The Dominance Tournament Method of Monitoring Progress in Coevolution

Kenneth O. Stanley and Risto Miikkulainen

The University of Texas at Austin

kstanley,risto@cs.utexas.edu

July 9, 2002

What Happened?

- A central issue with coevolution experiments
- We want to know if an arms race took place
- Master Tournament (Cliff and Miller 1995; Floreano and Nolfi 1997) is the most common analysis method
  - Every generation champion is compared to every other generation champion
  - Results show whether wins increase over generations
- But does that demonstrate an arms race?
- Does it provide insight into what happened?

Master Tournament Accuracy

- Monitoring n generations of progress requires $\frac{n^2-n}{2}$ comparisons
- 500 generations require 124,750 comparisons
- Accurate comparison between two strategies requires multiple trials
- A single multi-trial comparison then may require 1 minute
- 124,750 comparisons then take almost 3 months!
- $\Rightarrow$ We must focus on relevant comparisons

Master Tournament can be Misleading

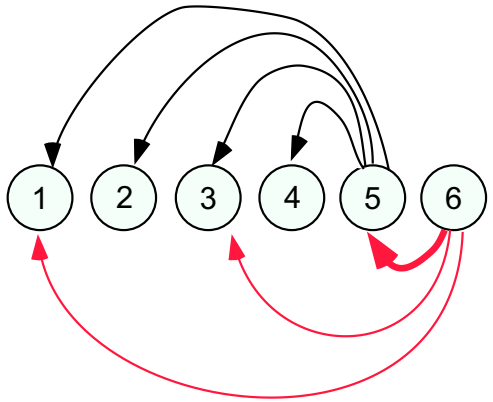

- Defeating more champions does not establish superiority

The Dominance Tournament

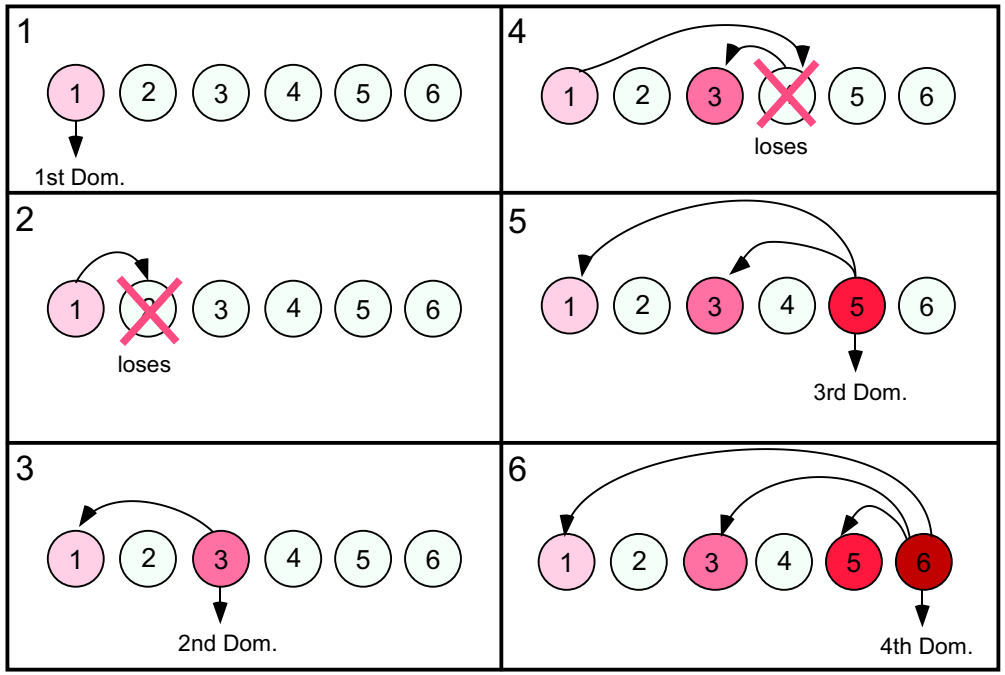

- The first dominant strategy d_1 is the generation champion of the first generation;
- dominant strategy d_j, where $j > 1$, is a generation champion such that for all $i < j$, d_j is superior to (wins a multi-trial comparison with) d_i.

Dominance Tournament Advantages

- Fewer comparisons allows more games per comparison
- Worst case 500 dom. levels would be perfect run
- Prohibits circularities
- Well-defined winner
- Reveals specific generations where transitions occurred

What if Something Non-dominant Can Defeat a Dominant Strategy?

- No reason for concern: Happens in the natural world (e.g. parasites)
- Only shows some idiosyncratic strategy is tuned to defeat a specific dominant strategy
- Such strategies cannot defeat the entire dominance ranking

Experimental Comparison

- Monitoring Progress with Dominance Tournament and Master Tournament
- Analysis of two methods: Fixed-Topology and Complexifying Coevolution of Neural Networks
- Complexifying coevolution does much better (Stanley and Miikkulainen 2002a)
- What differences do the analysis methods reveal?

Experimental Platform: Robot Duel

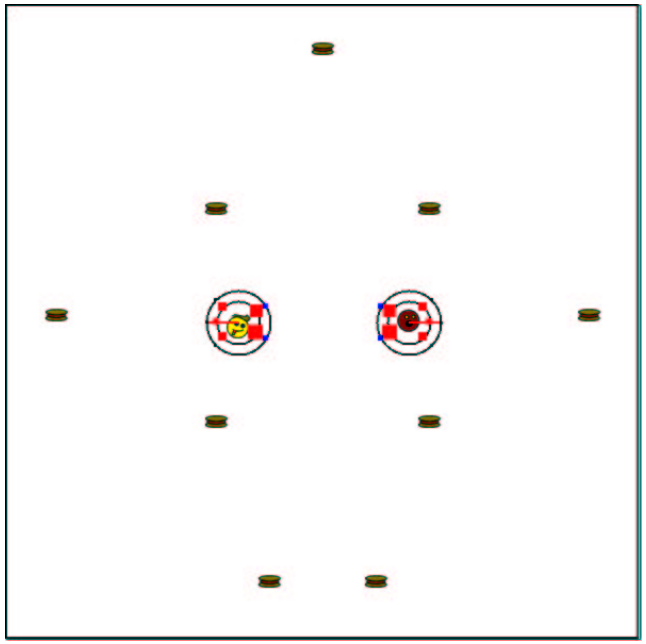

- Robot with higher energy wins by colliding with opponent
- Moving costs energy
- Collecting food replenishes energy
- Complex task: When to forage/save energy, avoid/pursue?

Complexifying System: NEAT

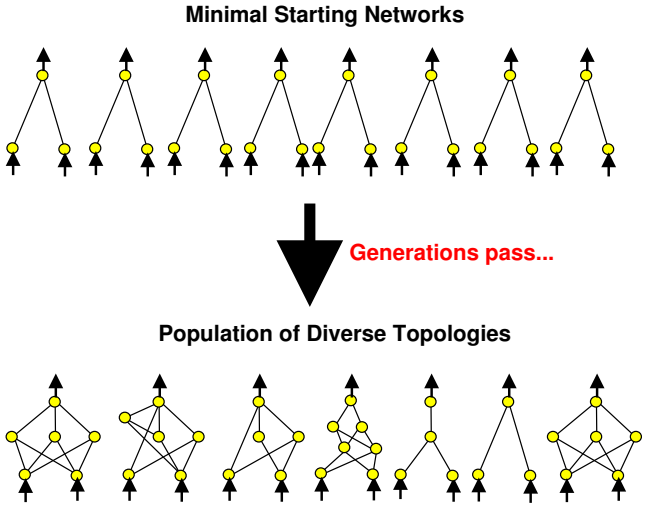

- Evolution of increasingly optimal and complex neural nets
- Increasing complexity allows elaboration
- Begins with networks with no hidden nodes
- New structure is introduced as necessary (Stanley and Miikkulainen 2002b)
- $\Rightarrow$ Implements *complexifying coevolution*

Experimental Setup

- 1 complexifying run and 1 fixed-topology run
- 500 generations per run
- host/parasite coevolution with hall of fame
- Master Tournament and Dominance Tournament analysis

Master Tournament Results

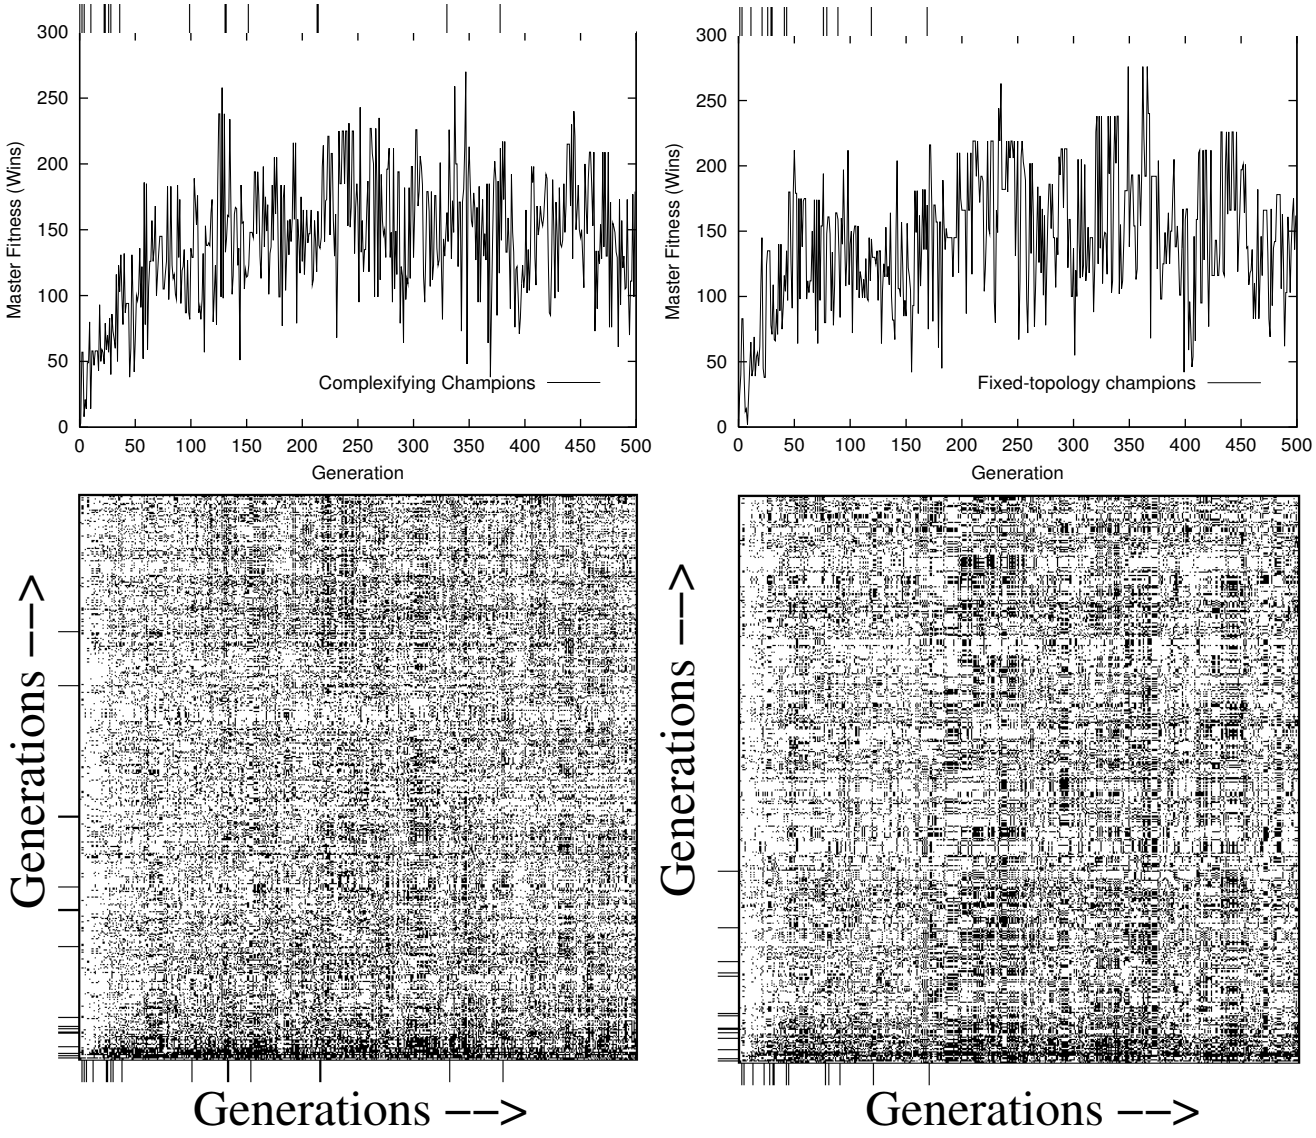

(a) Complexifying Coevolution

(b) Fixed−topology Coevolution

Dominance Tournament Results

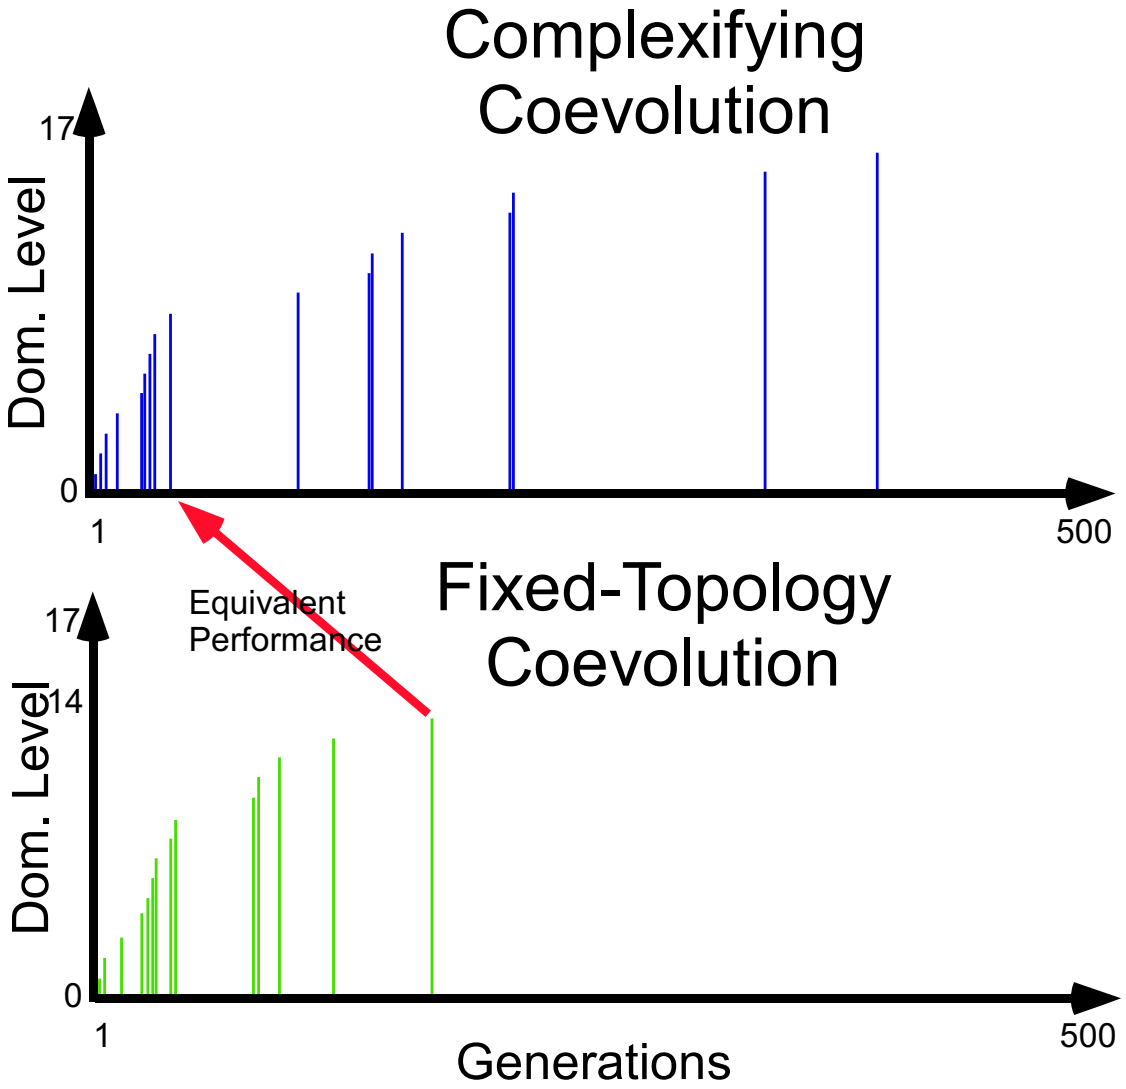

Dominance Tournament Results Summary

- Complexifying coevolution:
  - Achieves 3 more levels of dominance than FT
    8 levels beyond the FT winner
  - Innovates for over 200 more generations
  - Wins 221 of 288 trials over FT winner
  - DT provides a detailed comparison of performance

Discovering Strategic Circularities

- Contending strategies that defeat *some* but not all the ranking starting from the highest dominant strategies
- If a, b are dominant strategies and c came later, then $c > b$ and $c < a \Rightarrow a < b < c < a$
- We want to avoid such circularities in the arms race
- Only DT can detect them since it maintains a ranking
- 48 circularities in complexifying coevolution
- 93 circularities in fixed-topology coevolution, 63 after the last dominant strategy evolved

Comparing Analysis Complexities

- DT: 738 comparisons, 288 trials each = 212,400 total trials
- MT: 124,750 comparisons, 2 trials each = 249,500 total trials
- DT allows much better comparisons for the same computation time

Discussion: DT with Differing Roles

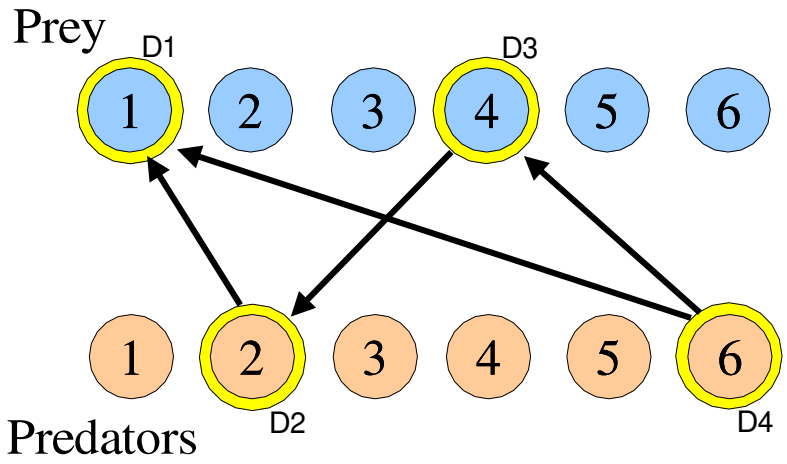

- Example: Predators cannot directly play themselves
- Solution: Modified criteria for entering ranking
  - New dominant strategy must defeat all previous dominant strategies from *opposing population*
  - Ranking alternates between predators and prey
- $\Rightarrow$ DT applies to wide range of competitions

Discussion: Applying DT over Multiple Runs

- *Average highest level of dominance*: compare across different methods
- *Equivalent dominance level*: highest ranking that one method can defeat from another's ranking
- *Equivalent generation*: Average number of generations for superior method to defeat highest dominant of inferior method
- Correlation of any statistic with dominance

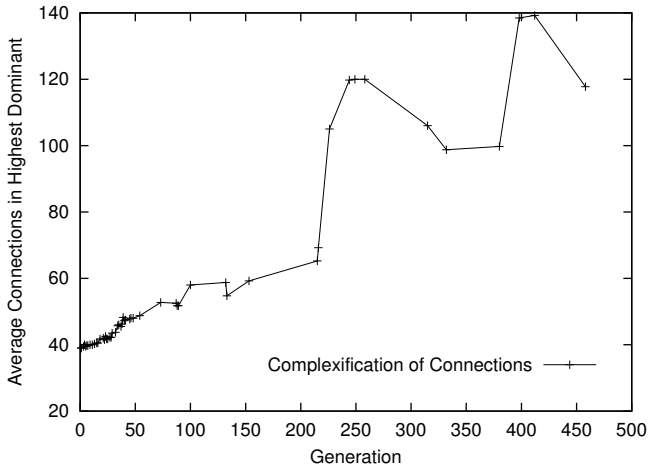

Conclusion

- Dominance Tournament provides specific details for drawing strong conclusions:
  - Indicates best individual from run, allowing direct comparisons
  - Measures duration of and level reached in arms race
  - Locates specific transition points
  - Detects strategic circularities
  - Requires significantly fewer comparisons
- $\Rightarrow$ Useful new method for analyzing coevolution

Pareto Coevolution and the Dominance Tournament

- Pareto coevolution (Ficici and Pollack 2001; Noble and Watson 2001) also uses a notion of dominance
- Uses *Pareto dominance* to rank a single generation into layers
- x Pareto dominates y if x performs better than y against at least one member of the opposing population, and x performs at least as well y against every other opponent
- Different conception of dominance from DT
- However, it may offer its own potential for analysis