EMOSOR: Evolutionary Multiple Objective Optimization Guided by Interactive Stochastic Ordinal Regression

Michał Tomczyk, Miłosz Kadziński

Laboratory of Intelligent Decision Support Systems

Institute of Computing Science

Poznań University of Technology

June 25, 2019

Evolutionary Multiple-objective Optimization (EMO)

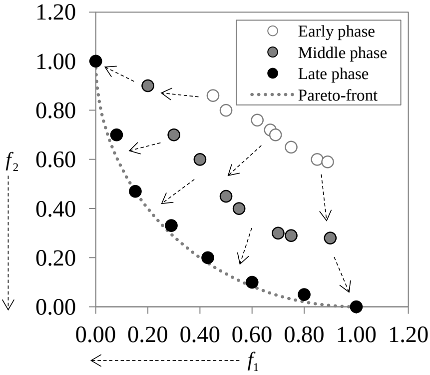

Evolutionary Algorithms for MOO

Mimic the process of naturall evolution to solve optimization problems

Advantages of EMO:

– can be applied to problems having complex fitness landscapes

– the computational complexity can be reduced since solutions are optimized in an interrelated manner

Preference vs. non preference-based EMOAs

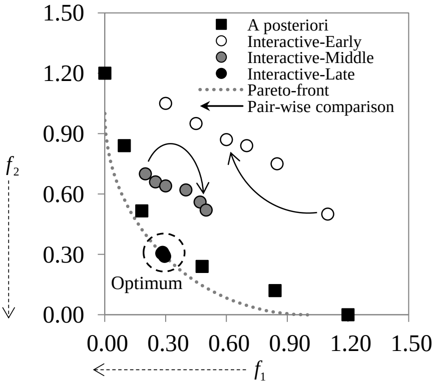

Preference-based EMOAs

Observation: it is not practical to approximate an entire PF since the DM is interested in finding only relevant solutions to him or her

Incorporation of DM's preferences

Preference information can be used to **constrain** the search space, thereby reducing the complexity of the problem.

The preference information can be used to impose an additional selection pressure, driving population of solutions towards **region in the PF, being highly preferred to the DM**

Scheme of an interactive EMOA

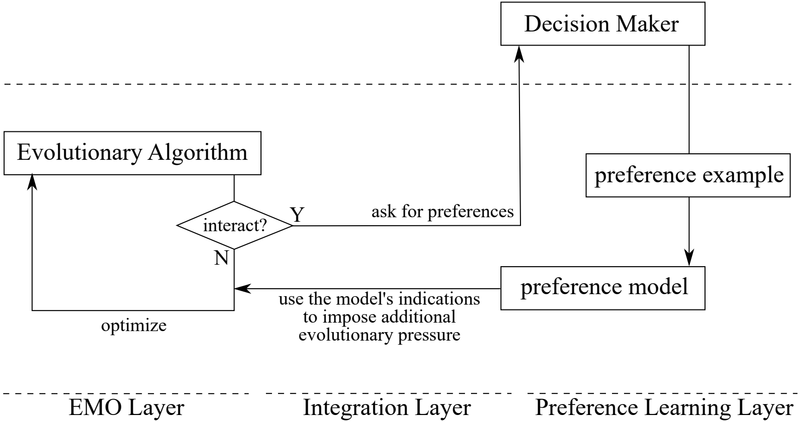

The proposed method: EMOSOR

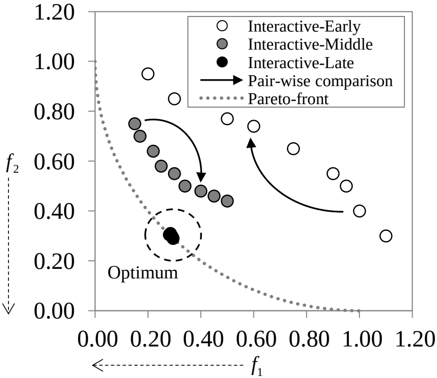

Properties of EMOSOR
▷ Interactive
▷ Based on pairwise comparisons
▷ Uses either: – a Chebyshef Function (CF) or – an Additive Value Function (AVF) to model the DM's preferences
▷ Uses stochastic indicators to: - impose an evolutionary pressure and - select solutions for pairwise comparisons

Preference modeling in EMOSOR

Preference model:

Chebyshev Function:

$$f_{CF}(s) = \max_{i=1,\ldots,M} w_i s_i.$$

parameters: weights

Additive Value Function:

$$f_{AVF}(s) = \sum_{i=1}^{M} u_i(s).$$

parameters: shapes of marginal value functions; the functions are piece-wise linear (characteristic points), monotonic, normalized.

Preference information

The DM is asked to compare two solutions selected from the current population:

$$s^a \; ? \; s^b.$$

This information is used to constrain the parameter space of the assumed preference model:

$$\forall_{s^a \succ^{DM} s^b \in \mathcal{H}} f_{CF}(s^a) < f_{CF}(s^b),$$

$$\sum_{i=1}^{M} w_i = 1.$$

Compatible model instances

How to exploit the set of compatible model instances in order to model the DM's preferences?

How to use these indications during the evolutionary search to direct the optimization towards the region in the PF containing highly preferred solutions?

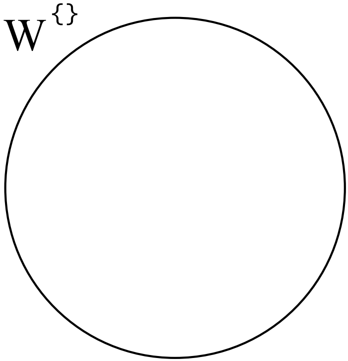

Compatible model instances

How to exploit the set of compatible model instances in order to model the DM's preferences?

How to use these indications during the evolutionary search to direct the optimization towards the region in the PF containing highly preferred solutions?

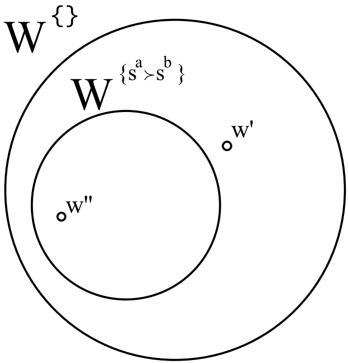

Compatible model instances

Representative model instance
Some methods select only one representative model instance, according to some policy. For instance, they may select the most discriminative model instance:

$$\max \epsilon$$
$$\forall_{s^a \succ^{DM} s^b \in \mathcal{H}} f_{CF}(s^a) < \epsilon + f_{CF}(s^b),$$
$$\sum_{i=1}^{M} w_i = 1.$$

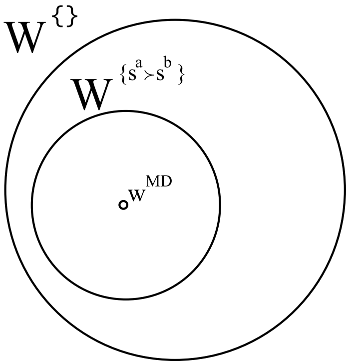

Compatible model instances

Some methods select only one representative model instance, according to some policy. For instance, they may select the most discriminative model instance.

Example: NEMO-0[1]

Sorting criteria of NEMO-0
1) Non-dominated fronts
2) Representative function

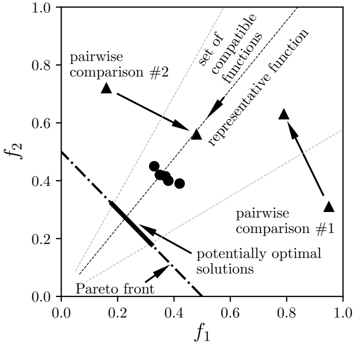

[1] J. Branke, S. Greco, R. Słowiński, and P. Zielniewicz, "Learning valuefunctions in interactive evolutionary multiobjective optimization," IEEE Transactions on Evolutionary Computation, vol. 19, no. 1, pp. 88-102, 2015.

Compatible model instances

Robustness preoccupation Some methods concern a whole set of compatible model instances. In this regard, they are prudent since they do not neglect any compatible model instance. Furthermore, they approximate a set of Pareto optimal solutions being potentially the most relevant (optimal) to the DM.

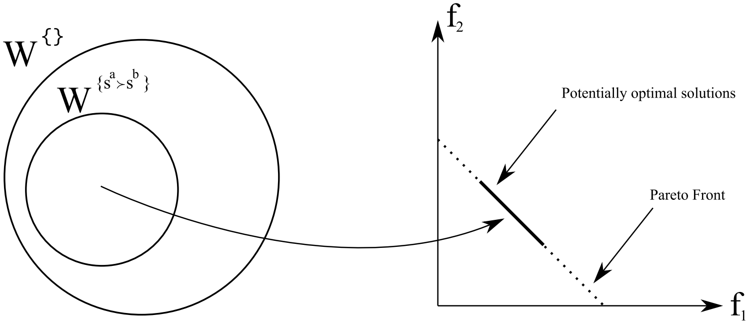

Compatible model instances

Robustness preoccupation Some methods concern a whole set of compatible model instances. In this regard, they are prudent since they do not neglect any compatible model instance. Furthermore, they approximate a set of Pareto optimal solutions being potentially the most relevant (optimal) to the DM.

Example: NEMO-II[2]

Sorting criteria of NEMO-0
1) Fronts of potential optimality
2) Crowding-distance

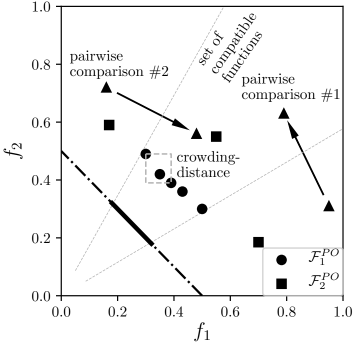

[2] J. Branke, S. Greco, R. Słowiński, and P. Zielniewicz, "Learning valuefunctions in interactive evolutionary multiobjective optimization," IEEE Transactions on Evolutionary Computation, vol. 19, no. 1, pp. 88-102, 2015.

Representative model instance vs. Robustness Preoccupation

Representative model instance

- ▷ Imposes a strong evolutionary pressure,
- ▷ Naïve approach.

Robustness preoccupation

- ▷ Imposes a weak evolutionary pressure,
- ▷ Prudent approach.

Representative model instance vs. Robustness Preoccupation

Representative model instance

- ▷ Imposes a strong evolutionary pressure,
- ▷ Naïve approach.

Robustness preoccupation

- ▷ Imposes a weak evolutionary pressure,
- ▷ Prudent approach.

Stochastic approach

Aggregates the acceptability indices derived from the stochastic analysis.

Stochastic Ordinal Regression

Pairwise Winning Index $PWI(s^j, s^k)$

$PWI(s^j, s^k)$ is a share of compatible preference model instances confirming that s^j is preferred to s^k.

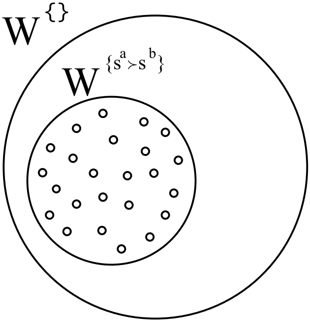

Stochastic Ordinal Regression

Rank Acceptability Index $RAI(s^j, r, \mathcal{P})$

$RAI(s^j, r, \mathcal{P})$ is a share of compatible preference model instances which assign s^j to the r^{th} rank in $\mathcal{P}$.

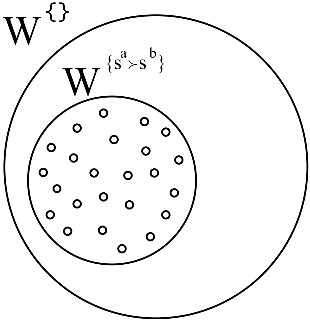

Stochastic Ordinal Regression

We use the stochastic indices to:

- ▷ impose an additional evolutionary pressure during the evolutionary search, i.e., promote these solutions which are – probably – highly preferred to the DM,
- ▷ select solutions to be compared by the DM.

FUN = FRAI

ranks the solutions according to their first rank acceptability indices, i.e., the probability of being the most preferred solution when taking into account a set of compatible preference model instances:

$$f^{FRAI}(s^j, \mathcal{P}) = FRAI(s^j, \mathcal{P}) = RAI(s^j, 1, \mathcal{P}).$$

$FUN = HA$

evaluates the solutions according to their holistic acceptabilities, being defined as weighted sums of rank acceptability indices for all possible ranks; in this regard, we distinguish three different weighting schemes, called linear ($FUN = HA\text{-}LIN$), inverse ($FUN = HA\text{-}INV$), or centroidal ($FUN = HA\text{-}CNT$):

$$f^{HA\text{-}LIN}(s^j, \mathcal{P}) = HA\text{-}LIN(s^j, \mathcal{P}) = \sum_{r=1}^{|\mathcal{P}|} \frac{|\mathcal{P}| - r}{|\mathcal{P}| - 1} \cdot RAI(s^j, r, \mathcal{P}),$$

$$f^{HA\text{-}INV}(s^j, \mathcal{P}) = HA\text{-}INV(s^j, \mathcal{P}) = \sum_{r=1}^{|\mathcal{P}|} \frac{1}{r} \cdot RAI(s^j, r, \mathcal{P}),$$

$$f^{HA\text{-}CNT}(s^j, \mathcal{P}) = HA\text{-}CNT(s^j, \mathcal{P}) = \sum_{r=1}^{|\mathcal{P}|} \frac{\sum_{i=r}^{|\mathcal{P}|} 1/i}{\sum_{k=1}^{|\mathcal{P}|} 1/k} \cdot RAI(s^j, r, \mathcal{P}).$$

$FUN = NFS\text{-}PWI$

derives for each solution a balance between its comprehensive strength and weakness, being defined as the shares of compatible preference model instances for which it is ranked, respectively, better or worse than other solutions:

$$f^{NFS\text{-}PWI}(s^j, \mathcal{P}) = NFS\text{-}PWI(s^j, \mathcal{P}) = \sum_{s^k \in \mathcal{P}, s^j \neq s^k} \left(PWI(s^j, s^k) - PWI(s^k, s^j) \right).$$

$FUN = MD$

ranks the solutions according to their scores for the most discriminative preference model instance, which maximizes the difference in scores for pairs of solutions compared by the DM (case for CF model):

$$f^{MD} = argmax_{d \in \mathcal{S}^d(\mathcal{H})} \{ min \left\{ d(s^k, w, z) - d(s^j, w, z) : (s^j \succ_{DM} s^k) \in \mathcal{H} \right\} \};$$

$FUN = MS$

ranks the solutions according to their scores for the preference model instance, which minimizes – in case of CF – the scores assigned to the pairs of solutions compared by the DM:

$$f^{MS} = argmax_{d \in \mathcal{S}^d(\mathcal{H})} \left\{ \sum_{(s^j \succ_{DM} s^k) \in \mathcal{H}} - \left(d(s^j, w, z) + d(s^k, w, z) \right) \right\}.$$

$REL = SOR\text{-}t$

instantiates the truth of a binary stochastic preference relation for each pair of solutions (s^j, s^k) for which s^j is preferred to s^k for at least $t\%$ of compatible preference model instances, i.e.:

$$SOR\text{-}t : s^j \succ^t_{SOR} s^k \iff PWI(s^j, s^k) \geq t;$$

$REL = PO$

instantiates the truth of a unary relation for each solution that is ranked first for at least one compatible preference model instance derived from the Monte Carlo simulation:

$$PO : PO(s^j) = true \iff RAI(s^j, 1, \mathcal{P}) > 0. \quad (1)$$

Active learning procedures

for selecting pairs of solutions

Traditionally, the interactive evolutionary hybrids select pairs of solutions to be compared by the DM randomly.

We use the results of SOR for choosing a pairwise elicitation question that contributes to the greatest information gain. Since the DM's answer to any preference elicitation question is unknown a priori, the questioning procedures need to aggregate the gains after the two possible answers corresponding to indicating either of the solutions.

Active learning procedures

for selecting pairs of solutions

$AL = DVF$

maximization of the worst case volume of the remaining subspace of preference model instances once the question is answered, which corresponds to the greatest reduction of $\mathcal{S}(\mathcal{H})$ irrespective of the DM's response, i.e.:

$$(s^j, s^k) \leftarrow argmax_{s^j, s^k \in \mathcal{P}} \min\{PWI(s^j, s^k), PWI(s^k, s^j)\};$$

Active learning procedures

for selecting pairs of solutions

$AL = MAX\text{-}PO$

minimization of the worst case number of potentially optimal solutions $\mathcal{F}_1^{PO}(\mathcal{H})$ after answering the question, which corresponds to the greatest reduction of $|\mathcal{F}_1^{PO}(\mathcal{H})|$ irrespective of the DM's response, i.e.:

$$(s^j, s^k) \leftarrow argmin_{s^j, s^k \in \mathcal{P}} \max\{|\mathcal{F}_1^{PO}(\mathcal{H} \cup (s^j, s^k))|, \\ |\mathcal{F}_1^{PO}(\mathcal{H} \cup (s^k, s^j))|\};$$

Active learning procedures

for selecting pairs of solutions

$AL = E\text{-}PO$

minimization of the expected number of potentially optimal solutions $\mathcal{F}_1^{PO}(\mathcal{H})$, when assuming that the probabilities of DM's responses are consistent with the respective *PWI*s (for a detailed justification of this assumption, see, i.e.:

$$(s^j, s^k) \leftarrow argmin_{s^j, s^k \in \mathcal{P}} \Big(PWI(s^j, s^k) \cdot |\mathcal{F}_1^{PO}(\mathcal{H} \cup (s^j, s^k))| + PWI(s^k, s^j) \cdot |\mathcal{F}_1^{PO}(\mathcal{H} \cup (s^k, s^j))| \Big).$$

Summary of the proposed method(s)

EMOSOR

The primary and secondary sorting criteria used by different variants of EMOSOR.

Algorithm	`primary-sort`	`secondary-sort`
EMOSOR-$0_{MODEL\text{-}FUN\text{-}AL}$	$\mathcal{F}^{\succ^\Delta}$	f^{FUN}
EMOSOR-$\mathrm{II}_{MODEL\text{-}REL\text{-}AL}$	$\mathcal{F}^{REL}$	crowding-distance

$FUN \in \{FRAI,\ HA\text{-}LIN,\ HA\text{-}INV,\ HA\text{-}CNT,\ NFS\text{-}PWI,\ MD,\ MS\}$
$REL \in \{SOR\text{-}1.00,\ SOR\text{-}0.85,\ SOR\text{-}0.70,\ PO\}$
$MODEL \in \{CF, AVF\}$
$AL \in \{RAND, DVF, MAX\text{-}PO, E\text{-}PO\}$

NEMO

The primary and secondary sorting criteria used by NEMO methods

Algorithm	`primary-sort`	`secondary-sort`
NEMO-0	$\mathcal{F}^{\succ^\Delta}$	a representative function
NEMO-II	$\mathcal{F}^{PO}$	crowding-distance

Experimental setting

1. Pregenerate 100 artificial DM's
   - using either a CF or an AVF
   - for each model instance, we find an optimal solution

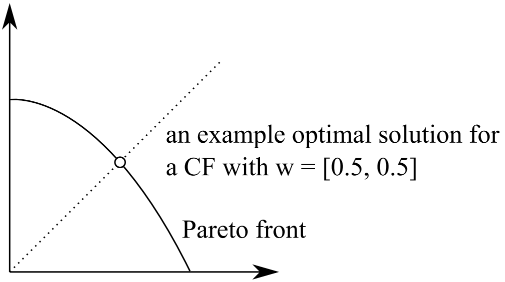

Experimental setting

1. Pregenerate 100 artificial DM's
   - using either a CF or an AVF
   - for each model instance, we find an optimal solution (benchmark)
2. Each method was run 100 times, each time interacting with a different DM

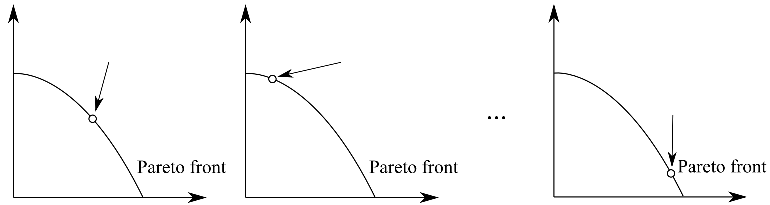

Experimental setting

1. Pregenerate 100 artificial DM's
   - using either a CF or an AVF
   - for each model instance, we find an optimal solution (benchmark)
2. Each method was run 100 times, each time interacting with a different DM
3. To assess the performance, we computed relative score differences:

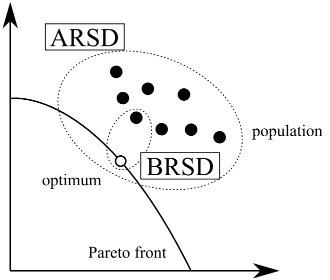

Experimental setting

1. Pregenerate 100 artificial DM's
    - using either a CF or an AVF
    - for each model instance, we find an optimal solution (benchmark)
2. Each method was run 100 times, each time interacting with a different DM
3. To assess the performance, we computed relative score differences
4. During the evolutionary run, each method
    - interacted with the DM 10 times
    - the solutions to be compared were selected either randomly (benchmark) or using indications derived from the stochastic analysis

We compared all variants of EMOSOR and we find out which sorting criteria derived from Stochastic Ordinal Regression are the most (least) advantageous in terms of the performance of interactive evolutionary optimization algorithms.

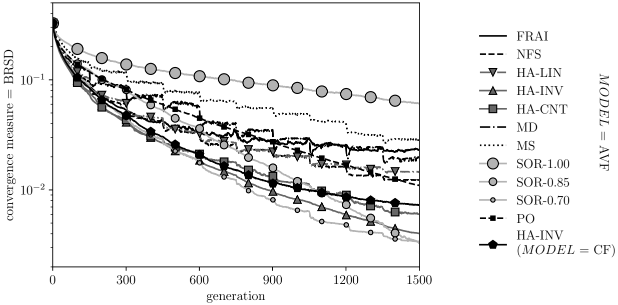

Figure: Averaged BRSD throughout evolutionary search attained by different variants of EMOSOR$_{MODEL=AVF}$ applied to WFG1 and $DM = WS$.

Performed Experiments

Comparison of different variants of EMOSOR

Inconsistency analysis

We compared all variants of EMOSOR and we find out whether the (in)compatibility between the assumed preference models and the decision making model influence the performance of the algorithms.

Table: The average numbers of pairwise comparisons removed throughout the evolutionary search to reinstate consistency by the variants of EMOSOR with either $MODEL = AVF$ or $MODEL = CF$ for the DLTZ(C)2 and WFG1 problems with different numbers of objectives M and various models (DM) of the simulated Decision Maker.

		$DM = WS$			$DM = CF$		
M	$MODEL$	Max	Mean	StD	Max	Mean	StD
2	AVF	0.00	0.00	0.00	5.09	1.12	1.23
	CF	2.59	0.18	0.46	0.00	0.00	0.00
3	AVF	0.00	0.00	0.00	3.71	0.47	0.76
	CF	1.77	0.11	0.30	0.00	0.00	0.00
4	AVF	0.00	0.00	0.00	2.86	0.29	0.54
	CF	1.55	0.07	0.24	0.00	0.00	0.00
5	AVF	0.00	0.00	0.00	2.45	0.30	0.50
	CF	1.00	0.04	0.15	0.00	0.00	0.00

Performed Experiments

Comparison of EMOSOR and NEMO methods

We selected 6 methods for the comparison, which differed in terms of the incorporated preference model and the way of exploiting this model to direct the evolutionary search towards preferred region in the objective space.

Category	Based on AVF	Based on CF
Based on the representative function	NEMO-0	EMOSOR-0_{CF-MD}
Based on the fronts of potential optimality	NEMO-II	EMOSOR-II_{CF-PO}
Based on the holistic acceptabilities *HA-INV*	EMOSOR-$0_{AVF\text{-}HA\text{-}INV}$	EMOSOR-$0_{CF\text{-}HA\text{-}INV}$

Performed Experiments

Comparison of EMOSOR and NEMO methods

We showed that the performance of a prefernce-based EMOA may be improved when:
– SOR is incorporated,
– the incorporated preference model is in alignment with the DM's decision policy.

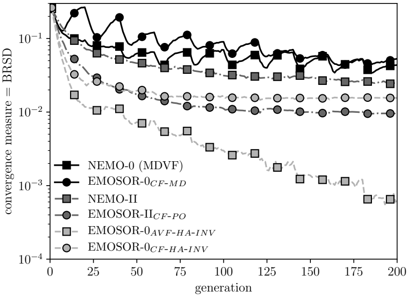

Figure: Averaged BRSD throughout evolutionary search attained by different methods applied to DTLZ2C and $DM = WS$.

Performed Experiments

Visualization of convergence

Solutions constructed by EMOROR_{CF-PO} after 20, 50, 100, and 200 generations, when applied to DTLZ2 with $M = 3$ and $w^{DM} = [1/3, 1/3, 1/3]$.

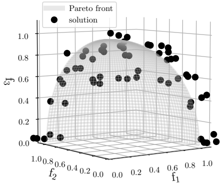

Figure: 20^{th} generation

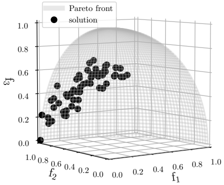

Figure: 50^{th} generation

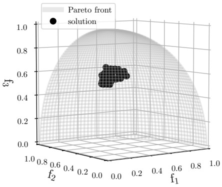

Figure: 100^{th} generation

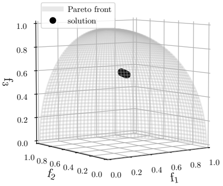

Figure: 200^{th} generation

Performed Experiments

EMOSOR with different active learning procedures

We showed that the performance of an interactive EMOA can be improved in terms of the quality of generated solutions and the required number of interactions with the DM as well.

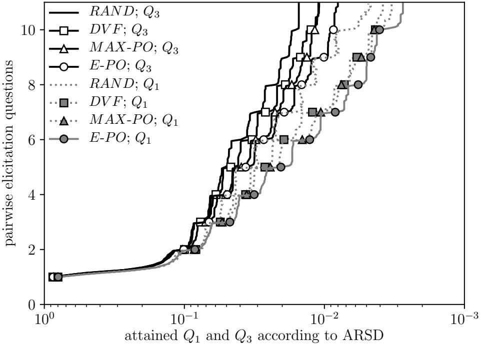

Figure: The quality of solutions (Q_1 and Q_3 for ARSD) constructed by EMOSOR-0$_{AVF\text{-}HA\text{-}INV}$ with $AL \in \{RAND, DVF, MAX\text{-}PO, E\text{-}PO\}$ after different numbers of interactions with the DM.

Conclusions

▷ We proposed a novel interactive preference-based EMOA[3], called EMOSOR, based on a stochastic ordinal regression.

▷ The proposed method uses the indications derived from the stochastic analysis to:

drive a population of solutions towards a highly preferred region of the PF, select a pair of solutions to be compared by the DM in order to maximize the information gain of the received answer.

▷ We evaluated the proposed method on a large number of benchmark problems and we:

▷ determined which sorting criteria are the most advantageous in the course of the evolutionary search,
▷ we performed the inconsistency analysis, showing that the performance of the interactive EMOA may be improved when the preference model used by the methods aligns with the DM's decision policy,
▷ we compared EMOSOR with some existing state-of-the art EMOAs, proving its competitiveness,
▷ evaluated EMOSOR with difference active learning procedures, showing that the total number of interactions with the DM may be reduced when suitably selecting paris of solutions to be compared.

[3] M. Tomczyk, M. Kadziński, "EMOSOR: Evolutionary multiple objective optimization guided by interactive stochastic ordinal regression," Computers & Operations Research, vol. 108, pp. 134-154, 2019.