Genetic Algorithms and Genetic Programming Lecture 13

Gillian Hayes

10th November 2006

Ant Colony Optimisation

- Ant Colony Optimisation - nature
- Artificial Ant Systems
- Applying Ants to TSP
- Local versus global updating
- Adding local search
- Next: Ants applied to other problems

Ant Colony Optimisation

Biological inspiration: ants find the shortest path between their next and a food source using **pheromone trails**.

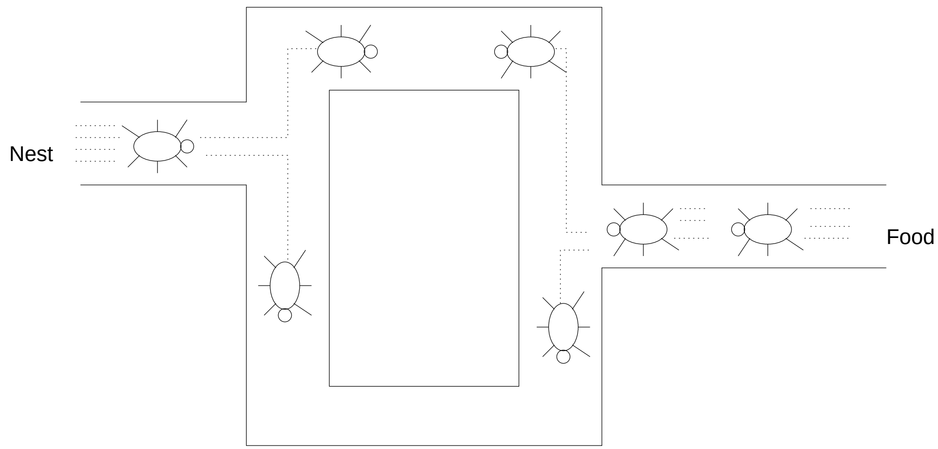

Ant Colony Optimisation is a population-based search technique for the solution of combinatorial optimisation problems which is inspired by this behaviour.

Marco Dorigo. PhD thesis 1992. Ant Systems Based on theoretical biology work of Deneubourg 1983 (and earlier)

What, how?

- Real ants find shortest routes between food and nest
- They hardly use vision (almost blind)
- They lay pheromone trails, chemicals left on the ground, which act as a signal to other ants – **STIGMERGY**
- If an ant decides, with some probability, to follow the pheromone trail, it itself lays more pheromone, thus reinforcing the trail.
- The more ants follow the trail, the stronger the pheromone, the more likely ants are to follow it.
- Pheromone strength decays over time.
- Pheromone builds up on shorter path faster, so ants start to follow it.

Pheromone Builds Up on Short Paths

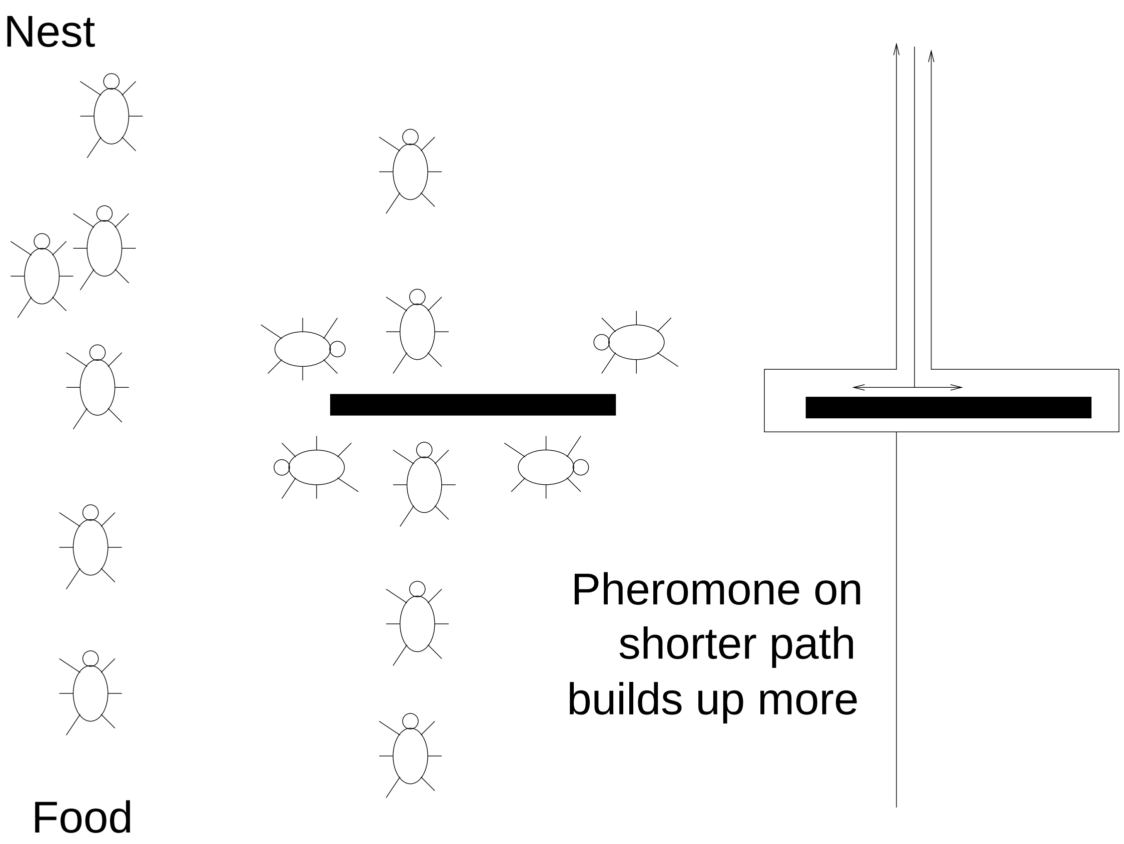

Artificial Ant Systems

- Do have some memory
- Are able to sense "environment" if necessary
- Use discrete time
- Are optimisation algorithms

So can we apply them to an optimisation problem: Travelling Salesman Problem

Ant System for the TSP: An Example

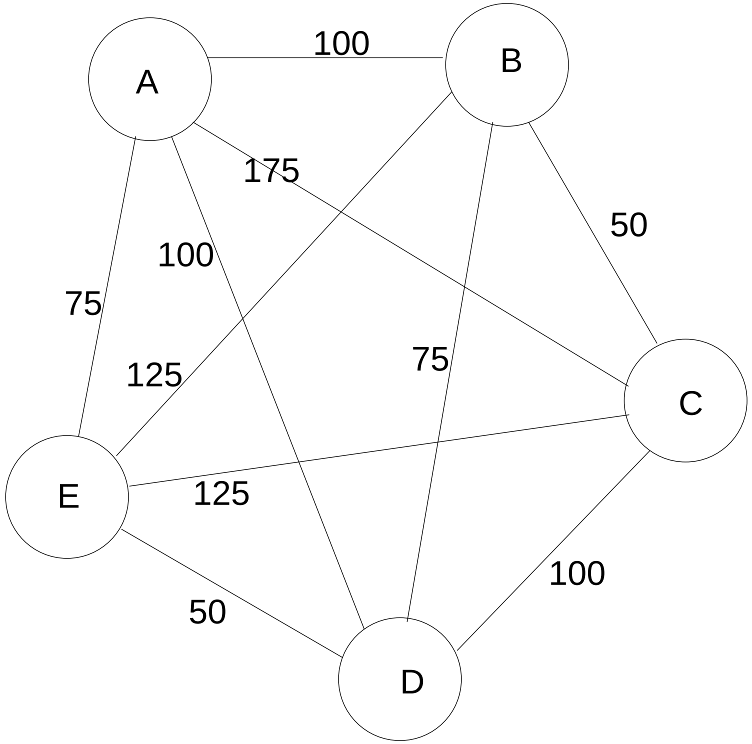

Find the tour that minimises the distance travelled in visiting all towns.

Ant System for the TSP

- Each ant builds a tour from a starting city
- Each ant chooses a town to go to with a probability: function of the town's distance and the amount of pheromone on the connecting edge
- Legal tours: transitions to already visited towns disallowed till tour complete (keep a tabu list)
- When tour completed, lay pheromone on each edge visited
- Next city j after city i chosen according to Probability Rule

Probability Rule

$$p(i,j) = \frac{[\tau(i,j)].[\eta(i,j)]^{\beta}}{\sum_{g \in \text{allowed}} [\tau(i,g)].[\eta(i,g)]^{\beta}}$$

- Strength of pheromone $\tau(i,j)$ is favourability of j following i, β is a constant, e.g. 2. Emphasises "**global** goodness"

- Visibility $\eta(i,j) = 1/d(i,j)$ is a simple heuristic guiding construction. In this case it's greedy – the nearest town is the most desirable (seen from a **local** point of view)

- $\sum_{g \in \text{allowed}}$: normalise over all the towns g that are still permitted to be added to the tour, i.e. not on the tour already

Pheromone

- Pheromone trail evaporates a small amount after every iteration

$$\tau(i,j) = \rho.\tau(i,j) + \Delta\tau_{ij}$$

  where $0 < \rho < 1$ is an evaporation constant

- The density of pheromone laid on edge (i,j) by the m ants at that timestep is

$$\Delta\tau_{ij} = \sum_{k=1}^{m} \Delta\tau_{ij}^{k}$$

- $\Delta\tau_{ij}^{k} = Q/L_k$ if kth ant uses edge (i,j) in its tour, else 0. Q is a constant and L_k is the length of k's tour. Pheromone density for k's tour.

- Initialise: set pheromone strength to a small value
- Transitions chosen to trade off visibility (choose close towns with high probability – greedy) and trail intensity (if there's been a lot of traffic the trail must be desirable).
- In one iteration all the ants build up their own individual tours (so an iteration consists of lots of moves/town choices/timesteps) and pheromone is laid down once all the tours are complete

Algorithm

- Position ants on different towns, initialise pheromone intensities on edges.
- Set first element of each ant's tabu list to be its starting town.
- Each ant moves from town to town according to the probability $p(i, j)$
- After n moves all ants have a complete tour, their tabu lists are full; so compute L_k and $\Delta\tau_{ij}^k$. Save shortest path found and empty tabu lists. Update pheromone strengths.
- Iterate until tour counter reaches maximum or until *stagnation* – all ants make same tour.

Can also have different pheromone-laying procedures, e.g. lay a certain quantity of pheromone Q at each timestep, or lay a certain density of pheromone Q/d_{ij} at each timestep.

Tweaks

This works well for short tours, say up to about 30 cities

BUT try including exploration and exploitation

ACS algorithm

- ϵ of the time (say 0.1), from city i choose a random allowed city j to move to.
- Rest of time choose greedily among cities using $\max_j p(i,j)$.
- Global updating rule: only globally best ant gets to deposit pheromone at the end of each iteration (i.e. the ant with the shortest tour so far)
- Local updating rule: while building a solution ants change pheromone on edges visited in a similar fashion but with $\Delta\tau(i,j) = \tau_0$, a constant. (Also Ant-Q, with a Q-learning-like update rule.)

Performance

Problem	ACS (avge)	SA (avge)	EN (avge)	SOM (avge)
50-city set 1	5.88	5.88	5.98	6.06
50-city set 2	6.05	6.01	6.03	6.25
50-city set 3	5.58	5.65	5.70	5.83
50-city set 4	5.74	5.81	5.86	5.87
50-city set 5	6.18	6.33	6.49	6.70

ACS – ant colony system, SA–simulated annealing, EN–elastic net, SOM–self-organising map

From Dorigo and Gambardella: Ant Colony System: A cooperative learning approach to the TSP. IEEE Trans. Evol. Comp 1 (1) 53–66 1997.

Can do larger problems, e.g. finds optimal in 100-city problem KroA100, close to optimal on 1577-city problem fl1577.

See TSPLIB for problems:

`http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95/`

But, for any given tour problem, a tailored (ad hoc) heuristic algorithm can often do better. So....

Improvement Through Local Search

- While building tour, apply an improvement heuristic at each step to each ant's partial tour.
- Use 3-opt: cut the tour in three places (remove three links) and attempt to connect up the cities in alternative ways that shorten the path.
- Reduces time, almost always finds optimal path.

Points

Encourage ants to follow best tour by getting best ant to lay pheromone.

Local Updating (the ants lay pheromone as they go along without waiting till end of tour). This “eats away” pheromone, visited edges seen as less desirable, encourages exploration. (Because the pheromone added is quite small compared with the amount that evaporates.)

Heuristic (meta-)improvements like 3-opt – not really “ant”

“Guided parallel stochastic search in region of best tour” [Dorigo and Gambardella]

Reading: 2 Dorigo papers linked to web page

Next:: ACO applied to other problems