

Unpicking Tartan CIAO Plots: Understanding Irregular Coevolutionary Cycling

John Cartlidge & Seth Bullock

School of Computing, University of Leeds, Leeds, LS2 9JT

Tel: +44 (0)113 233 5322, Fax: +44 (0)113 233 5468, Email: {johnc, seth}@comp.leeds.ac.uk

June 4, 2004

Abstract

We report results from a series of studies coevolving players for simple Rock–Paper–Scissors games. These results demonstrate that “Current Individual versus Ancestral Opponent” (CIAO) plots, which have been proposed as a visualization technique for detecting both coevolutionary progress and coevolutionary cycling, suffer from ambiguity with respect to an important but rarely discussed class of cyclic behavior. While regular cycling manifests itself as a characteristic banded plot, irregular cycling produces an irregular tartan pattern which is also consistent with random drift through strategy space. Although this tartan pattern is often reported in the literature on coevolutionary algorithms, it has received little attention or analysis. Here we argue that irregular cycling will tend to be more prevalent than regular cycling, and that it corresponds to a class of coevolutionary scenario that is important both theoretically and in practice. As such, it is desirable that we improve our ability to distinguish its occurrence from that of random drift, and other forms of coevolutionary dynamic.

1 Introduction

The majority of adaptive systems on this planet are biological, and as a result share several characteristics. Foremost among these is that they are all a product of natural evolutionary processes. More specifically, since populations of adaptive systems have not tended to evolve in isolation, the selective pressures responsible for shaping them have been significantly influenced by the presence and behavior of the other organisms in their environment: competitors, predators, prey, offspring, mates, etc. Hence it is *coevolution*, rather than evolution *per se*, that is responsible for adaptive behavior.

In an effort to either engineer artificial adaptive systems or better understand natural systems, adaptive behavior researchers have employed *coevolutionary* algorithms as design tools and simulation models. However, while this research remains promising, it has raised several problematic issues. Perhaps the most pressing of these concerns our ability to understand the dynamics of coevolutionary systems (e.g., Watson & Pollack, 2001).

In particular, there are inherent difficulties associated with detecting coevolutionary progress. Often, it is virtually impossible for an observer to know whether populations are improving over time. A short-term improvement in fitness, relative to contemporary competitors, does not necessarily lead to long-term improvement in some objective sense, since it is possible for coevolutionary systems to cycle: coevolving populations may follow a repeating sequence of adaptive transitions.

For instance, consider coevolving two populations of players for the simple parlor game Rock–Paper–Scissors (RPS). In each bout of the game, two players simultaneously choose one of the three possible moves (Rock, Paper, Scissors), the winner being decided according to an intransitive superiority relationship (Rock blunts Scissors, Scissors cuts Paper, Paper covers Rock—draws are shared). Although the global optimum for either player is to play each move with equal probability, any deviation from this strategy on the part of either population encourages a complementary counter-deviation.

If one population is biased in favor of playing Paper, for example, the other will benefit from playing Scissors more often than one third of the time. As the frequency with which Scissors is played increases, the first population is under selection pressure to increase the frequency with which its members play Rock. In response, the opponent population will tend to favor Paper. In this way, while both populations are continually favoring offspring better able to compete against their current opponents, in the longer term, the populations are continually and repetitively cycling through a sequence of globally sub-optimal strategies as they seek to exploit the temporary biases of their opponents. Consequently there is no objective improvement—one cannot guarantee that an individual will outperform its ancestors.

In standard evolutionary algorithms, because the fitness of an individual is measured against a static function, continual progress can be detected as an improvement in fitness over time. However, this is not true of coevolutionary systems, where an individual's fitness is measured *relative* to its contemporary opponents. When this type of relative fitness measure is plotted against time, individuals from different generations are being compared using a metric which is itself varying unpredictably over time, since the opponents against which they were assessed will not, in general, have been the same. These considerations suggest that such plots are not merely “difficult to interpret”, but are effectively meaningless. Unfortunately, this ensures that detecting the occurrence of cycling in either natural or artificial systems is problematic, since an external observer cannot determine, on the basis of relative fitness measures, whether a coevolutionary system is progressing, cycling, or drifting randomly.

In an attempt to circumvent this problem, Cliff and Miller (1995) proposed the “Current Individual versus Ancestral Opponent” (CIAO) plot as a visualization tool for detecting coevolutionary progress. At the conclusion of a coevolutionary run, CIAO plots can be constructed by pitting the elite (i.e., best-scoring) individual from every generation against the elite opponent from each ancestral generation and plotting the results as shaded cells in a matrix. In this way individuals are directly assessed against the ancestors of their opponents. If, over many generations, most individuals can beat their ancestral opponents, the matrix will exhibit a consistent gradation in shading from dark cells at the origin to light cells at the leading edge. This pattern suggests that there has been

continual progress over the course of coevolution, since individuals from later generations are outperforming their ancestors. By contrast, coevolutionary cycling manifests itself as a diagonal “banded” pattern (see section 2).

In coevolutionary research, CIAO plots are a widely accepted problem-independent visualization technique with few alternatives. However, in the studies reported here we explore whether they are as easy to interpret as has previously been implied. We demonstrate that CIAO plots can be misleading even in the simple coevolutionary domain of Rock–Paper–Scissors. We show that coevolutionary cycling can fail to produce a characteristic banded CIAO plot, resulting instead in a “tartan” pattern. This class of CIAO plot is commonly reported but has received little attention. Here we establish that tartan-like CIAO plots can result from cyclic coevolution that is *irregular*, or even random drift. We argue that CIAO plots are thus vulnerable to ambiguity and that, as a result, their use should be accompanied by more problem-specific analysis.

2 The Red Queen & CIAO Plots

Appearing in Lewis Carroll’s *Through The Looking Glass*, the Red Queen must continually run in order to maintain her position. No matter how fast she moves, the surrounding landscape always keeps up with her.

van Valen (1973) made an analogy between the Red Queen and biological coevolution after discovering a surprising trend concerning the probability of species extinction. After analyzing huge data sets collected across a wide range of biological taxa, van Valen (1973) noticed that, counter to intuition, “all groups for which data exist go extinct at a rate that is constant for a given group”. Assuming that species evolve in a relatively static environment, one would expect beneficial adaptations to accumulate over evolutionary time, enabling progressive generations to be better equipped at defending against extinction. To explain his findings, van Valen proposed a new evolutionary law of extinction, with the Red Queen Hypothesis as its central tenet: “biotic forces provide the basis for self-driving . . . perpetual motion of the effective environment”. Any beneficial adaptation by a particular species is inevitably detrimental to other species inhabiting the same effective environment. Coevolutionary forces will, in turn, select for specific counter-adaptations in these species. In this way, adaptive advantage is continually eroded. In Red Queen fashion, species continually struggle to maintain their relative fitness.

One might expect Red Queen dynamics to drive a run-away process of continual counter-adaptation—a coevolutionary *arms race*. Such arms races are considered a profound force driving evolutionary adaptation in the natural world, and have sometimes been characterized as a source of strong selection for novel adaptations capable of accelerating evolutionary progress (Dawkins & Krebs, 1979).

In order to take advantage of arms-race dynamics, the field of evolutionary computation has seized upon coevolution as an attractive alternative to standard evolutionary optimization. Living up to expectation, artificial coevolution has had success in several domains (e.g., Hillis, 1990; Juillé, 1995; Pollack, Blair, & Land, 1996; Pollack & Blair, 1998). However, as an optimization technique, coevolutionary search suffers from the relative nature of fitness assessment—coevolutionary systems can be difficult to drive in a consistent, objectively “progressive” direction.

There are several ways in which an arms race can unfold. One side may “win” the race, discovering an adaptation to which there is no available counter-adaptation. If coevolution is within a species, evolutionary stasis may be reached, but during between-species coevolution the disadvantaged species may be driven extinct. Alternatively, one population may temporarily win by “outstripping” the other to such an extent that the opponents are no longer discriminated by selection: whilst a counter-adaptation may exist, it has yet to be discovered. Termed “coevolutionary disengagement” (and discussed further in section 4.2), this type of outcome has been studied elsewhere (Watson & Pollack, 2001; Cartlidge & Bullock, 2002; Cartlidge & Bullock, to appear). Finally, an arms race may cycle, as populations follow repeated trajectories through strategy space, discovering strategies that enjoy only a temporary advantage over their opponents. From the perspective of coevolutionary optimization, cyclic coevolution has gained the most attention. Since these cyclic trajectories waste computational resources,

it would be useful to detect (and ultimately prevent) their occurrence; or, at least, to realize that the optimization problem as-stated has a possibly unanticipated set of unstable equilibria with no escape trajectory.

In general, one would like to know how coevolution is progressing. Has there been steady and continuous improvement, or transient bursts of progress amidst long periods of stasis? Is the system cycling, and, if it is, are the cycles regular? Frustratingly, the Red Queen renders standard approaches to answering these questions (by plotting individual fitness as it changes over time) obsolete.

Consider a predator-prey arms race. As predators *improve* over evolutionary time, one might expect them to catch more prey. The prey, however, are improving too. “There is no general reason to expect the average success of animals at out-running or out-witting *contemporary* enemies, victims, prey or competitors, to improve over evolutionary time” (Dawkins & Krebs, 1979, our italics). This is a direct consequence of the Red Queen. An improvement in any one species is countered by each coevolving species, resulting in a deterioration of the effective environment (van Valen, 1973).

However, one might expect a progressive coevolutionary arms race to result in an advantage for modern predators and prey over their *ancestral* adversaries. In nature, it is difficult to perceive how a competition across evolutionary time may arise, without the use of cloning or time-travel. Yet, in simulation it is quite feasible to carry out such ancestral opponent competitions. Cliff and Miller developed this technique during a series of papers in which they attempted to coevolve pursuer and evader strategies in continuous-time neural-network controllers (Miller & Cliff, 1994a, 1994b; Cliff & Miller, 1995, 1996).

Once a coevolutionary run has terminated, ancestral opponent contests (competitions against ancestral opponents) are carried out between the highest scoring individual of each generation: the “elite”. The elite of population A (the A-elite) is pitted against the elite of population B (the B-elite) in a series of contests. For each generation g of coevolution, the A-elite(g) competes with the B-elite of the current, and each ancestral, generation. Hence, A-elite(5) plays B-elite(5), B-elite(4), …, B-elite(0). The resulting scores of each contest are normalized and converted into gray-scale values and plotted on a 2-dimensional grid: the greater the victory in favor of the A-elite, the heavier the shading of the relevant matrix cell (see figure 1, adapted from Cliff & Miller, 1995).

Along with the invention of CIAO plots, Cliff and Miller left a lasting legacy—the idealized plot. Idealized CIAO plots demonstrate the patterns that they predicted one would find, given either perfect continual coevolutionary progress, or perfect coevolutionary cycling. Example schematics are shown in figure 2. Continuous progress appears as continuous diagonal *gradation* from dark to light across the plot. Each individual beats its ancestral opponents—the earlier the ancestor, the greater the victory. In contrast, coevolutionary cycling produces diagonal *banding*. Whilst recent ancestral opponents are easily beaten, elite opponents from a few generations earlier have the upper hand. As the contests span increasing periods of evolutionary time, competitive advantage repeatedly transfers between novel and ancestral populations—the system is cycling.

Since Cliff and Miller (1995) first introduced CIAO plots for pursuit and evasion, they have become widely accepted as a standard tool for visualizing coevolutionary progress. As such, CIAO plots have received extensive use in the fields of evolutionary robotics, (Floreano & Nolfi, 1997a, 1997b; Nolfi & Floreano, 1998; Stanley & Miikkulainen, 2002), the games of Go (Lubberts & Miikkulainen, 2001), 3-D Tic-Tac-Toe and Nim, (Rosin & Belew, 1997), and in the coevolution of string generators and predictors (Ficici & Pollack, 1998).

However, the results obtained in these relatively complex domains rarely resemble the idealized plots of Cliff and Miller. Even discounting noise, real CIAO plot visualizations are often qualitatively different from the ideal schematics. Many CIAO plots exhibit a tartan pattern—a patchwork of unpredictable lines and rectangles across the plot (figure 3 presents a hypothetical example). Of 22 plots found in the literature, 10 are tartan in nature, 8 show progress (smooth gradation) and 4 show no progress (a largely homogeneous plot). There are no examples of regular banding. In a tartan plot, blocks of uniform shade represent periods of stasis bounded by adaptive innovations. Consider Rock–Paper–Scissors, for example: if the elites in both populations play Rock from generation 50 onwards, until population A’s elite adopts Paper at generation 55, and the B-elite adopts

Scissors at generation 60, then the resulting CIAO plot will display a gray rectangle with sides 5×10 generations, cornered (50, 50).

Tartan CIAO plots have sometimes been interpreted as supporting evidence for cycling (Floreano & Nolfi, 1997a, 1997b; Nolfi & Floreano, 1998). The principle reason appears to be that, unlike the idealized progressive CIAO plot, there is little correlation between gray-scale value and time. It is difficult to imagine how a jumble of dramatic jumps between winning and losing can occur across evolutionary time in a continuously progressing coevolutionary system. How such a pattern relates to cycling, however, appears unclear. As we shall demonstrate, although tartan plots of this kind are consistent with coevolutionary change that is cyclic but irregular, it is dangerous to infer irregular cycling from them uncritically. In the following section we demonstrate that tartan plots are, in fact, ambiguous with respect to showing irregular cycling and/or random drift.

3 Rock–Paper–Scissors

In this section we explore coevolution (and resulting CIAO plots) in a simple domain specifically chosen to exhibit cyclic coevolutionary trajectories. To better understand the underlying coevolutionary dynamics, several different visualizations are presented. However, whilst each individually offers insight into the underlying coevolutionary dynamics, one cannot interpret them with confidence until information from the whole suite is considered. This is particularly true if some are ambiguous.

3.1 Study 1: A Simple Encoding

In this baseline study the inherent intransitivity of the Rock–Paper–Scissors (RPS) game, in combination with the smooth fitness landscape resulting from a simple genetic encoding, produces regular coevolutionary cycling. The visualizations that reflect this include CIAO plots that closely resemble the idealized banded plot (figure 2). For the most part this section introduces fairly straightforward results and visualizations which are intended to contrast with the more complex results from the subsequent studies.

For the simple encoding, each genome consists of three positive integers that sum to 100. Each integer represents the probability of choosing one of the three game moves, Rock (R), Paper (P), or Scissors (S). For example, genome {50, 0, 50} represents an individual that chooses to play Rock and Scissors with equal likelihood, but will never play Paper. Genome {33, 33, 34} represents a near-optimal individual that plays each move randomly with almost equal probability.

During reproduction, mutation occurs with probability 0.1 per locus. A mutated gene is incremented by an integer drawn at random from the uniform distribution $[-30 \dots 30]$. Following mutation, the genome is *normalized* to once again sum to 100.

Two reproductively isolated (initially randomized) populations, each containing 20 RPS players, are coevolved for 256 generations. Each generation, every individual is pitted against every member of the opponent population. Each game consists of 10 bouts, with each individual choosing R, P, or S probabilistically on the basis of their genes. Throughout a generation, an individual accumulates a score from its 200 competitions. Individuals reproduce asexually, with tournaments biasing selection in favor of high-scoring individuals. The winner of each randomly assembled 5-member tournament is chosen to reproduce.

3.1.1 Feature Detection

To assist the detection of features within the CIAO plots reported in this paper, a three stage processing of the raw images is employed. The method chosen (described below) is generalizable and easy to implement. Further, Marr (1982) has suggested that it may have biological implications as a model for the very first stages of visual processing. For further details, one should refer to Marr (1982, chapter 2).

1. **Gaussian Blur:** A 2-dimensional Gaussian filter (radius r) is initially applied to the image; effectively removing structures smaller than the standard deviation of the Gaussian distribution. As it is smooth and omnidirectional, a Gaussian distribution is appropriate for blurring as it is unlikely to introduce structure that was not present in the original image.
2. **Laplacian of Gaussian (LoG):** The Laplacian (∇^2) is an isotropic second-order differential operator that can be used to detect intensity changes in a blurred image, as seen at the scale of the particular Gaussian employed (determined by the standard deviation σ). The Laplacian operator is used to locate the zero-crossings of an image—changes from dark to light, or light to dark—thus generating a contour map.
In mathematical notation, the blur of an image function $I(x, y)$ with a Gaussian function G is denoted by $G * I$, read as “ G convolved with I ”. The Laplacian of this is denoted $\nabla^2(G * I) = (\nabla^2G) * I$, read as “the Laplacian of Gaussian (LoG) convolved with I ”.
3. **Binarize:** Finally, the direction of each zero-crossing is highlighted by “binarizing” the image. Areas with negative values are colored black whilst areas with positive values are left white.

Whilst stages 1 and 2 of the feature detection process naturally combine into one operation, throughout this paper they will be performed independently. This enables the result of blurring to be viewed before zero-crossings are detected. Throughout, r and σ values are chosen through trial and error: values are reduced to detect smaller features, for larger features they are increased.

3.1.2 Results

Throughout this section, all data is the result of one representative coevolutionary run. Figure 4 displays four CIAO plots. From left to right, plot 1 displays the raw data set, with subsequent CIAO plots presenting the same data after each stage of the image processing routine described in section 3.1.1. In plot 1, regular patterns and possible diagonal banding are discernable, potentially indicating regular coevolutionary cycling. After blurring the image with a large Gaussian of radius 20 pixels, small features are removed and diagonal banding becomes more pronounced (plot 2). A contour map is produced using a Laplacian of Gaussian (LoG) with standard deviation 8 pixels: diagonal banding is clear (plot 3). Plot 4 shows the result of binarizing the contour map in order to highlight the direction of any zero-crossings. The regularity and clarity of diagonal banding in the fully processed plot is striking. A comparison with the idealized banded CIAO plot of figure 2 allows plot 4 to be easily classified as demonstrating coevolutionary cycling with regular period: the feature detection routine has clearly enhanced the interpretability of the original CIAO plot data.

However, to have confidence in the nature of cycling, further visualizations were constructed. For this purpose, an “event plot” is used to highlight each occurrence of specific events during coevolution: evidence for (or against) cycling is generated by the resulting sequence of events. Assuming events are chosen adequately, regular cycling should manifest as a repeated sequence of events with fixed period. The event plot of figure 5 records the occurrence of events associated with the play of the best scoring individual of each population: the A-elite and B-elite. The plot is marked for each generation in which the A-elite plays P , or the B-elite plays R , with probability $p > 0.66$. Figure 5 clearly shows a repeated sequence. As expected, once B-elites regularly play Rock, A-elites quickly evolve to play Paper. After approximately 40 generations, the system completes a full cycle with B-elites once again predominantly playing Rock. The coevolutionary system appears to be following a cyclic trajectory with regular period. In general, if mutation rate or selection pressure (in the form of tournament size) increases, the cycle period decreases.

In figure 6, the elite (i.e., best-scoring) and average (the mean of all individuals) strategy of each population are depicted for a representative period of evolutionary time. The vertical axis represents the probability of playing each of the three possible moves. One full cycle takes place during the 45 generations depicted. At

generation 120, the elite member of population B (bottom-left) always plays Paper (genotype $\{0, 100, 0\}$). Within 5 generations, Scissors dominates the elite strategy in population A (top-left). As a counter-adaptation, population B converges on Rock, which in turn drives the elite of population A to play Paper with 90% probability (genotype $\{10, 90, 0\}$) by generation 140. This completes a half-cycle. Scissors becomes dominant for population B elites around generation 150, followed by Rock for population A, and finally the cycle is completed around generation 165 as the elite strategy of population B returns to Paper (genotype $\{3, 91, 6\}$). The mean population strategies (right) lag their elite counterparts slightly, but demonstrate the same trend. Notice that each population is able to counter-adapt to its opponents in a smooth, regular manner.

In summary, study 1 has (predictably) demonstrated that simple regular cycling can manifest itself on CIAO plots as cyclic banding, closely resembling the idealized banded CIAO plot in figure 2. In conjunction, it has also been demonstrated that greater insight into cyclic coevolutionary behavior can be obtained through the additional use of alternative, perhaps problem-specific, visualizations. In the following section, the complexity of the genetic encoding is increased to demonstrate that, whilst cycling may persist, the ease with which it is visualized is considerably reduced.

3.2 Study 2: A Complex Encoding

In this study a more complex genetic encoding scheme is employed. Although the strategic structure of the RPS game remains unchanged, by increasing the complexity of the strategies available to players, more complex courses of coevolutionary adaptation are available. It is intended that this new encoding will influence the system's coevolutionary dynamics, ensuring that the search for counter-adaptations more closely resembles that experienced in more realistic models of coevolutionary competition.

Here, genomes code for a deterministic finite state machine (FSM). Example genomes are presented in figure 7. The start node defines the choice of play during the first bout of a contest. In figure 7, both individuals initially play Scissors. From each node there are exactly three out-edges (which may be self-connecting), with each transition representing an opponent's move. In figure 7, since both individuals begin by playing Scissors, each follows the state transition S . A returns to the start node (0) and continues to play Scissors whilst B transfers to a new node (1) with state Paper. In the second bout of the game, A 's Scissors beats B 's Paper. A follows the P transition associated with B 's play and once again returns to state S (node 0). In response to A 's Scissors, B follows the S transition to a new node (2) with state Rock. B is victorious in bout three, as Rock beats Scissors. A moves to state P (node 1) and B remains at state R (node 2). Figure 7 details the full results of a five-bout contest.

Each node contains a unique integer identifier, a play state, three in-edges and three out-edges, each associated with an opponent's play and each connecting to a legal node. Both populations are initialized with random, self-connected, single-node FSMs. During reproduction, three mutation operators are employed (described below).

Node Mutation: With probability 0.03 per genome, a node is either added or removed. Any edges previously connected to a removed node become self-connections. FSMs were constrained to have between one and one hundred nodes. In practice, however, FSMs rarely grew above 10 nodes.

State Mutation: With probability 0.02 per node, a node state is mutated to one of the other two states, chosen at random.

Edge Mutation: With probability 0.02 per edge, edges are mutated by randomly changing either of the nodes to which the edge connects.

Note that the deterministic FSM encoding employed in this study (and study 3) does not allow individuals to reach the global optimum (available in study 1) of playing each move randomly with equal probability, irrespective of opponent play: to achieve this a non-deterministic FSM encoding is necessary. However, since we are not

interested here in a direct comparison between the results of studies 1, 2, and 3, this is not problematic. What we are interested in is the impact of regular and irregular cycling on CIAO plot data. As we will see, the deterministic encoding is a good choice in this respect.

Two reproductively isolated (initially randomized) populations, each containing 25 RPS players, are coevolved for 256 generations. Each generation, every individual is pitted against every member of the opponent population. Each game consists of 10 bouts, with each individual choosing R, P, or S deterministically on the basis of their genes. Throughout a generation, an individual accumulates a score from its 250 competitions. Individuals reproduce asexually, with tournament selection biasing selection in favor of high-scoring individuals. The winner of each randomly-assembled 3-member tournament is chosen to reproduce.

3.2.1 Results

Throughout this section, all data is the result of one representative coevolutionary run. Figure 8 displays four CIAO plots. From left to right, plot 1 displays the raw data set, whilst subsequent CIAO plots display the same data after each stage of the image processing routine described in section 3.1.1. Plot 1 (far left) presents a tartan pattern qualitatively similar to the schematic shown in figure 3: irregular blocks of uniform shading indicate periods of competitive stasis separated by rapid transitions in competitive advantage. Plot 1 resembles CIAO plots obtained from more complex (practical) problem domains (e.g., Floreano & Nolfi, 1997a, 1997b). The patchwork effect suggests a lack of progress and may imply cycling. However, the image is difficult to interpret with confidence.

After carrying out the same image-processing process employed in study 1, the plot displays some vertical banding and one diagonal band to the right of the image, but remains difficult to interpret. Whilst the vertical banding may be a result of bias in the original CIAO plot design—adaptive mutations will necessarily appear as banding on the horizontal (if A-elite) or vertical (if B-elite)—the black diagonal band to the right of the plot is potentially interesting. Perhaps the diagonal reflects a competitive advantage enjoyed by each population over its immediate ancestral opponents, or even regular cycling that can be observed over short evolutionary time scales, but that is lost over longer periods. Alternatively, perhaps the diagonal is an artifactual edge-effect of the CIAO plot: a result of the dark plot meeting a white background. In either case, whilst the processed plot has highlighted some potentially interesting structures, ideal banding is not exhibited. Considering the intransitive dominance relationships inherent within the RPS domain, however, one might suspect cycling to be taking place (particularly when the genetic encoding does not allow global stability to be reached). To investigate the underlying coevolutionary dynamics, further analysis is necessary.

The simplicity of the RPS domain enables us to gain insights into its coevolutionary dynamics by direct observation of genotypes and phenotypes as populations change over generations. At the beginning of a run, when each player's FSM has only one node, all strategies are simple. Each player repeatedly chooses the same play irrespective of its opponent's behavior. Over evolutionary time, more complex strategies arise. These multi-node FSMs change state in response to an opponent's play. However, as more time passes, simple strategies, such as “always play Scissors”, again begin to dominate. This process often repeats several times throughout a coevolutionary run.

For a coevolutionary system to be described as cycling, it must repeatedly enter the same or similar states in the same or similar order. Regular cycling repeats with fixed period, irregular cycling does not. Either class may fail to enter *exactly* the same set of states during each cycle, or maintain *exactly* the same ordering over these states during each cycle. However, each is clearly distinguishable from random drift. Whilst a randomly drifting system may return to previous states over evolutionary time, unless there is a heavy bias influencing the “random” walk (perhaps due to some bias within the genetic encoding or genetic operators, e.g., Bullock, 1999, 2001), the trajectory of change is unpredictable. Any useful method for detecting coevolutionary cycling should be able to distinguish cases of regular or irregular cycling from stochastic repetition that may arise due to random drift. As

we will show, when used in isolation, this is something CIAO plots struggle to do.

Choosing suitable visualizations for complex problems is often challenging. The space of variable length FSMs is difficult to represent graphically. In order to re-apply some of the graphing techniques used for the simple encoding scheme, some extra work is necessary.

FSM networks can be re-described simply using three probabilities: the likelihood of playing each of the three possible moves. Since each FSM is a directed graph, an m -bout contest against an opponent can be described by a path of length m through the graph. By traversing each of the 3^m possible m -length paths through the network—the equivalent of playing every possible game—the *effective* probability of a network playing each of the three possible moves can be determined. Since these probabilities are comparable to the three probabilities encoded in the simple genomes of study one, they can be plotted in a similar fashion. However, whilst the effective probabilities of FSMs are comparable to the probabilities of study 1, they are not equivalent: deterministic FSMs do not behave probabilistically but are sensitive to opponent play.

Figure 9 shows an event plot for study 2. The plot is marked for each generation in which the A-elite plays P, or the B-elite plays R, with effective probability $p > 0.66$. The system repeatedly moves through the same states in a similar sequence. “B-elite plays Rock” is always followed by “A-elite plays Paper”. However, on 2 out of 8 occasions, the system enters “A-elite plays Paper” without passing through “B-elite plays Rock”. The repeating event plot sequence provides some evidence for irregular cycling: if cycling is occurring, it is not as regular as that exhibited in figure 5.

Cycling can be directly observed by plotting the effective probabilities of each population’s elite and mean strategies during a particular (representative) coevolutionary period (see figure 10). Whilst it is not perfectly predictable, sequential cycling can be observed between generations 22-57. The B-elites (bottom-left) and B-means (bottom-right) complete one cycle with the expected sequence SPRS. In response, the A-elites (top-left) and A-means (top-right) also complete one cycle, but with the less predictable sequence RSRP. However, the “unexpected” shift from S to R (rather than directly to P) at generation 49 can be explained by considering the system as a whole at generation 48. With every individual in population A playing S (top-right) and every individual in population B playing R (bottom-right), the system is disengaged in generation 48: every member of population B is beating every member of population A. At this point, a novel mutation of *either* Paper or Rock is beneficial to population A: whilst Paper gives victory, Rock allows a draw. Thus, the move from Scissors to Rock is a direct result of disengagement and is not entirely “unexpected” (for a further discussion of disengagement, refer to section 4.2).

Comparing the graphs of elite strategies with those of the population mean strategies in figure 10 shows that each population exhibits several variable-length periods of strategic stasis. Typically, during these periods, the majority of population A (top) are being beaten by the majority of population B (bottom). Unlike in figure 6, population A appears unable to easily discover counter-adaptations to the successful adaptations of population B. This is a result of the deterministic FSM encoding and the limitations of the associated mutation operators. A desired mutation may be difficult to achieve. The initial 57 generations of study 2 are shown in the expanded CIAO plot of figure 11. The rapid changes in strategy depicted in figure 10 clearly reflect the internal structure of the CIAO plot. The patchwork nature appears because the FSM mutation operators tend to produce either no change in player strategy, or very rapid changes in behavior.

The different nature of the coevolutionary dynamics generated in study 2 most likely result from the ruggedness and/or neutrality introduced into the fitness landscape by the more complex genetic encoding scheme and associated genetic operators. Single mutations, such as “remove node”, can result in large changes in player behavior. By contrast, it can often be difficult to make a small change to a player’s behavior (e.g., increasing the tendency to follow Rock with Paper) without altering several parts of the FSM, which requires several separate mutation events. The simple genetic encoding scheme and associated mutation operator of study 1 ensured that single mutations tended to have modest phenotypic impact, yet most phenotypes were only a few mutations apart.

In general, for any population in study 2 some parts of strategy space will be less attainable than others to a larger degree than was the case in study 1. The resulting CIAO plot is difficult to interpret. However, alternative visualizations have shown that cycling is occurring, but that it is irregular.

3.3 Study 3: Random Drift

Do banded and tartan CIAO plots directly imply regular and irregular cycling, respectively? Can other coevolutionary trajectories lead to banded or tartan CIAO plots? In particular, what kind of CIAO plot is generated by the kind of walk through strategy space produced by random drift? As previously discussed, random drift can generate trajectories that revisit earlier states despite being non-cyclic. Can this kind of repetition be distinguished from regular behavior using CIAO plots?

Random drift is simulated for the same complex encoding scheme and mutation operators employed in study 2 (256 generations with 25 individuals in each population). Evolutionary selection pressures are removed, allowing each individual an equal chance of reproduction irrespective of ability. Individuals reproduce asexually and at random, irrespective of score. Figure 12 displays four CIAO plots resulting from typical random drift. From left to right, plot 1 displays the raw data set, whilst subsequent CIAO plots display the same data after each stage of the image processing routine described in section 3.1.1. Plot 1 (far left) presents a patchwork pattern of a tartan nature, but less regular than that shown in figure 8. The irregular blocks of shading suggest a lack of progress, but may imply cycling of some sort. The plot is difficult to interpret with confidence. After image processing, the binarized plot (far right) displays no diagonal banding. The most prominent features are the horizontal bands across the center of the image. It is likely that these bands are an artifact of the underlying grid structure of the raw plot, similar to the horizontal banding seen in the binarized plot of study 2 (figure 8). The most significant difference between the binarized plots of study 2 and 3 is the diagonal band to the right of figure 8 that is missing in figure 12. The binarized plot of random drift shows no sign of cycling (diagonal banding) even in the short term. However, whilst the processed plots highlight this feature, they still remain difficult to interpret. Can irregular cycling and random drift be correctly classified on this basis alone?

Figure 13 shows an event plot for the random drift of study 3. Once again, the plot is marked for each generation in which the A-elite plays P, or the B-elite plays R, with effective probability $p > 0.66$. Unlike the event plot of study 2 (figure 9), random drift shows no obvious sequence repetition. Events occur without regularity. The event plot strongly suggests that the system is not cycling.

Figure 14 displays the effective probabilities of each population's elite strategies during a particular (representative) period. The state transitions of each population's elites appears random, suggesting that neither cycles. Further, there is no obvious correlation between the two populations: knowing the state of one population does not improve our ability to predict the state of the opponent population. Thus, both populations vary *independently* and at random: it is *not* the case that one population changes randomly, with the other adapting to it in a regular way.

Random drift allows populations to return to previously evolved strategies, but not systematically—the system is *not* cycling. Despite this, random drift produces a CIAO plot with a tartan nature, not dissimilar to that produced by the irregular cycling of study 2. While we have not quantitatively measured the differences between the CIAO plots generated in studies 1, 2, and 3, the very fact that we might have to rely upon some such measure applied to CIAO plots in order to discover what they have to tell us about coevolutionary dynamics is disappointing. To anybody using CIAO plots as their only method of coevolutionary visualization, the results reported here are unfortunate. Not only can coevolutionary cycling result in tartan, rather than banded, plots, but these tartan patterns can also occur in the absence of any cycling at all. If used alone, CIAO plots are potentially ambiguous.

4 The Nature of Cycling

By highlighting the prevalence of cycling in natural systems, this section demonstrates that cycling, in general, is a ubiquitous phenomenon of great importance to the study of adaptive behavior: cycling is *not* confined to “toy” artificial coevolutionary systems, such as those studied above. More specifically, this section considers the likely prevalence of *irregular* (as opposed to *regular*) cycling. Using the previous RPS studies as evidence, particular attention is given to coevolutionary systems.

4.1 Cycling in Nature

Ever since coevolution emerged as an independent discipline in the 1970s, cyclic coevolutionary trajectories have been anticipated (see Maynard Smith, 1982; Futuyma & Slatkin, 1983). Dawkins and Krebs (1979) suggest that (in asymmetric systems) cycling may be common both *between* and *within* species. They illustrate this by highlighting the genetic model of parent-offspring conflict postulated by Parker and Macnair (1979) and Parker (1979). A dominant “conflictor” gene causes offspring to demand more investment than the parental optimum. This is countered by the spread of a “suppressor” gene in parents which causes them to invest equally in offspring, irrespective of demand. Assuming that the conflictor behavior has a cost, then the direction of selection on children is reversed once suppressor genes are frequent amongst parents: non-conflictor genes spread and the cycle starts again.

Dawkins and Krebs (1979) also present an example in the context of predator-prey coevolution. Each individual pays a cost of size: predators pay the cost of growing big enough to swallow prey and prey pay the cost of growing big enough to prevent being swallowed. Hence, there ensues a coevolutionary race as each species’ body size grows until their upper limit is reached; the size at which cost outweighs reward. Given the “life-dinner principle”,¹ prey are likely to invest more in getting bigger (and thus have a higher limit) than predators, which are equally likely to divert scarce resources towards other adaptations (sexual attractiveness, for example). The side that can afford to pay the highest cost will do so; at which point selection will favor a rapid reduction in the cost paid by the other side. Any return to the start state may result (after some possibly unpredictable period of time) in a repeated bout of escalation. Maynard Smith (1996) describes such cycles as “sawtooth” oscillations and suggests that they can occur in general when a variable, such as size, can vary continuously with no ESS: hence, size may increase gradually until a threshold is reached, when the population can be invaded by much smaller individuals.

Sawtooth oscillations have been observed in *Anolis* lizard populations of the Caribbean islands (Roughgarden, 1983). The insects on which *Anolis* lizards feed have food value proportionate to their size (the bigger the better, as long as they can be swallowed), and abundance inversely proportional to size (smaller insects are more common). A relationship between lizard size and insect size exists such that larger lizards on average take larger prey. On all islands inhabited by only one species, lizards have an equilibrium “solitary size” of approximately 50mm length in females and 60mm length in males. Where two species exist, however, one population is generally much larger than the solitary size whilst the other is *always* smaller. Where two species compete, it is assumed that the larger species exerts a stronger pressure on the smaller species than *vice versa*: large lizards take more food away; disputes over territory favor large lizards. To fit this data, Roughgarden (1983) proposed the “coevolution-invasion turnover hypothesis”, suggesting that through a series of invasions and extinctions, the lizard populations cycle in body length (see figure 15, adapted from Roughgarden, 1983). On islands containing only one species, lizard lengths approximate the optimum solitary size; because of the asymmetry in competition, only a larger species can invade (figure 15.A). After invasion, the species coevolve as competitors. The resident’s body size reduces to avoid competition from the invader; an example of character displacement. The invader’s body size

¹The life-dinner principle states that coevolving species may be subject to asymmetric evolutionary pressures. For instance, consider foxes and rabbits: “The rabbit runs faster than the fox, because the rabbit is running for his life while the fox is only running for his dinner” (Dawkins & Krebs, 1979, after Aesop).

also reduces to take advantage of the greater resources left by the retreating residents (figure 15.B). Finally, the resident species is driven to extinction by competitive exclusion. The invaders approach the optimum solitary size and the system has completed one full cycle (figure 15.C).

Since the example of the *Anolis* lizards requires species replacement (an intraspecific cycle of the same nature would require invasion by an implausibly large mutant), Maynard Smith (1996) suggests that the first example of a population cycle arising from intraspecific interactions is that reported relatively recently in side-blotched lizards (Sinervo & Lively, 1996). Side-blotched lizards exhibit three alternative male mating strategies, each associated with a distinctive phenotypic trait: blue-throated males mate-guard females and are territorial; yellow-throated “sneaker” males are non-territorial and roam about freely, looking to copulate with the females of others; and aggressive orange-throated males are polygynous and maintain large territories. Whilst blue-throated lizards avoid cuckoldry by yellow-throated sneakers, they are easily overpowered by orange-throated males, which cosire offspring with their females. Yellow-throated males are able to sire offspring via secretive copulations with the females of orange-throated males and often share paternity of offspring within a female’s clutch (Zamudio & Sinervo, 2000). Sinervo and Lively (1996) showed that the frequencies of the three male morphs were found to oscillate over a six year period.

The fitnesses of each morph relative to other morphs were non-transitive in that each morph could invade another morph when rare, but was itself invadable by another morph when common. Concordance between frequency-dependent selection and the among-year changes in morph fitnesses suggest that male interactions drive a dynamic ‘rock-paper-scissors’ game (Sinervo & Lively, 1996).

Using quantitative measures of the reproductive success of males adopting each strategy, Zamudio and Sinervo (2000) confirmed that the morphs were indeed playing RPS. The relative fitness of each strategy during dyadic interactions confirmed this.

Intransitive RPS dynamics have also been observed in populations of the bacteria *Escherichia coli* (Kerr, Riley, Feldman, & Bohannan, 2002). Colicinogenic bacteria (C) possess a ‘col’ plasmid, a toxin that kills colicin-sensitive (S) bacteria. A third strain (R) is resistant to the colicinogenic bacteria. In some cases, the growth rate of R cells will exceed that of C cells, but be less than the growth rate of S cells.

In such cases, S can displace R (because S has a growth-rate advantage), R can displace C (because R has a growth-rate advantage) and C can displace S (because C kills S). That is, the C-S-R community satisfies a rock-paper-scissors relationship (Kerr et al., 2002).

Confirming the predictions of their simulation model, empirical observations of *E. coli* (constrained to local interactions in a petri dish) demonstrated the cyclical coexistence of all three strains, with R outperforming C, C outperforming S and (a suggestion of) S chasing R across the plate.

Using the public goods game (applicable to theoretical biology), cycling has also been demonstrated in experimental economics; both in theory (Hauert, De Monte, Hofbauer, & Sigmund, 2002) and in practice (Semmann, Krambeck, & Millinski, 2003).

Semmann et al. (2003) demonstrate that voluntary participation in the game can lead to cooperation amongst sizable groups; despite anonymity, random assortment and non-repetition of interactions. Three strategies exist within a population: defectors (D) and cooperators (C), both willing to engage in the public goods game and speculate on the success of a joint enterprise; and low-risk loners (L) who choose to reject participation and settle for a small, but guaranteed, return. From time to time, groups of individuals are offered the choice to compete in a public goods game; loners will always refuse.

In every group, defectors outperform cooperators, but if all cooperate, they are better off than if all defect. Whilst it is better to be a loner than in a group of defectors, it is better still to be in a group of cooperators. Hence, in a well-mixed population, strategies display a RPS cycle. If most play C, then it is better to play D, but if most play D, then it is better to play L. However, if most play L, then small groups can form, increasing the chance of

mutual cooperation. Thus, C dominates if group size is small, D dominates if group size is large, and the option to be a loner preserves a balance between the two options: the system cycles (Hauert et al., 2002). An empirical study involving 280 students playing a 57-round strategy game confirmed these results (Semmann et al., 2003).

4.2 Irregular Cycling

The three RPS studies of section 3 demonstrated that although one particular kind of coevolutionary cycling is easily detected using CIAO plots, a second class of cyclic behavior is much harder to detect without resorting to alternative visualizations. Whilst CIAO plots *can* give valuable insights into coevolutionary dynamics, they should preferably be used amongst a suite of techniques in order to enhance their interpretability.

In general, these results contribute to a growing realization that our understanding of coevolutionary dynamics in artificial systems is far from complete. More specifically, their significance hinges, to some extent, on how important irregular coevolutionary cycling turns out to be. Here we argue that there are good reasons to suppose that this class of dynamic will be more frequently encountered than regular cycling, and that for many kinds of interesting system, when these irregular dynamics are exhibited, they will often be of both theoretical and practical significance.

First, as evidenced by the different results of studies 1 and 2, as search problems become increasingly complex, their search spaces are increasingly structured by the genetic encodings and genetic operators employed, in addition to the strategic advantages of different phenotypes. This ensures that some (perhaps most) adaptations will only be discovered after a period of evolutionary exploration. The stochastic nature of this exploration coupled with the rugged and/or neutral structure of the search space ensures that the time that this takes is variable. Under such conditions, if a coevolving system finds itself cycling, and hence repeatedly rediscovering the same or similar adaptations in the same or similar order, there is no guarantee that the period of coevolutionary cycling will be constant—indeed it is likely not to be. For these reasons, regular cycling should be regarded as a rarely attainable special case of the more general class of irregular cycling. Perhaps unsurprisingly, of the 22 CIAO plots found in the literature, none display regular cycling (Cliff & Miller, 1995; Floreano & Nolfi, 1997a, 1997b; Rosin & Belew, 1997; Ficici & Pollack, 1998; Nolfi & Floreano, 1998; Lubberts & Miikkulainen, 2001; Stanley & Miikkulainen, 2002).

Second, many coevolutionary phenomena currently of interest to adaptive behavior researchers are characterized by irregular cycling. Neutrality and disengagement are each discussed below.

Search-space neutrality occurs when many genotypes share the same fitness, perhaps as a result of redundancy in the genetic encoding. A neutral set contains all the genotypes that achieve a particular fitness score, while a neutral network comprises evolutionarily adjacent members of a neutral set. It has been argued that the presence of neutral networks may have profound consequences for the dynamics of evolutionary search. For instance, the neutrality exhibited by natural RNA search spaces has been demonstrated to be of a potentially useful kind, allowing more efficient search (Huynen, Stadler, & Fontana, 1996). More generally, neutrality of the right kind is thought to reduce the chance of premature convergence (Harvey & Thompson, 1996; Barnett, 1998). However, since evolving populations tend to drift at random across neutral networks (but see Bullock, 2002, for analysis of the biases that this drift is subject to) it is difficult to predict how long a population will spend on each one. If a coevolutionary system cycles through a repeated sequence of neutral networks each population will spend some time drifting across each neutral network, before transitioning to the next. As such, although a particular sequence of phenotypes may be generated over evolutionary time, this repetition is unlikely to exhibit a constant period.

Coevolutionary disengagement occurs in a competitive coevolutionary system when one population *outperforms* the other to the extent that different individuals are not discriminated from their contemporaries in fitness terms, i.e., floor or ceiling effects (Watson & Pollack, 2001; Cartlidge & Bullock, 2002; Cartlidge & Bullock, to appear). For instance, in the coevolution of pursuit and evasion, disengagement could occur if evaders discover a simple hiding strategy that defeats all contemporary opponent pursuers—each hiding evader would score 100%,

while all opponent pursuers would score 0%, despite variation in their strategies. When disengagement occurs, selective pressure disappears, leaving populations free to drift until such time as populations happen to re-engage. Re-engagement takes place when mutant strategies arise that achieve distinctive fitness scores (e.g., a pursuer able to discover hidden evaders, or an evader unable to hide successfully). The time taken for such mutants to arise via neutral drift, and hence the duration of disengagement, is variable for the reasons described above. As a result, a cycling coevolutionary system suffering from disengagement is also likely to cycle with non-fixed period. In short, irregular cycling is likely to be evolutionarily typical because useful evolutionary innovations and counter-innovations are not typically discovered at a constant rate.

The data collected during study 2 exhibits these variable periods of coevolutionary disengagement and neutral drift (in coevolutionary systems the two concepts are closely related). At generation 23, for example, figure 10 shows the elite strategy in population one plays Rock, whilst the elite strategy of population two plays Paper. In contrast with the immediate and smooth coevolutionary responses reported for study 1, it is not until generation 37 that the elite strategy of population one discovers the Scissors counter-adaptation. The intervening period is one in which the populations have disengaged, and are drifting across neutral networks of equivalent strategies. Disengagement can be directly observed in figure 16, where mean fitness (left) and fitness variance (right) in both populations is plotted over time. During generations 28-26 and 44-49, fitness diversity is very low and occasionally falls to zero (right): here population B is easily outperforming population A (left).

In the following section, CIAO plots are used (for the first time, as far as the authors are aware) to analyze simulation data from a biological model: a replication of the Rock–Paper–Scissors *E. coli* experiments of Kerr et al. (2002, refer to section 4.1). It is demonstrated that CIAO plots, whilst specifically developed for evolutionary computation, may benefit evolutionary biology (and the modeling of adaptive behavior more generally).

5 Study 4: A Simulation of *E. coli*

Rock–Paper–Scissors cycling has been demonstrated in *E. coli* populations (Kerr et al., 2002). However, practical problems forced experiments on real bacteria to end after short time periods. To collect more “data” (and gain a better understanding of the coevolutionary dynamics) Kerr et al. (2002) ran a simulation. This was shown to behave realistically and suggested that RPS dynamics characterized the system. However, it is our belief that more effective visualization techniques could have afforded greater insight into the underlying dynamics. To test this, the simulation of Kerr et al. (2002) is replicated in this section. The resulting data is then investigated using the visualizations discussed in this paper (CIAO plots, event plots and probability graphs). The reason for this is two fold: firstly, the lessons learned in previous studies can be applied to a biological system; secondly, this is a proof by example that techniques designed for evolutionary computing can have relevance in evolutionary biology.

5.1 Replication

Three strains of *E. coli* exist on a (toroidal) lattice grid of 100^2 cells. At the start of each run, each cell is initialized at random (equal probability) to one of four states: occupied by C, S, R, or “empty”. Cells are asynchronously updated. A focal cell, c , is randomly chosen and updated probabilistically based upon local interactions (the relative states of the 8 nearest neighboring cells). If c is empty a bacteria of strain $i \in \{C, S, R\}$ is chosen to occupy the cell. The probability of choosing i is given by f_i , the fraction of the local neighborhood occupied by each strain. If c is occupied, the bacteria is killed with probability Δ_i . Throughout the simulation, Δ_C and Δ_R are fixed. However, Δ_S varies with f_C (the fraction of neighboring C cells) such that $\Delta_S = \Delta_{S,0} + \tau f_C$, where $\Delta_{S,0}$ is the probability of death for S cells with no neighboring C cells, and τ is the toxicity of C cells.

To set up a Rock–Paper–Scissors intransitivity, it is necessary for $\Delta_{S,0} < \Delta_R < \Delta_C < \Delta_{S,0} + \tau$, which ensures S displaces R, R displaces C, and C (if sufficiently dense) displaces S. Following Kerr et al. (2002), the

following parameter values were chosen: $\Delta_C = \frac{1}{3}$; $\Delta_{S,0} = \frac{1}{4}$; $\Delta_R = \frac{10}{32}$; and $\tau = \frac{3}{4}$.

An “epoch” is defined as the mean turnover time across all cells. In the 100^2 lattice used here, an epoch occurs every 10^4 updates. We consider each epoch to be 1 timestep. The *E. coli* simulation was run for 5000 timesteps.

5.2 Results

Adopting the visualization scheme employed by Kerr et al. (2002), figure 17 shows four instantaneous “snapshots” taken of the simulation at timesteps (from left to right) 1000, 1200, 1400, and 1600. The state of each cell in the lattice is denoted by its color. Whilst empty cells are white, those occupied by C, R, or S strains are colored light gray, dark gray, and black, respectively. It is possible to observe areas of black replaced by light gray, light gray by dark gray, and dark gray by black. The bacteria are following the relationship S>R>C>S (analogous to RPS). Between timesteps 1000-1600, one full cycle takes place: the system returns to a similar state.

These results qualitatively map those of Kerr et al. (2002) who showed the same succession sequence both in simulation and in real bacteria populations. Whilst the continued coexistence of all three strains and the S>R>C>S relationship are adequately demonstrated, the visualizations are not wholly satisfying. One cannot determine the nature of cycling from a series of snapshots.

To further visualize the *E. coli* data, the local density of each strain was recorded for one 5×5 portion of the lattice. Since the three strains coexist without large fluctuations in global density, plotting CIAO plots for the entire lattice is uninformative. Due to the Rock–Paper–Scissors intransitivity of *E. coli* strains, each local density can be considered equivalent to the probability of playing a particular move in RPS. For example, equal densities of each strain $\{\text{Col} = 33\frac{1}{3}, \text{Resistant} = 33\frac{1}{3}, \text{Sensitive} = 33\frac{1}{3}\}$ is equivalent to the optimum RPS strategy of playing each move with equal probability, whilst a neighborhood containing only Resistant bacteria $\{0, 100, 0\}$ is equivalent to the strategy “always play Paper”. Figure 18 shows four CIAO plots of the *E. coli* simulation. The CIAO plot data is calculated by comparing the local density at each timestep against the local density at each previous timestep. To reduce the size of the CIAO plot, data is sampled every 50 timesteps (hence the plot is 100 pixels wide and deep). Each density comparison is evaluated as the expected result in a Rock–Paper–Scissors contest: results in favor of the later timestep (y axis) are dark. For example, if local density $D_1 = \{100, 0, 0\}$ is compared with the *later* local density $D_2 = \{50, 50, 0\}$, the expected result will be 0.75 in favor of D_2 ; thus giving pixel (D_1, D_2) a gray value 75% of maximum darkness.

From left to right, the CIAO plots display the raw data, Gaussian blur, zero-crossings, and fully processed image. The raw data plot (far left) exhibits both diagonal banding and tartan structures. The binarized image (far right) is much easier to interpret. The clear diagonal banding suggests that the system is exhibiting fairly regular cycling in this region of the lattice. The vertical “fault” line down the center of the plot is an artifact of sampling: without sampling this feature does not exist. However, unsampled CIAO plots are too large to appear here.

The event plot of figure 19 highlights the points (sampled at a 50-timestep resolution) at which the local density of strain R, or strain C, exceeds 66%. The sequence “Resistance follows Col toxic” occurs regularly enough to again suggest fairly regular cycling in this region of the lattice.

Finally, figure 20 displays the local density of each strain over the first 1000 timesteps (plotted at a 50-timestep resolution). Regular cycling is clear.

In summary, a suite of visualization techniques, including (for the first time) CIAO plots, have enabled us to explore the nature of the cycling suggested in the initial system snapshots of a biological model. In particular, we have been able to confirm the regularity of cycling and determine the period of this cycling. The image processing techniques introduced in section 3.1.1 have been very useful in achieving this.

Given the discussion in section 4.2 of reasons why we would not expect regular cycling to be exhibited particularly often by coevolutionary systems, why is it that our model of *E. coli* coevolution generates such regular, periodic behavior? There are two candidate explanations. First, the *E. coli* simulation employs a very simple representation of the space of possible strategies. Even the simple RPS system analyzed in study 1 employed

a discrete three-dimensional space of 100^3 possible strategies, whereas the *E. coli* system implements what is basically a simple set containing three strategies. There is no equivalent of the genotype space employed in our previous studies.

However, even if the *E. coli* model involved a more complex treatment of genotype space, coevolutionary cycles could remain regular, since the three bacterial morphs are typically always present in the evolving population. While the strategy being played in each cell of the 100×100 world varies over time, the global frequency of each strategy remains non-zero and roughly constant.

The update rules for the model ensure that if, for instance, all S individuals were to die, this strategy could never re-enter the population—only strategies that are already present may reproduce and spread to unoccupied cells since the model includes no analogue of mutation. Since, for the particular parameter values reported here, all three possible strategies are ever present, the system never spends time searching genotype space for adaptations and counter-adaptations; they are already present somewhere in the population. The cycling evident in the population time-series data reflects the time taken for a strategy to “migrate” across the grid, exploiting inferior competitors. At the relatively high population densities employed in the model, this leads to regular, periodic waves of succession.

The game-theoretic Evolutionarily Stable Strategy models of evolutionary dynamics typically employed in theoretical biology (Maynard Smith, 1982) also have no explicit representation of genotype space or mutation. In effect, such models assume that every possible strategy is always present at some non-zero (but perhaps infinitesimal) frequency within the population. All that varies in such models are the rates of reproduction enjoyed by these strategies. This ensures that the factors here identified as responsible for the irregular period of coevolutionary cycling (disengagement, neutrality, etc.) are not typically considered. The very different nature of the finite (co)evolving populations simulated in the adaptive behavior literature offers an important opportunity to explore these factors.

6 Conclusions

Cycling occurs in many adaptive systems. Mutational stochasticity, neutrality, disengagement, and rugged fitness landscapes each contribute to the irregularity of these cycles. As a result, irregular cycling may characterize the adaptive behavior of many unstable coevolutionary systems and may contribute to the failure of coevolutionary optimization. Detecting and characterizing this type of dynamic is difficult within modeling paradigms that privilege stable states (e.g., evolutionary game theory). As a result, coevolutionary adaptive behavior simulation and visualization are good candidates for improving our understanding of these types of irregular cycling.

We have demonstrated that CIAO plots, a widely accepted tool for visualizing coevolutionary progress, are difficult to interpret with respect to irregular coevolutionary cycling. Further, we claim that there are reasons to believe that irregular coevolutionary cycling is a common and significant category of coevolutionary dynamics. Hence, CIAO plots should preferably be used in conjunction with other visualizations: problem-specific analysis methods can usefully complement CIAO plots and can aid in their interpretation.

Finally, we have applied CIAO plots to biological simulation data, affording valuable insights into the underlying coevolutionary dynamics that agree with empirical observations. This exemplifies the potential cross-over between techniques designed for evolutionary computation and adaptive behavior more generally.

Acknowledgements

The authors would like to thank Steven Walker for preliminary studies and also Geoffrey Miller, Peter Todd, Dave Cliff and an anonymous reviewer for useful comments and criticisms on earlier versions of this paper. Thanks also to Tom Barnes-Lawrence and Neil Currum for help with some figures.

References

Barnett, L. (1998). Ruggedness and neutrality — the NKp family of fitness landscapes. In Adami, C., Belew, R., Kitano, H., & Taylor, C. (Eds.), *Artificial Life VI*, pp. 18–27. MIT Press, Cambridge, MA.

Bullock, S. (1999). Are artificial mutation biases unnatural? In Floreano, D., Nicoud, J.-D., & Mondada, F. (Eds.), *Fifth European Conference on Artificial Life*, pp. 64–73. Springer, Berlin.

Bullock, S. (2001). Smooth operator? Understanding and visualising mutation bias. In Kelemen, J., & Sosík, P. (Eds.), *Sixth European Conference on Artificial Life*, pp. 602–612. Springer, Berlin.

Bullock, S. (2002). Will selection for mutational robustness significantly retard evolutionary innovation on neutral networks? In Standish, R., Bedau, M., & Abbass, H. (Eds.), *Artificial Life VIII*, pp. 192–201. MIT Press, Cambridge, MA.

Cartlidge, J., & Bullock, S. (2002). Learning lessons from the common cold: How reducing parasite virulence improves coevolutionary optimization. In Fogel, D. (Ed.), *Congress on Evolutionary Computation*, pp. 1420–1425. IEEE Press.

Cartlidge, J., & Bullock, S. (to appear). Combating coevolutionary disengagement by reducing parasite virulence. *Evolutionary Computation*, 12(2).

Cliff, D., & Miller, G. F. (1995). Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In Morán, F., Moreno, A., Merelo, J. J., & Chacón, P. (Eds.), *Third European Conference on Artificial Life*, pp. 200–218. Springer, Berlin.

Cliff, D., & Miller, G. F. (1996). Coevolution of pursuit and evasion II: Simulation methods and results. In Maes, P., Matarić, M., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.), *Fourth International Conference on Simulation of Adaptive Behavior*, pp. 506–515. MIT Press, Cambridge, MA.

Dawkins, R., & Krebs, J. R. (1979). Arms races between and within species. *Proceedings of the Royal Society of London, Series B*, 205, 489–511.

Ficici, S. G., & Pollack, J. B. (1998). Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In Adami, C., Belew, R. K., Kitano, H., & Taylor, C. (Eds.), *Artificial Life VI*, pp. 238–247. MIT Press, Cambridge, MA.

Floreano, D., & Nolfi, S. (1997a). Adaptive behaviour in competing co-evolving species. In Husbands, P., & Harvey, I. (Eds.), *Fourth European Conference on Artificial Life*, pp. 378–387. MIT Press, Cambridge, MA.

Floreano, D., & Nolfi, S. (1997b). God save the Red Queen! Competition in co-evolutionary robotics. In Koza, J. R., Deb, K., Dorigo, M., and Max Garzon, D. B. F., Iba, H., & Riolo, R. L. (Eds.), *Proceedings of the Second Annual Conference on Genetic Programming*, pp. 398–406. Morgan Kaufmann, , San Mateo, CA.

Futuyma, D. J., & Slatkin, M. (1983). *Coevolution*. Sinauer Associates, Sunderland, MA.

Harvey, I., & Thompson, A. (1996). Through the labyrinth evolution finds a way: A silicon ridge. In Higuchi, T., Iwata, M., & Wexlin, L. (Eds.), *Frist International Conference on Evolvable Systems*, pp. 406–422. Springer, Berlin.

Hauert, C., De Monte, S., Hofbauer, J., & Sigmund, K. (2002). Volunteering as Red Queen mechanism for cooperation in public goods games. *Science*, 296, 1129–1132.

Hillis, W. D. (1990). Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D*, 42, 228–234.

Huynen, M., Stadler, P., & Fontana, W. (1996). Smoothness within ruggedness: The role of neutrality in adaptation. *Proceedings of the National Academy of Science USA*, 93(1), 397–401.

Juillé, H. (1995). Incremental co-evolution of organisms: A new approach for optimisation and discovery of strategies. In Morán, F., Moreno, A., Merelo, J. J., & Chacón, P. (Eds.), *Third European Conference on Artificial Life*, pp. 246–260. Springer, Berlin.

Kerr, B., Riley, M. A., Feldman, M. W., & Bohannan, B. J. M. (2002). Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. *Nature*, 418(11), 171–174.

Lubberts, A., & Miikkulainen, R. (2001). Co-evolving a go-playing neural network. In Belew, R. K., & Juillé, H. (Eds.), *Proceedings of the GECCO-01 Workshop on Coevolution: Turning Adaptive Algorithms upon Themselves*, pp. 14–19.

Marr, D. (1982). *Vision*. Freeman, San Francisco, CA.

Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.

Maynard Smith, J. (1996). The games lizards play. *Nature*, 380(6571), 198–199.

Miller, G. F., & Cliff, D. (1994a). Co-evolution of pursuit and evasion I: Biological and game-theoretic foundations. Technical report CSRP311, University of Sussex, School of Cognitive and Computing Sciences.

Miller, G. F., & Cliff, D. (1994b). Protean behaviour in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. W. (Eds.), *Third International Conference on Simulation of Adaptive Behavior*, pp. 411–420. MIT Press, Cambridge, MA.

Nolfi, S., & Floreano, D. (1998). How co-evolution can enhance the adaptive power of artificial evolution: Implications for evolutionary robotics. In Husbands, P., & Meyer, J.-A. (Eds.), *First European Workshop on Evolutionary Robotics*, pp. 22–38. Springer, Berlin.

Parker, G. (1979). Sexual selection and sexual conflict. In Blum, M. S., & Blum, N. B. (Eds.), *Sexual Selection and Reproductive Competition in Insects*, pp. 123–166. Academic Press, New York.

Parker, G., & Macnair, M. (1979). Models of parent-offspring conflict IV. Suppression: evolutionary retaliation by the parent. *Animal Behaviour*, 27, 1210–1235.

Pollack, J., Blair, A., & Land, M. (1996). Coevolution of a backgammon player. In Langton, C. G., & Shimohara, T. (Eds.), *Artificial Life V*. MIT Press, Cambridge, MA.

Pollack, J. B., & Blair, A. D. (1998). Co-evolution in the successful learning of a backgammon strategy. *Machine Learning*, 32(3), 225–240.

Rosin, C. D., & Belew, R. K. (1997). New methods for competitive coevolution. *Evolutionary Computation*, 5(1), 1–29.

Roughgarden, J. (1983). Coevolution between competitors. In Futuyma, D. J., & Slatkin, M. (Eds.), *Coevolution*, pp. 383–403. Sinauer Associates, Sunderland, MA.

Semmann, D., Krambeck, H.-J., & Millinski, M. (2003). Volunteering leads to rock-paper-scissors dynamics in a public goods game. *Nature*, 425(6956), 390–393.

Sinervo, B., & Lively, C. (1996). The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 380, 240–243.

Stanley, K. O., & Miikkulainen, R. (2002). The dominance tournament method of monitoring progress in coevolution. In Barry, A. M. (Ed.), *GECCO 2002: Proceedings of the Bird of a Feather Workshops, Genetic and Evolutionary Computation Conference*, pp. 242–248. AAAI.

van Valen, L. (1973). A new evolutionary law. *Evolutionary Theory*, 1, 1–30.

Watson, R. A., & Pollack, J. B. (2001). Coevolutionary dynamics in a minimal substrate. In Spencer, L., Goodman, E. D., Wu, A., Langdon, W., H., Gen, M., Sen, S., Dorigo, M., Pezeshk, S., Garzon, M. H., & Burke, E. (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 702–709. Morgan Kaufmann, , San Mateo, CA.

Zamudio, K. R., & Sinervo, B. (2000). Polygyny, mate-guarding, and posthumous fertilization as alternative male mating strategies. *Proceedings of the National Academy of Sciences USA*, 97(26), 14427–14432.

List of Figures

- 1 Schematic of a “Current Individual versus Ancestral Opponents” (CIAO) plot. At each square (x, y) the result of competition between the A-elite from generation y and the B-elite from generation x is plotted as a gray-scale value, with increasingly heavy shading representing an increasing margin of victory in favor of A. The leading diagonal (far-right) plots A-elite(g) against B-elite(g) for all generations, $0 \leq g < N$. The diagonal immediately to the left plots each generation’s A-elite against the *previous* generation’s B-elite. Horizontal rows plot the results of A-elite(y) against *all* ancestral B-elites($0 \dots y$) (adapted from Cliff & Miller, 1995). 21
- 2 Idealized CIAO plots showing smooth progress (left) and cycling with regular period (right). With each horizontal row depicting the performance of one A-elite against each ancestral B-elite, smooth progress produces a gradation in intensity from dark to light: A-elites beat ancestral opponents, the more ancestral the greater the victory (i.e., the further left, the darker the cell). In contrast, regular cycling manifests as diagonal banding. Whilst the latest A-elites beat recent ancestral B-elites (diagonals to the right are dark), they perform less-well against more ancestral B-elites (middle diagonals are light), but outperform even earlier ancestors (diagonals further left are dark), etc. 22
- 3 Schematic of a “tartan” CIAO plot. Periods of competitive stasis (blocks of uniform shade) are separated by sharp transitions in competitive advantage (block boundaries). Such an irregular pattern is referred to as “tartan” throughout this paper. 23
- 4 Study 1 results. CIAO plots depicting one representative run over $N = 256$ generations. From left to right: (1) Diagonal banding in the raw data plot suggests regular cycling; (2) A large Gaussian ($r = 20$) removes fine detail; (3) LoG ($\sigma = 8$) produces a contour map; (4) The image is binarized. The fully processed plot displays clear diagonal banding. 24
- 5 Study 1 results. Event plot for a single representative run, labeling the generations in which the A-elite plays Paper or the B-elite plays Rock with probability 0.66 or greater. The repeated “Paper follows Rock” event sequence suggests regular cycling. 25
- 6 Study 1 results. Graphs plotting the probability values of playing each RPS move. Top-left: the A-elites. Top-right: the mean probabilities for all individuals in population A. Bottom-left: the B-elites. Bottom-right: the mean probabilities for all individuals in population B. The system completes one full cycle between generations 120 and 165. 26
- 7 Complex genotypes encode a deterministic finite state machine (FSM) that governs player behavior. State transitions are dictated by an opponent’s play (represented by transition labels). Here we see that machine *A* beats machine *B* over a five-bout contest. The final score is 3:2 in favor of A. As both machines have returned to their initial states, further bouts will produce repeated play. 27
- 8 Study 2 results. CIAO plots depicting one representative run over $N = 256$ generations. From left to right: (1) The raw data plot exhibits a tartan pattern that is difficult to interpret; (2) A large Gaussian ($r = 20$) removes fine detail; (3) LoG ($\sigma = 8$) produces a contour map; (4) The image is binarized. Whilst the fully processed plot does not exhibit diagonal banding, some vertical banding is clear. However, one cannot predict regular cycling from this image. 28
- 9 Study 2 results. Event plot for a single representative run, labeling the generations in which the A-elite plays Paper or the B-elite plays Rock with effective probability 0.66 or greater. The strong correlation between events suggests (perhaps irregular) cycling. 29
- 10 Study 2 results. Graphs plotting the effective probabilities associated with playing each move. Top-left: the A-elites. Top-right: the mean probabilities for all individuals in population A. Bottom-left: the B-elites. Bottom-right: the mean probabilities for all individuals in population B. The populations evolve in response to each other, resulting in an irregular cycle. 30

11	Enlarged CIAO plot depicting the coevolutionary period graphed in figure 10. The rapid strategy transitions depicted in figure 10 clearly reflect the boundaries in the CIAO plot pattern.	31
12	Study 3 results. CIAO plots depicting one representative run over $N = 256$ generations. From left to right: (1) The tartan pattern of the raw data plot is difficult to interpret; (2) A large Gaussian ($r = 20$) removes fine detail; (3) LoG ($\sigma = 8$) produces a contour map; (4) The image is binarized. Whilst there is no diagonal banding, some horizontal banding is clear. However, one can neither predict nor rule out cycling on the basis of the fully processed plot.	32
13	Study 3 results. Event plot labeling the generations of one representative run in which the A-elite plays Paper or the B-elite plays Rock with effective probability 0.66 or greater. Given the lack of correlation between events, the likelihood of cycling appears small.	33
14	Study 3 results. Graphs plotting the effective probability values of playing each move during random drift. Left: the A-elites. Right: the B-elites. The populations do not appear to evolve in response to each other. Cycling is not observed.	34
15	Invasion-extinction cycles in <i>Anolis</i> lizards of the Caribbean islands. (A) Islands containing one resident population evolve to the equilibrium solitary size. (B) Lizard populations with larger body size are able to invade. Both populations evolve towards smaller body size. (C) The resident population becomes extinct, leaving the invaders to evolve to the solitary size. The cycle is complete (adapted from Roughgarden, 1983).	35
16	Study 2 results. Graph showing the mean fitness (left) and fitness variation (right) within populations A and B. Periods of zero variance signify disengagement. In both populations, there is very little fitness variation during generations 28-36 and 44-49. During these periods, population B is easily outperforming population A. As expected, disengagement coincides with stasis (homogeneous shading) on the enlarged CIAO plot of figure 11.	36
17	Study 4 results. Adopting the visualization scheme employed by Kerr et al. (2002), we plot the locations of each bacterial strain during one representative run of the <i>E. coli</i> simulation (toroidal grid size 100^2). Sensitive (S) bacteria are colored black, Resistant (R) are dark gray and Col toxic (C) are light gray. Empty cells are white. From left to right: (1) At timestep 1000 the three strains coexist across the grid; (2) By timestep 1200, C (light gray) have moved into areas occupied by S (black), S have moved to areas originally occupied by R (dark gray), and R have replaced C; (3) S continue to replace R, R replace C and C replace S; (4) By timestep 1600, the three strains have returned to roughly the same locations they occupied at timestep 1000. One full cycle is complete.	37
18	Study 4 results. CIAO plots depicting one representative run of the <i>E. coli</i> simulation over $N = 5000$ timesteps (sampled at a resolution of 50 generations). From left to right: (1) Some diagonal banding is observable in the raw data plot; (2) A large Gaussian ($r = 5$) removes fine detail; (3) LoG ($\sigma = 2$) produces a contour map; (4) The image is binarized. The fully processed plot exhibits clear diagonal banding.	38
19	Study 4 results. Event plot labeling the points (sampled at a 50-timestep resolution) at which the local density of Resistant (R) bacteria or Col toxic (C) bacteria exceeds 66% during one representative run. The repeated “R follows C” event sequence suggests regular cycling.	39
20	Study 4 results. Graph plotting the local density of each of the three bacterial strains during the first 1000 timesteps (at a 50-timestep resolution). Cycling can be observed with regular sequence CRSCRS.	40
21	John Cartlidge.	41
22	Seth Bullock.	41

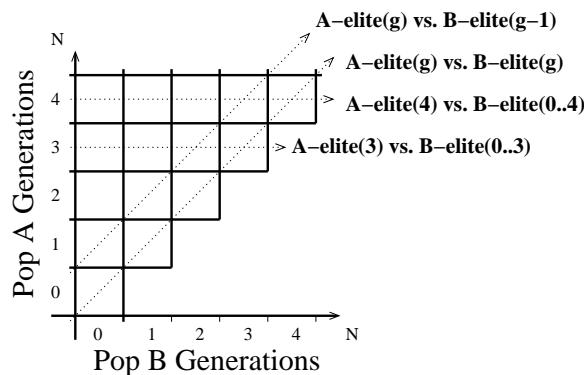


Figure 1: Schematic of a “Current Individual versus Ancestral Opponents” (CIAO) plot. At each square (x, y) the result of competition between the A-elite from generation y and the B-elite from generation x is plotted as a gray-scale value, with increasingly heavy shading representing an increasing margin of victory in favor of A. The leading diagonal (far-right) plots A-elite(g) against B-elite(g) for all generations, $0 \leq g < N$. The diagonal immediately to the left plots each generation’s A-elite against the *previous* generation’s B-elite. Horizontal rows plot the results of A-elite(y) against *all* ancestral B-elites($0 \dots y$) (adapted from Cliff & Miller, 1995).

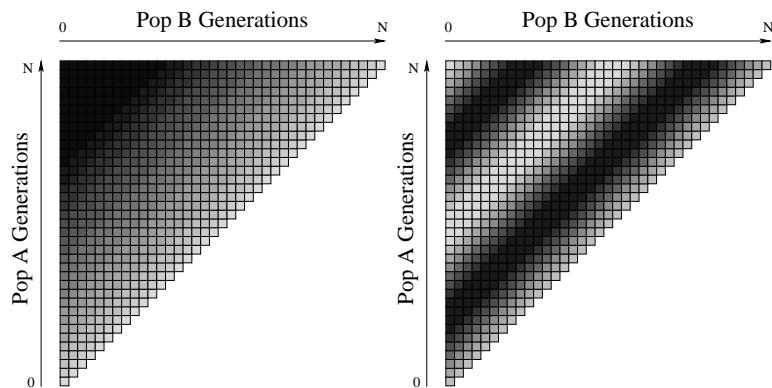


Figure 2: Idealized CIAO plots showing smooth progress (left) and cycling with regular period (right). With each horizontal row depicting the performance of one A-elite against each ancestral B-elite, smooth progress produces a gradation in intensity from dark to light: A-elites beat ancestral opponents, the more ancestral the greater the victory (i.e., the further left, the darker the cell). In contrast, regular cycling manifests as diagonal banding. Whilst the latest A-elites beat recent ancestral B-elites (diagonals to the right are dark), they perform less-well against more ancestral B-elites (middle diagonals are light), but outperform even earlier ancestors (diagonals further left are dark), etc.

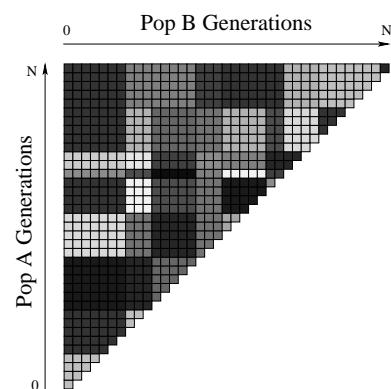


Figure 3: Schematic of a “tartan” CIAO plot. Periods of competitive stasis (blocks of uniform shade) are separated by sharp transitions in competitive advantage (block boundaries). Such an irregular pattern is referred to as “tartan” throughout this paper.

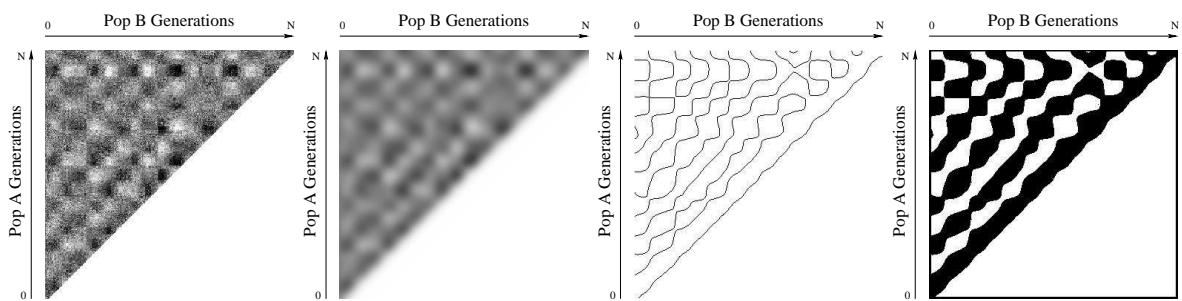


Figure 4: Study 1 results. CIAO plots depicting one representative run over $N = 256$ generations. From left to right: (1) Diagonal banding in the raw data plot suggests regular cycling; (2) A large Gaussian ($r = 20$) removes fine detail; (3) LoG ($\sigma = 8$) produces a contour map; (4) The image is binarized. The fully processed plot displays clear diagonal banding.

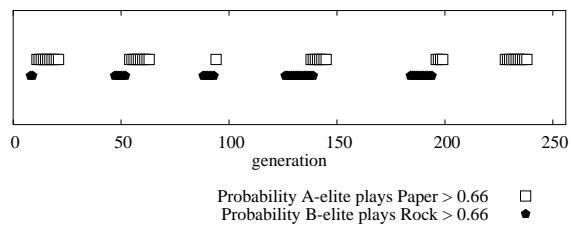


Figure 5: Study 1 results. Event plot for a single representative run, labeling the generations in which the A-elite plays Paper or the B-elite plays Rock with probability 0.66 or greater. The repeated “Paper follows Rock” event sequence suggests regular cycling.

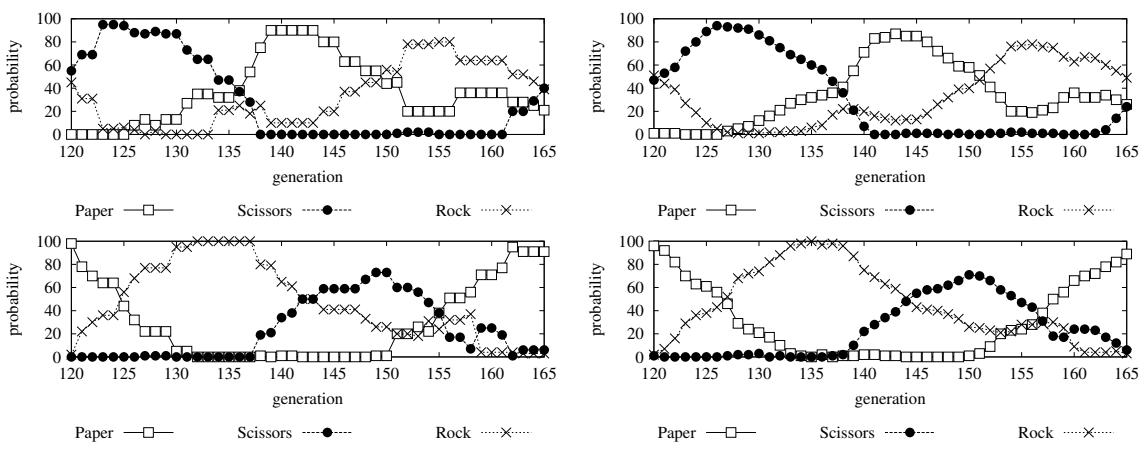


Figure 6: Study 1 results. Graphs plotting the probability values of playing each RPS move. Top-left: the A-elites. Top-right: the mean probabilities for all individuals in population A. Bottom-left: the B-elites. Bottom-right: the mean probabilities for all individuals in population B. The system completes one full cycle between generations 120 and 165.

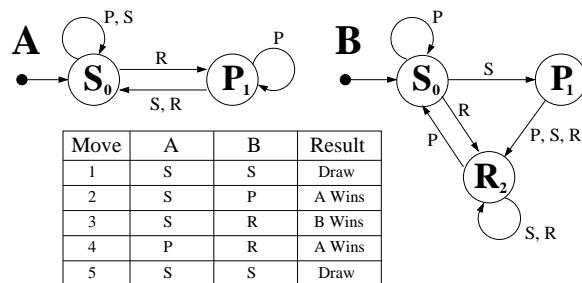


Figure 7: Complex genotypes encode a deterministic finite state machine (FSM) that governs player behavior. State transitions are dictated by an opponent's play (represented by transition labels). Here we see that machine *A* beats machine *B* over a five-bout contest. The final score is 3:2 in favor of *A*. As both machines have returned to their initial states, further bouts will produce repeated play.

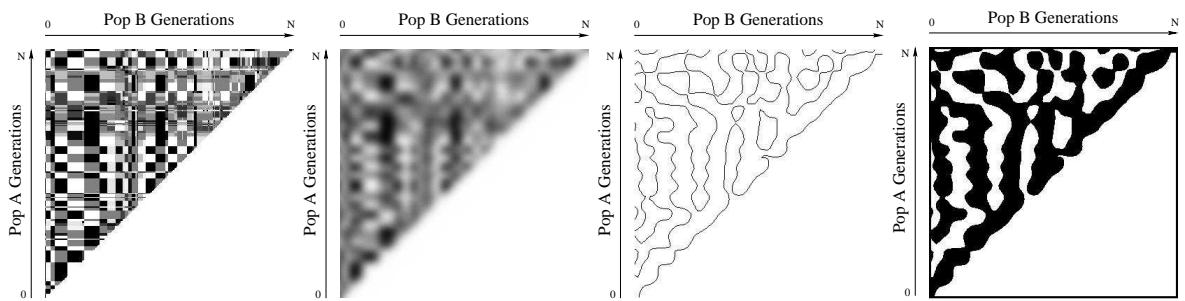


Figure 8: Study 2 results. CIAO plots depicting one representative run over $N = 256$ generations. From left to right: (1) The raw data plot exhibits a tartan pattern that is difficult to interpret; (2) A large Gaussian ($r = 20$) removes fine detail; (3) LoG ($\sigma = 8$) produces a contour map; (4) The image is binarized. Whilst the fully processed plot does not exhibit diagonal banding, some vertical banding is clear. However, one cannot predict regular cycling from this image.

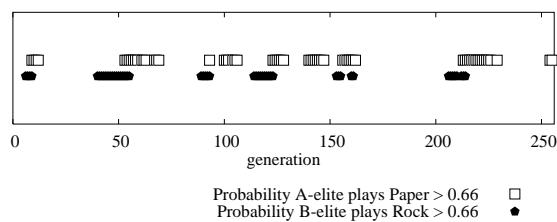


Figure 9: Study 2 results. Event plot for a single representative run, labeling the generations in which the A-elite plays Paper or the B-elite plays Rock with effective probability 0.66 or greater. The strong correlation between events suggests (perhaps irregular) cycling.

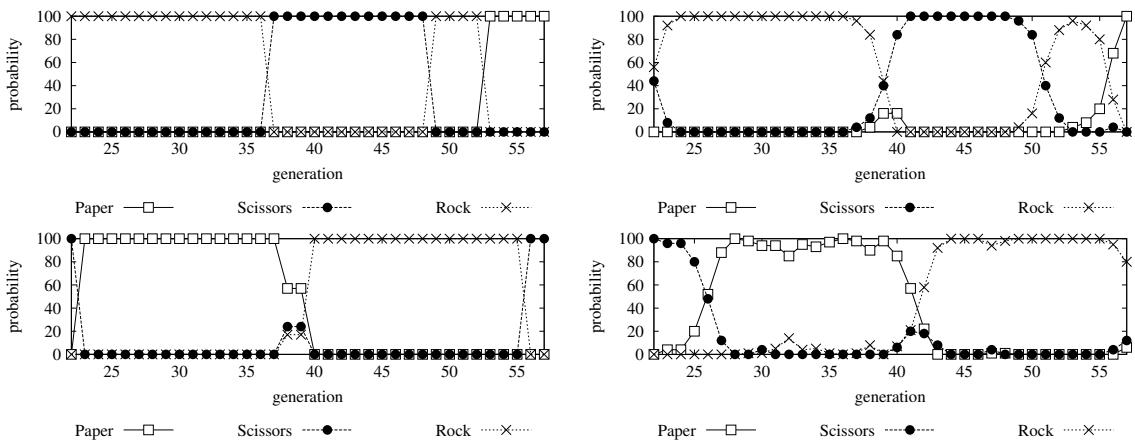


Figure 10: Study 2 results. Graphs plotting the effective probabilities associated with playing each move. Top-left: the A-elites. Top-right: the mean probabilities for all individuals in population A. Bottom-left: the B-elites. Bottom-right: the mean probabilities for all individuals in population B. The populations evolve in response to each other, resulting in an irregular cycle.

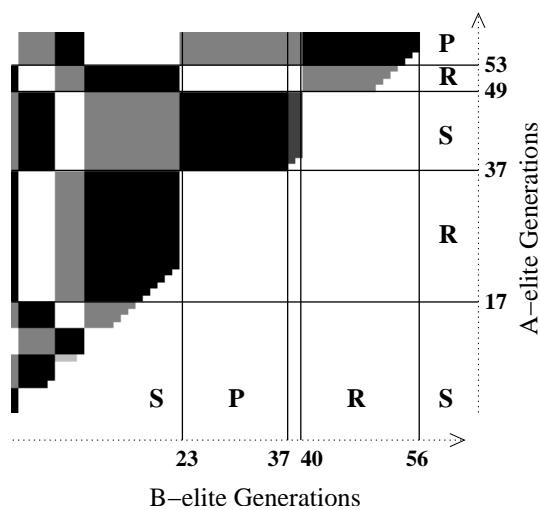


Figure 11: Enlarged CIAO plot depicting the coevolutionary period graphed in figure 10. The rapid strategy transitions depicted in figure 10 clearly reflect the boundaries in the CIAO plot pattern.

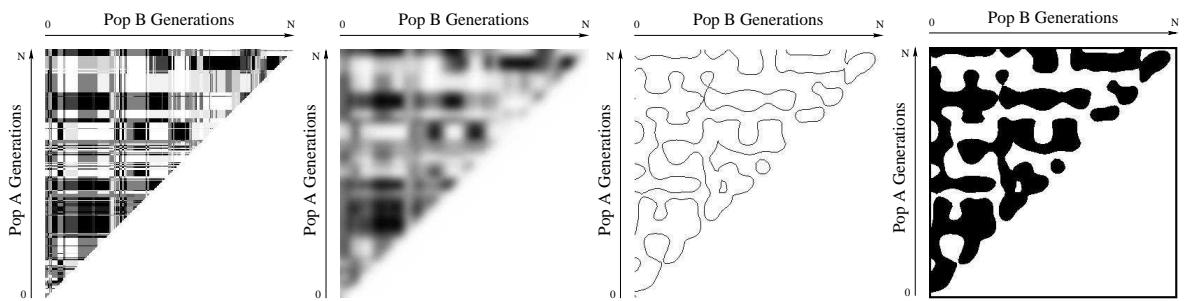


Figure 12: Study 3 results. CIAO plots depicting one representative run over $N = 256$ generations. From left to right: (1) The tartan pattern of the raw data plot is difficult to interpret; (2) A large Gaussian ($r = 20$) removes fine detail; (3) LoG ($\sigma = 8$) produces a contour map; (4) The image is binarized. Whilst there is no diagonal banding, some horizontal banding is clear. However, one can neither predict nor rule out cycling on the basis of the fully processed plot.

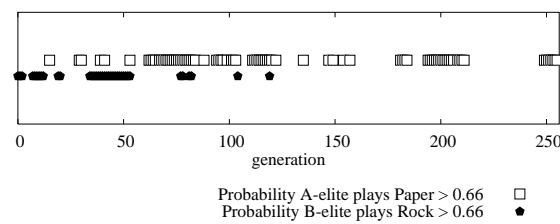


Figure 13: Study 3 results. Event plot labeling the generations of one representative run in which the A-elite plays Paper or the B-elite plays Rock with effective probability 0.66 or greater. Given the lack of correlation between events, the likelihood of cycling appears small.

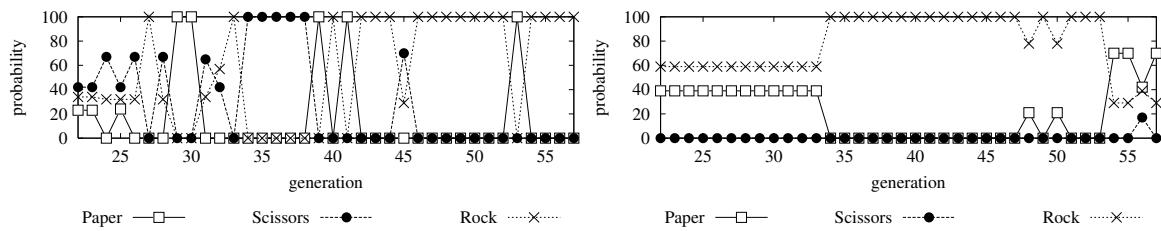


Figure 14: Study 3 results. Graphs plotting the effective probability values of playing each move during random drift. Left: the A-elites. Right: the B-elites. The populations do not appear to evolve in response to each other. Cycling is not observed.

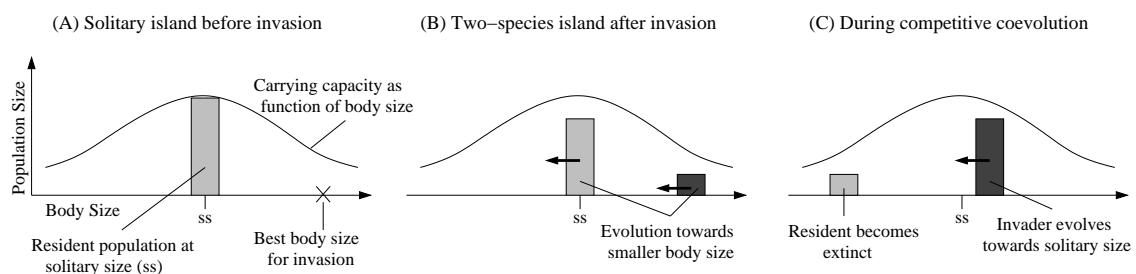


Figure 15: Invasion-extinction cycles in *Anolis* lizards of the Caribbean islands. (A) Islands containing one resident population evolve to the equilibrium solitary size. (B) Lizard populations with larger body size are able to invade. Both populations evolve towards smaller body size. (C) The resident population becomes extinct, leaving the invaders to evolve to the solitary size. The cycle is complete (adapted from Roughgarden, 1983).

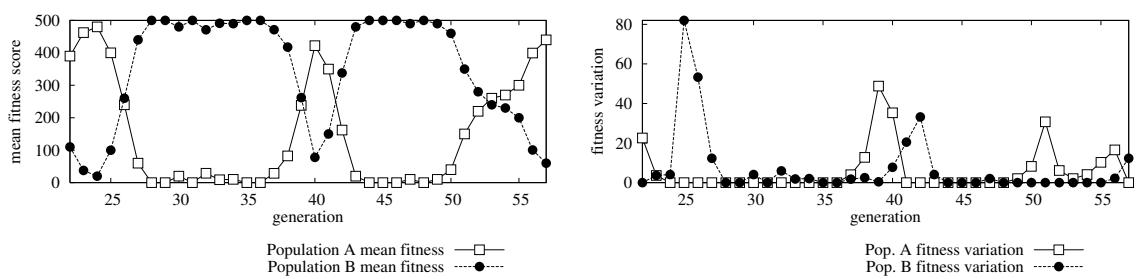


Figure 16: Study 2 results. Graph showing the mean fitness (left) and fitness variation (right) within populations A and B. Periods of zero variance signify disengagement. In both populations, there is very little fitness variation during generations 28-36 and 44-49. During these periods, population B is easily outperforming population A. As expected, disengagement coincides with stasis (homogeneous shading) on the enlarged CIAO plot of figure 11.

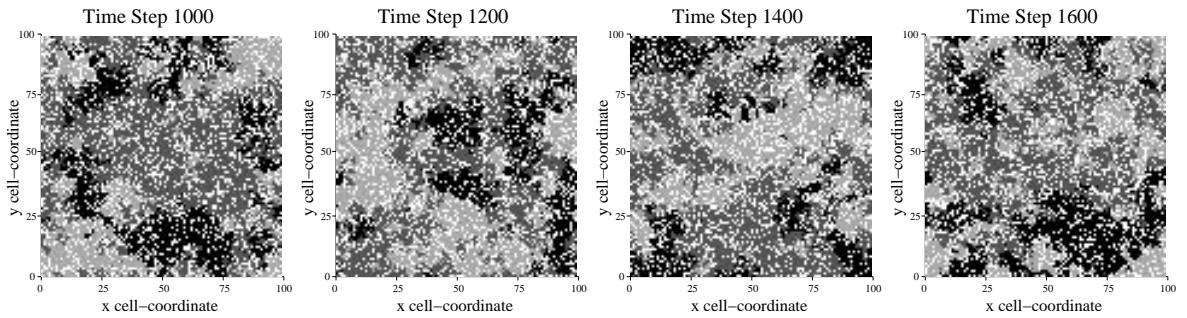


Figure 17: Study 4 results. Adopting the visualization scheme employed by Kerr et al. (2002), we plot the locations of each bacterial strain during one representative run of the *E. coli* simulation (toroidal grid size 100^2). Sensitive (S) bacteria are colored black, Resistant (R) are dark gray and Col toxic (C) are light gray. Empty cells are white. From left to right: (1) At timestep 1000 the three strains coexist across the grid; (2) By timestep 1200, C (light gray) have moved into areas occupied by S (black), S have moved to areas originally occupied by R (dark gray), and R have replaced C; (3) S continue to replace R, R replace C and C replace S; (4) By timestep 1600, the three strains have returned to roughly the same locations they occupied at timestep 1000. One full cycle is complete.

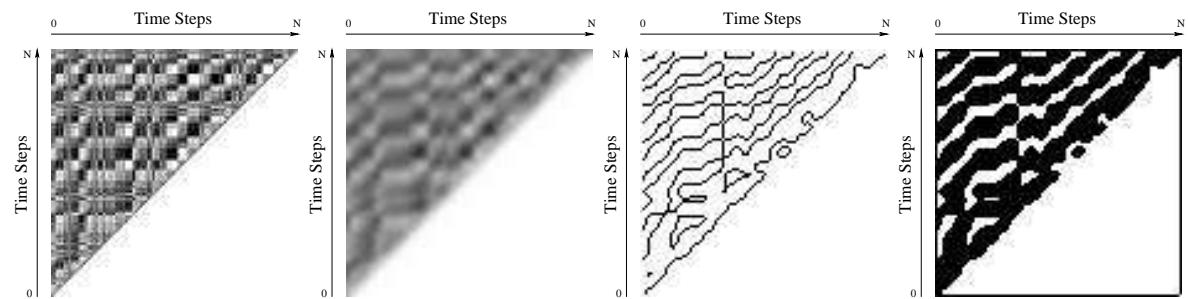


Figure 18: Study 4 results. CIAO plots depicting one representative run of the *E. coli* simulation over $N = 5000$ timesteps (sampled at a resolution of 50 generations). From left to right: (1) Some diagonal banding is observable in the raw data plot; (2) A large Gaussian ($r = 5$) removes fine detail; (3) LoG ($\sigma = 2$) produces a contour map; (4) The image is binarized. The fully processed plot exhibits clear diagonal banding.

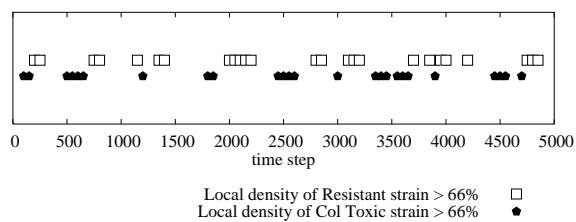


Figure 19: Study 4 results. Event plot labeling the points (sampled at a 50-timestep resolution) at which the local density of Resistant (R) bacteria or Col toxic (C) bacteria exceeds 66% during one representative run. The repeated “R follows C” event sequence suggests regular cycling.

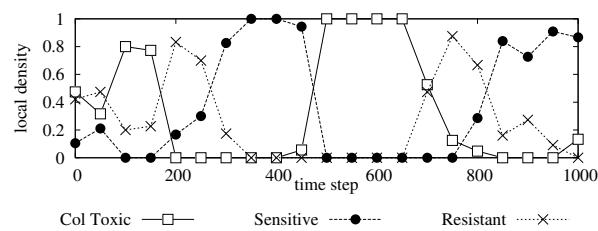


Figure 20: Study 4 results. Graph plotting the local density of each of the three bacterial strains during the first 1000 timesteps (at a 50-timestep resolution). Cycling can be observed with regular sequence CRSCRS.

Biography of Authors

John Cartlidge graduated (first class honours) in Artificial Intelligence & Mathematics at the Centre for Joint Honours in Science, University of Leeds, in 2000. He subsequently became a Ph.D. candidate at the School of Computing, University of Leeds, under the supervision of Dr. Seth Bullock. His research interests primarily focus upon the application of artificial coevolutionary algorithms to complex problem solving and his recently submitted doctoral thesis is entitled *Rules of Engagement: Competitive Coevolutionary Dynamics in Computational Systems*.



Figure 21: John Cartlidge.

Seth Bullock has a B.A. in Psychology & Computer Models (1993) and a D.Phil. in Computer Science & Artificial Intelligence (1997), both from the University of Sussex, UK. He currently works as a University Research Fellow at the University of Leeds, where he leads the Biosystems research group within the School of Computing. His research has focused on interdisciplinary problems at the boundary between the computational, biological and psychological sciences, developing evolutionary and adaptive computation techniques, game theory, and agent-based evolutionary simulation models.



Figure 22: Seth Bullock.