

Swarm Intelligence

Jon Timmis
jon.timmis@york.ac.uk



Collective Behaviour



Group coordination

- Collective motion
- Flocking, shoaling, herding, crowd movement, locust plague, etc.
- Example: shoaling
 - The group appears to move in a wave, as if it were a single organism
 - How do they coordinate their behaviour?
- With coordination, local dependence is good (and necessary)

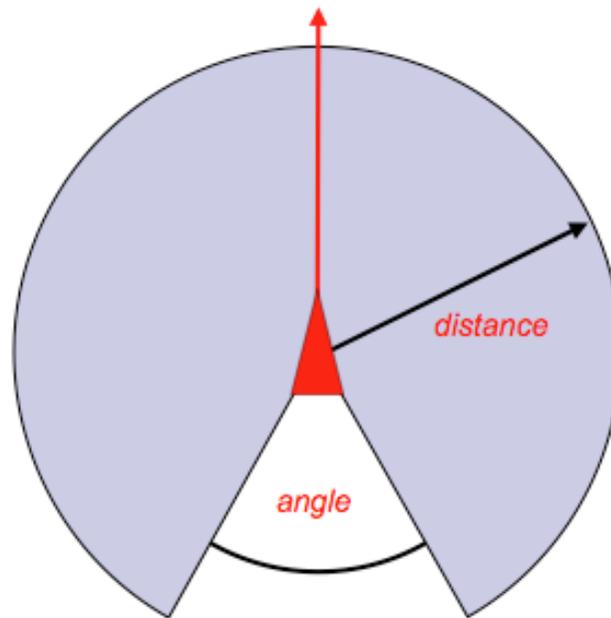


Explanations

- ✗ • Centralised
 - All individuals have knowledge of where all other individuals are going (complete knowledge)
 - The group has a leader who the rest ultimately follow
- ✓ • Reality: decentralised
 - Individuals only know about individuals close to them
 - There is no leader
 - The group coordinates through individuals quickly responding to the movements of a few close neighbours
 - Similar to the way a Mexican wave works

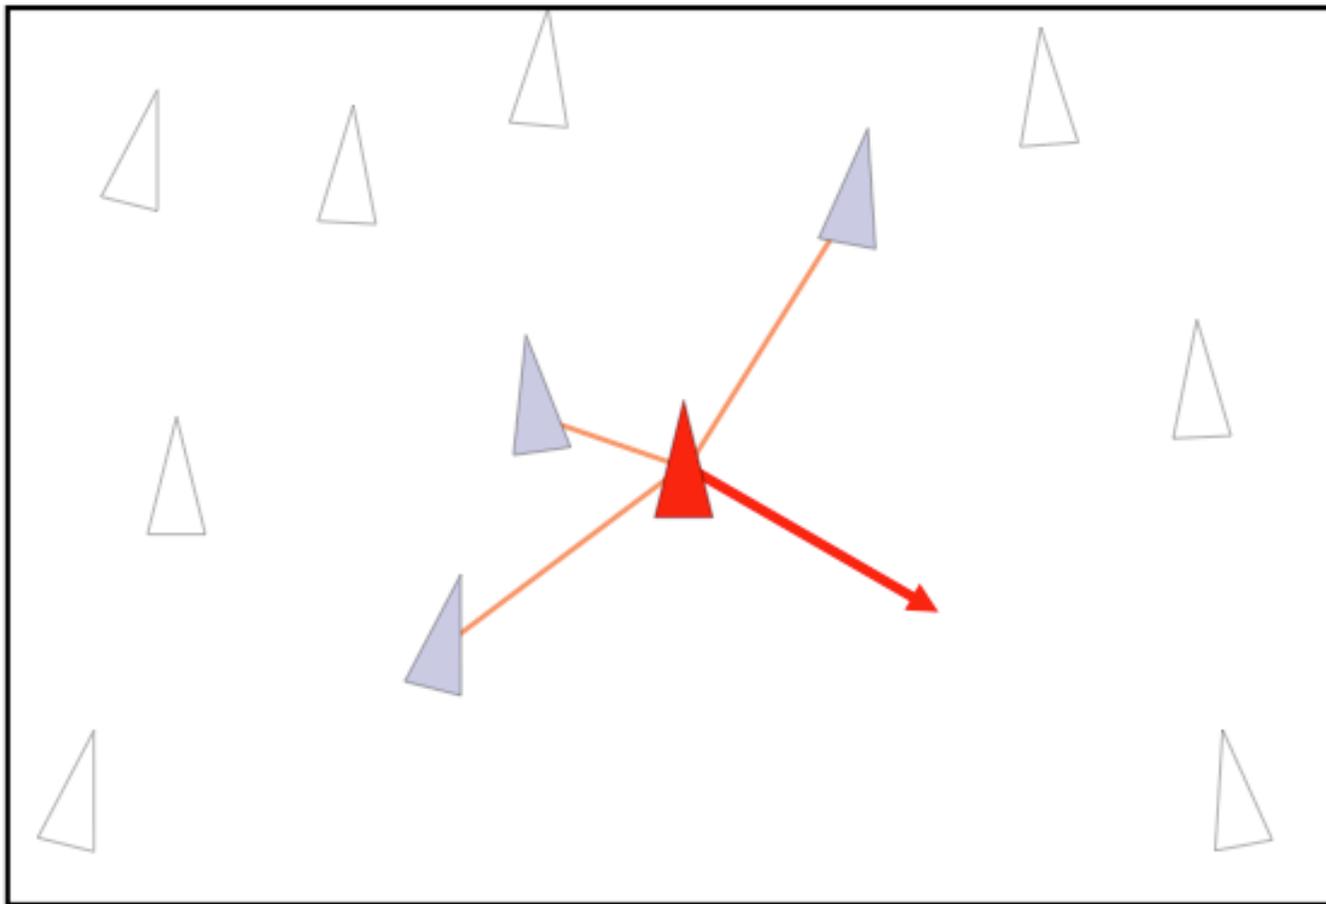
Boids (Reynolds, 87)

- Boids can only see individuals within a certain radius
 - Only local information
 - They have no idea where the group in general is heading



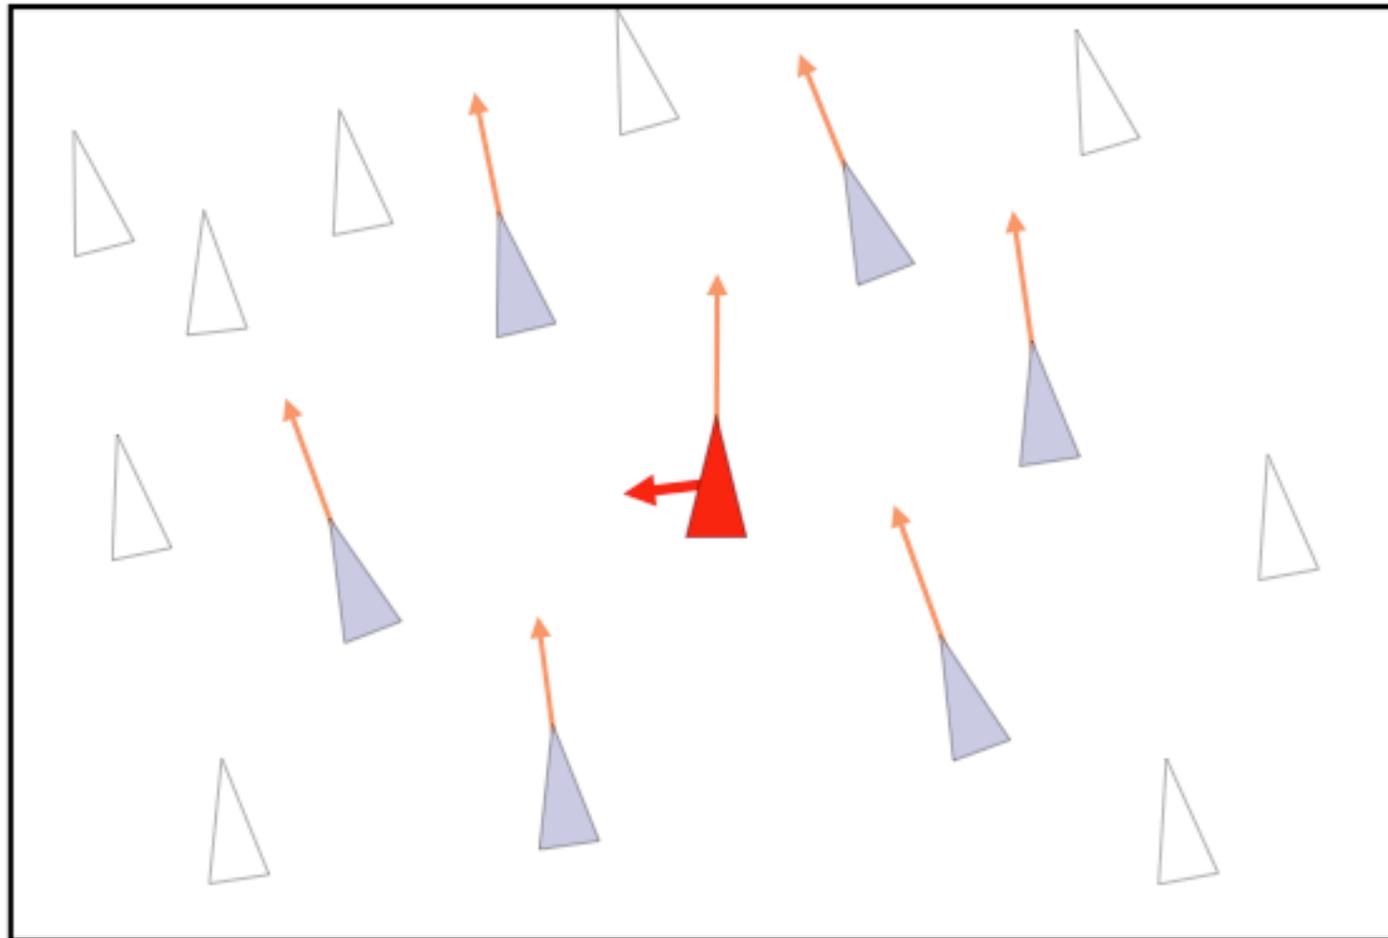
Rule 1: Collision avoidance

- Steer to avoid crowding others



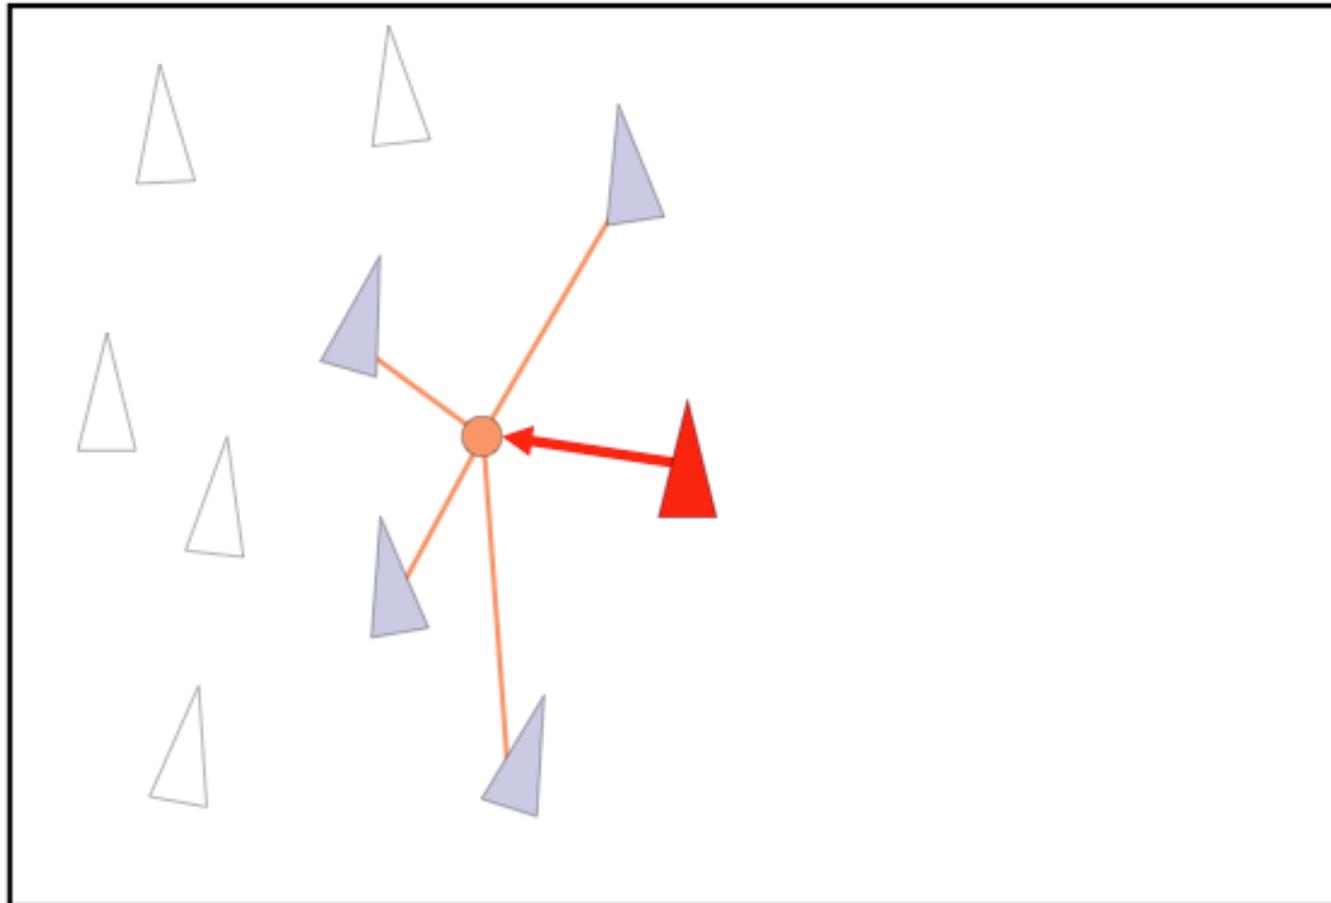
Rule 2: Alignment

- Turn to go towards average direction of flockmates



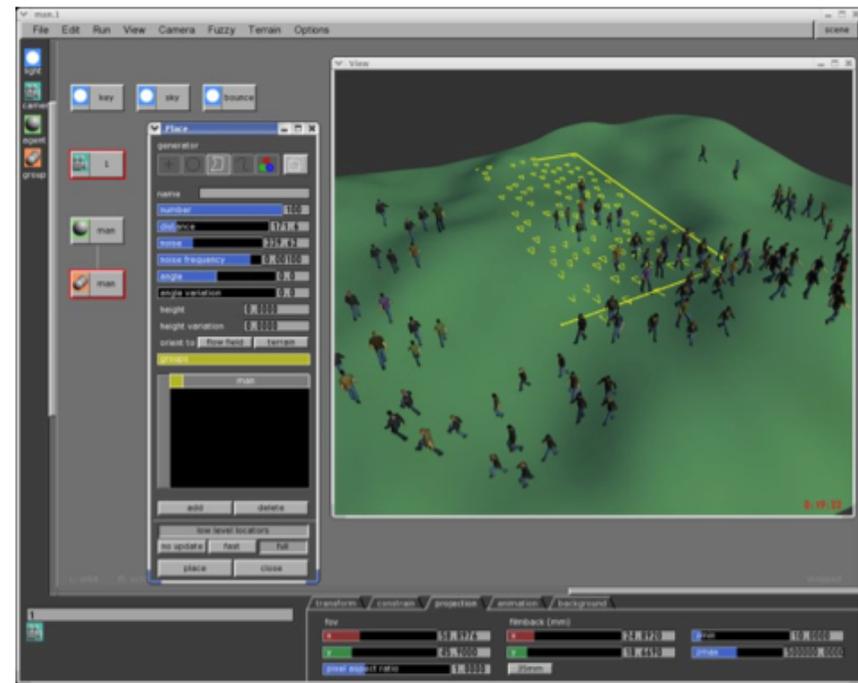
Rule 3: Cohesion

- Steer to move towards others if too far away



Applications

- Jurassic Park
- Bats in Batman Returns
- Lord of the rings



Massive, Massive Software

Swarm Intelligence

- Inspired by insects such as bees and ants
- Simple agents of limited individual intelligence exhibiting collective intelligence
 - The swarm is smart, but its members are not necessarily
- Remember ‘emergence’?
- The swarm is sometimes called a ‘superorganism’



Termites

- Termites
 - Build complex nest structures
 - How do they do it?
 - Does each termite have a complex plan?
- How they do it
 - 1) Move in the direction of the strongest pheromone
 - 2) Deposit what they carry where it smells the most
 - Who is in control?
 - Who plans? (are there any plans)?



Ant Foraging

- Ants
 - Collect food
 - Find a fast route to the food
- Queen does not control behaviour!
- How they do it
 - 1) Drop pheromones as they search for food
 - 2) Stochastically follow stronger pheromone concentrations
 - Again no single individual has a clue what is going on!

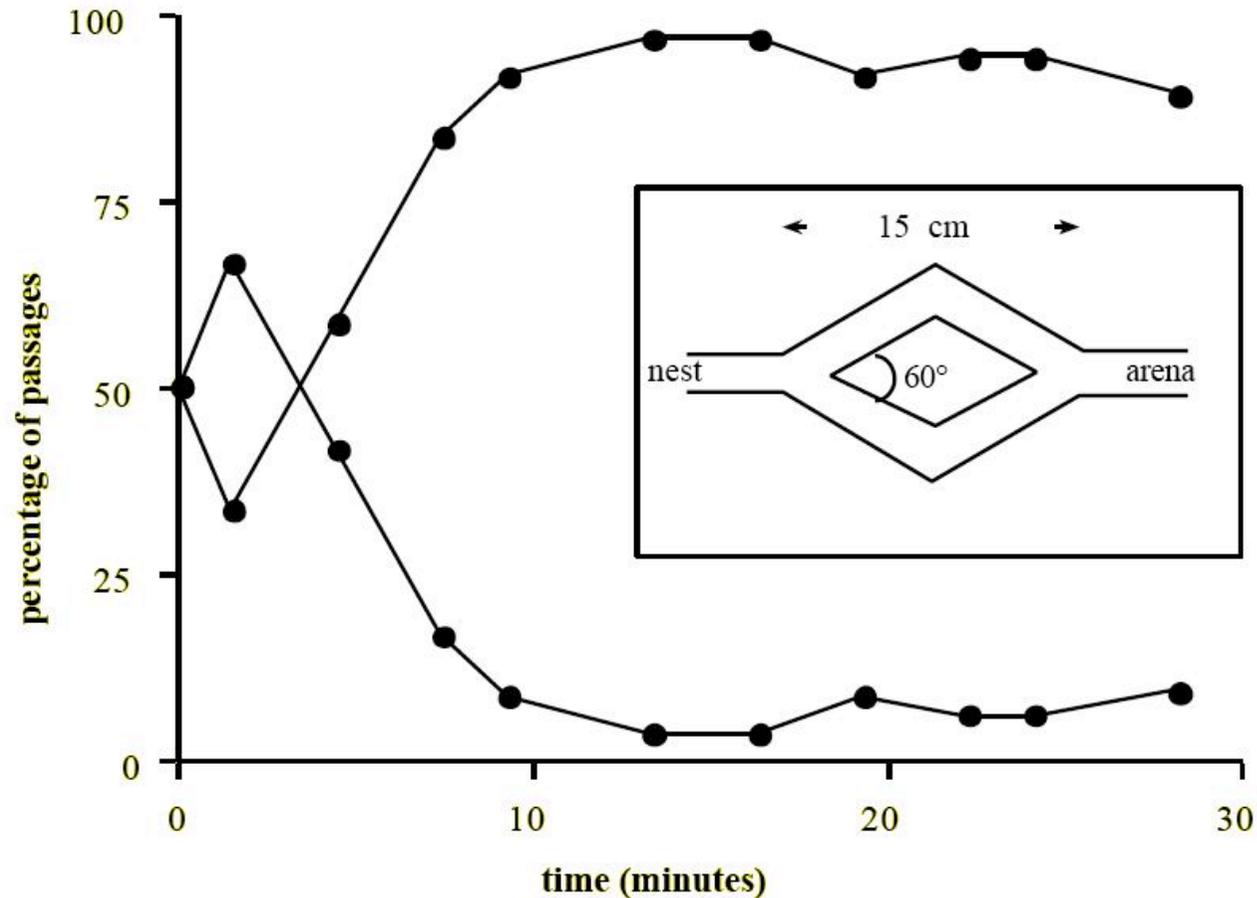


Communication

- Direct
 - Touch, food exchange, carrying, recruiting
- Stigmeric (stigmergy)
 - A type of indirect, local, communication
 - An individual modifies the environment, and the environment modifies the behaviour of individuals
 - For example: laying pheromones, building structures, removing obstacles
 - An example in ants is the dropping of pheromones

Ant foraging: The beginnings

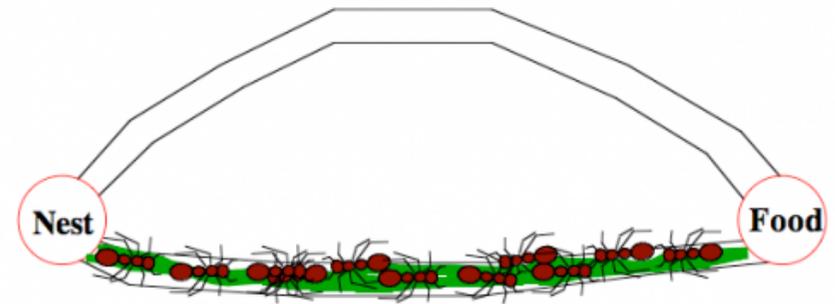
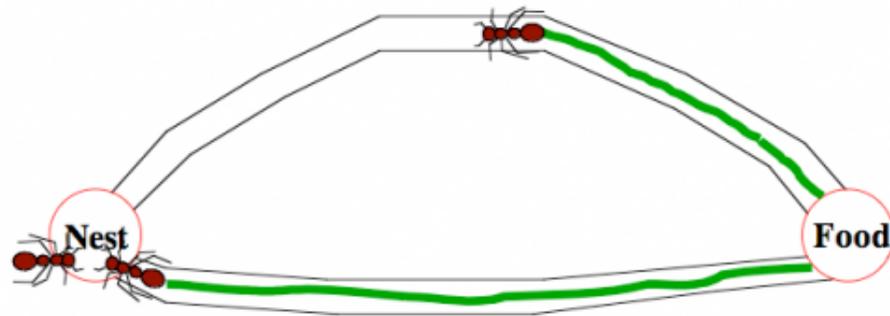
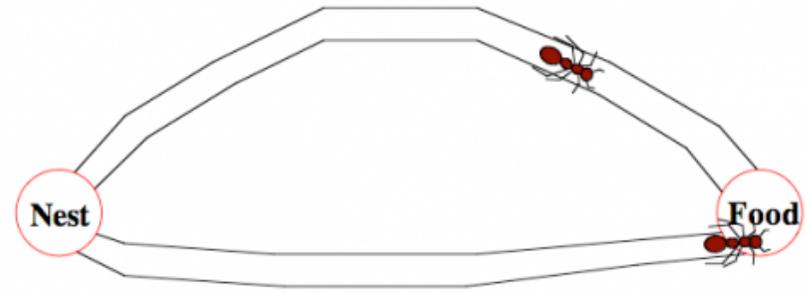
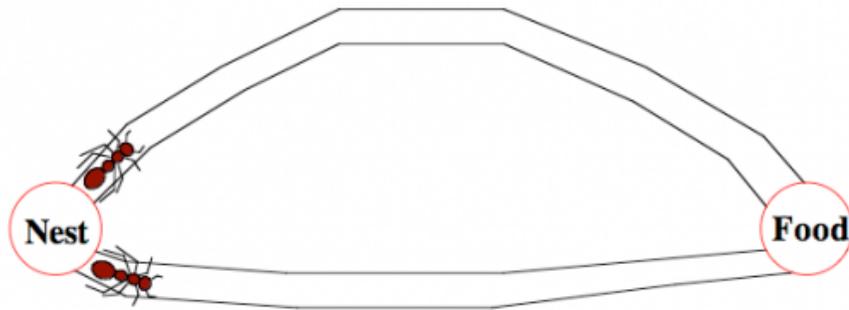
- How do ants find the nearest food source?
- Binary bridge experiment



Ant foraging

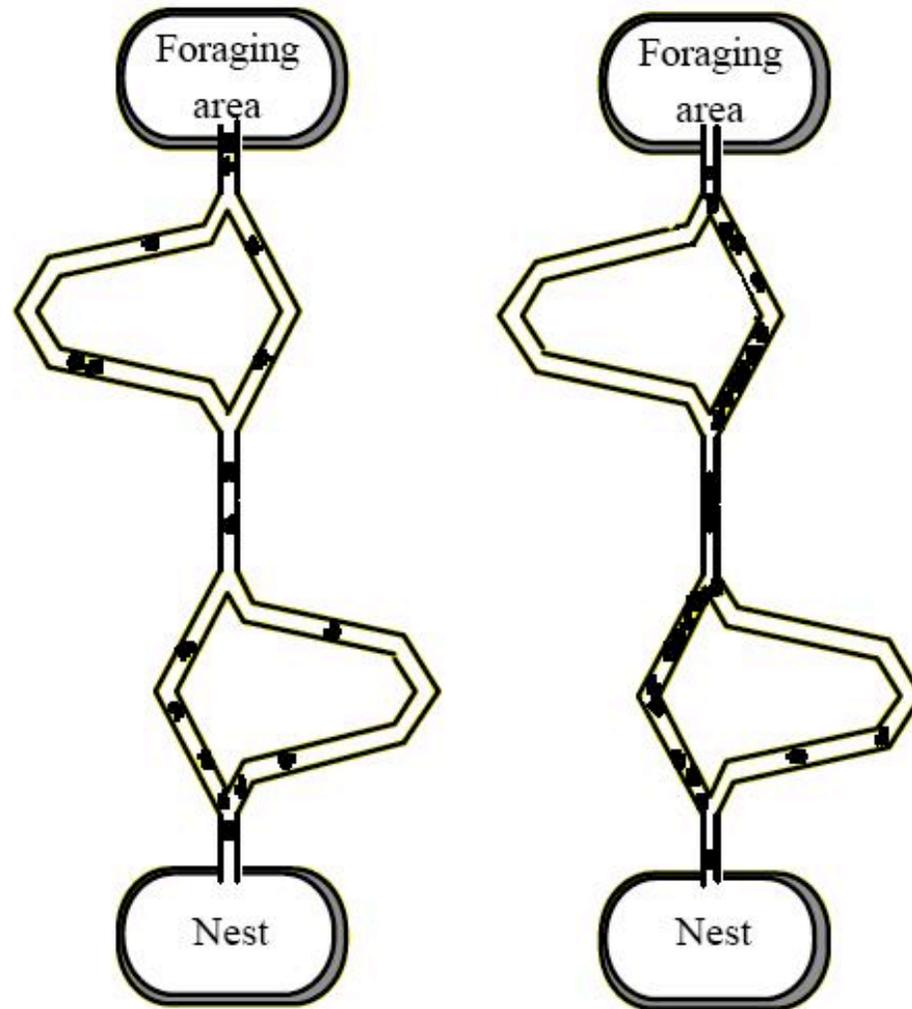
- Simple individual rules self organise at the level of the colony:
 - 1) Ants move out randomly from the nest
 - 2) When ants find food they return to the nest while adding pheromones to the trail
 - 3) When ants next move out to look for food they are more likely to follow stronger traces of pheromones
- Because individuals will return from closer nests at a higher rate the pheromone concentration on that route will be stronger
- This means that there is a higher probability that ants will take the route to the nearest nest
- The process is self-reinforcing

Ant Foraging: How it works



See Ant netlogo

Ant Foraging: More complex routes

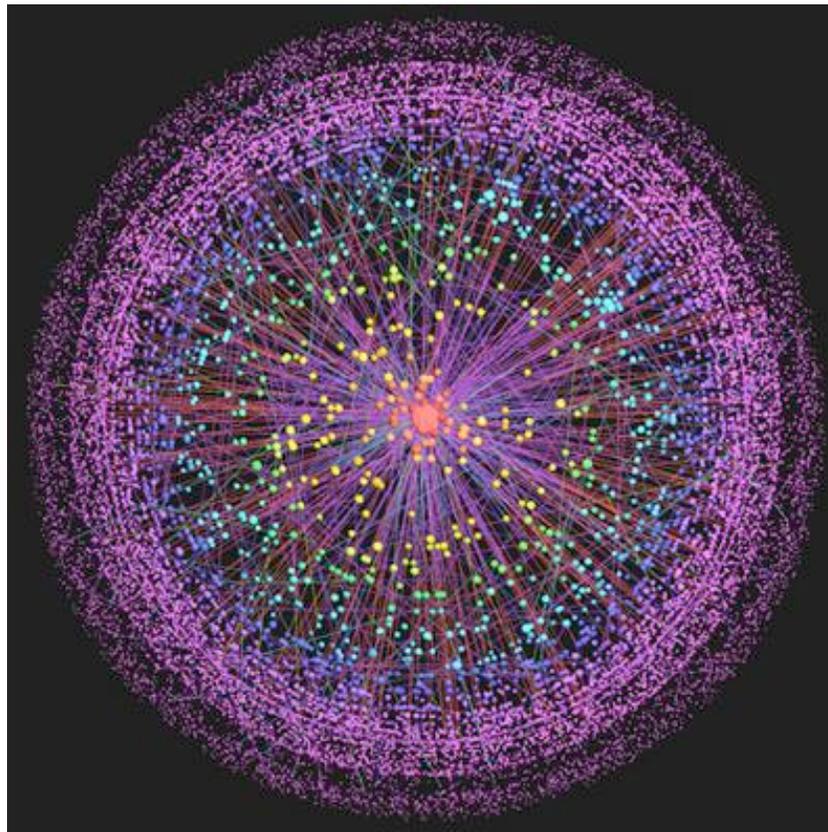


Ant Colony Optimisation (ACO)

- ACO is a good example of computer science taking inspiration from nature
- Ants solve the problem of finding a good route from one point to another
 - There are many computational tasks for which this could be a useful general-purpose algorithm

Ant Colony Optimisation (ACO)

- Finding the shortest path is a common optimisation problem
 - Internet routing



Abstract ACO Algorithm

input : The route to be optimised

output: Shortest route located

begin

 Initialisation (For every edge create random trail intensity place m ants on the n nodes, place starting town of the k -th ant on their search list)

repeat

 Choose the next node j with probability p

 Comment: p , for each node, is proportional to the number of ants that have visited that

node

 Move k -th ant to that node

 Insert node in list for that ant

until each ant's list is full

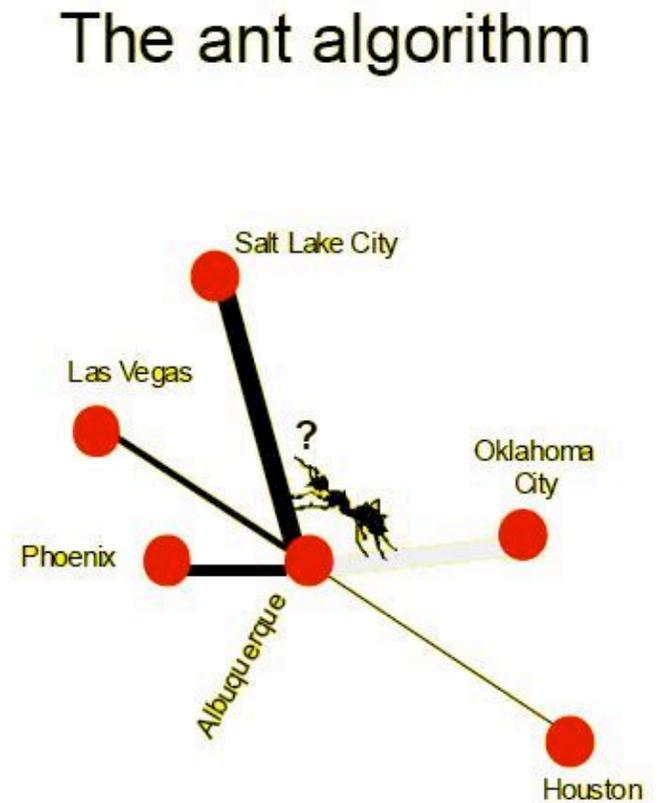
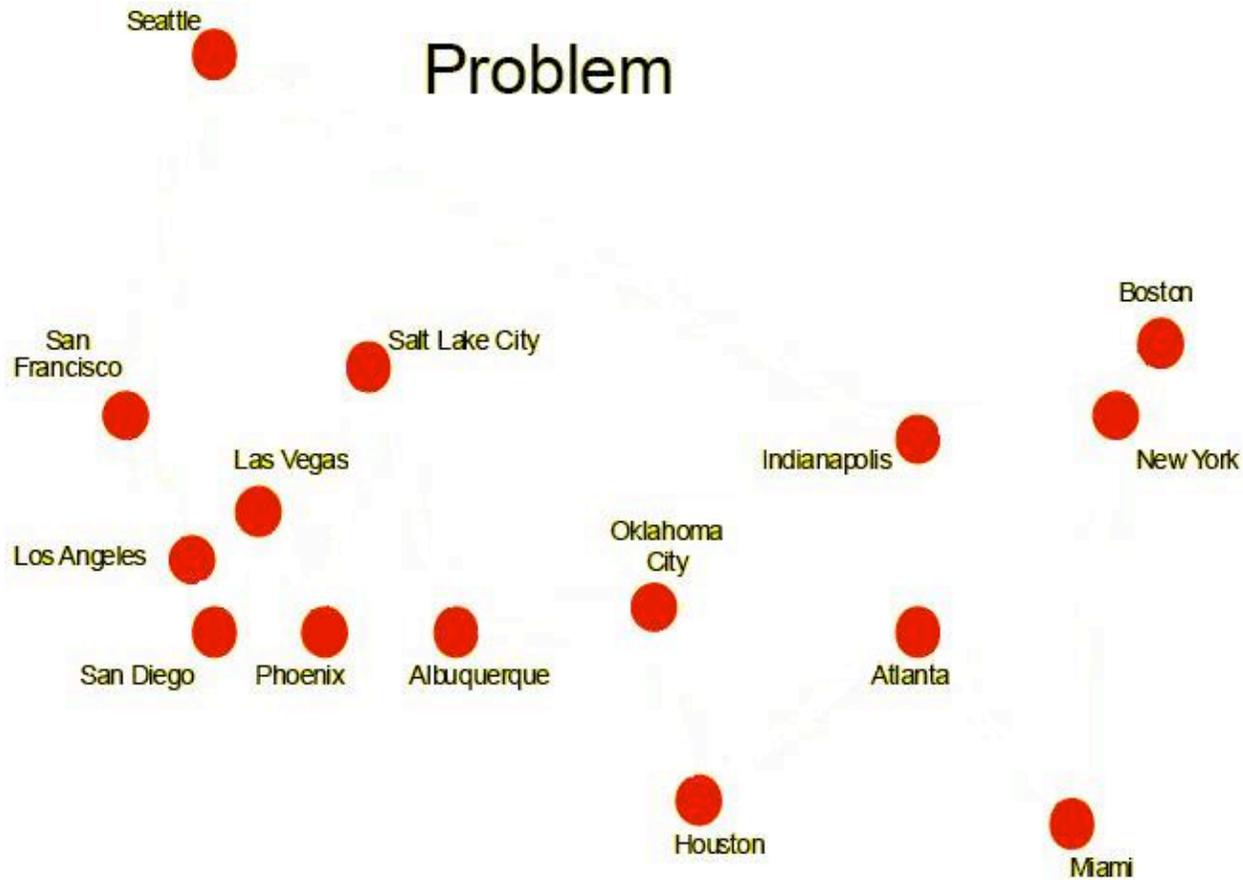
 Select the shortest route found during the algorithm

end

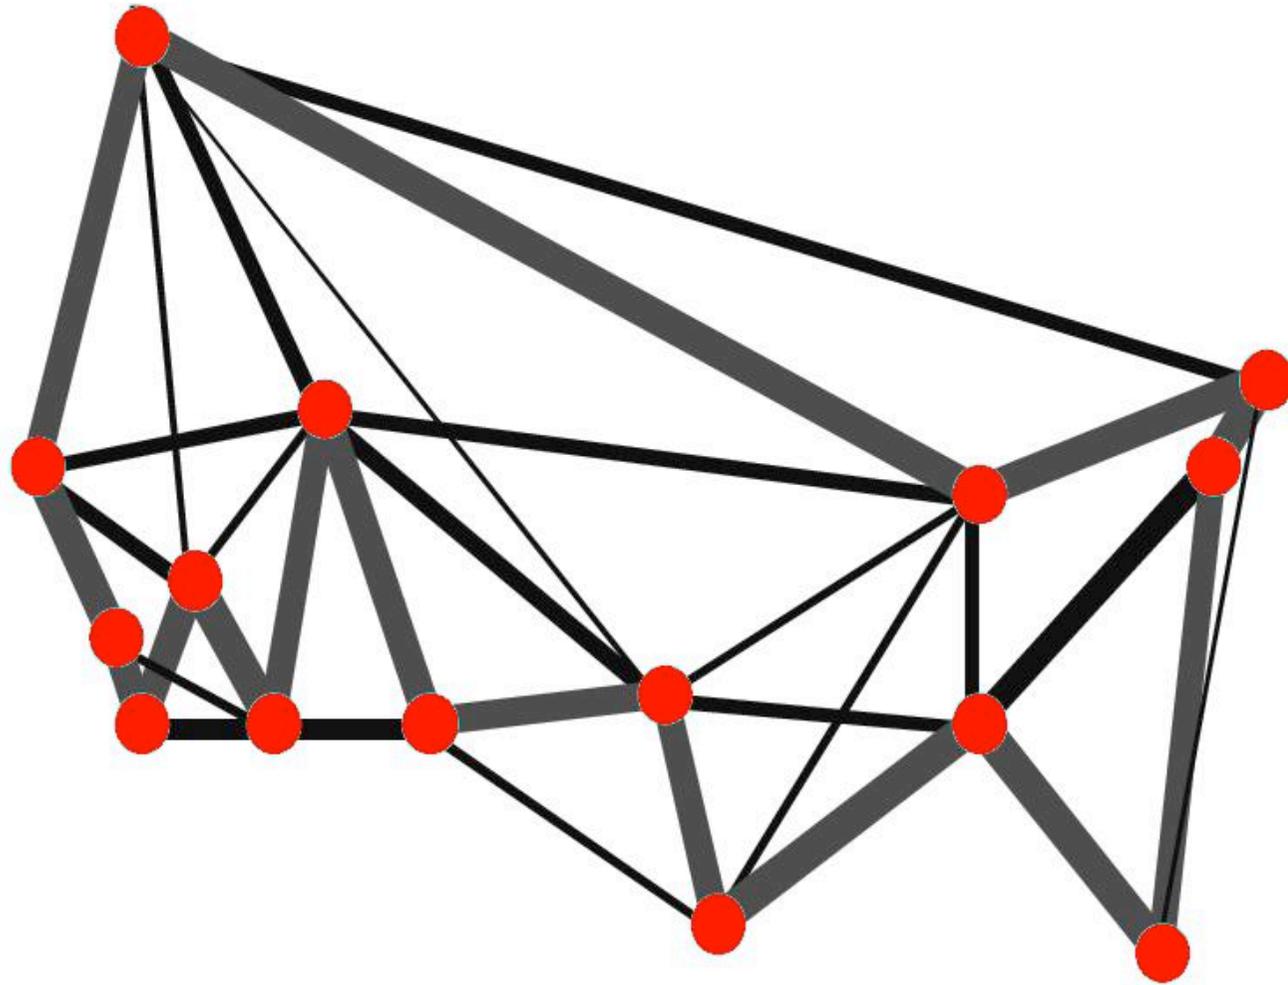
Ant Colony Optimisation (ACO)

- Key advantages:
 - Decentralised
 - No leading individual
 - The task is divided between many interacting individuals
 - Robust
 - Centralised systems are vulnerable
 - What if the central system is lost (breakdown or attack)
 - The system is robust to random and targeted loss
 - Scalable
 - Individuals can be added and removed without changes to the algorithm
 - Adaptive
 - If the problem changes the algorithm can adapt without needing knowledge of the change

The Traveling Salesman Problem



The Traveling Salesman Problem



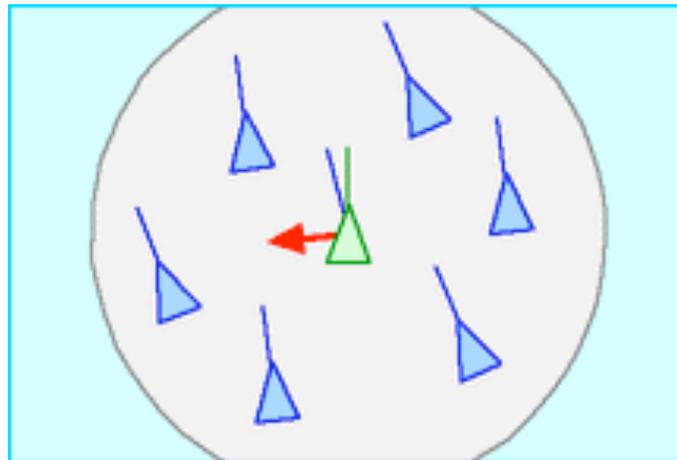
Inspiration

- Early work on simulation of bird flocking aimed at understanding the underlying rules
- Individuals modelled in a social context
- Kennedy & Eberhart stumbled upon the possibility that it can be used as an optimisation algorithm
 - Solutions represented as individuals in flocks of solutions
 - Individuals ‘fly’ through solution space

PSO Background

- Suppose the following scenario:
 - a group of birds are randomly searching food in an area.
 - There is only one piece of food in the area being searched.
 - All the birds do not know where the food is. But they know how far the food is in each iteration.
- So what's the best strategy to find the food?
 - They found that an effective one is to generally fly around the bird which is nearest to the food.

Return of the Boids



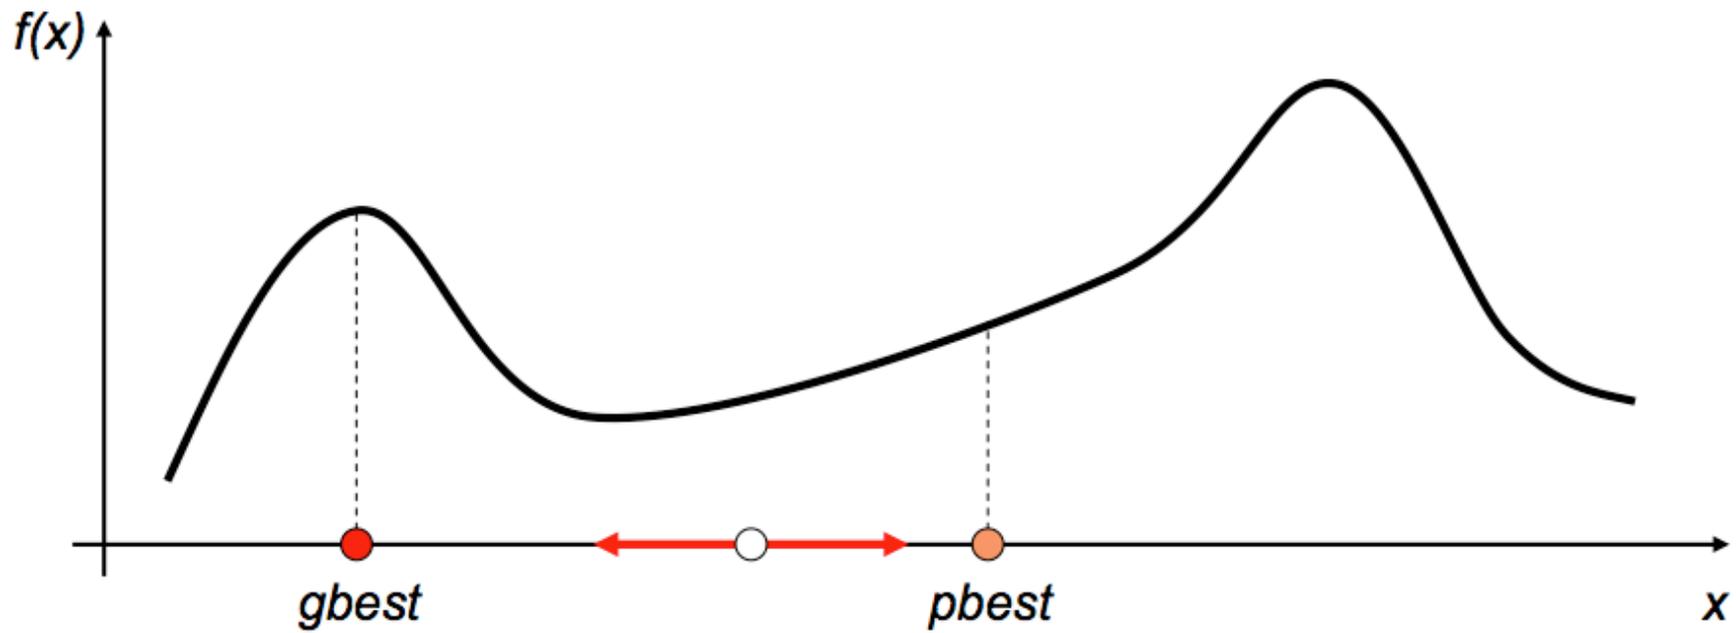
Sketch of PSO

- PSO was developed from the scenario and used to solve the optimization problems
- In PSO, each single solution is a "bird" in the search space. Named a "particle"
- All particles have:
 - Fitness values which are evaluated by the fitness function to be optimized
 - Velocities which direct the flying of the particles
- The particles fly through the problem space by generally following the current optimum particles

Sketch of PSO

- PSO is initialized with a group of random particles and then searches for optima by updating generations
- In every iteration, each particle is updated by following two "best" values
 - The best solution *it* has achieved so far (pbest)
 - The best value obtained so far by *any* particle in the population. (gbest)
 - So there is a level of global information in this system

The PSO



Velocity update

$$v_i(t) = \omega v_i(t-1) + c_1 \phi_1(p_i - x_i(t-1)) + c_2 \phi_2(p_g - x_i(t-1))$$


where:

ω : inertial constant (keep things moving)

$C_{1,2}$: constants that affects how much each best affects overall particle

ϕ_1 random scalar (learning rate)

Check back with the flocks idea - how similar is it?

Velocity update (simplified version)

$$(a) \quad v_x = v_x + 2 \cdot \text{rand}() \cdot (pbest_x - present_x) + 2 \cdot \text{rand}() \cdot (gbest_x - present_x)$$
$$0 \leq \text{rand}() \leq 1$$

$$(b) \quad present_x = present_x + v_x$$

- Individuals often overshoot the target
- Balance of exploration and exploitation
- v_x is limited by v_{max}
- (2 is the typical learning factor and can be changed)

Paper

- R. Mendes, J. Kennedy, and J. Neves (2004) “The fully informed particle swarm: Simpler, maybe better” IEEE Transactions on Evolutionary Computation, 8(3):204, 210, 2004.

Pseudo Code

Initialize each particle

Do

For each particle

 Calculate fitness value

 If the fitness value is better than the best fitness (pBest) in history

 set current value as the new pBest

End

Choose the particle with the best fitness of all the particles as the gBest

For each particle

 Calculate particle velocity according equation (a)

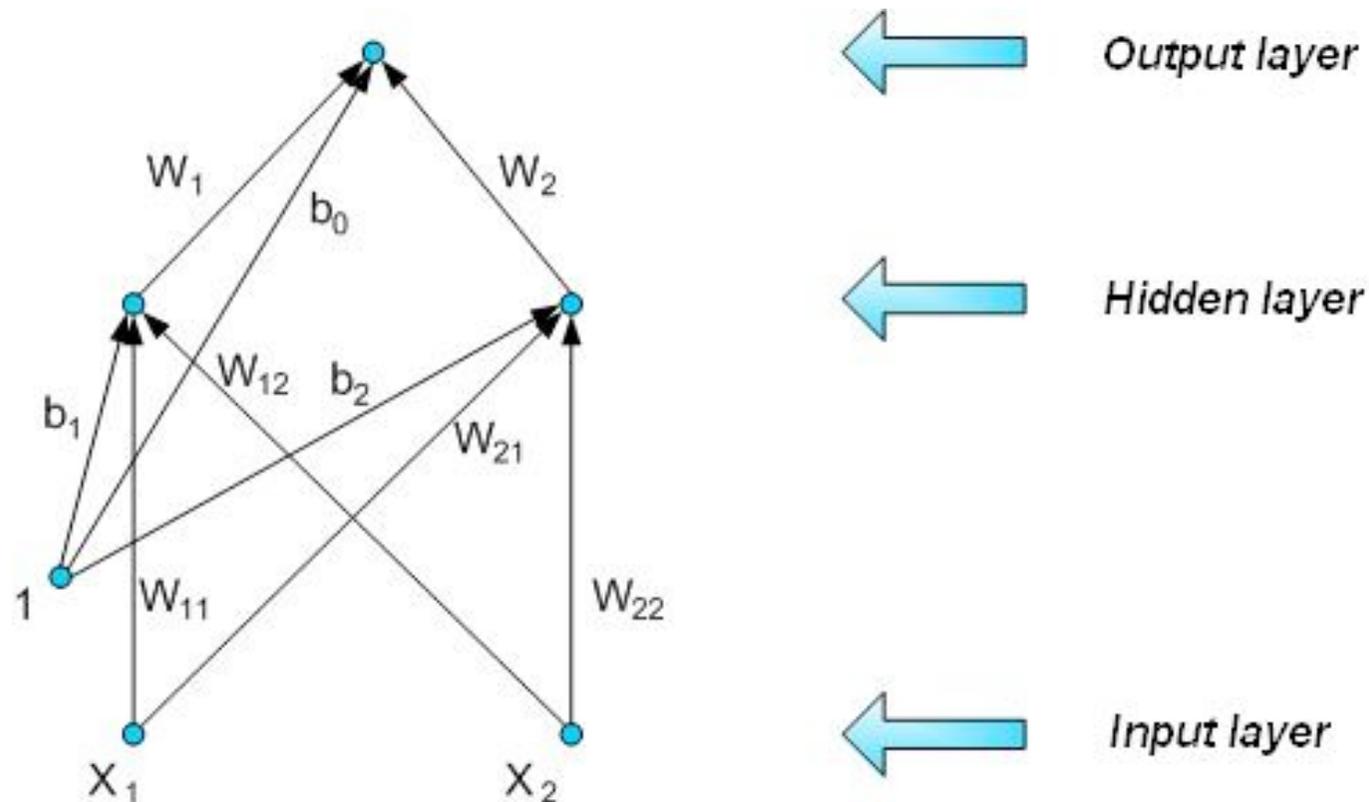
 Update particle position according equation (b)

End

While maximum iterations or minimum error criteria is not attained

Example application

- ANN weight optimiser
 - E.g. simple Xor problem (feed-forward ANN)



Summary

- PSO can be a good technique for solving some optimisation problems
- We can take inspiration from biology to produce generic problem-solving techniques
- Often a combination of bio-inspired and engineered is best

References

- Kennedy, J. and Eberhart, R. C. **Particle swarm optimization**. Proc. IEEE int'l conf. on neural networks Vol. IV, pp. 1942-1948. IEEE service center, Piscataway, NJ, 1995.
- Eberhart, R. C. and Kennedy, J. **A new optimizer using particle swarm theory**. Proceedings of the sixth international symposium on micro machine and human science pp. 39-43. IEEE service center, Piscataway, NJ, Nagoya, Japan, 1995.
- Eberhart, R. C. and Shi, Y. **Particle swarm optimization: developments, applications and resources**. Proc. congress on evolutionary computation 2001 IEEE service center, Piscataway, NJ., Seoul, Korea., 2001.
- Shi, Y. and Eberhart, R. C. **Parameter selection in particle swarm optimization**. Evolutionary Programming VII: Proc. EP 98 pp. 591-600. Springer-Verlag, New York, 1998.