

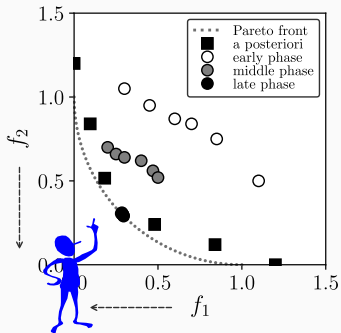


Decomposition-based co-evolutionary algorithm for interactive multiple objective optimization

Michał Tomczyk Miłosz Kadziński

Institute of Computing Science
Poznan University of Technology, Poland

michal.tomczyk@cs.put.poznan.pl
www.cs.put.poznan.pl/mtomczyk/



Keywords:

- EMO {
- Multiple objective optimization
 - Evolutionary Algorithms
- MCDAs {
- Preference Learning

Preference learning:



1. It is a cooperation between the algorithm and the DM where one participant **interactively** learns from the other.
2. The DM's preferences are inferred via **preference disaggregation** – deriving a global model from some preferential structures, e.g., **pairwise comparisons**.

Reminder on IEMO/D

- IEMO/D uses a functional preference model to **represent the DM's preferences mathematically** (L-norms):

$$L_{\alpha}^w(s) = \begin{cases} \left[\sum_{i=1}^M (w_i s_i)^{\alpha} \right]^{1/\alpha} & \text{for } \alpha < \infty \\ \max_{1, \dots, M} \{w_i s_i\} & \text{for } \alpha = \infty \end{cases}$$

- α – compensation level – is provided *a priori*. Weight vector is **unknown**.
- Interactively provided pairwise comparisons are used to **constrain** the model parameter space.
- A fine representation of the compatible weight vectors is used to **instantiate goals** in the decomposition-based evolutionary framework.



M. K. Tomczyk and M. Kadziński. Decomposition-based interactive evolutionary algorithm for multiple objective optimization. *IEEE Transactions on Evolutionary Computation* 24, 2 (2020), 320–334

Assumptions and inconsistency

Assumptions:

- The assumed preference model is compatible with the DM's value system
- DM behaves rationally

violation

Inconsistency

There does not exist a set of parameter values that makes instantiated in such a way model feasible.

Potential solutions to the problem:

- replacing the incorporated model
- revising the set of maintained preference examples
- identifying irrational decisions
- maintaining a set of different preference models






The proposed algorithm – CIEMO/D

Key features:

- CIEMO/D is based on **IEMO/D** algorithm.
- It **co-evolves** a set of different **species (subpopulations)**, each associated with a different preference model.
- Solutions can migrate between different **subpopulations**, hence implementing the **cooperative paradigm**.
- Sizes of sub-populations are **dynamically adjusted** according to their level of consistency with the DM's pairwise comparisons, implementing in this way the **competitive paradigm**.

Each species incorporates a different model: α for the L-norm, for instance:



If implemented naïvely, CIEMO/D could be perceived just as IEMO/D, run several Times for different α -values

Resource distribution strategies

Each species implements the preference learning independently, i.e., it maintains its own copy of the history of preference elicitation (pairwise comparisons) and constructs a fine representation of the space of compatible model parameters.

The inconsistency may occur when the DM's value system does not align with the incorporated L-norm model. To reinstate consistency, CIEMO/D follows the procedure implemented in NEMO algorithms:

1. **remove** the oldest pairwise comparisons, one by one, until the consistency is reinstated;
2. **bring back**, in the reversed order, these preference examples that do not violate consistency

Inconsistency level

→ the number of pairwise comparisons the species had to remove throughout the evolutionary process to maintain consistency:

0 : model is compatible so far with the DM's pairwise comparisons

> 0 : model is inconsistent; the greater is the number of removed pairwise comparisons, the greater is the inconsistency

J. Branke, S. Greco, R. Słowiński, P. Zielniewicz, Learning Value Functions in Interactive Evolutionary Multiobjective Optimization, IEEE Transactions on Evolutionary Computation, 19(1):88–102, 2015.

Resource distribution strategies

TPS – Population Size
 PS_j – j -th subpopulation size
 RPC_j – the number of pairwise comparisons removed by j -th species
 $\mathcal{L} = [a_1, a_2, \dots, a_l]$ – model definitions

- **Equal Distribution (EQ)**: a benchmark strategy – the computational resources are distributed equally, i.e., $PS_j = TPS \div |\mathcal{L}|$
- **Promoting the most compatible models (TMCM- β)**: prioritizes species according to two criteria – (1) consider only the most compatible models; (2) among these promote the most **compensatory** models such that: $PS_{J_k} \div PS_{J_{k+1}} = \beta$ ($J = \{j: RPC_j = \min(RPC), j \in \{1, \dots, |\mathcal{L}|\}\}$).

Example ($TPS = 100$)

j	1	2	3
α	1	5	∞
RPC	1	0	0
β	PS		
1	0	50	50
2	0	67	33
3	0	75	25
∞	0	100	0

Resource distribution strategies

TPS – Population Size
 PS_j – j -th subpopulation size
 RPC_j – the number of pairwise comparisons removed by j -th species
 $\mathcal{L} = [a_1, a_2, \dots, a_l]$ – model definitions

- **Proportional distribution (PROP)**: the subpopulation sizes are set inversely proportional to the numbers of pairwise comparisons removed by them to reinstate the consistency:

$$PS_j = TPS \cdot \frac{\sum_{k=1}^{|\mathcal{L}|} (RPC_k + 1) - (RPC_j + 1)}{(|\mathcal{L}| - 1) \sum_{k=1}^{|\mathcal{L}|} (RPC_k + 1)}$$


Example ($TPS = 100$)

j	1	2	3
α	1	5	∞
RPC	PS		
[2,1,0]	25	33	42
[0,3,2]	44	25	31
[4,0,0]	14	43	43




Visualization of convergence

- WFG3 with $M = 2$ objectives
- $TPS = 30$
- $\mathcal{L} = [a_1 = 1, a_2 = 5, a_3 = \infty]$
- 10 interactions distributed evenly throughout the evolution,
- the method was run for 1000 generations
- The artificial DM's value system was modeled as L_∞ -norm with $w = [0.7, 0.3]$
- Distribution policy = *TMCM-2*



Visualization of convergence

- WFG3 (non-degenerated variant) with $M = 3$ objectives
- $TPS = 60$
- $\mathcal{L} = [a_1 = 1, a_2 = 5, a_3 = \infty]$
- 10 interactions distributed evenly throughout the evolution,
- the method was run for 1000 generations
- The artificial DM's value system was modelled as L_∞ -norm with $w = [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]$
- Distribution policy = $TMCM-2$



- **Evolutionary setting:** follows the standards in the literature on EMO.
- **Decision-making layer:**

Interactions: triggered 10 times during a single run, evenly distributed

Simulating the DM's answers: the DM's value system was modeled using an L-norm


Comprehensiveness & reliability of the experiments: for each setting, the run was repeated 100 times, each time involving a different artificial DM (these were pre-generated by generating uniformly distributed weight vectors).

Performance evaluation: solutions in the population were compared against the optimal solution identified in advance using exact or heuristic methods. Specifically, we computed the Best/Average Relative Score Differences (BRSD/ARSD) between the most favored (average for all solutions) and the optimum, where scores were assessed using the artificial DM's internal function.

Statistics: mean, standard deviation, average rank

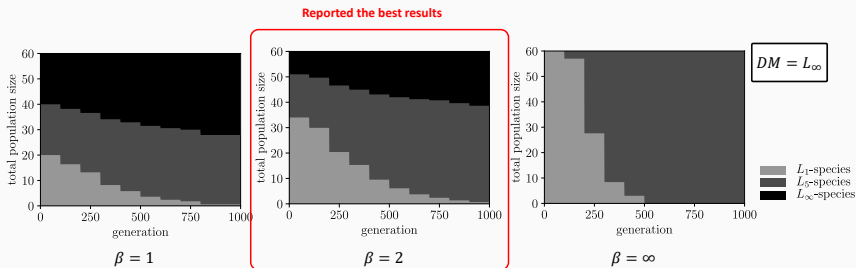
Impact of different values of β on the performance of CIEMO/D with a distribution policy = $TMCM-\beta$

Non-compensatory models approximate well the compensatory ones... therefore, inconsistencies rarely occur.
In this way, the greater the β , the more dominant is the L_1 -species in the population



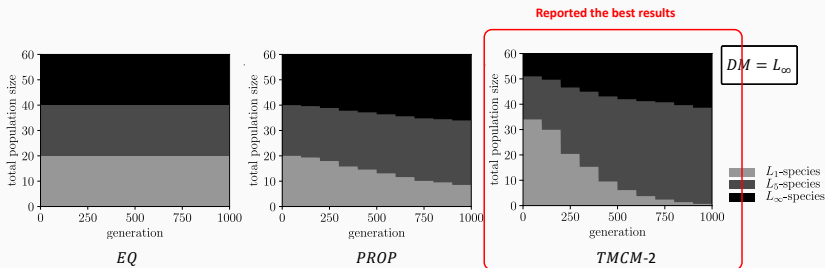
Average population size of each subpopulation maintained by CIEMO/D with $\mathcal{L} = [L_1; L_5; L_\infty]$ and distribution policy = $TMCM$ -throughout the evolutionary search when applied to (c)WFG4 with $M = 3$. The colored areas illustrate a ratio between sizes of respective subpopulations.

Impact of different values of β on the performance of CIEMO/D with a distribution policy = $TMCM-\beta$



Average population size of each subpopulation maintained by CIEMO/D with $\mathcal{L} = [L_1; L_5; L_\infty]$ and distribution policy = $TMCM$ throughout the evolutionary search when applied to (c)WFG4 with $M = 3$. The colored areas illustrate a ratio between sizes of respective subpopulations.


Comparison of CIEMO/D with the following distribution policies: *EQ*, *PROP*, and *TMCM-2*




Average population size of each subpopulation maintained by CIEMO/D with $\mathcal{L} = [L_1; L_5; L_\infty]$ with different distribution procedures throughout the evolutionary search when applied to WFG4 with $M = 3$. The colored areas illustrate the ratio between sizes of respective subpopulations.

Evaluation of CIEMO/D incorporating different procedures for resource allocation

Average BRSD and ARSD throughout 1000 generations for CIEMO/D with $\mathcal{L} = [L_1; L_5; L_\infty]$ and different distribution procedures applied to (c)WFG4 with $M = 3$ and $DM \in \{L_1, L_\infty\}$.



$DM = L_1$



$DM = L_\infty$


- ▲— EQ BRVD
- ▲·· EQ ARVD
- ▲— PROP BRVD
- ▲·· PROP ARVD
- TMCM-2 BRVD
- TMCM-2 ARVD

CIEMO/D with *TMCM-2* outperformed the *EQ* and *PROP* variants

Comparison of CIEMO/D with its counterpart IEMO/D not involving co-evolution




The total number of function evaluations was set to the same value in all algorithms



Our research hypothesis was: "can CIEMO/D perform no worse than its main competitor – a dedicated IEMO/D?"

Comparison of CIEMO/D with its counterpart IEMO/D not involving co-evolution

Average BRSD and ARSD throughout 1000 generations for different algorithms procedures applied to (c)WFG4 with $DM \in \{3,5\}$



CIEMO/D performed similar to the dedicated IEMO/D_{L1} algorithm



Other performer experiments

- Comparison with existing state-of-the-art methods: one-model NEMO algorithms and two-model NEMO-II-Choquet method ✓
- Evaluation of CIEMO/D on a real-world green logistics problem ✓
- Analysis of the performance of CIEMO/D not including the DM's true model ✓
- Analysis of the computational complexity and the execution times ✓



Avenues for future research

- Determining the level of inconsistency based on other factors than the number of removed pairwise comparisons ?
- Using different models than L-norms to represent the DM's preferences ?

Thank you for your attention! 😊

M. K. Tomczyk, M. Kadziński, Decomposition-based co-evolutionary algorithm for interactive multiple objective optimization, Information Sciences, 549:178–199, 2021.

michal.tomczyk@cs.put.poznan.pl