The Fourteenth International Symposium on Artificial Life and Robotics 2009 (AROB 14th '09),
B-Con Plaza, Beppu, Oita, Japan, February 5 - 7, 2009

The Pareto Operating Curve for Risk Minimization in Life and Robotics

Hussein A. Abbass
School of Information Technology and Electrical Engineering,
University of New South Wales, Australian Defence Force Academy,
Canberra, ACT 2600, Australia

Abstract

The use of non-dominance in multi-objective search has traditionally focused on generating the set of non-dominated solutions and choosing an element of this set to implement. In this paper, I will show the richness of the non-dominated set when the objectives (in the multi-objective search problem) represent complexity measures. I will present the concept of Pareto Operating Curves, whereby a system operates along these operating curves based on the risk, complexity and required trade-off it encounters in the environment. Key fundamental features these systems possess are robustness and the ability to adapt in different environments.

1 introduction

The concept of non-dominance has been associated in the evolutionary multi-objective computation (EMO) literature with multi-objective optimization problems (MOP). The topology of the set of non-dominated solutions in the objective space shapes up a curve that is known as the Pareto curve. An optimization problem is traditionally seen in terms of objectives - representing the performance measures of the system - and a set of constraints.

In this paper, we wish to expand the use of the Pareto curve from being a set of trade-off "independent" solutions to an operating curve, where the environment will dictate which solution from this set will be used. Risk is traditionally defined as the impact of uncertainty on objectives. The uncertainty that this paper is concerned with, is the uncertainty in the operating environment. In this case, the objective functions need to reflect the performance of a solution in an operating environment. We will call this operating curve as the Pareto Operating Curve (POC). This paper is the first to discuss the concept of POC.

In the rest of this paper, we will introduce some basic definitions in MOP, followed by discussions of some of my work where the POC was used - although not necessarily discussed explicitly.

2 Multi-objective optimization

Consider a *multi-objective optimization problem* (MOP) as presented below:-

$$\textbf{Optimize } F(\vec{x} \in \Upsilon) \quad (1)$$

$$\textbf{Subject to: } \Upsilon = \{\vec{x} \in R^n | G(\vec{x}) \leq 0\} \quad (2)$$

Where $\vec{x}$ is a vector of decision variables $(x_1, \ldots, x_n)$ and $F(\vec{x} \in \Upsilon)$ is a vector of objective functions $(f_1(\vec{x} \in \Upsilon), \ldots, f_K(\vec{x} \in \Upsilon))$. Here $f_1(\vec{x} \in \Upsilon), \ldots, f_K(\vec{x} \in \Upsilon)$, are functions on R^n and Υ is a nonempty set in R^n. The vector $G(\vec{x})$ represents a set of constraints.

The aim is to find the vector $\vec{x}^* \in \Upsilon$ which optimizes $F(\vec{x} \in \Upsilon)$. Without any loss of generality, we assume that all objectives are to be minimized. We note that any maximization problem can be transformed to a minimization one by multiplying the former by -1.

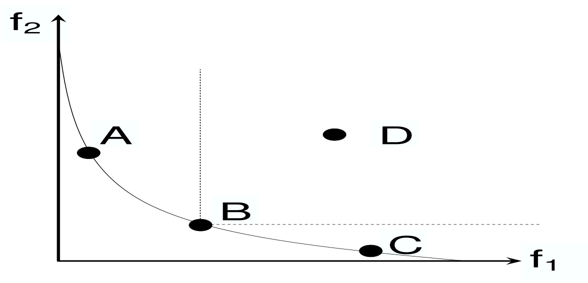

Figure 1: The concept of dominance in multi-objective optimization. Assuming that both f_1 and f_2 are to be minimized, D is dominated by B since B is better than D when measured on all objectives. However, A, B and C are non–dominated since none of them is better than the other two when measured on all objectives.

The principle of dominance (Figure 1) in *multi-objective optimization problem* (MOP) allows a partial order relation that works as follows: a solution does not have an advantage to be included in the set of optimal solutions unless there is no solution that is better than the former when measured on all objectives. A non–dominated solution is called Pareto. A MOP can be solved in different ways. Evolutionary algorithms (EAs) [8, 10], being population based, they are able to generate a set of near-Pareto solutions in a single run. In addition, they do not require assumptions of convexity, differentiability, and/or continuity as traditional optimization problems do. EAs with local search are usually used to improve the performance of EAs to get closer to the actual optimal or, the Pareto set in the case of MOPs.

3 The Pareto Operating Curve

In many situations, the Pareto curve can be seen as the single solution to the problem. Take for example a problem where there is a need to evolve controllers for a robot. The objective functions can potentially be to minimize energy consumption and minimize the robot's performance error. In this case, a solution on the Pareto curve for this problem is just one possible trade-off that can be made between the previous two objectives. However, this robot is likely to encounter a number of situations where it needs to trade-off differently between these two objectives over time. As such, the Pareto Curve can be seen as an Operating Curve, as the level of trade-off needed changes over time, a solution moves from one location to another on that curve (See Figure 2).

Definition 3.1 Pareto Operating Curve *A Pareto Operating Curve (POC) is a Pareto Curve for a problem where the trade-off between the objectives to be optimized varies over time; thus a solution selected along this curve at one point of time needs to move to a different solution at another point of time to minimize the impact of uncertainty on objectives (i.e. risk).*

We need to differentiate between adaptive feedback control with the concept of Pareto Operating Curve. In traditional adaptive feedback control, a controller adjusts its parameters in response to changes in the environment. The Pareto Operating Curve provides the most efficient set of models to be operated in different environments to minimize the risk. Each member in this set is optimal in a particular environment in the sense that each environment represents a specific level of trade-off and there is a solution in the efficient set which is optimal on that required level of trade-off. One can then imagine the existence of a switch or a decision maker that senses the environment, determines the optimal level of trade-off needed, then selects the corresponding non-dominated solution from the

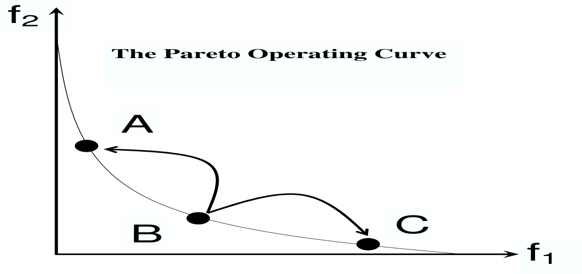

Figure 2: The Pareto Operating Curve. Solution B can be the best solution for a specific level of trade-off between objective function f_1 and f_2 at time t. When the required level of trade-off changes, this solution may need to move along the curve to become, for example, solution A or C.

efficient set. Each time a solution is selected from the non-dominated set, it defines a movement on the Pareto curve. This movement may be constrained in terms of its cost or characteristics, thus bounds the impact of the risk mitigation strategy. We now provide examples where this concept is successfully demonstrated.

4 The Pareto Operating Curve and Evolution

The majority of research in decision making and engineering has focused on selecting a single solution. Recent research showed the benefits in viewing problems in the eyes of multi-objective search. For example, in single objective optimization, one can simply benefit from transforming it into multi-objective as being demonstrated in [6].
To discuss the concept of POC in an artificial life context, it would be less attractive to do so without discussing its biological roots and impact. Although this is the first time this concept is introduced in this paper, we can trace some seeds for this concept in the literature. Darwin wrote:

> It seems clear that organic beings must be exposed during several generations to new conditions to cause any great amount of variation; and that, when the organisation has once begun to vary, it generally continues varying for many generations" (from [9] P25).

What is interesting about Darwin's quote is the emphasis he placed on variations. According to Darwin, any organization is in a constant state of flux. But we know from

common sense and decision sciences that each state of flux is likely to require different levels of trade-offs. The Pareto curve represents the optimal set of solutions in the sense that for any level of trade-off required, there is a solution in that set that is the optimal solution for the single objective optimization problem derived from the utilities associated with the required level of trade-off. Therefore, evolution does not necessarily need to be an optimizer, but for evolution to work, it needs to maintain diversity along the Pareto curve. In so doing, evolution can move from one trade-off to another. Evolution is not an optimizer from traditional optimization point of view, while from multi-objective optimization point of view, I will make the assertion that evolution is a multi-objective optimizer. In fact, I would claim that this is the evolutionary strategy for risk mitigation. The objectives that evolution optimizes include for example adaptive capacity, robustness, and survivability. In a number of situations, such as in viruses where the level of unpredictability in the change in the environment is high, we should accept that random selection can be an efficient strategy for risk minimization in such environment. Once the signal to noise ratio is high, random selection fails as a strategy and other types of selection mechanisms become more appropriate.
Not so long after Darwin's writing, Pareto wrote:

> If, as has generally been the case, it is held that, for a people, utility is coterminous with its material prosperity and its moral and intellectual development, then we have a criterion for making comparisons between different people. But there still remains a difficulty, deriving from the fact that society has to be considered as a complex whole, as a system, as an organism. [12].

Well said before its time, Pareto pointed us to the right direction, that an organism is a system of systems (SoS), evolution is, the mind is, and society is. As such, each sub-system (which is a system in its own right) has its own utilities which can be in conflict. For example, material prosperity can be in conflict with intellectual development. These competing objectives on the sub-system level, along with the different levels of trade-off possessed by each subsystem (representing their own individual biases) generate the diversity required for the system as a whole to operate and function.

5 The Pareto Operating Curve and Complexity

In [14, 15], we have shown for the first time the relationship between Pareto and Complexity. The essence of this work is that complexity of species is not a single measure. Moreover, combining many measures of complexity using some index into a single measure, not only hides information because of the strict order bias generated by a weighted sum approach, it also violates the essence of what complexity is. Thus, complexity should be defined as a strict partial order rather than a linear strict order. In these papers, we introduced the following definition of complexity:

- **Complexity** is a strict partial order relation.

This definition moves away from what the majority of literature in Engineering attempts to do; that is, to come up with a quantitative measure (single number) of complexity to establish a linear rank. This single number hides information of its constituent parts and the level of trade-offs required on the sub-system level. It also assumes that one must unify dimensionality and scale before combining the different complexity measures. Pareto's view to complexity, however, accepts the existence of many different quantitative measures of complexity but it rejects the idea of combining them as a single measure. Pareto optimality does not satisfy reflexivity; that is, a solution cannot dominate itself. It also acts as a filter of these measures since a measure is redundant if it is not in conflict (i.e. it produces identical order) with an existing measure. Pareto optimality, thus, imposes a complexity hierarchy on the set of objects/solutions.

6 The Pareto Operating Curve and Robotics

The use of Pareto-based Evolutionary Multi-objective (EMO) Search techniques in computational intelligence - particularly fuzzy inferencing and neural networks, is a relatively new literature. The work on Pareto-based EMO for fuzzy inferencing was pioneered in a number of papers, particularly [11], while the work on Pareto-based EMO for neural networks was pioneered in [1, 2, 5]. Work on neuro-ensemble was then introduced in a number of papers including [3, 4].
Traditionally, one would search for a learning machine - such as a neural network - that performs well on the average on all environmental conditions it may encounter. However, there are many applications where this average performance is not acceptable. For example, imagine a walk-gate performed by a neural network. Imagine that we want the robot to walk in different environments. Here, we can use the concept of Pareto optimality to optimize along different environmental conditions. The objective functions represent the robot performance in different environmental conditions. For example, one objective can represent the robot's speed while the second represents the friction in the

terrain. The set of non-dominated solutions for this problem represents a trade-off between performance (speed) and complexity (nature of the terrain). Every solution in this set represents the optimal solution on the corresponding terrain's friction. As such, the whole set can be used within a robot with a switch attached to a sensor that senses the friction between the robot's leg and the terrain. One can then imagine that at any particular point of time, the robot is operating in one area of the Pareto curve and as the environmental conditions change, it moves to other areas.
The previous concept was used in [13], where the two objectives were distance travelled and the size of the controller. The resultant robots trade-off, the size of the controller and the Pareto curve clearly demonstrated a smooth transition from no-walking behavior to a robot that jumps.
Another application of this concept was in the area of Air Traffic Management [7]. Different algorithms for conflict detection work better in specific environments. Once more, one can imagine the Pareto Operating Curve as the set of environmental conditions where an algorithm would fail. By combining these conditions using a switch/gate, one would minimize the overall failure of the aircraft detection mechanism and the risk associated with that by combining different detection algorithms.

7 Conclusion

In this paper, I introduced the concept of Pareto Operating Curve, whereby the decision making process is seen as a set of movements along the curve to minimize risk. The roots of this concept were traced in evolution, and its relationship with complexity and robotics were discussed. As a new concept, the doors are open to adopt it to many applications including data mining, robotics or decision theory.

Acknowledgement

This work is supported by an Australian Research Council (ARC) discovery grant No. DP0667123.

References

[1] H.A. Abbass. A memetic Pareto evolutionary approach to artificial neural networks. In M. Stumptner, D. Corbett, and M. Brooks, editors, *AI'01, LNAI 2256*, pages 1–12. Springer-Verlag, Berlin, 2001.

[2] H.A. Abbass. An evolutionary artificial neural network approach for breast cancer diagnosis. *Artificial Intelligence in Medicine*, 25(3):265–281, 2002.

[3] H.A. Abbass. Pareto neuro-ensemble. In *AI'03, LNAI 2903*, pages 554–566. Springer-Verlag, Berlin, 2003.

[4] H.A. Abbass. Pareto neuro-evolution: Constructing ensemble of neural networks using multi-objective optimization. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC2003)*, number 3, pages 2074–2080. IEEE-Press, 2003.

[5] H.A. Abbass. Speeding up back-propagation using multiobjective evolutionary algorithms. *Neural Computation*, 15(11):2705–2726, 2003.

[6] Hussein A. Abbass and Kalyanmoy Deb. Searching under multi-evolutionary pressures. In C.M. Fonseca et. al., editor, *EMO'03, LNCS 2632*, pages 391–404. Springer-Verlag, Berlin, 2003.

[7] S. Alam, K. Shafi, H.A. Abbass, and M. Barlow. An ensemble approach for conflict detection in free flight by data mining. *Transportation Research Part C*, In Press.

[8] C.A. Coello Coello, D.A. Van Veldhuizen, and G.B. Lamont. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Kluwer Academic, New York, 2002.

[9] C.H. Darwin. *The origins of species by means of natural selection*. London, Penguin Classics, 1859.

[10] K. Deb. *Multi-objective Optimization using Evolutionary Algorithms*. John Wiley & Sons, Chicester, UK, 2001.

[11] H. Ishibuchi, T. Murata, and IB Türkşen. Single-objective and two-objective genetic algorithms for selecting linguistic rules for pattern classification problems. *Fuzzy Sets and Systems*, 89(2):135–150, 1997.

[12] V. Pareto. *Cours d'ėconomie politique professė à l'universitė de Lausanne*, volume 1,2. F. Rouge, Laussanne, 1896.

[13] J. Teo and H.A. Abbass. Automatic generation of controllers for embodied legged organisms: A Pareto evolutionary multi-objective approach. *Evolutionary Computation*, 12(3):355–394, 2004.

[14] J. Teo and H.A. Abbass. Multi-objectivity and complexity in embodied cognition. *IEEE Transactions on Evolutionary Computation*, 9(4):337–360, 2005.

[15] Jason Teo, Minh Ha Nguyen, and Hussein A. Abbass. Multi-objectivity as a tool for constructing hierarchical complexity. In E. Cantu-Paz et. al., editor, *GECCO'03, LNCS 2723*, pages 483–494. Springer-Verlag, Berlin, 2003.