences. However, they are quite limited in their representational power, as they require a synthetic precision that does not reflect the true meaning of the DM’s preferences.” Alternatively, the use of linguistic or fuzzy preferences with local semantics (membership functions defining the meaning of the labels as fuzzy preferences) provides a more natural representation but only induces a partial order on the universe of preferences. Partial or global aggregations of these preferences must reflect tradeoffs that usually are not easy to elicit and might not be uniform in the solution space. These aggregations have been a research focus in the fuzzy logic community.

The fourth challenge, extending the third one, deals with the efficient **representation of domain knowledge** to improve search, preference aggregation, and interaction visualization of solutions. Domain knowledge can be embedded in evolutionary search, as described in [4], using internal representations, such as customized data structures and variational operators, or external representations, such as meta-heuristics. Knowledge in preference aggregation has been partially described in the third challenge. We could further add the use of local or global search to tune the parameters (term sets) and structures (rules) of this representation, based on past cases (selection and performance). Finally, domain knowledge can be used to configure the solutions visualization, via customized subsets of projections (or specific orders in a parallel coordinates plot), which could be initialized by a case-based reasoner.

The fifth challenge is created by the need of porting MCDM techniques to **real-time applications** using Field Programmable Gate Arrays (FPGA), ASICs, Cell Broadband Engine, etc. For example, in [5], the authors note that “since evolutionary algorithms work with a population of solutions, parallelizing the fitness computation has the benefit of significant speedup. For the efficient execution of applications requiring high-frequency multi-objective optimization constrained by the size of the computational unit, it is desirable to develop

a multi-objective evolutionary technique that enables high optimization speed-ups with a small computational footprint. An example of such an application is in unmanned vehicle control, where a high optimization speed is required, the computational hardware footprint and weight constraints are severe, and the domain demands simultaneous consideration and optimization of multiple conflicting objectives such as thrust and range given varying mission needs while operating with a finite fuel resource.” In the above reference the authors run an Evolutionary Multi Objective Optimization (EMOO) problem on an FPGA, achieving performance improvements over 300 times faster than on a 3GHz workstation.

The sixth challenge is the result of allowing **distributed MCDM systems** (e.g., adaptive solution evaluators running on multiple platforms) to **learn and adapt** during their deployment. The key issues are the management of local adaptation, knowledge sharing, local verification and validation (V&V), model updating, and version control of multiple, distributed, deployed learning systems.

Due to space constraints, we will focus on the first challenge, i.e., the presence of uncertainty, fuzziness, and imprecision in the inputs and in the solution evaluation methods. Then we will describe an illustrative example, containing some useful approaches to this challenge.

4. First Challenge: MCDM with uncertain solution evaluation

In many real-world applications, the solution evaluation is not a deterministic process. As such, the output is not a crisp value, but rather an interval, a distribution, or a fuzzy value. For instance, Paenke uses local approximate models to estimate expected fitness and variance [17]. One of the most common sources of uncertainty is the use of a function approximation model, such as a neural network (NN), to evaluate the fitness vector of a solution during the evolutionary search. Several researchers have studied the error bounds of NN [1], and special types NN, such as Radial Basis Functions [27]. If we use NN’s to compute each solution’s coordinates in performance space, we will have a hyper-rectangle (or a hyper-ellipsoid) rather than a point as the solution’s image under the NN mapping. A way to reduce this uncertainty is to develop multiple models (e.g., NNs),

trained on different subsets of the training set and of the feature space – to increase their diversity – and fuse their outputs to increase the overall accuracy [2], [28].

In any case, we need to redefine the concept of dominance. This can be achieved by a variety of methods, such as including a tolerance parameter in the dominance filter, e.g. ϵ -indistinguishability [13], extending the concept of dominance relations to fuzzy arguments [26], or treating the evaluations as crisp, expected values and then compensating for this simplification by limiting the magnitude of solution changes [23], [24].

Reference [19] exemplifies the impact of uncertainty in search. The authors modified an NSGA-II evolutionary algorithm to accommodate the uncertainty derived from imprecise fitness function evaluation. To this extent, they modified and extended the computation of three elements in the NSGA-II:

- 1) Precedence (dominance) operator
- 2) Non-dominated sorting of the individuals
- 3) Crowding distance – used to assess local solutions density and uniformly sample the front.

A second source of uncertainty is caused by extrapolation errors in the solution evaluation. This occurs when we operate a model outside of its competence region, i.e., if the image of the training set under the mapping results in interior points that are far from the Pareto front. This situation will be illustrated in the application example described in the following section.

An additional uncertainty source is the use of imprecise constraints, which define a fuzzy feasibility region in the solution space. A partial constraint satisfaction score could be computed and used to discount the performance metrics of the obtained solutions, in a manner similar to fuzzy scheduling [22].

5. Example of MCDM application with uncertain solution evaluation

An illustrative case for this first MCDM challenge is the management of a power plant, which integrates predictive modeling based on neural networks, optimization based on multi-objective evolutionary algorithms, and automated decision-making based on Pareto set down-selection techniques [23]. [24]. The

predictive models are adaptive, and continually update themselves to reflect with high fidelity the gradually changing underlying system dynamics. The integrated approach, embedded in a real-time plant optimization and control software environment has been deployed to dynamically optimize emissions and efficiency while simultaneously meeting load demands and other operational constraints in a complex real-world power plant.

We used nonlinear neural-network models to generate mappings between the inputs space of control variables and time-variable ambient uncontrollable variables, and the various outputs (objectives and constraints) of interest. We used first-principles-based methods and domain-knowledge to identify the relevant NN inputs.

The evolutionary multi-objective optimizer generates test inputs/set-points and receives as feedback the corresponding output performance metrics after transformation by suitable objective (performance) functions.

The multi-objective optimizer uses this feedback to generate and identify the Pareto-optimal set of input-output vector tuples that satisfy operational constraints.

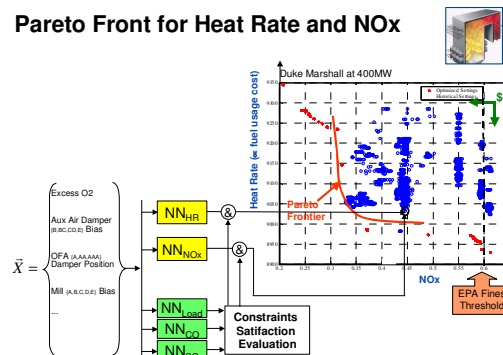


Figure 2. NN's mapping solution \vec{X} to (HR, NOx) plane.

A decision function is superimposed on this Pareto-optimal set of input-output vector tuples to identify a deployable input-output vector, which is then dispatched to the underlying plant control system, or recommended to the operator for execution. This is illustrated in Figure 2.

Uncertainty derived by the functional approximation.

A more accurate representation of the image of a solution \vec{X} under the NN mapping would be an

ellipsoid rather than a point in (HR, NOx). This would more accurately represent the uncertainty induced by the NN mapping.

To reduce this kind of uncertainty, we use a committee of predictive models (NN's) and an intelligent fusion of their predictions. Fusing the outputs from an ensemble of models in an effective way can often boost overall model accuracy. A prerequisite for a successful fusion is to create a strong diversity of the models to be fused [15]. This concept is further developed in [28], wherein we present a novel method called locally weighted fusion, which aggregates the results of multiple predictive models based on local accuracy measures of these models in the neighborhood of the probe point for which we want to make a prediction. This fusion method may be applied to develop highly accurate predictive models. The locally weighted fusion method boosts the predictive performance by 20~40% over the baseline single model approach for the various prediction targets. In [2], we further refine the performance by using CART algorithms to pre-compile a segmentation of the input space for each model. Using this approach we improve the predictive performance by 34~48% over the same baseline. Relative to these approaches, fusion strategies that apply averaging or globally weighting only produce a 2~6% performance boost over the baseline.

Uncertainty derived by extrapolation error. A Pareto-optimal front that jointly minimizes NOx and Heat Rate (inversely related to efficiency) for a 400MW target load demand in a 400MW power plant is shown in Figure 3. In this figure, the circles show the range of historical operating points from a NOx—Heat Rate perspective. The stars and inter-connecting line show the optimized Pareto frontier in the NOx—Heat Rate space. Each point not on this frontier is a sub-optimal operating point—the goal being the operation of the plant or process at a Pareto optimal point at all times. Moving the system operation from the interior of the decision space to the Pareto frontier results in a large operational savings opportunity. All these points are feasible solutions, satisfying load, CO, and SO emission constraints.

The Pareto frontier in NOx—Heat Rate space, identified by the multi-objective search and depicted in Figure 3, is clipped by the systematic application of profit-based and operational-need

constraints for each of NOx and Heat Rate. Next, a solution from this reduced frontier that is the closest in inputs space to the current plant state is selected and transmitted to the plant control system. Such an approach minimizes the state deviations while achieving Pareto-optimal operation and end-user acceptability. The decision-making approach further highlights the inherent flexibility of Pareto frontier techniques whereby the entire efficient set of solutions is first identified without regard to situation specific down-selection, and later a flexible decision function is superimposed to identify a deployable input set (or set point).

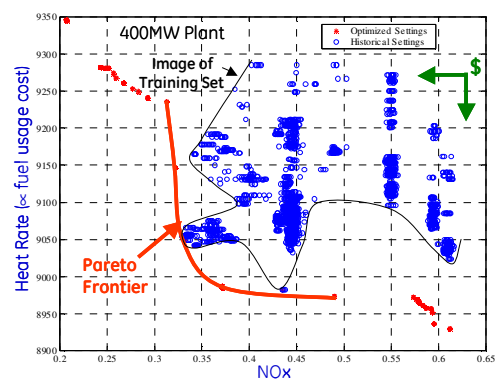


Figure 3. Pareto tradeoff between NOx emissions and Heat Rate for a 400MW power plant

The results of these experiments are reported in [23], where we describe the characteristics of this application within the framework proposed in section 2.

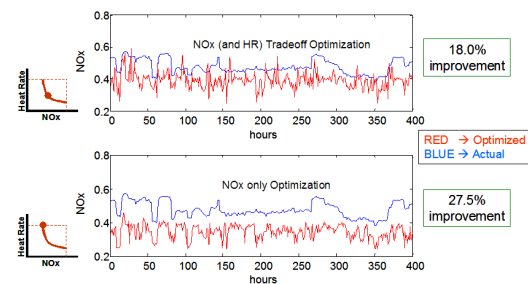


Figure 4: Pareto frontier tradeoff optimization of NOx and Heat Rate and NOx only optimization.

Figure 4 shows the performance gains that may be achieved in NOx emissions using this decision-making approach. When a decision-making function is used which simultaneously considers a tradeoff Pareto point at each instant, roughly 18% improvement in NOx emissions may be achieved (upper figure half). However, if the optimization favors a NOx minimization

that satisfies a given Heat Rate constraint, more significant NO_x emissions improvement is possible (lower figure half). Similarly, 1-2% improvements in Heat Rate are possible. In a typical power plant setting, such savings in NO_x and Heat Rate are very significant and could lead to operational savings of hundreds of thousands to millions of dollars per year.

6. Conclusions

Summary. This paper formalizes some of the author's ideas originally presented in oral form at the First IEEE Symposium of Computational Intelligence in Multicriteria Decision Making [3]. We described MCDM as the intersection of three components: 1) a search process over the space of possible solutions; 2) a preference tradeoff process to guide the down-selection; and 3) an interactive visualization process to understand the tradeoff impacts.

We presented a framework to represent different requirements for MCDM problems. Within this framework, we highlighted MCDM research challenges ranging from uncertainty in inputs and solution evaluations, to high-dimensional objective spaces and cognitive limitations, interactive decision-making processes, fuzzy preferences, leveraging domain knowledge to customize search, aggregations and visualizations, real-time, high throughput requirements, and deployment/maintenance of adaptive, distributed decision-making systems.

Finally, we focused on the impact of uncertainty in inputs and solution evaluations, and we illustrated it with a case study describing the development and deployment of a MCDM system to optimize power plant management.

Current and Future Work. We organized this special session on MCDM to explore various facets of the framework and research challenges defined above. These efforts are illustrated by [20], where Sanchez et al. address the first challenge, the presence of uncertainty in solution evaluation, by developing a fuzzy rule-based system for a diagnostics problem. In [16], Montero et al. address the second challenge, supporting the decision-maker cognitive limitations, by estimating objectives from preferences and interacting via graphical representations. Two other papers [8], [29] address the third challenge, preference aggregation, by leveraging goal interactions to structure the aggregation, and by analyzing the

properties of a lexicographic, prioritized aggregation, respectively. In [10], the authors address the fourth challenge, leveraging domain knowledge, by customizing the initial population and the operators of an evolutionary search to improve the legibility of fuzzy rule sets, while tuning their parameters for improved accuracy. Finally, in [7], the authors examine a multiple ant colony system as an alternative search method, illustrating it with a real-world problem.

In conclusions, we believe that we just started to identify and describe some key research challenges in MCDM. Addressing such challenges will result in the development of robust, efficient techniques that will extend MCDM applicability to complex, real-world problems that are currently outside its reach.

References

- [1] A. Barron (1994). "Approximation and estimation bounds for artificial neural networks", *Machine Learning* 14(1): 115-133
- [2] P. Bonissone, F. Xue, and R. Subbu (2008). "Fast Meta-models for Local Fusion of Multiple Predictive Models", *Applied Soft Computing Journal*, to appear
- [3] P. Bonissone (2007). "Multi-Criteria Decision-Making: The Intersection of Search, Preference Tradeoff, and Interaction Visualization Processes", *IEEE Symposium on Computational Intelligence in Multi-Criteria Decision-Making (MCDM 2007)*, page 1, Honolulu, HI, USA
- [4] P. Bonissone, R. Subbu, N. Eklund, and T. Kiehler (2006). "Evolutionary Algorithms + Domain Knowledge = Real-World Evolutionary Computation", *IEEE Transactions on Evolutionary Computation*, 10(3): 256-280
- [5] S. Bonissone and R. Subbu (2007). "Evolutionary Multiobjective Optimization on a Chip", *2007 IEEE Workshop on Evolvable and Adaptive Hardware*, April 1-5, 2007, Honolulu, HI, USA
- [6] S. Bonissone (2001). "Evolutionary Algorithms for Multi-Objective Optimization: Fuzzy Preference Aggregation and Multi-Sexual EA's", *SPIE 2001: Applications and Science of Neural Networks, Fuzzy Systems, and Evolutionary Computation IV*, pp. 157-164, San Diego, CA
- [7] M. Chica, O. Cordon, S. Damas, J. Bautista, J. Pereira (2008). "A Multiobjective Ant Colony Optimization Algorithm for the 1/3 Variant of the Time and Space Assembly Line Balancing Problem", *IPMU 2008*, Malaga, Spain

- [8] R. Felix (2008). "Multi Criteria Decision Making (MCDM): Management Of Aggregation Complexity Through Fuzzy Interactions Between Goals Or Criteria", *IPMU 2008*, Malaga, Spain
- [9] P. Fleming, R.C. Purshouse, R.J. Lygoe (2005). "Many-Objectives Optimization: An Engineering Design Perspective", *Int. Conf. on Evolutionary Multi-Criterion Optimization, EMO 2005*, pp. 14-32, Springer LNCS 3410, Guanajuato, Mexico
- [10] M. J. Gacto, R. Alcalá and F. Herrera (2008). "Multi-Objective Genetic Fuzzy Systems: On the Necessity of Including Expert Knowledge in the MOEA Design Process", *IPMU 2008*, Malaga, Spain
- [11] J. Horn (1997). Multicriterion Decision Making, *Handbook of Evolutionary Computation*, (T. Back, D. Fogel and Z. Michalewicz, eds.), pages F1.9:1-F1.9:15, IOP Pub. and Oxford Univ. Press
- [12] J. Hu, K. Seo, Z. Fan, R. Rosenberg, and E. Goodman (2003). "HEMO: A Sustainable Multi-Objective Evolutionary Optimization Framework", *2003 Genetic and Evolutionary Computation Conference*, Chicago, Springer, Lecture Notes in Computer Science, pp. 1029-1040
- [13] N. Iyer (2002). "A family of dominance rules for multiattribute decision making under uncertainty", *IEEE Trans. on Systems, Man and Cybernetics, Part A*, 33(4): 441 – 450
- [14] R. Kumar, P. Rockett (1998). "Multiobjective Genetic Algorithm Partitioning for Hierarchical Learning of High-Dimensional Pattern Spaces: A Learning-Follows-Decomposition Strategy", *IEEE Trans. on Neural Networks*, 9(5): 822-830
- [15] L. Kuncheva, C.J. Whitaker (2003). "Measures of Diversity in Classifier Ensembles", *Machine Learning*, 51:181–207
- [16] J. Montero, D. Gomez, S. Muñoz (2008). "Fuzzy information representation for decision aiding", *IPMU 2008*, Malaga, Spain
- [17] I. Paenke, J. Branke, Y. Jin (2006). "Efficient search for robust solutions by means of evolutionary algorithms and fitness approximation", *IEEE Trans. on Evolutionary Computation*, 10(4): 405-420
- [18] L. Rachmawati, D. Srinivasan (2006). "Preference incorporation in multi-objective evolutionary algorithms: a survey", *IEEE Congress on Evolutionary Computation*, pp. 962-968, Vancouver, BC, Canada, July 16-21, 2006
- [19] L. Sanchez, I. Couso and J. Casillas (2007). "Modeling Vague Data with Genetic Fuzzy Systems under a Combination of Crisp and Imprecise Criteria", *IEEE Symposium on Computational Intelligence in Multi-Criteria Decision-Making (MCDM 2007)*, pp. 30-37, Honolulu, HI, USA
- [20] L. Sanchez, A. Palacios, M. R. Suarez, I. Couso (2008). "Graphical exploratory analysis of vague data in the early diagnosis of dyslexia", *IPMU 2008*, Malaga, Spain
- [21] K. Sastry, M. Pelikan, and D.E. Goldberg (2005). "Decomposable problems, niching, and scalability of multiobjective estimation of distribution algorithms", (*IlligAL Report # 2005004*), Urbana, IL
- [22] W. Slany (1996). "Scheduling as a Fuzzy Multiple Criteria Optimization Problem", *Fuzzy Sets and Systems*, 78, pp. 197-222
- [23] R. Subbu, P. Bonissone, S. Bollapragada, K. Chalermkraivuth, N. Eklund, N. Iyer, R. Shah, F. Xue and W. Yan (2007). "A review of two industrial deployments of multi-criteria decision-making systems at General Electric", *IEEE Symposium on Computational Intelligence in Multi-Criteria Decision-Making (MCDM 2007)*, pp. 136-145, Honolulu, HI, USA
- [24] R. Subbu, P. Bonissone, N. Eklund, W. Yan, N. Iyer, F. Xue, and R. Shah (2006). "Management of Complex Dynamic Systems based on Model-predictive Multi-objective Optimization", *2006 IEEE International Conference on Computational Intelligence for Measurement Systems and Applications*, La Coruña, Spain
- [25] R. Subbu, P. Bonissone, N. Eklund, S. Bollapragada, and K. Chalermkraivuth (2005). "Multiobjective Financial Portfolio Design: A Hybrid Evolutionary Approach", *2005 IEEE CEC*, vol. 2: 1722-1729
- [26] R. Tong, P. Bonissone (1980). "A Linguistic Approach to Decisionmaking with Fuzzy Sets", *IEEE Trans. on Systems, Man, and Cybernetics*, 10(11): 716-723
- [27] N. W. Townsend, L. Tarassenko (1997). "Estimations of Error Bounds For RBF Networks", *Artificial Neural Networks*, 7-9 July 1997, Conf. Publication No. 440, IEE
- [28] F. Xue, R. Subbu, P. Bonissone (2006). "Locally Weighted Fusion of Multiple Predictive Models", *IEEE Int. Joint Conf. on Neural Networks (IJCNN'06)*, Vancouver, BC, Canada
- [29] R. Yager (2008). "Lexicographically Prioritized Multicriteria Decisions Using Scoring Functions", *IPMU 2008*, Malaga, Spain