An Introduction and Companion to the Third Workshop on Multiobjective Problem Solving from Nature

Joshua Knowles[1], David Corne[2] and Kalyanmoy Deb[3]

[1]University of Manchester, UK. Email: j.knowles@manchester.ac.uk
[2]Heriot-Watt University, UK. Email: dwcorne@macs.hw.ac.uk
[3]IIT Kanpur, INDIA. Email: deb@iitk.ac.in

September 7th 2006

Abstract. The participants will discuss the virtues and vices of an evolutionary and multiobjective approach to problem solving. To inform discussions, a number of talks (mostly invited) will present material on how multiple objectives are being used to solve particular problem types or problem issues.

1 Introduction

This workshop, held at PPSN IX in Reykjavik University, will have a slightly different format and focus from previous ones in the series. Rather than have an open call for new research material, we thought we would invite most of our speakers to give talks on specific topics around a common theme: the theme being, "the origins and benefits of multiple objectives", which is the subtitle of the workshop itself. In this short introduction paper, we will just expand a little bit on what this theme means and how the speakers will be focusing on particular aspects of it.

The ideas presented herein should be read with a background context in mind. That context is the growing popularity of multiobjective optimization in evolutionary computation, and a substantial existing body of technical literature. Much of the work done on evolutionary multiobjective optimization (EMO) over the past 15 years has focused on:

- algorithm design (efficiency, diversity, convergence)
- test problems and functions
- performance assessment
- theory (non-trivial convergence results),

and there are some excellent reviews that give a picture of this development [2, 6].

Arguably, though, there has been a little bit less attention on what the real virtues and vices of the whole multiobjective approach are, how multiple objectives are being used to solve practical problems, how helpful it is, and so on. In particular, we might consider whether there are niches where EMO is particularly suited: niches as application areas, or as ways of using the multiple objectives. Conversely, there may be niches where EMO doesn't really provide benefit. If that is the case, can we say why?

2 Vices and virtues of evolutionary and multiobjective problem solving

According to Fogel [5], evolutionary computation has a number of distinct advantages as a general problem solving technique. Among these are:

1. Conceptual simplicity of the method, and freedom to tinker with the algorithm/representation, etc.
2. Broad applicability: can handle ‘functions’ that are discrete, continuous, mixed-integer, multi-modal, dynamic, constrained and ”black-box”. Moreover, Fogel underlines the importance of ‘playing with’ (our phrasing) the objective function to arrive at a successful one.
3. Potential to hybridize
4. Parallelism
5. Robustness to dynamic changes.

A multiobjective approach maintains and potentially extends the above advantages in our opinion. For each of the points above, there are corresponding virtues from adding ‘multi-objectivity’:

1. Simplicity maintained but an extra dimension to the tinkering obtained. There are now many more ways to assign fitness, to think about diversity, and so on.
2. Broadens applicability because several such functions (maybe of completely different type, complexity etc.) can be optimized. We can also create additional functions to help the optimization find what we truly desire. This aspect is really key to the multiobjective agenda.
3. We can now have different subroutines being used to optimize different objective functions.
4. Each objective can be evaluated on a different processor, which adds to parallelism. Also other schemes are possible.
5. Robustness can be enhanced by making it an explicit objective.

On the side of vices, there is also a common pitfall with evolutionary computation (general) and multiobjective problem solving. This is that we can begin to see every problem as a potential application of the approach, even though other methods may be more appropriate. This may be particularly true in heavily researched areas, such as supervised learning, where other techniques are highly advanced.

Moreover, a multiobjective approach is subject to another pitfall more of its own; it resists simple evaluation because the subjectivity of a good solution to a problem is highlighted when several or many different objectives are considered simultaneously. This can lead to a sloppiness in analysing the real utility of solutions, or even in just stating how a “final” solution, or set of solutions, should be identified.

3 Solution concepts and decision making

There is a body of work in the management science and operations research literature (and also in mirco-economics) on the subject of ‘utility’, which states how a set of alternatives can be evaluated, starting often from an analysis of pair-wise preference information, through the formation of objectives, constraints, weightings and priorities, and leading finally to a rational choice, or decision (see **??** and references therein). This area is called MCDM or MCDA, and it is an important cousin to EMO for obvious reasons. However, while knowledge of this area is important and useful, it is also true that MCDM is not quite suitable as a method for selecting solutions in all the cases where an EA might be applied. Concepts might be adaptable from MCDM, but we might need methods that do not rely on human decision-makers; rather, we may need methods where the objectives themselves are formulated automatically, discarded or retained during search, changed or added to, and so on. An organic and adaptive set of criteria is in keeping with what evolutionary algorithms do best (see the previously mentioned paper by Fogel).

If MCDM is one useful component for arriving at solutions, or sets of solutions, another one might be the notion of a 'solution concept'. A solution concept is a rule that unambiguously partitions the search/problem space into solutions and non-solutions. E.g. a solution concept could be: solution iff phenotype is Pareto optimal. Another solution concept could be: solution iff population of phenotypes comprises complete Pareto optimal set. Another solution concept could be: solution iff population of phenotypes satisfices certain goals. This framework may be too restrictive but it would give some guard against the pitfall of not being able to evaluate progress objectively.

The recent adoption of the solution concept idea in the area of (artificial) co-evolution [4] seems to be useful; perhaps a similar use in EMO would be advised. Certainly, the solution concept idea allows us to formalize other notions of optimality beyond Pareto optimality. There are many ways in which solutions in a multiobjective space might be evaluated, individually or as an ensemble. Examples are *favour* [3], Bentley's weighted average rank (WAR) [1], and the hypervolume dominated by a set of points [7]. However, up to now, these are not really seen as the ultimate goal of search; often a different criterion is used to assess the product of the search than was used to conduct it, without making these criteria explicit. It may be the case that optimizing a goal which is seemingly orthogonal to the stated goal can help the search, but this can only be appreciated when the overall goal is stated unambiguously in the first place.

4 Niches in EMO: the talks

The schedule of the workshop is given in Figure 1. Each of the talks will inform us about how multiple objectives are being used for particular purposes, to achieve certain effects, or in certain application domains. A list of papers related to these talks is included in the bibliography. (The list is to be expanded).

5 Questions and themes for discussion

The workshop chairs will attempt to lead and minute discussions in the workshop. Some themes of interest are given below.

AM 9.00-12.30 (coffee break 10.30-11.00)		
9.00		Introduction to the workshop
9.15	Carlos Coello	A Survey of Constrained-Handling based on EMO
9.45	Carlos Fonseca	Preference Articulation in MOEAs
10.15		***** Mini-discussion and coffee break *****
11.00	Joshua Knowles	Bias, Trivial Solutions, Proxies and Solution Selection
11.30	Yaochu Jin	Modeling Regularity in Multi-Objective Optimization
12.00	Amiram Moshaiov	Multiobjective Cybernetics

PM 14.00-17.30 (coffee break 15.30-16.00)		
14.00		Welcome back, summary of morning session and discussion
14.30	Jonathan Fieldsend	Multiobjective Supervised Learning
15.00	Hisao Ishibuchi	Multiobjective Association Rule Mining
15.30		*********** Coffee ***********
16.00	Stefan Bleuler	Reducing Bloat in GP with Multiple Objectives
16.30	Arjun Chandra	NN Ensemble Construction using Multiple Objectives
17.00-17.30		Discussion and final summary

Fig. 1. Workshop schedule

1. MINIMAL INFORMATION: What is the minimal information that should be given in an applications paper on MOO?

   We need to know why/how it was formulated as a MOP, and what is the ultimate performance criterion of the problem-solving exercise. Anything else?
2. PROCESS: Can we describe the whole problem-solving process? What stages might it have? Defining the search space, defining some objectives, some constraints, solving the problem, realizing some objectives/constraints are not useful, and some are missing, iterate. Could this be formalized or algorithmatized? What machinery already exists for this? Perhaps the DM can't give objectives at all, or rejects at least some of them. How can their preferences be captured in that case?
3. SOLUTION SELECTION: If the goal of a MOP is to choose one single solution, what methods are there for selection of solutions from a Pareto front?
4. FOCUSING: If the goal of a MOP is to choose one single solution, why do Pareto optimization at all? [There are examples why] What alternatives are there?
5. ENSEMBLES: What applications call for whole sets of solutions as the ultimate answer? In these cases, is the required set the Pareto front, or something else (smaller)? What other types of relation lead to sets of solutions other than the dominance relation?
6. PURPOSE/CAUSE/ORIGINS of objectives: What happens if some objectives can't be measured directly or accurately or without some kind of bias? Can other objectives be used as proxies? What kind of multiobjective problem does this lead to? What happens if the search space induced by an objective function is not easy to search? Can we use other objectives - what kind of problem does this lead to?
7. NATURE/LIFE: In Natural systems — Evolution and its products: neural systems, ecosystems, etc, and non-evolved systems, is there a sense of optimization of mul-

tiple conflicting objectives? Is there any profit in seeing things in this way? What might the objections to such a view be?

References

1. P. J. Bentley and J. P. Wakefield. Finding Acceptable Solutions in the Pareto-Optimal Range using Multiobjective Genetic Algorithms. In P. K. Chawdhry, R. Roy, and R. K. Pant, editors, *Soft Computing in Engineering Design and Manufacturing*, Part 5, pages 231–240, London, 1997. Springer Verlag London Limited. (Presented at the 2nd On-line World Conference on Soft Computing in Design and Manufacturing (WSC2)).
2. Carlos A. Coello Coello. Evolutionary Multi-Objective Optimization: A Critical Review. In Ruhul Sarker, Masoud Mohammadian, and Xin Yao, editors, *Evolutionary Optimization*, pages 117–146. Kluwer Academic Publishers, New York, 2002. ISBN 0-7923-7654-4.
3. Nicole Drechsler, Rolf Drechsler, and Bernd Becker. Multi-objective Optimisation Based on Relation *favour*. In Eckart Zitzler, Kalyanmoy Deb, Lothar Thiele, Carlos A. Coello Coello, and David Corne, editors, *First International Conference on Evolutionary Multi-Criterion Optimization*, pages 154–166. Springer-Verlag. Lecture Notes in Computer Science No. 1993, 2001.
4. S.G. Ficici. *Solution Concepts in Coevolutionary Algorithms*. PhD thesis, Brandeis University, 2004.
5. D.B. Fogel. The Advantages of Evolutionary Computation. *Biocomputing and emergent computation: Proceedings of BCEC97 table of contents*, pages 1–11, 1997.
6. Carlos M. Fonseca and Peter J. Fleming. An Overview of Evolutionary Algorithms in Multi-objective Optimization. *Evolutionary Computation*, 3(1):1–16, 1995.
7. Eckart Zitzler. *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*. PhD thesis, Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, 1999.

A Other topics and speakers

We invited a number of other speakers to the workshop to talk about the themes we have outlined here and, specifically, the subjects here listed:

Aaron Sloman	Does Nature Solve Problems and Are They Multiobjective?
Marco Laumanns	Spatial Predator-Prey Models of Multiobjective Optimization
Sevan Ficici	Solution Concepts in Co-evolution
Edwin de Jong	Ideal Evaluation and Compression of Objectives
Ian Parmee	Ill-defined Problem Spaces
Peter Bentley	Modularity Causes Multiple Objectives in Natural and Computational Systems
Richard Watson	Problem Decomposition, Modularity and their relation to Multiple Objectives
Mikkel Jensen	Helper Objectives and Multiobjectivization
Justin Boyan	Learning Evaluation Functions for Global Optimization
Katya Rodriguez-Vazquez	Multiobjective GP for Human-Understandable Models
Eckart Zitzler	How Multiple Objectives Help in the Analysis of Biological Data
Kalyanmoy Deb	Unveiling Salient Insights in Engineering Designs with MOEAs

These also serve to understand the idea behind the workshop and our plans for a book on the same theme.