Advanced Local Search

CPSC 383: Explorations in Artificial Intelligence and Machine Learning
Fall 2025

Jonathan Hudson, Ph.D
Associate Professor (Teaching)
Department of Computer Science
University of Calgary

August 27, 2025

UNIVERSITY OF CALGARY

Outline

- Queueing
- Local Search
- Hill-Climbing
- Simulated Annealing
- Particle Swarm Optimization
- Genetic Algorithms

Queueing

The One Queue

- Algorithms' structures mentioned prior are similar!
  - Have a fringe to explore
  - Ordered by some eval of each item in fringe
  - We can use a priority queue
    - Most common data structure for this is a heap

  - Practically, for DFS and BFS, you can use stacks and queues to be even more efficient

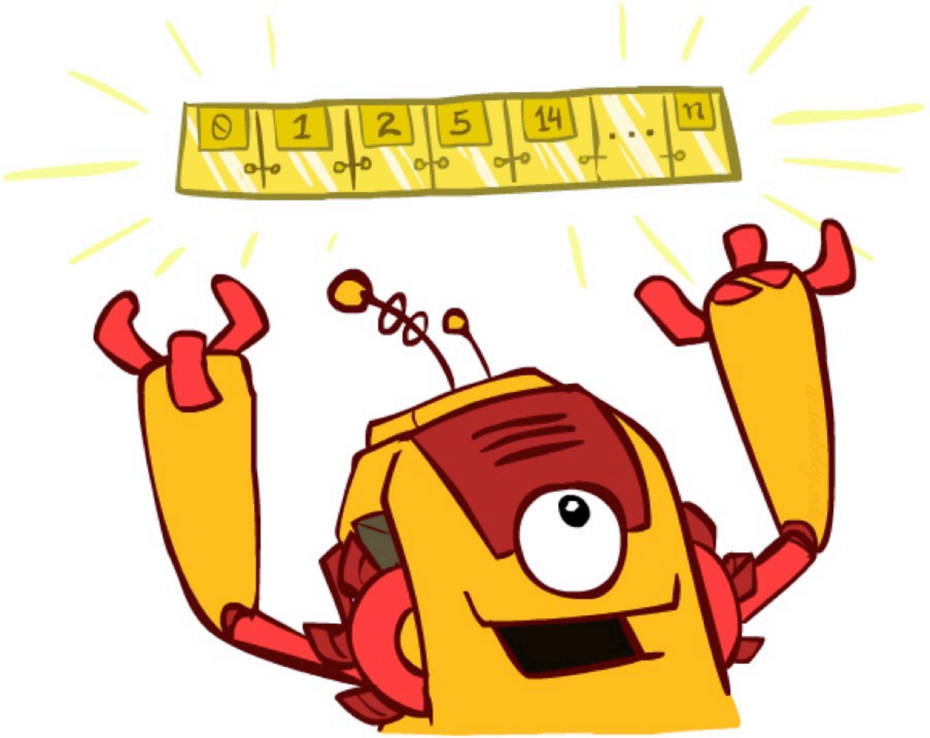

Local Search

Local Search (I)

General Idea:

After selecting a transition, do not consider any transitions that were possible in previous states

☞ “Never-look-back-Heuristic”

Example: trees (works for sets also ☞ one-element sets)

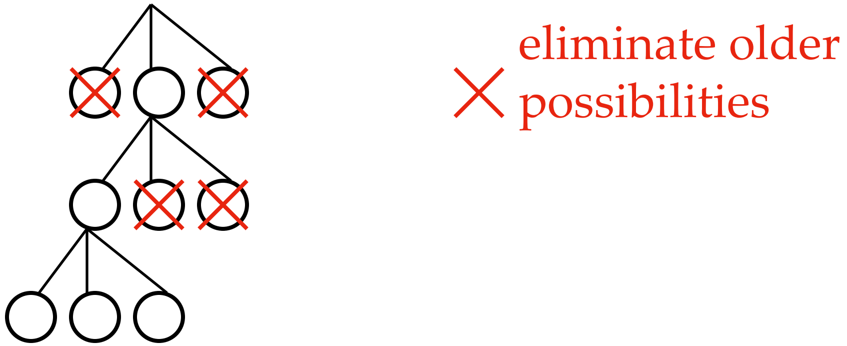

Local Search (II)

Advantages:

- Less decisions
- Complexity can be bound by depth of tree (number of solution steps)
- Each transition contributes to found solution
- Predictable behavior with regard to run time

Disadvantages

- No guarantee for optimality of solution
- No guarantee for optimality of number of necessary transitions

Simple Local Search

Hill Climbing

- Simple, general idea:
  - Start wherever
  - Look one-move neighbours (left-right for 2D example)
  - If better choice move to it
  - If no neighbours better quit

- What's bad about this approach?

- What's good about it?

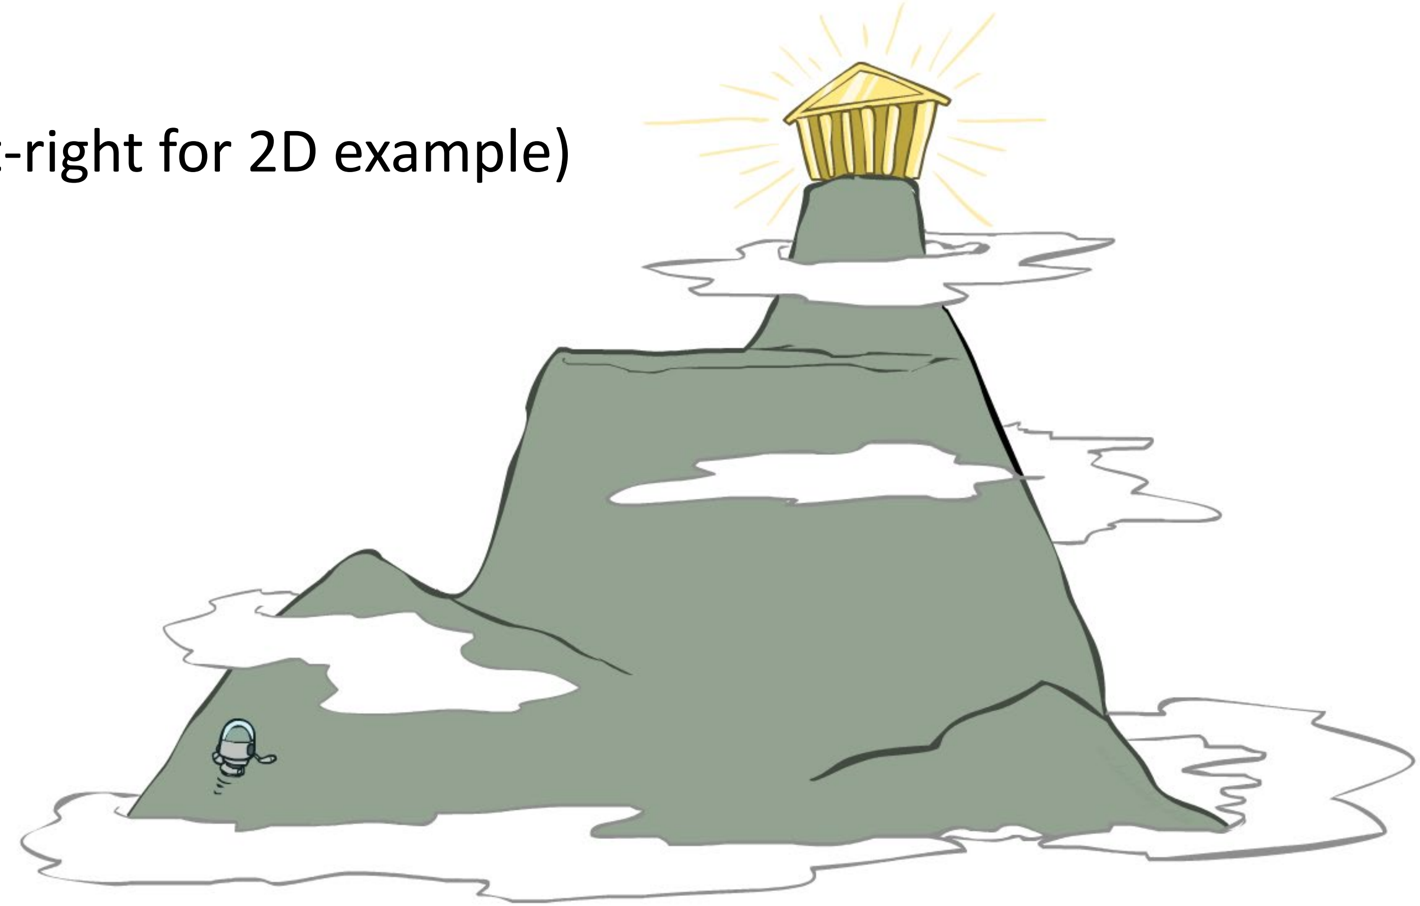

CALGARY

Hill Climbing Diagram

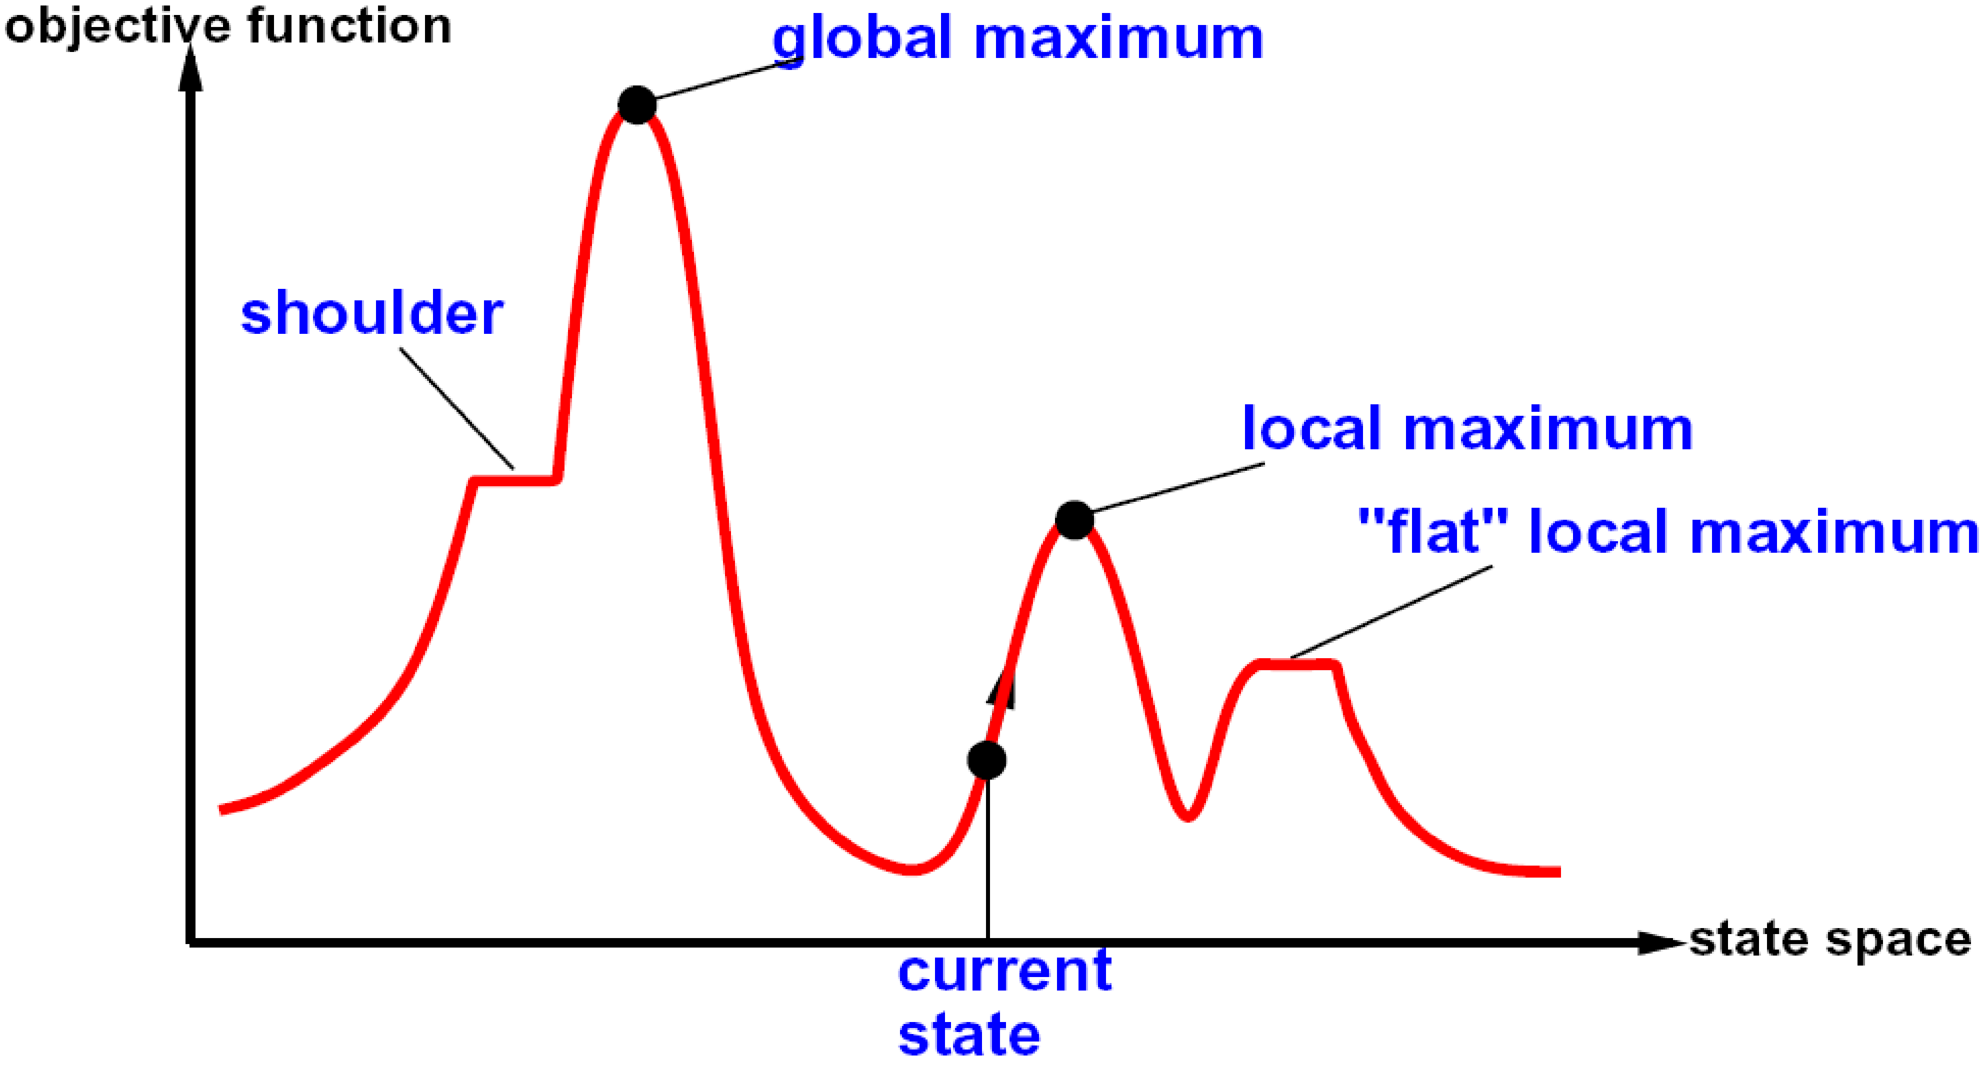

Hill Climbing Quiz

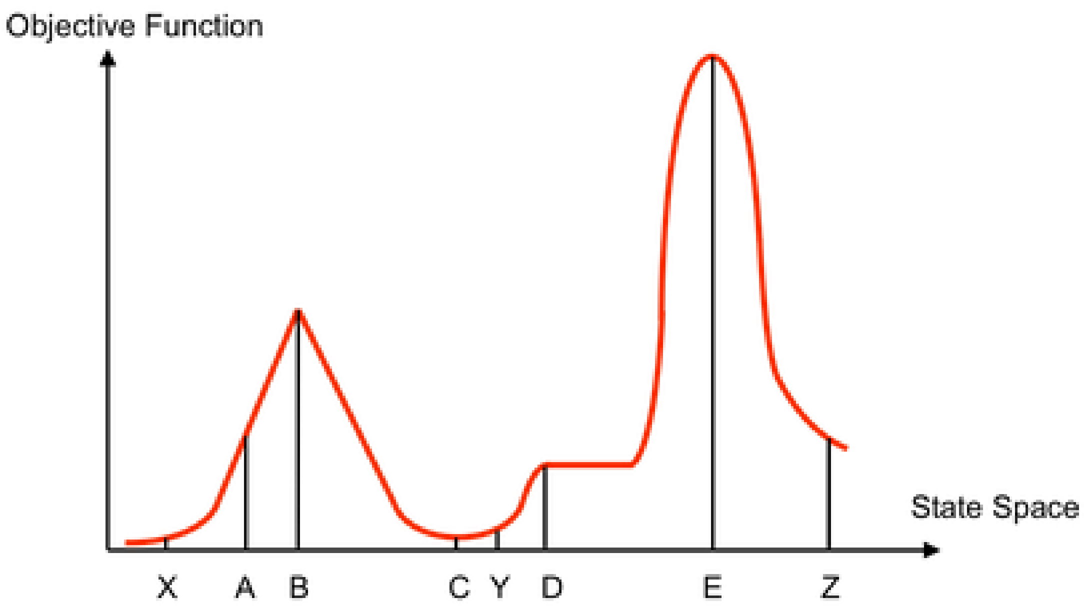

Starting from X, where do you end up ?

Starting from Y, where do you end up ?

Starting from Z, where do you end up ?

Advanced Local Search

Simulated Annealing

- Idea: Allow for movement farther away than immediate neighbours
  - Allow for negative eval choices as long as they are below a threshold T of change
- Use declining T (temperature) for controlling negative difference movements

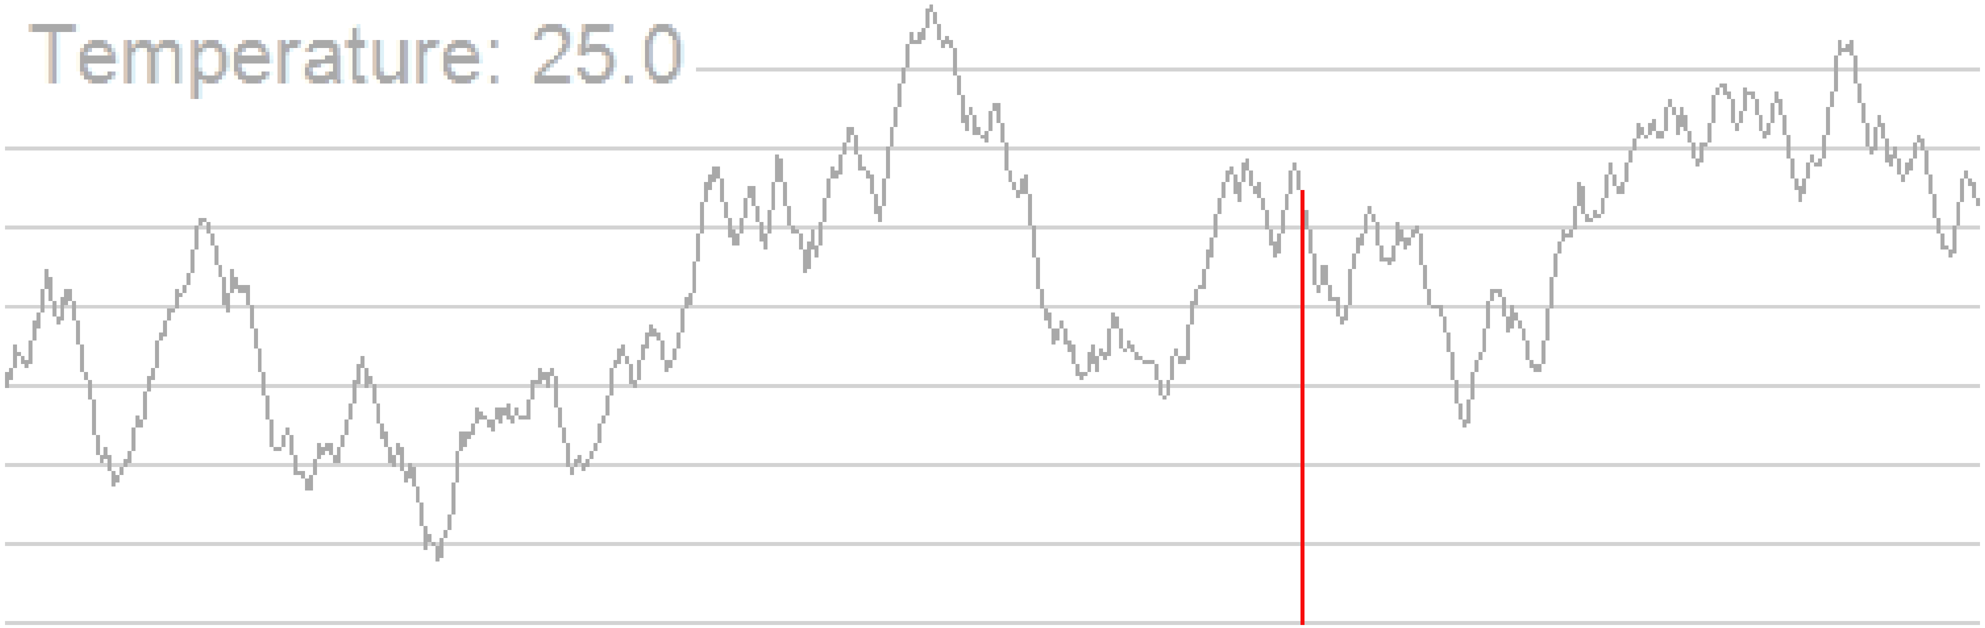

Simulated Annealing

- Theoretical guarantee:
  - If ‘Temperature’ decreased slowly enough, will converge to optimal state!

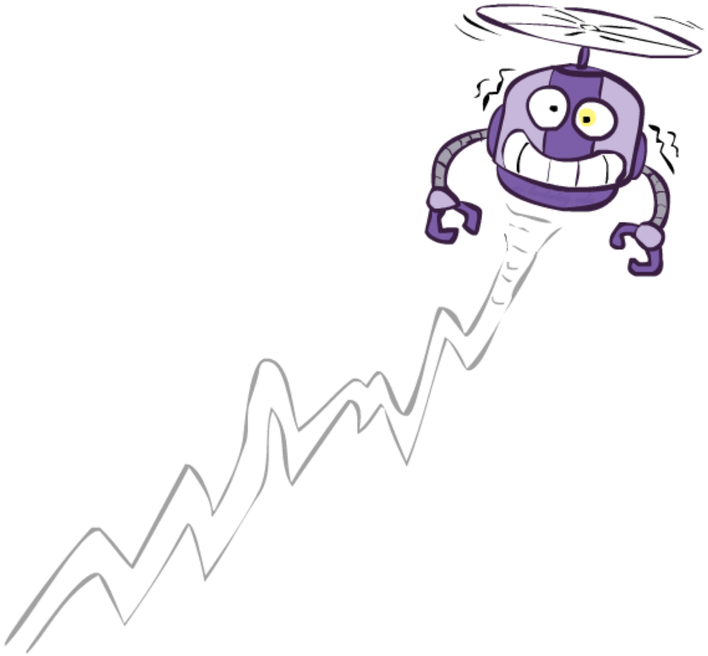

- Is this an interesting guarantee?
  - **Slowly** may be a long long long long long time!!!!!!
  - The more downhill steps you need to escape a local optimum, the less likely you are to ever make them all in a row

UNIVERSITY OF CALGARY

Particle Swarm Optimization

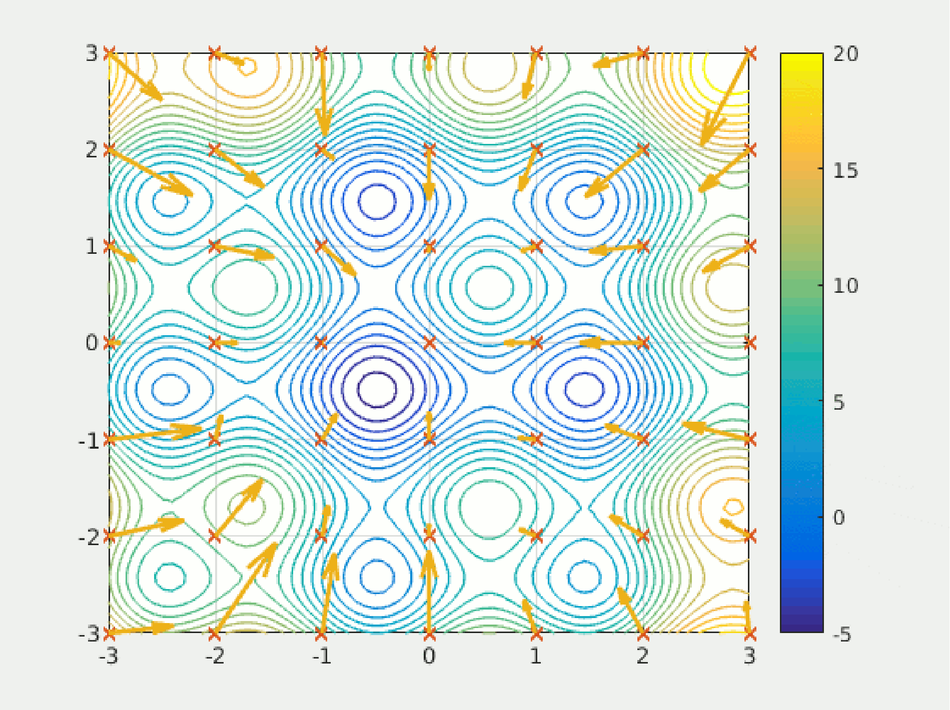

- Multiple solutions to problem (each is a particle)
- Start random created with a velocity in solution space
- Think of particles as having 'gravity' related to eval
- Each step, particles update based on their velocity and the influence of the 'gravity' of other particles with good solutions
- Often cooling principle included to help find best at end.
- Good at numeric problems with smooth evaluation spaces
- Challenges with discrete problems.

Genetic Algorithms

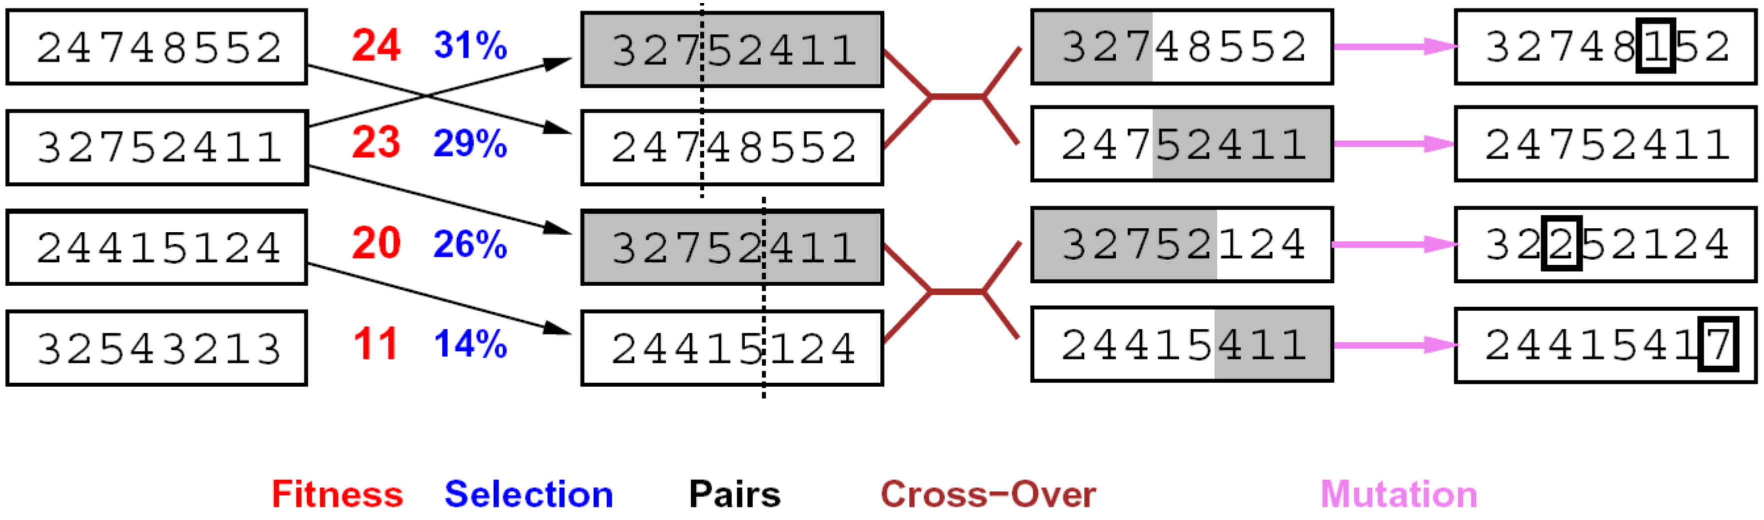

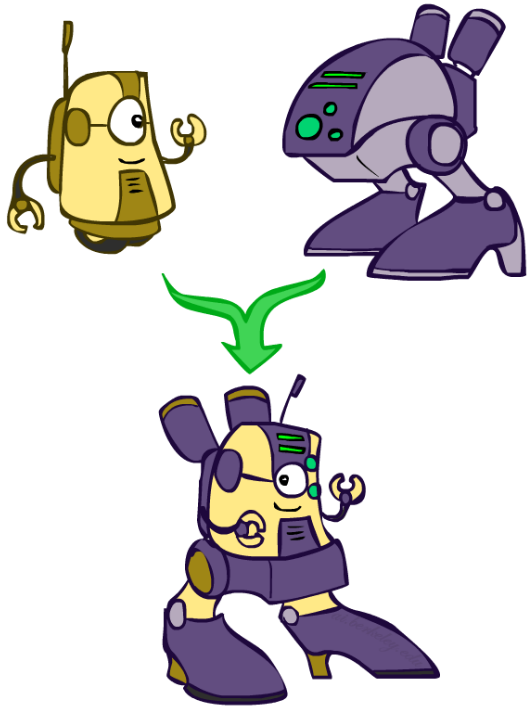

- Based on natural selection in populations
- Start with multiple solutions to problem
- Create next generation by combining ‘DNA’ of the previous selected by fitness quality
    - Crossover operators (two parents) and mutation operators
- More suited to discrete problems than particle swarm optimization

UNIVERSITY OF CALGARY

Search Summary

Search and Models

- Search generally operates over models of the world
  - The agent doesn't actually try all the plans out in the real world!
  - Planning is all "in simulation"
  - Your search is only as good as your models…
- Some search uses real-world evaluation
  - This can be fun but often really time consuming

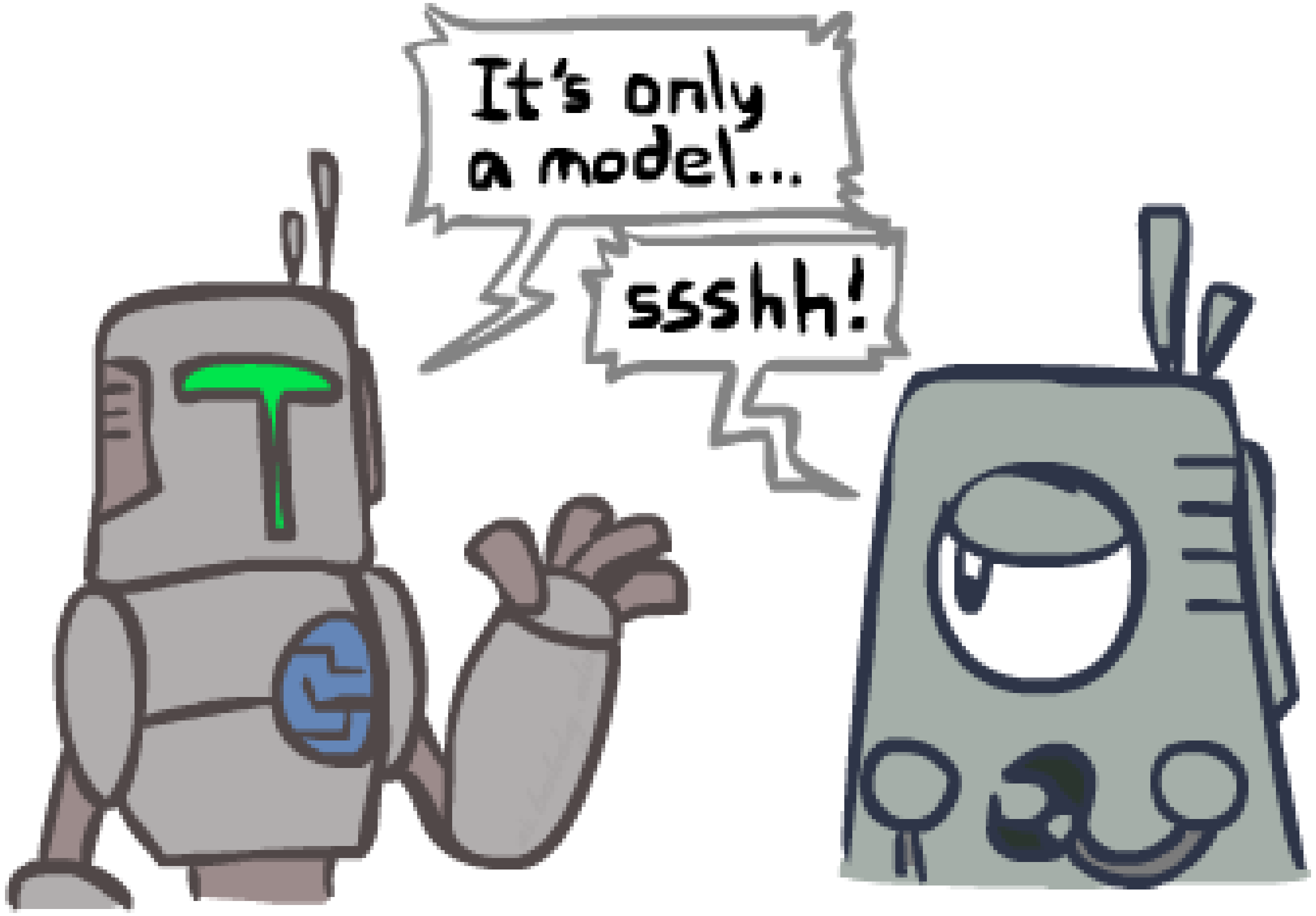

Search Gone Wrong?

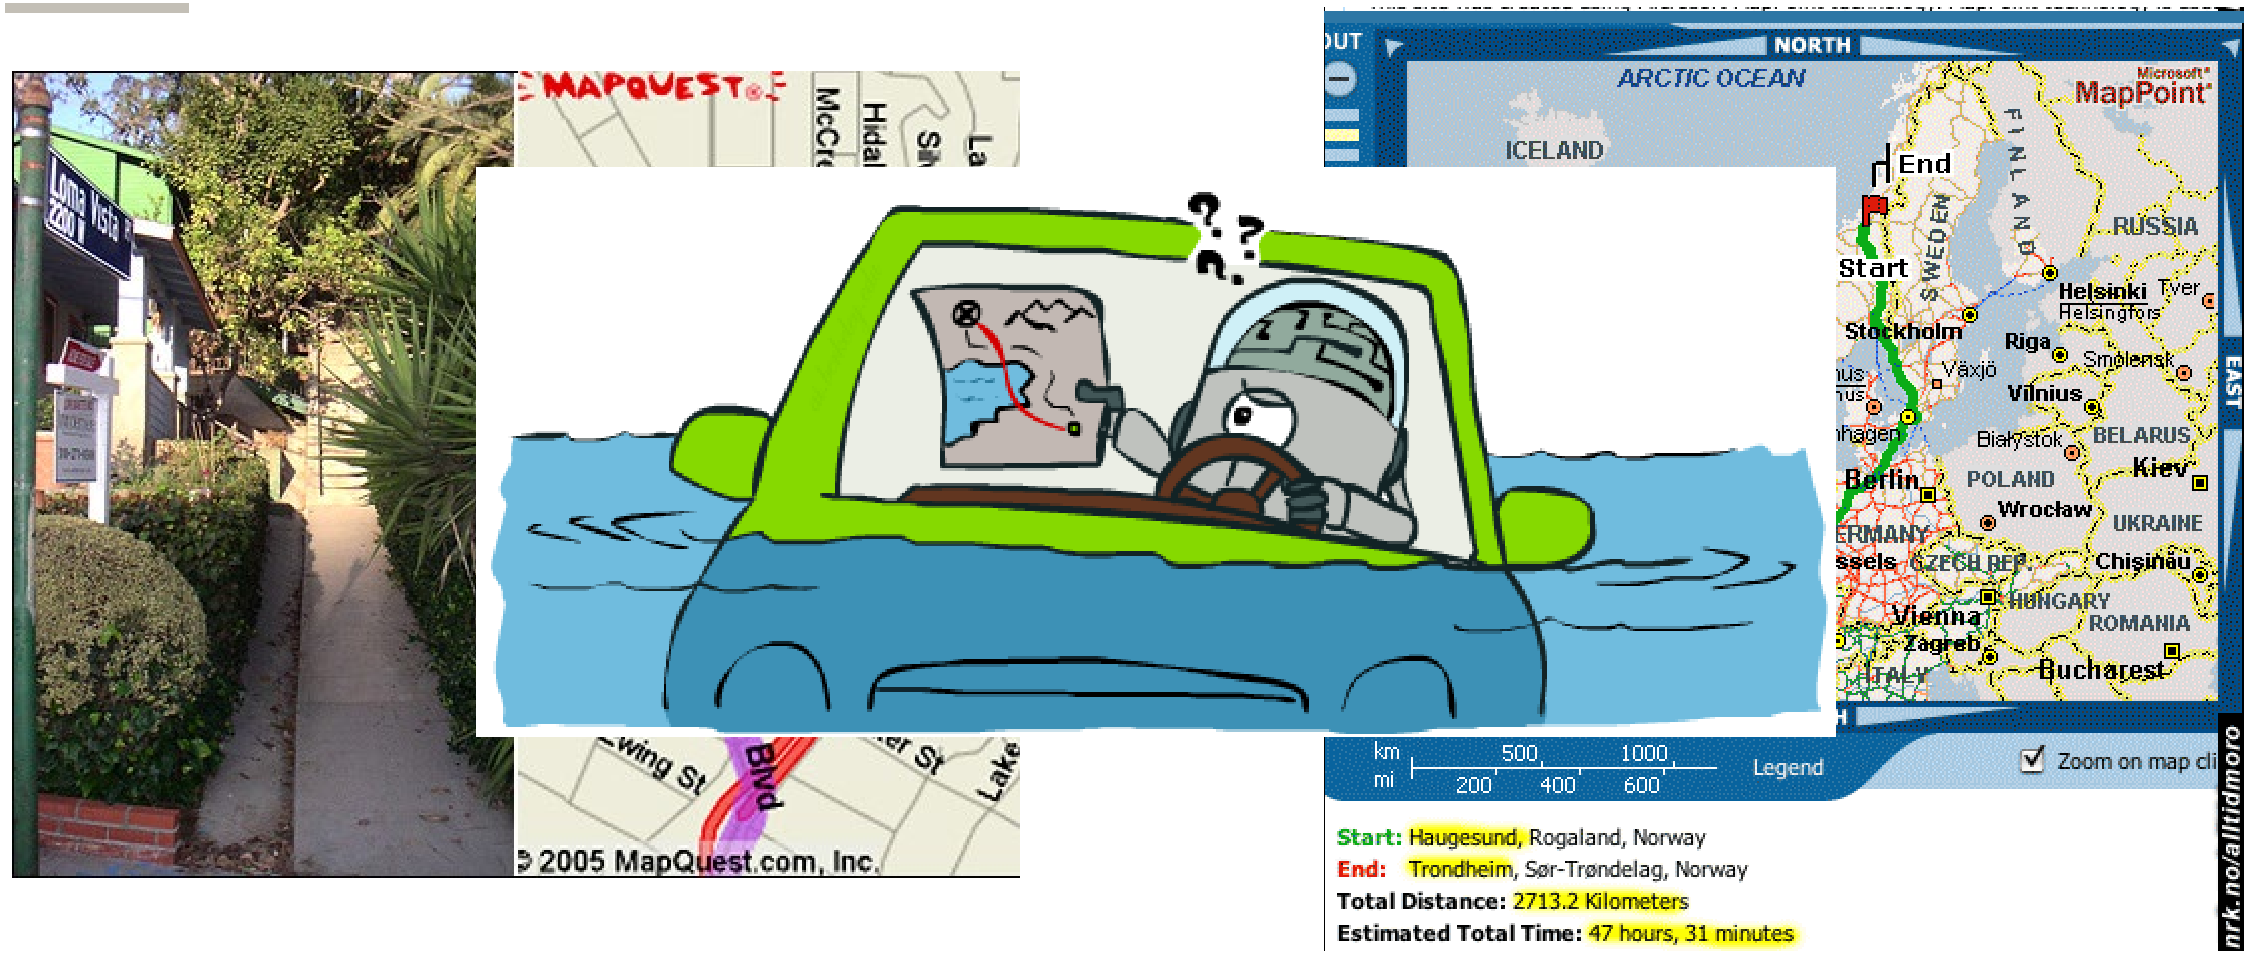

Next…ethics, legality & society

Jonathan Hudson, Ph.D.
jwhudson@ucalgary.ca
https://cspages.ucalgary.ca/~jwhudson/

UNIVERSITY OF CALGARY