





Competitive and Cooperative Co-Evolution




What you will learn in this class

- Models of competitive co-evolution
- Moving fitness landscape
- Methods for ensuring incremental progress in competitive co-evolution
- Altruism: the ultimate form of evolutionary cooperation
- Identifying the best algorithm for ensuring altruistic artificial intelligence
- Applications to biology and to engineering





Copyright © Pearson Education, Inc., publishing as Benjamin Cummings.




Competitive Coevolution

Competitive Co-Evolution is a situation where two different species co-evolve against each other. Therefore, fitness of a species depends on fitness of opponent species:

- Prey-Predator
- Host-Parasite


Potential advantages of Competitive Co-evolution:

- It may increase adaptivity by producing an evolutionary *arms race* [Dawkins & Krebs, 1979]
- More complex solutions may *incrementally* emerge as each population tries to win over the opponent
- It may alleviate the problem of designing fitness functions for evolution of complex behaviour
- Continuously *changing fitness landscape* may help to prevent stagnation in local minima [Hillis, 1990]



Formal models

Formal models of competitive co-evolution are based on the Lotka-Volterra set of differential equations: they describe variation in population size (not in population performance)




$$dN_1/dt = N_1 (r_1 - b_1 N_2)$$

$$dN_2/dt = N_2 (-r_2 + b_2 N_1)$$


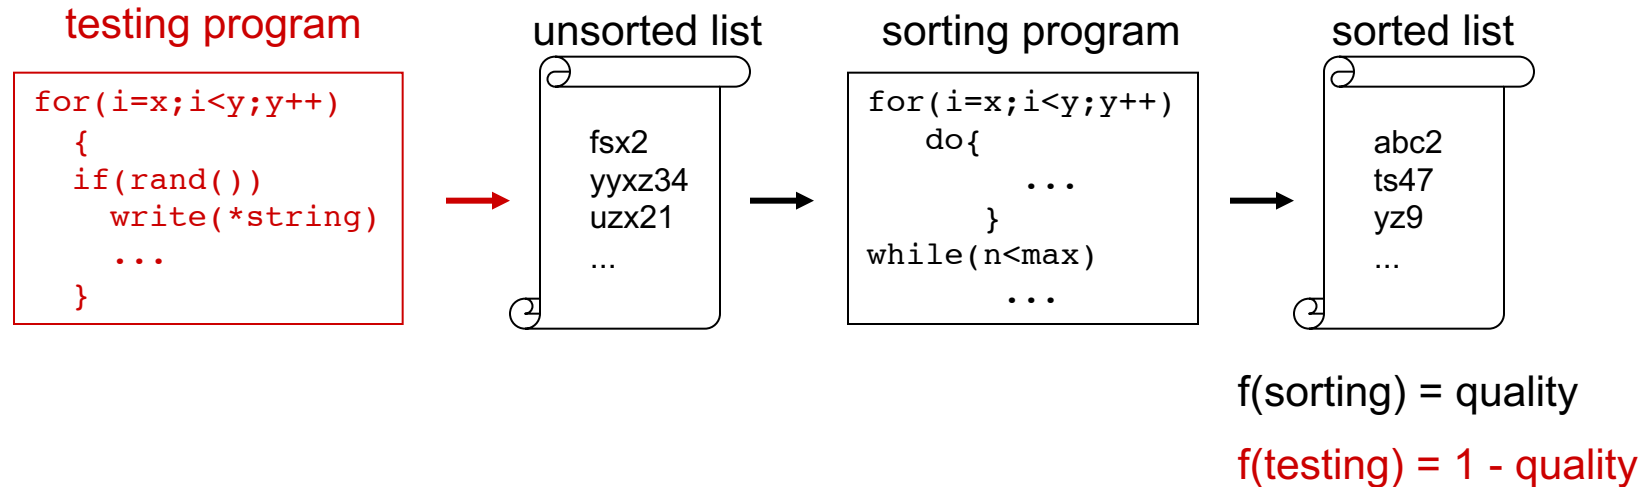
where:

- N_1 is the size of the prey population
- N_2 is the size of the predator population
- r_1 is increment rate of prey without predators
- r_2 is death rate of predators without prey
- b_1 is death rate of prey caused by predators
- b_2 is ability of predators to catch prey




Does it result in better performance over generations?

Hillis (1990) showed that co-evolution can produce more efficient sorting programs than evolution alone (or hand design).




Complication: Strategy recycling

The same set of solutions may be discovered over and over again across generations. After some initial progress, this cycling behavior may stagnate in relatively simple solutions.



Possible causes of strategy recycling:



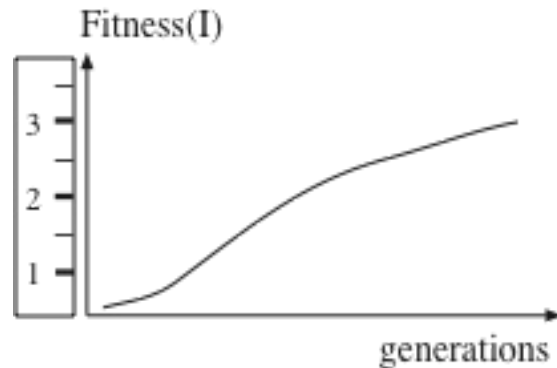
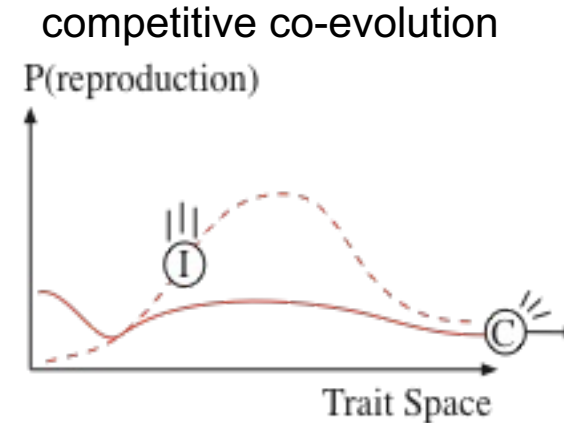
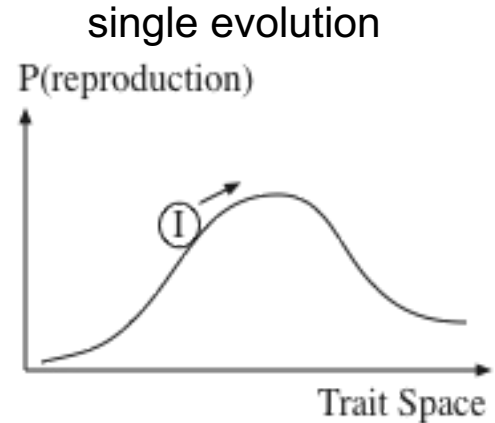
- Lack of « generational memory »
- Restricted possibility for variation
- Small genetic diversity



Complication: Dynamic fitness landscape

Whereas in single-species evolution the fitness landscape is static and fitness is a monotonic function of progress, in competitive co-evolution the fitness landscape is affected by the competitor.

Therefore, fitness is not an indicator of progress.



Investigation with robots


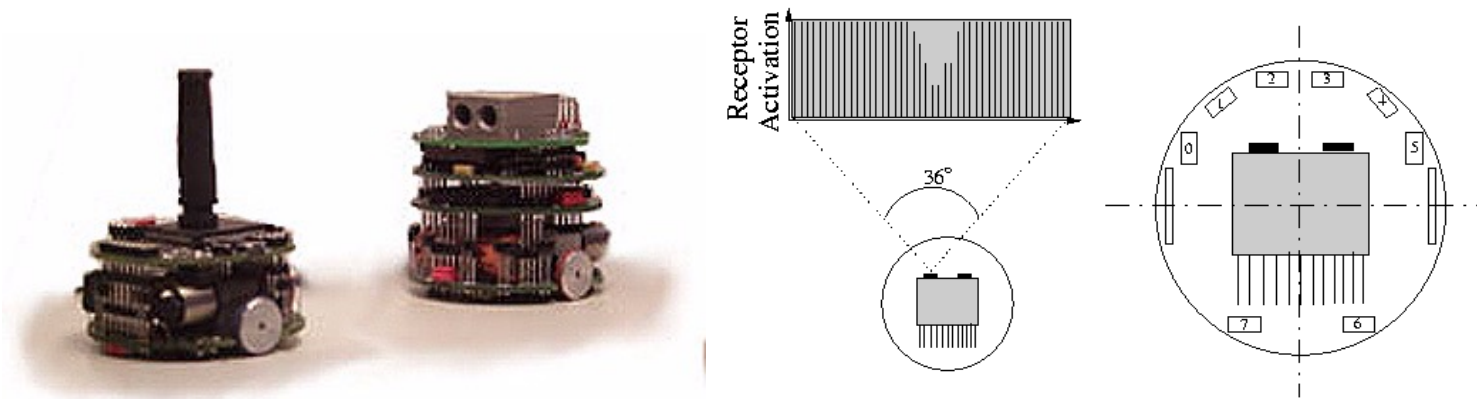
Let us consider the case of two co-evolutionary robots, a predator and a prey, that evolve in competition with each other. Questions:

- can we evolve functional controllers with simple fitness functions?
- what are the emerging dynamics?
- do we observe incremental progress?
- are co-evolved solutions better than evolved solutions?

Goal = Predator must catch the prey, prey must avoid predator

Prey = proximity sensors only, twice as fast as predator



Predator = proximity + vision, but half max speed of prey



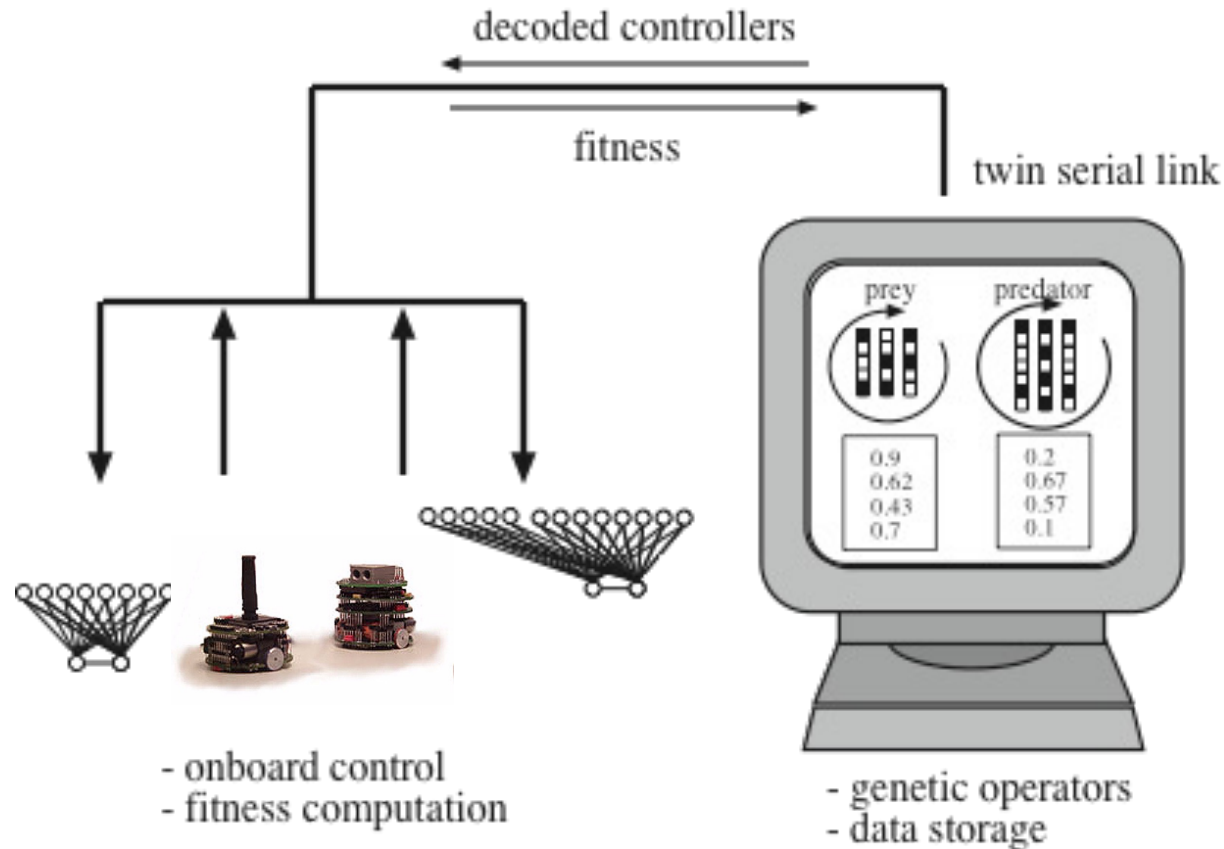
Experimental setup

The two robots are positioned in a white arena. Predator and prey are tested in tournaments lasting 2 minutes. Robots are equipped with contact sensors.


Fitness prey = TimeToContact Fitness predator = $1 - \text{TimeToContact}$



Co-evolutionary algorithm




Two populations, one for the prey and one for the predator, are maintained in the computer. Each individual of one population is tested against the best opponents of the previous 5 generations.





Experimental results

As expected, average and best fitness graph display oscillations.

with real robots



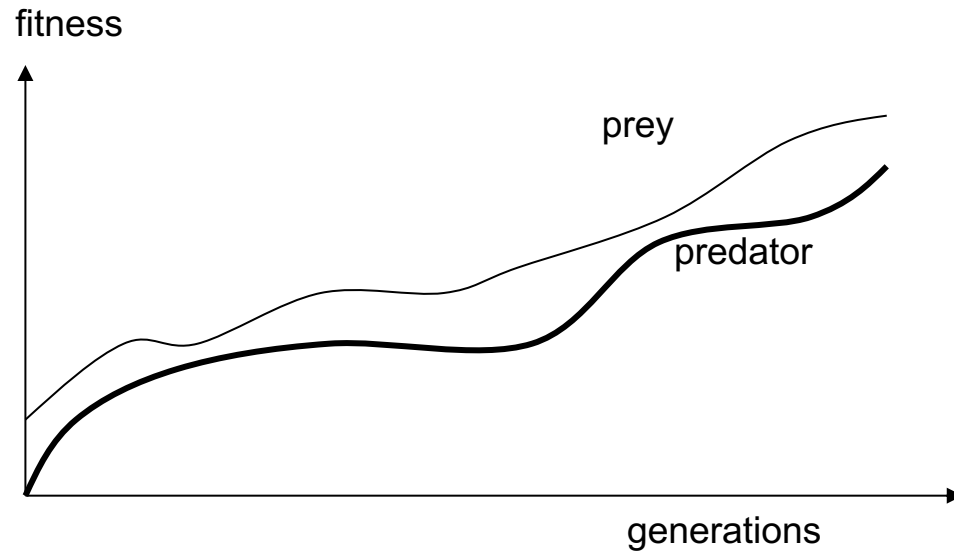
with simulated robots



How to measure progress?


Progress after evolution can be measured by testing evolved individuals against all best opponents of previous generations.

MASTER tournaments [Floreano & Nolfi, 1997]



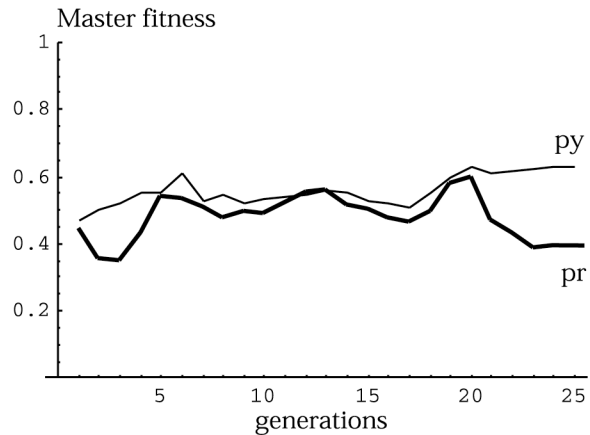
Average outcome (fitness) of tournaments of best individual at generation n against all best opponents of generations 0 to $n-1$.

Continuous growth of the master fitness indicates progress



Limited observed progress

with real robots




Master Tournament analysis shows some progress only during the initial 20 generations.

In other words, individuals born after 50 generations may be defeated by individuals that were born 30 generations earlier.

with simulated robots








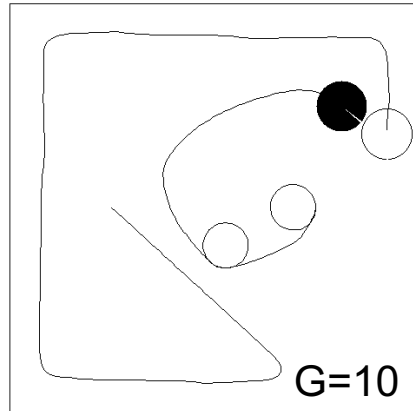



It suggests that co-evolution may have developed into re-cycling dynamics after 20 generations.




Evolved behaviors

Despite lack of progress measured against previous opponents, co-evolved individuals display highly-adapted strategies against their opponents, behavioral diversity, and fast behavioral switch




The Hall of Fame Selection Method

Master fitness of individuals evolved with Hall of Fame selection method




Individual fitness is measured against the best opponents of the previous generations (*Hall of Fame*) (Rosin and Belew, 1997). Problem: the number of tournaments increases with the number of generations.

It is sufficient to measure fitness of individuals only against a limited sample (10, e.g.) of best opponents randomly extracted from the Hall of Fame in order to produce continuous incremental progress, as shown by the Master Fitness graph measured after evolution.





Introducing learning in competitive coevolution



GENES ENCODE:

- weights
- hebb rules
- random values

- After 20 generations, predators always win
- Predators always evolve learning (Hebb rules)
- Prey most often evolve random change because its short-range sensors cannot benefit from learning




Evolution of cooperation

Simple cooperation easily evolves if there is an advantage but no cost in helping somebody else because the fitness of cooperator is increased

Photo: Pete Ellison





Altruistic cooperation is difficult to explain because it involves a cost for the individual. Example: Warrior ants that die to save the colony



Genetic relatedness

Cost


Genetic relatedness



Hamilton (1964)




$$\frac{C}{B} < r$$



Group selection

All individuals of a group that is better (at foraging, at land grabbing, at defense) than competing groups have the same chance to reproduce (independently of genetic relatedness (Wynne-Edwards, 1986; Michod, 1999)).




Criticism: The theory is not convincing because individual mutations during reproduction may destroy cooperation and make the group weaker




Four possible algorithms


HOMOGENEOUS





TEAM SELECTION





HETEROGENEOUS



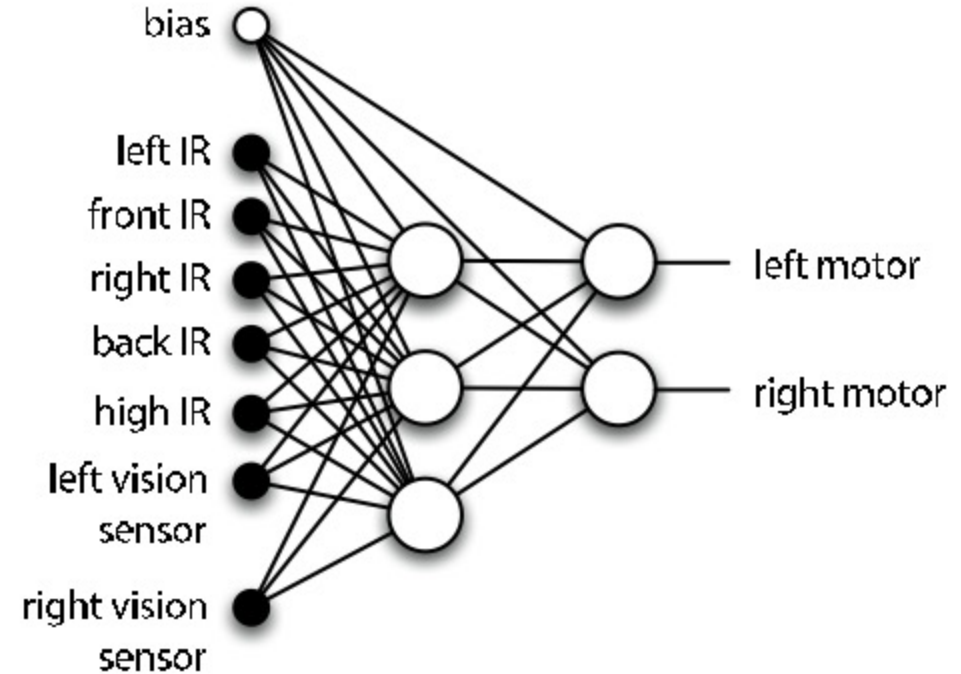


INDIVIDUAL SELECTION



Robot Foraging Task



Control structure




Connection weights of network are encoded in artificial genome

Each team is composed of 10 robots


The population is composed of 100 teams

Each team is evaluated 10 times




Types of tasks

INDIVIDUAL




1 fitness point per object to foraging robot

COOPERATIVE




1 fitness point to **all** robots for each object (2 robots necessary to push an object)

ALTRUISTIC




1 fitness point to **all** robots for each large object



1 fitness point per small object to individual robot






Genetically related, group selection



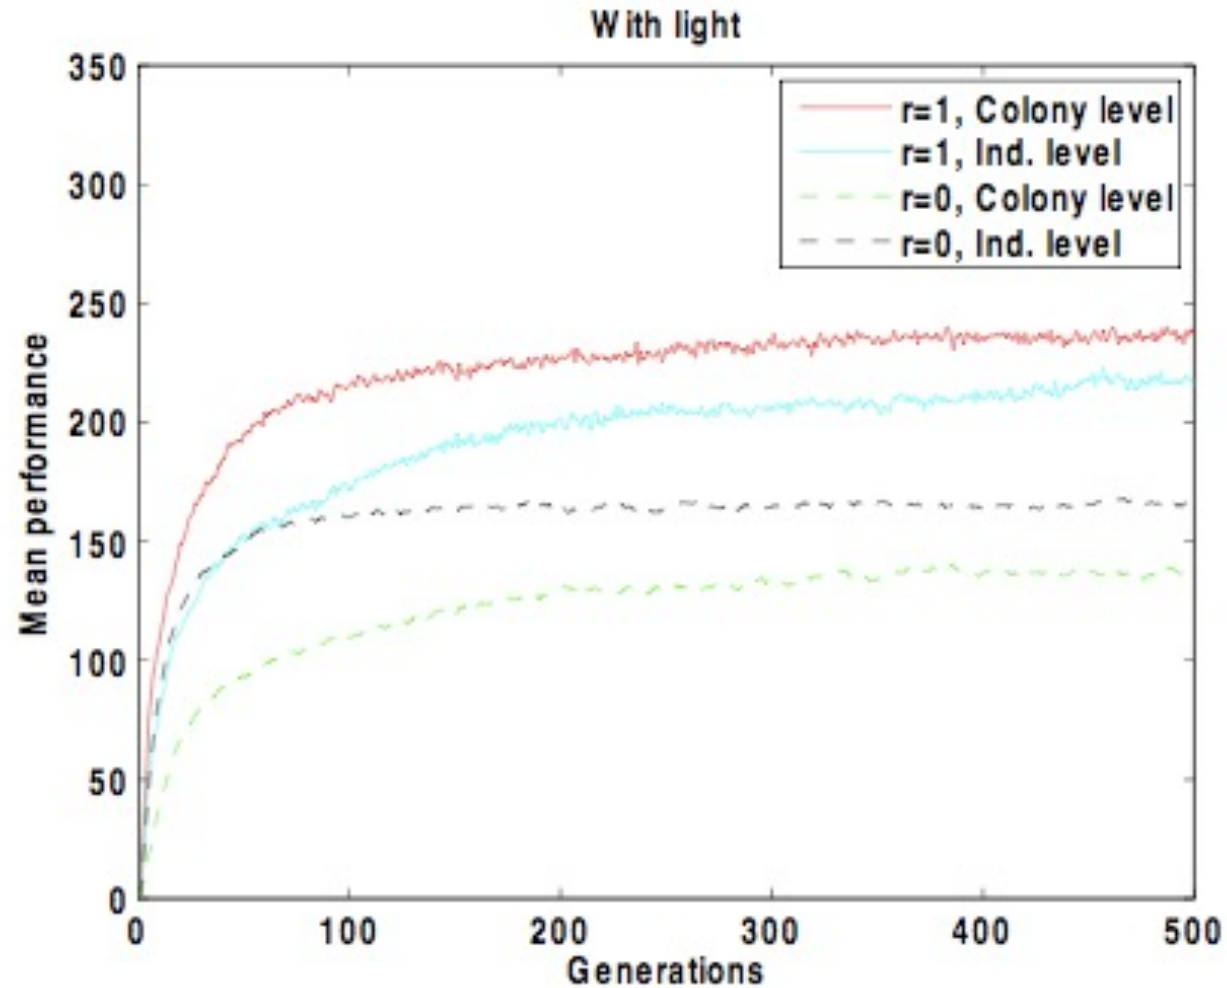
Foraging with Uncertainty




Control structure




Comparative Performance





Genetically related individuals obtain highest performance





FLYING RADIO NETWORKS



Aerial Platform

SMAV platform with control electronics (weight: 370g, speed: 10-15m/s, endurance: 30min)

μ C (dsPic33), 2 gyros + 2 pressure sensors ensure flight stabilization



Linux board calculates swarm control algorithms

Pitot tube


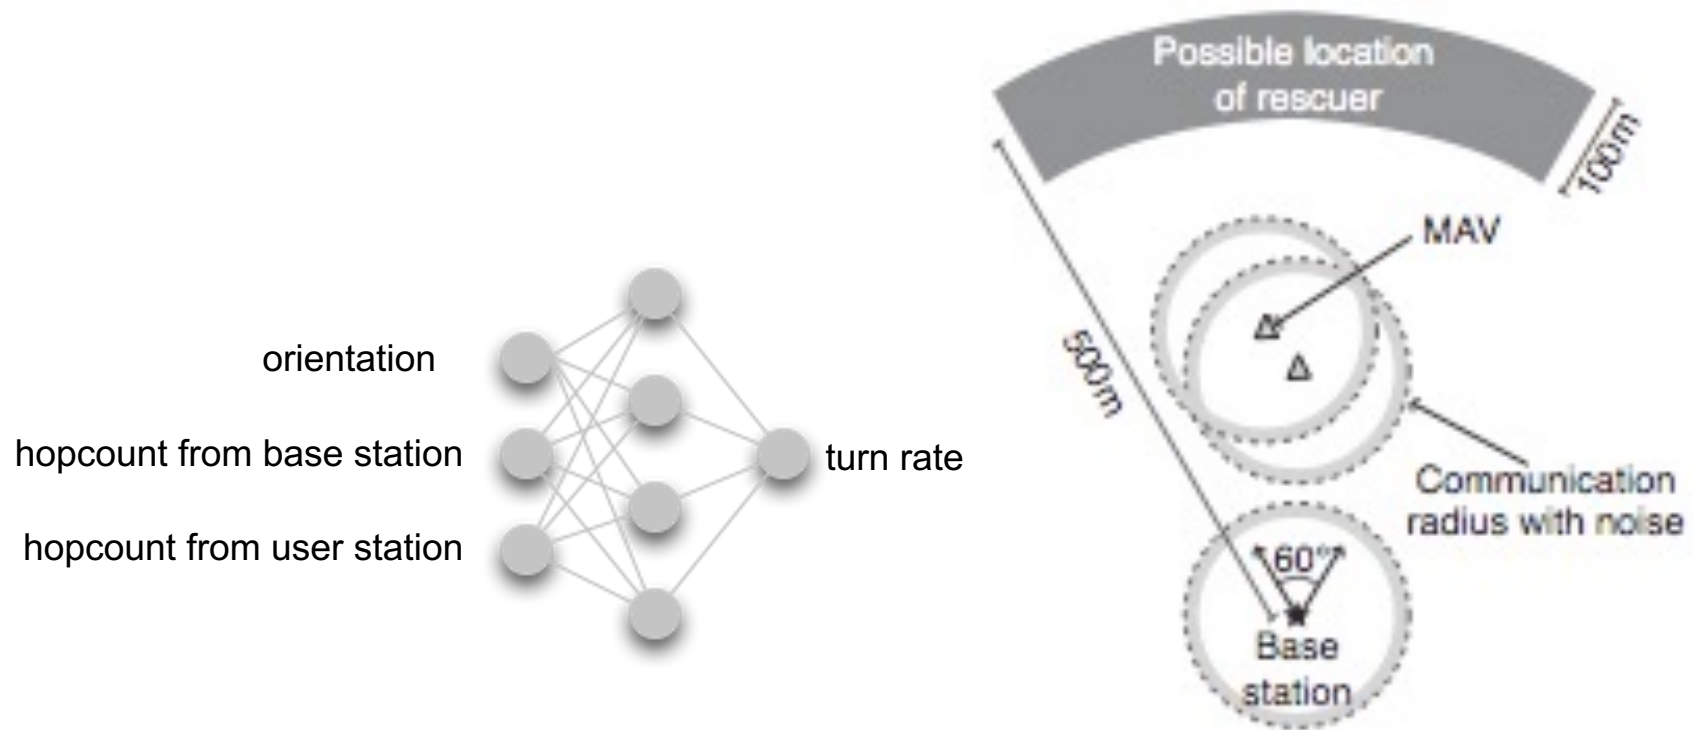
1.5g, 3-axis magnetic compass (inside wing)

standard WiFi-USB dongle

3-cell LiPo battery



Evolutionary Conditions






Evolved Swarm Control


SMAVNET project, EPFL

Sabine Hauert, Severin Leven, Dario Floreano, Jean-Christophe Zufferey

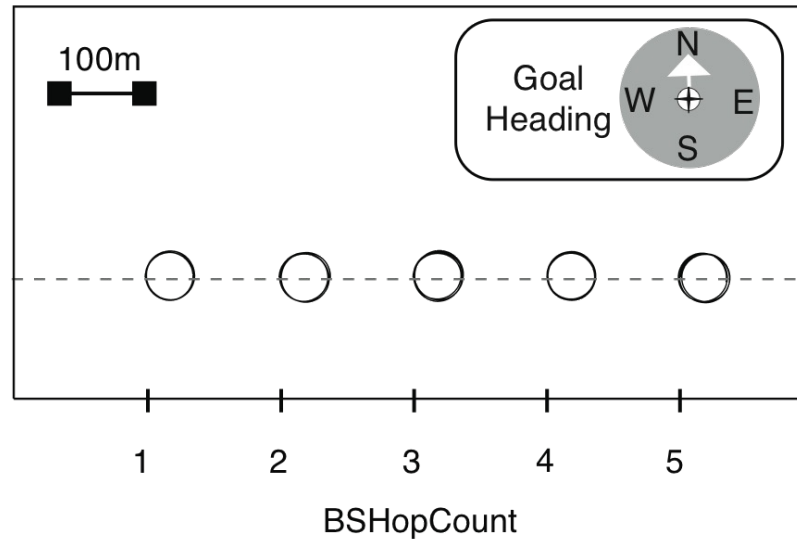


Evolved Control Strategy


Trajectory of MAV **disconnected** from user



Trajectory of MAV **connected** to user



Transfer to Reality



Summary

Competitive co-evolution can lead to increasingly better artificial intelligence (recently neural networks are trained by reinforcement learning to compete with each other to play games)

Moving fitness landscape can encourage new solutions, but also install cycling dynamics

Generational memory (Hall of Fame) is useful for preventing or retarding cycling dynamics

Altruistic cooperation evolves if individuals are genetically related or at least there is group-level selection

