Convergence Detection and Stopping Criteria for Evolutionary Multi-Objective Optimization

Luis Martí[1,2,3]   Nayat Sanchez-Pi[1,3]

1 Group of Research on Intelligence and Optimization (RIO Group)

2 Institute of Computing,
Universidade Federal Fluminense (UFF).

3 TAU Team, INRIA/Saclay.

4 Institute of Mathematics and Statistics,
State University of Rio de Janeiro (UERJ).

2017 IEEE Congress on Evolutionary Computation
Donostia/San Sebastián, Spain. June 5–9, 2017.

Talk outline

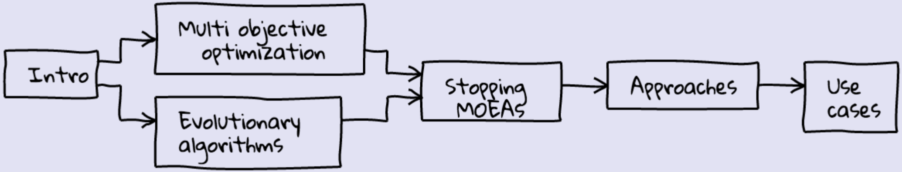

Do you know them?

Rudolph Diesel

Fred Duesenberg

Internal combustion ‘diesel’ engine

Hydraulic brakes

It seems that stopping things does not makes you famous.

What does “until end condition is met” really means?

Multi-objective optimization

Most *-if not all-* optimization problems involve more than one objective function to be optimized simultaneously.

- For example: optimize a given feature of an object while keeping under control the resources needed to elaborate that object.
- Sometimes those other objectives are converted to constraints or fixed to default values, but they are still there.
- Multi-objective optimization is also known as *multi-objective programming*, *vector optimization*, *multicriteria optimization*, *multiattribute optimization* or *Pareto optimization* (and probably by other names, depending on the field).

The multi-objective 'fruit selection problem'

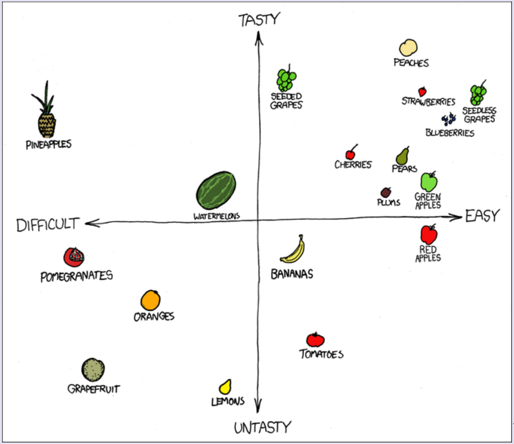

[1]

We do it all the time!

[1] from http://xkcd.com/388

Multi-objective optimization problem

$$\text{minimize } \boldsymbol{F}(\boldsymbol{x}) = \langle f_1(\boldsymbol{x}), \ldots, f_M(\boldsymbol{x}) \rangle \,,$$
$$\text{with } \boldsymbol{x} \in \mathcal{D} \,.$$

- $\mathcal{D}$: feasible set — can be defined as constraints;
- $\mathcal{O}$: objective set;
- optimality — *Pareto dominance*;
- $\mathcal{D}^*$: Pareto-optimal set;
- $\mathcal{O}^*$: Pareto-optimal front, and;
- $\mathcal{P}^*$: optimizer solution.

A decision maker selects elements of $\mathcal{P}_t^*$.

Pareto dominance relation

Usually, there is not a unique solution that minimizes all objective functions simultaneously, but, instead, a set of equally good *trade-off* solutions.

- Optimality can be defined in terms of the Pareto dominance relation.
- That is, having $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{D}$, **$\boldsymbol{x}$ is said to dominate $\boldsymbol{y}$** (expressed as $\boldsymbol{x} \prec \boldsymbol{y}$) iff $\forall f_j$, $f_j(\boldsymbol{x}) \leq f_j(\boldsymbol{y})$ and $\exists f_i$ such that $f_i(\boldsymbol{x}) < f_i(\boldsymbol{y})$.
- Having the set $\mathcal{A}$. $\mathcal{A}^*$, the *non-dominated subset* of $\mathcal{A}$, is defined as

$$\mathcal{A}^* = \left\{ \boldsymbol{x} \in \mathcal{A} \,\middle|\, \nexists \boldsymbol{y} \in \mathcal{A} : \boldsymbol{y} \prec \boldsymbol{x} \right\}.$$

- The *Pareto-optimal set*, $\mathcal{D}^*$, is the solution of the MOP. It is the subset of non-dominated elements of $\mathcal{D}$.
- Its image in objective set is called the *Pareto-optimal front*, $\mathcal{O}^*$.

Example: Dent problem

$$
\begin{aligned}
\text{minimize} \quad & f_1(\boldsymbol{x}), f_2(\boldsymbol{x}) \\
\text{such that} \quad & f_1(\boldsymbol{x}) = \frac{1}{2}\left(\sqrt{1+(x_1+x_2)^2}\sqrt{1+(x_1-x_2)^2} + x_1 - x_2\right) + d, \\
& f_2(\boldsymbol{x}) = \frac{1}{2}\left(\sqrt{1+(x_1+x_2)^2}\sqrt{1+(x_1-x_2)^2} - x_1 - x_2\right) + d, \\
\text{with} \quad & d = \lambda e^{-(x_1-x_2)^2}. \\
& \text{generally } \lambda = 0.85 \text{ and } \boldsymbol{x} \in [-1.5, 1.5]^2 .
\end{aligned}
$$

Example: Dent problem — Plots

Decision set

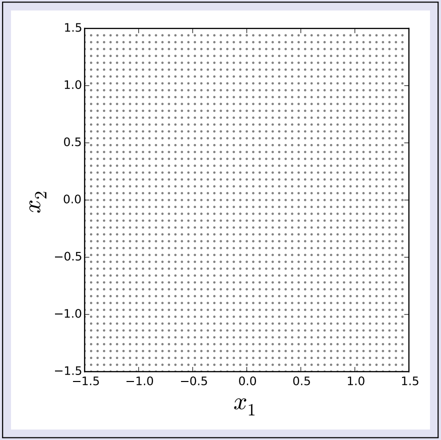

Objective set

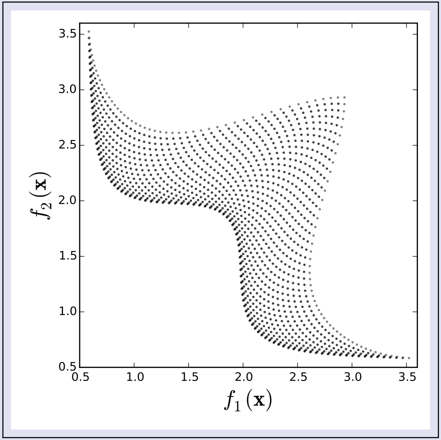

Example: Dent problem — the dominance relation

Decision set

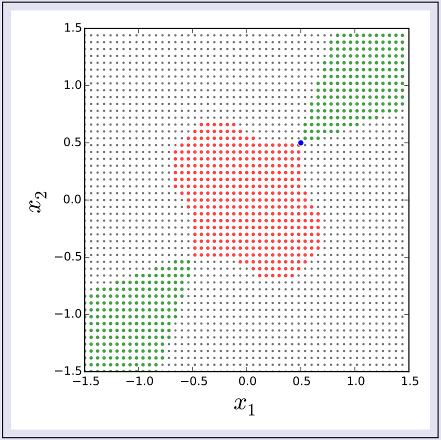

Objective set

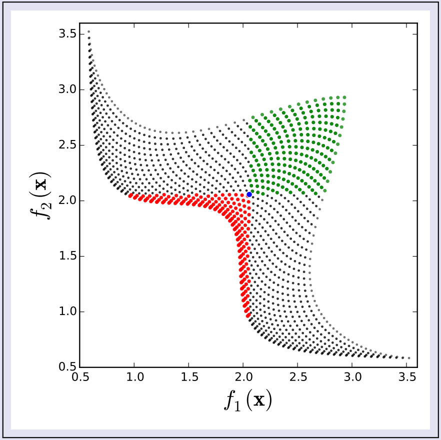

Example: Dent problem — non-dominated front

Decision set

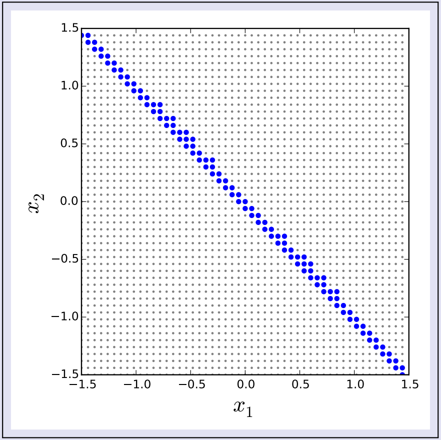

Objective set

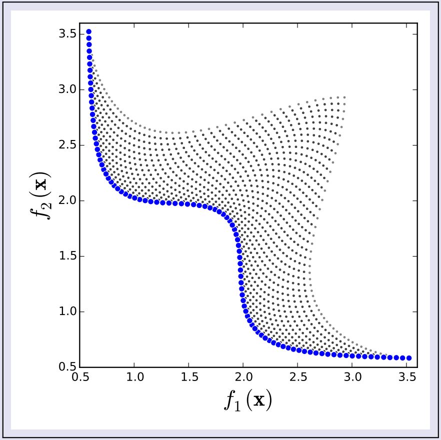

Performance indicators

How can we compare different (sets of) solutions?

- Hypervolume indicator;
- additive/multiplicative epsilon indicator;
- R1/R2 indicators;
- inverted generational distance, etc.

The hypervolume indicator

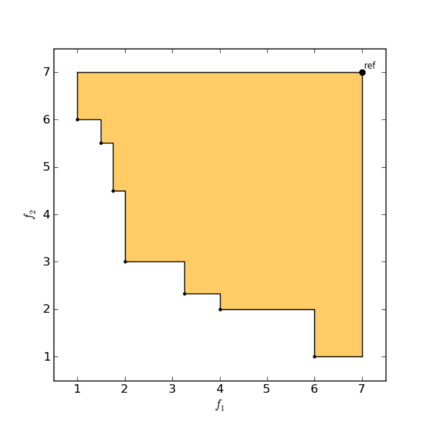

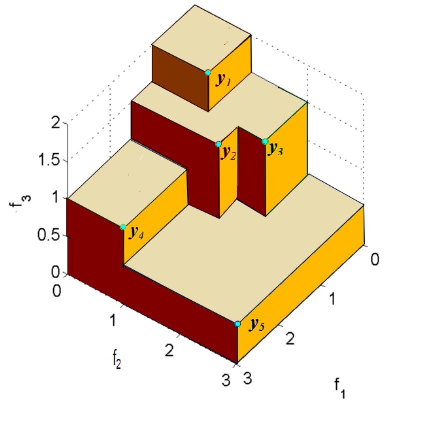

[2]

[2] From Günter Rudolph's site: `https://ls11-www.cs.uni-dortmund.de/rudolph/hypervolume/start`.

Formalization of the hypervolume

For a set of solutions $\mathcal{A}$,

$$I_{\text{hyp}}\left(\mathcal{A}\right) = \text{volume}\left(\bigcup_{\forall \boldsymbol{a} \in \mathcal{A}} \text{hypercube}(\boldsymbol{a}, \boldsymbol{r})\right).$$

- We need a *reference point*, $\boldsymbol{r}$.
- Hypervolume is Pareto compliant (Fleischer, 2003): for sets $\mathcal{A}$ and $\mathcal{B}$, $\mathcal{A} \prec \mathcal{B} \implies I_{\text{hyp}}(A) > I_{\text{hyp}}(B)$.
- Calculating hypervolume is **NP**-hard.

Evolutionary algorithms

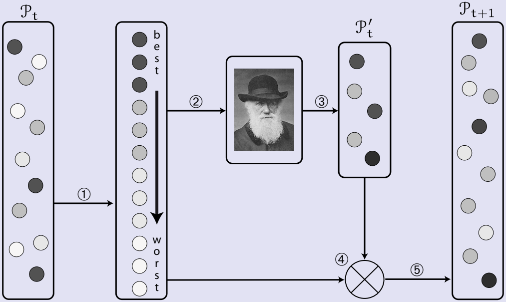

- A **population** of **individuals**;
- individuals are ranked using a **fitness assignment function**;
- **evolution-inspired** operators are applied;
- fittest individuals have a more active role.

Multi-objective evolutionary algorithms (MOEAs)

- One of the hottest topics in EA research.
- MOPs put EAs "to the limit".
- Succeeded in yielding relevant results.
- Do not make any assumptions about the problem.
- Parallel nature of the search process produces sets of solutions.

Cornerstone issue: fitness assignment.

Ranking a multi-objective population

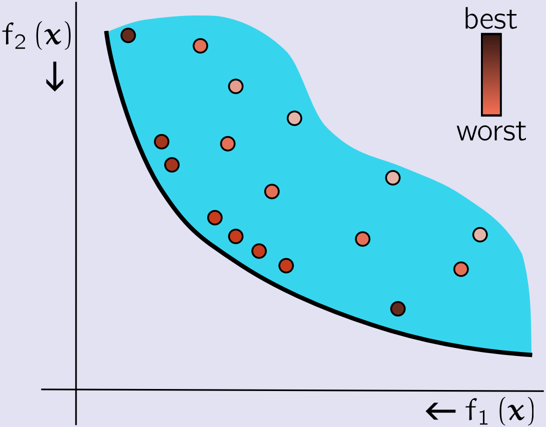

- Objective aggregation.
- Pareto-based ranking.
- Performance indicator-based ranking.

Many-objective problems

Problems with four or more objectives.

Challenges

- Visualization.
- Poor understanding of convergence and progress $\rightarrow$ stopping criteria.
- Scalability.

Scalability

- Exponential relation between the number of objectives and the amount of resources.
- Large populations are needed.

Stopping a multi–objective optimization

Stopping criteria

- detect when there is no sense in proceeding with the search;
- they are usually a heuristic.
- This matter have been overlooked in the EMO context, but;
- complex and real–world applications demand them.

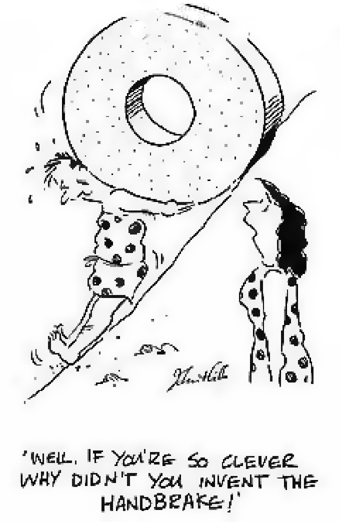

[3]

[3] from https://www.cartoonstock.com/cartoonview.asp?catref=jhin64.

Stopping criteria in general

A stopping criterion is invoked at the end of an iteration of the algorithm being controlled.

Scenarios

1. the solution obtained so far is satisfactory;
2. we have a feasible solution that is not satisfactory in terms of optimality, but a better one is unlikely to be produced;
3. the method is unable to converge to any solution, or;
4. the amount of computation performed so far is *sufficient*.

In brief…

A stopping criterion should detect "**success**" and "**failure**" scenarios.

A multi-objective optimization stopping criteria

- Judging the advance of the optimization can become as complex as the optimization itself;
- unlike other problems there is no “axis” to be used as reference;
- therefore any assessment must be carried out in a relative fashion, but;
- current performance indicators have a high computational cost.

Desirable properties

- An execution-wise criterion is required because of the nature of the problem;
- Resource requirements should be kept as low as possible, in particular;
- the criterion should be embedded in other processes.
- As few parameters as possible!

Karush-Kuhn-Tucker (KKT) optimality

If we encounter problems, where the feasible set is given implicitly by constraints,

$$\mathcal{S} = \{\boldsymbol{x} \in \mathbb{R}^n; (c_1(\boldsymbol{x}), \dots, c_C(\boldsymbol{x})) \leq \boldsymbol{0}\} \ ,$$

and $f_1, \dots, f_M$ and constraint functions $c_1, \dots, c_C$ are continuously differentiable.

Definition (Karush-Kuhn-Tucker proper optimality condition[4][5])

A solution $\boldsymbol{x} \in \mathcal{S}$ is said to be properly Pareto-optimal if it holds the Pareto-optimality condition and $\nexists \boldsymbol{b} \in \mathbb{R}^n$ such that

$$\forall i = 1, \dots, M : \quad \nabla f_i(\boldsymbol{x})^{\mathrm{T}} \boldsymbol{b} \leq \boldsymbol{0}\,; \tag{1}$$

$$\exists j = 1, \dots, M \text{ such that} : \quad \nabla f_j(\boldsymbol{x})^{\mathrm{T}} \boldsymbol{b} < \boldsymbol{0}\,; \tag{2}$$

$$\forall c_c(\boldsymbol{x}) = 0 : \quad \nabla c_c(\boldsymbol{x})^{\mathrm{T}} \boldsymbol{b} \leq \boldsymbol{0}\,. \tag{3}$$

[4] Karush, W. (1939). Minima of functions of several variables with inequalities as side constraints. Master's thesis, Department of Mathematics, University of Chicago, Chicago, IL, USA

[5] Kuhn, H. W. and Tucker, A. W. (1951). Nonlinear programming. In Neyman, J., editor, *Proceedings of Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 481–492, Berkeley, LA, USA. University of California Press

…from multi-criteria decision making?

1. initialize: calculate ideal and approximated nadir objective vectors and show them to the decision maker.
2. generate a Pareto optimal starting point (by using e.g. some no-preference method or solution given by the decision maker)
3. ask for preference information from the decision maker (e.g. aspiration levels or number of new solutions to be generated)
4. generate new Pareto optimal solution(s) according to the preferences and show it/them and possibly some other information about the problem to the decision maker
5. if several solutions were generated, ask the decision maker to select the best solution so far
6. stop, if the decision maker wants to; otherwise, go to step 3). $\leftarrow$

They even have a name for that: **psychological convergence**[6]

[6] Stewart, T. J. (1997). Convergence and validation of interactive methods in mcdm: simulation studies. In *Essays in Decision Making*, pages 7–18. Springer

Where can we look for ideas?

"Classical" multi–criteria decision making

- We have proper definitions of optimality, like the Kuhn–Tucker condition, but;
- perhaps they are more relevant from a theoretical point of view.
- *Ideal solution.*

Single–objective EAs

- Search space exploration[7]: all states visited with a certain probability;
- objective convergence[8];
- population convergence.

[7] Aytug, H. and Koehler, G. J. (2000). New stopping criterion for genetic algorithms. *European Journal of Operational Research*, 126(3):662–674

[8] Jain, B. J., Pohlheim, H., and Wegener, J. (2001). On termination criteria of evolutionary algorithms. In Spector, L., Goodman, E. D., Wu, A., Langdon, W., Voigt, H.-M., Gen, M., Sen, S., Dorigo, M., Pezeshk, S., Garzon, M. H., and Burke, E., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2001)*, page 768, San Francisco, California, USA. Morgan Kaufmann

Guidelines

Desirable properties

- An execution–wise criterion is required because of the nature of the problem;
- resource requirements should be kept as low as possible, in particular;
- the criterion should be embedded in (and profit from) other processes.
- as few parameters as possible!

At least two components

- A local **progress indicator**, and;
- An **evidence gathering process** that combines the local measurements.

Situation

- There have been few theoretical works[9] that deal with EMO convergence, and;
- there has been even more sparse attempts to deal with the stopping issue.

The importance of this matter has not been correctly underscored until recently.[10]

[9] Rudolph, G. and Agapie, A. (2000). Convergence properties of some multi–objective evolutionary algorithms. In *2000 IEEE Congress on Evolutionary Computation (CEC 2000)*, volume 2, pages 1010–1016, Piscataway, New Jersey. IEEE Press

[10] Fonseca, C., Gandibleux, X., Korhonen, P., Martí, L., Naujoks, B., Thiele, L., Wallenius, J., and Zitzler, E. (2009). Working group on EMO for interactive multiobjective optimization (1st round). In Deb, K., Greco, S., Miettinen, K., and Zitzler, E., editors, *Hybrid and Robust Approaches to Multiobjective Optimization*, number 09041 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany. Schloss Dagstuhl — Leibniz–Zentrum fuer Informatik

Approaches in chronological order

Deb and Jain: Running Metrics[11]

- Deb and Jain were the first authors who proposed the investigation of performance metrics over the run of MOEAs.
- They used two metrics, one for evaluating the convergence and one for measuring the diversity of the Pareto front.
- Convergence metric (CM) calculates the average of the smallest normalized euclidean distance from each individual in the Pareto front to a precomputed reference set.
- For the diversity metric (DVM), all objective value vectors of the Pareto front are projected onto a $m-1$–dimensional hyperplane which is then uniformly divided into discrete grid cells.
- The DVM tracks the number of attained grid cells and also evaluates the distribution by assigning different scores for predefined neighborhood patterns.

[11] Deb, K., Jain, S.: Running performance metrics for evolutionary multi-objective optimization. In: Simulated Evolution and Learning (SEAL), pp. 13–20 (2002) — 2002

Approaches in chronological order

Rudenko and Schoenauer[12] — 2004

- A stopping criterion to be used in conjunction with the NSGA–II algorithm.
- Measure the mean of the spread of the non–dominated individuals.
- Compute an average across iterations of the measurements.

Stoooooop! We found a leitmotiv!

To stop when the “indicator” is **zero** and has a **flat** tendency.

[12] Rudenko, O. and Schoenauer, M. (2004). A steady performance stopping criterion for pareto–based evolutionary algorithms. In *The 6th International Multi–Objective Programming and Goal Programming Conference*, Hammamet, Tunisia

Approaches in chronological order (II)

Martí, García, Berlanga and Molina (MGBM)[13][14] — 2007, 2009

- Iteration–wise measurements with the *mutual domination rate* (MDR) indicator, and;
- a simplified Kalman filter tracks the measurements in an *evidence gathering process.*
- It detects when MDR is close to zero and flat.

- The KF helps to pay less attention to measurements at beginning.
- MDR can be embedded in the EMO supervised.
- Shown to detect “failure” situations.

[13] Martí, L., García, J., Berlanga, A., and Molina, J. M. (2007). A cumulative evidential stopping criterion for multiobjective optimization evolutionary algorithms. In Thierens, D., Deb, K., Pelikan, M., Beyer, H.-G., Doerr, B., Poli, R., and Bittari, M., editors, *GECCO'07: Proceedings of the 2007 GECCO Conference Companion on Genetic and Evolutionary Computation*, page 911, New York. ACM Press

[14] Martí, L., García, J., Berlanga, A., and Molina, J. M. (2009). An approach to stopping criteria for multi–objective optimization evolutionary algorithms: The MGBM criterion. In *2009 IEEE Conference on Evolutionary Computation (CEC 2009)*, pages 1263–1270, Piscataway, New Jersey. IEEE Press

Mutual domination rate

Having…

- Non-dominated individuals of consecutive iterations: $\mathcal{P}^*_t$ and $\mathcal{P}^*_{t-1}$;
- $\Delta\left(\mathcal{A},\mathcal{B}\right)$: elements of $\mathcal{A}$ that are dominated by at least one element of $\mathcal{B}$;

Mutual domination rate indicator

- Contrasts how many non-dominated individuals of iteration t dominate the non-dominated individuals of the previous one $(t-1)$ and vice versa.
- Formally,

$$I_{\mathrm{mdr}}\left(\mathcal{P}^*_t,\mathcal{P}^*_{t-1}\right)=\frac{\left|\Delta\left(\mathcal{P}^*_{t-1},\mathcal{P}^*_t\right)\right|}{\left|\mathcal{P}^*_{t-1}\right|}-\frac{\left|\Delta\left(\mathcal{P}^*_t,\mathcal{P}^*_{t-1}\right)\right|}{\left|\mathcal{P}^*_t\right|}\,;$$

It is easy to embed MDR in Pareto-based fitness assignment computations.

Accumulating evidence via Kalman filters

Kalman filter

- Estimates the state of a discrete-time controlled process that is ruled by a *linear* stochastic difference equation;
- an efficient computational means to estimate the state of a dynamic system from a series of incomplete and noisy measurements.

For the stopping purpose...

- It is tracked the MDR indicator at iteration t, $I_{\text{mdr}}(t)$.
- The filter is set to predict that the indicator will remain the constant.
- The filter status is updated with the measurements of the indicator.

Stop when estimation and confidence interval bellow threshold.

MGBM at work

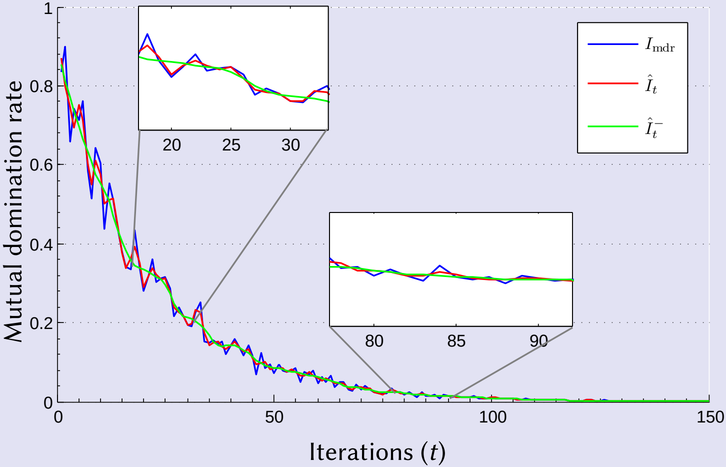

Evolution of the MDR indicator, $I_{\mathrm{mdr}}\left(\mathcal{P}_t^*, \mathcal{P}_{t-1}^*\right)$, and the *a priori* and *a posteriori* estimations $\hat{I}_t^-$ and $\hat{I}_t$ across iterations. Here the NSGA–II algorithm is supervised as it successfully solves the DTLZ3 problem.

Approaches in chronological order (II)

Offline convergence detection (Trautmann et al.[15]) — 2008

- A series of statistical hypothesis tests performed on concurrent algorithm executions.
- Different performance indicators are computed.
- This is information is used to decide when the best solutions of the algorithm was obtained.
- Sound and robust, but;
- resource demanding.

[15] Trautmann, H., Ligges, U., Mehnen, J., and Preuss, M. (2008). A convergence criterion for multiobjective evolutionary algorithms based on systematic statistical testing. In *Parallel Problem Solving from Nature — PPSN X*, pages 825–836. Springer, Heidelberg

Approaches in (...) order (III): The 2009 boom

Online convergence detection (Wagner et al.[16]) — 2009

- Tracks a number of progress indicators;
- stops when the variance is below a threshold, and;
- passes a linear trend statistical test.
- Less resource demanding than the previous one.

Dominance–based stability measure (Bui et al.[17]) — 2009

- Dominance–based quality indicator.
- Determines how many solutions in a given radius dominate current solutions.
- Only local measurements.
- The radius is hard to set a priori.

---

16 Wagner, T., Trautmann, H., and Naujoks, B. (2009). OCD: Online convergence detection for evolutionary multi–objective algorithms based on statistical testing. In Ehrgott, M., Fonseca, C. M., Gandibleux, X., Hao, J.-K., and Sevaux, M., editors, *5th International Conference on Evolutionary Multi–Criterion Optimization (EMO 2009)*, volume 5467 of *Lecture Notes in Computer Science*, pages 198–215. Springer

17 Bui, L. T., Wesolkowski, S., Bender, A., Abbass, H., and Barlow, M. (2009). A dominance–based stability measure for multi–objective evolutionary algorithms. In *2009 IEEE Congress on Evolutionary Computation (CEC '09)*, pages 749–756, Piscataway, New Jersey. IEEE Press

Approaches in (...) order (IV): The 2009 boom

Guerrero et al.[18] — 2009

- Improved formulation of filters by adding adaptation;
- combines information from progress indicators, and;
- voting is used to decide when to stop.
- More robust than MGBM.

[18] Guerrero, J. L., García, J., Martí, L., Molina, J. M., and Berlanga, A. (2009). A stopping criterion based on Kalman estimation techniques with several progress indicators. In Raidl, G., Alba, E., Bacardit, J., Bates Congdon, C., Beyer, H.-G., Birattari, M., Blum, C., Bosman, P. A. N., Corne, D., Cotta, C., Di Penta, M., Doerr, B., Drechsler, R., Ebner, M., Grahl, J., Jansen, T., Knowles, J., Lenaerts, T., Middendorf, M., Miller, J. F., O'Neill, M., Poli, R., Squillero, G., Stanley, K., Stützle, T., and van Hemert, J., editors, *GECCO'09: 11th Annual Conference on Genetic and Evolutionary Computation*, pages 587–594, New York, NY, USA. ACM Press

Recent developments

Goel and Stander[19] — 2010

- An indicator that measures how stable are solutions in an archive that contains the best solutions obtained.
- Tied to an specific version of NSGA–II.

Least square stopping criterion (Guerrero et al.[20]) — 2010

- Test if a sample of the indicator value can be modelled by a linear regression.
- Instead of using the variance (as in OCD) the slope is used.
- The value of the slope is independent of the scale of the indicators.

[19] Goel, T. and Stander, N. (2010). A non-dominance-based online stopping criterion for multi–objective evolutionary algorithms. *International Journal for Numerical Methods in Engineering*

[20] Guerrero, J. L., Martí, L., García, J., Berlanga, A., and Molina, J. M. (2010). Introducing a robust and efficient stopping criterion for MOEAs. In *2010 IEEE Conference on Evolutionary Computation (CEC), part of 2010 IEEE World Congress on Computational Intelligence (WCCI 2010)*, Piscataway, New Jersey. IEEE Press

Summary

Custom indicators:

- mutual domination rate;
- dominance–based quality indicator, and;
- Goel and Stander.

Evidence gathering:

- online convergence detection;
- Kalman filters;
- combination of filters, and;
- least square stopping criterion.

Salient issues: Local progress

Are performance indicators suitable?

- Applying them measure the improvement of the solutions, but;
- in "plateau" situations when evolution could be progressing at a slow pace?
- Dominance–based methods are not suitable for the recent indicator–based EMOs.

What are the characteristics of stagnated EMOs?

What if we analyze the "health" of the evolutionary process?

Fitness homogeneity

- It has been shown that a balance between dominated and non–dominated solutions is needed.
- If most of the population is non–dominated the search process becomes stagnated.

We have been working in this matter[21]

- Simple approach that computes the deviation of the fitness values;
- needs a transformation scheme, but;
- can be used with indicator–based EMOs.

[21] Martí, L., García, J., Berlanga, A., and Molina, J. M. (2010). A progress indicator for detecting success and failure in evolutionary multi–objective optimization. In *2010 IEEE Conference on Evolutionary Computation (CEC), part of 2010 IEEE World Congress on Computational Intelligence (WCCI 2010)*, Piscataway, New Jersey. IEEE Press

Salient issues: Comparing criteria

- Testing stopping criteria is an awkward task, but;
- comparing them is even more.

We need to come up with a set of problem+algorithm pairs.

- Different convergence rates;
- incorrect parameter setup, in particular;
- population sizes.

Development areas

- Testing stopping criteria is an awkward task, but;
- comparing them is even more.

We need to come up with a set of problem+algorithm pairs.

- Different convergence rates;
- incorrect parameter setup, in particular;
- population sizes.

Available software

- Matlab taxonomomy of stopping criteria:
  `http://github.com/lmarti/emo-stopping-criteria-taxonomy`.
- Python implementation of current state of the art:
  `http://github.com/lmarti/py-emostop`.
  - Can be used as in conjunction with the `DEAP` and `inspyred` modules.

Final remarks

- There is room for improvement;
- **indicators**: assume elitism, take into account diversity, etc.
- **evidence gathering**: many tools available: statistics, Markov, filters, etc.
- Approximate KKT points as proximity measure[22] (very interesting).
- Ideas from anomaly detection, time-series processing and/or outlier detection.
- *A posteriori* analysis.
- The most recent review on the topic is still valid[23].

[22] Dutta, J., Deb, K., Tulshyan, R., and Arora, R. (2013). Approximate KKT points and a proximity measure for termination. *Journal of Global Optimization*, 56(4):1463–1499

[23] Wagner, T., Trautmann, H., and Martí, L. (2011). A taxonomy of online stopping criteria for multi-objective evolutionary algorithms. In Takahashi, R. H. C., Deb, K., Wanner, E. F., and Greco, S., editors, *6th International Conference on Evolutionary Multi-Criterion Optimization (EMO 2011)*, volume 6576, pages 16–30, Berlin/Heidelberg. Springer

Eskerrik asko!

Thank you!

Obrigado!

¡Gracias!

- lmarti@ic.uff.br; nayat@ime.uerj.br
- http://lmarti.com/stopping

Bibliography I

Aytug, H. and Koehler, G. J. (2000).
New stopping criterion for genetic algorithms.
European Journal of Operational Research, 126(3):662–674.

Bui, L. T., Wesolkowski, S., Bender, A., Abbass, H., and Barlow, M. (2009).
A dominance–based stability measure for multi–objective evolutionary algorithms.
In *2009 IEEE Congress on Evolutionary Computation (CEC '09)*, pages 749–756, Piscataway, New Jersey. IEEE Press.

Dutta, J., Deb, K., Tulshyan, R., and Arora, R. (2013).
Approximate KKT points and a proximity measure for termination.
Journal of Global Optimization, 56(4):1463–1499.

Fonseca, C., Gandibleux, X., Korhonen, P., Martí, L., Naujoks, B., Thiele, L., Wallenius, J., and Zitzler, E. (2009).
Working group on EMO for interactive multiobjective optimization (1st round).
In Deb, K., Greco, S., Miettinen, K., and Zitzler, E., editors, *Hybrid and Robust Approaches to Multiobjective Optimization*, number 09041 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany. Schloss Dagstuhl — Leibniz–Zentrum fuer Informatik.

Goel, T. and Stander, N. (2010).
A non-dominance-based online stopping criterion for multi–objective evolutionary algorithms.
International Journal for Numerical Methods in Engineering.

Guerrero, J. L., García, J., Martí, L., Molina, J. M., and Berlanga, A. (2009).
A stopping criterion based on Kalman estimation techniques with several progress indicators.
In Raidl, G., Alba, E., Bacardit, J., Bates Congdon, C., Beyer, H.-G., Birattari, M., Blum, C., Bosman, P. A. N., Corne, D., Cotta, C., Di Penta, M., Doerr, B., Drechsler, R., Ebner, M., Grahl, J., Jansen, T., Knowles, J., Lenaerts, T., Middendorf, M., Miller, J. F., O'Neill, M., Poli, R., Squillero, G., Stanley, K., Stützle, T., and van Hemert, J., editors, *GECCO'09: 11th Annual Conference on Genetic and Evolutionary Computation*, pages 587–594, New York, NY, USA. ACM Press.

Bibliography II

Guerrero, J. L., Martí, L., García, J., Berlanga, A., and Molina, J. M. (2010).
Introducing a robust and efficient stopping criterion for MOEAs.
In *2010 IEEE Conference on Evolutionary Computation (CEC), part of 2010 IEEE World Congress on Computational Intelligence (WCCI 2010)*, Piscataway, New Jersey. IEEE Press.

Jain, B. J., Pohlheim, H., and Wegener, J. (2001).
On termination criteria of evolutionary algorithms.
In Spector, L., Goodman, E. D., Wu, A., Langdon, W., Voigt, H.-M., Gen, M., Sen, S., Dorigo, M., Pezeshk, S., Garzon, M. H., and Burke, E., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2001)*, page 768, San Francisco, California, USA. Morgan Kaufmann.

Karush, W. (1939).
Minima of functions of several variables with inequalities as side constraints.
Master's thesis, Department of Mathematics, University of Chicago, Chicago, IL, USA.

Kuhn, H. W. and Tucker, A. W. (1951).
Nonlinear programming.
In Neyman, J., editor, *Proceedings of Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 481–492, Berkeley, LA, USA. University of California Press.

Martí, L., García, J., Berlanga, A., and Molina, J. M. (2007).
A cumulative evidential stopping criterion for multiobjective optimization evolutionary algorithms.
In Thierens, D., Deb, K., Pelikan, M., Beyer, H.-G., Doerr, B., Poli, R., and Bittari, M., editors, *GECCO'07: Proceedings of the 2007 GECCO Conference Companion on Genetic and Evolutionary Computation*, page 911, New York. ACM Press.

Martí, L., García, J., Berlanga, A., and Molina, J. M. (2009).
An approach to stopping criteria for multi–objective optimization evolutionary algorithms: The MGBM criterion.
In *2009 IEEE Conference on Evolutionary Computation (CEC 2009)*, pages 1263–1270, Piscataway, New Jersey. IEEE Press.

Bibliography III

Martí, L., García, J., Berlanga, A., and Molina, J. M. (2010).
A progress indicator for detecting success and failure in evolutionary multi–objective optimization.
In *2010 IEEE Conference on Evolutionary Computation (CEC), part of 2010 IEEE World Congress on Computational Intelligence (WCCI 2010)*, Piscataway, New Jersey. IEEE Press.

Rudenko, O. and Schoenauer, M. (2004).
A steady performance stopping criterion for pareto–based evolutionary algorithms.
In *The 6th International Multi–Objective Programming and Goal Programming Conference*, Hammamet, Tunisia.

Rudolph, G. and Agapie, A. (2000).
Convergence properties of some multi–objective evolutionary algorithms.
In *2000 IEEE Congress on Evolutionary Computation (CEC 2000)*, volume 2, pages 1010–1016, Piscataway, New Jersey. IEEE Press.

Stewart, T. J. (1997).
Convergence and validation of interactive methods in mcdm: simulation studies.
In *Essays in Decision Making*, pages 7–18. Springer.

Trautmann, H., Ligges, U., Mehnen, J., and Preuss, M. (2008).
A convergence criterion for multiobjective evolutionary algorithms based on systematic statistical testing.
In *Parallel Problem Solving from Nature — PPSN X*, pages 825–836. Springer, Heidelberg.

Wagner, T., Trautmann, H., and Martí, L. (2011).
A taxonomy of online stopping criteria for multi-objective evolutionary algorithms.
In Takahashi, R. H. C., Deb, K., Wanner, E. F., and Greco, S., editors, *6th International Conference on Evolutionary Multi-Criterion Optimization (EMO 2011)*, volume 6576, pages 16–30, Berlin/Heidelberg. Springer.

Wagner, T., Trautmann, H., and Naujoks, B. (2009).
OCD: Online convergence detection for evolutionary multi–objective algorithms based on statistical testing.
In Ehrgott, M., Fonseca, C. M., Gandibleux, X., Hao, J.-K., and Sevaux, M., editors, *5th International Conference on Evolutionary Multi–Criterion Optimization (EMO 2009)*, volume 5467 of *Lecture Notes in Computer Science*, pages 198–215. Springer.