CONCERNING OPTIMIZATION PARAMETERS AND PREFERENCES

Nils Wilde Computer Science, Dalhousie University, nils.wilde@dal.ca

Dylan A. Shell Computer Science and Engineering, Texas A&M University, dshell@tamu.edu

The Dilemma

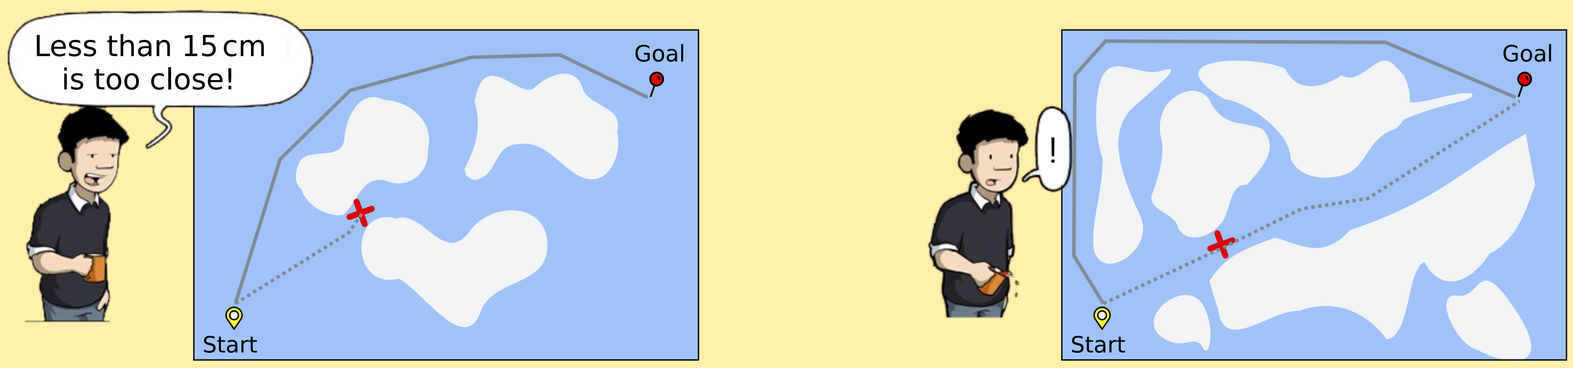

Optimization parameters are often chosen to align planning with user preferences.
To what extent do mechanisms like weightings / constraints actually express preferences?

Preferences *vs* Parameters *vs* Pareto fronts

- When doing (hyper)parameter tuning, you are implicitly navigating a space of closely related optimization problems to suit a preference.
- By picking a point from a Pareto front, multi-objective optimization explicitly treats identification of a preference. But then

$$\text{Parameters} \stackrel{?}{=} \text{Preferences}$$

Definitions: Constraints and Objectives

- Optimization problems often consider various criteria in form of functions $f_1(T), \ldots, f_n(T)$
- These are combined on various formulations:
  - Constraints: $\min_T f_1(T),\ s.t. f_i(T) \leq \epsilon_i$
  - Weighted sum (WS): $\min_T \sum w_i \cdot f_i(T)$
  - Weighted max (WM): $\min_T \max w_i \cdot f_i(T)$
- Despite being popular, WS can only find points on *convex* Pareto fronts

An illustrative example

Dubins car navigating around obstacles

Setup:

- Vary problem by changing obstacle placement
- Fix constraint values on clearance / closeness

Problem 1, Scenario A — obstacle placement.

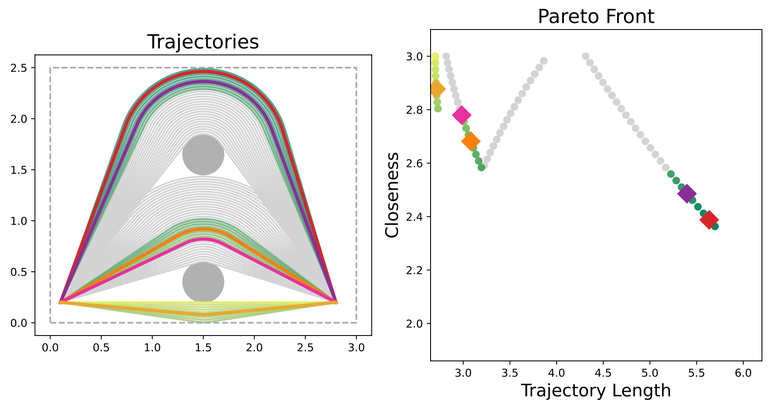

Problem 1, Instance B — *slightly* moved obstacle.

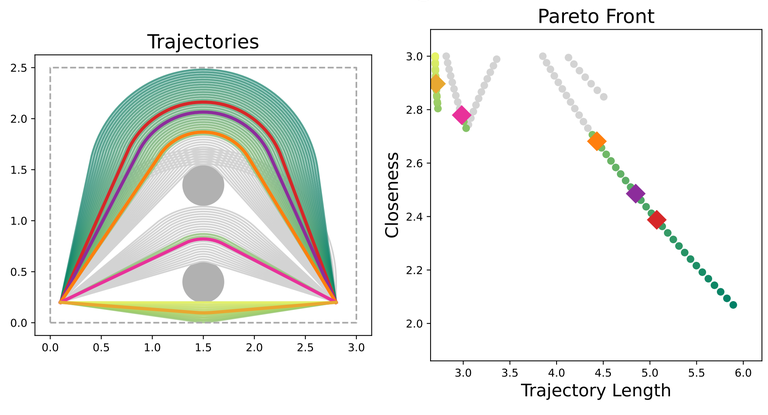

Observation: Fixed constraints can lead to major changes in paths, despite only small changes in the environment.

Case Study I — Fixed constraints

Setup:

- Vary problem by changing obstacle placement
- Fix constraint values ϵ on clearance / closeness

Solutions for fixed constraints across Pareto-fronts

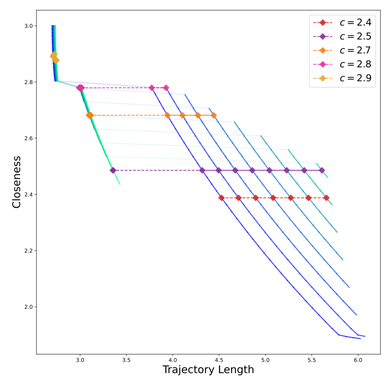

Observation: Constraints are always tightly met, when feasible; yet nearby solutions are disregarded even when large improvements are attainable for small constraint violation.

Case Study II — Fixed objective functions

Setup:

- Vary problem by changing obstacle placement
- Fix weights w in WS / WM objectives

Solutions for fixed WS weights across Pareto-fronts

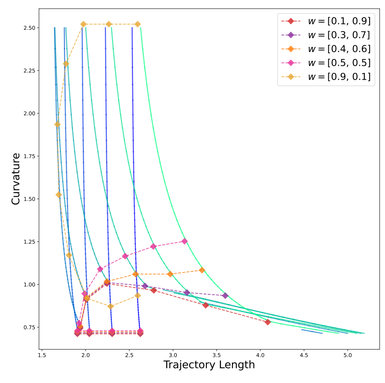

Solutions for fixed WM weights across Pareto-fronts

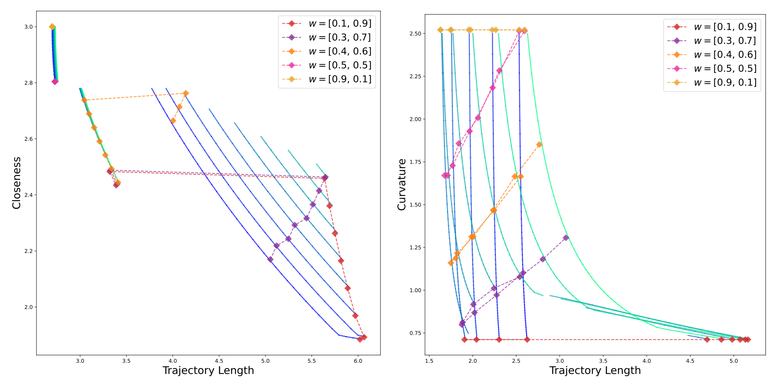

Observation: Weights allow for balancing objectives in each instance individually. However, trends across instances are inconsistent, and weights are often ambiguous. This can lead to **preference indefiniteness**— multiple weights can express the same preference in some instances, but not others.

Some interpretation

Pros of the multi-objective view:

+ Multi-objective optimization avoids imposing a total order where none had been expressed.
  → This is a form of *unbiasedness*.

+ The multi-objective approach never combines incommensurate quantities.
  → Incompatible or incomparable quantities highlight how ill-founded *any* such combination would be.

+ A Pareto front enables visibility, not just of individual solutions but their relationships.
  → Decisions informed by how problem constrains what can be attained.

Cons of the multi-objective view:

– The Pareto front for one instance visualizes nothing about other instances.
  → Whence any *generalization*?

– Pareto front *geometry* is altered by objective transformations.
  → One *biased* view of inter-relationships.

– Delay is appropriate if utilizing new information to resolve uncertainty.
  → A Pareto frontier expresses uncertainty poorly.

Open Questions

1. How can multi-objective optimization be better informed by their downstream uses?
2. Can we improve on Pareto frontiers?
   (a) For example, how to explicitly visualize inter- and intra-instance uncertainty?
   (b) How to remove arbitrariness such as scaling artifacts and visual aspect ratios?
3. What is the interplay between decision-theoretic axioms for preferences (e.g., antisymmetry and transitivity) versus the selection of options from a Pareto front?
4. Do learning theoretic concepts for generalization (e.g., empirical risk minimization) apply to preference selection?

Acknowledgments

Our cartoon of the dilemma is based upon an original illustration by Etienne Issartial of www.commitstrip.com. The Ron Swanson was obtained from meme checkpoint.