

ROB 537: Learning-Based Control

Week 2, Lecture 1
Search and Optimization

Announcements:
HW 1 Due TODAY
Project Background due on 10/9

Reading assignment:
EA survey paper (Fleming)



Search - Optimization

- Basic premise:
 - A set of variables $\mathbf{x} = \{x_1, \dots, x_N\}$ (or states)
 - A cost function $C(\mathbf{x})$
 $E(\mathbf{x}), J(\mathbf{x}), G(\mathbf{x}) \dots$
 - A set of constraints H , on the possible values of \mathbf{x}
- Find the variables \mathbf{x}' such that:
 - \mathbf{x}' satisfies H ; and
 - And $C(\mathbf{x}') > C(\mathbf{x})$ for all \mathbf{x} (for maximization)



Search - Optimization

- Two fundamentally different types of problems:
 - Cost is an analytic function of x
 - Example: Minimize $C(x) = 7x^4 + 3.4x^3 + 3x^2 + 4x$
 - Unconstrained: gradient descent
 - Constrained: Lagrange multipliers



Search - Optimization

- Two fundamentally different types of problems:
 - Cost is an analytic function of x
 - Example: Minimize $C(x) = 7x^4 + 3.4x^3 + 3x^2 + 4x$
 - Unconstrained: gradient descent
 - Constrained: Lagrange multipliers
 - Cost is function of final state
 - Example: Traffic congestion management
 - Many vehicles
 - Roads of given capacity
 - Problem: determine departure times and path to minimize congestion



Optimization is Easy

$$F(x)$$



Optimization is Easy

$$F(\quad x \quad)$$



Optimization is Easy

$$F(t \vec{x})$$



Optimization is Easy

$$F(t \vec{x} \vec{y} \vec{z})$$



Optimization is Easy

$$F(g(t, \vec{x}, \vec{y}, \vec{z}))$$

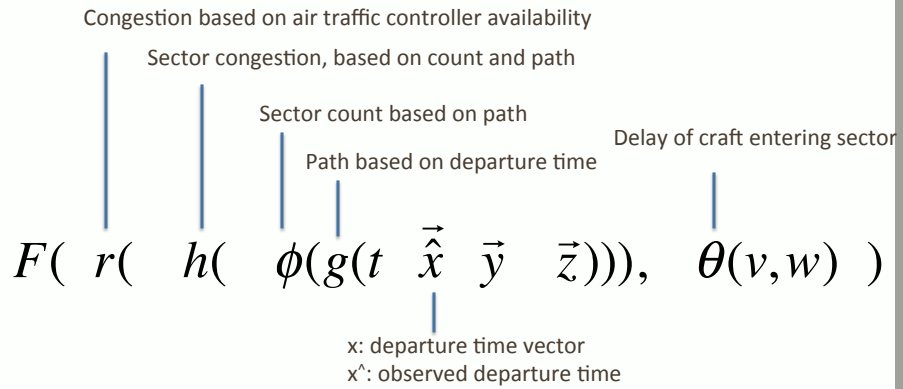


Optimization is Easy

$$F(h(\phi(g(t, \vec{x}, \vec{y}, \vec{z})), \theta(v)))$$



Optimization Maybe is not so Easy



Search Problem: Bin Packing

- Bin Packing
 - Given object sizes, bin sizes.
 - States: contents of each bin
 - Example: $\mathbf{x} = \{1.2, 0, 2.5, 4.1, 3.2, 0, 0.5\}$
 - Cost function: $C(\mathbf{x}) =$ number of nonzero bins as a function of \mathbf{x}
 - Constraints: each bin must be below size
 - $\{1.2, 0, 2.5, 4.1, 3.2, 0, 0.5\}$ is not valid if bin size is 4.
 - Find object distribution \mathbf{x}' such that: $C(\mathbf{x}') < C(\mathbf{x})$ for all \mathbf{x}



Search Problem: Traveling Salesman

- Traveling Salesman Problem:
 - You need to visit N cities. You want to minimize the distance traveled but still see each city once and return to the starting city.
 - Given distance matrix specifying distance between cities
 - States: order in which a city is visited
 - Example: $\mathbf{x} = \{1,2,5,4,3\}$
 - Cost function: $C(\mathbf{x})$ length of the trip, based on \mathbf{x} and distance matrix
 - Constraints: each city visited exactly once
 - $\{1,1,3,2,1\}$ is not a valid tour
 - Find shortest valid tour \mathbf{x}' such that: $C(\mathbf{x}') < C(\mathbf{x})$ for all \mathbf{x}



Search - Optimization

- Local search
 - Generally applicable to any search problem
 - Use single current state
 - Improves solution at each step
 - Susceptible to local minima
- Example: Traveling Salesman Problem
 - Start with a valid tour.
 - Improve tour by swapping city orders
 - Keep going



Discrete vs. Continuous Spaces

- Discrete Spaces
 - All states can be enumerated
 - 9-puzzle, chess, bin packing
- Continuous Spaces
 - Infinite number of states
 - Use algorithms that do not require enumerating states
 - Discretize neighborhood of each variable
- Most real world problems are continuous



Local Search in Discrete Spaces

- Single solution-based search:
 - Hill climbing
 - Stochastic hill climbing (pick among “good” moves)
 - ϵ -greedy search
 - Simulated annealing
 - Tabu Search
- Population-based search:
 - Beam search
 - Stochastic beam search
 - Simplex Search
 - Particle Swarm Optimization
 - Evolutionary algorithms
 - Genetic Algorithms



Hill Climbing in Discrete Spaces

- Problem Setup:
 - Objective function
 - States
- Algorithm:
 1. Start from an initial (random) state
 2. Generate (all) successor states from current state
 3. Select successor state with highest value
 4. Go to step 2
- Issues:
 - Local maxima
 - Stop/go when reaching a plateau?



Stochastic Hill Climbing in Discrete Spaces

- Problem Setup:
 - Objective function
 - States
- Algorithm:
 1. Start from an initial (random) state
 2. Generate (all) successor states from current state
 3. Select a random successor state of higher value than current state
 4. Go to step 2
- Issues:
 - Slower convergence than hill-climbing
 - Potentially better solution than hill-climbing



First-Choice Hill Climbing

- Problem Setup:
 - Objective function
 - States
- Algorithm:
 1. Start from an initial (random) state
 2. Generate successor state randomly from current state
 3. Select first successor state of higher value than current state
 4. Go to step 2
- Issues:
 - Implementation of stochastic hill climbing
 - Much more efficient than stochastic hill climbing when there are many successor states




Tabu Search

- Problem Setup:
 - Objective function
 - States
- Algorithm:
 1. Start from an initial (random) state
 2. Randomly generate many candidate successor states "close" to current solution
 3. Check Tabu status of each
 4. Choose the best admissible successor state (worst solutions allowed)
 5. Update Tabu Status
 6. Goto 2 or stop if criteria is met
- Issues:
 - What if all successors are Tabu?
 - Domain knowledge intensive

Tabu Status:
Solutions near recently
explored solutions are
not allowed



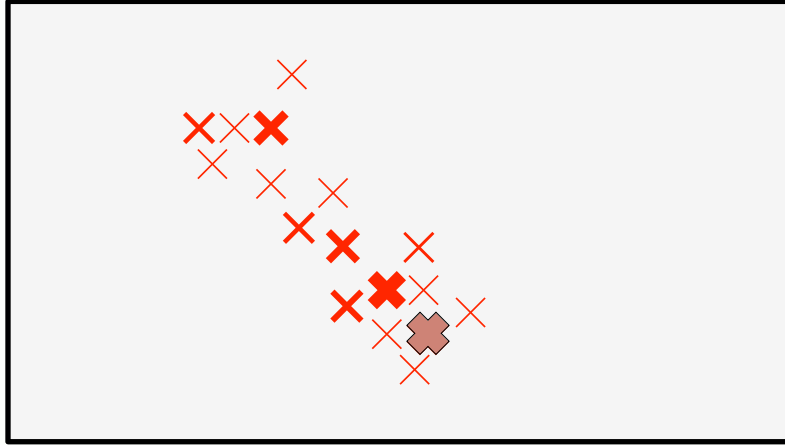
Tabu Search



Simulated Annealing

- Problem Setup:
 - Objective function
 - States
- Algorithm:
 1. Start from an initial (random) state
 2. Update Parameter T (temperature)
 3. Randomly generate a successor state from current state
 4. If successor is better than current state:
 - select it as the next state
 - Go to step 2
 5. If successor state is NOT better than current state:
 - Select it with probability p , where: $p \propto e^{-\Delta E / T}$
 - Go to step 2
- Issues:
 - Accept bad solution some of the time (based on how bad they are)
 - Accept bad solutions early
 - If T is reduced slowly enough simulated annealing finds global optimum

Simulated Annealing



ϵ -Greedy Search

- Problem Setup:
 - Objective function
 - States
- Algorithm:
 1. Start from an initial (random) state
 2. Randomly generate a successor state from current state
 3. If successor is better than current state:
 - select it as the next state
 - Go to step 2
 4. If successor state is NOT better than current state:
 - Select it with probability ϵ
 - Go to step 2
- Issues:
 - Accept bad solution some of the time (NOT based on how bad they are)
 - More efficient than simulated annealing (no exponential computation)
 - Can modify ϵ if desired

Population Based Search

- All previous algorithms were based on modifying a single solution
- What if you store and search through multiple solutions at once ?

Population based search



Beam Search in Discrete Spaces

- Problem Setup:
 - Objective function
 - States
- Algorithm:
 1. Start from k initial (random) states
 2. Generate all successor states for all current states
 3. Select k best states from the pool
 4. Go to step 2
- Issues:
 - Not the same as running hill climbing k times
 - “Information” passed among k states (not directly)
 - Lack of diversity among the k states
 - States quickly concentrate on specific areas of the state space



Stochastic Beam Search in Discrete Spaces

- Problem Setup:
 - Objective function
 - States
- Algorithm:
 1. Start from k initial (random) states
 2. Generate all successor states for all current states
 3. Select k states randomly from the pool, with value of state affecting its selection probability
 4. Go to step 2
- Issues:
 - Useful information passed among k states
 - Keeps more variation among k states
 - Most basic “population-based” search algorithm




Simplex Search

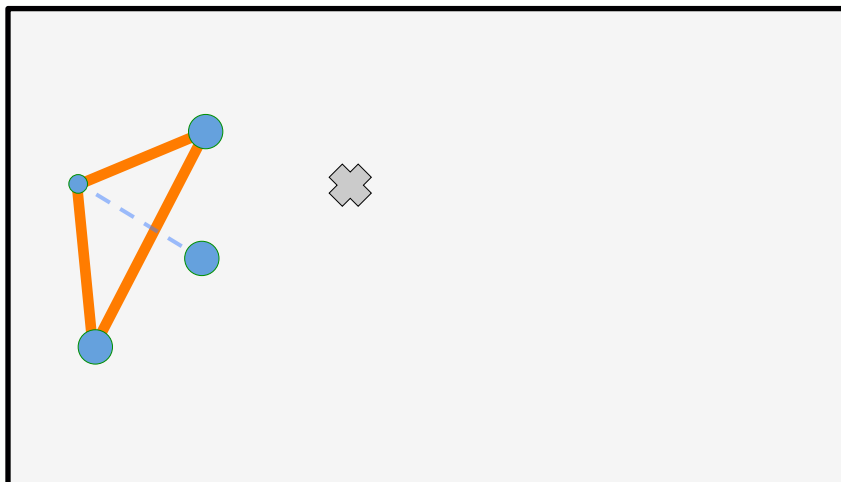
- Problem Setup:
 - Objective function
 - States
 - Scaling factor $F=[1:2]$
- Algorithm:
 1. Start from $\{k = size(x)+1\}$ initial (random) states
 2. Evaluate all k states
 3. Discover worst solution s^{\sim}
 4. Calculate difference between $x(s^{\sim})$ and centroid of all other solutions $x(s^k)$
 5. $\Delta = x(s^{\sim}) - x(s^k)$
 6. Create $x(s') = x(s^{\sim}) - F*\Delta$
 7. Stop after desired number of iterations
- Issues:
 - Scaling in decision variables
 - Deterministic solution




Simplex Search



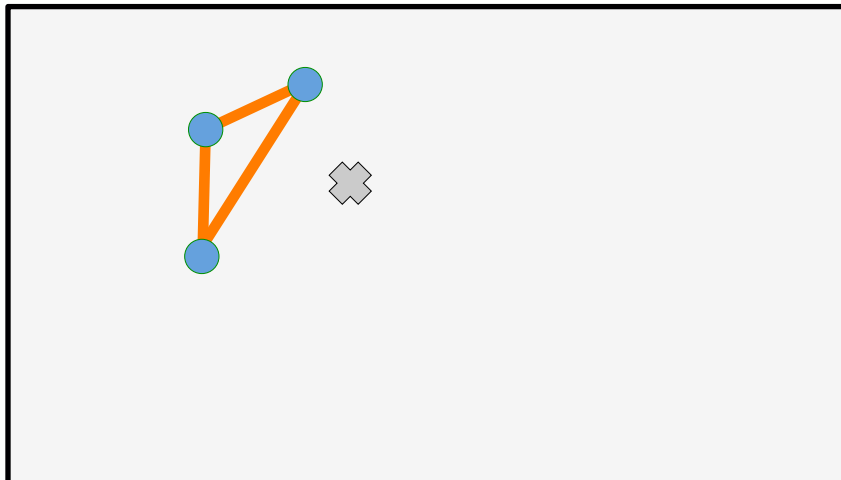
Simplex Search




Simplex Search




Simplex Search



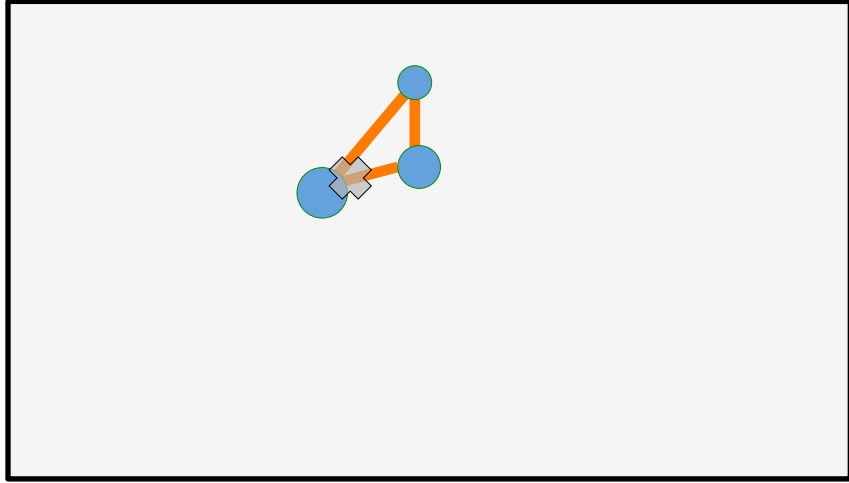
Simplex Search




Simplex Search



Simplex Search



Simplex Search




Particle Swarm Optimization

- Problem Setup:
 - Objective function
 - States
- Algorithm:
 1. Start from n initial (random) states (particles with velocities)
 2. Choose global, personal weights ϕ_g ϕ_p
 3. For each particle
 1. Pick random numbers p, g in $[0:1]$
 2. Update velocity: $v = \omega v + p \phi_p (x - b_p) + g \phi_g (x - b_g)$
 3. Update position: $x = x + v$
 4. Track best known position b_p
 4. Track best known global position b_g
 5. Repeat steps 3,4 until convergence or stopping criteria is met
- Issues:
 - Parameters: Velocity, Momentum



Particle Swarm Optimization




Evolutionary Algorithms

- Problem Setup:
 - Objective function
 - States
- Algorithm:
 1. Start from k initial (random) states
 2. Generate k successor states randomly
 3. Perturb those states (mutation)
 4. Select k states from the pool of 2k states, with value of state affecting its selection probability
 5. Go to step 2
- Issues:
 - Addition of perturbation generates “new” solutions
 - Slower convergence than stochastic beam search
 - Potential for finding better solution due to “mutation”
 - Steps 2 and 3 are indistinguishable in most implementations

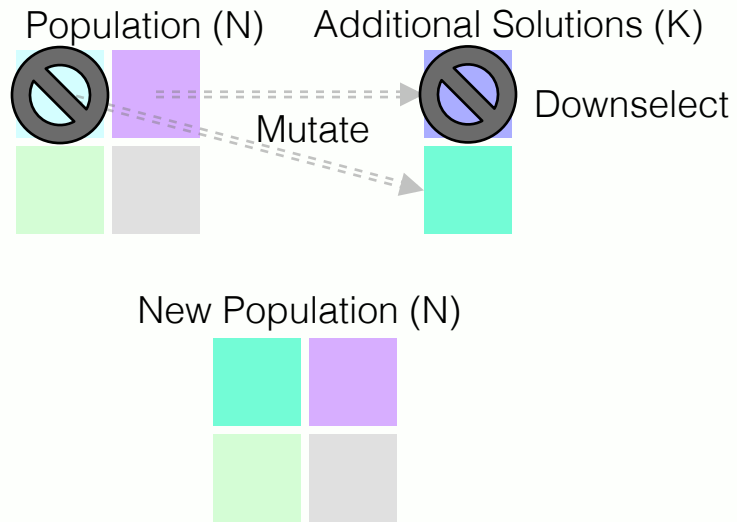


Evolutionary Algorithms

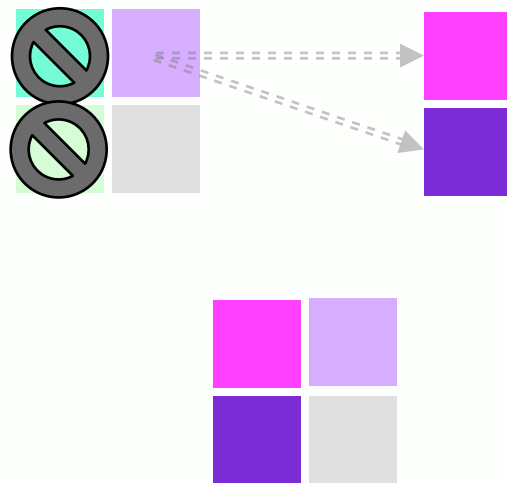
Population (N)





Evolutionary Algorithms





Evolutionary Algorithms





Evolutionary Algorithms





Evolutionary Algorithms





Evolutionary Algorithms




Evolutionary Algorithms



Evolutionary Algorithms



Evolutionary Algorithms



Genetic Algorithms

- Problem Setup:
 - Objective function
 - States
- Algorithm:
 1. Start from k initial (random) states
 2. Generate k new states using current states (crossover: “merge” parts of two states to generate a new one)
 3. Perturb those states (mutation)
 4. Select k states randomly from the pool, with value of state affecting its selection probability
 5. Go to step 2
- Issues:
 - Crossover aims to preserve and combine good “partial” states
 - Addition of perturbation generates “new” solutions
 - Potential for finding better solution due to “mutation”
 - Genetic algorithms provide no “mathematical advantage”
 - Genetic algorithms may provide an implementation advantage: preserving good partial solutions allows for a “hierarchical search”

