Artificial Intelligence

Fiona Yan Liu

Department of Computing

The Hong Kong Polytechnic University

Learning Outcomes of Search Agents

- Informed search
  - Also called blind search
  - The strategies have no additional information about states beyond that provided in the problem definition
  - Depth-first search
  - Breadth-first search
  - Uniform-cost search
- Informed search
  - Also called heuristic search
  - The strategies know whether one non-goal node is “more promising” than another
  - Greedy best-first search
  - A* search
- Reference reading
  - Chapter 3

Admissible Heuristic of A* Search

- The tree-search version of A* is optimal if $h(n)$ is admissible
  - Evaluation function: $f(n) = g(n) + h(n)$
  - $g(n)$: the cost so far to reach the node n
  - $h(n)$: the estimated cost to goal from the node n
  - $f(n)$: the estimated total cost of path through the node n to goal
- An admissible heuristic never overestimates the cost to reach the goal
  - A heuristic $h(n)$ is admissible if for every node n, $h(n) \leq h^*(n)$, where $h^*(n)$ is the true cost to reach the goal state from n.

Domination of Heuristic Functions

- If $h_2(n) \geq h_1(n)$ for all n (both admissible)
  - then h_2 dominates h_1, which indicates that h_2 is better for search
- Domination translates directly into efficiency
  - A* using h_2 will never expand more nodes than A* using h_1.

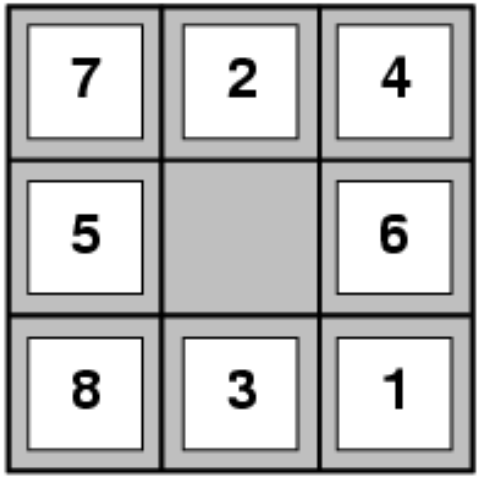

Start State

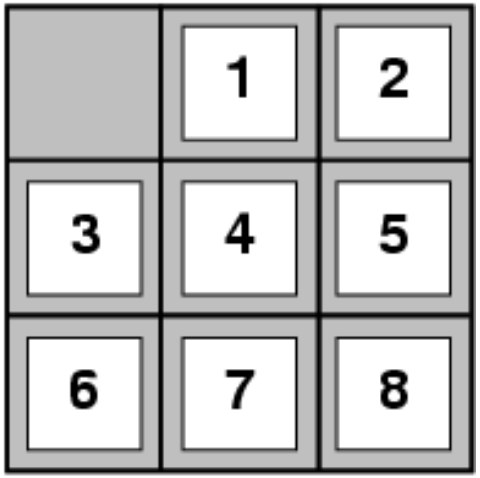

Goal State

- $h_1(S) = 8$
- $A^*(h_1) = 39{,}135$ nodes
- $h_2(S) = 18$
- $A^*(h_2) = 1{,}641$ nodes

N-queens Problem

- Put n queens on an $n \times n$ board with no two queens on the same row, column, or diagonal
- Move a queen to reduce number of conflicts

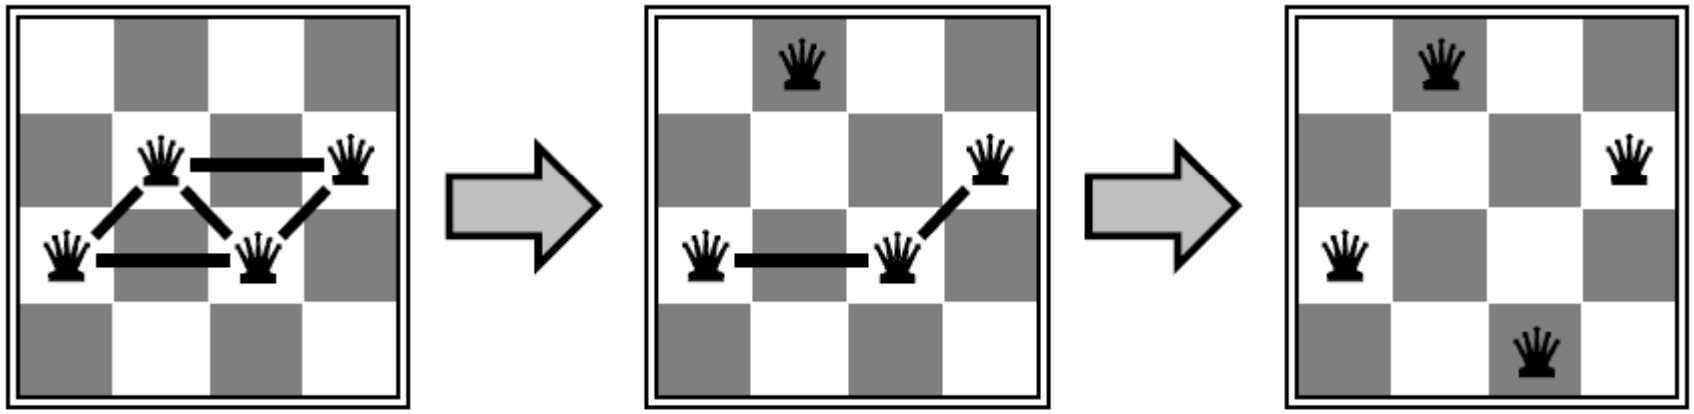

8-Queens Problem

- Place 8-queens on a chessboard such that no queen attacks any other
  - A queen attacks any piece in the same row, column, or diagonal

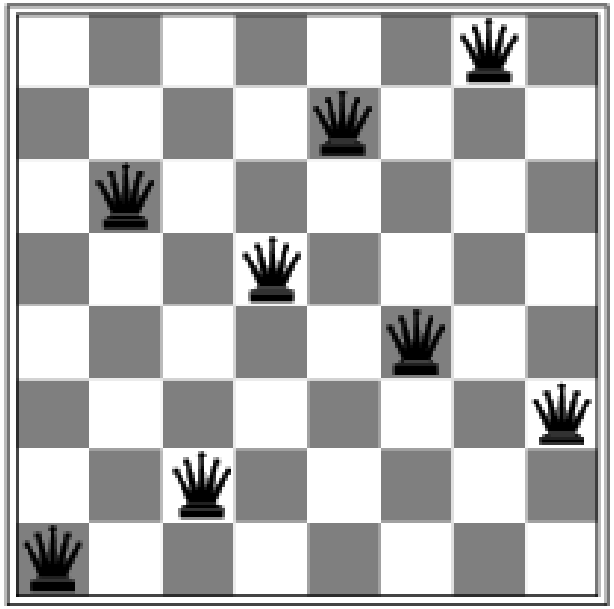

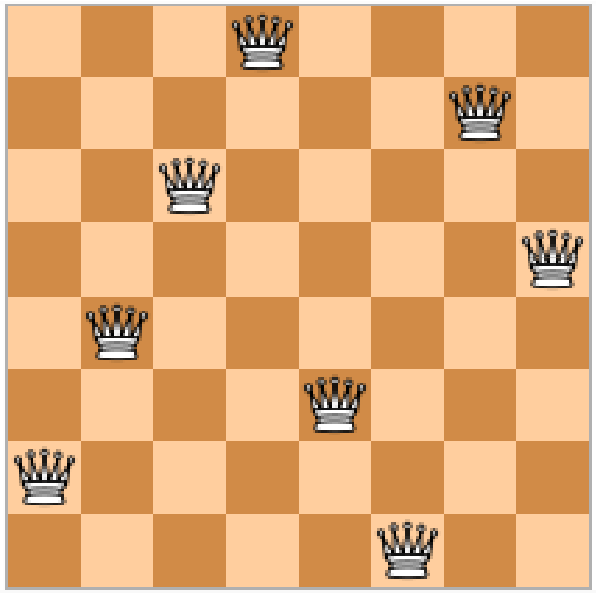

Local Search

- Problem: Romania
  - The search algorithms explore search space systematically
  - When a goal is found, the path to the goal also constitutes a solution to the problem
- Problem: N-Queens
  - The path to the goal is irrelevant
- Local search
  - Operate using a single current node
  - Generally move only to neighbors of that node
  - Useful to find the best state according to an objective function

State-space Landscape

- Location: defined by the state
- Evaluation: defined by the value of the objective function

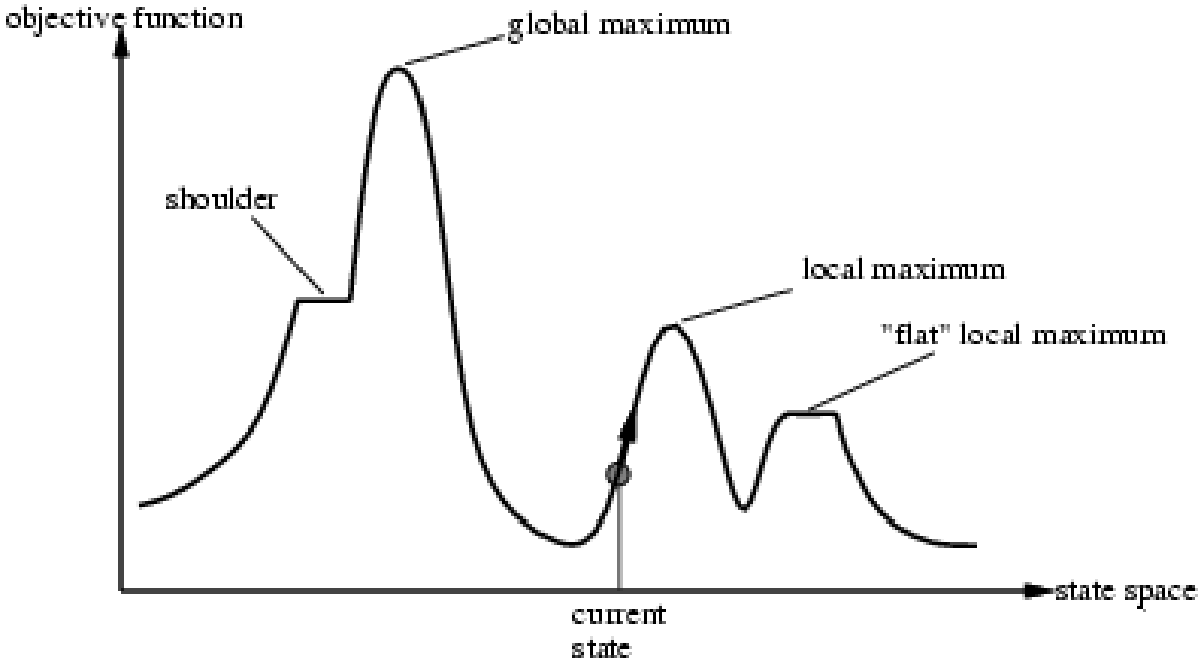

Hill-climbing Search

- Depending on initial state, can get stuck in local maxima

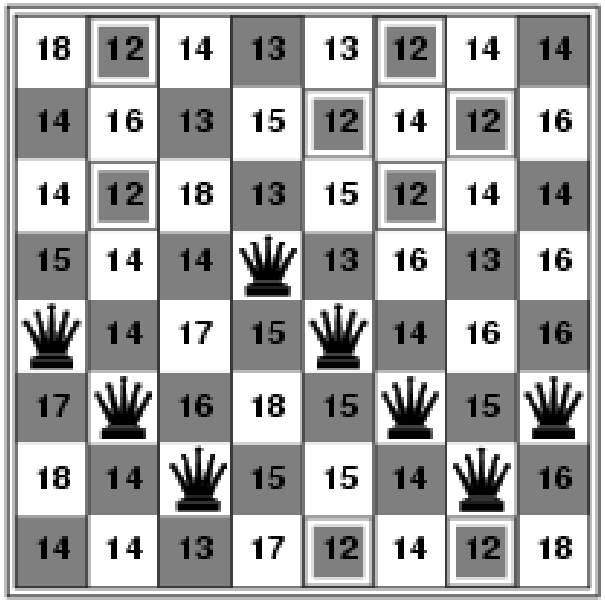

Arthur Samuel

- Professor at Stanford University
- Proposed hill-climbing search in 1963
  - Made the first computer checker program on IBM's first commercial
- Applications of hill-climbing search
  - Network flow problem
  - VLSI(Very Large Scale Integration) design

Simulated Annealing Search

- Hill-climbing
  - Never makes downhill moves
  - Get stuck on a local maximum
- Random walk
  - Moving to a successor chosen uniformly at random
  - Is complete but extremely inefficient
- Simulated annealing
  - Escape local maxima by allowing some "bad" moves but gradually decrease their frequency
  - In metallurgy, annealing is the process used to temper or harden metals and glass by heating them to a high temperature and then gradually cooling them
  - The algorithm is quite similar to hill-climbing
    - Instead of picking the best move, it picks a random move
    - If the move improves the situation, it is always accepted

Simulated Annealing Search

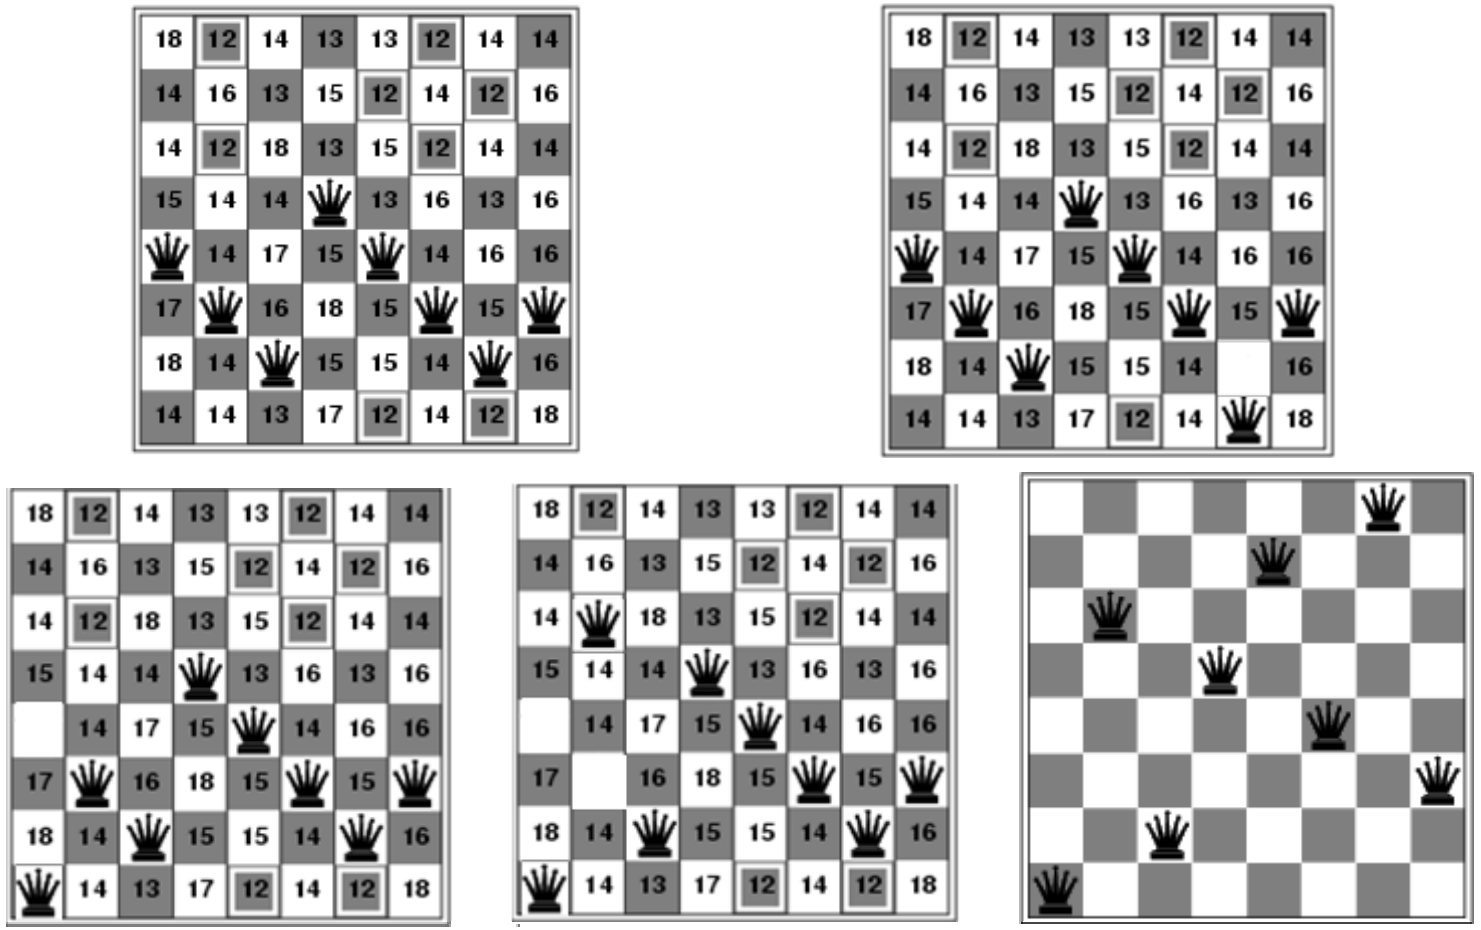

Simulated Annealing Search Proposed in 1983

- Scott Kirkpatrick
  - Ph.D. in Physics from Harvard University
- C. Daniel Gelatt, JR
  - PhD in Physics from Harvard University
- Mario P. Vecchi
  - PhD in Electrical Engineering from Massachusetts Institute of Technology