Social Learning: Potential Mechanisms Compared in a Simple Environment

Jason Noble

jn2@ecs.soton.ac.uk

Science and Engineering of Natural Systems Group

School of Electronics and Computer Science

University of Southampton

Outline

- Social learning: why and how.
- A simulation model with learning agents in a complex environment.
- Results suggest that imitation is not always the best mechanism for social learning.

Talk draws heavily on Noble and Franks (2002, 2004). Apologies for it not being fresher!

Motivation

- In AI & psychology, learning often treated as a solo pursuit.
- Multi-agent systems often just a collection of individual learners.
- But social agents are not Robinson Crusoe — they can learn from others.
- This logic applies as much to animals and to artificial agents.

Social learning in nature

- Many examples: e.g., rats, guppies, lions.
- We suspect they're doing it, but we often don't know how.
- Imitation is a popular theory.

Imitation

- Imitation is an appealing idea, but the mechanism must be complex:
  - One animal perceives the bodily movements of another.
  - Identifies changes achieved in environment.
  - Understands why behaviour brings about goal.
  - Successfully reproduces the movements in the right context.

Mirror neurons?

- Enthusiasm for mirror neurons because they appear to account for imitation.
- Exciting findings, but a correlation is not a mechanism.
- Several different mechanisms would give us neurons that fire both when seeing-X and doing-X.

Simpler mechanisms

- Noble and Todd (2002) argue that mechanisms simpler than imitation have been neglected.
- Examples: stimulus enhancement, contagious behaviour, emulation.
- Simulation aims to show that the success of a social learning mechanism is not a straightforward function of its complexity.
- Simpler mechanisms should be of inherent interest to multi-agent system designers.

The model

- 10×10 toroidal grid-world.
- Grid squares can contain agents, resources, and tools.
- Agents perceive local state and must select an action.
- Different payoffs awarded depending on appropriateness of action in context.

The model

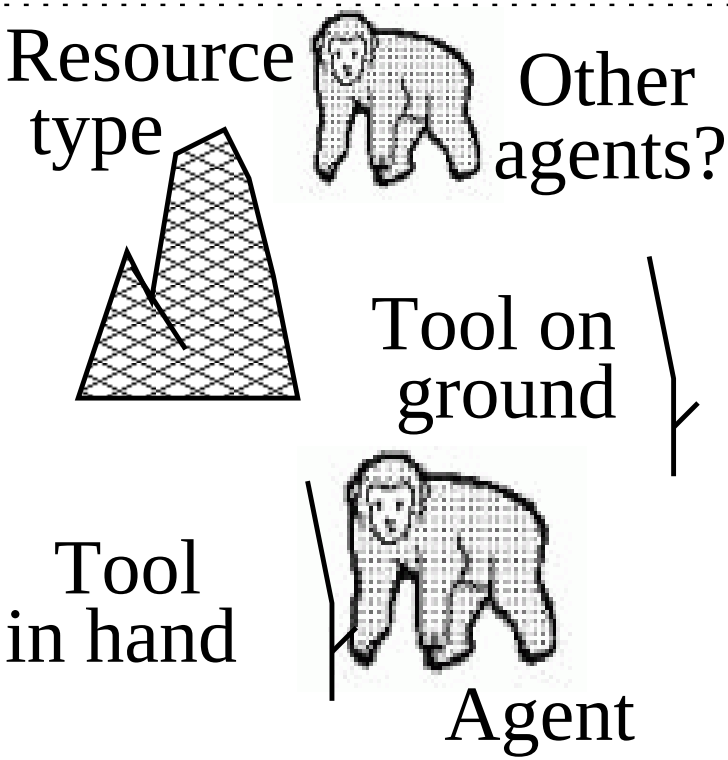

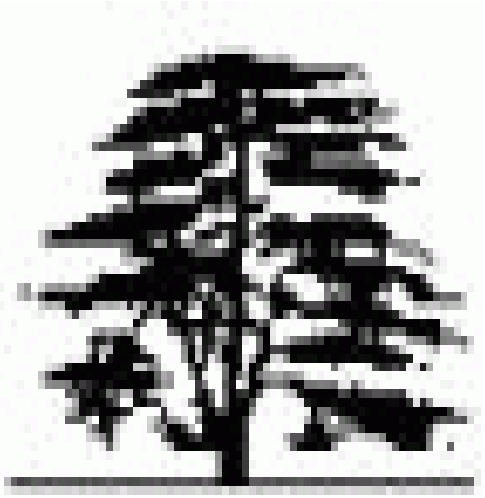

Next resource

The environment

- Designed to provide a range of learning challenges for the agents.
- Inspired by ape ecology, but *not* a serious model of any one species' environment.
- Possible grid squares: empty, tree, berries, nettles, termite nest, coconuts, monkeys, thorns, beehive, wasp nest.
- Possible tools: sticks, peeled sticks, small stones, large stones.

The agents

- Population of 25 held constant, oldest agent removed to make way for newest.
- Decisions based on history of reinforcement with 20% random actions to encourage exploration.
- Possible actions: do nothing, shake, fold, peel, poke, bash, chase, throw, rub, scratch, swap, travel.
- Agents powered by Q-learning with some generalization.

A range of behaviours

- Simple, easily learned behaviours: lying under the shade of a tree, eating berries.
- Medium-complexity behaviours: folding nettle leaves in order to safely eat them, or bashing a termite next to get inside.
- Difficult behaviours: peeling the bark from a stick to make it into a better poking tool for termite nests, lighting a fire in order to drive bees from their hive and safely retrieve the honey.

Implementing social learning

- Baselines: random behaviour & individual learning.
- Four strategies implemented: following, contagious behaviour, emulation, and imitation.
- Following also combined with other behaviours.

Implementing social learning

- Following: stick with your “parent” for the first 25% of your lifetime.
- Contagious behaviour: 10% chance of doing whatever your neighbour just did.
- Emulation: if you see another agent get a payoff, learn that this is a positive (or negative) situation.
- Imitation: if another agents gets a payoff, take its perspective and learn from that.

Data collection

- Collected data on the stable, long-run population.
- Measured only what ethologists would be able to record in a real population.
- Considered average payoffs as an overall measure of success.
- Looked at behavioural convergence as this is usually taken as a sign of cultural learning in nature.

Mean payoff per timestep

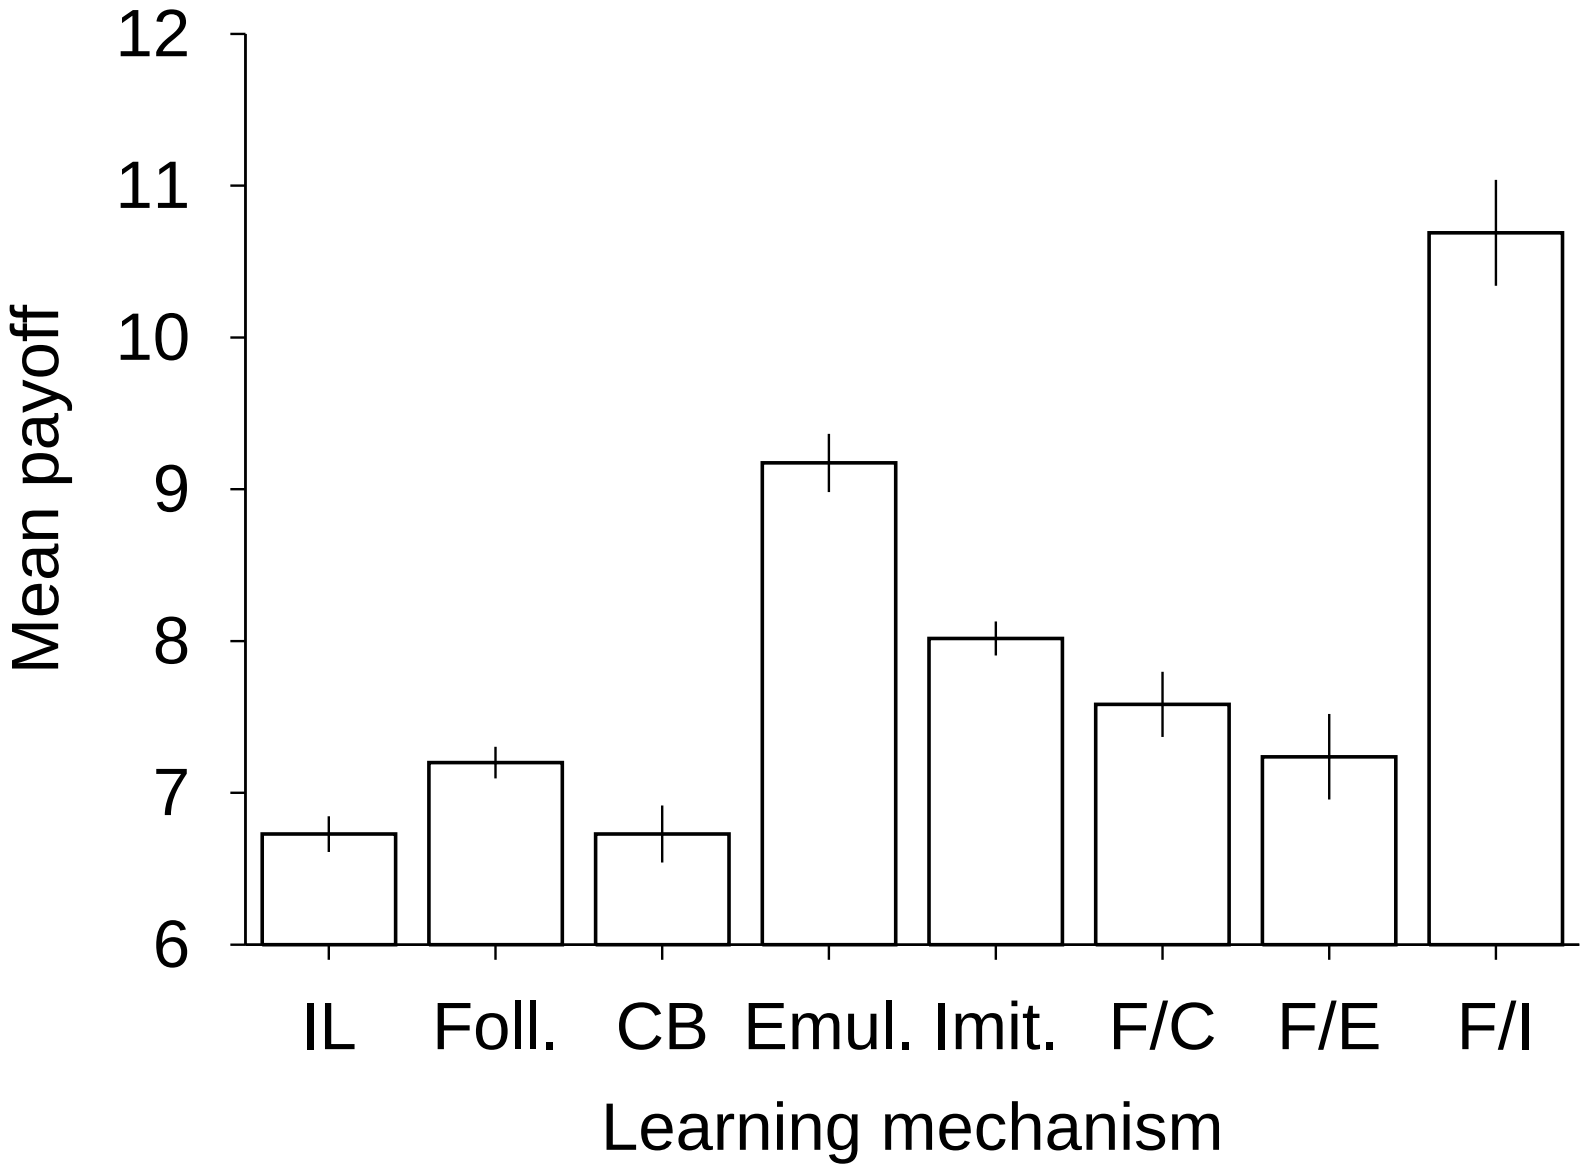

Mean payoff per timestep

- Best mean payoff for a single strategy is for emulation, not imitation.
- Following is useful, contagious behaviour is not.
- Following combined with other behaviours leads to complex results. Imitation best directed at a parent, emulation best directed at random individuals.

Behavioural convergence

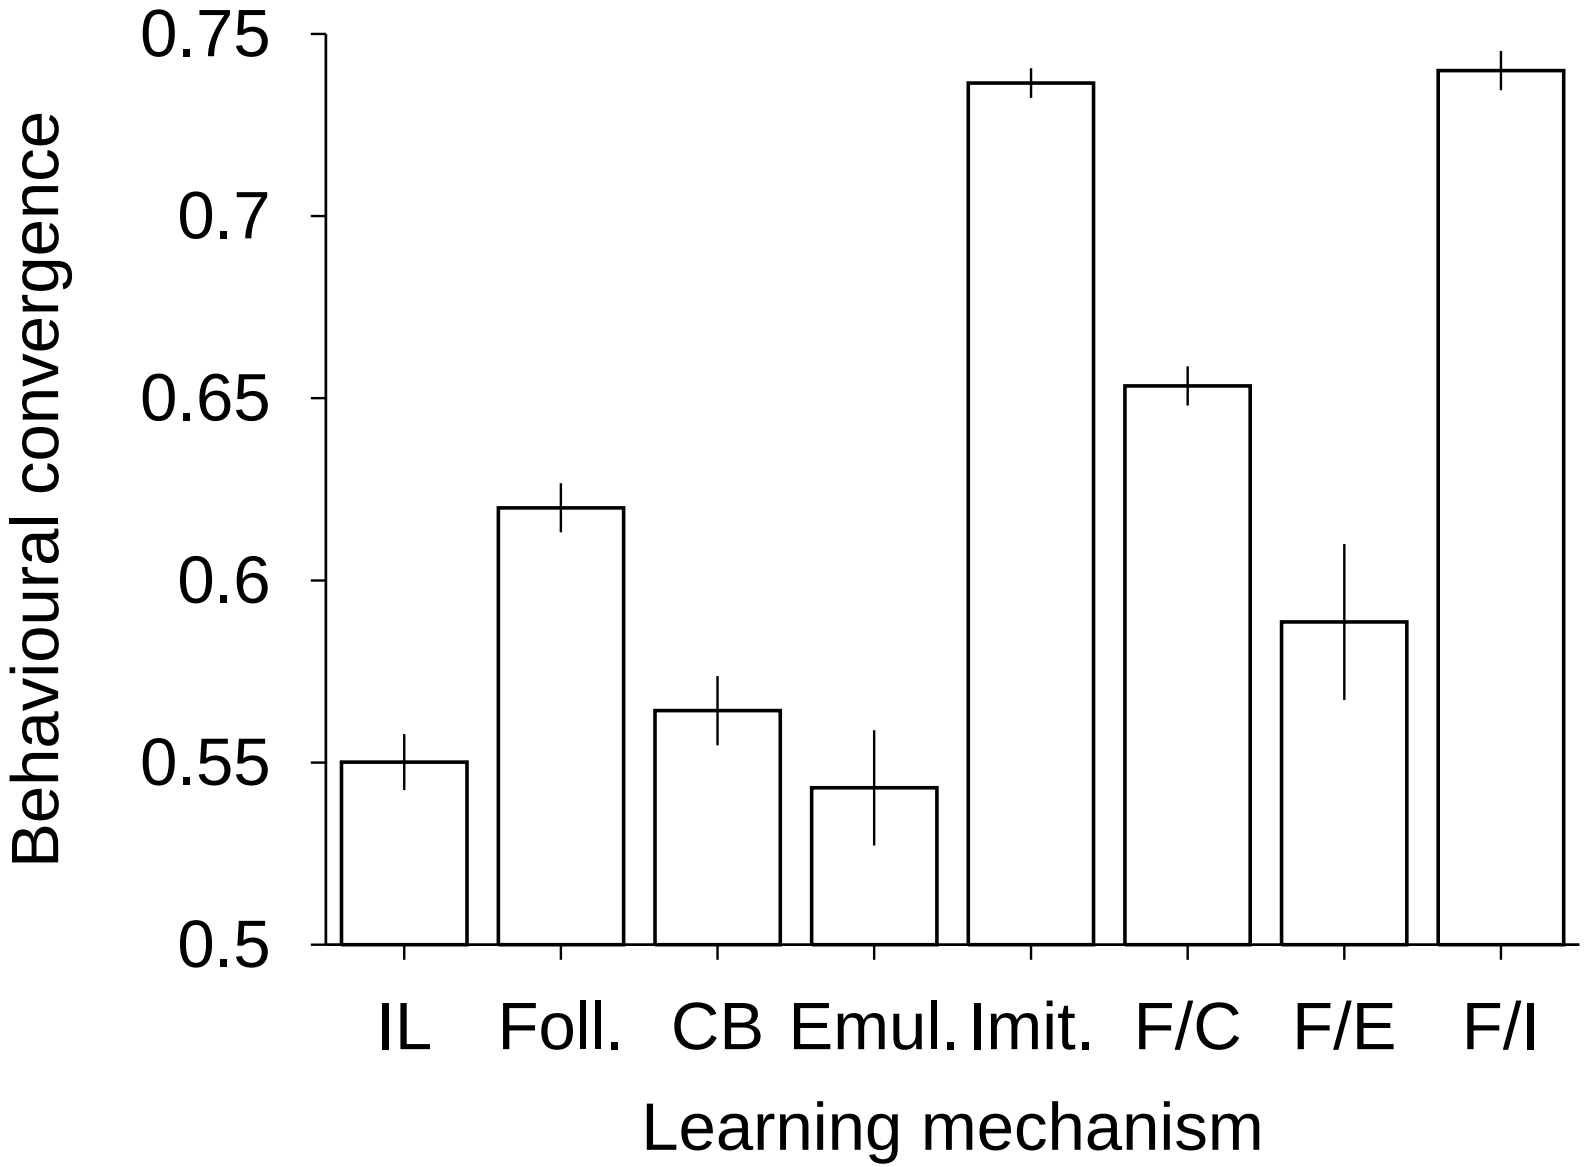

Behavioural convergence

- Score of 0.55 for individual learning (compare 0.18 for random behaviour). Similar behaviour sometimes just means similar experiences.
- Convergence high for imitation — unsurprising.
- Convergence low for emulation, despite its success.

Frequency of sample behaviours

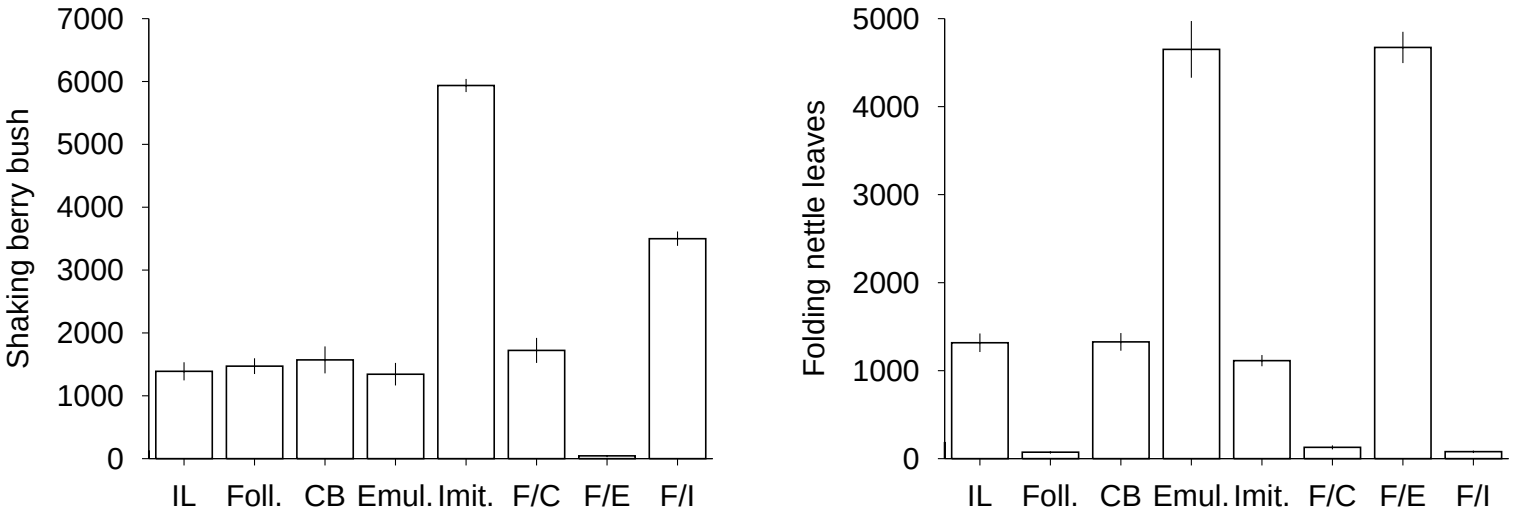

- The easy, low-payoff action of shaking a berry bush is common amongst imitators.
- The trickier action of learning to brave nettle stings in order to fold up the leaves and eat them is seen amongst emulators.
- Shows that imitators are very conservative.

Conclusions

- Complex imitative learning not always the best strategy.
- Effectiveness of any social learning strategy will depend on details of the agents and their environment.
- Noble and Todd (2002) argued that simple mechanisms were inherently more likely to evolve, and that proposing a simple mechanism is theoretically parsimonious.
- Stronger argument possible: simple mechanisms may actually perform *better*.

Conclusions

- Behavioural convergence not a reliable indicator of social learning.
- Evolution of social learning mechanisms not straightforward: a “following” mutation has negative fitness for emulators and positive for imitators.
- Focus on imitation may be misplaced, and multi-agent system builders should look to all corners of the animal kingdom for inspiration.

References

Noble, J., & Franks, D. W. (2002). Social learning mechanisms compared in a simple environment. In Standish, R. K., Bedau, M. A., & Abbass, H. A. (Eds.), *Artificial Life VIII: Proceedings of the Eighth International Conference on Artificial Life*, pp. 379–385. MIT Press, Cambridge, MA.

Noble, J., & Franks, D. W. (2004). Social learning in a multi-agent system. *Computing and Informatics*, *22*, 101–114.

Noble, J., & Todd, P. M. (2002). Imitation or something simpler? Modelling simple mechanisms for social information processing. In Dautenhahn, K., & Nehaniv, C. (Eds.), *Imitation in Animals and Artifacts*, pp. 423–439. MIT Press.