Analysing Co-evolution Among Artificial 3D Creatures

Thomas Miconi and Alastair Channon

University of Birmingham,
Edgbaston B152TT,
Birmingham, UK
t.miconi@cs.bham.ac.uk

Abstract. This paper is concerned with the analysis of coevolutionary dynamics among 3D artificial creatures, similar to those introduced by Sims [1]. Coevolution is subject to complex dynamics which are notoriously difficult to analyse. We introduce an improved analysis method based on Master Tournament matrices [2], which we argue is both less costly to compute and more informative than the original method. Based on visible features of the resulting graphs, we can identify particular trends and incidents in the dynamics of coevolution and look for their causes. Finally, considering that coevolutionary progress is not necessarily identical to global overall progress, we extend this analysis by cross-validating individuals from different evolutionary runs, which we argue is more appropriate than single-record analysis method for evaluating the global performance of individuals.

1   Introduction

Coevolution has been introduced in artificial evolution as an alternative to traditional evolutionary methods based on fixed, explicitly defined fitness functions such as the genetic algorithm. The use of coevolutionary methods is based on the assumption that constant mutual adaptation between evolving individuals will lead to ever-increasing levels of fitness. This assumption of progress through mutual adaptation is the basis for *arms race* hypothesis [3]. Rosin & Belew [4] summarise the transposition of the "arms race" concept to artificial evolution:

> Since the parasites are also evolving with a fitness based on a competition's outcome, the success of a host implies failure for its parasites. When the parasites evolve to overcome this failure, they create new challenges for the hosts; the continuation of this may lead to an evolutionary "arms race" (...) New parasite types should serve as a drive toward further innovation, creating ever-greater levels of complexity and performance by forcing hosts to respond to a wider range of more challenging parasite test cases.

The assumption which underlies artificial coevolution, therefore, can be stated as follows: coevolution is expected to lead to an "arms race" (formally defined as

a sequence in which newer individuals consistently outperform their ancestors), which is expected to result in superior individuals. Unfortunately the fundamentally local nature of natural selection (based on differential gene propagation within a given, current environment which is local both in space and time) means that several problems may hinder this intuitive mechanism.

First, it is well-known that the "arms race" metaphor begs the question of *intransitivity* in the global fitness landscape: if an organism A can be said to be superior to B, and B is superior to C, it is not necessarily the case that A should always be superior to C. This may lead to the appearance of "cycles" [2] [5] or "circularities" [6] in the dynamics of evolution.

Moreover, the arms race concept refers to a historical progress, in which newer individuals outperform ancestral ones against their ancestral opponents: performance and progress are evaluated against the history of a particular evolutionary trajectory. However, such a progress is not necessarily related to global, overall progress towards superior individuals in the wider context of the whole search space. Nolfi & Floreano [2] have shown that these two notions of progress are not as correlated as it may seem. They performed two coevolutionary experiments based on a predator-prey scenario, with one important difference: in one run, coevolution occurred in a straightforward manner, by pitting individuals of a given generation against the champion of the previous generation (a method inspired by Sims [1], which we also use in the present article). In the other run, however, individuals of a given generation were evaluated not only against the current opposing champion, but also against the *previous* champions of the opponent population, following the "Hall of Fame" technique suggested by Rosin & Belew [4]. Unsurprisingly, the second type of experiment led to a more robust arms race, in that newer individuals were significantly better at outperforming their own ancestors. However, in some circumstances, when the authors compared the results of coevolution with a Hall of Fame against "naked" coevolution, they found that individuals evolved using the Hall of Fame were defeated by individuals evolved without it. While progress had been more straightforward and unambiguous, it had also been more limited in scope. This difference between historical progress with regard to a given evolutionary history, and overall superiority, is an important topic in this article.

2   Monitoring and Analysis of Coevolution

If progress can occur in competitive coevolution, it is important that it be properly detected. Several types of statistics have been proposed for analysing the results of coevolutionary processes, with a stress on the identification of progress.

First, Cliff & Miller's "Current Individual vs. Ancestral Opponents" method (CIAO) [7] and Nolfi & Floreano's "Master Tournament" method [2] both pit the champions of each generation against each other, and displaying the result as a grid of coloured dots, in such a way that dot (n, m) is coloured if the champion of generation n in one population defeats the opposing champion of generation m, and left blank otherwise. CIAO pits the champion of a population at generation
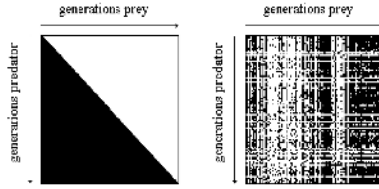
Fig. 1. Master Tournament matrices, taken from Nolfi & Floreano [2] in a predator-prey experiment. Black (resp. white) dots indicate a victory for the champion of the predator (resp. prey) population. The picture of the left represents an "ideal" situation of perfect progress, in which each champion is able to defeat all previous opponents. The picture on the right represents the results of a real experiment.

n against the champions of each *previous* generation in the opposing population, thus resulting in a triangle of dots. Master Tournament performs general confrontation between all champions of all generations, producing a square matrix of dots. The Master Tournament square can be seen as the collation of two CIAO triangles (one for each population) joined along their common hypotenuse, the diagonal of the square.

These methods have the advantage of providing reasonably complete information about an entire evolutionary run. However this completeness comes at a price. An obvious problem with these methods is their combinatorial complexity. Since N^2 evaluations are needed to obtain a complete table ($(N^2 - N)/2$ in the single-population case), as soon as N becomes even moderately large, calculating the figure is a time-consuming process. Of more concern to the analyst is the fact that the resulting figures are often somewhat obscure: although "ideal" conditions of progress lead to a very simple figure, these ideal conditions are rarely met in practice. Real experiments often produce disorderly arrangements of dots from which it may be difficult to extract any meaning at all.

A more recent technique for observing progress in coevolution has been proposed by Stanley & Miikkulainen under the name of Dominance Tournament [6]. Dominance Tournament was developed for single-population coevolution, but can be readily extended to multiple populations. In a dominance tournament analysis, one must keep track of every new individual that defeats all previously dominant individuals. Dominance is defined recursively: The first dominant strategy d_1 is the champion of the first generation; then, at every generation, the current champion becomes the new dominant strategy d_i if it can defeat all previous dominant strategies $d_{j<i}$. When two populations coevolve against each other, the method is adapted by specifying that a new dominant strategy must be able to defeat all dominant strategies from the other population.

Thus the dominance tournament method concentrates on a sequence of individuals which are seen as particularly important, due to their recursive superiority relationship. Dominance, in this context, is not synonymous with absolute superiority: some earlier individuals may be able to defeat the current dominant strategy. However such individual are seen as "idiosyncratic strategies", similar to parasites specialised against a (supposedly superior) host.

Dominance Tournament has the advantage of being much easier to compute than Master Tournament, since at any time the total number of dominant strategies against which candidates are to be tested is significantly lower than the total number of generation champions. It is also much easier to analyse, since it can be represented as a one-dimensional series of ticks along a time-coordinate axis, each tick corresponding to the appearance of a new dominant strategies. However, the massive simplification of the statistics eliminates a lot of information, and it is not clear exactly how precisely the Dominance Tournament captures the global trajectory of a given run.

Finally, both types of method must be applied to the history of a particular run: they essentially rely on "single record" analysis. They are useful in studying the trajectory of evolution and the presence (or absence) of coevolutionary progress. However it would be quite wrong to deduce anything from them about general progress in the sense of overall superiority over the whole search space.

3   Artificial Creatures

In the following sections we describe our own model for the evolution of artificial creatures in a physically realistic 3D environment. This model is broadly similar to the one introduced by Sims [1]. Besides minor technical modification, the most important difference between our system and Sims' is that our creatures are controlled by standard neural networks, based on classical McCulloch & Pitts neurons with sigmoid or radial activation functions, in contrast to Sims' creature which were controlled by functional networks, including arithmetic functions, tunable oscillators and logic operators (among others) as elementary building blocks. A complete description of (and justification for) the system can be found in a previous publication [8].

Morphology: As in Sims' model, the creatures are branching structures composed of rigid 3D blocks. The blocks (or "limbs") are connected to their parent limb by a hinge joint, except for the first ("root") limb. The genetic specification of a creature is given as a tree of nodes. Each of these nodes contain morphologic and neural information about one limb. Each node is responsible for storing the description of its limb's physical connection with its parent node's limb. The morphologic information in each genetic node specifies the dimensions of the limb (width, length and height), the orientation of this limb with regard to its parent (in the form of two parameters indicating polar angles with the xz and the xy planes, that is longitude and latitude, in the frame of reference of the parent limb), the direction of movement which may be either vertical or horizontal (that is aligned either with the y or with the z axis of the limb), and a boolean flag for reflection which governs symmetric replication along the xz plane of its parent. A limb also contains neural information, as described in the following paragraphs.

Creature control and neural organisation: Our creatures are controlled by neural networks. As in Sims' model, each limb contains a set of neurons. Genetic infor-
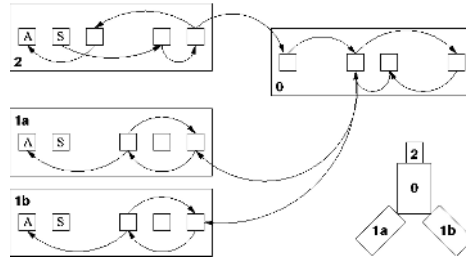
Fig. 2. Organisation of a fictional creature pictured in the bottom-right corner. Limb 0 has no sensor (S) or actuator (A). Limb 1 is reflected into two symmetric limbs 1a and 1b, which share the same morphologic and neural information.

mation about a given neuron specifies the *activation function* for this neuron, a threshold/bias parameter θ, and connection information. The activation function may be either a sigmoid ($\frac{1}{1+\exp^{-(\sigma+\theta)}}$) or the hyperbolic tangent $\tanh(\sigma+\theta)$ where σ is the weighted sum of inputs. The main difference between sigmoid and tanh is that the first has values in $[0,1]$ while the latter has values in $[-1,1]$. Connection information specifies, for each connection, the source of this connection (that is the neuron whose output is received through this connection) and a weight value. Neurons can only be connected with other neurons from the same limb, from adjacent limbs, or from the root limb. Each neuron may receive a variable number of connections, up to a maximum value (3 in the present experiments).

Sensor neurons and actuator neurons are handled specially. The first type of sensor neuron is a proprioceptive neuron, which measures the current angle formed by the hinge joint to which this neuron's limb is attached, scaled within the $[-1,1]$ range. Two other types of sensors exist, each of them measuring respectively the x and y coordinates of the centre of a specific object (an inert cube) in the frame of reference of the limb, squashed through the tanh function. Every limb has exactly one proprioceptor, and may have any number of other sensors (within the maximum number of neurons for each limb). Actuator neurons command the movement of each limb, that is, its rotation around its joint. The output of an actuator indicates the desired angular velocity at this joint. Their inputs are defined similarly as other neurons, but their activation function is always a scaled hyperbolic tangent of the form $\tanh(\sigma+threshold)$. Each limb has exactly one actuator.

Expression of the genome: The creatures are constructed according to the information contained in the genetic nodes. A very simple developmental system translates the genotype into a corresponding phenotype, and may introduce additional complexity if the genetic information dictates it. Our system uses only one developmental feature, adapted from Sims: bilateral symmetry. In our model, each genetic node (corresponding to a limb) may possess a "reflection" flag, which means that when this node is read and the corresponding limb attached to its parent, a symmetric copy of this limb will also be created. Any further sub-limbs

will similarly be duplicated in a symmetric fashion, which leads to the appearance of bilaterally symmetric branches. Our present design allows for only one type of symmetry, namely symmetry along the parent's xz plane. When a given limb is randomly generated, its reflection flag is set with probability P_{ref} (for this paper, $P_{ref} = 0.1$).

Genetic operators: We use three genetic operators, broadly similar to those used by Sims. *Crossover* is performed by simply aligning the genetic nodes of both parents in two rows, then building a new list of genetic nodes by concatenating the left part of one parent with the right part of the other. *Grafting* corresponds to the removal of a branch (that is a limb and all its sub-limbs), and its replacement by a branch taken from another individual. Connectivity information is adapted and maintained: the neurons of the trunk establish the same connections with the new branch as they had with the old one, and similarly the new branch has the same connection with its new trunk as it had with its previous trunk. *Mutation* occurs by sequentially modifying each parameter within a genome (from limb size to connection weight) with a given probability P_{mut}, and also removing a limb and adding a new, randomly generated limb, also with probability P_{mut} (in this paper, $P_{mut} = 0.04$).

4   Experiments and Results

4.1   The Evolutionary Algorithm

We use the same task as Sims [1]: two creatures compete for control of a single cube. The cube is placed in the center of the world, and the creatures start on each side of the cube. After a fixed amount of time has elapsed, distances d_1 and d_2 between the centre of the root limb of each competitor and the centre of the cube are computed. The score of each contestant is the difference between these distances, $d_1 - d_2$ for competitor 1 and $d_2 - d_1$ for competitor 2. Lower score correspond to superior creatures.

The evolutionary algorithm is also similar to Sims'. For every run, creatures are divided into two populations. At every generation, creatures of each population are evaluated against the current champion of the opposing population. The creature which obtains the best score becomes the new champion of this population. Survival rate is 50%, which means that half the population is replaced at every generation. Selection of parents occurs by direct tournament selection based on score. New individuals are created with equal probability by one of three operations: grafting between the two individuals, crossover between the two parents, or three successive applications of the mutation operator to one of the parents. Then the mutation operator is applied to the resulting creature and produces the final offspring. If the developed phenotype of an offspring creature contains two intersecting non-adjacent limbs, or too many limbs, the creature is deemed non viable, and the reproductive operation chosen is repeated as often as necessary until a viable creature is produced. Each run covers 500 generations.

The system produced a wide variety of behaviours, some of which are illustrated in Figure 3. In the top-left frame, one creature catches the cube in a
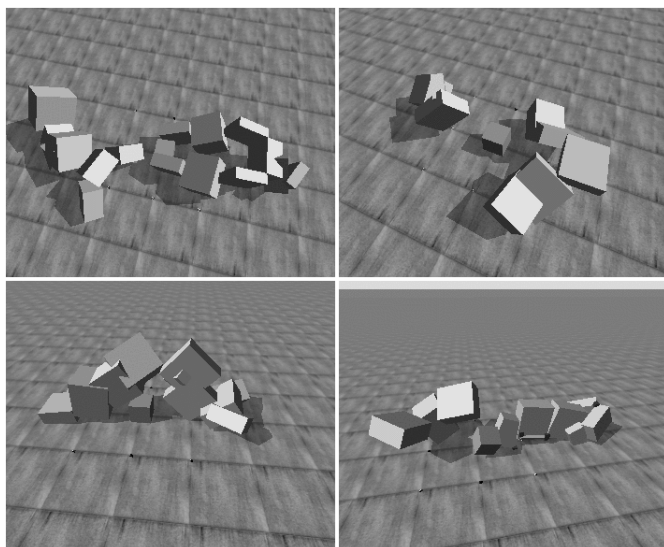
Fig. 3. Creatures evolved from four different runs. See text for details.

pinching motion and draws it towards its trunk limb before its opponent manages to reach it. In the top-right frame, two creatures use different methods to move toward the cube. In the bottom-left frame, one creature manages to push its opponent away from the cube, even though the other creature had reach the cube first. In the bottom-right frame, a two-armed creature is chasing the box that is being pushed aside by its opponent.

4.2 Coarse-Grained Master Tournament Matrices

To monitor the progress of evolution, we chose to introduce a modified version of the Master Tournament method. In our case the original method would be difficult to apply, since computing the whole Master Tournament matrix for 500 generations would be computationally prohibitive. Furthermore, as we mentioned in the introduction, Master Tournament grids are often difficult to read and analyse.

Our method consists in simplifying the Master Tournament by a "coarse-graining" operation. Instead of performing a full tournament between the champions of all N generations, we pick a fixed number k of champions and use this sample to produce a coarse-grain Master Tournament matrix. This means that we only perform tournaments between champions of generations which are integer multiples of the N/k fraction. In our example, we chose to condense our 500 generations into a 50x50 tournament grid, which means that we need to select the champions of every tenth generation (roughly) in each population. By sampling 10% of the generations, computing costs for head to head competitions are divided by 100.

Like any sampling process, coarse-graining incurs a loss of information. However, the information which is lost by coarse-graining is essentially short-term,

small-scale information. When analysing the results of a coevolutionary experiment, we are usually interested in long-term trends, especially regarding evolutionary progress Coarse-graining adequately preserves this type of larger-scale information. In particular, the question of whether or not a given individual can consistently outperform older ones, which is the crucial aspect of the "arms race" concept, is not affected by coarse-graining. Moreover, coarse-graining can actually make a Master Tournament matrix more descriptive by suppressing spurious, irrelevant information: as we make clear in the following paragraphs, coarse-grained Master Tournament matrices may exhibit discernible, informative features which are often difficult to observe in full matrices.

4.3   Reading a Coarse-Grained Master Tournament Matrix

Figure 4-left shows a coarse-grained Master Tournament matrix for a particular run. Each (m, n) location is marked with a dark square point if the champion of population 1 generation $10 * m$ defeats the champion of population 2 at generation $10 * n$, or with a light cross mark otherwise. The $y = x$ line, drawn in a lighter shade, provides a time axis for the actual run. Points on this line indicate how the actual run went along, indicating the victorious population at each generation. On a coarse-grained Master Tournament matrix, vertical patterns are related to individuals from population 1, while horizontal patterns are related to individuals from population 2.

A first observation for this run is that the $y = x$ line goes through several regions of different colour. This means that the champions of the two populations successively outperform each other, an indicator of healthy competition. However, the particular patterns of this alternation provide a better insight about the course of evolution in this run.

Identifying similar phenotypes from their competitive profiles: The graph in Figure 4-left contain many similar lines and columns. In particular, it may be seen that many columns offer strikingly similar patterns of dark and light marks, although with appreciable variation. Each column, however, corresponds to the competitive profile of a champion in population 1: it accounts for its successes
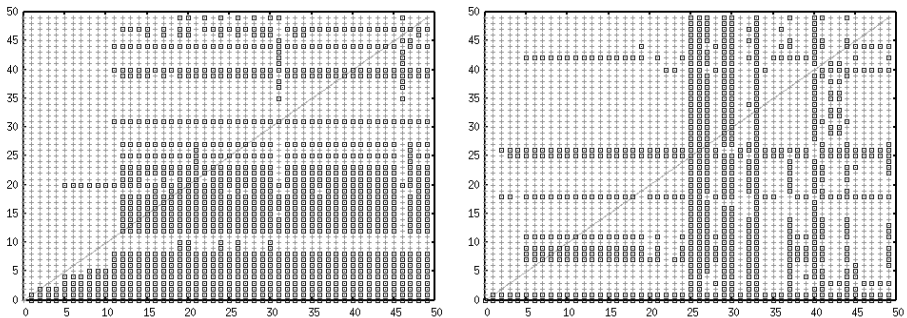

Fig. 4. Coarse-grained Master Tournament matrix for two different runs

and failures against every champion of population 2. Two identical columns denote two individuals that defeat the same opponents, and are defeated by the same opponents. It is not too far-fetched to assume that similarity in competitive profile is linked to similarity in phenotypes.

Similarity is not identity, and much variation can be seen. However there are at least two columns which offer a significantly different profile to the neighbouring columns, namely columns 31 and 46. These two columns can be said to represent different types from their neighbours, due to the difference in competitive profiles. In particular they are unique in being able to defeat the opponents in rows 41-43.

The significance of the high similarity in columns after 12 can be seen as an indication that, at least in population 1, evolution seems to have settled on a particular type of creature, which is marginally "fine-tuned" in the later course of coevolution. This capacity to indicate phenotypical convergence is an interesting property of coarse-grained Master Tournament matrices.

Evidence of breakdowns in arms race: The arms race concept implies that newer individuals are consistently able to outperform their ancestors. Breakdowns in the arms race are associated with a loss of adaptive function, since an ability (to defeat some individuals that could be defeated by ancestors) has been lost by the newer individual [9].

Breakdowns or interruptions in arms races are easy to locate on a Master Tournament matrix. Any dark mark occurring immediately above a light mark, or any light mark occurring immediately on the right of a dark mark, indicates such an interruption: it means that a given individual (from population 2 in the first case, from population 1 in the second case) was unable to defeat an opponent that could be defeated by its ancestor. Such breakdowns may be very short events, indicating a prompt recovery. Alternatively they may result in a long-term loss, or even permanent loss of the capacity to defeat some particular individuals.

Let us take the example of the first horizontal stripe of lighter marks, at rows 9-11. These rows correspond to particularly fit individuals which are able to defeat a large number of opponents (all of them for row 11). In particular, they have no difficulty defeating the champions of generations 12 to 18 in population 1, as can be seen from the fact that their rows are void of dark marks in the section between columns 12 and 18.

Yet the same graph shows that, from generation 12 onwards, the first bisectant encounters a series of dark marks, indicating superiority of the champions of population 1. This indicates that by generation 12, the current champion of population 2 had become unable to defeat individuals that earlier ancestors could defeat. How did this come to be ? If the champion of population 2 at generation 11 was good enough to defeat all opponents that population 1 would ever come up with, why was it displaced with one that would prove to be inferior?

Causes for breakdowns in arms races: This alternation between a lighter stripe and a sudden block of dark marks indicates a dramatic example of a breakdown

in the arms race. In this particular case, the cause can be identified as *over-specialisation*. While it is true that the champion of generation 12 in population 2 was potentially able to defeat a large number of opponents, population 2 managed to come up with a new individual which was even better (that is obtained a lower differential distance to the cube) against the current champion of population 1. Unfortunately this change, while beneficial in the short term, proved fatal when population 1 managed to evolve a counter-strategy which defeated this specialised opponent. This allowed the newly evolved type of individual in population 1 to take the lead, even though previous champions from population 2 would have been able to defeat it.

Figure 4-right shows the results of a different run. This figure exemplifies several other informative patterns. In particular, let us look at the centre of the matrix, at row and column 25. At that point, we see that the first bisectant encounters a kind of wedge, composed of two stripes of dark marks - one vertical, one horizontal. Can we infer some meaning from this pattern ? The wedge shape indicates that a successful change in population 1 (indicated by the appearance of a different competitive profile, leading to a distinct, darker series of columns) has led to a dramatic breakdown in the arms race on the side of population 2. The appearance of this new champion in population 1 has *upset* the hierarchy in population 2: the previous champion was no longer the best possible candidate against this new opponent. Confronted with the new, successful champion of population 1, population 2 has settled on a new champion, which happened to perform better, or at least less badly, than others against this particular new opponent (though not well enough to actually defeat it). This new "champion of fortune", however, was not particularly well-rounded and performed badly against a large range of opponents. Innovation in population 1 has caused a *confusion* in population 2.

This idea of new individuals breaking down the arms race by upsetting the hierarchy and voiding previous adaptations in their opponents is not necessarily linked to wedge-like patterns, but simply to the appearance of a new type of opposing champions. For example, Figure 4-left contains several dark horizontal lines, apparently isolated. In particular, the individual in row 31 indicates that this champion suddenly lost much of its ancestors' aptitudes against opposing champions. What is the cause of this loss ? If we track the point at which this new, poorly performing champion occurs (by locating its intersection with the first bisectant) and observe the corresponding column, we notice that the individual from population 1 at column 31 has a subtly different pattern from its predecessors. The poor performance of population 2 at generation 31 is thus caused by the emergence of a new opponent which upsets the hierarchy in population 2 and propels an apparently poor individual to the rank of "champion".

These interruptions in the arms race (temporary or long-term) that can be observed on the coarse-grained Master Tournament matrix are an indication of the local nature of co-evolution. Because co-evolution is only concerned about the immediate present, it may directly induce a loss of ability against past or

future opponents. This loss may occur spontaneously (as in over-specialisation) or may be provoked by a change in the opposing population (as in "confusion").

5   Cross-Validation of Coevolutionary Runs

Master Tournament matrices, however informative, can only describe performance within the context of a particular run. This is not necessarily sufficient to express the general level of performance of an individual in the larger context of the entire search space. In order to detect whether a given individual may really be called superior, it is not enough to confront it to the population against which it evolved. Such a test could be seen as a confusion between the training set and the test set. Given several evolutionary runs, if we want to obtain a more global view of each individual's performance, the most simple method we can use is simply to test each individual not only against its own opponents, but also against other populations of other runs. In other words, we expect that *cross-validating* individuals from different evolutionary runs would provide more reliable information about their global efficiency.
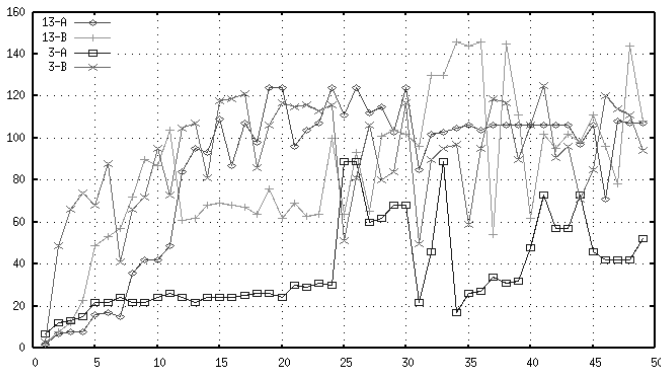

Fig. 5. Cross-validation of each individual in each of the four populations shown in Fig. 4. 13-A and 13-B are population 1 and 2 from the left-hand side matrix, while 3-A and 3-B are population 1 and 2 from the right-hand side matrix in Fig. 4.

Figure 5 shows, for each champion of all 50 generations in each population, the number of champions of all other opposing populations that it is able to defeat. 13-A and 13-B are population 1 and 2 from the left-hand side matrix in Fig. 4, while 3-A and 3-B are population 1 and 2 from the right-hand side matrix. This graph is interesting both for its similarities and its difference with the individual Master Tournament matrices in Figure 4.

Within this larger context, the best performing individuals are the champions of generations 34 and 36 from population 13-B, with a score coming close to the maximum 150, meaning that they can defeat almost all other champions. Looking at the corresponding rows in Fig. 4, we observe that they indeed

obtain 'clean sheets' against all their opponents. However, this is also the case with rows 11 and 28-30, yet these ones obtain a much lower score on the cross-validation graph. This indicates a difference in performance that could not have been deduced from Master Tournament (coarse-grained or not) or Dominance Tournament analysis, nor indeed from any single-record analysis method alone.

Similarly, we see that population 3-A seems to perform rather poorly when compared to others. Specifically, after generation 10, all champions of population 3-A obtain much lower performance that champions in population 13-A. This is in contrast with the corresponding Master Tournament matrices, in which it can be seen that some champions of population 3-A are able to defeat all opposing champions from population 3-B (columns 25-26 and 33), while no champion in population 13-A shows such a perfect record. Again, a single-record analysis could not have detected this apparent superior performance of individuals from population 13-A.

References

1. Sims, K.: Evolving 3d morphology and behavior by competition. In: ALife IV : Proc. of the 4th Conference on Artificial Life, MIT Press (1994) 28–39
2. Nolfi, S., Floreano, D.: Coevolving predator and prey robots: Do "arms races" arise in artificial evolution? Artificial Life **4** (1998) 311–335
3. Dawkins, R., Krebs, J.R.: Arms races between and within species. Procs of the Royal Society of London, Series *B* **205** (1979) 489–511
4. Rosin, C.D., Belew, R.K.: New methods for competitive coevolution. Evolutionary Computation **5** (1997) 1–29
5. Watson, R.A., Pollack, J.B.: Coevolutionary dynamics in a minimal substrate. In Spector, L., Goodman, E.D., Wu, A., Langdon, W.B., eds.: Procs GECCO 2001, Morgan Kaufmann (2001)
6. Stanley, K.O., Miikkulainen, R.: The dominance tournament method of monitoring progress in coevolution. In: Procs GECCO 2002 Workshop, Morgan Kaufman (2002)
7. Cliff, D., Miller, G.F.: Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In: Proceeding of the European Conference on Artificial Life (ECAL-95). (1995) 200–218
8. Miconi, T., Channon, A.: A virtual creatures model for studies in artificial evolution. In: Procs of the 2005 IEEE Congress on Evolutionary Computation. (2005)
9. Ficici, S.G., Pollack, J.B.: A game-theoretic memory mechanism for coevolution. In et al, C.P., ed.: Proc. GECCO 2003, Springer (2003)