Tracking The Red Queen: Measurements of adaptive progress in co-evolutionary simulations*

Dave Cliff[1] and Geoffrey F. Miller[2]

[1] School of Cognitive and Computing Sciences, University of Sussex
Brighton BN1 9QH, U.K.
Phone +44 1273 678754, Fax +44 1273 671320, Email davec@cogs.susx.ac.uk
[2] Department of Psychology, University of Nottingham,
Nottingham NG7 2RD, U.K.
Phone +44 115 9515364, Fax +44 115 9515324, Email gfm@psyc.nott.ac.uk

Abstract. Co-evolution can give rise to the "Red Queen effect", where interacting populations alter each other's fitness landscapes. The Red Queen effect significantly complicates any measurement of co-evolutionary progress, introducing *fitness ambiguities* where improvements in performance of co-evolved individuals can appear as a decline or stasis in the usual measures of evolutionary progress. Unfortunately, no appropriate measures of fitness given the Red Queen effect have been developed in artificial life, theoretical biology, population dynamics, or evolutionary genetics. We propose a set of appropriate performance measures based on both genetic and behavioral data, and illustrate their use in a simulation of co-evolution between genetically specified continuous-time noisy recurrent neural networks which generate pursuit and evasion behaviors in autonomous agents.

1 Introduction

Some biologists have suggested that the 'Red Queen effect' arising from co-evolutionary arms races has been a prime source of evolutionary innovations and adaptations [19, 5, 16]. The Red Queen was a living chess piece in Lewis Carroll's *Through the Looking Glass*, who ran perpetually without getting very far because the landscape kept up with her. Similarly, in co-evolution between predators and prey, hosts and parasites, males and females, or competitors within a species, traits in organisms evolve against traits in competitor organisms that are themselves evolving: each lineage's fitness landscape changes perpetually. Adaptive advantage is continually eroded under co-evolution.

Or so the theory goes. But does sustained competition really lead to smooth, directional evolutionary progress, or to noisy, unreliable, fits and starts, or to endless cycling through different evolutionarily unstable strategies? How important is tight co-evolution among two or a few competing lineages, versus diffuse co-evolution among many? These issues are critical to the debate between those

who view evolution as a smoothly running engine of adaptation (e.g. [5]), and other theorists who view it as a more contingent history of genetic drift, ad hoc modification, and developmental limitation (e.g. [9]). The Red Queen question is a microcosm of the ancient debate over the links between evolution, life, teleology, and progress.

Testing the significance of the Red Queen has proven difficult. The fossil record provides only ambiguous evidence of co-evolutionary progress [7], and fossils may not reveal the bodily and behavioral innovations that are important in most co-evolutionary scenarios. Simple population genetics models may over-estimate the smoothness of co-evolution by neglecting phylogenetic and developmental constraints that keep lineages stuck in local optima while their competitors surge ahead. Comparative studies across extant species reveal what adaptations exist, but not whether they were acquired through tight, synchronized co-evolution.

Evolutionary computer simulations are ideal for investigating co-evolution. They allow much more complex genotypes, phenotypes, behaviors, and interactions than population-genetic models or evolutionary game theory. And they allow researchers to make detailed measurements during and after co-evolution, revealing much more than could be inferred from fossils or comparative studies.

This paper focuses on developing measurement tools for such simulations. We are concerned with methods for measuring co-evolutionary 'progress', both to check that the evolutionary simulation is working properly, and to illuminate issues in theoretical biology. The difficulty in most interesting cases is that, unlike most genetic algorithms research, there is no pre-determined 'fitness landscape' against which progress can be measured. Lineages may evolve against each other with respect to certain domains of competition, but there may be no single correct solution (e.g. no single optimal stable strategy) for each domain.

The remainder of this paper is structured as follows. Section 2 reviews the the goals and methods of our experiments with simulated co-evolution. Section 3 then discusses the need for monitoring techniques in greater detail, and Section 4 describes several of the techniques we have developed.[3]

2 Co-evolution of Pursuit and Evasion

2.1 The Red Queen Effect in Co-evolutionary Simulations

Our interest in measuring co-evolutionary progress arises from our major ongoing research project: co-evolving things with eyes, brains, and wheels, that chase each other around. Or, more technically, using artificial co-evolution to develop neural-network sensory-motor architectures for controlling robot-like autonomous virtual pursuers which chase autonomous virtual evaders around a 2-dimensional (2-D) space, generating pursuit and/or evasion strategies on the basis of simulated visual input.

[3] An expanded version of this paper, with full illustrations, is available [4].

We have argued for the importance of studying these pursuit-evasion contests in previous papers [13, 12]. For the purposes of this methodological paper, we can simply note that pursuit-evasion contests offer a prime scenario for studying co-evolutionary dynamics. Pursuit-evasion contests are common in nature: predators pursue prey and prey evade predators, until the prey get eaten or the predator gets tired and abandons the chase. The success of strategies for pursuit and evasion are often mutually coupled: if a new strategy confers extra fitness on individuals, then that strategy should spread through the population, and its increased frequency makes it more likely that the opponent population will adapt to counteract it, thereby reducing its fitness benefits. It is this co-evolutionary coupling which underlies the Red Queen effect: the fitness landscape of one population is affected by the current strategies of any opponent populations; and the movements of one population over a fitness landscape can significantly alter the fitness landscapes of the other populations. Despite having fixed fitness *functions* for determining the reproductive success of individuals, the fitness *landscapes* will vary over time as adaptations in one population warp and shift and deform the fitness landscape of the other.

2.2 Simulation Methods

Brief details of our simulation system were given in [13], and full details are available in [3]. We will give here here only enough detail to establish the context for discussion of the problems of measuring co-evolutionary progress.

The simulation uses a conventional generational (as opposed to steady-state) genetic algorithm (GA). There are two separate populations which compete and co-evolve against each other: one undergoes selection for pursuit behaviors, the other for evasion. Each population is spatially distributed with local mating and local replacement. That is, each individual in the population is assigned a spatial location on a 2-D grid (with toroidal wrap-around at the edges). When a new generation is bred, each individual is only allowed to breed with other individuals from nearby grid locations, and the offspring is also placed in a nearby grid location. In principle, this spatial structuring of the population should allow for the emergence and maintenance of somewhat distinct *subpopulations*, that shade into each other across "clines".

Each individual has a genotype which is a string of approximately 1600 bits. A relatively complex 'morphogenesis' process translates the genotype into the agent's body morphology: the body has simple effectors, visual sensors, and a recurrent continuous-time dynamic artificial neural network whose parameters (connectivity, weights, thresholds, time-constants, etc) are determined by the genotype. Rather than using variable-length genotypes (e.g. [2]) to allow evolutionary control of the number of artificial neurons in the network, sequences of the fixed-length genotype are ignored in morphogenesis unless they are preceded by an appropriate 'marker sequence' which enables their expression: thus some portions of each genotype may be *active* (i.e. expressed in morphogenesis) while others may be *inactive* 'junk DNA'. Reproduction uses mutation, a stochastic multipoint crossover operator, and a translocation operator: see [3]. Elitism is

used in the breeding phase (i.e. the newly bred population receives an unadulterated copy of the most fit genotype from the previous generation, located at the same grid position as before).

Reproductive success is determined by fitness, and fitness is evaluated for each individual by taking the mean score of a number of noisy trials with differing initial conditions (i.e. individual positions and orientations). In each trial a pursuer and an evader are given fixed amounts of energy which is expended in movement. They compete until one of three termination conditions is met: (1) if there is a collision between the pursuer and the evader; (2) if both contestants have run out of energy and drifted to a halt; or (3) 15 seconds of simulated time have elapsed.[4] Significant noise affects the simulated sensors and effectors, and in the activities of the artificial neural units. For efficiency, we use the same technique as Sims [18], where each individual's fitness is evaluated only in trials against the elite (i.e. highest-scoring) individual from the previous generation of the opponent population: we refer to this technique as LEO (Last Elite Opponent) evaluation. At the end of each trial, the individual under evaluation is given a score. In the experiments discussed below, the score for evaders is simply the amount of simulated time before the trial ended; the score for pursuers is a temporal integral of the instantaneous rate of approach (which encourages the pursuer to approach the evader), plus a 'bonus' reward awarded if a collision occurs. The differences in the scoring techniques mean that the contests are not zero-sum. See [3] for further discussion of this and other fitness scoring methods.

At the end of each generation, various statistics are calculated in order to monitor progress. The genotype of the elite of each generation is saved for use in the LEO contests of the next generation, and acts as a representative of the population for several of our monitoring procedures. Because starting conditions can vary and behavior is noisy, there is uncertainty as to whether the genotype ranked as the elite really is the best in the population or is an unexceptional individual that was lucky in evaluations. We reduce this uncertainty, using the standard statistical techniques such as blocking and stratified sampling across initial conditions. Thus, it is *highly* unlikely that the same poor genotype will be incorrectly identified as a legitimate elite over several successive generations.

3 The Need for New Measurement Techniques

3.1 Instantaneous Fitness Tells Us Little

During a simulation experiment, the fitness values of each individual are available at the end of each generation. These values have to be calculated and stored for use in the breeding phase of the generational GA. These values usually form the basis of monitoring progress in non-co-evolutionary applications where the fitness landscape is fixed in advance. In such non-co-evolutionary scenarios, progress can be monitored by plotting summary statistics (such as the mean or best) of the

[4] The simulation approximates continuous time using Euler integration techniques at a temporal resolution of 100 steps per simulated second.

fitnesses of the population. A successful simulation experiment can be expected to show a gradual increase in a relevant population fitness measure. The fitness measure will asymptote as the population converges on one or more fitness peaks, signifying an end to the evolutionary search. In unsuccessful experiments (e.g. where the GA parameters have been poorly set), fitness may never increase above that of the initial random population, or it may increase and then hold at a comparatively low value (indicating convergence to a suboptimal local fitness peak), or it may climb to a high value and then subsequently fall away (indicating convergence followed by genetic drift, caused by an excessive mutation rate).

However, in a co-evolutionary scenario, the Red Queen effect makes it hard to monitor progress by taking instantaneous measures of fitness at the end of each generation. Because fitnesses are defined relative to a co-evolving set of traits in other individuals, the fitness landscape(s) for the co-evolving individuals can vary dynamically. Hence periods of comparative stasis in instantaneous measures could signify a corresponding evolutionary stasis or could disguise a period of tightly-coupled co-evolution where adaptive changes in one population (which would register as increases in instantaneous fitness if the opponent population was held fixed) are matched by adaptive counter-changes in the opponent population (thereby holding the instantaneous fitness measures of both populations close to the values exhibited prior to the change and counter-change in strategies). Similarly, if the instantaneous measures decrease over time, this may represent either a setback in progress due to genetic drift, or co-evolutionary progress where *both* populations have adapted to make the pursuit-evasion contests significantly more difficult for their opponents.

As illustration, Figures 1 and 2 show instantaneous fitness measures for pursuer and evader populations in an experiment lasting 700 generations. If each population got better, the lines should go up. Although the data for the pursuer population shows a steady climb in fitness over the first 200 generations, the mean fitness value at generation 205 is never improved upon, and the mean fitnesses over generations 500 to 600 are roughly the same as those over generations 50 to 100, which could be interpreted as a lack of progress in the intervening 400 generations. Superficially, the data for the evader population is even worse: mean population fitness is at its highest at the start of the experiment, showing a steady decline over the first 300 generations, followed by 400 generations of roughly constant fitness at a level around 70% of that exhibited at the start. Given that both populations are selected for maximizing fitness, these data could indicate that something is seriously wrong, with progress either not occurring or not being maintained. As we shall see in the following sections, real progress *is* occurring in this experiment, but other monitoring techniques are required to demonstrate this.

We use the term *fitness ambiguities* to refer to such cases where qualitative trends in time-series of instantaneous fitness measures could feasibly be interpreted as either continuing progress or as a breakdown of the co-evolutionary process. Fitness ambiguities introduce two problems:

First, how do we know when to terminate an experiment that is failing due to

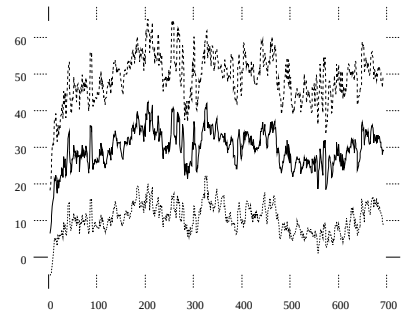

Fig. 1. Instantaneous fitness measures for pursuer population over 700 generations of co-evolution. Horizontal axis is generation number, vertical is fitness. Solid line is mean population fitness: dashed and dotted lines are mean plus and minus one standard deviation, respectively. Data smoothed by calculating rolling average over preceding five generations. See text for discussion.

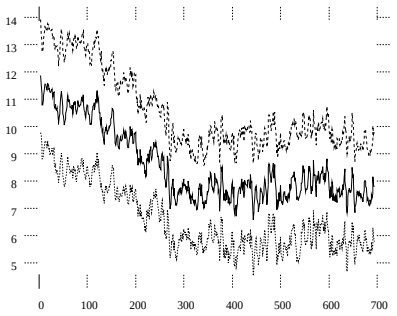

Fig. 2. Instantaneous fitness measures for evader population over 700 generations of co-evolution: format as for Figure 1. The range of fitness values is different from that of the pursuers, because of the difference in evaluation functions. See text for discussion.

bugs or poor choice of parameters? The importance of this problem can be appreciated when one considers the computational requirements of our co-evolutionary simulation experiments. In the experimental regime described in Section 2.2, running on an unladen Sun SPARC20, with two populations each of size 100, using LEO competitions with 9 trials per individual, our fully optimized C code manages to evaluate each individual in an average of about 10 seconds (evaluation time can vary greatly because of the multiple termination conditions described in Section 2.2). So one complete generation takes roughly 35 minutes, and a single experiment of 1000 generations takes a little over three weeks of continuous computation. As one of our aims is to study the effects of varying the experimental conditions (e.g. different evaluation functions, genetic encoding techniques, settings of the constants governing motion, or environmental circumstances such as worlds with obstacles or boundaries), it would be nice to be able to disambiguate fitness ambiguities at the earliest possible opportunity. In short, we need to know whether to kill a pointless experiment or allow hidden progress to continue.

The second problem stems from our concern to establish an informative and reliable characterization of the co-evolutionary dynamics exhibited by our experiments. The possibility of ambiguities in time-series of instantaneous fitness measures makes such characterization impossible without further analysis. Yet we believe that our experiments offer an opportunity to empirically explore issues of progress that are keenly debated in the evolutionary biology literature: notions of teleology, diffuse vs. tight co-evolution, smooth vs. punctuated equi-

libria, directional progress vs. cycling, etc. Resolving these issues has proven difficult because of the limitations of fossil records, genetics, comparative morphology, and comparative psychology: our simulations offer an opportunity to study the co-evolution of complex autonomous agents under experimental conditions that allow detailed measurements of genotypes, body morphologies, and behaviors.

3.2 (The Lack of) Related Work

These two problems have led us to explore techniques that allow us to reliably monitor the sometimes hidden co-evolutionary progress in our simulations. To our surprise, we found that no applicable techniques had been developed in the fields of artificial life, theoretical biology, behavioral ecology, or evolutionary genetics. The complexity of our simulations violates many of the simplifying assumptions on which theoretical studies of evolution, co-evolution, and population dynamics, are founded. Furthermore, although other artificial life research has employed co-evolution to develop autonomous agent architectures (e.g. [18, 15]), such work has concentrated mainly on the end results, rather than on the dynamics of the co-evolutionary process.

Co-evolving pursuit and evasion strategies may appear related to the long-established literature on theoretical modeling of predator-prey population dynamics. Yet all such work with which we are familiar, from the well-known deterministic Lotka-Volterra equations to the more recent spatially-distributed stochastic population models (see e.g. [14]), depends on monitoring fluctuations in the sizes of two competing populations. As the population size is constant in our experiments, this large body of theoretical work is of little use to us.

Studies in ethology and behavioral ecology (e.g. [11]), although acknowledging the importance of predator-prey arms races, focus on the functions of current behavior rather than the dynamics of co-evolution. Macroevolutionary theory (e.g. [6]) typically treats co-evolution as a phenomenon that is hard to observe outside the fossil record.

Finally, evolutionary genetics and almost all research in either theoretical biology or artificial life which could be relevant (e.g. [1]) studies (co-)evolution at the level of discrete genes for particular traits. W. Hamilton and his associates (e.g. [17]) have developed techniques for visualizing and analyzing simplified co-evolutionary systems. However, the genetic encoding used in our simulations is sufficiently complex that there is no clear method of identifying a gene for a given trait: sequences of bits in a genotype may affect the connectivity of the artificial neural network, or may determine a parameter for one of the artificial neurons, but the observable behavior of the phenotype is a complex emergent property of the network interacting with the environment, which itself includes another network (controlling the opponent). In fact, because the space of possible behaviors for our artificial agents is continuous in both time and space, a precise definition of what constitutes a 'trait' is problematic. Although Kauffman's work on co-evolution in NKC fitness landscapes (e.g. [10, Chapter 6]) appears relevant,

it is not clear how to determine N, K, or C for our simulations,[5] particularly as all three factors could vary dynamically as co-evolution progresses.

These issues have led us to conclude that our work is exploring largely uncharted territory: we are attempting to gain the insight offered by theoretical analyses in an artificial co-evolutionary system sufficiently complex that no established theoretical analysis tools are applicable. In the remainder of this paper we describe analysis and monitoring techniques we have developed to fill this gap.

4 Measurement Techniques

4.1 Ancestral Opponent Contests

In an earlier paper [13], we noted that one possible technique for monitoring co-evolutionary progress is to evaluate an individual I from generation g against representatives of I's opponent population from each previous generation $g - \Delta g : \Delta g \in \{0, 1, \ldots, g\}$. That is, I is entered into contests with the 'ancestors' of its current opponent. Current individuals should do well against outdated opponents; the more ancient the opponent, the better they should do. More technically, if progress has occured then we might expect that I's fitness will increase with Δg: the fitness scores for I will be positively correlated with Δg over some time-scale which we will refer to as the 'evolutionary time-lag' τ; I should in general do better in competitions against opponents drawn from earlier opponent generations, because co-evolutionary adaptation in I's population should have rendered these strategies less effective.

However, we needn't expect performance against ancestral opponents to improve all the way back. For example, ancient ancestors may have had tricks that more recent ancestral opponents have lost. So there are no strong reasons for expecting the positive correlation to be extended indefinitely, nor even for expecting the correlation to be monotonically increasing. For example: while I may reasonably be expected to do better in contests with individuals drawn from the opponent population at (say) generation $g - 10$ than in contests with opponents from generation g, it could be that when I competes with individuals from opponent generation $g - 100$ it fares much worse than it does against opponents from generation g. This statement may appear counterintuitive, but there are at least two possible explanations for such a result: the limits of 'evolutionary memory', and the possibility of cyclical trajectories through strategy space over evolutionary time-scales.

- Co-evolutionary adaptation in a population P_1 (e.g. pursuers) can be expected to render strategies from recent generations of an opponent population P_2 (e.g. evaders) less effective. But it is feasible that the genetic changes selected in P_1 to combat these recent P_2 strategies eliminates or reduces phenotypic traits in P_1 that contributed to counteracting P_2 strategies earlier

[5] N is the number of traits coded on a genotype, K the number of epistatically linked traits within a genotype, and C the number of epistatically linked traits in a co-evolving (opponent) species.

in evolutionary time. Such displacements will not reduce the fitness of P_1 individuals if the distant P_2 strategies are no longer employed in the current or recent generations of P_2. In a reciprocal manner, the P_2 population is less likely to have retained the genetic material responsible for the distant strategies if P_1's subsequent counter-adaptations rendered them ineffective. If bounds are imposed by limited resources or developmental constraints, displacement of out-dated genetic material is likely to form the basis of continuing adaptation. Even if such limitations are not significant, the escalating arms race may render distant P_2 strategies obsolete, so that they and their P_1 counter-strategies fade away through neutral mutations. Either way, the ultimate result is that current P_1 individuals fare badly when pitted against P_2 individuals from generations sufficiently distant that they are beyond the 'evolutionary memory' of the P_1 genomes.

– Cycling between strategies is possible if there is an intransitive dominance relationship between strategies. For example, suppose one attempted to co-evolve two populations that compete by playing each other at the childrens' game "rock-paper-scissors"[6] where each individual in the population is limited to one genetically determined choice which it uses in all contests in its lifetime (i.e., in game-theoretic terms, only pure strategies are allowed). More generally, consider a co-evolutionary competition between two populations P_1 and P_2 and let $P_i(f) \succ P_j(g)$ denote the fact that individuals from population P_i at generation f generally win competitions against individuals from population P_j at generation g; assume that $P_1(g) \succ P_2(g - \Delta g_a)$ and $P_1(g - \Delta g_a) \succ P_2(g - \Delta g_b)$ with $\Delta g_a < \Delta g_b$. A transitive dominance hierarchy would exist if $P_1(g) \succ P_2(g - \Delta g_b)$, but an intransitive dominance cycle could be established if $P_2(g - \Delta g_b) \succ P_1(g)$. LEO contests would tend to make cycles especially likely.

Of course, it is possible that particular co-evolutionary systems will not exhibit either of these two phenomena, but it is crucial to appreciate that, in general, co-evolutionary systems such as ours should not be expected to exhibit continuous progressive adaptation with each generation 'improving' on previous generations, toward some optimal state. Our qualitative notion of an evolutionary time-lag τ (which may itself vary over evolutionary time) serves to emphasize that the results from ancestral opponent contests need to be judged with care.

Figures 3 and 4 show results from ancestral opponent contests for the elite pursuer and elite evader from generation 700 of the run illustrated in Figures 1 and 2. In both cases there is a general trend towards higher fitness as Δg increases[7] (note that Δg decreases from left to right: in this case $\Delta g = 700$

[6] This is a two-player game where each player simultaneously announces a choice of 'scissors', 'paper', or 'rock': unless there is a tie, the player with the dominant choice wins. The dominance relationships are: scissors cut the paper; paper smothers the rock; and rock breaks the scissors.

[7] These data could perhaps be better characterized as periods of relative stasis either side of a period of significant improvement over generations 200 to 350: we return to discussion of these data in Section 4.2.

at $g = 0$ and $\Delta g = 0$ at $g = 700$). These ancestral fitness data support the claim made in Section 3.1 that the data presented in Figures 1 and 2 mask some underlying co-evolutionary progress.

Of course Figures 3 and 4 only compare generation 700 elites against their ancestral opponents. What about the corresponding data for all the previous elites against *their* ancestral opponents? To show that we would need as many graphs as there have been generations. However there is an efficient way to display all such comparisons: by representing data from Figures 3 and 4 as a row of intensities with high scores darker and lower scores lighter. Clearly for generation 1 we would have only one previous generation from which to draw ancestral opponents, yielding a single cell. As we compare elites from each successive generation against all of their ancestral opponents the rows will get longer and longer, and we can stack them one above the other, so the top row represents all the data from a single plot like Figures 3 or 4. We refer to the fitness data from these tests as CIAO data (from Current Individual vs. Ancestral Opponents).

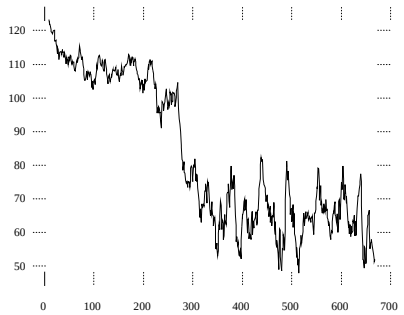

Fig. 3. Generation 700 elite pursuer scores better against earlier evaders. Graph shows pursuer ancestral fitness scores for the run illustrated in Figures 1 and 2. Horizontal axis is generation number g, from which an ancestral opponent is drawn; vertical axis is fitness scored by elite pursuer from generation 700 in contest with elite evader from generation g. Data smoothed by calculating rolling average over preceding ten generations. See text for discussion.

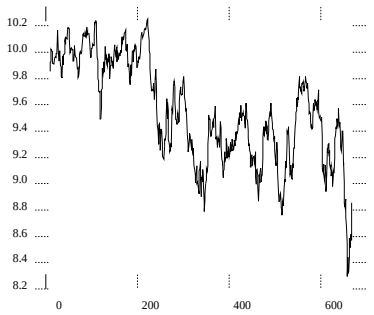

Fig. 4. Generation 700 elite evader scores better against earlier pursuers. Graph shows evader ancestral fitness scores for the run illustrated in Figures 1 and 2. Horizontal axis is generation number g, vertical is fitness scored by elite evader from generation 700 in contest with elite pursuer from generation g. Smoothing as for Figure 3. See text for discussion.

Figure 5 shows a simplified schematic of the CIAO display format, for visualizing the results of an experiment where two populations P_1 and P_2 have co-evolved. Essentially, the format is a triangle formed by stacking successive 2-D data-sets such as those shown in Figures 3 and 4. In order to spatially compress

the data, the fitness scores for P_2 determine the darkness of each cell on the grid: darker cells signify higher scores. Thus the top row of cells is the ancestral fitness data for P_2 at generation 5, the one below that is the ancestral fitness data for P_2 at generation 4, and so on. The cells along the diagonal edge therefore represent the score of the elite of P_2 from generation g in contest with the elite from P_1 at the same generation, and the shading in the *next* diagonal line represents the scores from the elite of P_2 at generation g in contest with the elite of $P1$ from generation $g-1$: a clear parallel with LEO contests; although these are the instantaneous fitness scores of the elites, they should be in close agreement with the instantaneous population-average fitness data such as Figures 1 and 2.

The utility of this CIAO display format is indicated by Figure 6, which shows idealizations of the patterns that would be present if the co-evolutionary process was affected by limited evolutionary memory or by cyclic trajectories through strategy space.

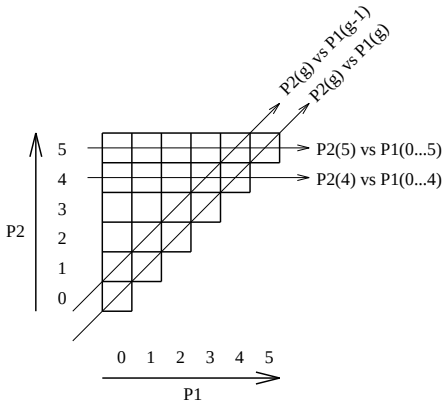

Fig. 5. Simplified Schematic of the CIAO fitness plots: the cells would be shaded to reflect the scores of individuals from P_2, with darker shades indicating higher fitness. See text for further details.

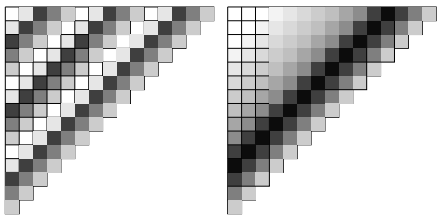

Fig. 6. Idealised illustration of patterns indicating intransitive dominance cycling (left, where current elites do well against opponents from 3 or 4 generations ago but not so well against those from generations further back); and limited evolutionary memory (right, where current elites do well against opponents from 3, 8, or 13 generations ago but not so well against generations inbetween). The presence of straight diagonal bands of intensity indicates a constant τ: if τ varies, the bands would follow curves.

The gray-scales in a CIAO plot could be set by normalizing all the scores in each data-set to the range of the entire data-set. Thus the darkest cell(s) in the figures would represent the highest scores in the entire data-set, and the brightest the worst. An alternative method of setting the gray-scales is to normalize all scores in each row in the image to the range of the data *in that row*. The effect on the image is similar to histogram-equalization used for contrast enhancement in image processing (e.g. [8]), but here the adjustment of gray-scales on each row

has a natural interpretation: the darkest cell in each row signifies the highest score for the elite in P_2 of the generation plotted on that row, and the lightest the worst. Figures 7 and 8 show CIAO data from the experiment discussed above, with gray-scales normalized across rows.

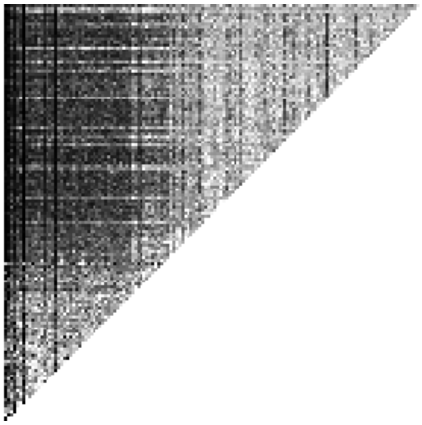

Fig. 7. Pursuer CIAO fitness scores for the run illustrated in Figures 1 and 2. Evader generations 0 to 700 run left to right in steps of 5, pursuer generations 0 to 700 run bottom to top in steps of 5. Darker cells represent higher fitness scores, gray-levels normalized to data range across each row. See text for discussion.

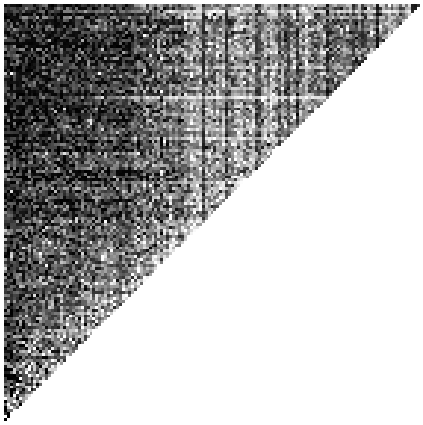

Fig. 8. Evader CIAO scores for the run illustrated in Figures 1 and 2. Pursuer generations 0 to 700 run left to right in steps of 5, evader generations 0 to 700 run bottom to top in steps of 5. Intensity range as for Figure 7. See text for discussion.

From these CIAO plots it is possible to give some account of the co-evolutionary dynamics of this particular experiment. Figures 7 and 8 both have the greatest density of (lighter) low-score cells towards the diagonal edge, with more of the (darker) high-score cells towards the left-hand edge. This indicates that in general there has been continuous progress in *both* populations, and that neither of the patterns shown in Figure 6 have occurred in the first 700 generations of this experiment.

In the pursuer scores there appears to be a significant improvement around generations 220 to 300: elite pursuers from all subsequent generations do well against the first 220 generations of elite evaders; moderately well against those in the range 220 to 300, and then fare fairly badly against evaders from generations 300 to 700. A similar pattern is revealed in the evader score, although there is evidence that the evaders improve slightly around generations 550 to 700.

Examination of CIAO data can clearly help identify major changes in the relative fitnesses of elites from the two populations, and helps to characterize

the dynamics of the co-evolutionary experiment to date. Not surprisingly, these benefits come at a (computational) cost: the number of evaluations required for a complete CIAO analysis of a particular simulation experiment can easily exceed the number of evaluations in the experiment itself. Consider a co-evolutionary simulation lasting n_g generations, with two populations each of size p. Then there will have been $2n_gp$ fitness evaluations over the duration of the experiment. A CIAO analysis of the resulting two sets of elite individuals requires n_g^2 evaluations. Thus, the full CIAO analysis will take longer than the experiment itself once $n_g^2 > 2n_gp$ (i.e. $n_g > 2p$), or when the number of generations exceeds twice the population size, as it does in this run. Of course, computational savings can be made by sub-sampling the CIAO space at an appropriate resolution Δg, e.g. evaluating once every $\Delta g = 10$ generations. The computational savings come with a corresponding loss of temporal resolution: the time-scale at which co-evolutionary interactions can be monitored is reduced (e.g. with $\Delta g = 10$, tightly coupled adaptation/counter-adaptation events occurring within 10 generations may be missed) and sub-sampling the space also introduces issues of spatial frequency aliasing which could be highly disruptive if intransitive strategy cycles are present in the CIAO data.

4.2 Genetic Distance Measures

Although the CIAO data indicates co-evolutionary interaction and progress at the phenotype level (i.e. by showing the fitness values resulting from evaluations of behavioral performance), it gives no indication of the corresponding dynamics at the genotype level; yet for a complete account of the co-evolutionary process, such genetic analysis is necessary. To this end, we have developed a set of simple monitoring procedures which gives good indication of significant co-evolutionary interactions at the genetic level.

Furthermore, because these genetic analysis techniques do not require (computationally intensive) evaluation of phenotypes, relevant data can be calculated on-line during the progress of an experiment without incurring a significant processing overhead. Therefore these techniques offer the advantage that they can be used to monitor progress in an experiment while it is running, and hence allow for identification of experiments which should be terminated due to lack of progress. Two techniques are introduced below: "elite bitmaps" and "ancestral Hamming maps"; a third technique, "consensus distance plots", is described in [4]. All three techniques analyze one population in isolation from any other co-evolving populations: the intention is to identify periods of significant genetic change which can clarify features present in the CIAO data. For brevity, we illustrate these techniques with examples generated using only the pursuer population from the experiment analyzed in the previous sections: in practice it is necessary to separately apply these monitoring techniques to both the pursuer *and* the evader populations: see [4] for illustration of the results of applying these techniques to the evader population. Another technique we are exploring is *chronospeciation analysis*: testing whether successful offspring result from crossing individuals of generation g with individuals from generation $g - \Delta g$. If not,

'chronospeciation' has occured: current individuals can no longer breed with their ancestors, indicating that significant evolutionary change has occurred. The chronospecies concept in theoretical biology can't be tested very easily in nature, but simulations allow us to perform these cross-generational breeding experiments.

Elite Bitmaps At the end of each generation in our experiments we record the genotype of the elite individual in each population. Figure 9 illustrates the genotypes for the elite pursuers over the 700 generations of the experiment. This 'elite bitmap' shows the elite genotype at each generation, stacked horizontally with the earliest generation at the top and the latest at the bottom. Such raw genotype data reveals some qualitative structures, three of which are worth attention:

- There are clear vertical bands of varying extent: these bands correspond to bits in the elite genotypes which were largely unchanged over a series of generations. The horizontal extent of the band is governed by the length of the gene sequence which is constant between elites of successive generations, and the vertical extent indicates the number of generations during which this sequence was maintained in the elite.
- There are several 'noisy' areas where no banding is present: In general these areas correspond to sequences on the elite genotypes which have no impact at the behavioral level: either because they are not expressed at the morphogenesis stage, or because the morphological features governed by these sequences has a negligible effect on the behavior of the individual. In either case, a high degree of genetic variance in such sequences can be expected on the elite genotypes, yielding little or no correlation between successive generations and hence no clear vertical banding in the bitmap.
- There are horizontal 'faultlines' at various locations. Where banding either commences, ceases, or continues but with a different pattern of bits in the band. These faultlines indicate significant changes in the genetic profile of the population elite. The horizontal extent of the faultline indicates the degree of change. If the faultline extends across the entire genotype, then a new genetic 'strain' of elite has emerged from the underlying population. Faultlines with a more limited horizontal extent indicate more gradual changes in the genetic constitution of the elite. Some faultlines will be caused by the translocation operator employed in our GA. If the faultline initiates a new pattern of banding, or marks a transition from 'noise' to banding, then it indicates a major change to the affected sequence of genotype which is retained over successive generations. If the new banding pattern fades into noise, then the sequence could be 'hitch-hiking': that is, the sequence does not itself contribute to the fitness of the elite, but is retained by virtue of its presence in a genotype which has *other* sequences that confer sufficiently high fitness to maintain the genotype as the elite.

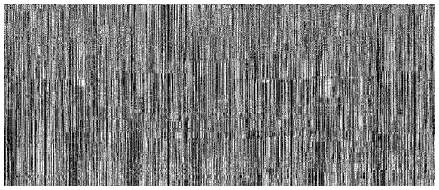

Fig. 9. Bitmap for elite pursuer genotypes. 1600-bit genotypes stacked horizontally, with generation 0 at the top and 700 at the bottom. Dark pixels represent a '0' in the genotype; light pixels represent '1'.

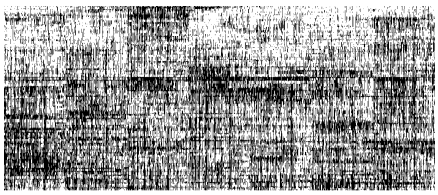

Fig. 10. Result of image-processing on the bitmap of Figure 10. Dark pixels indicate bits which are elements of vertical bands lasting for 9 generations or longer. See text for discussion.

These qualitative phenomena can be identified in a more objective manner by applying elementary image processing operations to the elite bitmap: convolving the bitmap with appropriate masks allows for the automatic highlighting of areas of banding, faulting, and noise. Figure 10 shows the result of convolving the elite bitmap with a simple one-dimensional mask which highlights vertical bands of 9 generations or more. Darker areas in the image indicate the presence of banding, and lighter areas indicate noise (or bands lasting less than 9 generations). The horizontal faultlines are also more prominent in the processed bitmap: Around generations 250 to 300 there is a clear group of faultlines with a large horizontal extent (indicating change in large sequences on the elite genotype). Several of these faultlines are followed by periods of strong banding (i.e. dark areas on the processed bitmap).

Ancestral Hamming Maps While the (processed) elite bitmap can help identify genotype sequences that are retained over successive generations and instances of significant change marked by faultlines, it is essentially a qualitative technique. To make meaningful comparative statements, quantitative measures of genotype-sequence retention and change are required.

In particular, it can be useful to quantify the degree to which a given elite genotype shares genetic material with the elites in the preceding generations. An obvious measure to use is the Hamming distance between the two bit-string genotypes. For brevity, we will use the following notation: let $E(g)$ denote the elite individual in a population at generation g; let $G(E(g))$ denote the genotype for $E(g)$; and let $H(f, g)$ denote the Hamming distance between $G(E(f))$ and $G(E(g))$ (note that if $f = g$ then $H(f, g) = 0$). In a manner similar to the CIAO plots, we can determine the Hamming distances $H(g, g - \Delta g) : \Delta g \in \{1, \ldots, g\}$. That is, the Hamming distances between $G(E(g))$ and the genotypes of the elites from each preceding generation of the same population. This "ancestral Hamming data" can be plotted as intensities on a 2-D grid, resulting in an *ancestral Hamming map*. Figure 11 shows a schematic ancestral Hamming map, and Fig-

ure 12 shows idealized qualitative patterns that indicate particular features of the underlying evolutionary dynamics. Further quantitative analysis of the ancestral Hamming data can be guided by searching for such qualitative features:

- If there is constant slight change in the genetic constitution of the elite over successive generations, then $H(g, g - \Delta g)$ should fall off smoothly as Δg increases, and all horizontal lines of cells in ancestral Hamming map will have roughly the same H (see the left-hand example in Figure 12).
- If a new elite genotype occurs at generation g and is sufficiently fit with respect to the opponent population that it or its immediate descendants also form the elite for subsequent generations, then this will show in the ancestral Hamming map as a 'wedge' of low-H cells on the grid: see the center example in Figure 12.
- Ancestral Hamming maps can give a useful indication of the genetic constitution of the underlying spatially distributed population in a GA. Because a spatially distributed population is capable in principle of sustaining separate 'subpopulations', as described in Section 2.2, there is no reason to expect that $H(g, g - \Delta g)$ will always decrease smoothly as Δg increases. In particular, it is possible that $H(g, g - \Delta g)$ is high for some small Δg, but low for a larger value of Δg (see the right-hand example in Figure 12). Such a situation would indicate that $E(g)$, is more strongly related to the earlier elites than to the more recent elites, and that the more recent elites came from a different 'subpopulation'. If ever $G(E(g))$ has a comparatively high H for all $G(E(g - \Delta g))$, then the indication is that $E(g)$ is a member of a subpopulation which shares comparatively little genetic material with the previous elites, and so g is the first generation in which members of that subpopulation have attained sufficient fitness to be selected as the elite.

Figure 13 shows an ancestral Hamming map for elites of the pursuer population introduced in Figure 1. For roughly the first 220 generations there is little evidence of any significant structure. Around generations 220 and 260 there are the first two of a number of dark 'wedges', indicating the emergence of genotype sequences which remain (partially) present in the population elite for many generations: even the column at generation $g = 400$ shows some darkening (i.e. lower H) around $\Delta g = 140$ and $\Delta g = 180$; this clearly (and quantitatively) indicates that the changes in genotype at $g = 220$ and $g = 260$ were retained for many subsequent generations: a fact that wasn't particularly clear in the processed bitmap of Figure 10.

Consensus Distance Plots Although monitoring genetic change in the elite genotypes is a valuable source of information for characterizing (co-)evolutionary dynamics, it is important to ensure that changes in the elite genotype are reflective of changes in the underlying population of genotypes. When the evaluation process is noisy or uncertain, there is always a possibility that changes in the elite genotype between one generation and the next are stochastic evaluation artifacts, rather than significant evolutionary events. To disambiguate the two,

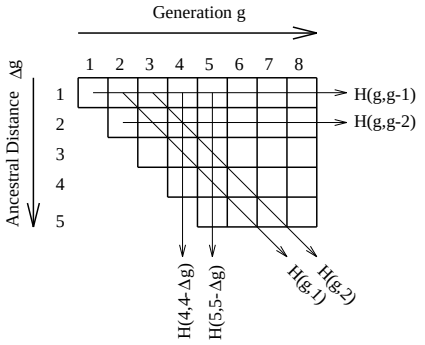

Fig. 11. Simplified schematic of Ancestral Hamming Maps. $H(f,g)$ denotes Hamming distance H between $G(E(f))$ and $G(E(g))$. Grid-cells would be shaded to reflect values of H (lower H given darker shading). Horizontal lines of cells indicate $H(g, g-\Delta g)$ for constant Δg while g varies. Vertical lines of cells indicate $H(g, g-\Delta g)$ for constant g while Δg varies. Diagonal lines of cells indicate $H(g,f)$ for constant f while g varies. This map terminates at $\Delta g = 5$, but could have been continued to $\Delta g = 8$.

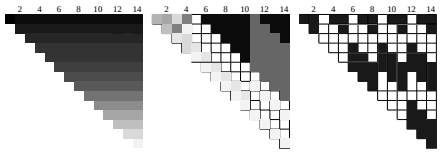

Fig. 12. Idealised illustration of patterns in Ancestral Hamming Maps, for three cases of generation 0 to 14. Left: steady genetic change (either through retention of adaptive mutations or through neutral genetic drift); the elite of each generation shares much genetic material with the elite of the previous generation (i.e. $H(g, g-1)$ is always low), but accumulated genetic change results in more distant ancestral elites having much less shared genetic material (i.e. $H(g, g-d) : d >\approx 10$ is always high). Center: at generation 5 a new elite genotype appears that is largely dissimilar from all previous genotypes. The $g = 5$ elite genotype (or its descendants) remains as the elite until generation 10, when another new genotype, sharing much material with the $g = 5$ elite, becomes the elite for the remainder of the generations shown on the map. Right: two converged subpopulations alternate in their role as the elite, with little or no genetic material shared between the subpopulations.

it is necessary to monitor a representative population statistic. We have found that significant evolutionary events in our experiments can be identified by monitoring, at each generation, the distribution of Hamming distances from the population's *consensus sequence* to each individual genotype in the population. The consensus sequence for the population can be thought of as the "average" genotype in the population. The rationale for these plots, and illustration of their use on data from the above experiment, are given in [4].

5 Conclusion

Co-evolutionary simulations for developing advanced artificial autonomous agents are so complex that new techniques for monitoring progress are required. Al-

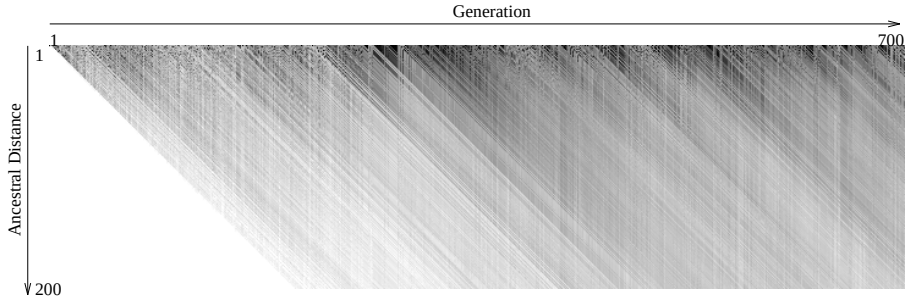

Fig. 13. Ancestral Hamming Map for elite pursuers. Generations 1 to 700 run from left to right. Ancestral distance (Δg) increases from top to bottom: top edge indicates Hamming distance from elite of generation $g - 1$: bottom edge indicates Hamming distance from elite of generation $g - 200$. Intensities indicate Hamming distance as a percentage of genome length: intensity increases linearly from black (0% distant) to white (> 50% distant). Note that the Hamming distance between two randomly generated bit-strings of the same length will be 50% of the length (on the average).

though this paper concentrated on our ongoing work in evolving pursuit and evasion strategies, we believe that both the problems and the solutions we have identified are general: it is likely that fitness ambiguities will occur in any co-evolutionary situation, and that these ambiguities can be resolved by combining CIAO tests with the various genetic analysis techniques described in Section 4.2. Thus, researchers interested in monitoring or analyzing evolutionary activity in either real or artificial systems where the fitness landscapes change over time (either through co-evolution or because the non-biotic environment is dynamic) should find use for the analysis methods we have described here. The open-ended nature of co-evolutionary simulations makes it difficult to detect the Red Queen's presence: the techniques developed in this paper now let us track her protean manifestations.

References

1. M. Bedau and N. Packard. Measurement of evolutionary activity, teleology, and life. In C. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, eds, *Artificial Life II*, pp.431–461. Addison Wesley, 1992.
2. D. Cliff, I. Harvey, P. Husbands. Explorations in evolutionary robotics. *Adapt. Behav.*, 2(1):71–108, 1993.
3. D. Cliff and G. F. Miller. Co-evolution of pursuit and evasion II: simulation methods and results. COGS Technical Report CSRP377, University of Sussex, 1995.
4. D. Cliff and G. F. Miller. Tracking the Red Queen: Measurements of co-evolutionary progress in open-ended simulations. COGS Technical Report CSRP363, University of Sussex, 1995.
5. R. Dawkins. *The Blind Watchmaker*. Longman, Essex, 1986.

6. N. Eldredge. *Macroevolutionary dynamics: Species, niches, and adaptive peaks.* McGraw-Hill, 1989.
7. D. J. Futuyama and M. Slatkin, editors. *Coevolution.* Sinauer, 1983.
8. R. C. Gonzalez and P. Wintz. *Digital Image Processing.* Addison-Wesley, 1977.
9. S. J. Gould. *Wonderful Life: The Burgess Shale and the Nature of History.* Penguin, 1989.
10. S. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution.* OUP, 1993.
11. J. R. Krebs and N. B. Davies. *An Introduction to Behaviuoral Ecology.* Blackwell Scientific, 1993.
12. G. F. Miller and D. Cliff. Co-evolution of pursuit and evasion I: Biological and game-theoretic foundations. Technical Report CSRP311, University of Sussex School of Cognitive and Computing Sciences, 1994.
13. G. F. Miller and D. Cliff. Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In D. Cliff, P. Husbands, J.-A. Meyer, and S. Wilson, editors, *Proc. Third Int. Conf. Simulation Adaptive Behavior (SAB94)*, pages 411–420. M.I.T. Press Bradford Books, 1994.
14. E. Renshaw. *Modelling Biological Populations in Space and Time.* Cambridge University Press, 1991.
15. C. Reynolds. Competition, coevolution, and the game of tag. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 59–69. M.I.T. Press Bradford Books, 1994.
16. M. Ridley. *The Red Queen: Sex and the evolution of human nature.* Viking, London, 1993.
17. J. Segers and W. D. Hamilton. Parasites and sex. In R. E. Michod and B. R. Levin, editors, *The evolution of sex: some current ideas*, pages 176–193. Sinauer, Sunderland, MA, 1988.
18. K. Sims. Evolving 3D morphology and behavior by competition. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 28–39. M.I.T. Press Bradford Books, 1994.
19. L. van Valen. A new evolutionary law. *Evolutionary Theory*, 1:1–30, 1973.

This article was processed using the LaTeX macro package with LLNCS style