

Interaction, Uncertainty, and the Evolution of Complexity.

Anil K Seth

Centre for Computational Neuroscience and Robotics
Innovations Centre, University of Sussex
Brighton BN1 9QG, UK
anils@cogs.sussex.ac.uk

*'a starlit or a moonlit dome distains,
all that man is,
all mere complexities,
the fury and the mire of human veins.'*

W.B. Yeats

Abstract

The evolution of complexity is investigated in the context of an 'iterated prisoner's dilemma' (IPD) co-evolutionary/game-theoretic ecology, populated by strategies determined by variable length genotypes. New evidence is found to support the dual hypotheses that both uncertainty, and interaction (by way of population stability), foster the evolution of progressively more complex entities. It is also argued that during periods of major evolutionary upheaval, complex entities suffer disproportionately and become less abundant in the population. The research is presented as an elaboration of the general principle that there is complexity in an organism by virtue of complexity in the environment, and has implications both for deepening understanding of the nature of biological evolution and for guiding the progress of artificial evolution.

1 Introduction

As Stephen Jay Gould [6] has consistently pointed out, the age of bacteria is not about to end anytime soon. Yet it can hardly be denied that over the course of biological evolution the complexity of the *most* complex things around has increased dramatically. It can indeed be said that in the earliest stages of life, there was nothing like the great variety of complex and wonderful creatures that now grace our world. Somewhere along the line complexity has evolved - not monotonically (witness the extinction of the dinosaurs), but it certainly has happened, and the mystery is *why*.

Two hypotheses are defended here. First, that *uncertainty* is a major drive for the evolution of complexity. The less predictable things are (up to a point), the more

complex an organism needs to be in order to behave effectively. Thus complexity affords robustness in the face of noise, and as such, may be expected to evolve in noisy environments. This line of argument appeals to the general principle that there is complexity in the organism by virtue of complexity in the environment (see [1], for example); and noise, up to a point, makes an environment more complex.

The second hypothesis is that *social interaction* is a major drive for the evolution of complexity. Most explanations of this effect would follow the above general principle in arguing that interaction augments environmental complexity, and subsequently agent complexity. However the argument here is that socialisation confers population *stability* which, in turn, promotes complexity. These two arguments are not mutually exclusive and may well be complementary - it is just that the model deployed in this study is designed to investigate the latter. By drawing this distinction, this study begins to focus on more *mechanistic* explanations for the evolution of complexity, and may also begin to answer questions such as why some organisms are more complex than others.

These hypotheses are investigated through the analysis of the dynamics of co-evolutionary artificial ecologies developed to allow their constituents to evolve from being simple to being complex. These ecologies were constructed on the basis of the 'iterated prisoner's dilemma', with complexity and simplicity being charted by the 'memory' of the strategies deployed by the constituent agents to play this game. The model developed is the first to incorporate both variable length genotypes *and* partner choice mechanisms.

Following the evaluation of these two hypotheses, it was further observed that in unstable (not just uncertain) circumstances, complex entities are more fragile than simple ones, and become less predominant in a population. The description of this (not entirely surprising) behaviour was not an explicit goal of the research, but some *a posteriori* interpretations are presented in section 6.2.

The remainder of this paper is structured as follows.

Section 2 discusses the theoretical motivation of the research. Sections 3 and 4 examine the model and implementation employed, and section 5 presents sample results, which are discussed in section 6. Section 7 presents proposals for further research.

2 Theoretical Background

2.1 What is complexity?

The term ‘complexity’ is abundantly used, in spite of the absence of a comprehensive and concise definition. Indeed the definitions employed by most researchers simply appear to reflect the particular task or model involved. For example, for the biologist Bonner, complexity is simply the number of different cell types in an organism [4], but the philosopher Shea defines complexity in much more abstract terms as ‘the number of parts and the irregularity of their arrangement’ [17].

One problem is that complexity only makes intuitive sense when considered in relative terms. Therefore, to enjoy any kind of intuitively valid quantitative basis, it is best to restrict the application of the term to representations within *finitely presented languages* (such as the genotypes used in genetic algorithms, see [8]). This outlaws, for our purposes, the more expansive definitions prevalent in the philosophical literature.

A standard measure of complexity, applicable to finitely presented languages, is that of ‘minimum description size’; otherwise known as ‘Kolmogorov complexity’. This definition suggests that the complexity of a given expression is determined by its minimum length following compression, (still retaining all the original information). This certainly allows that mere duplications of expressions do not augment complexity, but it also implies that a purely random expression will necessarily be maximally ‘complex’, by virtue of having maximal information content. This is not ideal, as it precludes the possibility of expressions displaying internal structure being classed as more complex than random expressions. And the relationship of complexity to a particular language is just as important as the relationship to an observer - different languages may express information and be compressible in many different ways.

Thus there is no universally accepted definition of complexity, nor any single measure appropriate to evolutionary investigations. In the present project, we take the *memory* of an agent (as indicated by genotype length) as a measure of complexity, being quite aware that it is a definition not immediately applicable to other domains of inquiry. The justification for the use of this metric is that the evolution of longer memories implies (by virtue of the structure of the model) that more complex behavioural strategies are being deployed within the population. Dues are nevertheless paid to Kolmogorov complexity by employing the Lempel-Ziv compression al-

gorithm [10] to determine the information content of the genotypes. For the reasons given above this measure is no more valid than a direct measure of memory, but is included for the sakes of completeness and comparison.

2.2 Existing theoretical explanations for the evolution of complexity

It has often been assumed that neo-Darwinism provides a direct and satisfactory explanation for the evolution of complexity. But taken at face value, neo-Darwinism explains only how organisms come to be increasingly adapted to their particular environments or niches. Certainly, if an environment encapsulates an escalating ‘predator-prey’ situation then Darwinism can predict the evolution of more complex capacities by way of ‘evolutionary arms races’ (see e.g. [5]), but it is not clear that this kind of explanation can be extended to all instances of complexity evolution¹. And just taking on board the general principle that organisms become complex through adapting to complex environments still leaves us in the dark over many questions. For example, why some things are more complex than others, what kinds of environmental complexities promote evolution, and how it is that organism complexity is vastly different from environmental complexity. More direct and satisfactory explanations are clearly required.

Bonner [4] describes a theory which argues that the evolution of ‘gross phenotypical complexity’ (such as the development of legs or lungs) can be attributed simply to organisms getting physically larger in order to exploit more areas of ecological niche space. Being larger requires being more internally complex in order to be self-sustaining, and canalization ensures that evolved internal complexity is retained should physical size reduce again over evolutionary time. More generally, it can be argued that greater structural complexity permits evolution to explore larger volumes of functional space, or even simply that, starting out simple, there is ‘nowhere to go but up’, and variation will more likely produce greater complexity than simplicity (but this has little to say about matters beyond the very simplest of organisms).

Other theorists have concentrated on trying to explain the evolution of neural/cognitive complexity, where physical size changes are not so important for their physical impact on the environment, as for their impact on extending the behavioural possibilities of the agent. A different class of theory has been developed to tackle this latter problem, many examples of which argue for various relationships between socialisation, language, and brain size, and as yet without any clear victor or objective methodology² (see e.g. [12]).

¹Indeed, not all predator-prey situations *do* lead to arms races - for example, algae have hardly evolved sophisticated ‘fish-avoidance mechanisms’.

²Ornstein, however, bucks the trend by suggesting that brain

The position adopted in the present paper is that a number of drives towards complexity undoubtedly existed, and may still exist. And since complexity is defined here in terms of ‘behavioural’ game-playing strategy, the present investigation falls more naturally into the context of neural/cognitive complexity rather than gross phenotypic complexity. However, since we are considering only general principles (neither the evolution of lungs nor the evolution of language), this is more an intuitive guide than a strict restraint. The dual hypotheses of noise and social stability as evolutionary motivational forces are therefore not to be considered as an exhaustive exploration of the subject, rather, as contributions to the aim of elaborating *just how* there is complexity in the organism in virtue of complexity in the environment.

2.3 Noise for robustness versus noise for complexity

Much emphasis has recently been placed on the importance of taking account of noise when evolving artificial systems in simulation - both for purposes of engendering real world robustness [15], and for disguising simulation inaccuracy [9]. The possibility of noise acting as a *driving force* for the evolution of complexity has, however, received much less attention, even though the concepts are closely related.

Lindgren [13] has used variable length genotype genetic algorithms ([8],[7]), to instantiate ecologies of agents interacting through the prisoner’s dilemma. He demonstrated the evolution, in noisy environments, of increasingly complex agents, but he didn’t rigorously address the possibility of a causal role for noise in the evolution of complexity, beyond suggesting:

...noise may disturb the actions performed by the players, which makes the problem of the optimal strategy more complicated. This increases the potential for having long transients showing evolutionary behaviour. [13], p.296.

Indeed, noise can be considered a good candidate for an evolutionary motivational force for the simple reason that uncertain environments are likely to punish very simple agents that rely on stability and regularity. Thus, if greater complexity can confer robustness against noise, we may expect to see evolution of complexity in noisy environments. In terms of our general principle, uncertainty could be considered as an important aspect of environmental complexity that leads to a corresponding complexity within the organism.

complexity evolved as a side-effect of some kind of neural air-conditioning mechanism during early hominid evolution in the hot African savannahs [16].

2.4 Partner selection and complexity

A general principle explanation would again have us believe that the inception of social interaction augments the complexity of the environments for the interacting agents, which then becomes reflected in the agents themselves. But this explanation, whilst conceptually attractive, is hard to assess in the context of the present paper because the actual mechanism for effecting social interaction is supplied, ready-made, to the agents³. Alternatively, it can be argued firstly that social interactions act to stabilise a population, by virtue of agents being able to choose suitable partners and refuse overtures from unsuitable agents. Then, secondly, echoing May [14], that this stability permits the evolution of complexity and an increase in population diversity. In this scenario, the very fact that non-trivial agent-agent interactions are happening at all provides the stimulus for evolution (with noise abetting this inherent tendency), and socially conferred stability provides a platform from which this evolution can really take off. Here we see how there might be mechanisms whereby different organisms may become differentially complex in similar environments, thanks to feedback mechanisms operating via processes such as social interaction.

3 Structure of the Model

3.1 The prisoner’s dilemma

The prisoner’s dilemma has long been established as a tool of great value in co-evolutionary investigations, [3],[11],[14],[20]. Essentially, it provides a framework for modelling non-trivial interactions between agents, where the maximisation of individual short term gain minimises the collective welfare, as illustrated by the following anecdote:

Imagine that you and an alleged accomplice have both been arrested, accused of a heinous crime. You are held in separate cells, and upon interrogation you can either *cooperate* by denying all knowledge, or *defect* by implicating your accomplice. You have no idea what your accomplice will do, but if you both cooperate, you will both be released (the reward, **R**), and if you both defect, then both of you will be jailed (the punishment, **P**). However, if you defect and she cooperates, then you will receive a payoff (the temptation, **T**) and she will go to jail for longer (the sucker, **S**). But if she defects and you cooperate, then you yourself are the sucker. The paradox is thus evident, - in a single meeting you will always do best to defect, in doing so either receiving the monetary payoff or avoiding being the sucker. But of course

³A worthy topic for future research would be to look at how ‘social mechanisms’ do evolve in social situations - the expanding literature on the evolution of communication is a useful step in this direction.

	<i>player 2 cooperates</i>	<i>player 2 defects</i>
<i>player 1 cooperates</i>	1:R=3 2:R=3	1:S=0 2:T=5
<i>player 1 defects</i>	1:T=5 2:S=0	1:P=1 2:P=1

Table 1: Prisoner’s Dilemma Scoring Table

the logic is the same for your alleged accomplice, and if you both defect then you will both do worse than if you had both cooperated (see table 1). Note that the actual scores don’t really matter so long as $T > R > P > S$ and $2R > T + S$.

Cooperation is thus unlikely to arise in a one-shot prisoner’s dilemma, but if players can meet time and time again, and retain some memory of previous interactions, then cooperation on any given move does become a rational strategy. It is this ‘iterated prisoner’s dilemma’ (IPD) that forms the core of the present study.

Many researchers have used genetic algorithms to evolve strategies to play the IPD (see e.g. [3],[11]), primarily for the purposes of investigating the evolution of cooperation. In these studies, as in the present model, the genotypes comprise of binary character strings representing policies for playing the IPD, with the length of the genotype determining the number of preceding moves (the game history) upon which each individual can base its strategy. Genotype length thus has a direct and valuable interpretation as ‘memory’. It has been repeatedly demonstrated that cooperative strategies can and do arise and persist in artificial ecologies populated by these evolving strategies, [3]. In the present study, the evolution of cooperation *per se* is not of primary interest. Rather, cooperation within a population is taken as a useful metric of population *stability*, and the evolution of cooperating population provides a good platform for the subsequent investigation of the evolution of complexity.

3.2 Variable length genotypes and the iterated prisoner’s dilemma

By introducing variable length genotypes (VLGs) into an IPD situation, Lindgren [13] demonstrated the evolution of longer and increasingly complex strategies (with longer memories) in noisy environments, (noise was introduced by invoking a certain probability for the opposite move to that specified by the genotype being made). In addition to the usual crossover and point mutations, Lindgren’s strategies could also undergo *splitting* and *doubling* mutations, which incremented or decremented the potential memory (and hence potential complexity) of the strategy in question by one game iteration.

The present study employs splitting and doubling mutations in the same way as Lindgren in order to provide a quantitative method for following the evolution of

complexity⁴. This method is particularly attractive because the phenotypical strategy is not directly affected by a doubling mutation. By itself, an increase in memory doesn’t change behaviour. Changes will only occur if the extra memory is subsequently used (through further mutations/crossover in the new genotype segment) to discriminate between possible courses of action (cooperation and defection in the current context). Thus, evolution of longer memories cannot simply be ascribed to some phenotypical side-effect of having a longer genotype, and must be attributed to some employment of the extra potential (or memory) provided; thus a more complex strategy.

3.3 Preferential partner selection

In real life, we normally exercise considerable discretion about who we interact with, rather than being forced to interact with everyone. This principle of partner selection allows new dimensions of emergent behavioural structure to develop, and has been introduced into a number of IPD models.

Stanley et al. ([19],[2],[18]) have published a series of papers looking at the formation of ‘social networks’ in an IPD context with choice and refusal permitted. Agents choose and refuse with reference to continuously updated expected payoffs that each agent maintains for every other agent in the population (they call this an IPD/CR mechanism). They demonstrate that cooperation is evolved rapidly under these conditions, and they discuss the emergence of a variety of metastable networks displaying distinct patterns of cooperativity.

The present study employs an IPD/CR mechanism similar to that developed by Stanley et al., but the innovation finds an original application in investigating the evolution of complexity, through the changes marked out by variable length genotypes.

4 Implementation

4.1 Genotype encoding scheme

At the heart of each individual in the ecology is a genotype, consisting of a string of *c*’s and *d*’s, determining the strategy of the individual for playing the IPD. The longer the genotype, the more it can be influenced by the history of the game, thus the longer the ‘memory’ of the individual.

Fig 1 illustrates how the genotype can code for a particular strategy. Each time a previous move in the game history (between two particular agents) is considered,

⁴VLGs have been considered to be very important in the evolution of complex entities by a number of researchers (see e.g. the SAGA genetic algorithm of Harvey [7]). This affirms the principle that in ALife GAs are not employed to solve a particular problem, but instead provide a substrate for open-ended evolution.

half of the genotype is (temporarily) thrown away (the non-shaded area in fig 1) - one half if the move had been cooperative, or the other if it had been a defection. In this way, the genotype in fig 1 (of length 16) can therefore encode a strategy with a memory of 4 prior interactions (after cutting a string of 16 characters in half 4 times, you are left with just a single character). The black square in fig 1 indicates which allele would be accessed for a [c,d,c,d] history.

However, the genotype must be made longer in order to specify the initial moves up until the memory limit. The genotype in fig 1 would require an extra 9 alleles to code for the initial 3 moves before the final 16 alleles can be used. The maximum genotype length employed is 127 alleles, providing a maximum memory of 6 iterations, although extensions to maximum memories of *any* integer value could be easily implemented.

Each time an iteration of the game is played, the moves made are stored in a history array so that they can be accessed by the strategies the next time the two players meet. This therefore constitutes the bones of the model; the agents interact by way of the IPD and their strategies can be influenced to different (and evolvable degrees) by what has happened in the past.

4.2 Code structure overview

The flow of the program is essentially very simple, as illustrated by the following pseudocode:

```

randomly initialise population
FOR EACH generation
  FOR EACH iteration
    if choice & refusal is enabled
      members choose and refuse partners
    else
      everyone chooses everyone (except self)
    one round of IPD is played (with or
                                without noise)
    history array and various scores are
                                updated
  ENDFOREACH
  new generation is created through breeding
  using a tournament GA
  every so often population statistics are
  calculated, and presented with graphics
ENDFOREACH

```

Noise is implemented in the same way as in the models of Lindgren [13]. That is, a certain probability (in this study, 1 percent)⁵ is set with which the actual move made during an IPD interaction is opposite to that specified by the genotype.

⁵Preliminary experiments indicated that with much greater noise levels (above 5 percent), population cooperation evolved far too infrequently to sustain further study.

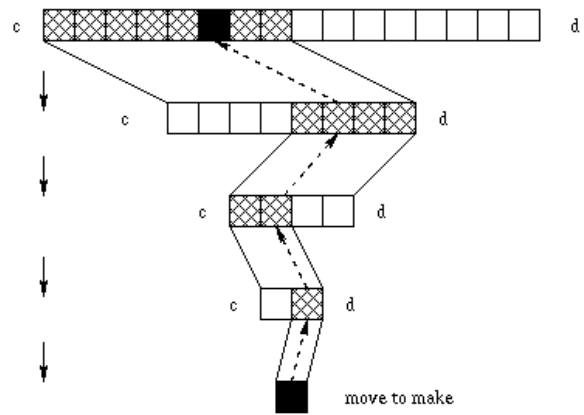


Figure 1: *genotype encoding scheme.*

4.3 Breeding procedures

A simple generational, tournament style GA is used, and the first part of the breeding procedure is to sort the population in order of fitness, which is a simple reflection of average score over all the IPD games played. If required, these fitnesses are adjusted so that longer genotypes, all else being equal, score less than short genotypes (a cost on complexity - see section 5.3).

Importantly, following crossover, splitting or duplication mutations can occur, with consequences for the memory of the new individuals. Duplication copies the latter half of the genotype twice again onto its own end, thereby preserving an appropriate section for initial move specification and also incrementing the memory by one. The actual strategy implemented by the genotype won't change after a memory increase - only further alterations through crossover or point mutation in the new genotype section can achieve this. Splitting the genotype reduces the memory by one, and of course, this mutation may well directly inflict a strategy change, if the discarded alleles were previously being used to dictate a more complex strategy utilising the full memory potential.

The crossover rate was set at 0.95, and all mutation rates were set at 0.005 (per bit for point mutations).

4.4 Preferential partner selection

If partner choice is allowed, then it proceeds on the basis that each member of the population maintains an 'expectation value' for every other member, reflecting the expected outcome of a round of prisoner's dilemma. This value, set initially to 3.0 (mutual cooperation) is allowed to alter on the basis of experience, as illustrated by the pseudocode overleaf, (based on [19]).

```

FOR EACH population member
  sort remainder on basis of expectation
  select the most preferable (up to a quota)
  and make offers; recipients thus form
  a list of offers
ENDFOREACH

FOR EACH population member
  FOR EACH offer received
    IF the offer is tolerable
      play one iteration and update
      scores and expectations
    ELSE
      refuse the offer (update expectation
      of refusee)
  ENDFOREACH
ENDFOREACH

FOR EACH member that has not played any games
  update fitness with 'wallflower' payoff

```

Given a standard IPD payoff matrix (as in table 1), the extra payoffs employed here are usually 1.0 for a refusal, and 1.6 for a ‘wallflower’. The tolerance limit, below which refusals will take place, is usually set to be the same as the wallflower payoff, thus 1.6.

The expectations are updated according to the following equation (based on [19]):

$$exp[i + 1] = \gamma \cdot exp[i] + (1 - \gamma) \cdot payoff$$

where ‘ γ ’ is a memory weight, and ‘*payoff*’ represents whatever payoff (IPD or otherwise) is awarded.

This mechanism is an example of reinforcement learning, and over the course of a generation, the population members will develop distinct expectations of others based on the outcomes from early iterations. This will guide their offers in the future, and allow networks of social interactions to develop.

4.5 Statistics

The progress of the model was charted with a number of statistical measures. Average score (on the IPD) was used to implement breeding. Average memory was the complexity metric, and average cooperativity (and hence performance) was taken to indicate population stability.

Population diversity (the average Hamming distance between each possible genotype pair) was also calculated, as was the average Kolmogorov complexity of the population. In the results section, however, memory is presented as the sole complexity metric, since Kolmogorov complexity was not calculated for all cases, and correlated very strongly with memory in those cases when it was known.

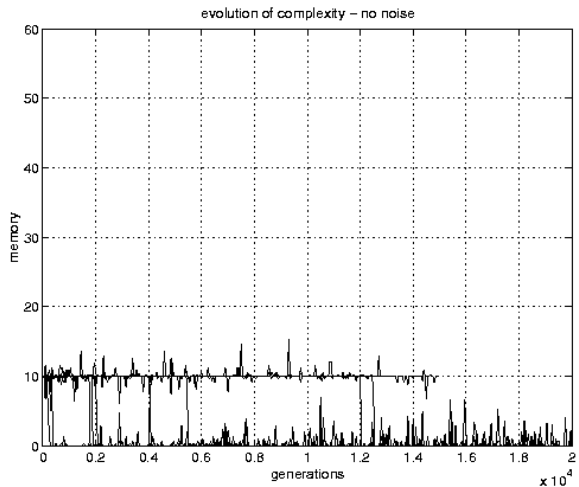


Figure 2: *evolution of complexity without noise.*

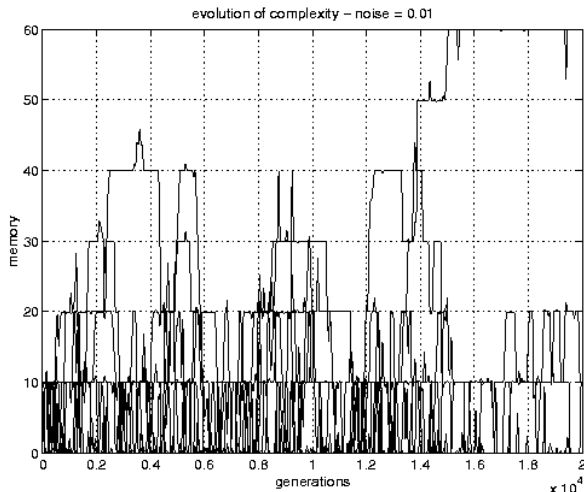


Figure 3: *evolution of complexity with noise.*

5 Results

5.1 Preliminary experiments

In general, the model was successful in evolving cooperative populations in non-noisy environments with or without partner choice, thus replicating the findings of Axelrod [3] and providing an adequate platform for further investigation. Further preliminary experiments looked at how strategies with *fixed* memories dealt with uncertain conditions, with results indicating that more complex strategies (of memory 6) were more robust and led to greater overall levels of cooperation than simple strategies (of memory 1). This provided some hope that complexity would therefore evolve in noisy environments.

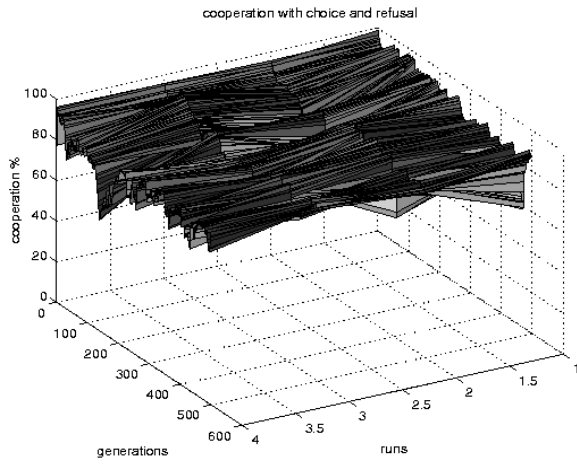


Figure 4: *cooperation with partner choice.*

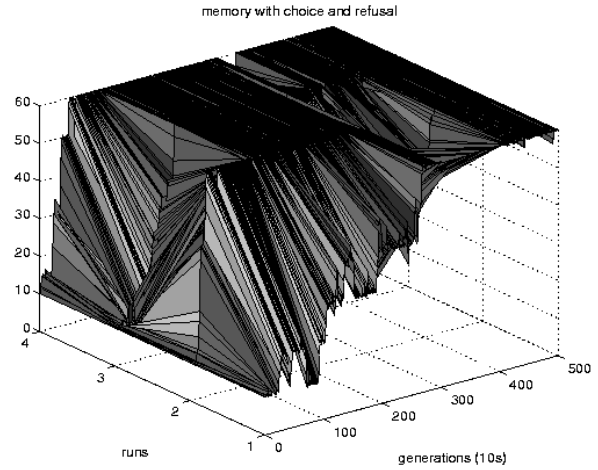


Figure 6: *evolution of complexity with partner choice.*

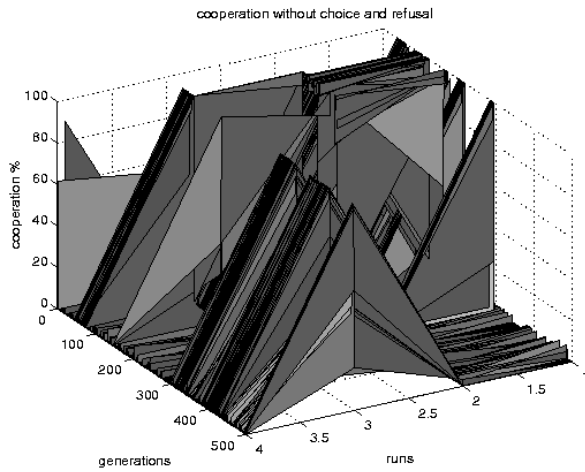


Figure 5: *cooperation without partner choice.*

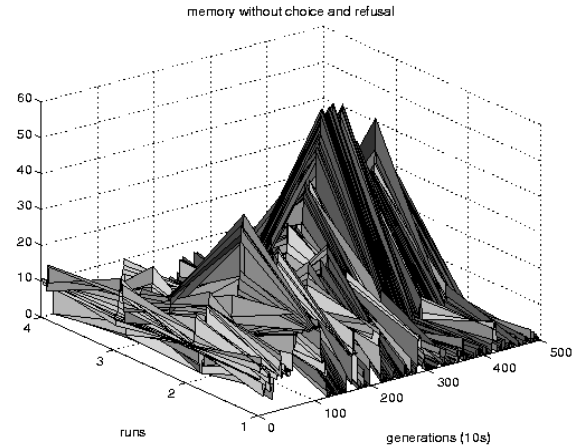


Figure 7: *evolution of complexity without partner choice.*

5.2 Results without partner choice; noise kick-starts complexity

Ecosystems were initially populated with 30 strategies of memory one, playing 20 rounds of IPD each generation (everyone playing everyone else), and were allowed to evolve over 20,000 generations. Fig 2 illustrates that, without noise, there is no tendency whatsoever for any kind of complexity to evolve, but fig 3 illustrates that, with noise, complexity does evolve, (12 runs are superimposed on these graphs). Not always, or even often, to the maximum. And not monotonically either. But the difference is clear; noise does promote complexity.

Also of interest is the observation that the evolving populations rarely sustain mixtures of complexities - this is particularly clear in fig 2 in which populations either completely comprise of memory 1 or memory 0 strate-

gies, but never both. And changes in population complexity occur very swiftly between stable, homogenously complex states - as can be seen in fig 3.

5.3 Results with partner choice; interaction promotes complexity

With the addition of the partner choice mechanism (and without noise), it is initially striking that the ecology evolves to full cooperation very quickly indeed, and is very stable over time, (the initial populations again consisted of 30 memory 1 strategies playing 20 IPD rounds each generation).

Fig 4 demonstrates that under partner choice, the level of cooperation (and population performance) achieved rises very quickly to near 100 percent and stays there. Without partner choice (fig 5), the evolved cooperation

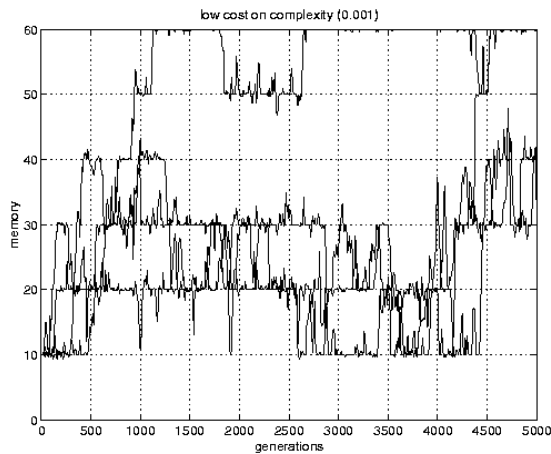


Figure 8: *a low cost on complexity.*

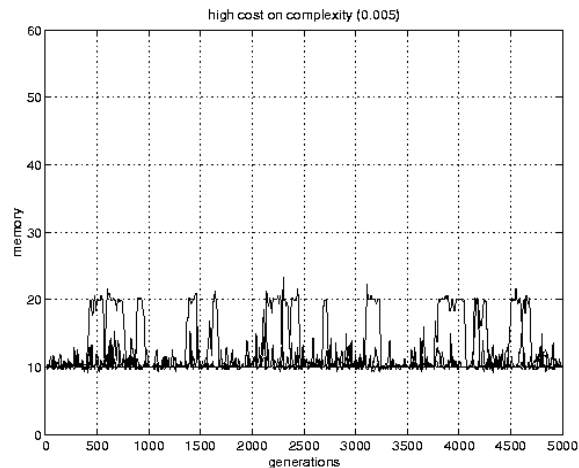


Figure 9: *a high cost on complexity.*

is much less stable and spends considerable time at zero. The case is similar for the evolution of complexity, as measured by memory. Fig 6 illustrates that, under partner choice, the initially simple memory 1 strategies evolve very quickly to a maximum memory of 6 and remain there. But without partner choice (fig 7), evolved complexity remains, almost without exception, at the initial level of 1 or lower. Thus, interaction (through partner choice) does promote stability and complexity.

5.4 A cost on complexity

Given that complexity evolves in an IPD/CR situation without any encouragement from noise, a series of experiments were performed with a *fitness cost* imposed on the possession of a long genotype. This fitness cost was compared over two values - high (0.5 percent) and low (0.1 percent). The values refer to the fitness percentage deducted per ‘memory unit’, and were kept low so that a long genotype which had scored well in the ecology would not score worse than a short genotype that performed poorly. It was only desired to assess the *ceteris paribus* effect of a cost on complexity.

Fig 9 overlays the results from 8 runs and illustrates that with a high cost, the evolution of complexity is abolished. With a low cost, (fig 8), this is not the case and, although complexity does not evolve with the rapidity and stability observed in the cost free cases, it certainly does evolve to a significant extent.

These experiments indicate that complexity is of sufficient value to an agent to offset a small incurred fitness cost, if not a large tariff. As a follow-up experiment, it was hypothesised that in ecologies where the ‘cost barrier’ precluded the evolution of complexity, the introduction of environmental noise might be able to kickstart the process and so allow complexity to develop after all.

A cost of 0.25 percent per memory unit was imposed,

and 12 runs were performed without noise, and 12 with 1 percent noise. Fig 10 and fig 11 clearly illustrate the beneficial influence of noise/uncertainty in the evolution of complexity.

5.5 Correlation between noise and complexity

A further interesting phenomenon observed in these experiments is the correlation observed between cooperation (or performance) and evolved complexity, (an effect that is only observable when both factors are sufficiently mobile). Two examples are presented here, fig 12 looking back at a non-partner-choice situation, with a rather high mutation rate (5 percent), and fig 13 looking at several superimposed runs with a high cost on complexity (0.5 percent) in a partner-choice situation. In both of these situations there are factors (high mutation, high cost) mitigating against long term stability of the population. Possible reasons for this correlation will be addressed below, in section 6.2.

6 Discussion

6.1 The evolution of complexity

It has been shown that both environmental uncertainty (noise) and social interaction can jointly promote the evolution of complexity.

In the case of noise, the suggestion is that since complexity confers robustness against noise (amongst other things), complex strategies will have a selective advantage over simple strategies in uncertain environments. The inherent uncertainty in the natural world may thus have been one of the catalysts for the evolution of the complexity of living organisms.

In the case of social interaction, the population stability evoked through partner choice is held to be in-




Figure 10: *evolution of complexity with noise, and a medium cost on complexity.*

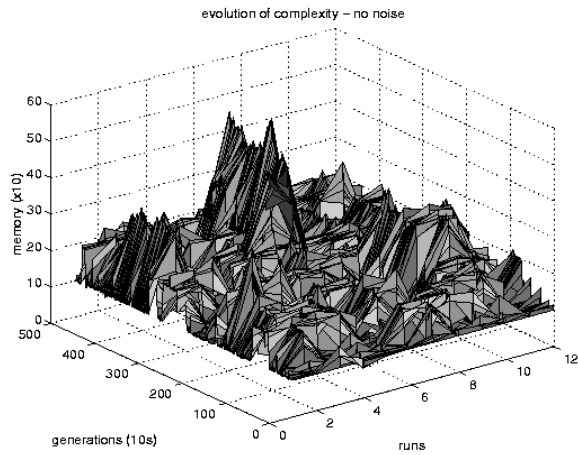


Figure 11: *evolution of complexity without noise, and a medium cost on complexity.*

strumental in allowing complexity to evolve. It is argued here that there is a pre-existing drive towards complexity based on the fact that the mode of interaction through the IPD is non-trivial and so more complex strategies have the potential to do better⁶. But this potential is best exploited in stable circumstances, and can be augmented by the presence of environmental noise. The evolution of complexity through social interaction is a process that may not be available to all evolving organisms, suggesting that differential complexity may exist in the world thanks to there being different drives towards complexity, not all of which are (or have been) available

⁶A good example would be the tit-for-tat strategy, which is highly stable in non-noisy environments (see [3]), but which breaks down into mutual defection very quickly with the introduction of noise.




Figure 12: *correlation between complexity and stability (high mutation).*




Figure 13: *correlation between complexity and stability (high cost).*

to all organisms.

These conclusions are presented as an elaboration on the overly general principle that ‘there is complexity in the organism in virtue of complexity in the environment’, (see [1]). This research begins to answer more specific questions about what *kinds* of complexity in the environment may promote organism complexity, and by what mechanisms. This has implications not only for understanding how complexity has evolved the way it has - why bacteria still flourish but now in a world also populated by human beings - but also for practioners of artificial evolution who wish to understand how to manipulate evolutionary forces in order to evolve artificial agents of increasing complexity, interest, and utility.

6.2 A correlation between stability and complexity

In unstable circumstances, complexity correlates strongly with population stability - an observation that makes good sense in the context of stability promoting complexity. The argument here is that although complexity delivers robustness against environmental noise, it incurs brittleness in the face of large scale ecological change (ie. population instability). This may be because with greater complexity, there is more potential to exploit in becoming appropriate for the ecological niche in question - even if that niche presents some uncertainty. However, should the nature of the niche change dramatically, the specialisations of the complex entity are more than likely to prove disadvantageous. No such problem for the simple entity, which will have to undergo a much smaller genotypic change in order to suit the new niche (to a less adapted extent).

It is suggested that this phenomenon can be discerned in natural history, where complex organisms have suffered disproportionately during the huge ecological upheavals of the extinction events⁷.

7 Future work

One obvious challenge would be to analyse the actual strategies that are deployed by the agents, to find out what kinds of complexities are being selected for. This would clearly deliver a deeper understanding of the nature and behaviour of the model.

As this project has been an exploratory study, the hypotheses presented would benefit from a more exhaustive testing and analysis process. A meta-study could combine the results across many different implementations, perhaps employing a variety of paradigms for founding the ecology (not just the IPD). This approach may also help in developing models that demonstrate evolved symbiosis between complex and simple organisms (bringing us back to our starting point - the age of bacteria is not over yet!).

A spatial element to the interaction of agents could be introduced, posing the question - would complexity evolve and thrive in isolated pockets before spreading across a population? And the mechanism for partner choice should be allowed to evolve as well, instead of being 'handed' to the agents in a ready-made form.

Acknowledgements

This work owes a great deal to my supervisors Phil Husbands and Inman Harvey, to my reviewers, and also to Tony White of Nortel Ltd. Financial support was provided by the EPSRC and by Nortel Ltd.

⁷The story told here is not going to be the only one. It is also undoubtedly important that the faster generational turnover of simple organisms abets their adaptation to changing ecologies.

References

- [1] W.R. Ashby. *Design for a brain: The origin of adaptive behaviour*. Chapman Hall, 1952.
- [2] D. Ashlock, M.D Smucker, E.A. Stanley, and L. Tesfatsion. Preferential partner selection in an evolutionary study of prisoner's dilemma. Economics report 35, Iowa State University, 1994.
- [3] R. Axelrod. *The Evolution of Cooperation*. New York : Basic Books, 1984.
- [4] J.T Bonner. *The Evolution of Complexity by means of natural selection*. Princeton University Press, 1988.
- [5] R. Dawkins. *The Blind Watchmaker*. Longman, 1986.
- [6] S.J. Gould. *Wonderful Life*. Penguin, 1989.
- [7] I.R. Harvey. *The Artificial Evolution of Adaptive Behaviour*. PhD thesis, Sussex University, 1995.
- [8] J.H. Holland. *Adaptation in Natural and Artificial Systems (2nd ed)*. Cambridge, MA : MIT Press, 1992.
- [9] N. Jakobi, P. Husbands, and I.R. Harvey. Noise and the reality gap: the use of simulation in evolutionary robotics. In F. Moran, A. Moreno, J.J. Merelo, and P. Chacon, editors, *Advances in artificial life; third european conference in artificial life*. Springer-Verlag, 1995.
- [10] F. Kasper and H. Schuster. Easily calculable measures for the complexity of spatio-temporal patterns. *Physical Review A*, 36(2):842-848, 1987.
- [11] C. Langton. *Artificial Life; an Overview*. MIT : Bradford Books, 1995.
- [12] R. Leakey. *The origin of humankind*. Weidenfeld and Nicholson, 1995.
- [13] K. Lindgren. Evolutionary phenomena in simple dynamics. In C. Langton, J.D. Farmer, S. Rasmussen, and C. Taylor, editors, *Artificial Life II*. Addison-Wesley, 1991.
- [14] R. May. *Stability and Complexity in model ecosystems*. Princeton University Press, Princeton, NJ, 1973.
- [15] J.M. Orbell, K. Runde, and H. Morikawa. The robustness of cognitively simple judgements in ecologies of ipd games. *Biosystems*, 37:81-97, 1996.
- [16] R. Ornstein. *The Evolution of Consciousness*. Touchstone, 1991.
- [17] M.C. Shea. Complexity and evolution: what everybody knows. *Biol. Phil.*, 6(3):303-324, 1991.
- [18] M.D. Smucker, E.A. Stanley, and D. Ashlock. Analyzing social network structures in the iterated prisoner's dilemma with choice and refusal. Technical report cs-tr-94-1259, University of Wisconsin-Madison, Dept of Computer Science, 1994.
- [19] E.A. Stanley, D. Ashlock D., and M.D. Smucker. Iterated prisoner's dilemma with choice and refusal of partners: Evolutionary results. In F. Moran, A. Moreno, J.J. Merelo, and P. Chacon, editors, *Advances in Artificial Life : Lecture Notes in Artificial Intelligence*. Springer-Verlag, 1995.
- [20] M. Taylor. *Anarchy and Cooperation*. Wiley Press, 1976.