Interactive co-evolutionary multiple objective optimization algorithms for finding consensus solutions for a group of Decision Makers

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/

Introduction

Keywords

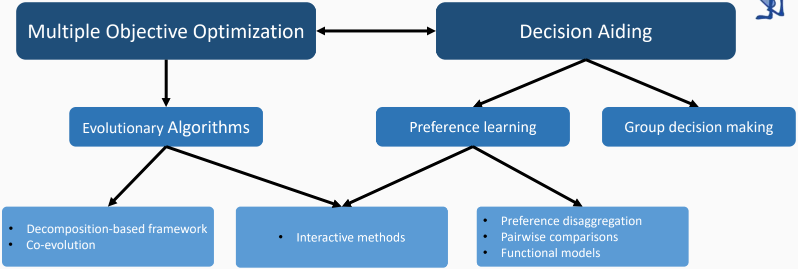

Introduction
Multi-objective optimization

Optimization perspective: class of optimization problems in which solutions are evaluated by at least two objectives.

MCDA perspective: multiple criteria decision problem in which solutions are not available at hand. Instead, they are expressed by some mathematical formulas.

Optimization vs. decision problem: (optimization) adds additional level of complexity (uncertainty) into the decision problem

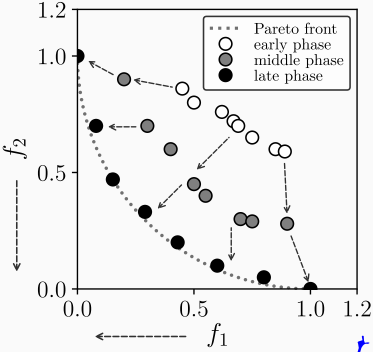

Evolutionary algorithms for MOO:

- „mimic" natural evolution: maintaining a population of solutions, creating new solutions using genetic operators, etc.
- Iterative algorithms
- Fast, efficient, flexible (heuristics)

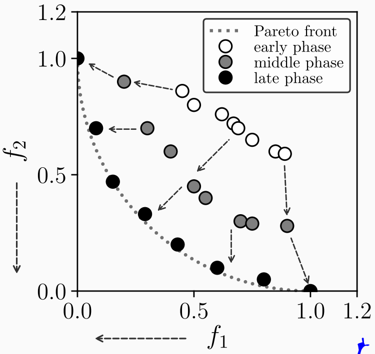

Decomposition-based evolutionary framework:

- **Uses a set of scalar optimization functions (goals) to steer the evolutionary search.**
- Supports efficient evolutionary mechanisms like, e.g., restricting mating pools or running in a steady-state mode.
- Fast, efficient, flexible.

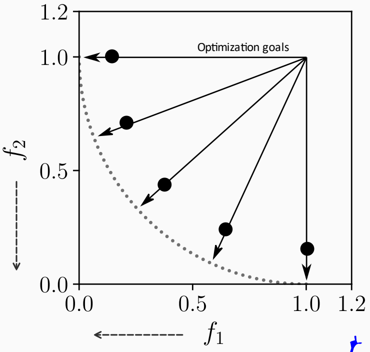

Interactive preference-based EMO:

Since the ultimate goal is to pick one solution that satisfies the DM's the most, it is not necessary to approximate an entire Pareto front. Instead, **the method may be interactively supplied with DM's preferences** and use thus derived knowledge **to bias the evolutionary search towards DM's preferred solutions in the Pareto front**.

Benefits:

- The convergence towards the Pareto front is faster due to the additional bias imposed.
- The DM-perceived qualities of obtained solutions can be improved.

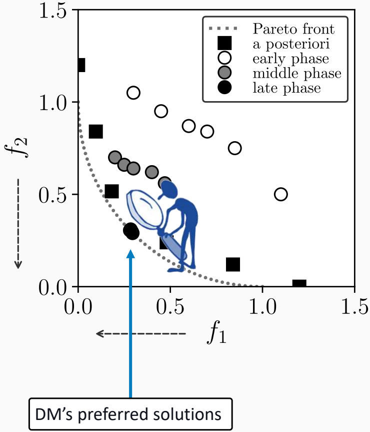

IEMO/D

- IEMO/D is based on **MOEA/D algorithm.**
- Interactive, based on pairwise comparisons of solutions
- Represents the DM's by using a fine representation of the space of compatible L-norms:

$$\forall_{s^j \succ s^k \in H} L_\alpha^w(s^j) < L_\alpha^w(s^k) \rightarrow w \text{ is compatible (feasible)}$$

- After each preference elicitation, this set is incorporated into the decomposition-based framework as optimization goals **(robustness concern).**

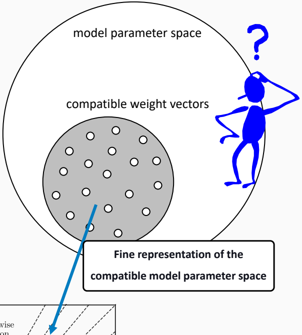

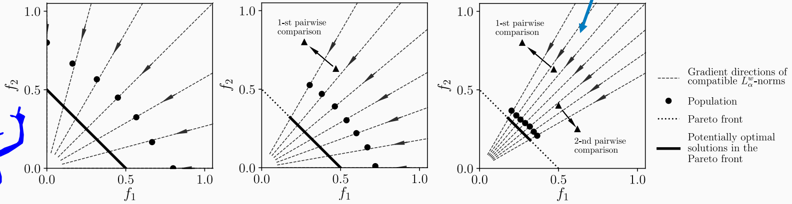

IEMO/D: *M. K. Tomczyk and M. Kadziński. Decomposition-based interactive evolutionary algorithm for multiple objective optimization. IEEE Transactions on Evolutionary Computation 24, 2 (2020), 320–334*

Decision problems often involve **multiple Decision Makers**. They may have

- different preferences,
- importance in the committee.

Therefore, **unanimity** may not be possible, and a **consensus** must be elaborated.

In the context of **evolutionary multi-objective optimization**, it means that the algorithm **will not** be able to find a **single** solution in the Pareto front **simultaneously satisfying all the DM's the most**.

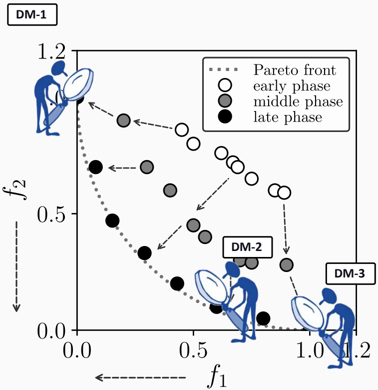

IEMO/D_G
Properties

IEMO/D_G extends IEMO/D to handle

<u>Group Decision Making</u>

As with IEMO/D, **IEMO/D_G**:

- is based on **pairwise comparisons** of solutions,
- is **interactive**,
- implements the **decomposition-based** evolutionary framework,
- implements **robust preference learning** techniques,
- uses **penalty-based boundary intersection** (PBI) function as a **preference model**:

Reference point $\boldsymbol{r}$ is the main model parameter

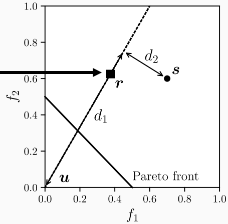

IEMO/D_G learns DMs' preferences **individually in the same way as IEMO/D** learns the preferences of a **single DM**. So, during the preference elicitation step, **for each DM**:

1. A suitable pair of solutions is presented to the DM for comparison
2. His or her feedback is used to constrain (individually for every DM) model parameter space (reference points).
3. A fine representation of model instances is sampled from thus constrained model parameter space.

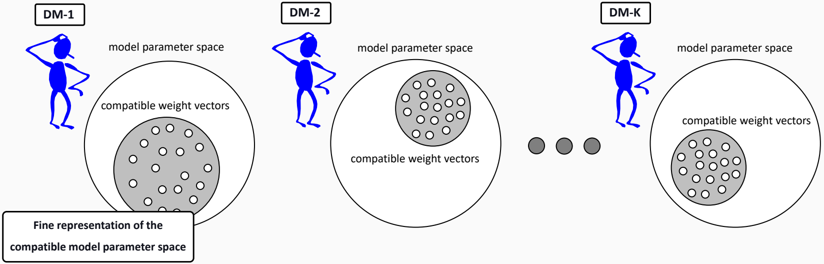

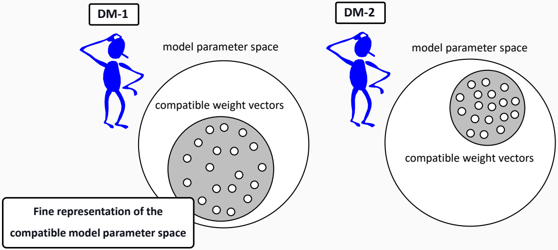

Okay, so we can obtain, for every DM, a set of PBI functions well-representing their preferences. **But how can thus obtained knowledge be incorporated within the decomposition-based evolutionary framework**? Should all these sets, as a joint set, be incorporated? **No, because in this way, the DMs' individual best solutions will be obtained instead of the consensus.**

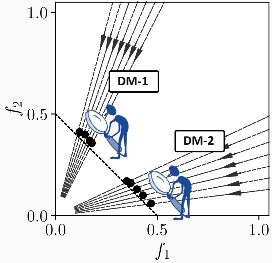

IEMO/D_G
Consensus modeling

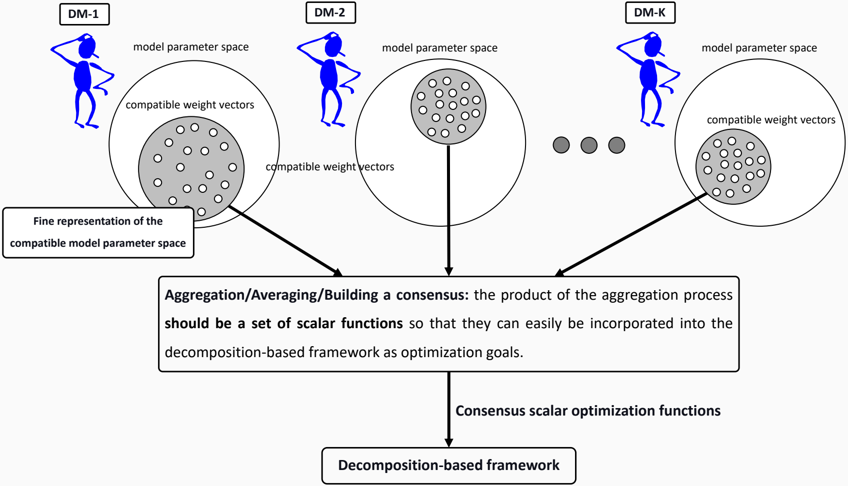

IEMO/D_G

Consensus modeling

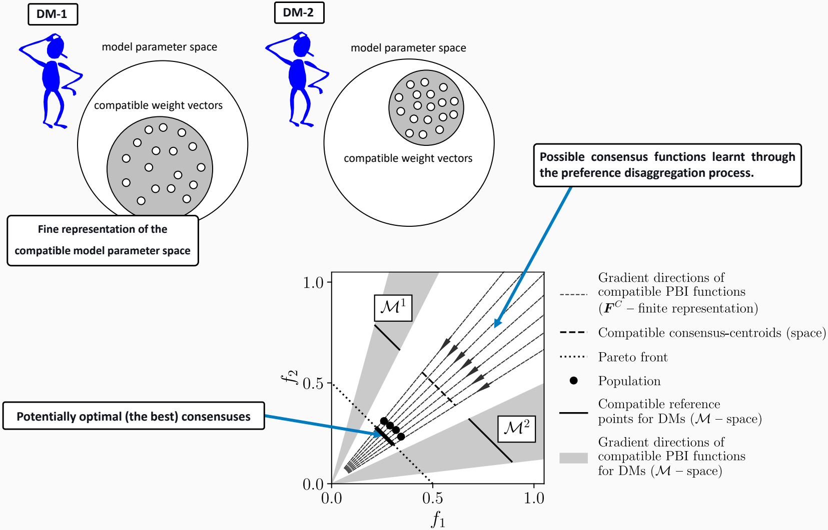

IEMO/D_G

Consensus modeling

<u>Model-level aggregation</u>: Having K (the number of DMs) compatible **PBI functions** for (one per DM) and their respective reference points, we build a new PBI function parameterized with the **centroid of these reference points** (CONS = CENTR policy).

<u>Evaluation-level aggregation</u>: Having K (the number of DMs) compatible PBI functions for (one per DM), we construct an **auxiliary function** that evaluates an input solution by calculating the **average score attained by this solution according to these PBI functions** (CONS = UTILIT policy).

For both consensus policies, the DMs' contributions to solutions' overall score can **be weighted** to reflect DMs' importance in the committee.

Okay, compared to IEMO/D, we have just **duplicated the preference learning module, for each DM**, and, during the search, **incorporate its averaged outcomes into the decomposition-based framework**. Is that all?

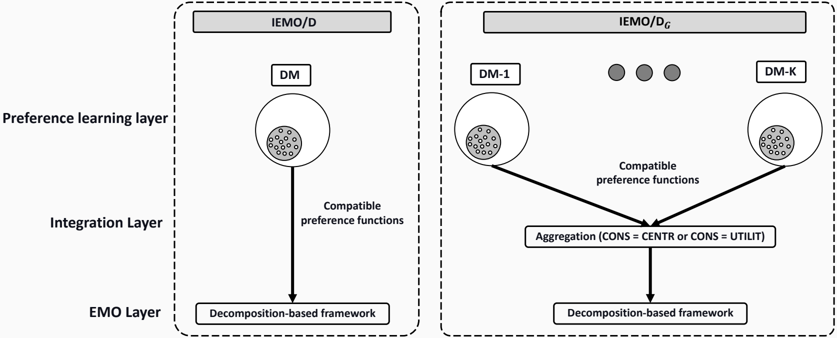

Critical problem (in optimization scenarios): The efficiency of the preference learning process vastly **depends on the solutions presented to DMs for evaluation**. However, when building the consensus, the **evolved solutions may not align with the DMs' preferences individually**. Hence, they **may not be helpful** when learning DMs' preferences, **which may cause stagnation**.

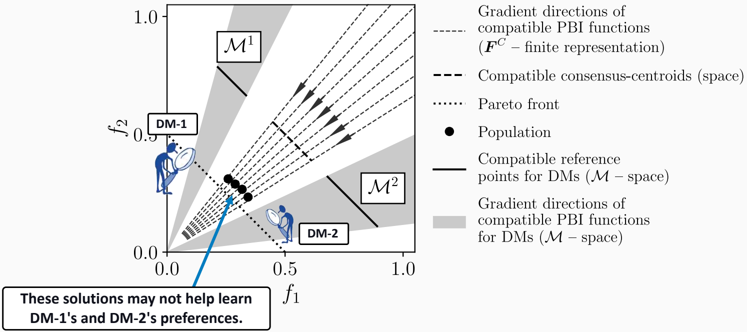

To solve this problem, we propose **to co-evolve two populations**, named primary and the supportive:

- **Primary population $\boldsymbol{P^P}$** : aims at discovering the best consensus solutions.
- **Supportive population $\boldsymbol{P^S}$**: aims at approximating the Pareto front, **offering a rich set of solutions representing various trade-offs** between the objectives.

Solution pairs for comparisons are selected from joints set $\boldsymbol{P^P}$ and $\boldsymbol{P^S}$, **for every DM individually**. For this reason, we are using a **PWIT** selection strategy that maximizes the potential information gain from the DMs' answers.

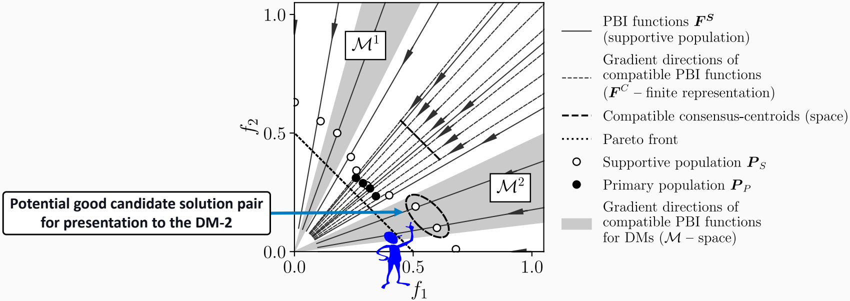

PWIT: *M. K. Tomczyk and M. Kadziński. On the elicitation of indirect preferences in interactive evolutionary multiple objective optimization.* Proceedings of the 2020 Genetic and Evolutionary Computation Conference

IEMO/D_G

Dealing with a diminished variability of solutions – example results

Setting: 3 DMs, WFG3 test problem, 2 objectives, DMs' answers simulated using artificial PBI functions, CONS=CENTR

ViPEMO method for averaging many independent runs

IEMO/D_G without co-evolution

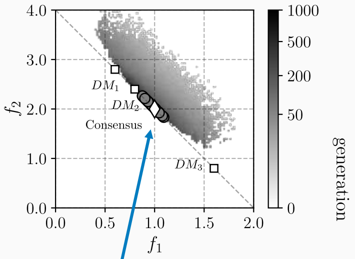

IEMO/D_G with co-evolution

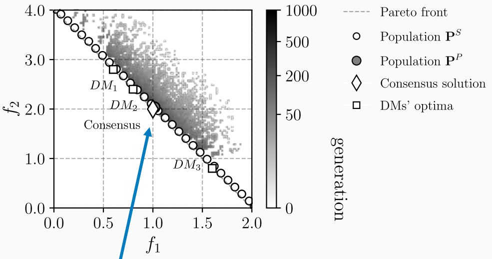

Greater spread of final solutions indicates a stagnation

Due to the co-evolving two species, the primary population was able to better concentrate around the best consensuses

ViPEMO: *M. Kadziński, M. K. Tomczyk, and R. Słowiński, Preference-based cone contraction algorithms for interactive evolutionary multiple objective optimization. Swarm and Evolutionary Computation 52, 100602, 2020.*

IEMO/D_G

Dealing with a diminished variability of solutions – example results

Setting: 3 DMs, WFG3 test problem (non-degenerate variant), 3 objectives, DMs' answers simulated using artificial PBI functions, CONS=CENTR

IEMO/D_G without co-evolution

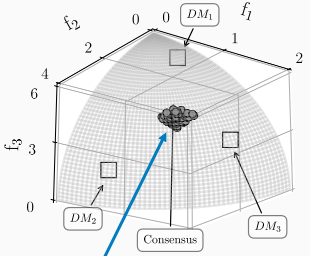

Greater spread of final solutions indicates a stagnation

IEMO/D_G with co-evolution

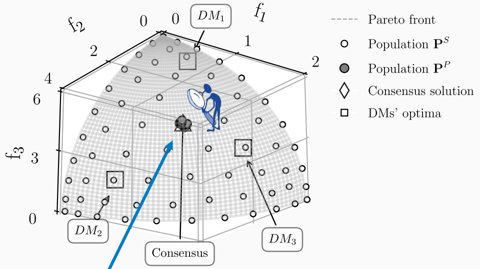

Due to the co-evolving two species, the primary population was able to better concentrate around the best consensuses

IEMO/D_G

Properly distributing the computational effort

Another problem: Co-evolving two populations **proves to handle the stagnation problem**. However, adding the new species to the evolution is **not costless**. Assuming a total constant population size, we are **lowering** the number of solutions converging towards consensuses, hence **diminishing the convergence speed**, in favor of additionally approximating the Pareto front.

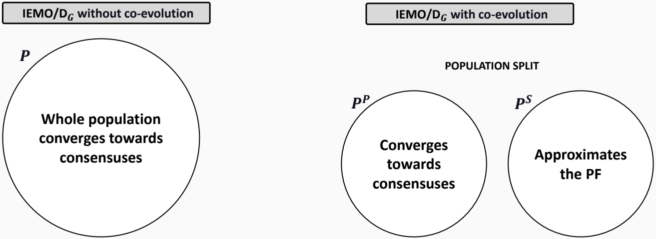

Solution → Dynamic Reallocation of Resources: we propose to check if **some of the optimization goals** (and thus corresponding solutions) in the **supportive population become** of **no use** when it comes to efficient preference learning. If so, **we remove them from the supportive population in favor of increasing the primary population size**. To verify the goal's usability, we check if it is compatible with pairwise comparisons of at least one DM. If not, we assume that this goal can be removed.

These ● goals are potentially useless for efficient preference learning so they can be removed

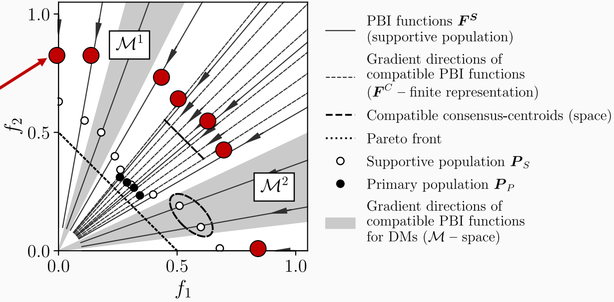

Properly distributing the computational effort – example results

Setting: 3 DMs, WFG3 test problem, 2 objectives, DMs' answers simulated using artificial PBI functions, CONS=CENTR

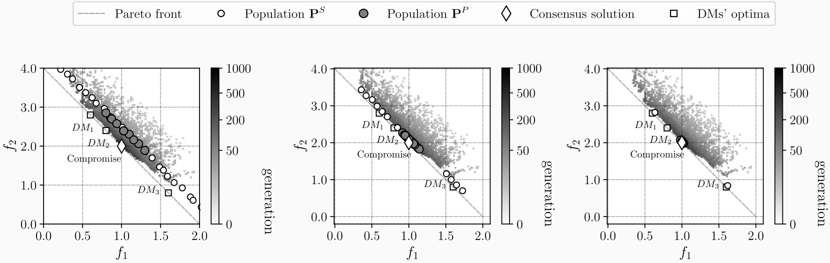

50-th generation

250-th generation

1000-th generation

IEMO/D$_G$

Properly distributing the computational effort – example results

Setting: 3 DMs, WFG4 test problem, 2 objectives, DMs' answers simulated using artificial PBI functions, CONS=CENTR

Black lines: average-in-population solution quality in view of DMs' true consensus function (the lower the better) **(only primary population counts)**

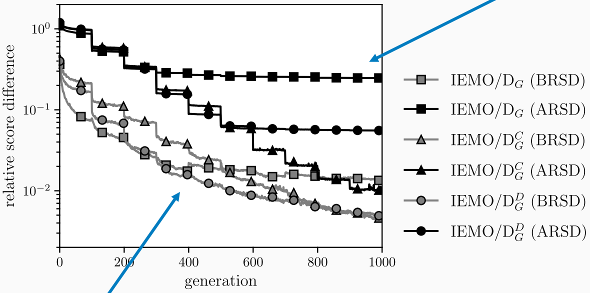

Grey lines: best-in-population solution quality in view of DMs' true consensus function (the lower the better) **(only primary population counts)**

Example results

Setting: 3 DMs, WFG4 test problem, 2 objectives, DMs' answers simulated using artificial PBI functions, CONS=CENTR

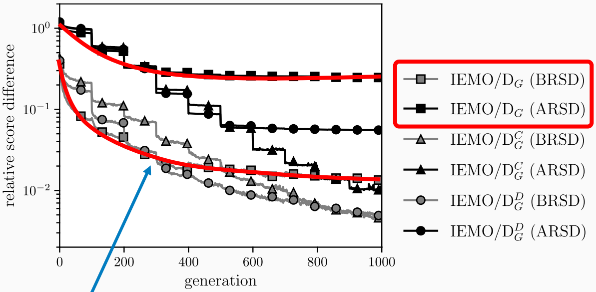

Non co-evolutionary and non-dynamic variant: initial fast convergence but prone to stagnation

Example results

Setting: 3 DMs, WFG4 test problem, 2 objectives, DMs' answers simulated using artificial PBI functions, CONS=CENTR

Co-evolutionary but non-dynamic variant: the best average-in-population solutions' qualities, but this result is rather positively affected by maintaining a small and very concentrated primary population

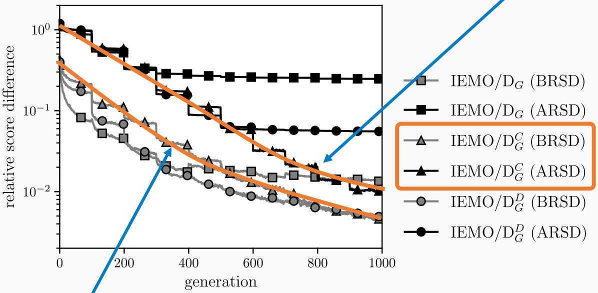

Co-evolutionary but non-dynamic variant: good final solutions' qualities, but overall the convergence speed is decreased

Example results

Setting: 3 DMs, WFG4 test problem, 2 objectives, DMs' answers simulated using artificial PBI functions, CONS=CENTR

Non co-evolutionary and non-dynamic variant: averge-in-population quality is worse than when compared with the non-dynamic variant, but mainly because of maintaining a primary population of greater size (hence the solutions are more spread)

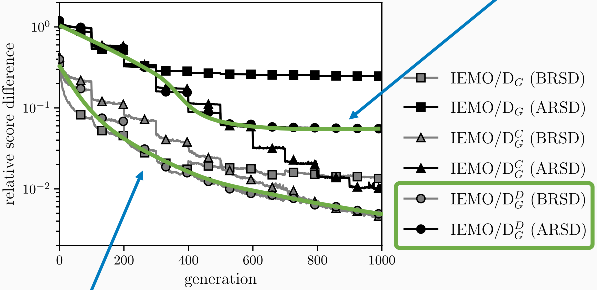

Non co-evolutionary and non-dynamic variant: boost in convergence speed can be observed when compared with the non-dynamic variant.

Example results

- the execution times are in **reasonable bounds**
- the n**on-co-evolutionary** + **non-dynamic method** attained **the worst execution times** (prolonged due to maintaining overall the greatest number of solutions in the primary population)
- the **co-evolutionary** + **non-dynamic** attained **the best execution times** (due to maintaining the smallest number of solutions in the primary population)
- the **co-evolutionary** + **dynamic method** attained **moderate execution times**

	$M=2$		$M=3$		$M=4$		$M=5$		$M=2$		$M=3$		$M=4$		$M=5$	
	$K=2$								$K=3$							
Algorithm	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
IEMO/D_G	0.40	0.12	1.35	0.18	6.14	0.38	27.07	2.84	0.49	0.18	1.35	0.23	8.13	0.60	37.02	4.59
IEMO/D_G^C	0.35	0.02	1.12	0.02	3.85	0.09	12.72	0.21	0.33	0.01	1.18	0.03	4.34	0.14	15.38	0.32
IEMO/D_G^D	0.33	0.01	1.34	0.19	5.65	0.49	23.83	2.17	0.36	0.01	1.36	0.21	7.16	0.44	29.37	5.86
	$K=4$								$K=5$							
IEMO/D_G	0.40	0.01	1.60	0.39	9.66	1.08	42.26	11.18	0.42	0.02	1.76	0.42	12.24	0.18	57.46	9.38
IEMO/D_G^C	0.35	0.01	1.29	0.04	4.77	0.14	18.00	0.40	0.38	0.07	1.34	0.06	5.83	0.39	20.79	0.52
IEMO/D_G^D	0.36	0.01	1.52	0.37	8.06	1.26	36.66	5.01	0.38	0.01	1.91	0.42	9.93	0.34	42.35	7.71

Note that the execution times were measured when the methods were applied to **benchmark** problems that take **almost no time for solution evaluation**. However, for **real-world problems**, the solution evaluation phase may take a **long** time. In that case, **the total number of evaluated solutions** required to reach some **quality threshold** is a more critical factor than pure execution time. In this view, the **co-evolutionary + dynamic variant can be favored over its non-dynamic counterpart**.

Example results

- We performed **extensive visualization experiments** revealing that our method can reach consensus solutions when applied to problems of different classes.
- We also **compared** our algorithm with some **existing methods in this stream**, proving its **high competitiveness**.
- Additionally, we **applied** our method to a **real-world optimization problem** of constructing an environmentally friendly supply chain in Southeastern Europe, **proving its usability**. In this experiment, we also **demonstrated** the method's performance when the **DMs involved** are of **different importance in the committee**.

$$\boldsymbol{w^{DM}} = \left[\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right] \rightarrow$$

$$\boldsymbol{w^{DM}} = [0.3, 0.6, 0.1] \rightarrow$$

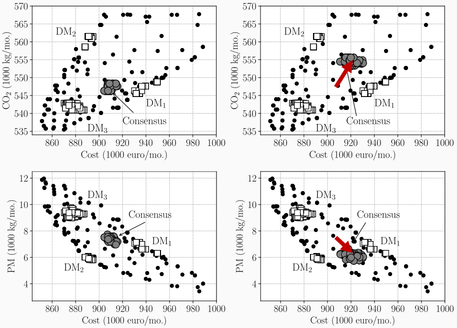

Conclusions & avenues for future research

Conclusions

- We introduced a novel preference-based evolutionary multi-objective algorithm **dedicated to group decision problems**
- **IEMO/D$_G$** is based on **the state-of-the-art concepts in EMO and MCDA**, i.e., it uses an efficient evolutionary framework, is based on preference learning, and it implements robustness concerns
- **IEMO/D$_G$** can be run in the **co-evolutionary** run to **handle the potential stagnation in learning DMs; preferences** that may happen due to maintaining a poorly diversified population
- **IEMO/D$_G$** can **dynamically redistribute solutions in the co-evolved** species to ensure a **proper balance** between the efficiency of the preference learning and the convergence speed

Avenues for future research

- further investigation on how to maintain a proper diversity of solutions to ensure efficient preference learning
- further investigation on how to dynamically re-allocate resources
- implementation of the negotiation system & possibility of altering previously expressed preferences interactively

Thank you for your attention!

michal.tomczyk@cs.put.poznan.pl