NDBI021, Lecture 9

User preferences, 2/1 ZK+Z,

Wed 12:20 – 13:50 S8

Wed 14:00 – 15:30 SW2 (odd weeks)

https://www.ksi.mff.cuni.cz/~peska/vyuka/ndbi021/2022/

matfyz

siret research group

https://ksi.mff.cuni.cz

Multi-criteriality in RS

matfyz

siret research group

https://ksi.mff.cuni.cz

Tutorial:
Multi-Objective Recommendations

Yong Zheng, Illinois Institute of Technology, USA
David (Xuejun) Wang, Morningstar, Inc., USA

Virtual Conference
KDD2021
August 14th – 18th

Different Types of Recommender Systems

Context-Aware RecSys

Incorporate context info (time, location, etc) into RecSys

Group RecSys

Recommend items to a group of users, e.g., group dinning

Multi-Stakeholder RecSys

Produce recommendations by considering multiple stakeholders, e.g., buyers and sellers on eBay

Multi-Task RecSys

Build joint learning model by considering multiple tasks, e.g., RecSys + opinion texts

Why MOO in RecSys

- There is an emerging demand in MOO
  - Traditional RecSys
    - Example: RecSys balancing multiple metrics, e.g., news

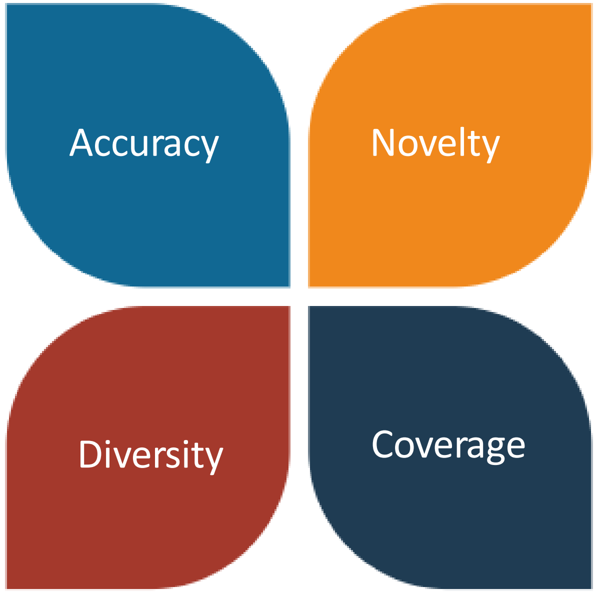

Why MOO in RecSys

- There is an emerging demand in MOO
  - New Types of RecSys
    - Example: Multi-stakeholder RecSys, e.g., marketplace

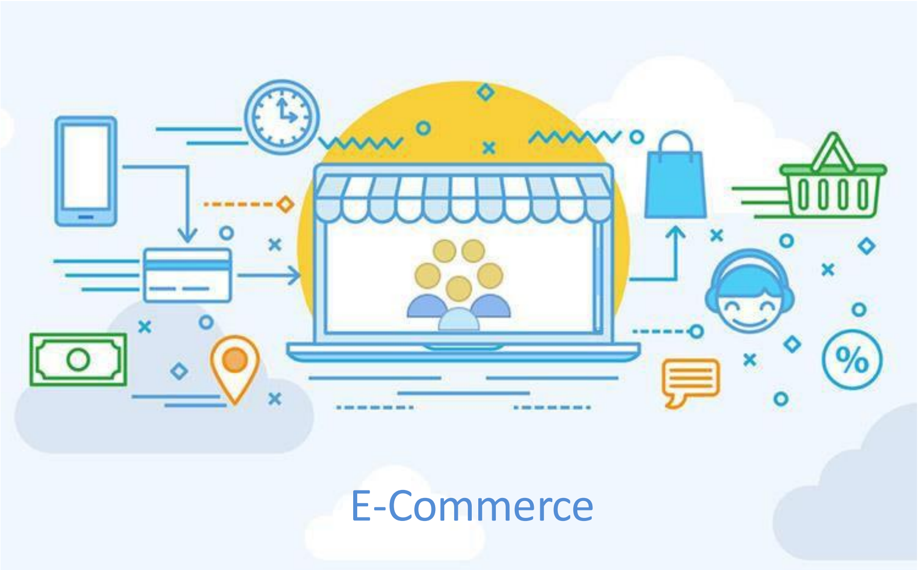

RecSys with MOO

- Contexts in which we need MOO in RecSys

- RecSys Balancing Multiple Metrics
- Group RecSys
- User-Based Collaborative Filtering
- Multi-Stakeholder RecSys
- Preprocessing for RecSys
- Multi-Task RecSys

Contents

Background and Some History

- Multi Objective Optimization in Finance

Financial Planning

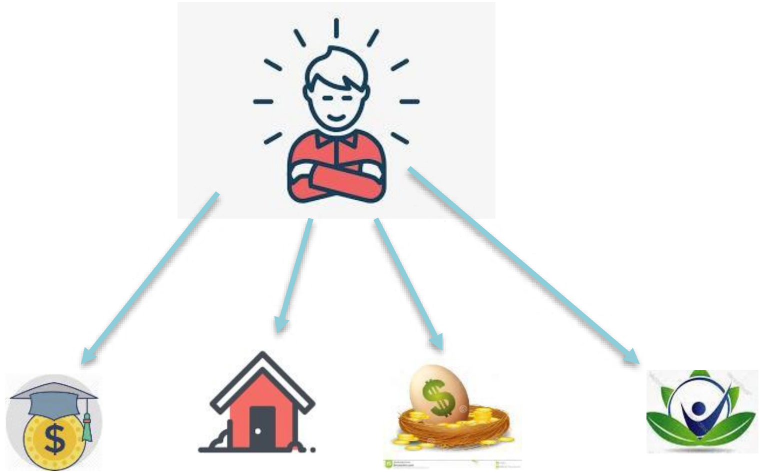

education house retiremen t Life quality

Portfolio Optimization

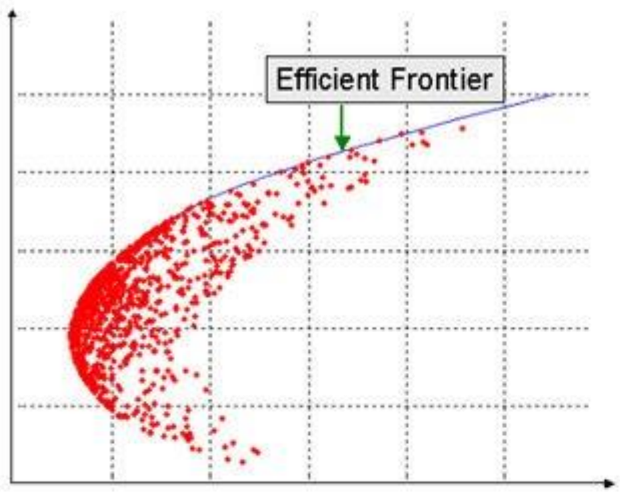

Expected Return

Risk

Background and Some History

- Multi Metrics in Recommendation Systems
  - Goal: Meets user's need
  - Objectives:
    - Maximize Accuracy
    - Maximize Diversity
    - Maximize Novelty
  - Challenge
    - Increase Diversity may decease Accurac
    - Increase Novelty may decease Accuracy

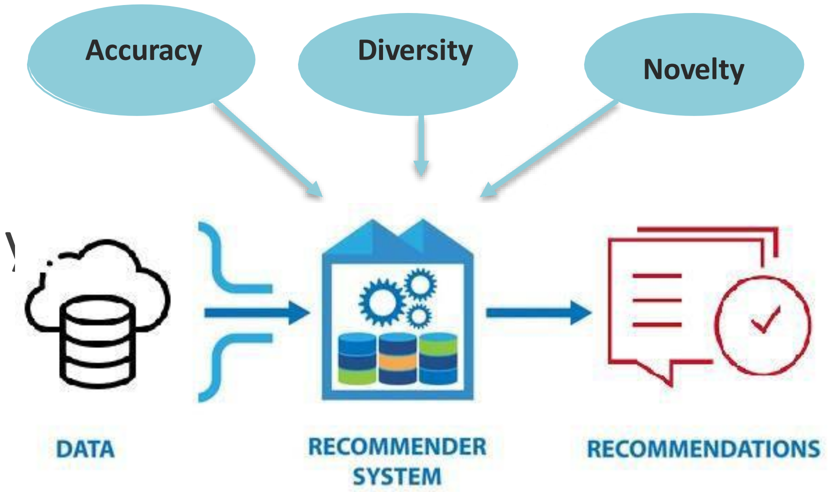

Background and Some History

- Multi Stakeholder Recommendation Systems
  - Goal: Meet interests of all stakeholders
  - Objectives: Maximize Three Item Utilities
    - In respect of End User
    - In respect of Provider
    - In respect of Platform Owner
  - Challenge
    - Utilities regarding to three stakeholders may conflict each other

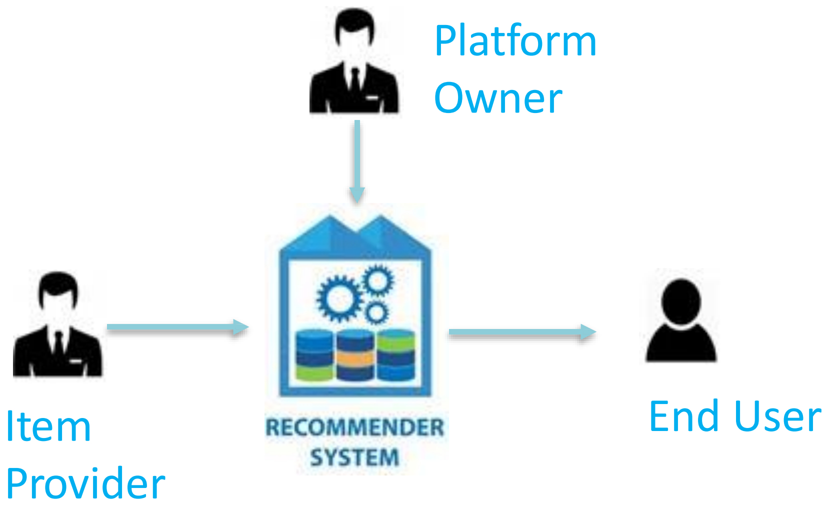

Background and Some History

- Francis Ysidro Edgeworth (1845-1926)

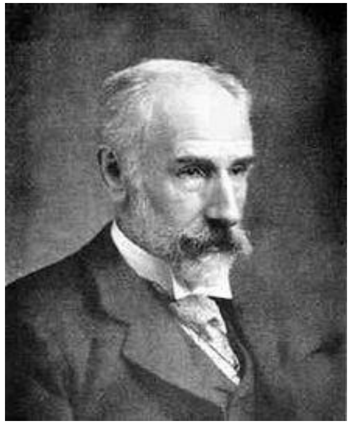

  – ***Mathematical Psychics**: **An Essay on the Application of Mathematics to the Moral Sciences***, published in 1881

  – "*It is required to find a point (x, y) such that, in whatever direction we take an infinitely small step, P and Π do not increase together, but that, while one increases, the other decreases*"

- Vilfredo Pareto (1848-1923)

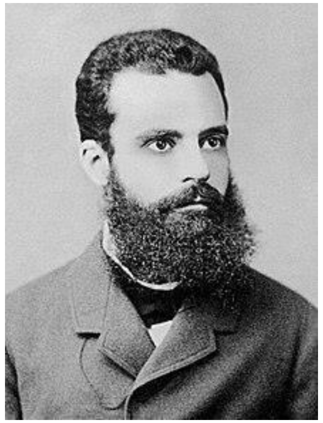

  – ***Manual of Political Economy***, published in 1906

  – "*The optimum allocation of the resources of a society is not attained so long as it is possible to make at least one individual better off in his own estimation while keeping others as well off as before in their own estimation.*"

Contents

Multi Objective Optimization (MOO)

- Multi Objective Optimization (MOO) Problem

$$\min_{\boldsymbol{x}}(f_1(\boldsymbol{x}), f_2(\boldsymbol{x}), \ldots, f_M(\boldsymbol{x}))$$

Objectives

Subject to:

$$\begin{aligned} g_j(\boldsymbol{x}) &\geq 0, & j &= 1, 2, \ldots, J \\ h_k(\boldsymbol{x}) &= 0, & k &= 1,2, \ldots, K \\ x_i^L \leq x_i &\leq x_i^U, & i &= 1, 2, \ldots, n \end{aligned}$$

Possible constraints format

Decision variable: $\boldsymbol{x} \in \boldsymbol{R}^n$

Objective Functions: $f_i,\ i = 1,2, \ldots, M$

Feasible Solutions: S

$$S = \{x \mid x_i^L \leq x_i \leq x_i^U, g_j(x) \geq 0, h_k(x) = 0, j = 1, 2, \ldots, J, k = 1,2, \ldots, K, i = 1, \ldots, n\}$$

Multi Objective Optimization (MOO)

- Example: Two Objectives in Recommender Systems

Name	Symbol	Meaning
Decision Variable	x	Top N recommendation list
Feasible Solution Set	S	All top N recommendation list
First Objective	$f_1(x)$	1 – accuracy
Second Objective	$f_2(x)$	1 – diversity

- Find recommendation list that maximize accuracy and diversity

$$\min_{x \in} (f_1, f_2) \quad \text{or} \quad \min_{x \in S} F(x), \text{ where } F(x) = (f_1, f_2)$$

Multi Objective Optimization (MOO)

- Special Characters of MOO
  - Objectives may be conflict each other
  - Cannot determine which solution is better
  - Example:

$$\min_{x}(f_1, f_2)$$

Where $f_1(x) = 2(x - 1) + 1,$

$$f_2(x) = 2(x-3)^2 + 1$$

Subject $x \in [0, 6]$

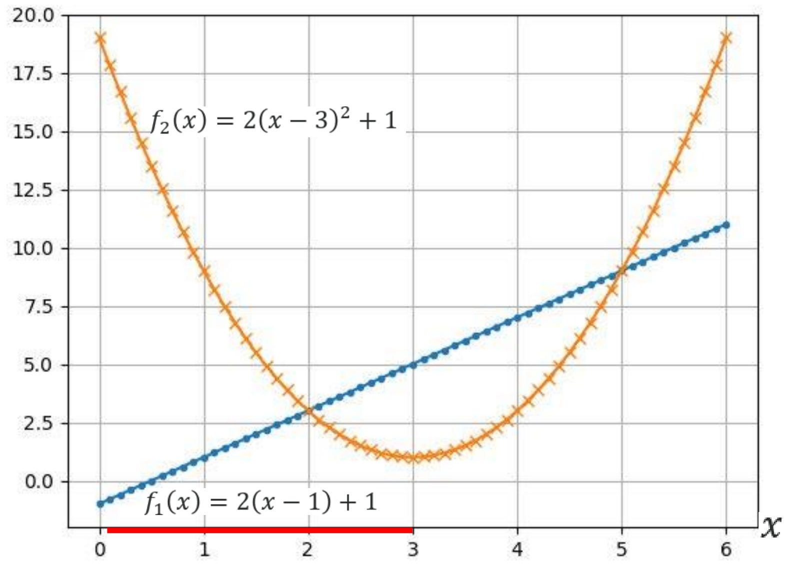

Multi Objective Optimization (MOO)

- Dominance Relation

A solution x is said to be Dominated by x^* if and only if

$$\min_x(f_1, f_2)$$

$$f_m(x^*) \leq f_m(x) \text{ for all } m = 1,2,\dots,M$$

and there exists at least one m' such that:

$$f_{m'}(x^*) < f_{m'}(x)$$

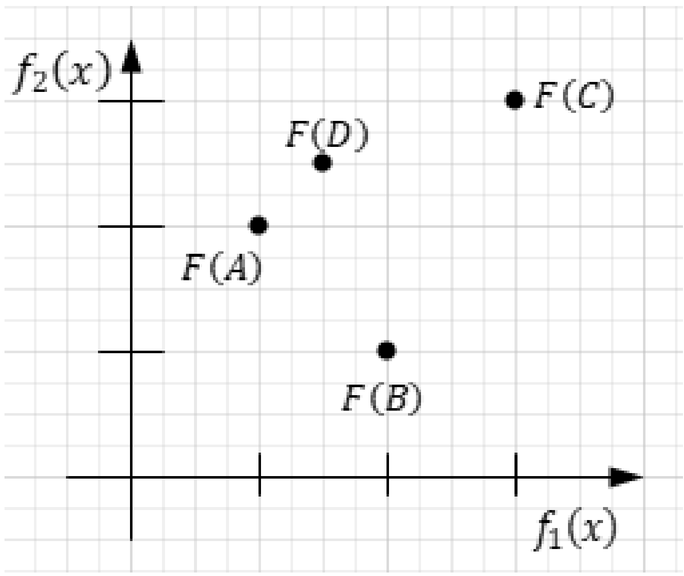

A and B dominate C, D is only dominated by A.

A and B: no dominance relationship

D and B: no dominance relationship

Multi Objective Optimization (MOO)

- Non-Dominated Solution (Pareto Optimal Solution)
  - Not dominated by any other solutions
  - Solution A, B, E and G are Pareto Optimal
- Pareto Optimal Set:
  - All x such that $F(x)$ is on curve from $F(E)$ to $F(G)$
- Pareto Front:
  - All $F(x)$ on curve from $F(E)$ to $F(G)$

$$\min_x (f_1, f_2)$$

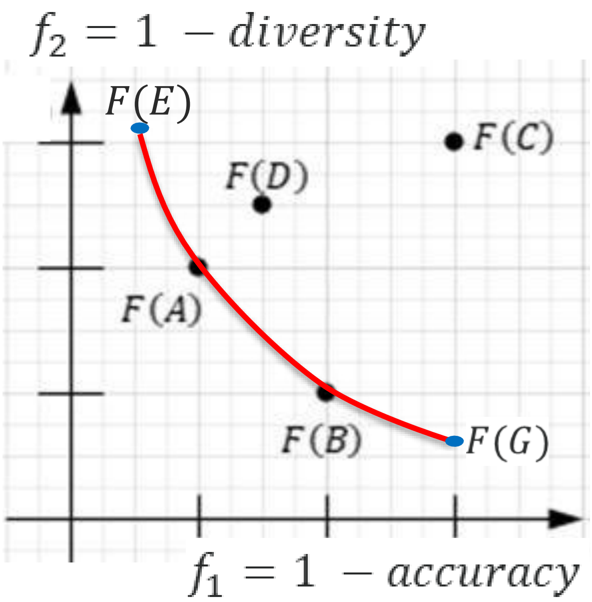

Multi Objective Optimization (MOO)

- Example:

$$\min_{x}(f_1, f_2)$$

Where $f_1(x) = 2(x-1)+1,$

$$f_2(x) = 2(x-3)^2 + 1$$

Subject $x \in [0, 6]$

- Analysis
  - Feasible solutions: $S = [0,6]$
  - Pareto Set: $\{x \mid x \in [0,3]\}$
  - Pareto Front: $\{(f_1, f_2) | x \in [0, 3]\}$

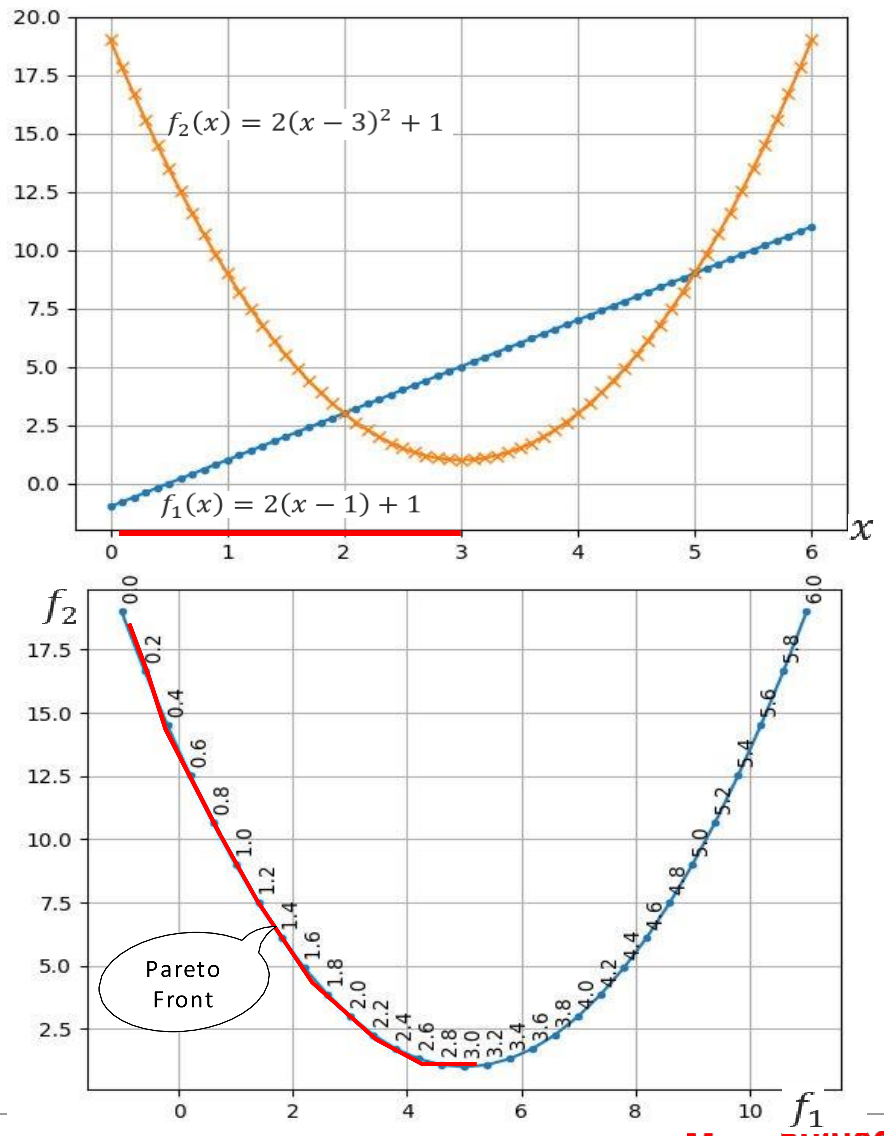

Multi Objective Optimization (MOO)

- Solving Multi Objective Optimization (MOO) Problem

$$\min_{\boldsymbol{x} \in S} F(\boldsymbol{x})$$

where $F(\boldsymbol{x}) = (f_1(\boldsymbol{x}), f_2(\boldsymbol{x}), \ldots, f_M(\boldsymbol{x}))$

$\boldsymbol{x} \in S$

S is set of all feasible solutions

- Outputs
  - Find a Non-Dominated Solution
  - Find All Non-Dominated Solutions (Pareto Set)
  - Find a representative subset of Non-Dominated Solutions

MORNINGSTAR®

Multi Objective Optimization (MOO)

- MOO Decision Making Process

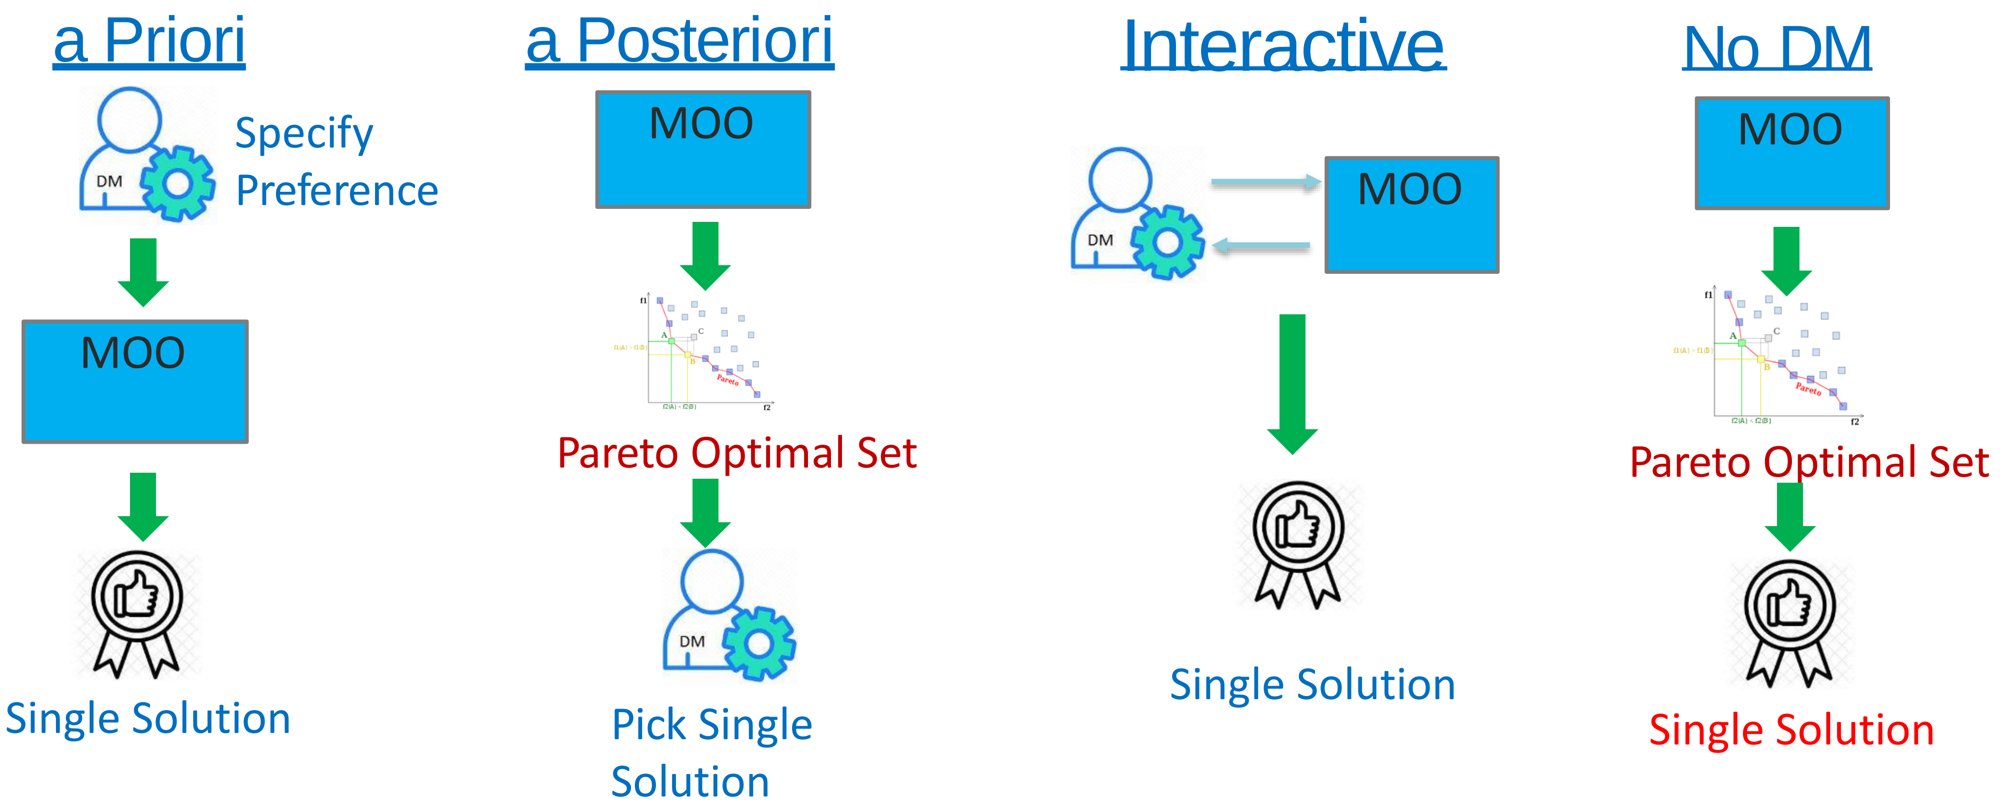

Multi Objective Optimization (MOO)

- Scalarization Algorithms
  - Transform multi-objectives into a single objective
  - Solve it by single objective optimizer
  - Find one Pareto optimal solution in one run
  - Find Pareto Set in multiple run
- Multi Objective Evolutionary Algorithms (MOEA)
  - Follow natural evolution process such as gene evolution, a flock of birds seeking food and other resources, a cooling process of melted crystal, ...
  - Find multiple Pareto optimal solutions in one run

Contents

Scalarization Algorithms

- Weighting Methods
- ϵ-Constraint Method
- Normal Boundary Intersection (NBI) & Normal Constraint (NC)
- Goal Programming
- Physical Programming
- Lexicographic Method

- Weighted Sum Method
  - A weight vector based on DM preference of each objectives:

$$\min_{x} \sum_{i=1}^{M} w_i f_i(x)$$

subject to $x \in S$

Where $\sum_{i=1}^{M} w_i = 1$ and $w_i > 0$

  - The condition of the weights guarantees Pareto optimal

Scalarization Algorithms: Weighted Sum Methods

- Example: Two Objective Metrics Recommender Systems

$$\min_{x \in S}(f_1, f_2), \qquad f_1 = 1 \text{ - accuracy}, \; f_2 = 1 \text{ - diversity}$$

Solve: $\min_{x \in S}(w_1 f_1 + w_2 f_2)$

$$w_1 + w_2 = 1, w, w_2 > 0$$

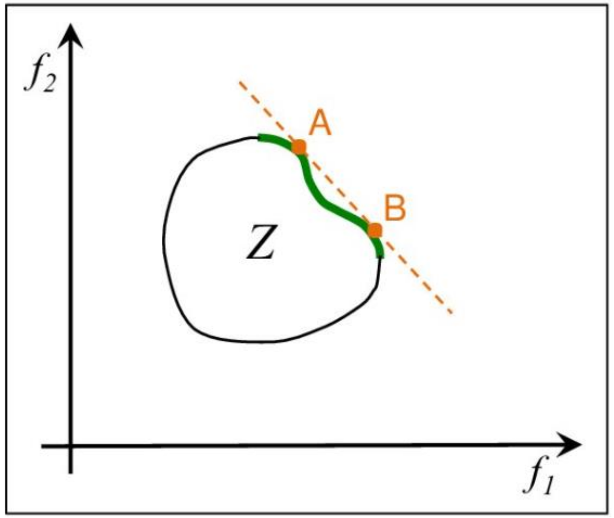

- Each (w_1, w_2) gives one Pareto solution
- Question: Can we get all Pareto solutions in this way? Spoiler: NO

Scalarization Algorithms: Weighting Methods Summary

- Conditions of Pareto optimal solution

Method	Formula	Pro	Con
Weighted Sum	$\sum_{i=1}^{M} w_i f_i(x)$	Simple	Require convex condition for Pareto set
Weighted Exponential Sums	$\sum_{i=1}^{M} w_i [f_i(x)]^p$	Increase p to approximate Pareto set	Bigger p may give non-Pareto solution
Weighted Metric Methods	$[\sum_{i=1}^{M} w_i^p \, \lvert f_i(x) - f_i^* \rvert^p]^{\frac{1}{p}}$	Different choice of ideal points	Bigger p may give non-Pareto solution
Weighted Chebyshev method	$\max_i \{w_i \lvert f_i(x) - f_i^* \rvert\}$	Can find complete Pareto set	Some solution may not Pareto optimal
Exponential Weighted Criterion	$\sum_{i=1}^{K} (e^{p w_i} - 1) e^{p f_i(x)}$	Can find complete Pareto set	May lead computation over low
Weighted Product Method	$\prod_{i=1}^{K} \lvert f_i(x) \rvert^{w_i}$	Deal with different magnitude of objectives	Rarely used

Scalarization Algorithms

- ϵ-Constraint Method[1]

$$\min f_l(x)$$

$$\text{subject to } f_i(x) \leq \epsilon_i, \text{ for all } i \neq l$$

ϵ_i is a known the upper bound of f_i

1) Choose different ϵ_i may produce all Pareto solutions

2) No convex requirement

3) May not Pareto optimal

1. Haimes, Lasdon, Wismer, *On a Bicriterion Formulation of the Problems of Integrated*, System Identification and System Optimization, IEEE Transactions on Systems, Man, And Cybernetics, July 1971

Scalarization Algorithms: ϵ-Constraint Method

- A sufficient condition of Pareto optimal[1]
  - If optimal solution x^* is unique
- Two objective example

solve

$$\min_{x \in S} f_1(x)$$

subject to $f_2 \leq \epsilon$

Constraint	Unique Solution	Pareto Optimal
ϵ^1	No solution	NA
$\epsilon^2, \epsilon^3, \epsilon^4, \epsilon^5$	Yes	Yes
ϵ^6	No	Not necessary

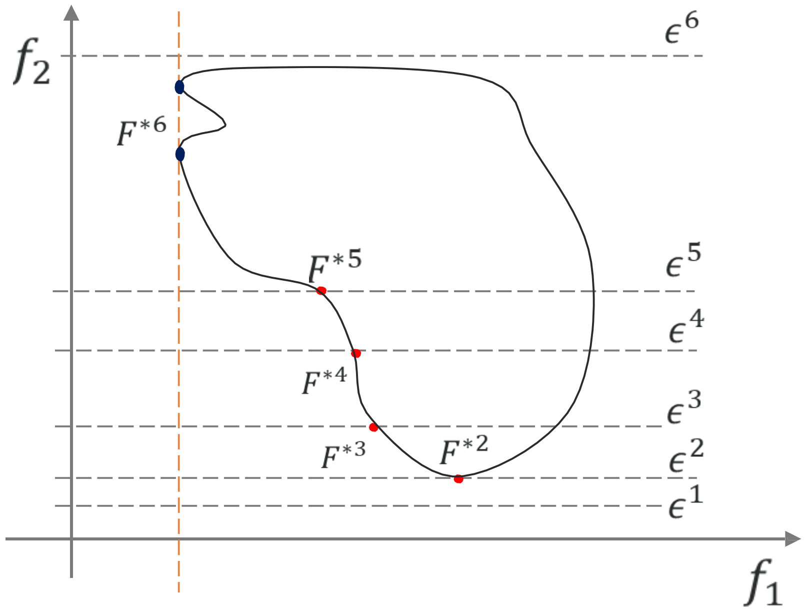

1. V. Chankong, Y. Haimes, Multiobjective Decision Making, Dover Publication, 1983

Contents

- Background and Some History
- Multi Objective Optimization (MOO)
- MOO Solutions
  - Scalarization Algorithms
  - Multi Objective Evolutionary Algorithms
- Selection of the best solution in Pareto set
- MOO libraries
- Summary & QA

Multi Objective Evolutionary Algorithms

- Evolutionary Algorithms inspired by natural evolutionary process:

Genetic Algorithm (GA)

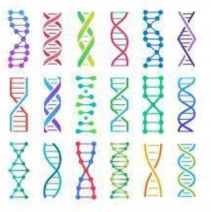

Particle Swam Optimization (PSO)

Simulated Annealing (SA)

Ant Colony Optimal (ACO)

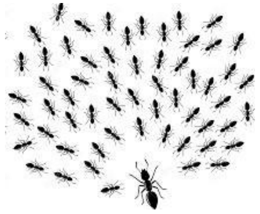

Martin Pilát: Evoluční algoritmy 1 a 2;
Přírodou inspirované algoritmy
NAIL 025, 086, 119

Multi Objective Evolutionary Algorithms

- Benefits:

Objective can be any function

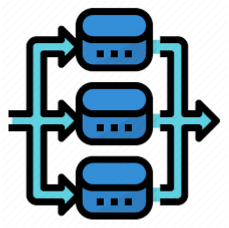

Parallel computing

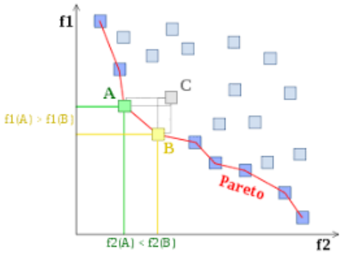

Get Pareto Set in one run

- Terminologies of Solutions:
  - Individual : a feasible solution x
  - Population: a set of individuals
  - Parents: selected from Population
  - Children: produced from Parents

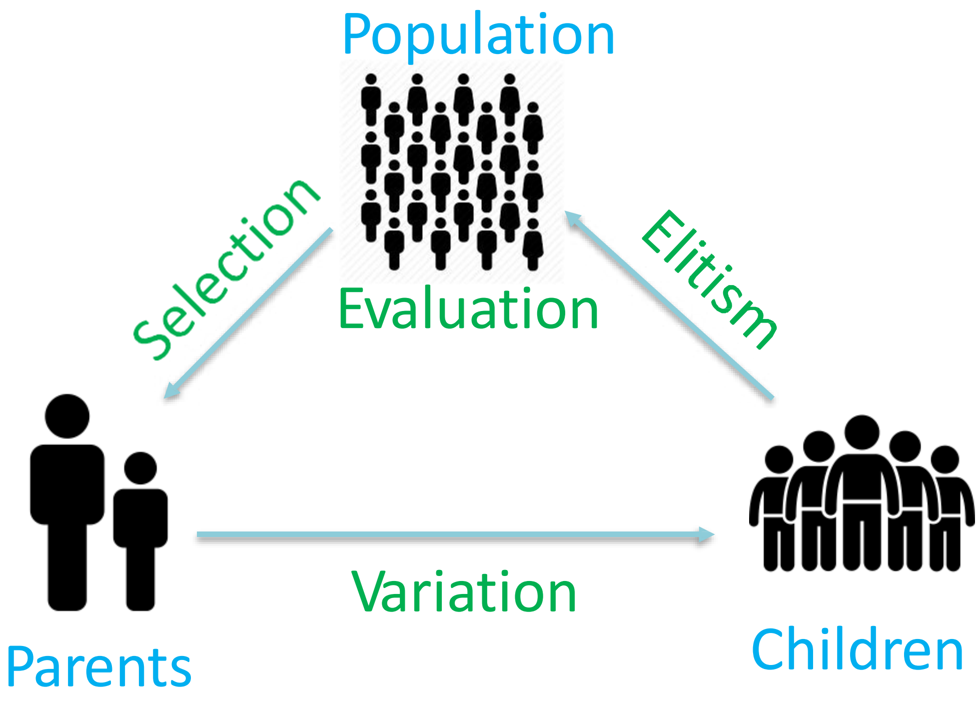

- Operators of Genetic Algorithm
  - Evaluation: measure how 'good' each solution is
    - assigning fitness value (or order): $z(x)$
  - Selection: find Parents
    - Random process
    - Tournament process
  - Variation: produce children
    - Crossover
    - Mutation
  - Elitism:
    - maintain 'better' solution in each iteration

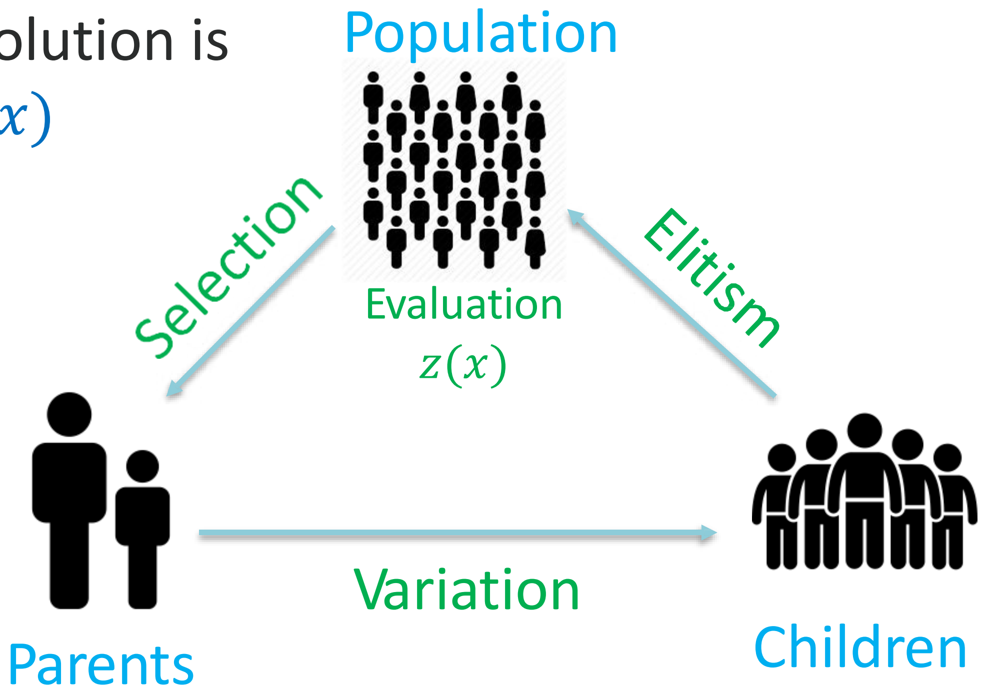

Multi Objective Evolutionary Algorithms : Encoding in Genetic Algorithm

- Binary Encoding in Recommender
  - Chromosome: a recommendation list (decision variable x)
  - Length of Chromosome = total number of available items
  - Each Gene position is corresponding an item
    - 1: item is in recommender list ,
    - 0: item is not in recommender list

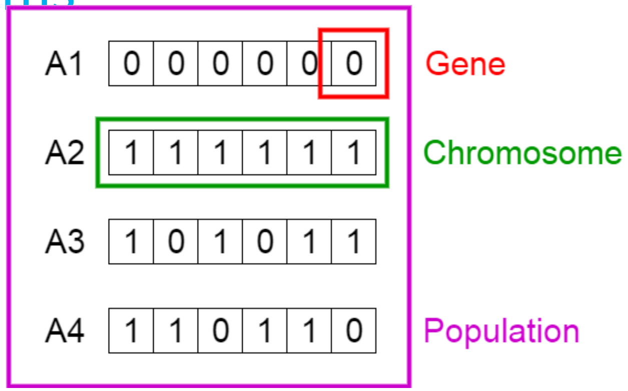

- Permutation Encoding in Recommender
  - Length of Chromosome = N in top N recommendation list
  - The value of each Gene: index of an item
    - 6: the 6th item
    - Total 9 available items
    - $N = 6$

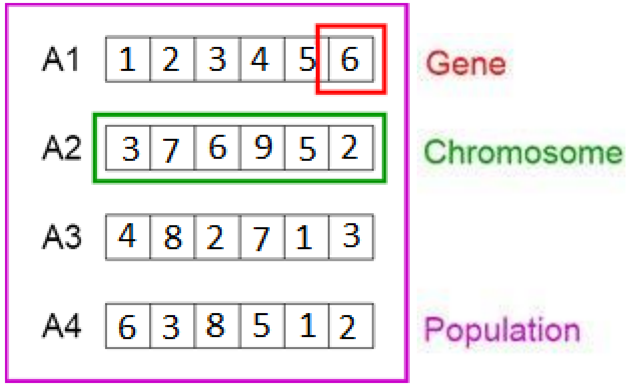

From Single Objective to Multi Objective Evolutionary Algorithms

- Solve a MOO problem
  - Find non-dominated solutions (Pareto optimal)
- Where are multi objective considered? Where are dominance relations applied?
  - Fitness value $z(x)$ evaluation
  - Parent selection
  - Elitism

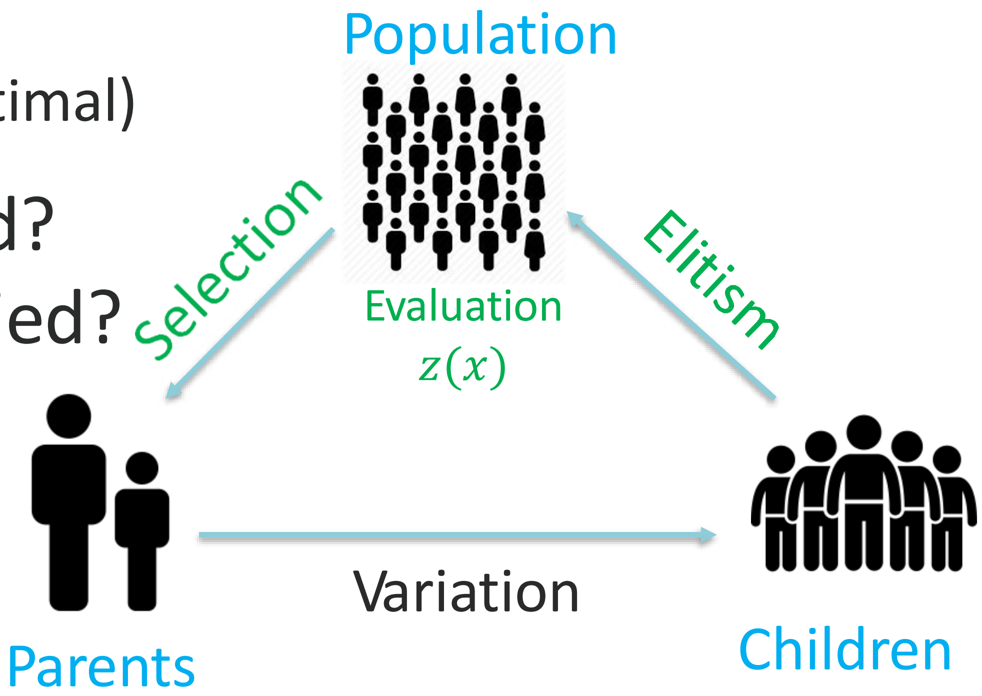

Multi Objective Evolutionary Algorithms:

- Major MOO Genetic Algorithm Methods

Method	Fitness Evaluation	Parent Selection	Elitism
VEGA (Schaffer,1985)	Single objective	Probability distribution	Dominance relation
MOGA (Fonseca & Fleming,1993)	Dominance relation	Probability distribution	NA
NSGA (Srinivas and Deb,1994)	Dominance relation	Probability distribution	NA
NSGA-II (Debb, etc., 2002)	Dominance relation	Probability distribution	Dominance relation
NPGA (Horn, etc. ,1994)	No Fitness	Tournament method (dominance relation)	NA
PAES (Knowles and Corne, 1999)	No Fitness	Local search (dominance relation)	Dominance relation

Classification of Multi Objective Genetic Algorithms

- Dominance relation in Elitism
  - VEGA (Schaffer,1985)
- Dominance relation in fitness function
  - MOGA (Fonseca & Fleming,1993),
  - NSGA (Srinivas and Deb,1994) ,
  - NSGA-II (Debb, etc., 2002)
- No fitness values, but dominance relation in selection process
  - NPGA (Horn, etc. ,1994),
  - PAES (Knowles and Corne, 1999)

- Vector Evaluated Genetic Algorithm (VEGA[1])
  - First genetic algorithm applied to MOO: $\min_{x \in S}(f_1, f_2)$
  - Fitness value, $z_i(x)$, is based on objective function $f_i(x)$

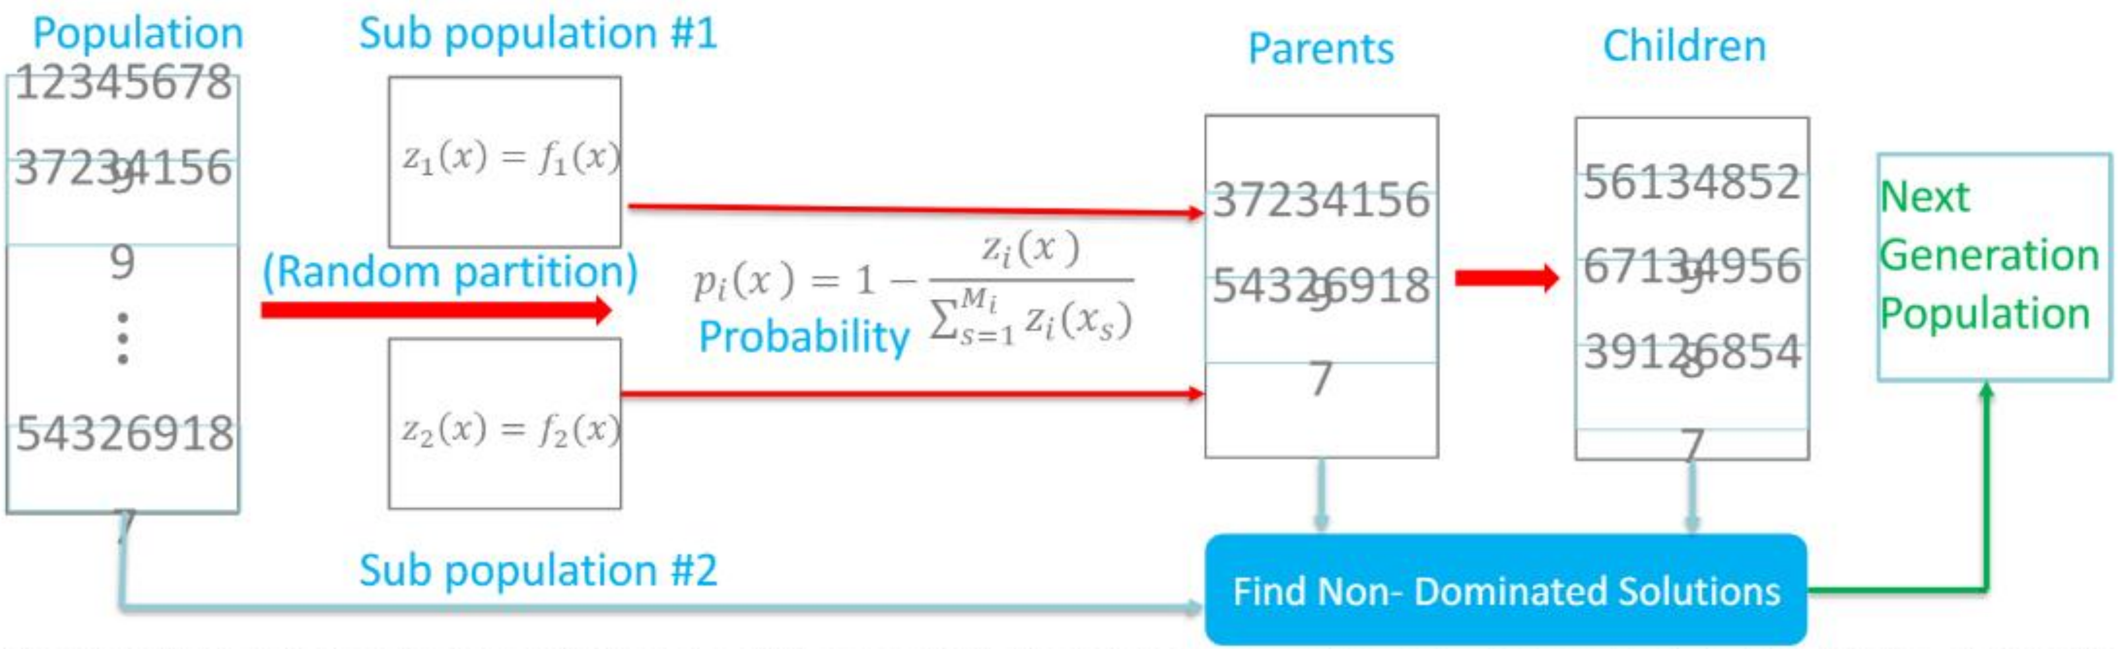

1. Schaffer, 1985: ***Multiple Objective Optimization with Vector Evaluated Genetic Algorithms***, The First International Conference on Genetic Algorithms and their Applications (held in Pittsburgh), pp. 93–100

Classification of Multi Objective Genetic Algorithms

- Fitness value determined by single objective
  - VEGA (Schaffer,1985)
- Fitness value determined by dominance relations
  - MOGA (Fonseca & Fleming,1993),
  - NSGA (Srinivas and Deb,1994) ,
  - NSGA-II (Debb, etc., 2002)
- No fitness value needed:
  - NPGA (Horn, etc. ,1994),
  - PAES (Knowles and Corne, 1999)

Multi Objective Evolutionary Algorithms

- Nondominated Sorting Genetic Algorithm (NSGA[1])
  - Sorting by selecting non-dominated solutions without replacement each time:

    $P_0 = \{A, B\},\ P_1 = \{C, D\}, P_2 = \{E\}$

    P_0 dominates P_1 dominates P_2
  - Assign same fitness value $z(x)$ for each x in P_i based on order

    P_0: $z(A) = z(B) = 10$
    P_1: $z(C) = z(D) = 8$
    P_2: $z(E) = 5$

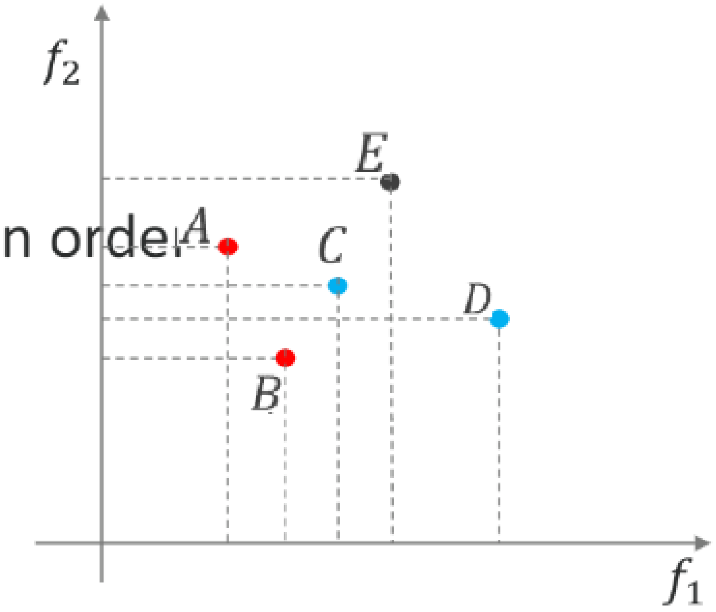

1. Srinivas and Deb, 1994: *Multiobjective optimization using nondominated sorting in genetic algorithms*. Journal of Evolutionary Computing 1994;2(3):221–48.

Multi Objective Evolutionary Algorithms

- MOEA methods based on other EAs
  - Particle Swam Optimization (PSO)

James Kennedy and Russell C. Eberhart. *Particle swarm optimization*. Proceedings of the 1995 IEEE International Conference on Neural Networks, Piscataway, New Jersey, 1995

M. Reyes-Sierra and C. Coello, *Multi-Objective Particle Swarm Optimizers: A Survey of the State-of-the-Art*, International Journal of Computational Intelligence Research, Vol.2, No.3 (2006), pp. 287–308

Contents

- Background and Some History
- Multi Objective Optimization (MOO)
- MOO Solutions
  - Scalarization Algorithms
  - Multi Objective Evolutionary Algorithms
- Selection of the best solution in Pareto set
- MOO libraries
- Summary & QA

Selection of the best solution in Pareto set

- When do we need to produce Pareto set?
  - Most MOEA method (A posterior, No DM available)
  - Scalarization without DM preference (A posterior, No DM available)

- Example: Recommender Systems balancing multi metrics
  - End user (DM) cannot choose from Pareto set
  - A single best recommendation list needs to be produced

Selection of the best solution in Pareto set

- No information from decision maker (DM)
  - All solutions in Pareto set are ‘equally good’
  - Need to make the ‘best’ guess
- Best guess method
  - Knee point method
  - Hypervolume Method
  - Multiple-criteria decision-making (MCDM) methods

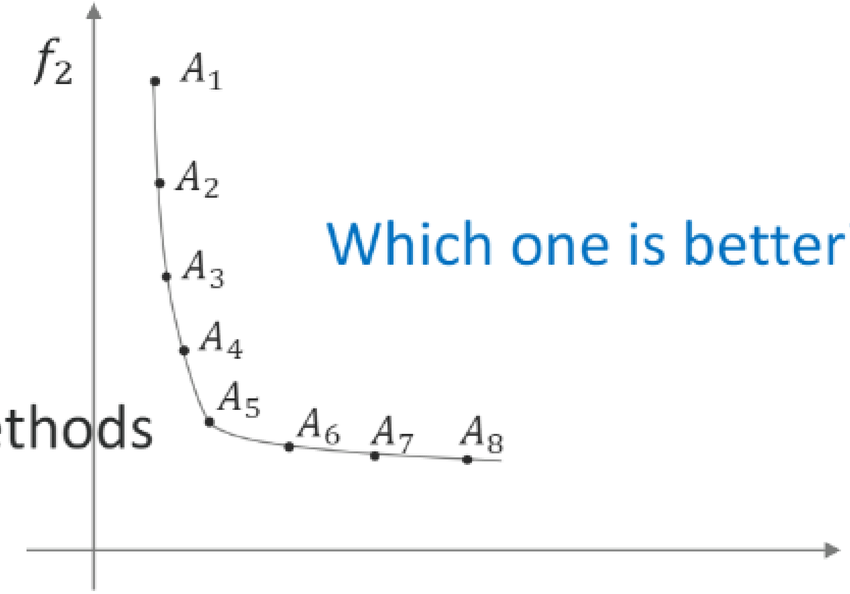

Which one is better?

Selection of the best solution in Pareto set

- Knee Point
  - A special point (A) on Pareto Front
  - small improvement in either objective will cause a large deterioration in the other objective,

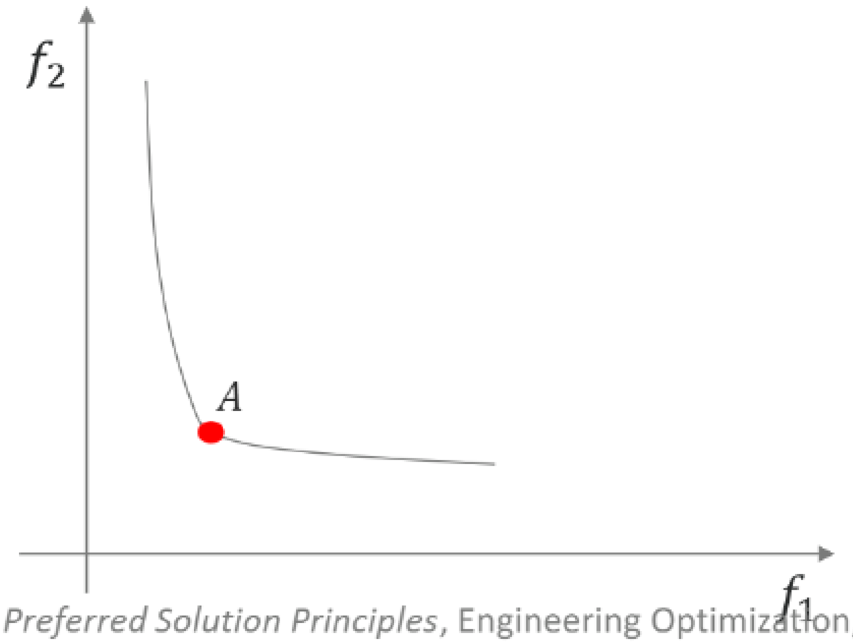

- Find Knee Point
  - Angle Based Method[1]
  - Marginal Utility Method[1]
  - Hyperplane Normal Vector Method[2]

1. Deb and Gupta, *Understanding Knee Points in Bicriteria Problems and Their Implications as Preferred Solution Principles*, Engineering Optimization, 43(11)
2. Yu, Jin, Olhofer, A Method for a Posteriori Identification of Knee Points Based on Solution Density, 2018 IEEE Congress on Evolutionary Computation (CEC)

Selection of the best solution in Pareto set

- Angle Based Knee Point[1]
  - Only work in two objectives MOO
  - Pareto front: $\{A_1, A_2, \ldots, A_8\}$
  - Calculate reflex angle α for each point:

    $$\alpha_i = \angle A_{i-1} A_i A_{i+1}$$

  - Find the point with $\max \alpha_i$

    $$\alpha_5 = \max\{\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6, \alpha_7\}$$

    A_5 is the Knee point

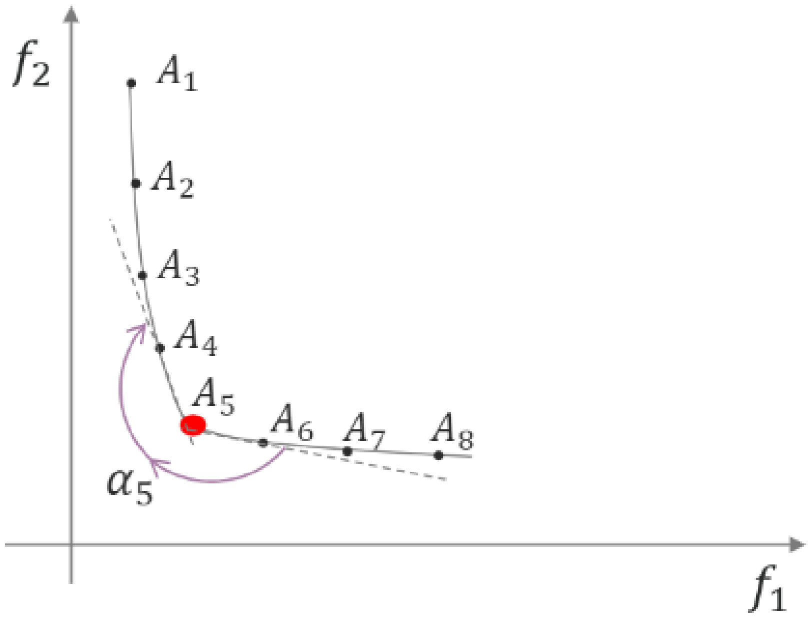

1. Deb and Gupta, *Understanding Knee Points in Bicriteria Problems and Their Implications as Preferred Solution Principles*, Engineering Optimization, 43(11)

RecSys with MOO

- Contexts in which we need MOO in RecSys

RecSys balancing multiple metrics

- Why we need MOO in this context
  - Relevance or accuracy is not the only focus
    - For example, news and music recommendations
      - Boring if always recommending the same types of items
      - Diversity: try something different
      - Novelty: try something never experienced before
    - For example, item recommendations in e-commerce
      - Co-sales
      - Bundle sales

RecSys balancing multiple metrics

- Goals
  - Improve other metrics at no loss or acceptable loss on accuracy
  - Challenges
    - No clear rules to define the "acceptable" loss

RecSys balancing multiple metrics

- Case Study 1: Hybrid Recommender

  Ribeiro, M. T., Lacerda, A., Veloso, A., & Ziviani, N. (2012). Pareto-efficient hybridization for multi-objective recommender systems. In ACM RecSys 2012.

  – Application: balancing accuracy, novelty, diversity

  – Recommendation Framework

Step 1

Build multiple recommendation models to predict rating for a user on an item

Step 2

Fuse predictions by weighted sum

$$\hat{r}_{ui} = w_1 \times \hat{r}_{ui,1} + w_2 \times \hat{r}_{ui,2} + \ldots + w_n \times \hat{r}_{ui,n}$$

Step 3

Produce recommendation list by the aggregated rating

Step 4

Learn the optimal weights (decision variables) by considering accuracy, novelty, diversity

RecSys balancing multiple metrics

- Case Study 1: Hybrid Recommender

  Ribeiro, M. T., Lacerda, A., Veloso, A., & Ziviani, N. (2012). Pareto-efficient hybridization for multi-objective recommender systems. In ACM RecSys 2012.

  – MOEA as the MOO Method
    - Consider accuracy, diversity, novelty as objectives
    - Use Strength Pareto Evolutionary Algorithm as MOEA optimizer
      – Encoding/Decision variables: the weights in the hybrid model
      – Output: a Pareto optimal set
    - Select the best single solution from Pareto set
      – Use a weighted sum on the three objectives
      – Try different set of weights (Q_j) manually

$$\arg\max_{i \in P} \sum_{j=1}^{|O|} Q_j O_{ij}$$

  – Results: balancing multiple metrics

RecSys with MOO

- Contexts in which we need MOO in RecSys

Multi-Stakeholder RecSys

- Why we need MOO in this context
  - The end user is not the only stakeholder

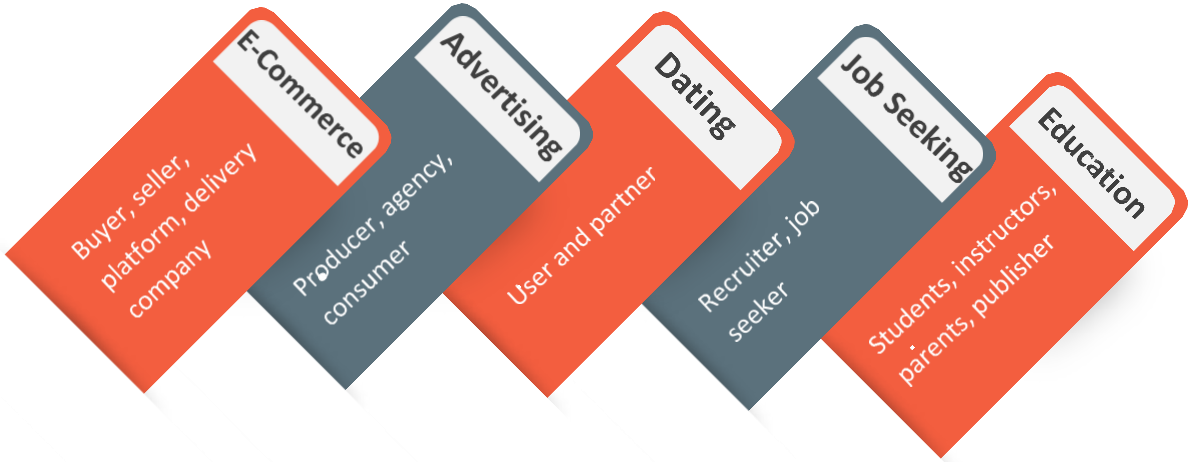

  - RecSys should be built by considering the item utility from the perspective of different stakeholders

Multi-Stakeholder RecSys

- Objective definitions
  - It varies from domains to domains
  - For each stakeholder, there's at least one objective
    - E-Commerce or Marketplace
      - Buyer: user preferences on items, budget
      - Seller: profits
      - Platform: commission fees
      - Delivery company: costs and profits
    - Job seeking
      - Job seeker: user preferences
      - Recruiter: talent requirements

Multi-Stakeholder RecSys

- Goals
  - Deliver item recommendations by balancing the needs of multiple stakeholders
  - With acceptable loss on the consumer side
  - Challenges
    - Which stakeholders should be considered
    - How to define and achieve the “balance”
    - No clear rules to define the acceptable loss

Multi-Stakeholder RecSys

- Case Studies
  - Using scalarization as the MOO method

    Lin, X., Chen, H., Pei, C., et al. (2019). A pareto-efficient algorithm for multiple objective optimization in e-commerce recommendation. In ACM RecSys, 2019.

  - Using MOEA as the MOO method

    Zheng, Y., Ghane, N., & Sabouri, M. (2019). Personalized educational learning with multi-stakeholder optimizations. In Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization (pp. 283-289).

Multi-Stakeholder RecSys

- Case Study 1: Using scalarization in E-Commerce
  - Objectives
    - CTR (Click Through Rate)
    - GMV (Gross Merchandise Volume)

Lin, X., Chen, H., Pei, C., et al. (2019). A pareto-efficient algorithm for multiple objective optimization in e-commerce recommendation. In ACM RecSys, 2019.

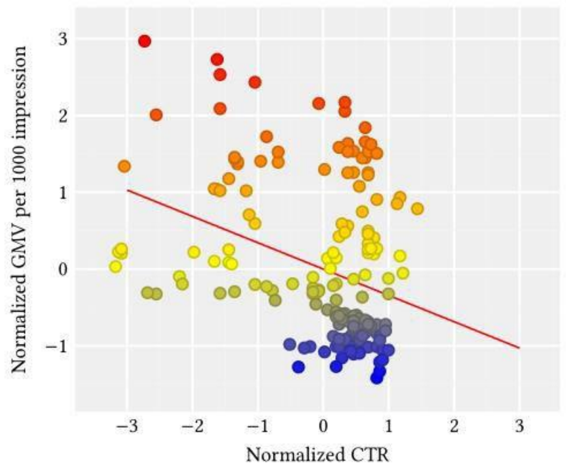

Figure 1: The trade-off between CTR and GMV. The Pearson Correlation Coefficient is -0.343086, with $p < 0.01$.

Multi-Stakeholder RecSys

- Case Study 1: Using scalarization in E-Commerce
  - MOO Method
    - Define a loss function for each objective

      CTR $\mathcal{L}_{CTR}(\theta, x, y, z) = -\frac{1}{N}\sum_{j=1}^{N} log(P(y_j|\theta, x_j))$, i.e., point-wise learning-to-rank

      GMV $\mathcal{L}_{GMV}(\theta, x, y, z) = -\frac{1}{N}\sum_{j=1}^{N} h(price_j) \cdot log(P(z_j = 1|\theta, x_j))$

      x: impression, y: clicks, z: purchases

    - Use weighted sum as the scalarization
      Joint Loss = $\omega \cdot L_{CTR} + (1 - \omega) \cdot L_{GMV}$

Multi-Stakeholder RecSys

- Case Study 1: Using scalarization in E-Commerce
  - MOO Method
    - Use weighted sum as the scalarization
      Joint Loss = $\omega \cdot L_{CTR} + (1 - \omega) \cdot L_{GMV}$
    - Try different weights to get Pareto Set
    - Select a single best solution by using Least Misery strategy, i.e., minimizing the highest loss function of the objectives

$$\min \max\{\mathcal{L}_1, \mathcal{L}_2, \ldots, \mathcal{L}_K\}$$

Summary

- Suggested Workflow

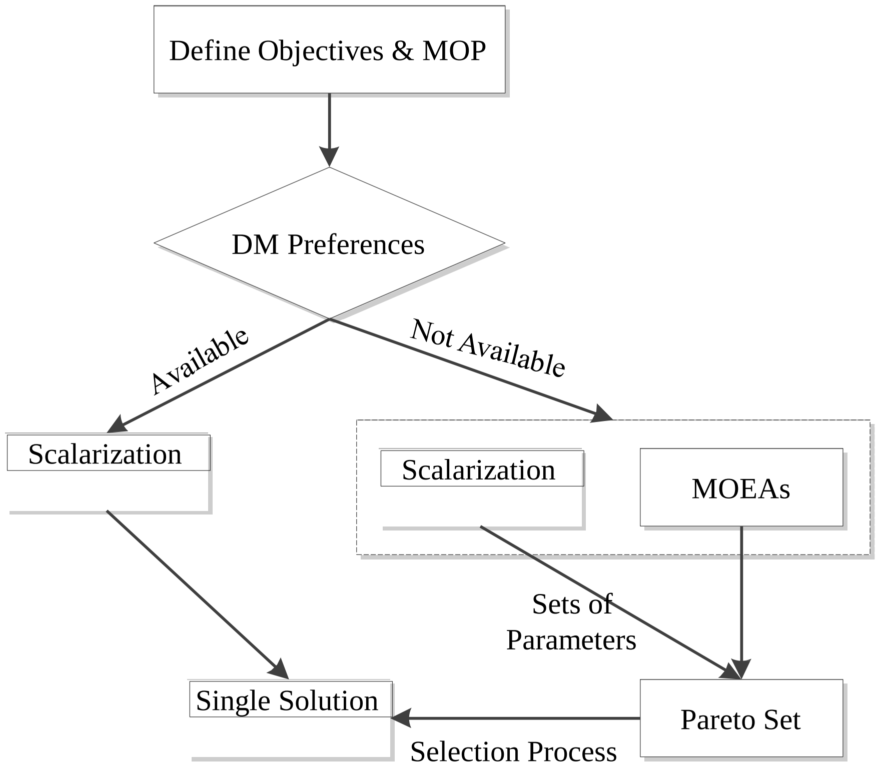

So, who is the decision maker in RecSys?

- The end user
- The developer

Summary

- Our Tutorial
  - Website: https://moorecsys.github.io/
  - Slide: https://github.com/moorecsys/moorecsys.github.io
  - Supplementary materials: "Multi-Objective Recommendations: A Tutorial" on arXiv.org (will be available soon on the github above)