Agent-Based Modelling in NetLogo

Evolution of Cooperation

David Hales

www.davidhales.com/abm-netlogo

Evolving Agents

- Agents can be programmed with fixed behaviours (e.g. segregation model)
- Agents can learn individually from their experience (e.g. El Farol bar model)
- Alternatively agents can evolve their behaviour
  - Agents that do well at some task reproduce more
  - Agents that perform poorly reproduce less
  - Random changes in behaviour introduce new behavioural variants

Evolving Agents

- Any system implements evolution if there is:
  - *Replication* (agents reproduce copies of themselves)
  - *Variation* (agents differ in some way)
  - *Selection* (based on their characteristics some agents are more likely to reproduce than others)

There are many ways to implement these functions in computer models. We could spend a whole course on it. But we will focus on a couple of examples.

Evolving Agents

- Agents can be viewed as biological entities evolving through reproduction and death
- Or cultural entities evolving their strategies through imitation (copying) and random innovation (trying something new)

Evolving Agents

- the reproductive success of an agent is often termed *fitness*
- in a model
  - Fitness may be implicit or explicit
    - implicit fitness arises from agent behaviour
    - explicit fitness is a value calculated explicitly by the model based on some evaluation of the agent characteristics
  - Selection may be implicit or explicit
    - Implicit selection arises from agent behaviour
    - Explicit selection is implemented by the model based on fitness

Example – wolf sheep

- In the wolf sheep model we looked at:
  - Implicit fitness based on survival and giving birth
  - Implicit selection based on being eaten or finding things to eat
  - Variation: two species, wolves and sheep
  - Note: evolution stopped when variation ended

Sample Models > biology > evolution > wolf sheep predation

Example – bug hunt speeds

- In the bug hunt speeds model:
  - Agents are bugs with different speeds (variation)
  - User hunts them by clicking on them (selection)
  - New bug born when user catches one (replication)
  - implicit fitness: longer bug can stay alive
  - Hence fast moving bugs tend to be selected because harder for user to catch

Sample Models > biology > evolution > bug hunt speeds

Explicit fitness & selection

- Many evolutionary models include explicit fitness, selection and variation functions
- Such models include a way to:
  – calculate a fitness value for each agent
  – select agents to reproduce based on the fitness
  – Produce agents with new characteristics

An Evolutionary Algorithm

- Initialise N agents with some initial characteristics
- For some number of generations
  - Calculate fitness of each agent (somehow)
  - Reproduce a new generation of N agents based fitness (somehow)
  - With some low probability randomly change characteristics of each reproduced agent (mutate)

Fitness

- Way fitness calculated depends on what model is about
- For optimisation problems fitness is often calculated with an objective function applied to each agent (which is a candidate solution to a problem)
- In agent systems fitness often depends on *interactions between agents* (not optimisation of objective function)
- One way to capture this abstractly is to have agents play simple "games" with each other
- We will look at an example game called the **Prisoner's Dilemma** (taken from **game theory**)

Reproduction based on fitness

- Assuming we have calculated a fitness for each agent
- How to generate the new population of N agents based on their fitness?
- Many possible ways as long as those with higher fitness have more chance of making copies of themselves into the next generation
- We will use a very simple method called **Tournament Selection**
- Another popular method is called **Roulette Wheel Selection**

Tournament Selection

- Many variants of tournament selection but here is a very simple variant:
- Repeat until next generation is full (N)
  - Select a random pair of agents (with replacement)
  - Reproduce the one with the highest fitness
  - Or a random one if both have the same fitness

Game Theory

- Game theory analyses the behaviour of agents playing games
- The games aim to capture, abstractly, interaction possibilities between agents
  - Where agents have choices of actions and different preferences over outcomes
  - And the outcomes are determined by the actions several agents take

Game Theory

- The actions agents can take (the "moves" in the game) are called *strategies*
- The outcomes for each agent are called "payoffs" and can be used as fitness values
- We will look at a simple two player game
  - called the "Prisoner's Dilemma"
  - in which each player (agent) has a choice of one of two strategies then the game ends
- When strategies are evolved based on payoffs then this is called **evolutionary game theory**

The Prisoner's Dilemma (PD) - "payoff matrix"

Game is a PD when: T > R > P > S and 2R > T + S

Player 1 / Player 2	C	D
C	(3) R R (3)	(5) T S (0)
D	(0) S T (5)	(1) P P (1)

The Prisoner's Dilemma (PD) - example games

Players =>	P1	P2	P1	P2	P1	P2	P1	P2
Moves =>	C	C	C	D	D	C	D	D
Payoffs =>	R	R	S	T	T	S	P	P
Values =>	3	3	0	5	5	0	1	1
Total =>	6		5		5		2	

A contradiction between collective and individual
interests: a so-called "social dilemma"

Story behind the PD game

- Two people are arrested for robbing a bank. The police interrogator can not prove it was them. He puts each in a different room and offers them a deal: "say the other did it and you go free. Keep quiet and you will be charged with a lesser crime"
- Keeping quiet is playing C (cooperate)
- Betraying the other is playing D (defect)
- P (punishment), R (reward), T (temptation) S (sucker)

Another interpretation of PD

- Imagine agents wander in an environment gathering resources and interacting
  - When an agent meets another it can either:
    - Share / trade resources with the other agent (C) OR
    - Steal the resources from the other agent (D)
  - If both share both benefit (mutual Cooperation)
  - If both steal neither does well (mutual Defection)
  - If one steals and another shares then the thief does really well and the victim does very badly

Evolving PD strategies

- Many disciplines use PD and evolution to gain insights into:
  - Cooperation and conflict in animal behaviour, Evolution of Life (Biology)
  - Societies, institutions, conflict, collective action, "social contracts" (Sociology, Political science)
  - Morality, rationality (Philosophy)
  - How to program open distributed software (P2P)

Robert Axelrod (1984) Evolution of Cooperation. Basic Books
Matt Ridley (1996) The Origins of Virtue. Penguin Books
Bram Cohen (2003) Incentives build robustness in Bittorrent. 1st Workshop on the Economics of Peer-2-Peer Systems.

A **mean-field** interaction model

- N agents initialised randomly to C or D
- Repeat some number of generations:
  - Game interaction: each agent is paired with another randomly chosen agent, plays PD
  - Reproduction: Generate new generation of N agents where fitness = average payoff
  - Mutate newly reproduced agents by flipping strategy with some small probability

Mean-field model

- ***What do you think will happen if we run this model?***

Mean-field model

- From *any* initial starting condition
- Evolution will quickly lead to Defect dominating the population and stay there
- This is called an **Evolutionary Stable Strategy** (or ESS)*
- A strategy is an ESS if a population all using it can resist “invasion” by other “mutant” strategies

***Book: John Maynard Smith (1982) Evolution and the theory of games. Oxford University Press**

Fixed lattice model (spatial PD)

- When interactions are non-random but structured in some way
- Then dynamics often get more complex
- An example of this is a classic model that situates agents on a 2D lattice (grid)
- Constraining their interaction *and* reproduction

Nowak & May (1992) Evolutionary games and spatial chaos. Nature 359, 826-829.
Nowak & May (1993) The Spatial Dilemmas of Evolution. Int. J. of Bifurcation and Chaos, Vol. 3, No. 1. 35-78

Lattice model

- Situate agents on a 2D lattice
- Randomly Initialise strategy of agent (C or D)
- Repeat some number of generations:
  - Interaction: each agent plays a PD game with its 8 neighbours (and itself)
  - Reproduction: each agent copies the strategy of the best performing (total payoff) agent in its neighbourhood (including itself)*

Note: In this model there is no mutation

* Make a random selection if several have the same best payoff

Lattice model

- Using PD payoffs: T = b, R = 1, P = 0, S = 0
- Outcomes explored for different b values
- This is a so-called “weak PD” since P=S
- ***What happens when we run this model?***

Lattice model observations

- Different values of b give different dynamics:
  - Often dynamical patterns in which freq. of cooperators change over time
  - Various threshold values for b leading to different dynamical regimes
- Generally, from *most* starting conditions:

  $b < 1$ cooperators take over (no longer PD)

  $1 < b < 1.8$ between 0.7 and 0.95 cooperators

  $1.8 < b < 2$ around 0.3 cooperators (chaotic)

  $b > 2$ defectors take over (no longer PD)

Lattice model observations

- Through analysis it can be shown that when:

  b < 1.8 only C clusters can grow

  b > 2 only D clusters can grow

  1.8 < b < 2 both C and D clusters can grow
- the latter interesting region produces complex dynamics
- hard to capture analytically because depends on interactions between C and D clusters

Growth of a Defector in an infinite sea of Cooperators

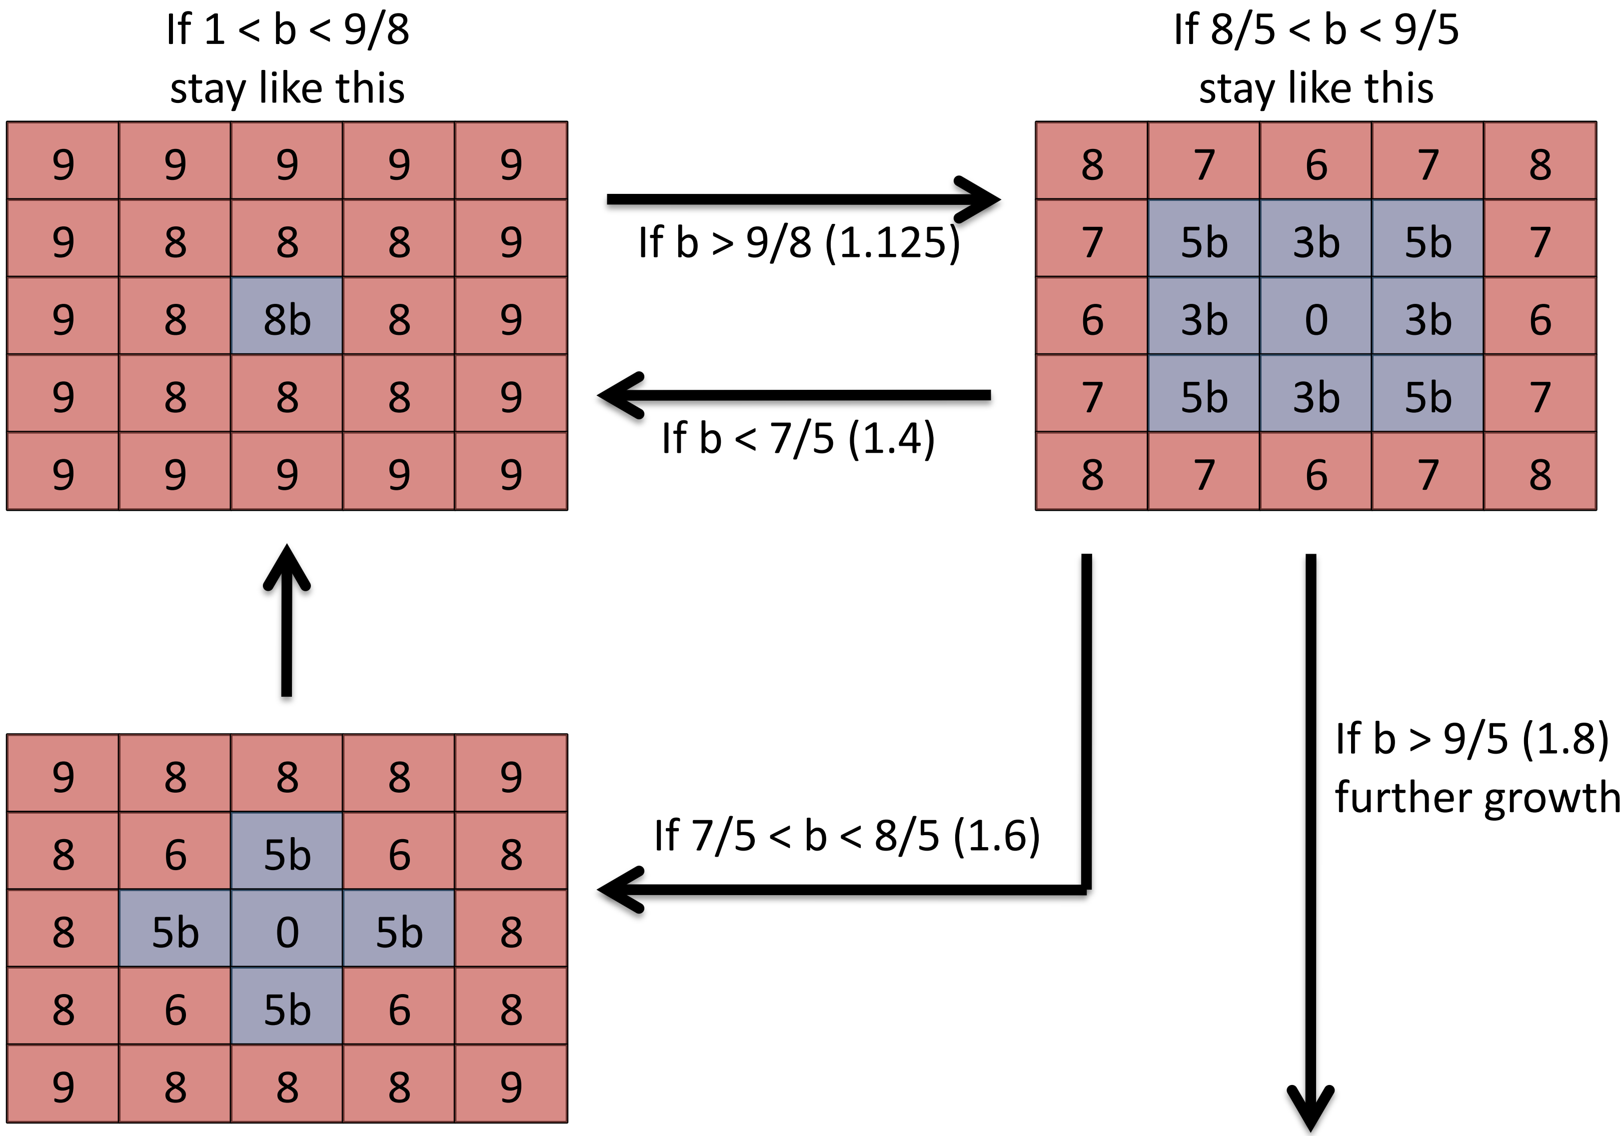

Growth of Cooperators in an infinite sea of Defectors

For 1 < b < 2

0	0	0	0	0	0
0	b	2b	2b	b	0
0	2b	4	4	2b	0
0	2b	4	4	2b	0
0	b	2b	2b	b	0
0	0	0	0	0	0

b	2b	3b	3b	2b	b
2b	4	6	6	4	2b
3b	6	9	9	6	3b
3b	6	9	9	6	3b
2b	4	6	6	4	2b
b	2b	3b	3b	2b	b

4	6	6	6	6	4
6	9	9	6	9	6
6	9	9	9	9	6
6	9	9	9	9	6
6	9	9	9	9	6
4	6	6	6	6	4

Aside: other fixed graph topologies

- Several works examine cooperation evolving on different fixed graph topologies
- See overviews in Szabo & Fath and Allen & Nowak *
- Overall general conditions are not "strong" and contingent on specifics

*** Gyorgy Szabo, Gabor Fath (2006) Evolutionary games on graphs.**
http://arxiv.org/abs/cond-mat/0607344

*** Allen, B. & Nowak, M. (2014) Games on graphs. EMS Surv. Math. Sci. 1. 113–151**

Task 1 – game on graphs

- Modify the mean-field PD model such that:
  - Agents are situated within a fixed network (such as random, ring, small-world, scale-free)
  - Each generation each agent plays one game with each of its neighbours in the network
  - Reproduction involves each node selecting as it's next strategy that of the highest scoring neighbour

Example solution will be given briefly next week!

Task 2 – lattice games

- Modify lattice model by:
  - Introducing mutation into the model
  - Allow user to set all payoffs (T, R, P, S)
  - Experiment with different mutation rates and payoffs (games)

Example solution will be given briefly next week!

Two further PD models

- Tag model – where a simple interaction structure evolves
- Network rewiring model – where agents evolve a dynamic network interaction structure

Tag based interaction model

- Where the interaction structure is itself evolved

Evolving interaction structure

- So far considered static interaction structures
- What if we evolve the interaction structure?
- A simple method uses “tags”
- Each agent stores a strategy AND a tag
- The tag can be observed by other agents
- It can be thought of as a genetic marker (e.g. eye colour) or a label or recognisable social cue (e.g. accent, dress)

John Holland (1993) The effect of labels (tags) on social interactions. Technical Report Working Paper 93-10-064, Santa Fe Institute.

Agents – a tag and a PD Strategy

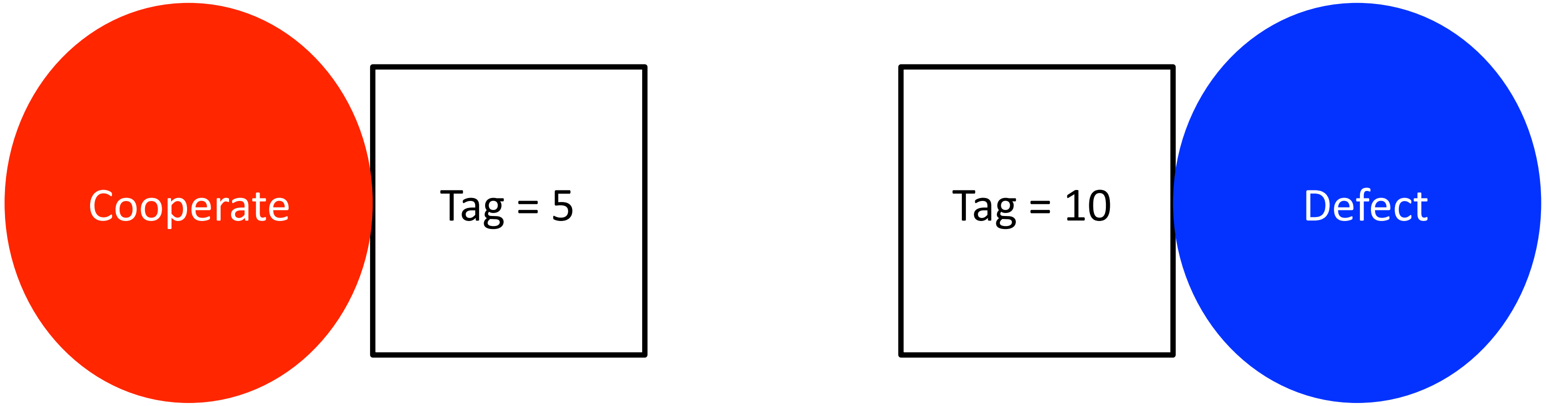

Tag can be represented by an integer, a bitstring, a colour a real number etc.

There should be many possible unique tags.

A **Tag** model

- Initialise agents with random strategies and tags
- Repeat some number of generations
  - Interaction: each agent plays a game of PD with a randomly selected other *with matching tags* (or random if no match exists)
  - Reproduction: generate next generation based on average payoff
  - Apply mutation to tag and strategy with small probability

Hales, D. (2000) Cooperation without space or memory: Tags, groups and the prisoner's dilemma. Multi-Agent-Based Simulation. Springer, Berlin

Tag model

- ***What happens if we run this model?***

Tag model observations

- Over a broad range of PD payoffs
- Very high C > 0.9 levels of cooperation quickly emerge from *any* initial starting condition
- The proportion of different tags in the population continually changes
- Those sharing the same tag can be thought of as dynamic interaction groups (tag groups)
- Too few tags leads breakdown of high C

Visualising the process

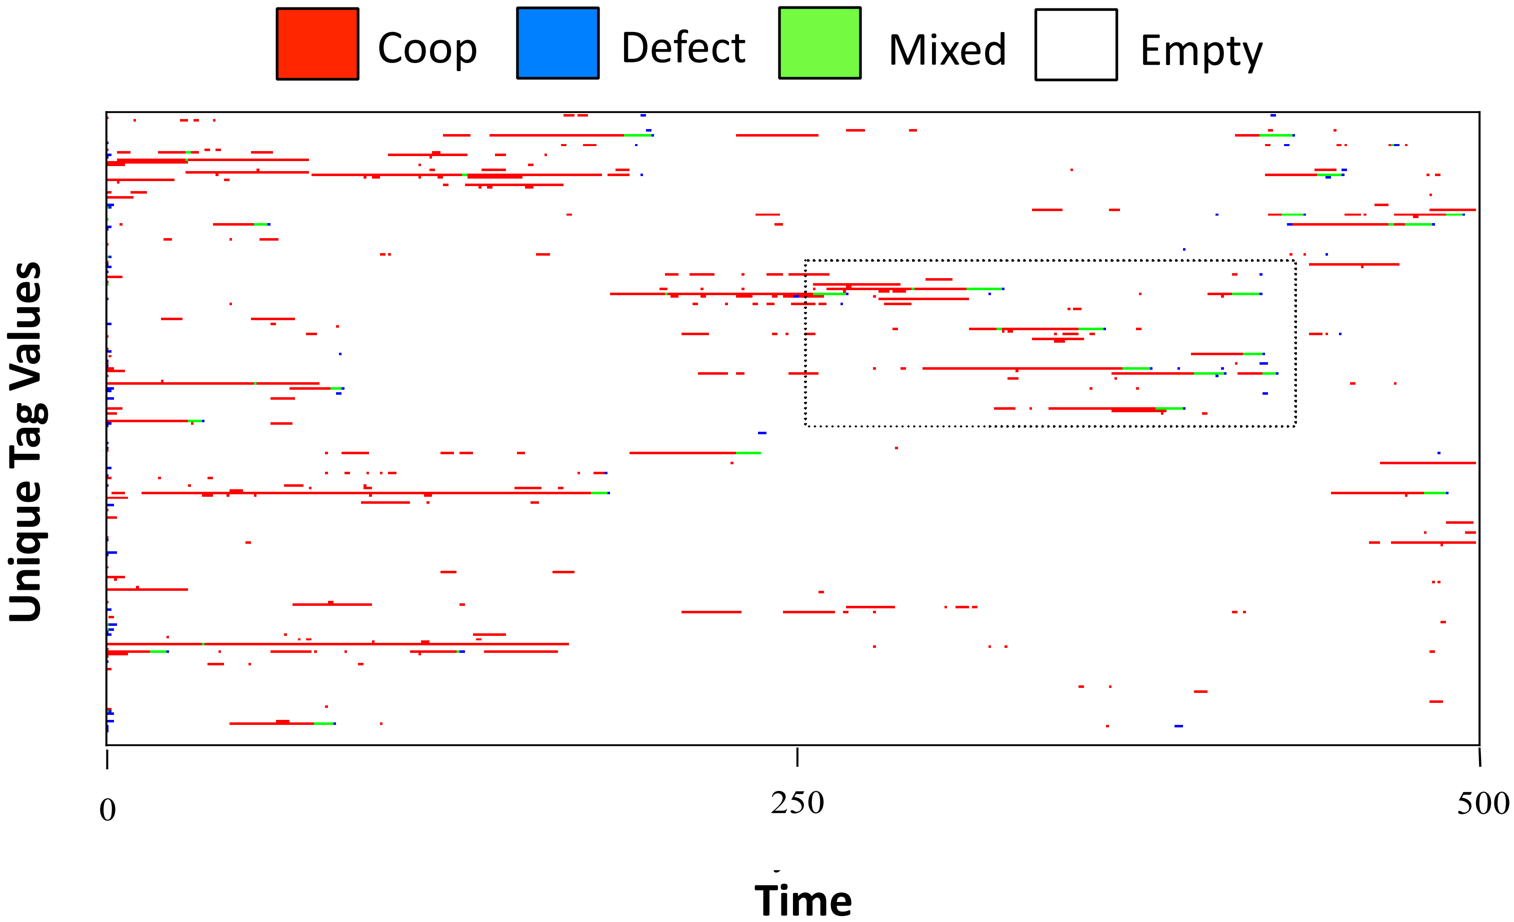

Visualising the process

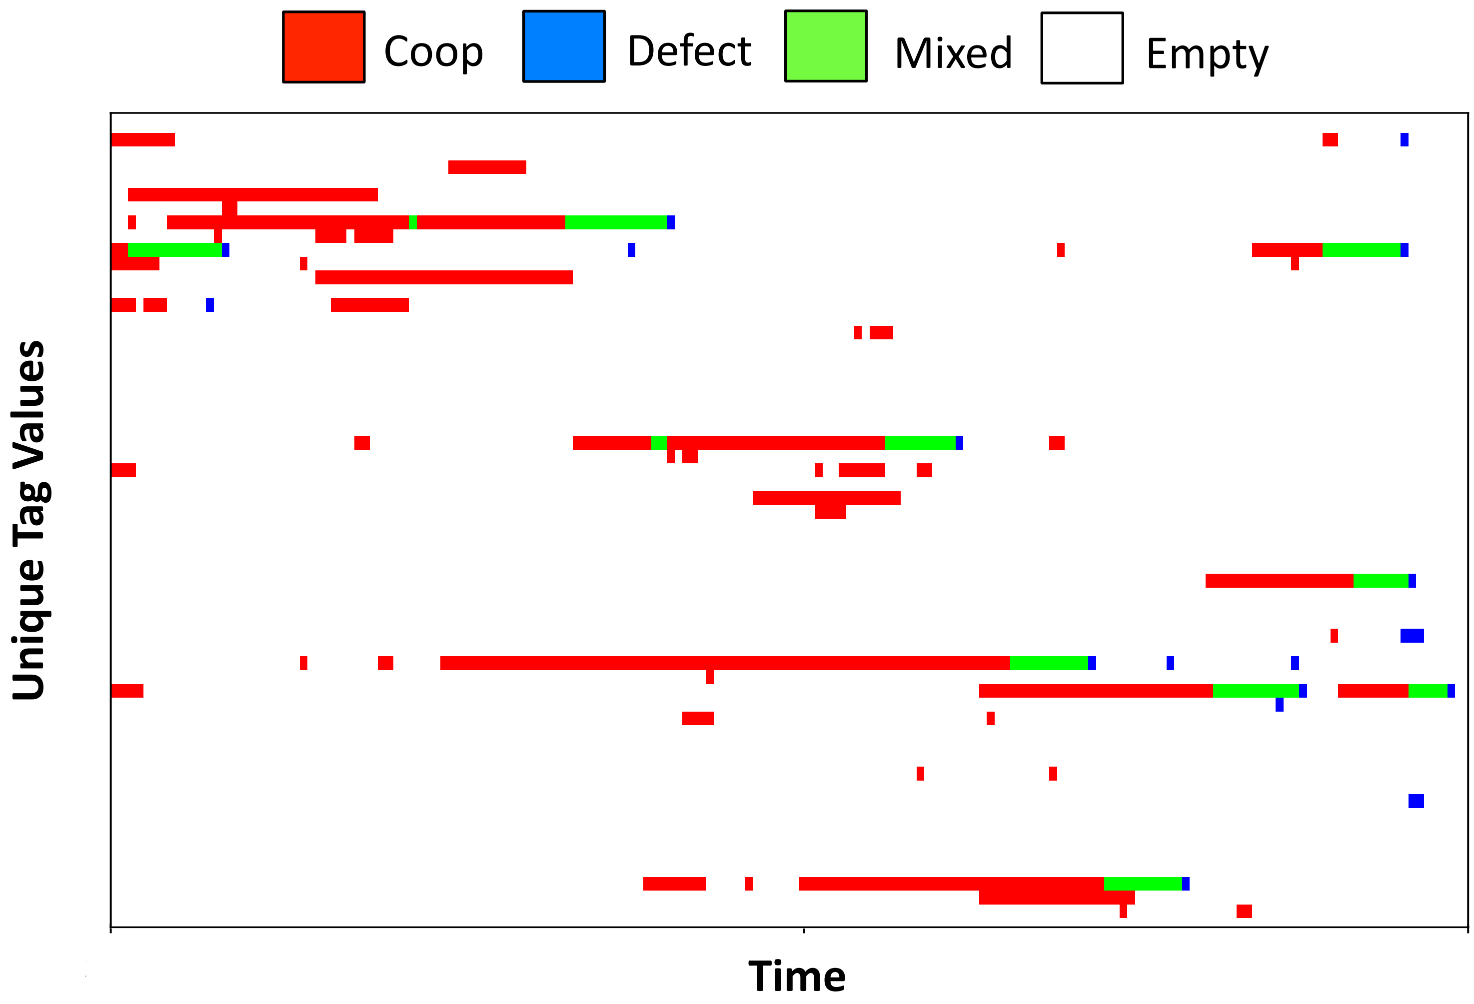

How tags work

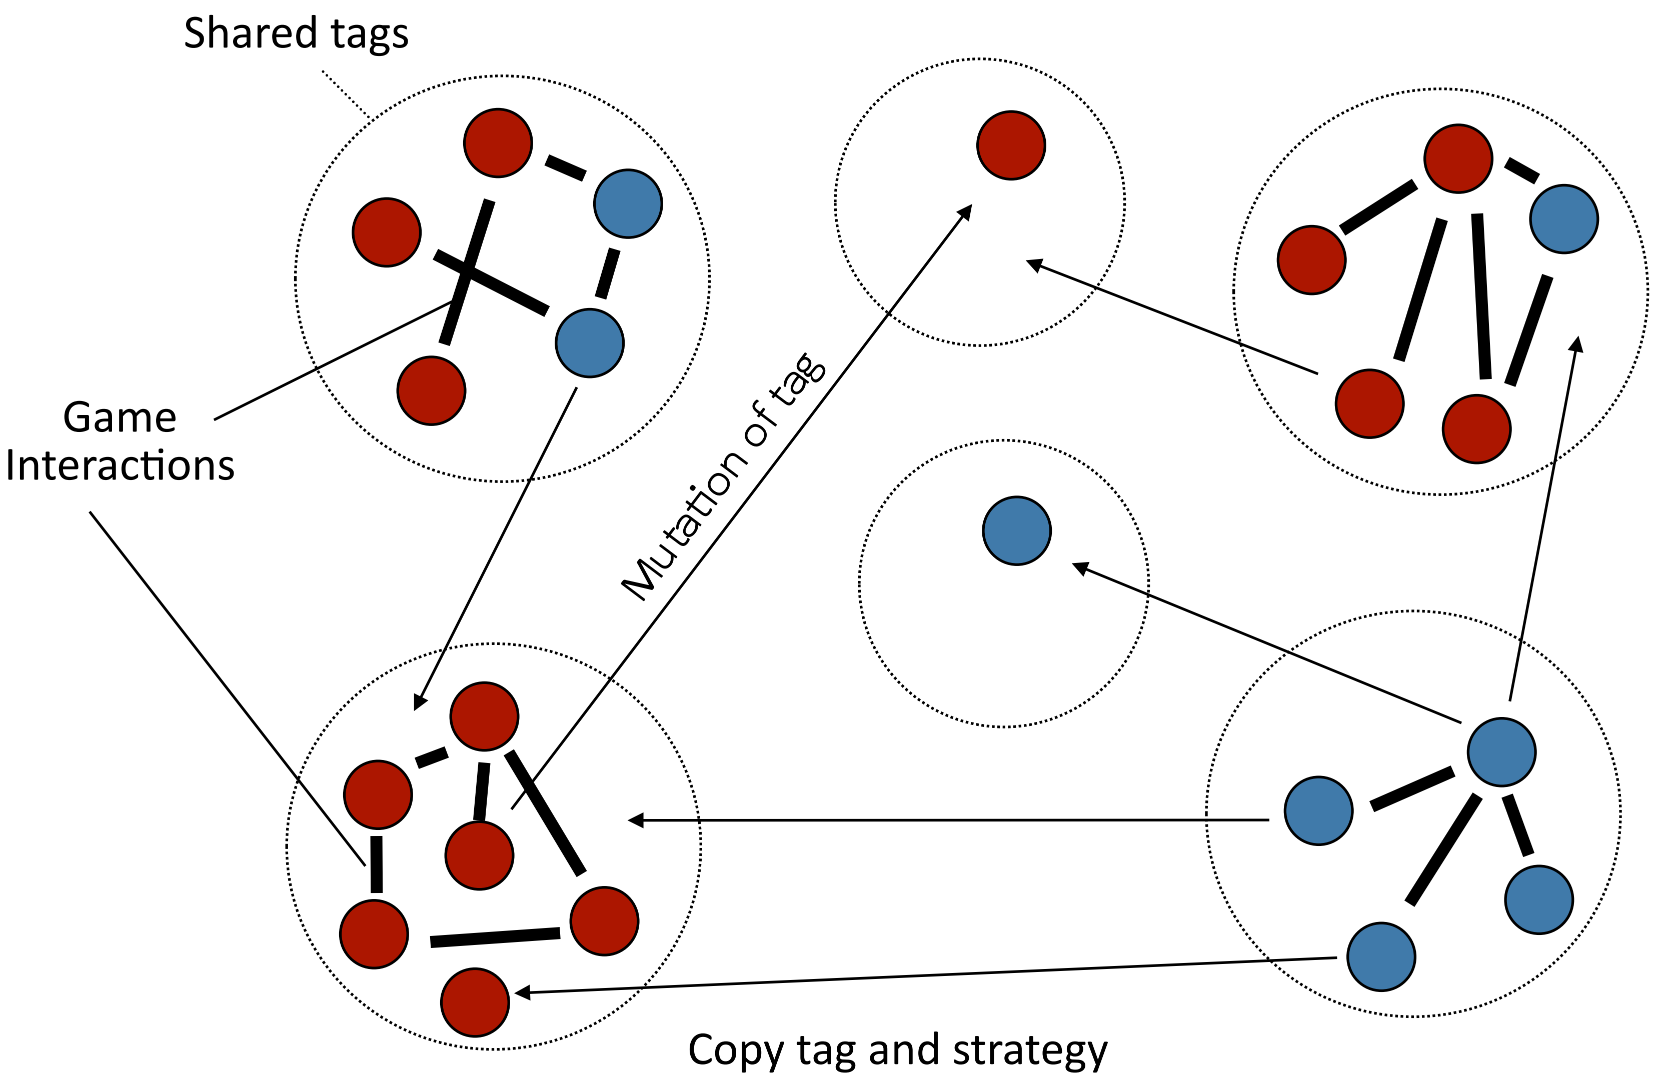

Tag models variants

- There are many variants of tag models:
  - Tags may effect the strategy played rather than the interaction structure only
  - tags may be combined with spatial interaction structures
  - Tags may be combined with a tolerance value allowing for similar tags to match

Riolo, R. L., Cohen, M. D., Axelrod, R., (2001) Evolution of cooperation without reciprocity. Nature 414, 441–443.

Jansen, V. A. A., Baalen, M., (2006) Altruism through beard chromodynamics. Nature 440, 663–666.

Traulsen, A., Nowak, M. A., (2007) Chromodynamics of cooperation in finite populations. PLoS ONE 2 (3), e270.

Network rewiring model

- Where the interaction structure is a network that evolves

Evolving interaction networks

- By evolving a network interaction structure in a similar way to tags..
- High levels of cooperation can emerge in dynamic evolving networks
- This equates to network rewiring
- Where nodes copy both strategies and links from each other

Hales, D. (2004) From Selfish Nodes to Cooperative Networks – Emergent Link-based Incentives in Peer-to-Peer Networks. The 4th IEEE Int. Conf. on Peer-to-Peer Computing. IEEE Computer Society Press

Santos, F. C., Pacheco, J. M., Lenaerts, T. (2006) Cooperation prevails when individuals adjust their social ties. PLoS. Comput. Biol. 2, 1284–1290

Network rewire model

Each node periodically (interaction):

plays PD with each of its network neighbours

Each node *p* periodically (reproduction):

q = select a random node

IF $fitness_q$ > $fitness_p$ (where fitness = average payoff)

drop each link with probability d

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

mutate (with low probability) strategy and links*

*Mutate links = drop each link with probability d and connect to a randomly selected node

Network rewire model

- ***What happens if we run this model?***

rewire model observations

- For a broad range of PD payoffs:
  - High C > 0.9 emerges from *any* initial starting condition
  - Network rewires into a dynamic highly clustered topology
  - When d = 1 clusters are disconnected components
  - When d < 1 network forms a small-world topology
  - Similar process to tags

Iterated PD

- In which the same two agents play repeated games of the PD with each other

Iterated variant of PD game

- There are many models that explore different variants of PD game and strategies
- We only considered two player, single round PD (sometimes called a one shot PD)
- In which agents played either C or D
- Another variant involves agents playing the same opponents repeatedly (the so-called Iterated PD)

Iterated PD strategies

- If the same two agents play repeated games of PD with each other (called Iterated PD or IPD)
- This allows for more complex strategies than only cooperate and defect
- Because agents can select their next move (C or D) based on previous moves of their opponent
- For example an agent could cooperate in the next round if their opponent cooperated in the last round

Iterated PD strategies

- In a famous book Robert Axelrod (a political scientist) describes "tournaments" in which he asked people to submit agents (as computer programs) that played the IPD
- He played them off against each other to see which did best (highest payoff)
- More details can be found on these slides:
  - www.davidhales.com/msiis/szeged-course4f-2015-coop.pptx.pdf

Robert Axelrod (1984) Evolution of Cooperation. Basic Books