MULTI-OBJECTIVE OPTIMIZATION

Pareto dominance, NSGA-II, MOEAs for RL

Dennis G. Wilson

Part of the 2021 RLVS.

MULTI-OBJECTIVE OPTIMIZATION

Optimizing **more than one** objective function **simultaneously**.

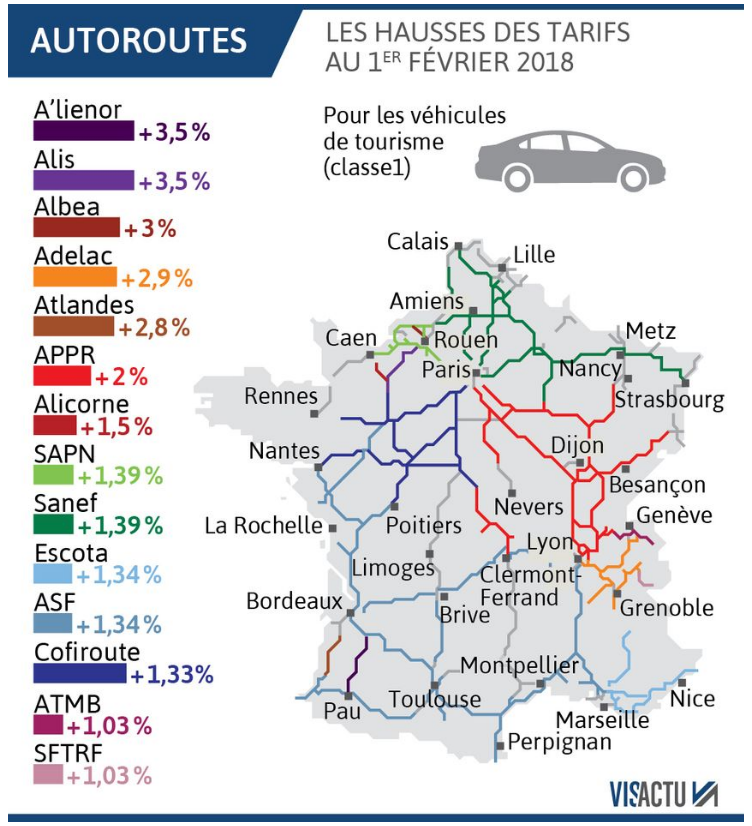

For example, when planning a trip, we want to minimize total distance travelled and toll fare.

MOEAS

MULTI-OBJECTIVE EVOLUTIONARY ALGORITHMS

- **NSGA**: Srinivas, Nidamarthi, and Kalyanmoy Deb. "Muiltiobjective optimization using nondominated sorting in genetic algorithms." Evolutionary computation 2.3 (1994): 221-248.
- **SPEA2**: Zitzler, Eckart, Marco Laumanns, and Lothar Thiele. "SPEA2: Improving the strength Pareto evolutionary algorithm." TIK-report 103 (2001).
- **NSGA-II**: Deb, Kalyanmoy, et al. "A fast and elitist multiobjective genetic algorithm: NSGA-II." IEEE transactions on evolutionary computation 6.2 (2002): 182-197.
- Deb, Kalyanmoy (2001) Multi-objective optimization using evolutionary algorithms. John-Wiley, Chichester
- **MOEA/D**: Zhang, Qingfu, and Hui Li. "MOEA/D: A multiobjective evolutionary algorithm based on decomposition." IEEE Transactions on evolutionary computation 11.6 (2007): 712-731.
- Emmerich, Michael TM, and André H. Deutz. "A tutorial on multiobjective optimization: fundamentals and evolutionary methods." Natural computing 17.3 (2018): 585-609. pdf

PARETO DOMINANCE

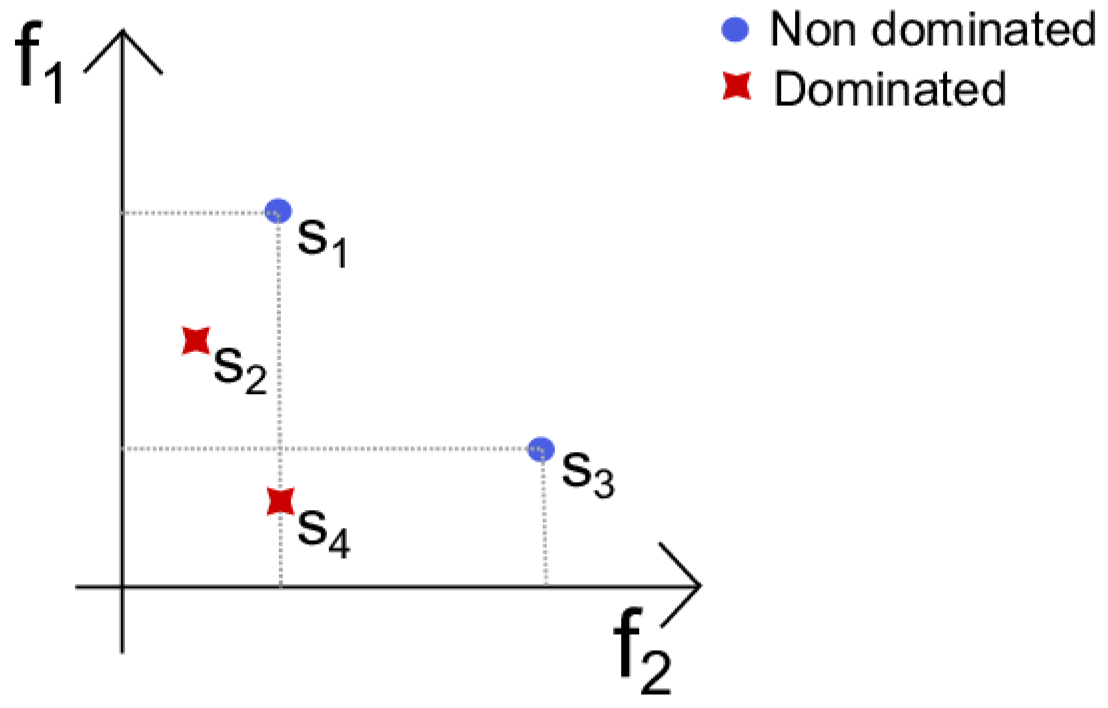

A solution is said to Pareto dominate another if it is more optimal in all dimensions.

Solutions which are not dominated by any other are called "non-dominated".

PARETO FRONT

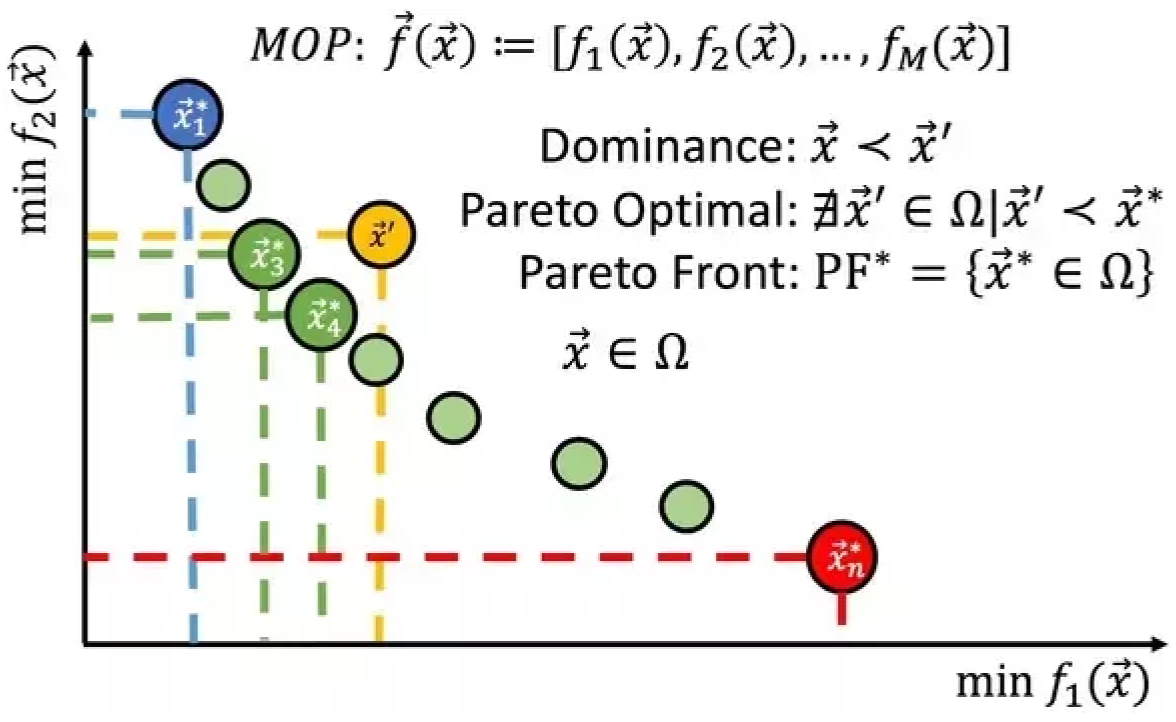

The Pareto Front is the set of Pareto Optimal solutions.

In Multi-Objective Optimization, we will search for the Pareto Front.

NSGA-II

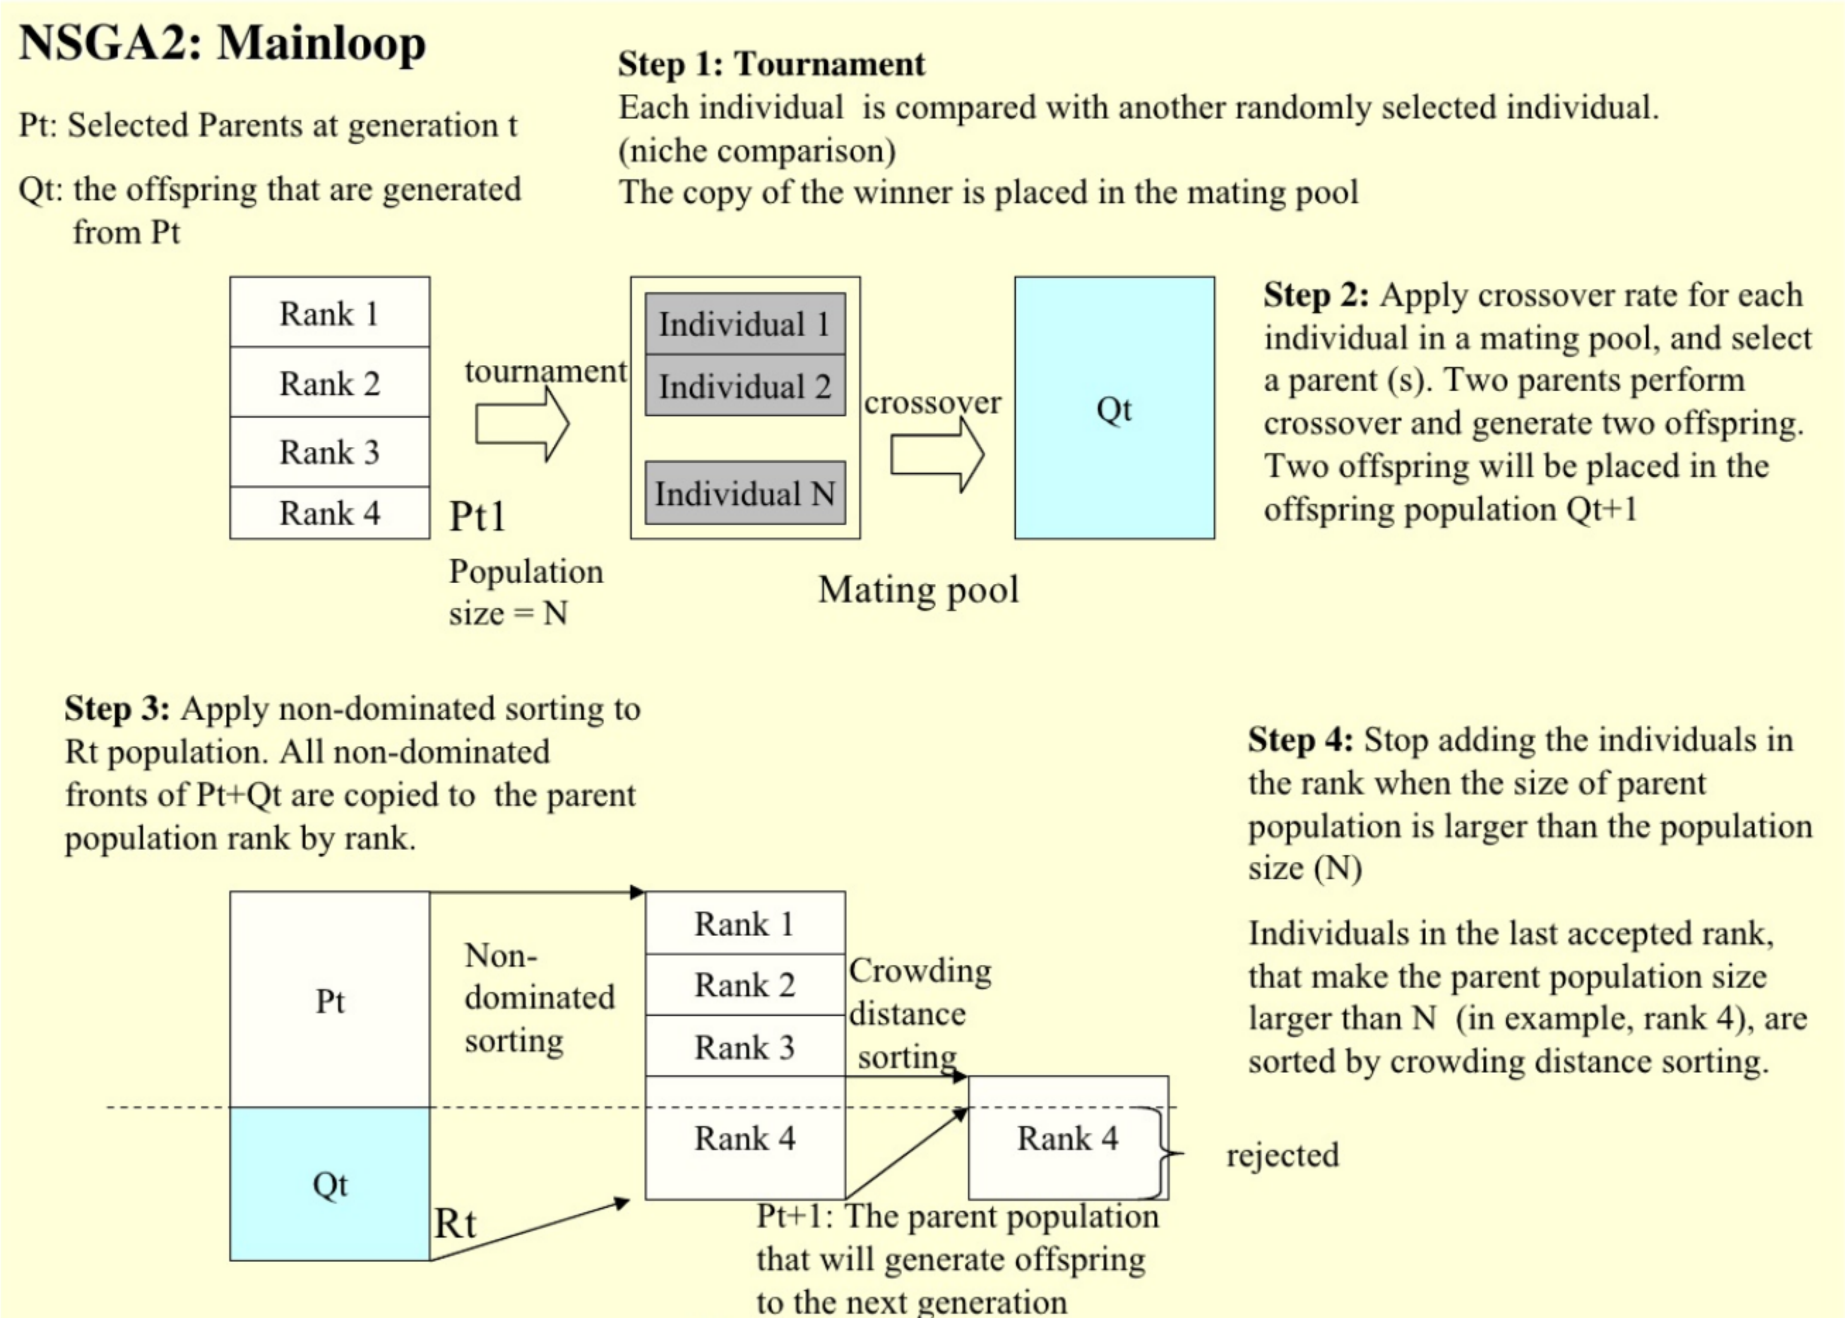

Deb, Kalyanmoy, et al. "A fast and elitist multiobjective genetic algorithm: NSGA-II." IEEE transactions on evolutionary computation 6.2 (2002): 182-197. pdf

NON-DOMINATED SORTING

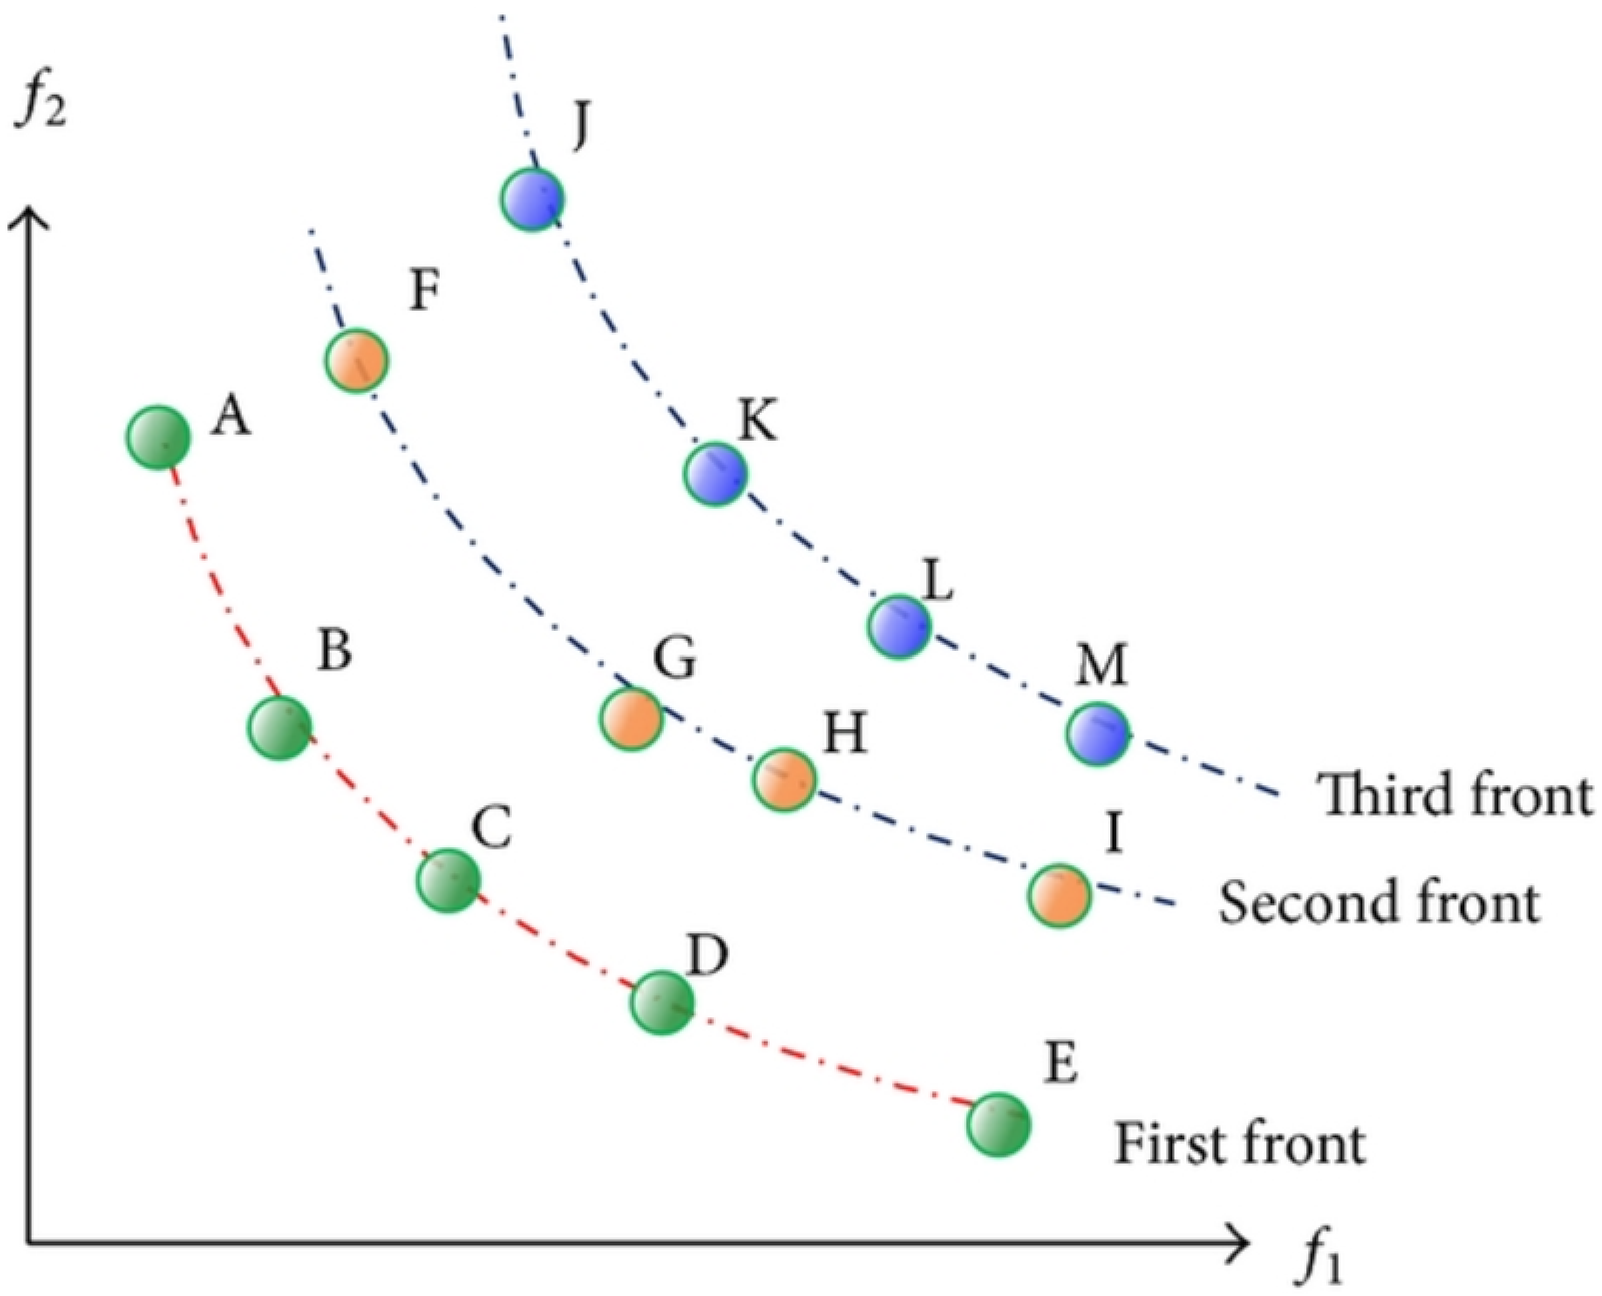

Wang, H. S., C. H. Tu, and K. H. Chen. "Supplier selection and production planning by using guided genetic algorithm and dynamic nondominated sorting genetic algorithm II approaches." Mathematical Problems in Engineering 2015 (2015).

FAST NON-DOMINATED SORT

`fast-nondominated-sort` (P)	
for each $p \in P$	
for each $q \in P$	
if $(p \prec q)$ then	if p dominates q then
$S_p = S_p \cup \{q\}$	include q in S_p
else if $(q \prec p)$ then	if p is dominated by q then
$n_p = n_p + 1$	increment n_p
if $n_p = 0$ then	if no solution dominates p then
$\mathcal{F}_1 = \mathcal{F}_1 \cup \{p\}$	p is a member of the first front
$i = 1$	
while $\mathcal{F}_i \neq \emptyset$	
$\mathcal{H} = \emptyset$	
for each $p \in \mathcal{F}_i$	for each member p in $\mathcal{F}_i$
for each $q \in S_p$	modify each member from the set S_p
$n_q = n_q - 1$	decrement n_q by one
if $n_q = 0$ then $\mathcal{H} = \mathcal{H} \cup \{q\}$	if n_q is zero, q is a member of a list $\mathcal{H}$
$i = i + 1$	
$\mathcal{F}_i = \mathcal{H}$	current front is formed with all members of $\mathcal{H}$

Deb, Kalyanmoy, et al. "A fast and elitist multiobjective genetic algorithm: NSGA-II." IEEE transactions on evolutionary computation 6.2 (2002): 182-197. pdf

CROWDING DISTANCE ASSIGNMENT

`crowding-distance-assignment` $(\mathcal{I})$

$l = \lvert\mathcal{I}\rvert$	number of solutions in $\mathcal{I}$
for each i, set $\mathcal{I}[i]_{distance} = 0$	initialize distance
for each objective m	
$\quad \mathcal{I} = \text{sort}(\mathcal{I}, m)$	sort using each objective value
$\quad \mathcal{I}[1]_{distance} = \mathcal{I}[l]_{distance} = \infty$	so that boundary points are always selected
$\quad$ for $i = 2$ to $(l - 1)$	for all other points
$\quad\quad \mathcal{I}[i]_{distance} = \mathcal{I}[i]_{distance} + (\mathcal{I}[i+1].m - \mathcal{I}[i-1].m)$	

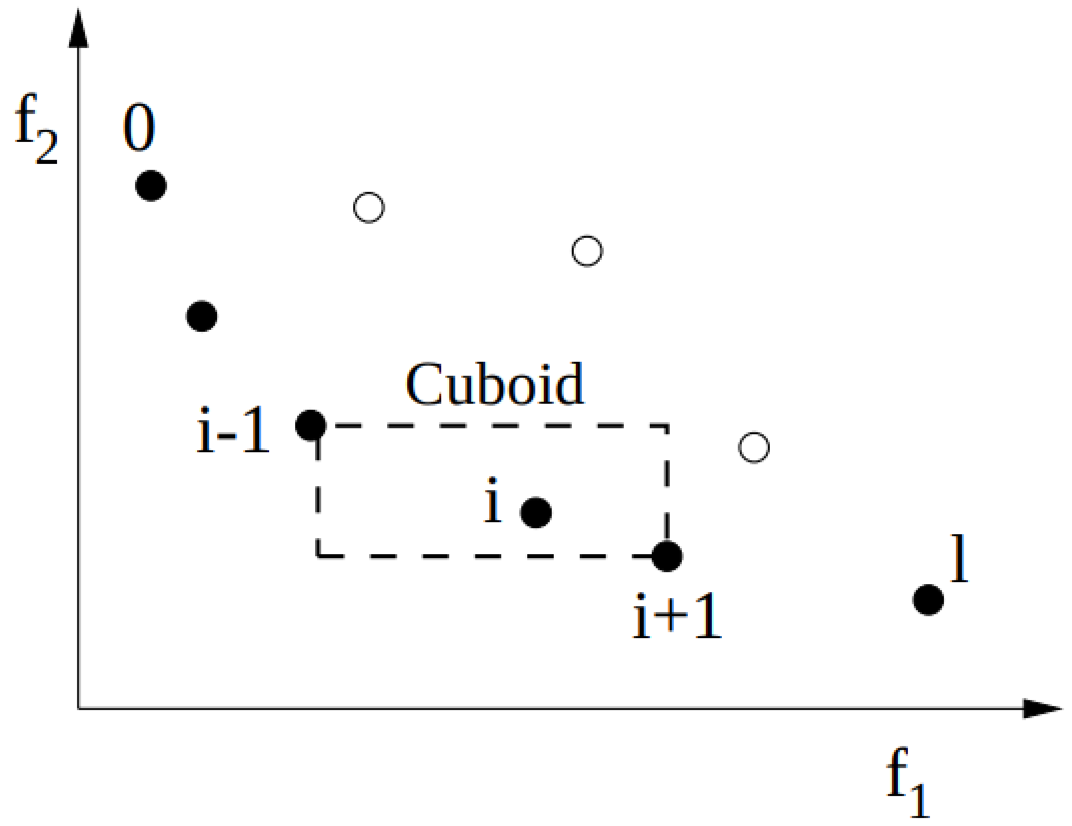

Deb, Kalyanmoy, et al. "A fast and elitist multiobjective genetic algorithm: NSGA-II." IEEE transactions on evolutionary computation 6.2 (2002): 182-197. pdf

NSGA-II OVERVIEW

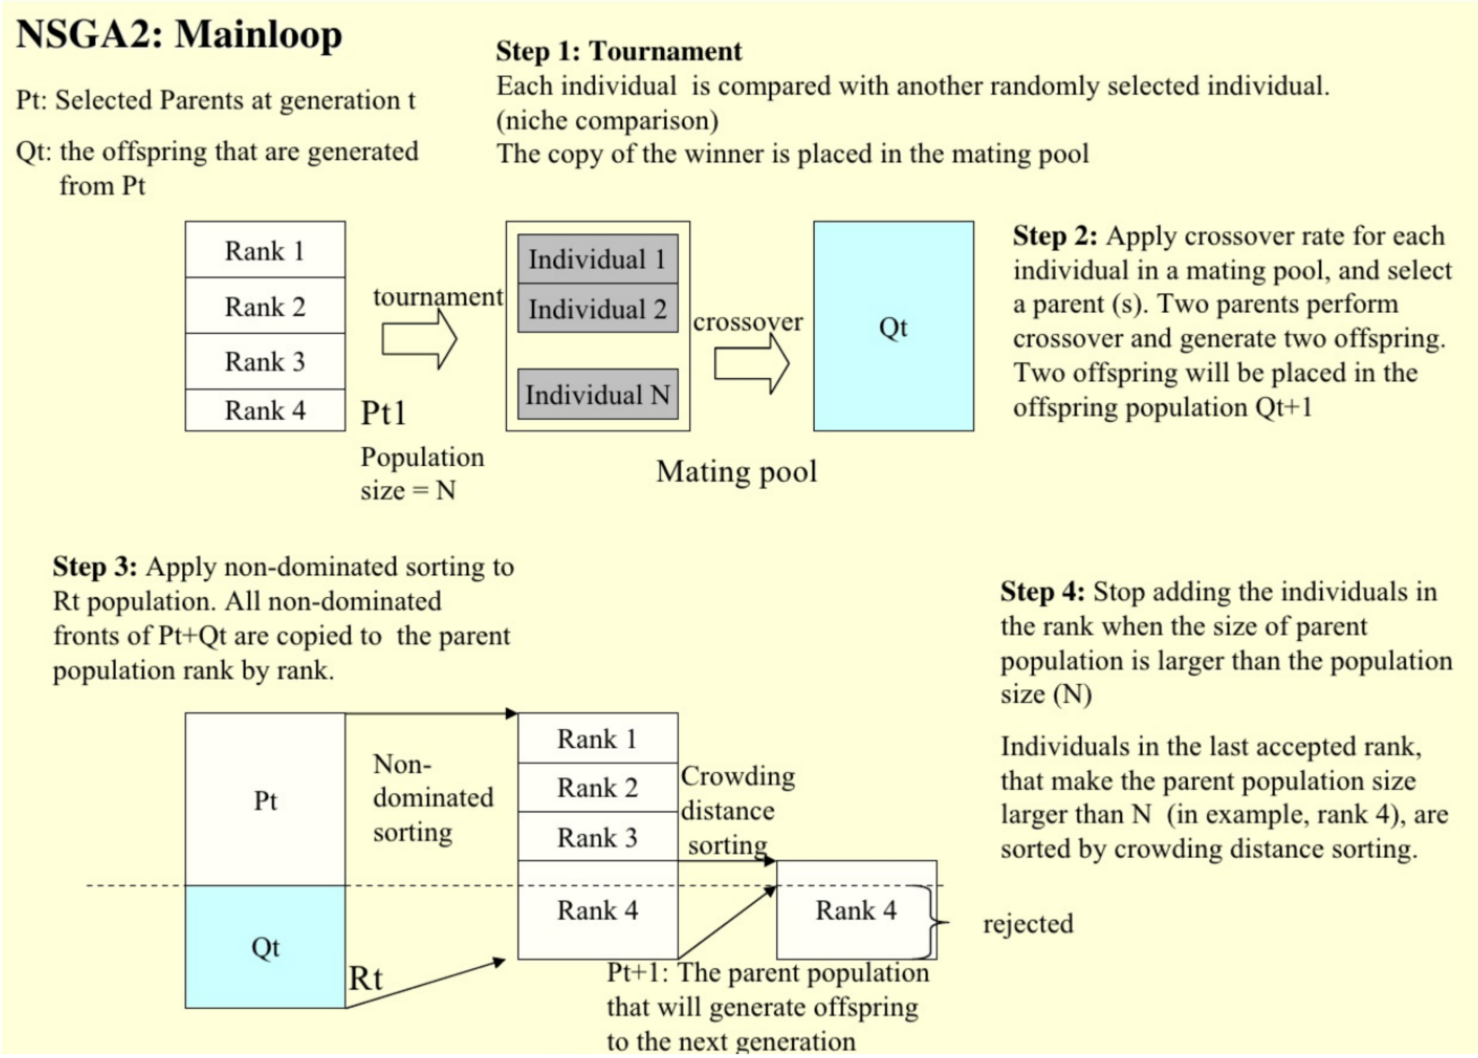

$$i \geq_n j \quad \text{if } (i_{rank} < j_{rank}) \text{ or } ((i_{rank} = j_{rank}) \text{ and } (i_{distance} > j_{distance}))$$

Deb, Kalyanmoy, et al. "A fast and elitist multiobjective genetic algorithm: NSGA-II." IEEE transactions on evolutionary computation 6.2 (2002): 182-197. pdf

PROBLEMS WITH NON-DOMINANCE

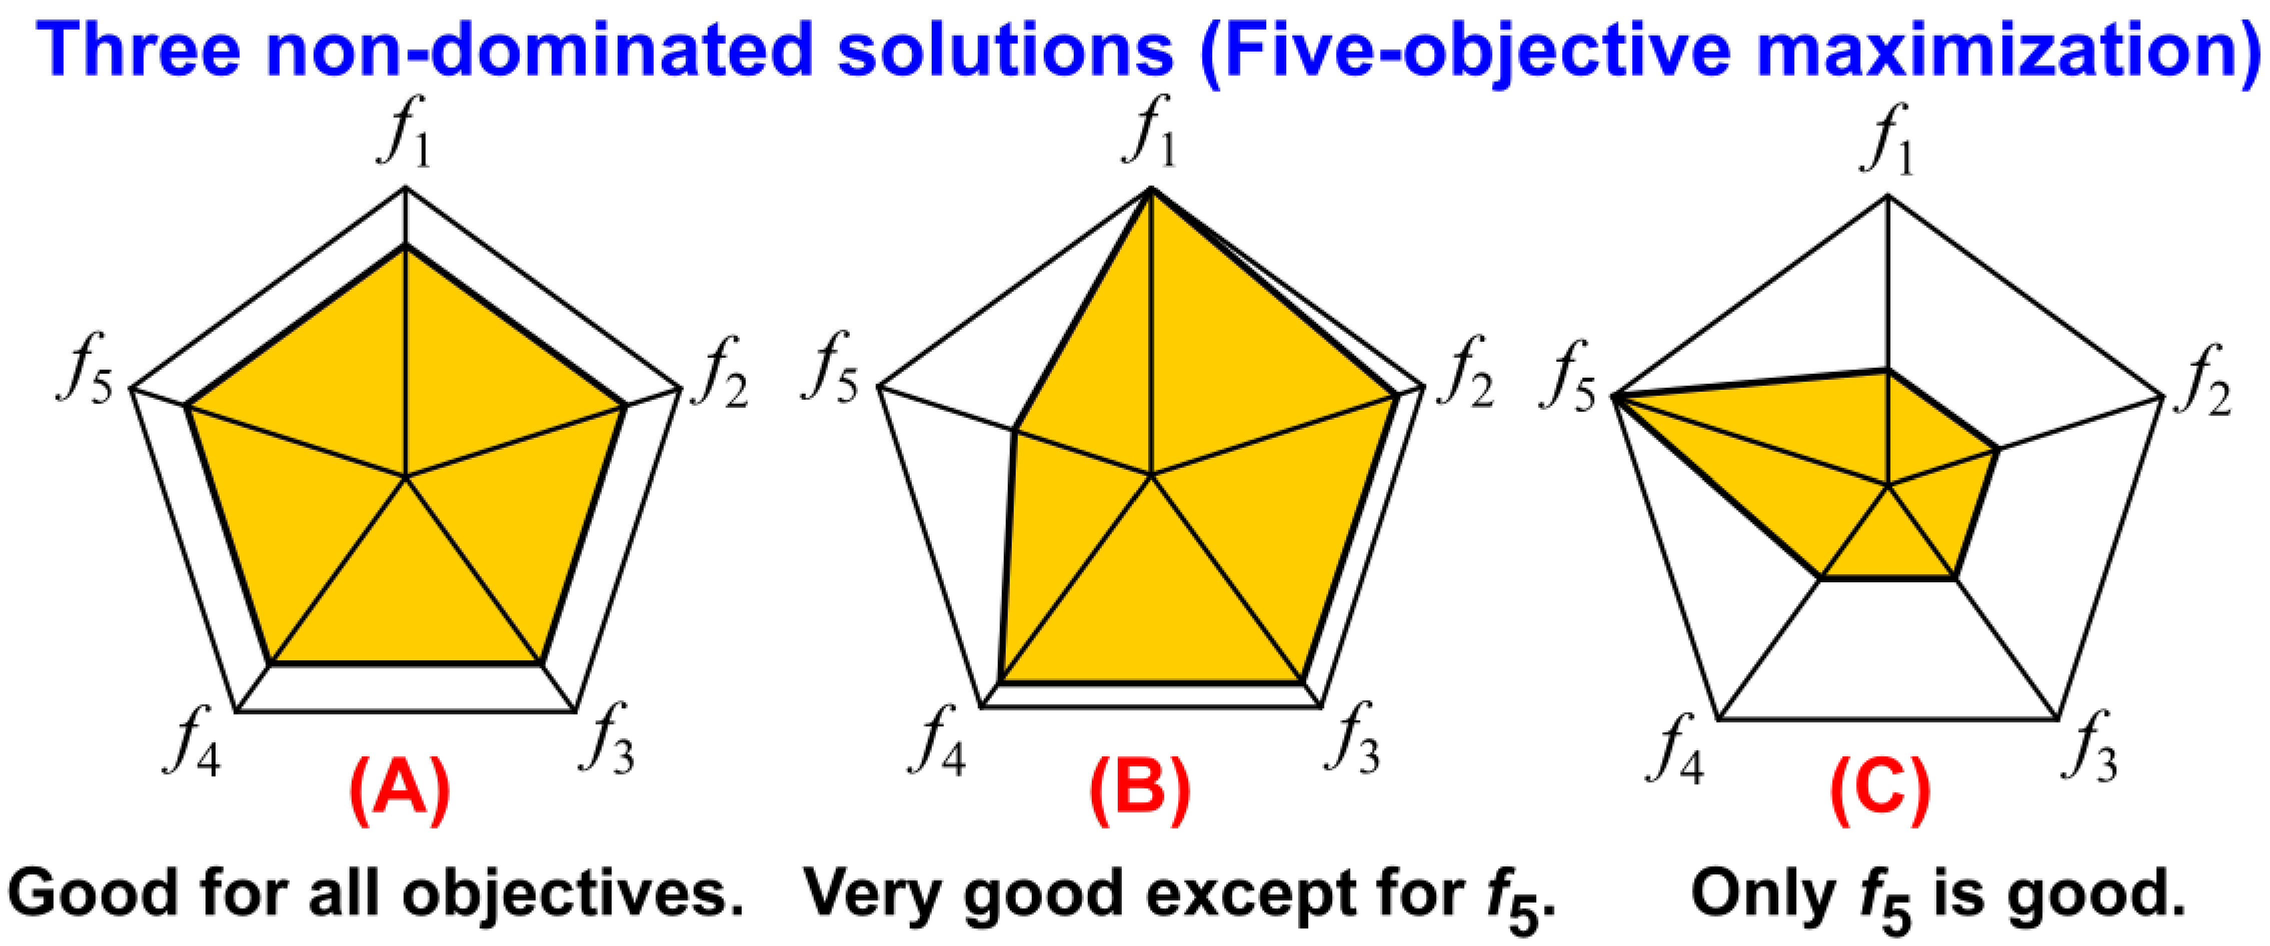

With more objectives, some objectives may be overrepresented in the non-dominated set.

Ishibuchi, Hisao, and Hiroyuki Sato. "Evolutionary many-objective optimization." Proceedings of the Genetic and Evolutionary Computation Conference Companion. 2019.

MANY-OBJECTIVE OPTIMIZATION

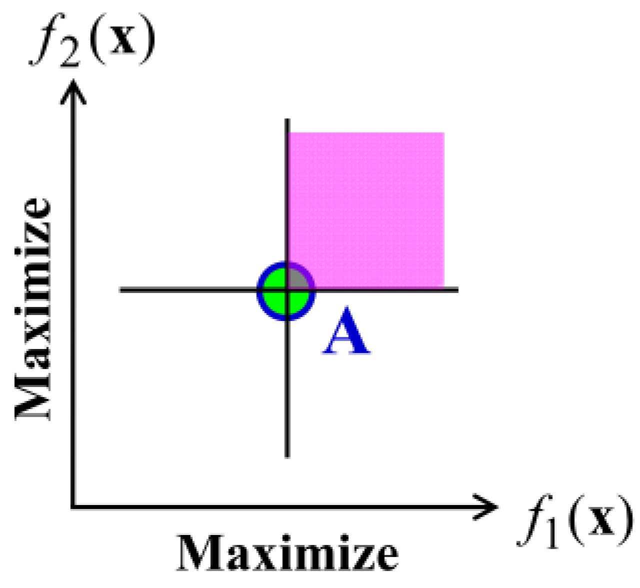

Percentage of the better region

2 objectives	1/4	25%
3 objectives	1/8	13%
4 objectives	1/16	6%
5 objectives	1/32	3%
10 objectives	1/1024	0.1%
15 objectives	1/32768	0.003%
20 objectives	1/1048576	0.0001%

When increasing beyond a small number (2-4) of objectives,

the chance of fully non-dominated solutions decreases.

Different algorithms, visualization methods, convergence metrics are needed.

Ishibuchi, Hisao, and Hiroyuki Sato. "Evolutionary many-objective optimization." Proceedings of the Genetic and Evolutionary Computation Conference Companion. 2019.

MOEAS FOR RL

An advantage of evolutionary RL:
multiple solutions along the Pareto front

Primary objective:
maximize total reward

secondary objective:
optimize different behaviors (use of specific robotic limb), increase efficiency.

Example: efficient robot locomotion

Xu, J., Tian, Y., Ma, P., Rus, D., Sueda, S., & Matusik, W. (2020, November). Prediction-Guided Multi-Objective Reinforcement Learning for Continuous Robot Control. In International Conference on Machine Learning (pp. 10607-10616). PMLR. http://pgmorl.csail.mit.edu/