Multi-Criteria Optimization with Preferences

Ulrich Junker
ILOG

Lost in Constraints

Pam, a project planner, is facing a decision-making nightmare with the automated planning system APS.

Pam: How many extra-hours do we need for a delivery in May? Find a project plan assigning everybody to his usual job

APS: 100

Pam: Too much! When can we deliver if we do it without extra-hours?

APS: December

Pam: Too late! When can we deliver if each project member does any task?

APS: October if nobody does his usual job.

Pam: Too bad! I am running out of constraints ...

In Search of the Preferences

Dan, a decision analyst, recommends Pam to relax the constraints into preferences.

Dan: What are your preferences about delivery time if you ignore everything else?

Pam: Early delivery is preferred

Dan: What are your preferences about extra-hours if you ignore everything else?

Pam: Fewer extra-hours are preferred

Dan: You may encounter conflicts between these objectives. Is any criterion more important?

Pam: That depends: for hot projects, it is earliest delivery time; for low-budget projects, it is fewest extra hours; and for all others it is a fair compromise

Problem Type

combinatorial problems where preferences are central

e.g. product configuration, project planning, trip planning

decision space	combinatorial
outcome space	combinatorial
preferences	incomplete & local
uncertainty	no

Combinatorial Problem

- Problem space: $X_1 \times \ldots \times X_n$

  e.g., for each task $i = 1, \ldots, t$ we introduce
  - set X_i of project members who can do task i;
  - set X_{t+i} of time periods for performing task i;

- Constraints: $C \subseteq X_{i_1} \times \ldots \times X_{i_k}$

  Local constraints of small scope $\{i_1, \ldots, i_k\}$, e.g.
  - precedence constraint between tasks i, j

$$x_{t+i} < x_{t+j}$$

  - resource constraint for each project member:

if $x_i = x_j$ then $x_{t+i} < x_{t+j} \vee x_{t+j} < x_{t+i}$

where $x \in X_1 \times \ldots \times X_n$

Combinatorial Decision Space

- Decision space: $D \subseteq X_1 \times \ldots \times X_n$

  s.t. $x \in D$ iff $(x_{i_1}, \ldots, x_{i_k}) \in C$ for all constraints with scope $\{i_1, \ldots, i_k\}$

Combinatorial Outcome Space

- Outcome space: $\Omega_1 \times \ldots \times \Omega_m$
- Criteria: $z_j : X_{j_1} \times \ldots \times X_{j_{k_j}} \rightarrow \Omega_j$

  global criteria of large scope $\{j_1, \ldots, j_{k_j}\}$:
  - delivery time
  - extra hours

  local criteria of small scope $\{j_1, \ldots, j_{k_j}\}$:
  - task of project member l in period p for each l, p
- Assumption:

  global criteria make the problem difficult

Incomplete & Local Preferences

- Preferences are viewpoint specific

  each viewpoint is defined by one or more criteria

  - Marketing: prefer earlier delivery dates all else ignored
  - Administration: prefer less extra-hours all else ignored
  - Project member i: prefer task A over B all else ignored

  rationality principles are restricted to viewpoints!

- Preferences may be incomplete
  - Project member i prefers task A over B, but has no opinion about C

Questions about Preferences

- Modelling:

  How to aggregate viewpoint-specific preferences?

  Which preference models can do this?

- Solving:

  How to solve combinatorial problems under those preferences?

- Explaining:

  How to explain the results while allowing user critics?

  Can we use the original user preferences for this?

Which Preferences? (outline)

- Single viewpoint
  - complete preference orders
  - incomplete preference orders
- Multiple independent viewpoints
  - lexicographic optimization
  - Pareto-optimization
- Importance preferences between viewpoints
  - unconditional importance
  - conditional importance
- Multiple overlapping viewpoints
  - ceteris-paribus
  - reversible

Atomic Optimization Problems

- Preference model:
  - single criterion $z : \mathcal{X} \rightarrow \Omega$
  - total order $\geq$ on Ω
- Problem:

  $$Max_{z,>}(D) := \{x \in D \mid \nexists x^* \in D : z(x^*) > z(x)\}$$
- Classic combinatorial optimization
  - represent order by utility u s.t. $\omega_1 \geq \omega_2$ iff $u(\omega_1) \geq u(\omega 2)$
  - solve $max\{u(z(x)) \mid x \in D\}$ and let x^* be a solution
  - can be solved by existing optimizers
- Solved form:
  - let ω^* be the value $z(x^*)$ of z in solution x^*
  - then $Max_{z,>}(D) = \{x \mid x \in D \wedge z(x) = \omega^*\}$

The Answer is 42 ...

but what was the question?

- Optimization proble:m

  let ω^* be the optimal value of z under constraints $\mathcal{C}$

- Explanation questions:
  - Why is ω^* optimal?
  - Why isn't ω chosen instead?

- Explanation of optimality: $(>, \omega^*, E)$

  where E is a simplest subproblem (minimal subset) of $\mathcal{C}$ s.t. ω^* is the optimal value of z under E
  - ω^* is optimal as E defeats all better values
  - ω is not chosen since $\omega^* > \omega$ or ω is defeated by E

How to compute explanations?

- Reduce to conflicts:
  - find a minimal unsatisfiable subset E' (“conflict”) of $\mathcal{C} \cup \{z > \omega^*\}$
  - $(>, \omega^*, E' \setminus \{z > \omega^*\})$ is an explanation of optimality
- How to compute conflicts?
  - perform a sequence of satisfiability checks
  - QUICKXPLAIN accelerates the basic method by divide-and-conquer [AAAI-04]
  - QUICKXPLAIN provides a well working explanation technology with a growing list of successful applications.

Incomplete Preference Orders

- Assumption:
  - the decision maker has given only some preferences
  - hence, the complete preference relation is a superset of the given preferences
  - the given preferences define a space of possible complete preference relations
- Preference model:
  - Single criterion $z : \mathcal{X} \rightarrow \Omega$
  - space of complete orders on Ω that are supersets of a given (Partial) preorder $\succsim$ on Ω

Alternative Optimizations

- Problem:
  - let $\tau(\succ)$ the set of complete extensions of $\succ$
  - $Max_{z,\succ}(D) := \bigcup_{>\in\tau(\succ)} Max_{z,>}(D)$
- Optimization under Partial Orders:

$$Max_{z,\succ}(D) = \{x \in D \mid \nexists x^* \in D : z(x^*) \succ z(x)\}$$

- Solved form:
  - let Ω^* be the optima to be found by the optimizer
  - $Max_{z,\succ}(D) = \{x \mid x \in D \wedge \bigvee_{\omega^* \in \Omega^*} z(x) = \omega^*\}$

Optimization under Partial Orders

- Inner branching
  - divide decision space by a constraint α
  - $Max_{z,\succ}(\{x \mid x \in D \wedge \alpha\}) \cup Max_{z,\succ}(\{x \mid x \in D \wedge \neg\alpha\})$
  - results into a relaxation of the original problem
  - used in multi-objective branch-and-bound
- Outer branching
  - divide space of optional solutions by a constraint α
  - $\{x \mid x \in Max_{z,\succ}(D) \wedge \alpha\} \cup \{x \mid x \in Max_{z,\succ}(D) \wedge \neg\alpha\}$
  - subproblems can be simplified for certain branching constraints
  - see [MOPGP'06]

Outer Branching

- Branching constraint $z(x) = \omega^*$
  - choose one linear extension $>$ of $\succ$
  - solve $Max_{z,>}(D)$ and let ω^* be the optimum
- Left-branch
  - set of all optimal solutions x of D s.t. $z(x) = \omega^*$
  - this reduces to $\{x \mid x \in D \wedge z(x) = \omega^*\}$
- Right-branch
  - set of all optimal solutions x of D s.t. $z(x) \neq \omega^*$
  - this reduces to $Max_{z,\succ}(\{x \mid x \in D \wedge z(x) \not\preceq \omega^*\})$
- Property

  each atomic optimization problem, except for the last one, produces a new optimal value

Explanations under Partial Orders

- Outer branching
  - chooses a linear extension $>$ of $\succ$ and
  - generates the optima $\omega_1^*, \ldots, \omega_k^*$ in decreasing $>$-order
- Explanation with dominance constraints
  - each ω_i^* has an explanation of optimality $(>, \omega_i^*, E_i)$
  - but E_i contains dominance constraints $z(x) \not\preceq \omega_j^*$
- Explanation without dominance constraints
  - define extension $\succ_i$ of $\succ$ s.t. $\omega_i^* \succ_i \omega_j^*$ for $j \neq i$
  - choose a linear extension $>_i$ of $\succ_i$
  - find new explanation $(>_i, \omega_i, E_i')$, namely for the optimality of ω_i^* w.r.t. to $>_i$

Which Preferences? (outline)

- Single viewpoint
  - complete preference orders
  - incomplete preference orders
- Multiple independent viewpoints
  - lexicographic optimization
  - Pareto-optimization
- Importance preferences between viewpoints
  - unconditional importance
  - conditional importance
- Multiple overlapping viewpoints
  - ceteris-paribus
  - reversible

Multiple criteria

- Assumptions:
  - each criterion represents an independent viewpoint
  - rationality principles are restricted to this viewpoint
  - DM answers questions about viewpoint-preferences promptly
  - DM needs to reason, make judgements and choices when being asked for preferences involving multiple viewpoints together; the rationality of those answers is not guaranteed
- Preference model:
  - multiple criteria $z_i : \mathcal{X} \rightarrow \Omega_i$
  - (partial) preorder $\succsim_i$ on Ω_i

Lexicographical Optimality

- Preference aggregation:
  - $\succ_{lex}$ on $\Omega_1 \times \ldots \times \Omega_m$
  - $(\omega_1^*, \ldots, \omega_m^*) \succ_{lex} (\omega_1, \ldots, \omega_m)$ iff there is a k s.t. $\omega_j^* = \omega_j$ for all $j = 1, \ldots, k-1$ and $\omega_k^* \succ \omega_k$
- Problem:

$$Lex_{\langle z_1, \succ_1 \rangle, \ldots, \langle z_m, \succ_m \rangle}(D) := \{x \in D \mid \nexists x^* \in D : z(x^*) \succ_{lex} z(x)\}$$

- Optimization:
  - $Lex_{\langle z_1, \succ_1 \rangle}(D) = Max_{z_1, \succ_1}(D)$
  - $Lex_{\langle z_1, \succ_1 \rangle, \ldots, \langle z_m, \succ_m \rangle}(D) = Lex_{\langle z_2, \succ_2 \rangle, \ldots, \langle z_m, \succ_m \rangle}(Max_{z_1, \succ_1}(D))$

Lexicographical Optimality

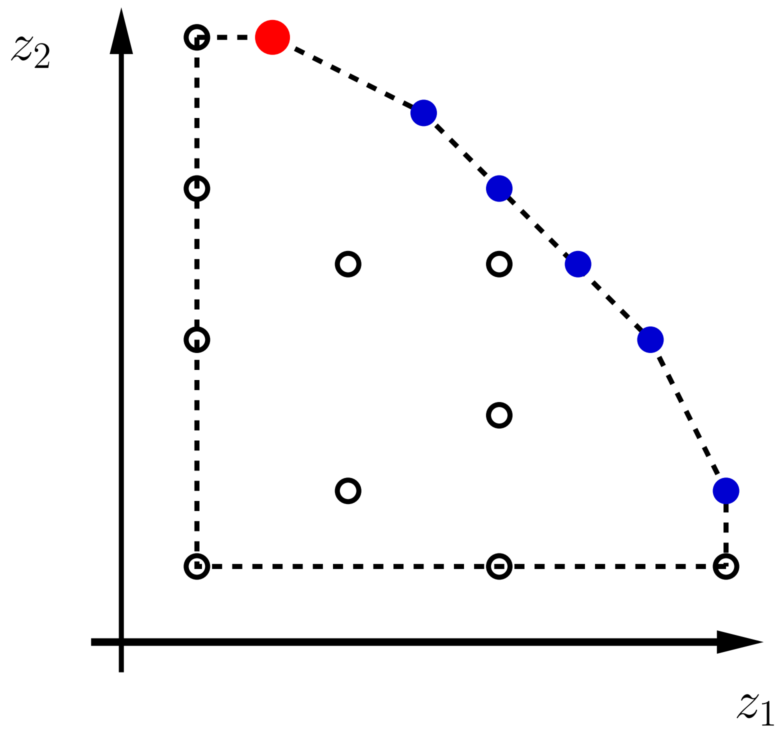

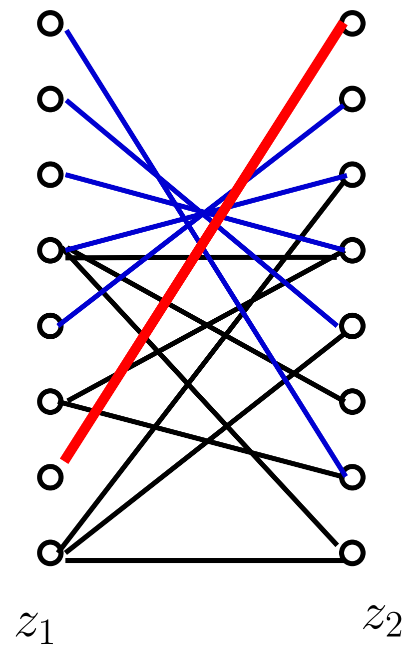

Lex-Optimality Explained

- Lexicographic Optimization:
  - solved by a sequence of atomic optimization problems
  - if $(\omega_1^*, \ldots, \omega_m^*)$ is a lex-optimal solution, then ω_i^* is an optimal value of z_i in $Max_{z_i, \succ_i}(\{x \mid x \in D \wedge z_1(x) = \omega_1^* \wedge \ldots \wedge z_{i-1}(x) = \omega_{i-1}^*\})$
- Explanations of Lexicographic Optimality:
  - sequence $(\xi_1, \ldots, \xi_n)$ of explanations
  - $\xi_i = (>_i, \omega_i^*, E_i)$ is an explanation of optimality of ω_i^*
  - E_i may contain constraints $z_j(x) = \omega_j^*$ for $j < i$
  - these constraints indicate that the optimal values of more important criteria defeated better values of z_i

Alternative Sequentializations

- Preference aggregation:
  - let Π be the set of permutations of $1, \ldots, m$
  - $Extreme_{\langle z_1, \succ_1 \rangle, \ldots, \langle z_m, \succ_m \rangle}(D) :=$
    $\bigcup_{\pi \in \Pi} Lex_{\langle z_{\pi_1}, \succ_{\pi_1} \rangle, \ldots, \langle z_{\pi_m}, \succ_{\pi_m} \rangle}(D)$

Alternative Sequentializations

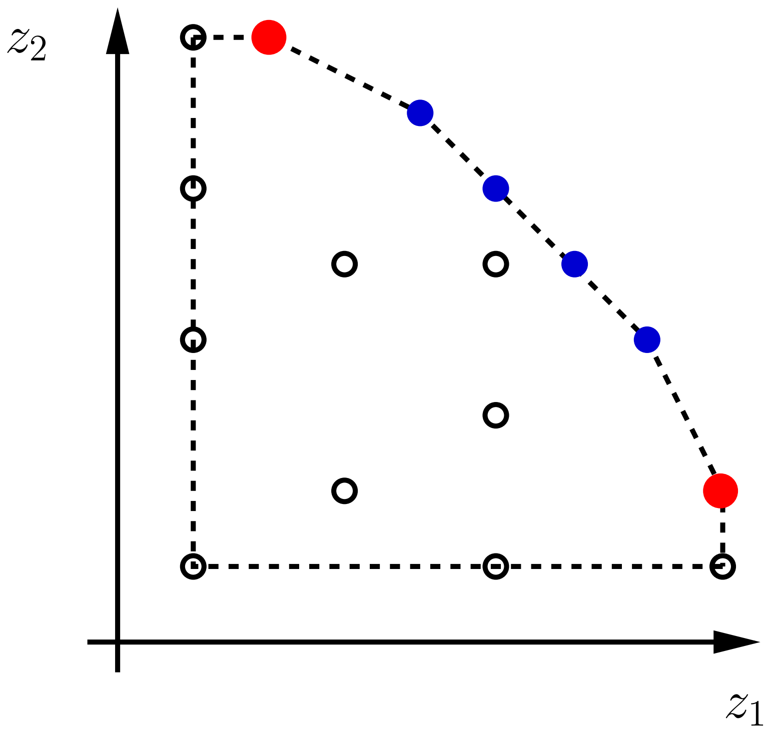

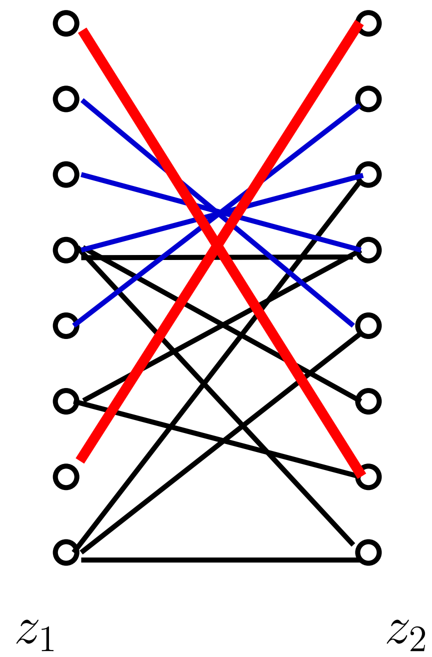

Extreme Solutions Explained

- A sequentialization:
  - is characterized by a permutation π
  - let $(\omega^*_{\pi_1} \ldots, \omega^*_{\pi_m})$ be a solution of $Lex_{\langle z_{\pi_1}, \succ_{\pi_1}\rangle, \ldots, \langle z_{\pi_m}, \succ_{\pi_m}\rangle}(D)$
- Explanation of Optimality:
  - take an explanation $(\xi_{\pi_1} \ldots, \xi_{\pi_m})$ of the lex-optimality of $(\omega^*_{\pi_1} \ldots, \omega^*_{\pi_m})$
  - it lists the criteria in the chosen order of importance

Pareto-Optimality

- Preference aggregation:
  - define $\succ_{Pareto}$ on $\Omega_1 \times \ldots \times \Omega_m$
  - $(\omega_1^*, \ldots, \omega_m^*) \succ_{Pareto} (\omega_1, \ldots, \omega_m)$ iff $\omega_i^* \succeq \omega_i$ for all i and $\omega_i^* \succ \omega_i$ for one i
- Problem:

$$Pareto_{\langle z_1, \succ_1 \rangle, \ldots, \langle z_m, \succ_m \rangle}(D) := \{x \in D \mid \nexists x^* \in D : z(x^*) \succ_{pareto} z(x)\}$$

- Optimization:
  - use outer branching
  - lexicographic order is used as linear extension of Pareto-dominance

Pareto Optimality

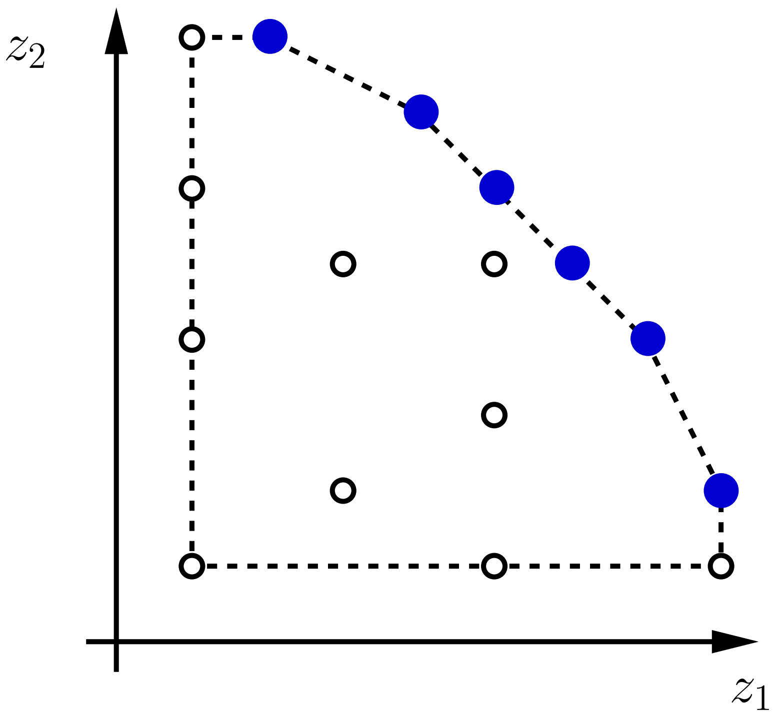

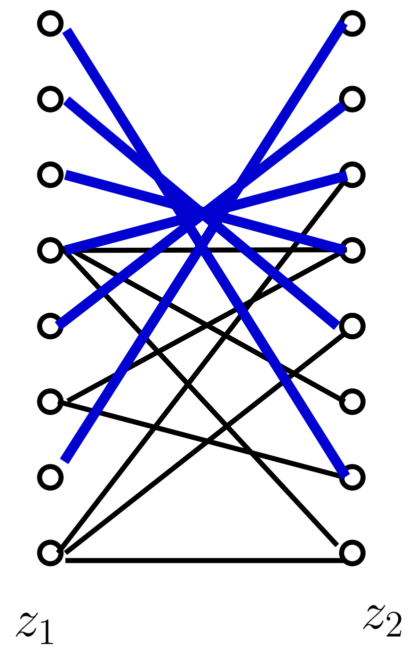

Which Preferences? (outline)

- Single viewpoint
  - complete preference orders
  - incomplete preference orders
- Multiple independent viewpoints
  - lexicographic optimization
  - Pareto-optimization
- Importance preferences between viewpoints
  - unconditional importance
  - conditional importance
- Multiple overlapping viewpoints
  - ceteris-paribus
  - reversible

Importance of criteria

- Assumption:
  - DM is easily able to decide the relative importance of criteria
- Preference model:
  - multiple criteria $z_i : \mathcal{X} \rightarrow \Omega_i$
  - (partial) preorder $\succsim_i$ on Ω_i
  - a strict partial order I on the set of criteria indices
- Aggregation principles:

  a preorder $\succsim$ on Ω is an aggregation under this model iff
  - if $(i, j) \in I$ then $\omega^* \succ \omega$ for all $\omega^*, \omega \in \Omega$ that satisfy $\omega_i^* \succ_i \omega_i$, and $\omega_k^* = \omega_k$ for $k \in \{1, \ldots, m\} \setminus \{i, j\}$

Constrained Sequentializations

- Preference aggregation:
  - a permutation π respects I iff
  $$(\pi_i, \pi_j) \in I \text{ implies } i < j$$
  - let $\Pi(I)$ be the set of permutations of $1, \ldots, m$ respecting I
  - $Extreme^I_{\langle z_1, \succ_1 \rangle, \ldots, \langle z_m, \succ_m \rangle}(D) := \bigcup_{\pi \in \Pi(I)} Lex_{\langle z_{\pi_1}, \succ_{\pi_1} \rangle, \ldots, \langle z_{\pi_m}, \succ_{\pi_m} \rangle}(D)$
  - see [ANOR 04]

Constrained Pareto-Optimality

- Preference aggregation:
  - define $\succ^{I}_{Pareto}$ on $\Omega_1 \times \ldots \times \Omega_m$
  - $(\omega_1^*, \ldots, \omega_m^*) \succ^{I}_{Pareto} (\omega_1, \ldots, \omega_m)$ iff
    1. $\omega_i^* \neq \omega_i$ for some i and
    2. if $\omega_i^* \neq \omega_i$ then (i) $\omega_i^* \succ \omega_i$ or (ii) there is a j that is more important than i and $\omega_j^* \neq \omega_j$

Which Preferences? (outline)

- Single viewpoint
  - complete preference orders
  - incomplete preference orders
- Multiple independent viewpoints
  - lexicographic optimization
  - Pareto-optimization
- Importance preferences between viewpoints
  - unconditional importance
  - conditional importance
- Multiple overlapping viewpoints
  - ceteris-paribus
  - reversible

Reversible Preferences

- Assumption:
  - there are different overlapping viewpoints
  - some viewpoints are more specific than others (supersets)
  - preferences on more specific viewpoints may reverse preferences obtained from more general preferences by Pareto-aggregation
- Approach:
  - extension principle: more general preferences are extended to more specific viewpoints
  - conservation principle: best choices of the more specific viewpoint need to be preserved under certain conditions

Conclusion

- Modelling
  - viewpoint specific semantics
  - importance preferences constraint Pareto-aggregation
- Optimzation
  - complex problems are reduced to atomic problems solvable
  - preference and constraint handling is decoupled
  - existing optimizers can be used
- Explanation
  - are important to justify a solution
  - exhibit critical preferences for changing the solution