7. Evolution, Co-evolution (and Artificial Life) Part 2

Modelling Social Interaction in Information Systems

http://davidhales.com/msiis

David Hales, University of Szeged

dave@davidhales.com

Summary

- Last time we looked at:
  - Genetic Algos. (GA) Genetic Prog. (GP)
  - Concept of co-evolution
  - Co-evolution of strategies in games (PD example)
  - General concept of an ESS, relation to Nash Eq.
  - Evolution of cooperation on fixed spatial grid (CA)

Games on graphs

- Games played on various fixed graph topologies
- Large diverse area of research *
- Capture wide range of interaction (or social) structures:
  - A fully connected graph = “mean field” random mixing
  - A regular lattice = fixed spatial structure such as a 2D CA
  - Random / Small-World / Scale-Free
- Analysis difficult so simulation often used
- In general however, like the CA model we have seen, it’s all to do with clusters of cooperators forming

***György Szabó & Gábor Fáth (2007) Evolutionary games on graphs. http://arxiv.org/pdf/cond-mat/0607344v3.pdf [large comprehensive review of area]**

Games on graphs

- An rule on *some* fixed graphs:
  - Cooperation *may* be sustained when b/c > k
  - Where k = average links per node (degree)
  - c = cost of cooperation, b = benefit of cooperation (PD payoffs being: T=b, R=b-c, P=0, S= -c)
  - Shown for 1D line, 2D lattice, random and scale-free topologies
  - Called "network reciprocity"

Ohtsuki et al (2006) A simple rule for the evolution of cooperation on graphs and social networks. Nature, 441(7092), 502-505.

Aside: cost / benefit

- Cost / benefit parameterisation of PD payoffs often used
- It reduces the number of parameters
- Makes clear the cost and benefits of cooperation
- Often interpreted as a "donation game":
  - An agent donates to another or not (C or D)
  - Donation involves a cost (c) to the donor and a benefit (b) to the receiver
  - Not donating leads to no cost or benefit (nothing happens)
  - Asynchronous game playing (no need for both agents to select a move at the same time)

Aside: other (non-graph) rules

- Kin selection: r > c/b
  - Where r is "coefficient of relatedness" (proportion of same genetic material)
  - Called: Inclusive fitness, Hamilton's rule
- Direct reciprocity: w > c/b
  - Where w is prob. of future interaction between same pair of individuals
  - Note: tit-for-tat works by "punishing in future"
- Indirect reciprocity: q > c/b
  - Where q is prob. of knowing another's reputation
  - Hence can punished / reward those only meet once

Nowak, M. (2006) Five rules for the evolution of cooperation. Science, 314(5805), [good short summary includes "network reciprocity"]

Evolving social structures

- So far we have only considered fixed social structures
- Strategies evolve but the structures are static
- However, cooperation can be facilitated by evolving the structures in addition to strategies
- Class of simple models use “tags” to evolve a dynamic structure

What are tags?

- Tags = observable labels, markings or social cues
- Agents display and can observe tags
- Tags evolve like any other trait (or "gene")
- Agents select game partners based on them having the same tag as their own
- John Holland (1993) => tags powerful "symmetry breaking" function in "social-like" processes
- In GA-type interpretation, tags = parts of the genotype reflected directly in the phenotype
- In a cultural interpretation tags are copyable observable traits like style of dress, accent etc.

HOLLAND, J. (1993). The Effect of Labels (Tags) on Social Interactions. Working Paper 93-10-064. Santa Fe Institute. Sante Fe, New Mexico.

Tag models

- There are many variants of tag models
- Tags may be represented as bit strings, integers, real numbers or any enumeration type (such as colours)
- Strategies may be simple C or D or more complex kinds of interaction
- However, all models follow a similar pattern
- We illustrate with a simple tag model

A simple tag model

Agents – a tag and a pure PD Strategy

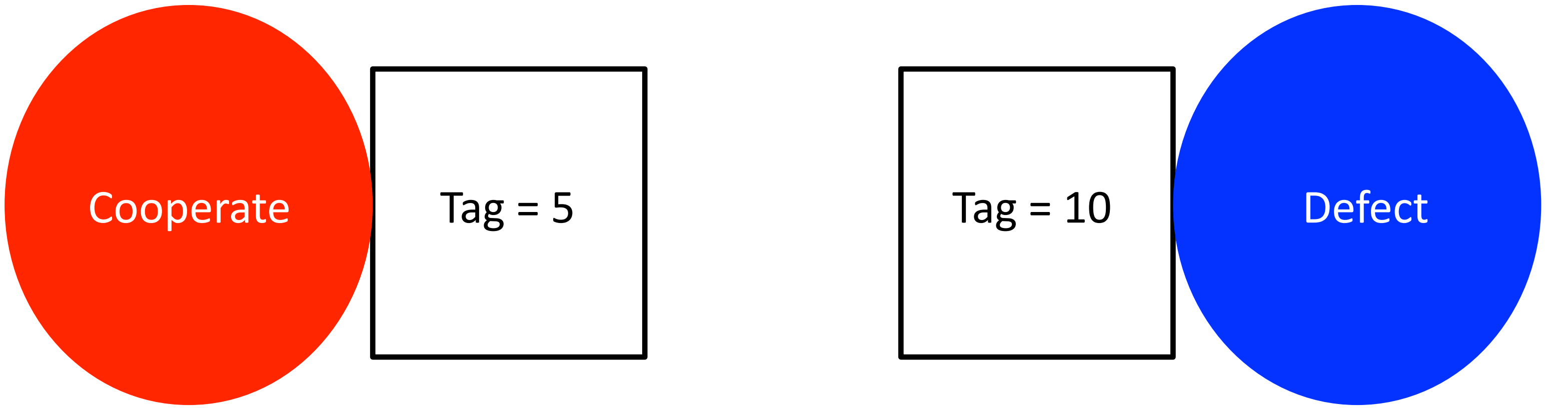

Tag = (Say) some integer
Game Interaction between those with same tag (if possible)

HALES, D. (2001). Cooperation without memory or space: Tags, groups and the prisoner's dilemma. In S. Moss & P. Davidsson (Eds.), Multi-Agent-Based Simulation, Springer, pp. 157-166.

Algorithm

```
Initialise all agents with randomly selected strategies and tags
LOOP some number of generations
    LOOP for each agent (a) in the population
        Select game partner (b) from the population
            with matching tag or randomly if no match exists
        Agent (a) and (b) invoke their strategies
            receiving the appropriate payoff
    END LOOP
    Reproduce agents in proportion to their average payoff
            with some small probability of mutation
END LOOP
```

How tags work

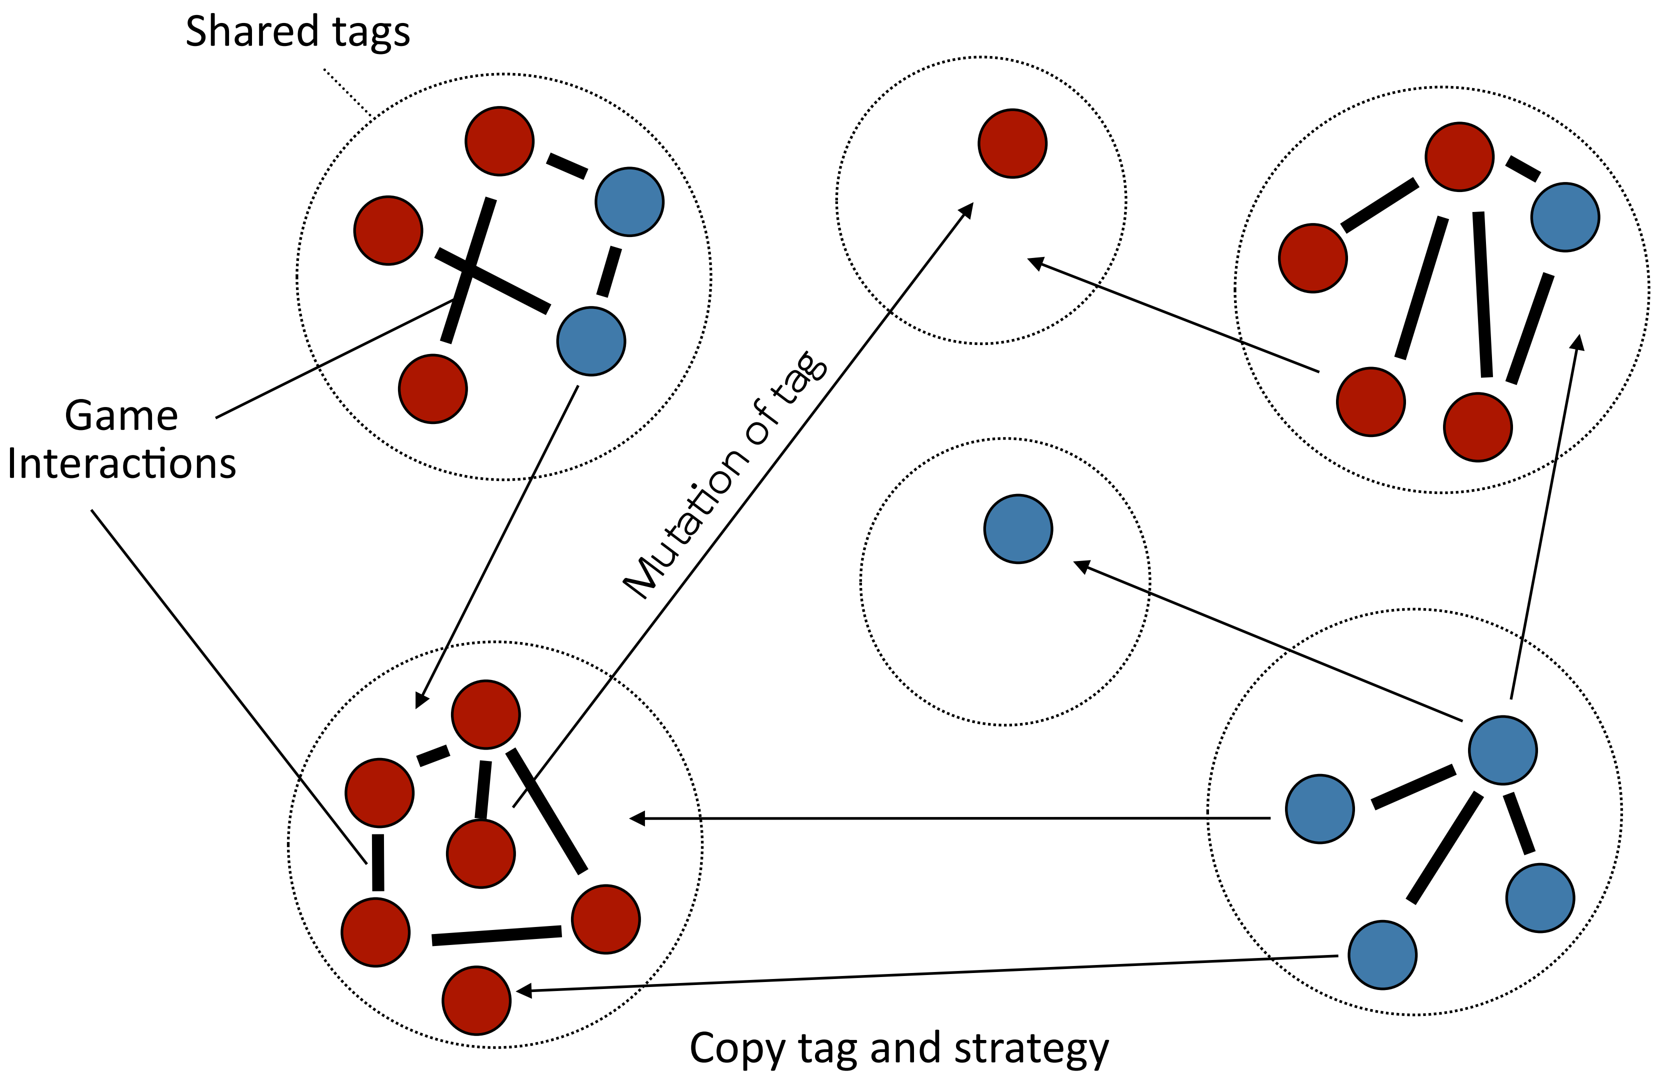

Visualising the process

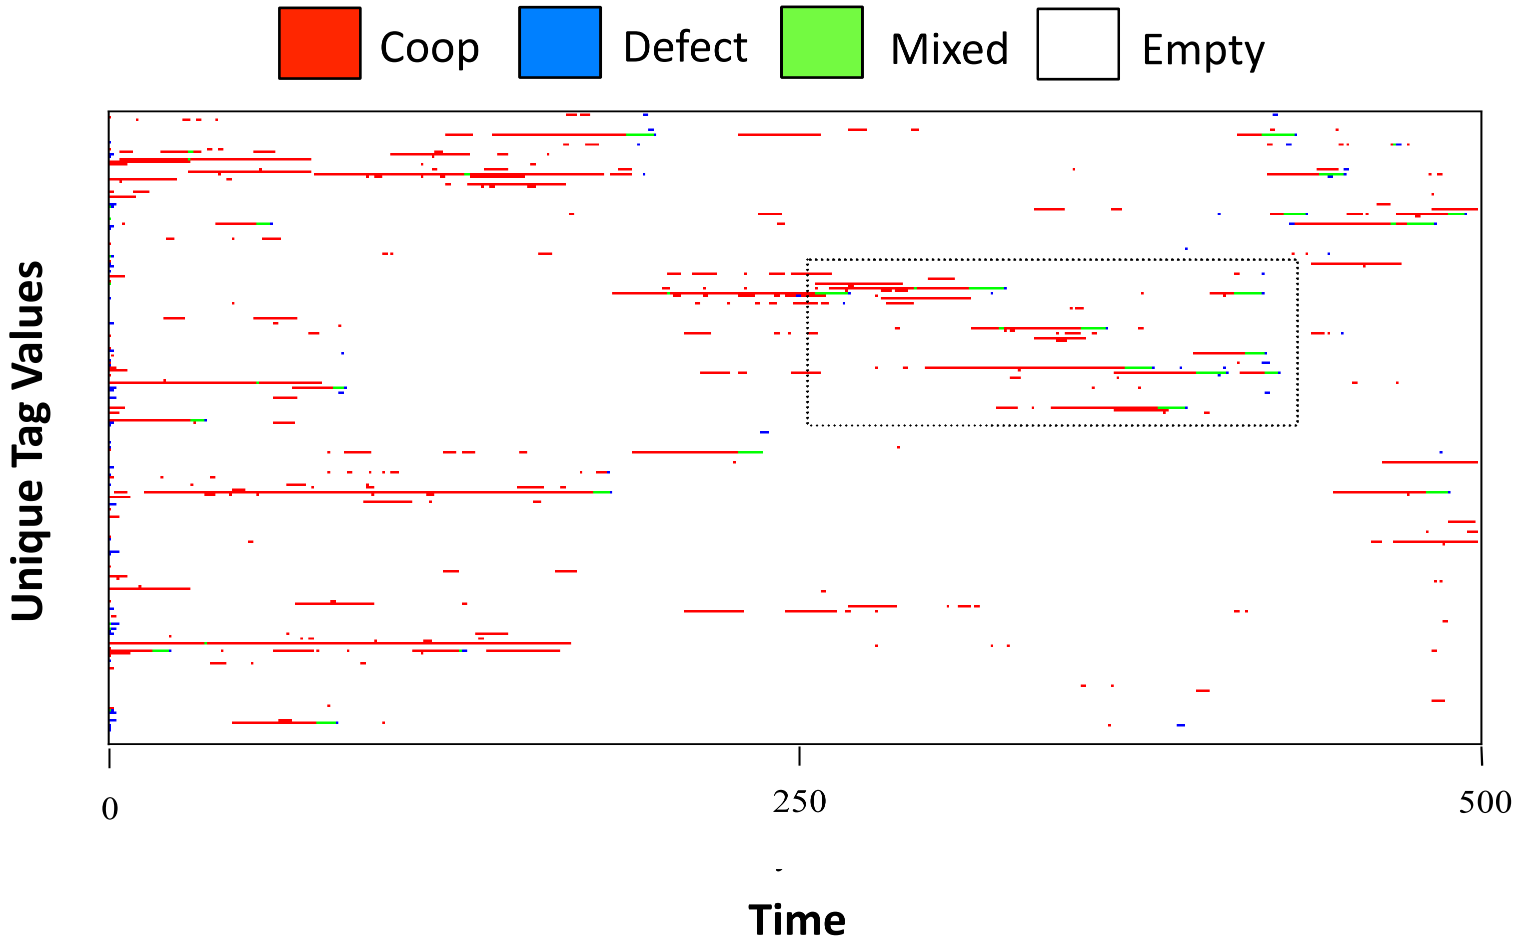

Visualising the process

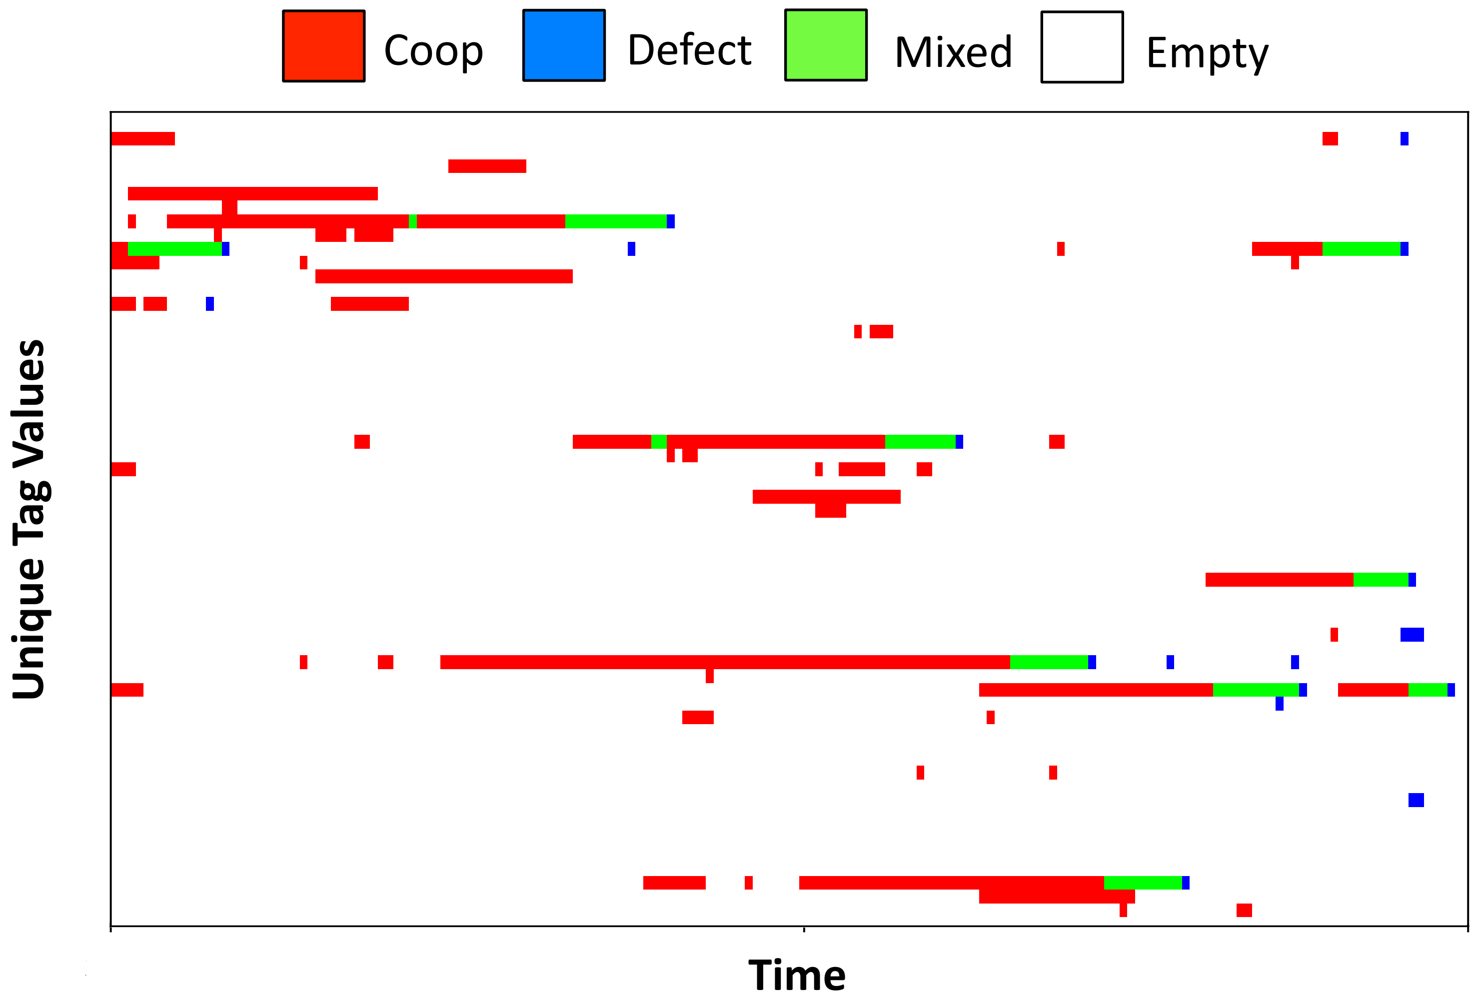

Change your tags fast…

- Groups have to be formed more quickly than they invaded and killed
- New groups are formed by mutation on the tag
- Old groups are killed by mutation on the strategy
- So if tag mutation > strategy mutation this should promote cooperation?
- Test it by looking at the existing models and implementing a new one

Paper: Traulsen A, Nowak MA (2007) Chromodynamics of Cooperation in Finite Populations. PLoS ONE 2(3): e270.

Tag / strategy mutation rate

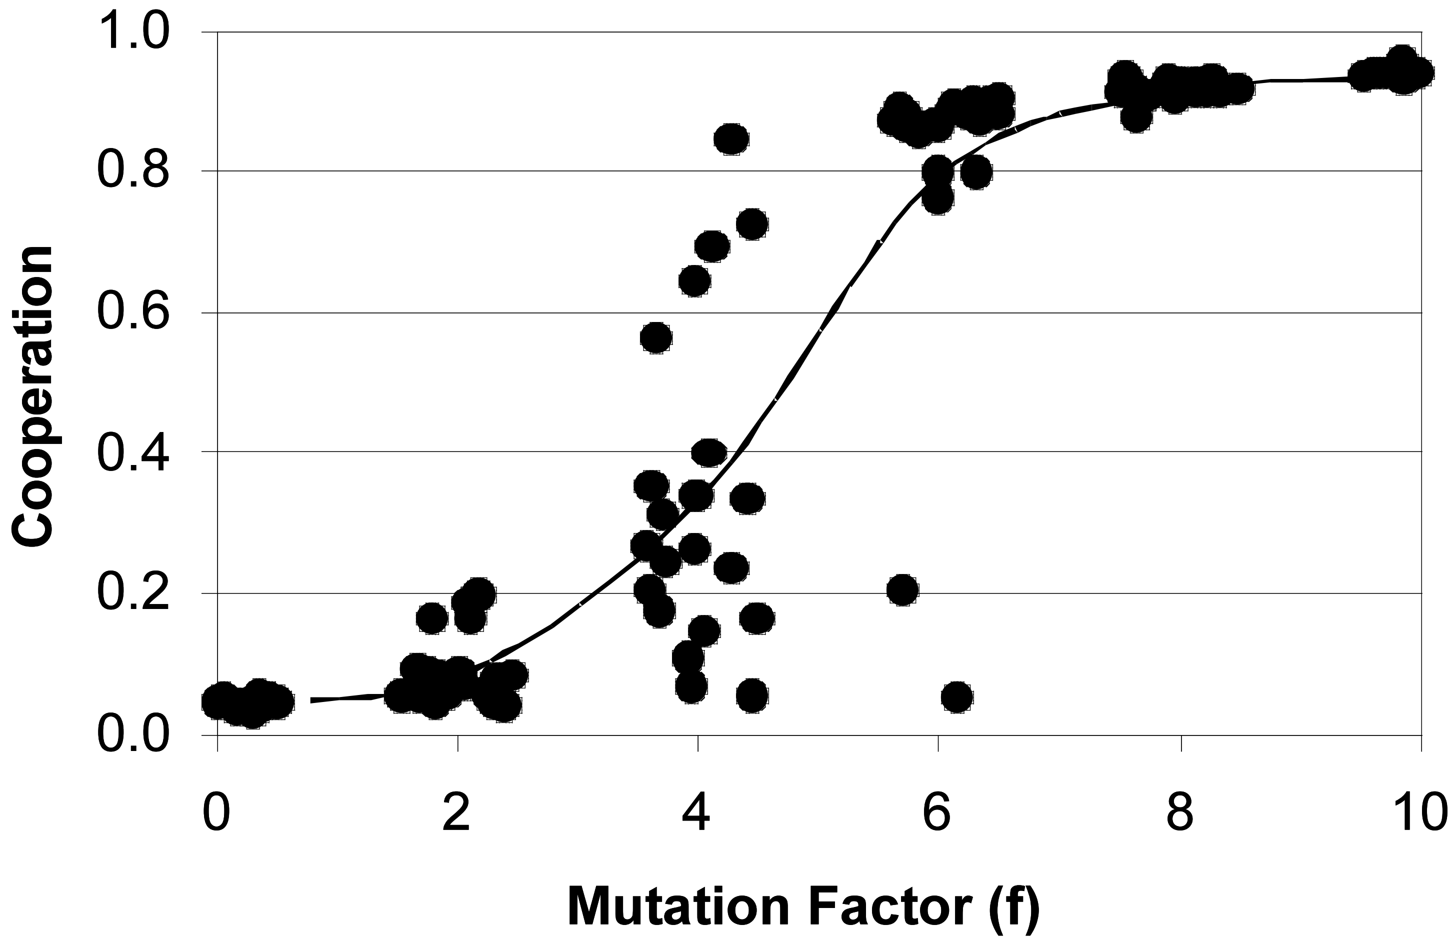

Network rewire - dynamic graphs

- Suppose we have a set of nodes in a network
- Each stores some maximum number of undirected links to other nodes (a graph)
- Each stores a PD strategy (C or D)
- Periodically nodes play PD games with their neighbours (those linked to) accumulating an average fitness

Paper: Hales, D. & Arteconi, S. (2006) Article: SLACER: A Self-Organizing Protocol for Coordination in P2P Networks. IEEE Intelligent Systems, 21(2):29-35

Paper: Santos F. C., Pacheco J. M., Lenaerts T. (2006) Cooperation prevails when individuals adjust their social ties. PLoS Comput Biol 2(10)

Network rewire model

Each node *p* periodically (Play game):

    if *p* has no neighbours

        link to a random node from population

    play PD game with a randomly chosen neighbour

Each node *p* periodically (Replication):

    q = random node from entire population

    If $\text{fitness}_q > \text{fitness}_p$

        drop each current link with high probability

        link to node *q* and copy its strategy and links

        mutate strategy and links

Network rewire movies

Video: SLACER: https://youtu.be/3X0JSuqCxkg
Video: SLAC: https://youtu.be/UkfDejU7-a0

What is this showing?

- Can be viewed as a form of emergent "cultural group selection"
- with nodes using tournament selection
- Except it's not *really* group selection in the the traditional meaning of the term
- It's individual selection that dynamically creates social interaction structures that support cooperative groups
- Term "group selection" is controversial in evolutionary biology!

Overview paper: Hales, D., Shutters, S. (2012). Cooperation through the endogenous evolution of social structure. Proc. of the Complex 2012 conference in Santa Fe, NM.

Video: me talking about model (2007) @ TU Dresden, http://youtu.be/z9H5FqDsJ24

Cultural Evolution

- When behaviours, strategies, beliefs etc. are copied between agents within their lifetime and / or outside of genetic inheritance
- This is sometimes called cultural transmission or cultural evolution
- Hence, if there is some mechanism for replication (imitation), innovation (mutation) and selection (selective imitation) then:
- Evolutionary algorithms can be viewed as cultural evolution rather than genetic evolution
- Analysis has looked at gene-culture co-evolution based on this view (or dual inheritance theory DIT)

Book: Boyd, R. and P. J. Richerson. 1985. Culture and the Evolutionary Process. Chicago: University of Chicago Press.

Aside: Cultural Evolution (memes)

- Richard Dawkins introduced the term “meme” in his famous book “the selfish gene”
- His idea was to show that evolution can occur outside of biology in a similar Darwinian way
- The “meme” is the cultural equivalent of the “gene” in biology
- Can be thought of as any unit of culture that can be copied accurately – such as a song, a joke, a way of throwing stones far etc.
- A group of people started to create an area called “memetics” that would study this form of cultural evolution
- But it died out as a serious area of study (that’s evolution for you!)
- The idea of cultural evolution in a memetic sense is still sometimes used in scientific work but other terms are often used to describe it
- The term “meme” seems now to be reserved for “internet memes”

Book: Dawkins, Richard (1976). The Selfish Gene. Oxford University Press

Artificial Life

- So far we have assumed:
  - Agents exist
  - Bounded space of possible behaviours (but not in GP!)
  - Interact in some fixed well defined way (game)
  - Get an explicit fitness value from interaction
  - Evolve through variation, replication, selection based on fitness
- can such processes emerge in the first place without “programming them in”?

Artificial Life

- ALife is an area that attempts to understand life by creating versions of it:
  - In computer simulations (software – *in silico*)
  - In robots (hardware)
  - In test tubes (wetware – *in vitro*)
- How would we know if we created life?
- What is “life” as apposed to “non-life”?
- ALife tries to throw light onto these kinds of philosophical questions
- Similar method to AI – understand it by building it

Evolving self-replication

- Incorporate the replication function into agents
- Evolve the function like other agent behaviour
- Hence evolving agents that are good at making copies of themselves
- Remember von Neumanns self-replicating CA?
- How could we do this in a computer simulation?
- Tom Ray created a system (early 1990’s) called “Tierra” that did just that to explore what has been termed **open-ended co-evolution**
- This is considered formative work in **artificial life**

Open-ended evolution

- Given a very large (infinite?):
  - Space of agent behaviours (universal computer?)
  - Set of possible agent interactions
  - Sufficiently complex environment
- Could this lead to:
  - Co-evolutionary processes leading to perpetual novelty, increasing complexity, “new games”, explosion of diverse replicating forms?

(rather vague concept – but so is life!)

Paper: Soros &. Stanley (2014) Identifying Necessary Conditions for Open-Ended Evolution through the Artificial Life World of Chromaria. In: Proc. of the 14th Int. Conf. on the Synthesis and Simulation of Living Systems (ALIFE 14) Cambridge, MA: MIT Press, 2014

Tom Ray's Tierra

- A simulated virtual machine with fixed size memory
- Agents are programs written in the assembly language of that machine (designed to be not too "brittle"):
  - Compete for processor and memory
  - Evolve through mutation, death and *self-replication*
  - Evolve novel ways to parasitise other programs in memory to reproduce
- Ray seeds system with an initial hand-coded self-replicating program (self inspection of code allowed)
- Complex dynamic parasitic ecologies emerge
- Ray was a biologist / ecologist trying to understand complex interactions between living things

Paper: Tom Ray (1992) Evolution, Ecology, and Optimization of Digital Organisms. Santa Fe Working Paper: 1992-08-042

Short documentary: http://youtu.be/Wl5rRGVD0QI

Aside: Core War game

- A simulation game “core war” was developed long ago
- People design algorithms to fight other algorithms to take over the memory space in a vm running “redcode”
- Some recent contestants have been *evolved*
- On some level, could one view the entire internet as something similar to this?
- Yet winning algorithms will need to make people run them too – social engineering
- Viral programming was researched early in networking / computing but has got a bad name due to malware / botnets using the technique for nefarious purposes

See: http://en.wikipedia.org/wiki/Core_War

Emerging agents?

- How do we get “agents” in the first place?
- What are they?
  - Autonomous: act to achieve their own goals independently
  - Have some boundary between themselves, other agents and the environment
  - Can sustain (and replicate?) themselves over time
- How do we identify some “string of bits” in a computer is an agent and some others as not?

Emerging agents

- Associated concepts (metabolism):
  - Autocatalysis (self-reinforcing chemical interactions)
  - Homeostasis (system stability)
  - Autopoiesis (self-sustaining, reproducing, complex process) [rather vague / controversial concept]
- Some of these are explored in an area called "Artificial Chemistry"
  - Simple "reactions" producing complex outcomes
  - Various ways of modelling / various results

Paper: Dittrich et al. (2001) Artificial Chemistries - A Review. Journal of Artificial Life 7: 225–275.

bootstrapping evolution?

- How does evolution get started?
- Such that it can evolve metabolism like processes?
- Can it start with selection of chemical reactions?
- That come to form higher-level replicators?
- Which came first the metabolism or the replicator?
- In vitro: make life from base chemicals
- In vivo hack: remove genes until you have the simplest reproducing cell?

Artificial Life

- If you are interested in ALife see recent overview of the area:
  - **Aguilar W, Santamaría-Bonfil G, Froese T and Gershenson C (2014) The past, present, and future of artificial life. Front. Robot. AI 1:8.**
- Chris Langton (remember his work with CA?) named the area (mid 1980's) and there is a journal and conferences.
- Some great introductory videos (by ALife pioneer Dave Ackley):
  - https://www.youtube.com/user/DaveAckley/

People: John Holland

- Interdisciplinary thinker, background in physics, maths, comp. sci, psychology.
- Invented modern genetic algorithms
- Interested in "complex adaptive systems" (CAS)
- Worked with models similar to Ray's Tierra. Developed "Echo" model
- Produced the "classifier" agent which evolves it's internal (production) rule set
- Has written widely on self-organisation and emergence in easy to understand language

Book: John Holland (1998) Emergence: from chaos to order. Perseus Books. [popular science introduction]
Book: Holland, John (1975). Adaptation in Natural and Artificial Systems. Cambridge, MA: MIT Press [Mainly GA's but later editions contain Echo model]
Video: Holland 2008 lecture: http://youtu.be/6aN6PlsvkpY

People: Richard Dawkins

- Background zoology / biology
- Great populariser of evolution as a way to explain life and even culture (meme)
- Takes “gene’s eye view” which is not accepted by all evolutionary biologists (e.g. Stephen J. Gould)
- Sees evolution as applying to individual genes not organisms or larger collectives
- Often debates those who believe in God saying that evolution is a better explanation for life
- Argues (like Daniel Dennett) that the “abstract evolutionary algorithm” can explain complex biological, social and cultural phenomena

Book: Dawkins, Richard (1976). The Selfish Gene. Oxford University Press

Readings and Questions

- Readings
  - Flake (1998) Chapter 5 – Adaptation
  - Gilbert et al (2005) Chapter 10 – Learning and Evo.
- Questions
  - Could evolution help design practical P2P protocols?
  - Is abstract evolution tautological? If so can it tell us anything about the “real world”?
  - What is life? How can we know something is alive?
  - Does evolution have a direction?
  - Why do you think models of cooperation / altruism are prominent in biology and artificial life?