Evolving cooperation in the non-iterated prisoners dilemma: A distributed, Small World Network inspired approach.

Nicolas Höning

June 6th, 2008

vrije Universiteit amsterdam

The Problem

In complex and dynamic markets, most of the opponents have never interacted before and it is hard to estimate for them wether they can trust each other.

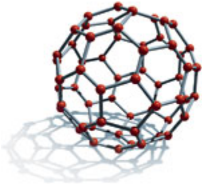

We propose that agents can ask other agents they already trust if the current opponent can be trusted.

vrije Universiteit amsterdam

The (iterated) Prisoner's Dilemma

- A non-zero-sum game with two players
- **C**ooperate (long-term) or **D**efect (short-term benefit) ?

	A	B	payoff: A	payoff: overall
traitor benefit (t)	D	C	5	5
reward (r)	C	C	3	6
punishment (p)	D	D	1	2
sucker payoff (s)	C	D	0	5

In the non-iterated version, there is no guarantee that the players will ever see each other again. How can cooperation evolve nonetheless?

Small World Networks

A recent observation about the connectedness of our world ...

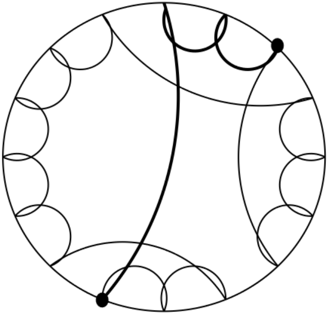

- Short average path lengths (like random graphs)
- High clustering index (# of direct neighbours higher than random graphs)

vrije Universiteit amsterdam

2. Framework

World Model

General theme: If cooperation would be good in the long run (PD), and if everyone is pretty well-connected (SW) - then gossip might be informative for cooperators.

- No central authority
- Agents are arranged on a likeness metric
- They play the Prisoner's Dilemma with random opponents
- Agents remember how opponents acted upon them in a trust table
- Evolutionary updates replace poor agents with copies of richer ones

Managing Trust

- Every agent (B) can be asked by some other agent (A) about the trustability of a particular opponent (C).
- If agent B doesn't have information on agent C in his trust table, then he will in turn ask another agent (D).
- From the trust table of B, agent D is selected by maximising the formula

$$distance\_bias * dist(D, C) + trust\_bias * trust(B, D)$$

- When agents report back along the chain, the trust values are multiplied to compute the trust in the chain.

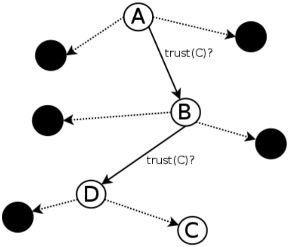

Being Nice

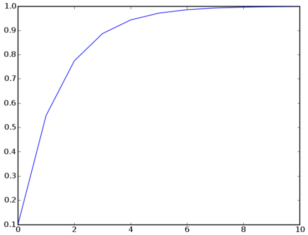

Figure: trust of 0.1, followed by positive interactions

- In the early 1980s, Axelrod held tournaments of competing strategies, from which the simple Tit-For-Tat strategy emerged as the most successful.
- Our strategy is also 'nice' in general, but will retaliate defection very soon.

vrije Universiteit amsterdam

Building A Network

- The default way of building a Small World Network (Watts and Strogatz, 1998) structure is to mostly connect neighbouring agents, and sometimes (according to some probability) remote agents.
- In this model, we follow a more precise approach by Kleinberg to evolve it using the distance metric.
- This metric works in the sense that agent A should have a higher probability to play against or know agent B than agent C if the distance $d(A, B)$ is lower than $d(A, C)$. Kleinberg calculated that probability with $\frac{1}{d(A,B)^{\alpha}}$ (where α is around 2).

Independent Variables

variable	possible values
N: # of individuals	[150, 300]
K: # of neighbours	$[0.7 * logN * \sqrt{logN}$, $1.1 * logN * \sqrt{logN}]$
α: network	[0.9, 1.2]
β: network	[1.9, 2.4]
strategy scenarios	$[(\frac{1}{2}$AD, $\frac{1}{2}$nuts), ($\frac{1}{3}$AC, $\frac{1}{3}$AD, $\frac{1}{3}$nuts), ($\frac{1}{3}$AD, $\frac{1}{3}$nuts, $\frac{1}{3}$nuts_dummy)]
bias to distance when asking	[0.3, 0.7]

In effect, we had $2 * 2 * 2 * 2 * 3 * 2 = 96$ different settings. *nuts* is our stratetgy.

vrije Universiteit amsterdam

Dependent Variables

variable	comment
L	average path length
NP	# of not connected pairs of agents
C	clustering index of network
SN	# of agents of particular strategies
SP	average payoff of the strategy

Hypothesis 1: The underlying structure will evolve to a small-world network.

Holds. In every simulation we ran, the network parameters achieved acceptable values. Small World Networks have $L \simeq L_{random}$, but $C >> C_{random}$. The table below shows the network parameters at the end, averaged over all our simulations.

	L	L_{random}	C	C_{random}
$N = 150, K = 13$	2.46	2.19	0.30	0.08
$N = 150, K = 21$	2.03	1.89	0.29	0.13
$N = 300, K = 16$	2.62	2.33	0.29	0.05
$N = 300, K = 25$	2.20	2.01	0.26	0.08

Hypothesis 2: Informed strategies will survive in various populations.

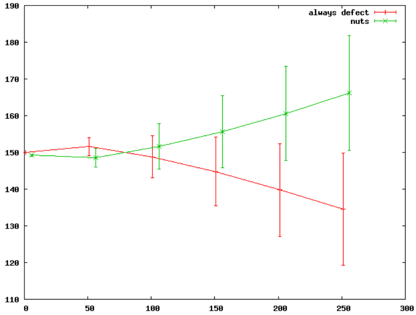

(a) N = 150

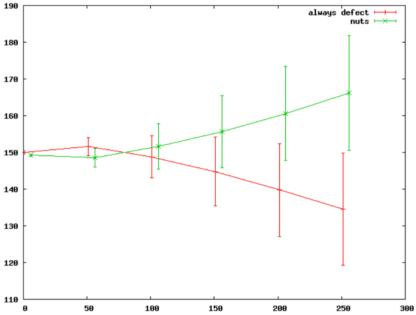

(b) N = 300

AlwaysDefect vs *nuts* in scenario 1 (AD, nuts)

vrije Universiteit amsterdam

Hypothesis 2: Informed strategies will survive in various populations.

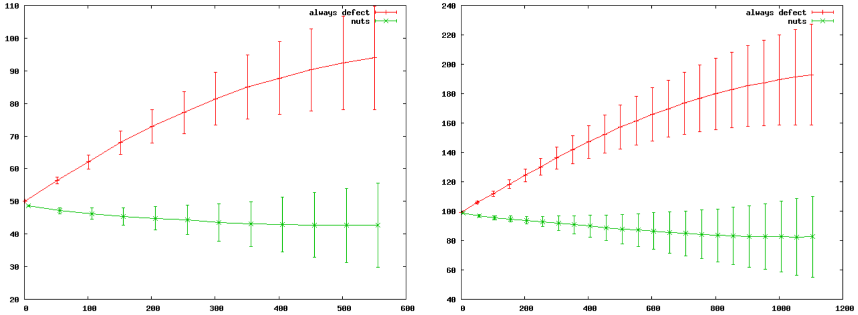

(c) N = 150 (d) N = 300

AlwaysDefect vs *nuts* in scenario 2 (AC, AD, nuts)

Hypothesis 2: Informed strategies will survive in various populations.

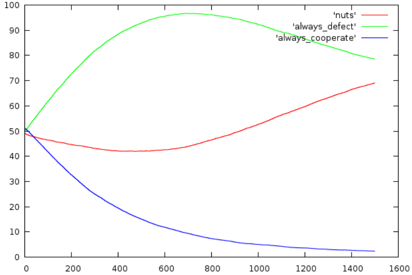

(e) N = 150

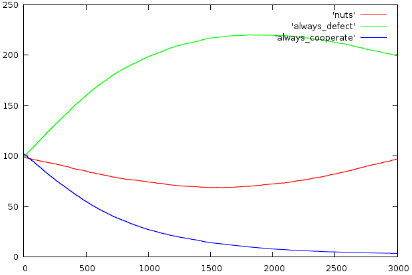

(f) N = 300

AlwaysDefect vs *nuts* in scenario 2 (AC, AD, nuts) with 2.5 times longer simulations.

Hypothesis 3: Information chains increase the success of an informed strategy.

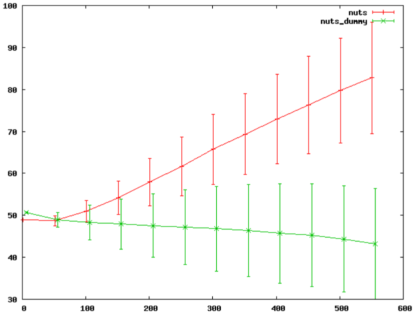

(g) N = 150

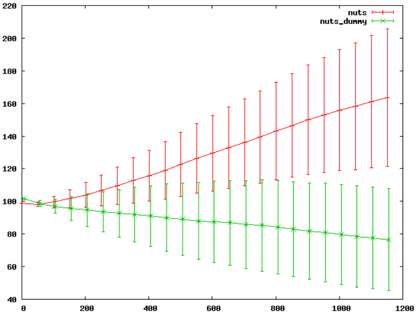

(h) N = 300

nuts vs *nuts_dummy* in scenario 3 (AD, nuts, nuts_dummy)

Outlook

- Evaluate Hypothesis 2 on the long run
- Develop collusion strategy (and beat them?)
- Develop smart defector strategy that cooperates when opponent is very likely to have good connections