To appear in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2002) Workshop Program*. San Francisco, CA: Morgan Kaufmann

The Dominance Tournament Method of Monitoring Progress in Coevolution

Kenneth O. Stanley
Department of Computer Sciences
University of Texas at Austin
Austin, TX 78712
kstanley@cs.utexas.edu

Risto Miikkulainen
Department of Computer Sciences
University of Texas at Austin
Austin, TX 78712
risto@cs.utexas.edu

Abstract

In competitive coevolution, the goal is to establish an "arms race" that will lead to increasingly sophisticated strategies. The existing methods for monitoring progress in coevolution are designed to demonstrate that the arms race indeed occurred. However, two issues remain: (1) How can progress be monitored efficiently so that every generation champion does not need to be compared to every other generation champion? (2) How can a monitoring method determine whether strictly more sophisticated strategies are discovered as the evolution progresses? We introduce a new method for tracking progress, the dominance tournament, which provides an answer to both questions. The dominance tournament shows how different coevolution runs continue to innovate for different periods of time, reveals the precise generation in each run where stagnation occurs, and identifies the best individuals found during the runs. Such differences are difficult to detect using standard techniques but are clearly distinguished in a dominance tournament, which makes this method a highly useful tool in understanding progress in coevolution.

1 INTRODUCTION

In competitive coevolution, two or more populations of individuals evolve simultaneously in an environment where an increased fitness in one population leads to a decreased fitness for another. Ideally, competing populations will continually outdo one another, leading to an "arms race" of increasing sophistication (Dawkins and Krebs 1979; Van Valin 1973). An important question is: How can we establish that the strategies indeed become more sophisticated over time?

Master tournament is currently the most common method for monitoring coevolution. The champion of every generation is compared to the champions of all prior generations, or to the champions of all generations in the entire run (Cliff and Miller 1995; Floreano and Nolfi 1997). The results can be depicted graphically by plotting the number of other champions defeated by each generation champion. If evolution makes progress, the graph should show that the higher the generation, the more opponents the champion can beat.

Although this method for monitoring progress can reveal a general trend in increasing capacity to defeat prior strategies, it has several shortcomings. The most obvious problem is computational complexity. If every champion is to be compared to every other champion, the time to complete an analysis can take weeks, or even years. For example, if a multi-trial comparison takes one minute, monitoring 500 generations of progress would take almost three months! Therefore, it is often only possible to make single trial comparisons between champions, which are often inaccurate.

Second, a master tournament does not clearly indicate whether an arms race took place. Although the master tournament may show that the champion of a given generation can defeat many other generation champions, such results do not take into account *which* strategies can be defeated and thereby does not identify circularities. For example, although a strategy from late in a run may defeat more generation champions than an earlier strategy, the later strategy may not be able to defeat the earlier strategy itself! Conversely, it is possible that although two strategies can defeat an equal proportion of generation champions, one may still easily defeat the other. To demonstrate an arms race, it is necessary to verify that as more sophisticated strategies are developed, they can indeed defeat the most sophisticated strategies that were found earlier. Because the master tournament does not provide such a verification, its results can be misleading.

To address these two issues, we developed the *dominance tournament* method of tracking progress in competitive coevolution. The method does not require comparing every champion with every other champion, thereby saving computation time and allowing more accurate comparisons. Moreover, the dominance tournament identifies a ranking of increasingly sophisticated strategies, such that every newly identified dominant strategy defeats *all* prior dominant strategies. With such a guarantee, the dominance tournament provides proof of progress.

In this paper, we monitor progress in a complex competitive robot duel domain using both the dominance tournament and master tournament. Neural network controllers for the robots were evolved using NeuroEvolution of Augmenting Topologies (NEAT; Stanley and Miikkulainen 2002b,c), which allows complexification of neural networks over the course of evolution. We examine two runs of evolution, each with a different coevolution method, so that the monitoring techniques have a chance to detect differences between the two runs. The first evolution run used fully-functional NEAT, while the second used a disabled version of NEAT that evolves only fixed-topologies. In the first run, the dominance tournament demonstrated progress in evolution that the master tournament was unable to detect. In the second run, the dominance tournament identified the specific generation where progress began to stagnate, whereas the master tournament could not.

We begin by describing the NEAT neuroevolution method, followed by a description of the robot duel domain and a detailed discussion of monitoring progress.

2 NEUROEVOLUTION OF AUGMENTING TOPOLOGIES (NEAT)

The NEAT method of evolving artificial neural networks combines the usual search for appropriate network weights with complexification of the network structure, allowing competitive coevolution to evolve both increasingly optimal and increasingly complex structures. Increasing the complexity of networks over evolution allows NEAT to elaborate on evolved strategies by adding new functionality to them. In this section a brief overview of NEAT is provided; see Stanley and Miikkulainen (2002c) for a complete description. In this paper, we will use the comparison of NEAT and fixed-topology neuroevolution as a test case. By using NEAT, we are able to compare evolution where structures complexify with evolution where they do not. Methods for monitoring progress should reveal differences between the two kinds of evolution.

NEAT begins evolution with neural networks with no hidden nodes. This way, NEAT searches the smallest, most efficient topologies first. If those topologies are insufficient, new nodes and connections are added through mutation, providing additional space for more complex solutions.

Crossover is possible in NEAT because the historical origin of every new gene is tracked by the system. Every new gene resulting from a mutation is assigned an *innovation number.* Offspring inherit the same innovation numbers on genes as were present in their parents. When genomes cross over, genes with matching innovation numbers are lined up, allowing NEAT to match up different network structures without expensive topological analysis.

Adding new structure to existing networks can initially reduce fitness, causing innovative structures to disappear from the population prematurely. In order to protect new structure, NEAT speciates the population based on topological similarity among networks: Networks with similar topologies are grouped into the same species. NEAT is able to measure topological similarity using the historical markings which indicate which genes represent the same structural building blocks.

Explicit fitness sharing (Goldberg and Richardson 1987) is used for reproduction in the speciated population. In explicit fitness sharing, networks in the same species share the fitness of the niche, thereby preventing crowding around single solutions. Because of fitness sharing, no single species can take over the population. Thus, a large variety of topologies can evolve simultaneously with little interference from each other. The result is that innovative topologies are protected in their own niches where they have a chance to optimize before competing with the population at large.

The NEAT system implements *complexifying* coevolution because networks become more complex as they become more sophisticated. Complexifying coevolution should result in more sustained progress and more sophisticated final strategies than evolving fixed-topology networks. In this paper, the dominance tournament method will be compared against other analysis methods in verifying this hypothesis.

3 THE ROBOT DUEL DOMAIN

To demonstrate the utility of different monitoring techniques, a domain is needed where it is possible to develop increasingly sophisticated strategies and where the sophistication can be readily measured. Pursuit and evasion tasks have been utilized for this purpose in the past (Gomez and Miikkulainen 1997; Jim and Giles 2000; Miller and Cliff 1994), and can serve as a benchmark domain for competitive coevolution as well. While past experiments evolved either a predator or a prey, an interesting coevolution task can be established if the agents are instead equal and engaged in a duel. To win, an agent must develop a strategy that outwits that of its opponent, utilizing structure in the environment.

In the robot duel domain, two simulated robots try to overpower each other (figure 1). The two robots begin on opposite sides of a rectangular room facing away from each other. As the robots move, they lose energy in proportion to the amount of force they apply to their wheels. Although the robots never run out of energy (they are given enough

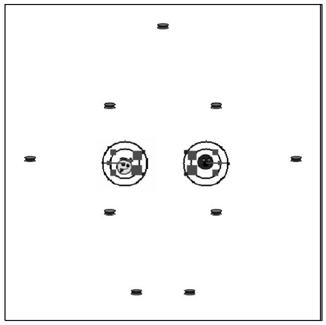

Figure 1: **The Robot Duel Domain.** The robots begin on opposite sides of the board facing away from each other as shown by the lines pointing away from their centers. The concentric circles around each robot represent the separate rings of opponent sensors and food sensors available to each robot. Each ring contains five sensors, which appear larger or smaller depending on their activations. From this initial position, neither robot has a positional advantage. The robots lose energy when they move around, yet they can gain energy by consuming food (shown as black dots). The food is placed in a horizontally symmetrical pattern around the middle of the board. The objective is to attain a higher level of energy than the opponent, and then collide with it. Because of the complex interaction between foraging, pursuit, and evasion behaviors, the domain allows for a broad range of strategies of varying sophistication. Animated demos of the robot duel domain are available at **www.cs.utexas.edu/users/nn/pages/research/neatdemo.html**.

to survive the entire competition), the robot with higher energy can win by colliding with its competitor. In addition, each robot has a sensor indicating the difference in energy between itself and the other robot. To keep their energies high, the robots can consume food items, arranged in a symmetrical pattern in the room.

The robot duel task supports a broad range of sophisticated strategies that are easy to observe and interpret without expert knowledge. The competitors must become proficient at foraging, prey capture, and escaping predators. In addition, they must be able to quickly switch from one behavior to another. The task is well-suited to competitive coevolution because naive strategies such as forage-then-attack can be complexified into more sophisticated strategies such as luring the opponent to waste its energy before attacking.

The simulated robots are similar to Kheperas (Mondada et al. 1993). Each has two wheels controlled by separate motors. Five rangefinder sensors can sense food and another five can sense the other robot. Finally, each robot has an energy-difference sensor, and a single wall sensor.

The robots are controlled with neural networks evolved with both NEAT and fixed-topology neuroevolution. The networks receive all of the robot sensors as inputs, as well as a constant bias that can be used to change the activation thresholds of neurons. They produce three motor outputs: Two to encode rotation either right or left, and a third to indicate forward motion power.

This complex robot-control domain allows competitive coevolution to evolve increasingly sophisticated and complex strategies, and can be used to benchmark coevolution methods. Thus, it serves as a useful testbed for different methods of monitoring progress.

4 EXPERIMENTS

We ran one run of evolution with full NEAT, and one run with NEAT's complexification capability turned off. In the latter run, fixed-topology networks were evolved with 10 hidden nodes. The master tournament and dominance tournament were used to identify any differences between the complexifying coevolution and fixed-topology coevolution runs. The experimental methodology is described below.

4.1 COMPETITIVE COEVOLUTION SETUP

In each evolution run, 2 populations, each containing 256 genomes, were evolved simultaneously. In each generation, each population was evaluated against a sample of networks from the other population. The population currently being evaluated is called the *host* population, and the population from which opponents are chosen is called the *parasite* population (Rosin and Belew 1997). The parasites are chosen for their quality and diversity, making such host/parasite evolution more efficient and more reliable than random or round robin tournament.

A single fitness evaluation included two trials, one for the east and one for the west starting position. That way, networks needed to implement general strategies for winning, independent of their starting positions. Host networks received a single fitness point for each win, and no points for losing. If a competition lasted 750 time steps with no winner, the host received no points.

In selecting the parasites for fitness evaluation, good use can be made of the speciation and fitness sharing that already occur in NEAT (the fixed-topology run was also able to speciate based on weight differences). Each host was evaluated against the champions of four species with the highest fitness. They are good opponents because they are the best of the best species, and they are guaranteed to be diverse because their compatibility must be outside the threshold for being grouped into the same species (section 2). Another eight opponents were chosen randomly from a Hall of Fame (Rosin and Belew 1997) that contained population champions from all generations. Together, speciation, fitness sharing, and Hall of Fame comprise a state of the art competitive coevolution methodology. The next section describes how progress in competitive coevolution can be monitored.

4.2 MONITORING PROGRESS

We will track progress in coevolution with the master tournament and dominance tournament techniques. First, the master tournament method is reviewed , followed by a discussion of how the superiority of one strategy over another can be established. Using this definition of superiority, a dominance tournament is used to reveal a ranking of increasingly sophisticated strategies.

4.2.1 Master Tournament

In *master tournament*, the champion of each generation is compared to all other generation champions (Floreano and Nolfi 1997). The master tournament is an extension of *CIAO* (*current individual vs. ancestral opponents*; Cliff and Miller 1996), in which every champion plays against all preceding champions. By counting the number of wins by each generation champion against every other generation champion, it is possible to see whether progress is being made over time.

In order to track progress in either the master tournament or dominance tournament, we need to be able to tell whether one strategy is better than another. Because a single trial evaluation may not be accurate, it is necessary for each champion to play multiple trials to reduce error. In the master tournament, the error for a particular generation champion is reduced by playing against every other generation champion. However, the master tournament does not reveal whether a specific ranking of increasingly sophisticated strategies exists. Thus, the question remains, what if we wanted specific information about which strategy is superior to which in a single pairing?

It is possible to reduce the error in specific comparisons using multiple trials between the same two strategies with slightly different initial conditions. However, this approach can take too much time for the master tournament. Given n generations and t trials for each of the $\frac{n^2-n}{2}$ pairings, a complete master tournament would take a total of $t(\frac{n^2-n}{2})$ trials. If both t and n are high, the duration of the tournament can take weeks or months. Thus, for the master tournament, it is necessary to keep t down to one or two. For the robot competition, we use $t = 2$, such that one trial is played from each of the two starting positions.

4.2.2 Dominance Tournament

Unlike the master tournament, the dominance tournament, as explained below, does not examine all $\frac{n^2-n}{2}$ possible pairings of champions, and therefore can afford $t > 100$. Thus, for the dominance tournament, networks were compared on 144 different food configurations from each side of the board, or 288 total. The food configurations included the same 9 symmetrical food positions used during training, plus an additional 2 food items, which were placed in one of 12 different positions on the east and west halves of the board. Some starting food positions give an initial advantage to one robot or another, depending on how close they are to the robots' starting positions. We say that network a *is superior* to network b if a wins more comparisons than b out of the 288 total. The high number of trials affords an accurate measure of superiority.

The dominance tournament is based on the philosophy that in order to track strategic innovation, we need to identify *dominant strategies*, i.e. those that defeat *all previous* dominant strategies. This way, we can make sure that evolution proceeds by developing a progression of strictly more powerful strategies, instead of e.g. switching between alternative ones.

Let a *generation champion* be the winner of a 288 game comparison between the two population champions of a single generation. Let d_j be the jth dominant strategy to appear in the evolution. Then dominance is defined recursively:[1]

- The first dominant strategy d_1 is the generation champion of the first generation;
- dominant strategy d_j, where $j > 1$, is a generation champion such that for all $i < j$, d_j is superior to d_i (i.e. wins the 288 game comparison).

This strict definition of dominance prohibits circularities. For example, d_4 must be superior to strategies d_1 through d_3, d_3 superior to both d_1 and d_2, and d_2 superior to d_1. We call d_n the nth dominant strategy of the run. The entire process of deriving a dominance ranking from a population is a *dominance tournament*, where competitors play all previous dominant strategies until they either lose a 288 game comparison, or win every comparison to previous dominant strategies, thereby becoming a new dominant strategy. Dominance tournaments require significantly fewer comparisons than the master tournament or CIAO techniques, because there are usually far fewer dominant strategies than total strategies in the population.

Dominant strategies must only defeat prior dominant strategies. It is possible that some prior strategy can defeat a dominant strategy. However, such strategies would not invalidate the dominance tournament: They simply indicate that some idiosyncratic strategies exist that are able to defeat specific dominant strategies. They do not themselves belong in the ranking, because they cannot defeat the entire ranking. Such a result is hardly surprising; many cases exist in natural evolution where an unsophisticated organism (such as a parasite) is optimized to defeat a specific higher organism (Hotez and Pritchard 1995; Walden 1991). Such cases can be viewed as a natural consequence of evolution even when rankings of increasing sophistication emerge.

Note that although the dominance tournament in this paper is applied to a competition in which opponents play equivalent, interchangeable roles, it is also applicable when populations are coevolved for different roles, such as predators

[1] Our definition of dominance is similar though not identical to the definition of a *transitive chain* (Rosin 1997, p.19).

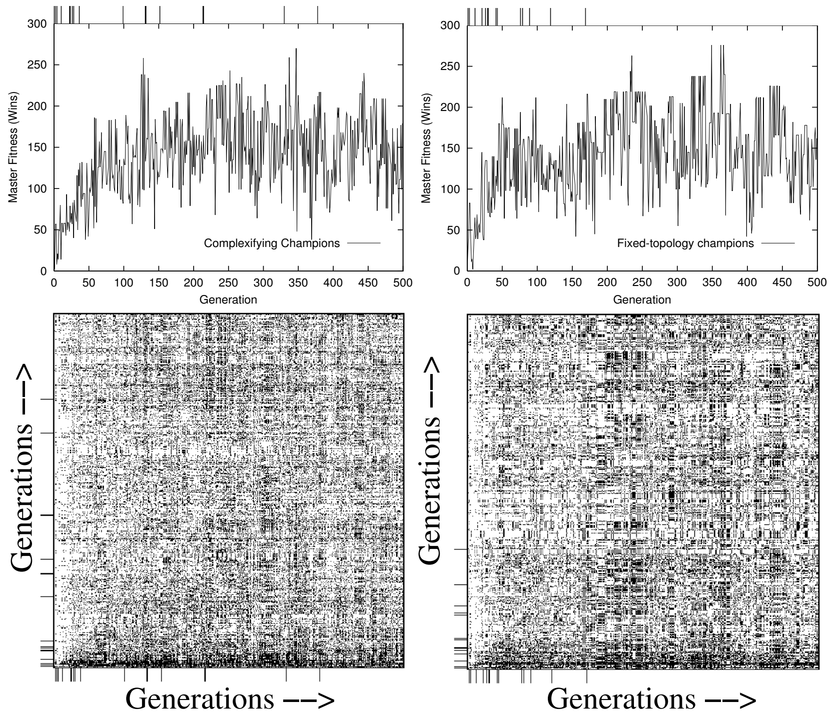

(a) Complexifying Coevolution

(b) Fixed−topology Coevolution

Figure 2: **Master Tournament and Dominance Tournament Results.** The graphs on the left side depict results from complexifying coevolution. On the right are results from fixed-topology coevolution. The upper graphs are plots of master fitness, which is the total number of wins for each generation champion against all other generation champions. The lower graphs are shaded for every point where a generation champion on the x axis defeated another generation champion in a two-trial comparison. From these graphs, it is difficult to tell whether either coevolution methodology produced better results. In contrast, dominance tournament results, represented as tick marks on the graphs for every generation in which a new dominant strategy appeared, reveal that new dominant strategies continued to evolve for significantly longer in complexifying coevolution. It is difficult to identify the best individual from the master tournament, whereas it is well defined in the dominance tournament.

and preys. Instead of requiring each new dominant strategy to defeat all previous dominant strategies, a strategy would only have to defeat all previous dominant strategies from the *opposing population*. For example, the first dominant strategy would be a prey from the first generation, and the second dominant strategy would be the first predator to defeat that prey. The third one would be the first prey that the second dominant (predator) strategy could not defeat. The fourth one would be the first predator that can defeat *both* the first dominant (prey) strategy and the third dominant (prey) strategy, and so on. This way, although e.g. a predator cannot be directly compared to another predator, a ranking of dominant strategies can still be constructed by alternating predators and preys in the ranking. Thus, dominance tournament analysis is applicable to a wide range of competitive scenarios.

The next section addresses a key question about monitoring techniques: Does the dominance tournament demonstrate a ranking when other methods do not?

5 RESULTS

Through extensive head-to-head comparisons, we found that complexifying coevolution consistently produces better strategies than fixed-topology coevolution (Stanley and Miikkulainen 2002a). The question for this paper is: How well do the master tournament and dominance tournament methods illustrate this result? In this section, we answer this question by analyzing two typical runs of robot duel evolution.

Figure 2 shows the results of the analyses. Because the opponents play interchangeable roles in the robot duel, the master tournament depicts progress of generation champions over *both* populations. Although the master fitness of the complexifying runs increases for slightly longer to a slightly higher level, it is difficult to judge whether one evolution run was more effective than the other at producing increasingly sophisticated strategies.

In contrast, dominance tournament analysis shows that complexifying coevolution innovated longer and maintained higher performance. Complexifying coevolution produced 17 levels of dominance, in generations 1, 3, 5, 11, 23, 24, 27, 29, 37, 100, 132, 133, 153, 215, 216, 332, and 380 (identified as tick marks in figure 2). In contrast, fixed-topology coevolution produced fewer total levels of dominance and stagnated significantly earlier in evolution. Only 14 levels of dominance arose in this run in generations 1, 3, 11, 21, 26, 29, 30, 41, 43, 76, 79, 89, 119, and 169.

Several important conclusions can be drawn from the dominance tournament that the master tournament does not reveal. First, complexifying coevolution continued finding better strategies for over 200 more generations than fixed-topology evolution. Second, both runs show that in general the higher the level of dominance, the more generations it takes to reach the next level. Thus, the 15th, 16th, and 17th levels of dominance represent more significant steps in sophistication than the 12th, 13th, and 14th, implying that the 17th level of dominance reached by complexifying coevolution is significantly higher than the 14th level reached by fixed-topology coevolution. Indeed, when we ran a 288-trial comparison between these strategies, the 17th dominant strategy won 221 of the 288 trials. Such a comparison, determining the absolute best strategy over two runs, could not be made using data from the master tournament, because it does not identify any specific transition points, and those competitions it does record are prone to error because they only comprise two trials.

A third interesting aspect of evolution revealed by the dominance tournament is the number of times a champion was able to defeat *some* but not all of the dominance ranking. Such failures to enter the ranking represent strategic circularities. It is just such circularities that we hope to avoid in establishing the arms race in competitive coevolution, so being able to identify and count them is a crucial gauge of the efficacy of a method. In the complexifying coevolution run, such circularities occurred 48 times, while they occurred 93 times, almost twice as often, in fixed-topology coevolution. Of those 93, 63 took place *after* the last dominant strategy evolved, indicating that fixed-topology coevolution was much more likely to suffer from circularities, to the point where they stifled its ability to continue to innovate.

Fourth, the master tournament analysis used 124,750 comparisons, while the dominance tournament averaged only 738 comparisons to analyze a single run of 500 generations. By allocating 288 trials per comparison, the dominance tournament utilized a total of 212,400 trials, while the master tournament utilized 249,500 trials at 2 trials per comparison. This comparison makes sense because there is usually a fixed amount of computation time available for the analysis. The time saved in comparisons can be used to run more trials per comparison, thereby achieving significantly better accuracy.

6 DISCUSSION AND FUTURE WORK

Given that the dominance tournament reveals important contrasts between the two runs, why did the master tournament fail to do so? One of the reasons is that although evolution may produce a generally increasing trend in strategic sophistication throughout a run, the trend over a small window of generations suffers from enormous variance. Such a variance makes the graphs in figure 2 particularly difficult to follow. Specific individuals may temporarily exploit some weakness in more general and sophisticated strategies, thereby obfuscating the underlying ranking for a while. In contrast, the dominance tournament identifies generations where lasting strategic transitions took place, as well as those where circularities occurred, giving a more reliable picture of the progress of evolution.

An interesting comparison can be made between the dominance tournament method and Pareto coevolution (Ficici and Pollack 2001; Noble and Watson 2001). Pareto coevolution is a method for guiding selection based on dominance within one generation. A ranking of *Pareto layers* is constructed by iteratively finding all non-dominated individuals in the population. These individuals are removed from the population before the next Pareto layer is found. Thus, the Pareto layers partition a *single generation* into progressively lower rankings, providing a gradient for evolution to climb: The highest-ranking layers reproduce in proportion to their ranking to form the next generation. In contrast, in the dominance tournament, dominance is a ranking between *individual* champions from *different generations* and describes progress throughout evolution. Thus, the two notions of dominance are significantly different, in addition to being used for different purposes.

It may be possible to extend Pareto-layer ranking *across generations* using the data returned by a master tournament. The highest ranking Pareto layer would correspond to the final generations of evolution, giving evidence of progress. However, Ficici and Pollack (2001) report that in some cases the highest ranking layer may contain as much as 75% of the population. Thus, such analysis would be quite coarse-grained compared to the dominance tournament. In addition, like the master tournament, Pareto-layer ranking requires every strategy to play every other strategy, making it computationally expensive. Third, a Pareto-layer ranking of champions would not identify specific transition points that could be used for further analysis, unless the partitioning happens to perfectly correspond to contiguous generations. Still, Pareto coevolution and the dominance tournament both demonstrate that dominance rankings are well-suited to analysis of coevolutionary problems.

An important question for any analysis method is how it can be applied across many runs in order to make results more reliable. Information about progress returned by the dominance tournament can be combined in four ways (Stanley and Miikkulainen 2002a):

- The *highest level of dominance* achieved can be averaged across all trials of a particular method. The average level reached by one method can be compared to that of another, suggesting that one method continues to innovate longer than the other.

- The highest dominant strategy of each run can be directly compared to the entire dominance ranking in every run of a competing method by playing the strategies against each other. The result indicates the *equivalent dominance level* i.e. the highest level that one method can defeat from the dominance rankings of another method on average. If one method can defeat only the first few tiers of another method's dominance ranking, then the second method is better.

- Complementing the equivalent dominance level, *equiv-*

alent generation represents the average number of generations it takes for a superior method to find a strategy that defeats the highest level of dominance achieved by an inferior method.

- Any population statistic can be correlated with increasing levels of dominance. For example, it is possible to graph the average number of nodes and connections in the highest dominant strategies from each generation of NEAT evolution, revealing that as dominance increases, so does network complexity. Thus, the dominance tournament affords a variety of ways to analyze performance across multiple trials.

Another important question is how the method should be initialized, i.e. how the first dominant strategy should be chosen. In this paper, the champion of the first generation is used for this purpose. This choice is natural because its strategy is poor enough that almost any generally superior strategy can defeat it. Any such poor strategy could be used without changing the results. However, it may be better to make the initialization more general. For example, several champions from the first few generations could be used as the first tier of strategies that must be defeated in order to enter the ranking. This way, the first step is harder to achieve, possibly making the lower levels of the ranking more meaningful. Such choices will be evaluated in future work.

7 CONCLUSION

The dominance tournament and master tournament provide complementary information. The master tournament gives a general feeling for the progress of a run, and the dominance tournament gives the specific details necessary for drawing strong conclusions. Most importantly, whereas it is difficult to identify the best individual from a master tournament, the result is well defined in a dominance tournament. Had we used the master tournament alone, it is likely we would have concluded that neither methodology led to a better evolution run, even though the best individuals from complexifying coevolution were much better than the best of the fixed-topology run. From the dominance tournament, we were able to conclude further that the arms race continued for significantly longer and to a significantly higher level of sophistication in complexifying coevolution compared to fixed-topology coevolution. Third, the dominance tournament does not require comparing every champion to every other champion, and it allows quickly testing specific claims about levels of sophistication by making specific comparisons between different runs. This way, the dominance tournament allows us to gain significant new insight into how progress is made in coevolution.

Acknowledgments

This research was supported in part by the NSF under grant IIS-0083776 and by the Texas Higher Education Coordinating Board under grant ARP-003658-476-2001. Thanks to Michael Shao for pointers to the biology literature.

References

Cliff, D., and Miller, G. F. (1995). Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In Moran, F., Moreno, A., Merelo, J. J., and Chacon, P., editors, *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*, Lecture Notes in Computer Science 929. Berlin; New York: Springer-Verlag.

Dawkins, R., and 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., and Pollack, J. B. (2001). Pareto optimality in coevolutionary learning. In Kelemen, J., editor, *Sixth European Conference on Artificial Life*. Berlin; New York: Springer-Verlag.

Floreano, D., and Nolfi, S. (1997). God save the red queen! Competition in co-evolutionary robotics. *Evolutionary Computation*, 5.

Goldberg, D. E., and Richardson, J. (1987). Genetic algorithms with sharing for multimodal function optimization. In Grefenstette, J. J., editor, *Proc. of the Second Intl. Conf. on Genetic Algorithms*, 148–154. San Francisco, CA: Morgan Kaufmann.

Gomez, F., and Miikkulainen, R. (1997). Incremental evolution of complex general behavior. *Adaptive Behavior*, 5:317–342.

Hotez, P. J., and Pritchard, D. I. (1995). Hookworm infection. *Scientific American*, 272(6):68–74.

Jim, K.-C., and Giles, C. L. (2000). Talking helps: Evolving communicating agents for the predator-prey pursuit problem. *Artificial Life*, 6(3):237–254.

Miller, G., and Cliff, D. (1994). Co-evolution of pursuit and evasion i: Biological and game-theoretic foundations. Technical Report CSRP311, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK.

Mondada, F., Franzi, E., and Ienne, P. (1993). Mobile robot miniaturization: A tool for investigation in control algorithms. In *Proceedings of the Third International Symposium on Experimental Robotics*, 501–513.

Noble, J., and Watson, R. A. (2001). Pareto coevolution: Using performance against coevolved opponents in a game as dimensions for parerto selection. In et al, L. S., editor, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*. San Francisco, CA: Morgan Kaufmann.

Rosin, C. D. (1997). *Coevolutionary Search Among Adversaries*. PhD thesis, Univ. of California, San Diego, San Diego, CA.

Rosin, C. D., and Belew, R. K. (1997). New methods for competitive evolution. *Evolutionary Computation*, 5.

Stanley, K. O., and Miikkulainen, R. (2002a). Continual coevolution through complexification. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2002)*. San Francisco, CA: Morgan Kaufmann.

Stanley, K. O., and Miikkulainen, R. (2002b). Efficient evolution of neural network topologies. In *Proc. of the 2002 Congress on Evolutionary Computation (CEC'02)*. IEEE. In press.

Stanley, K. O., and Miikkulainen, R. (2002c). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2). In press.

Van Valin, L. (1973). A new evolutionary law. *Evolution Theory*, 1:1–30.

Walden, J. (1991). Parasitic diseases. Other round worms. Trichuris, hookworm, and strongyloides. *Primary Care*, 18(1):53–74.