Coevolution of rules and topology in cellular automata

(Article begins on next page)

25 April 2024

Darabos C. , Mackenzie C.O., Tomassini M., Giacobini M., Moore J.H.
Coevolution of rules and topology in cellular automata

in

Coevolution of Rules and Topology in Cellular Automata

Christian Darabos
The Geisel School of Medicine
at Dartmouth College
Computational Genetics Lab.
Lebanon, NH, USA
Christian.Darabos@
dartmouth.edu

Craig O. Mackenzie
Institute for Quantitative
Biomedical Sciences
Dartmouth College
Hanover, NH, USA
Craig.O.Mackenzie.gr@
dartmouth.edu

Marco Tomassini
Information Systems Dept.
Faculty of Business and
Economics
University of Lausanne,
Switzerland
Marco.Tomassini@unil.ch

Mario Giacobini
Computational Epidemiology
Group at Department of
Veterinary Sciences Complex
Systems Unit at Molecular
Biotechnology Center
University of Torino, Italy
Mario.Giacobini@unito.it

Jason H. Moore
The Geisel School of Medicine
at Dartmouth College
Computational Genetics Lab.
Lebanon, NH, USA
Jason.H.Moore@
dartmouth.edu

ABSTRACT

Coevolution is nature's response to highly complex and rapidly changing conditions. Biological systems are able to have multiple traits evolving concurrently to adapt to their environment. For many years, evolutionary algorithms have been successfully used on cellular automata (CA) to produce performant update functions. The resulting CAs are, however, much slower and more sensitive to perturbations than CAs with an evolved topology and fixed uniform update rule. Unfortunately, these are not nearly as performant, and suffer from scaling up the number of cells. We propose a hybrid paradigm that simultaneously coevolves the supporting network and the update functions of CAs. The resulting systems combine the high performance of the update evolution and the robustness properties and speed of the topology evolution CAs. Coevolution in CAs a viable tradeoff between the two single trait evolutions.

Categories and Subject Descriptors

F.1.1 [**Models of Computation**]: Unbounded-action devices—*cellular automata*; I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search—*Heuristic methods*; H.3.4 [**Data Structures**]: Graphs and networks

Keywords

cellular automata, artificial evolution, network, topology

1. INTRODUCTION

In biology, coevolution refers to the concurrent or sequential mutation in one or more organisms driven by one or more changes in a related biological object. All members/features taking part in coevolution exert mutual selective pressure on each other, influencing each other's evolution. When taking place within a single biological entity, coevolution is beneficial for the entire organism, where traits coevolve in order to produce individuals with a higher degree of "fitness".

Cellular automata (CAs) have been used for years as a rudimentary proxy for biological organisms. In a prominent study, Mitchell *et al.* have successfully used genetic algorithms (GAs) to artificially evolve a single feature, the update function shared by all cells of small radius linear CAs [1] to perform the *density classification* task. More recently, we have conducted a study evolving a different property of CAs, the underlying network topology of CAs, with comparable success [2]. The resulting evolved topologies are general graphs that exhibit social network properties. In this work, we study a new framework for CA evolution consisting of the simultaneous evolution of both a single update function shared by all cells (uniform CA) and the supporting network topology of the CAs. We hypothesize that evolutionary algorithms (EA) will generate individuals with a high capacity to solve the task at hand, and develop network topologies supporting higher robustness and resilience to transient failures than that of strictly regular CAs [2].

In this work, we propose a new framework allowing simultaneous evolution of the update function and the underlying topology of the CA. Coevolution starts with a population of uniform regular CAs, each with a randomly generated lookup table. Selected parents at each generation will produce mutated offspring, which may replace its parents in the subsequent generation. Mutations affect the topology of the CA, by rewiring edges, and the update function, by modifying the lookup table. Rewiring causes one end of the edge to connect to a different cell, thus changing the neighborhood. The size of the lookup table might need to be adapted to the growing sizes of the neighborhoods. Indeed, every time the size of the largest neighborhood increases by 1, the table size doubles, and the new half of the lookup table is completed with randomly selected values with equal probability.

CAs are evolved for the density task. The initial population is made of uniform regular (i.e. ring) CAs with a radius $r = 2$. The termination criterion is satisfied when

the entire population has reached an optimal fitness or a *fitness plateau*, where improvement becomes marginal or null. The fitness is defined as the fraction of instances for which the CA converges to the known density of the initial configuration (IC). At each generation a different set of ICs is generated for each individual. We use a combination of binary tournaments over pairs of individuals picked by *roulette wheel* selection to generate the next generation. The loser has a small probability of $1 - p$ that it will the next population, which helps to maintain the diversity, explore larger portions of the solution landscape, and escape local maxima. At the end of each evolutionary process, we select the CA in the 95^{th} percentile of *performance* (elite population, EP) by evaluating the entire population on 1,000 instances on ICs with a density of $\rho \approx 0.5$. We replicate all experiments 100 times.

2. DISCUSSION & CONCLUSIONS

In order to compare results with previously proposed evolutionary CAs frameworks, we conduct parallel simulations for all possible combinations of frameworks (rule only, topology only, and coevolution) and sizes $N \in \{99, 199, 299\}$. Figure 1 shows the results of the performance evaluation of the elite population (EP). Each column represents a combination of frameworks (UFE, TE, or CE) and CA sizes. Performance, just like fitness, is normalized to represent the fraction of correctly classified IC after at most $2N$ time steps. We show the consistency of the EP results by showing the absolute best performance, the average performance, and the lowest performance. The number in parentheses is the size of the EP for each case, and gives an estimate of how rich the solution space is in "good individuals".
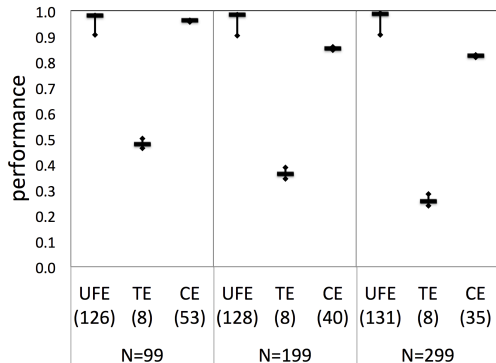

Figure 1: Performance of Best Evolved CAs.

CAs with a UFE only consistently show the highest average and top performance above 0.98, they are also the most steady across all studied CA sizes. However, UFE also shows the widest deviation between the lowest and the highest performances within the EP. CE CAs come close to performing as well as UFE, and their deviation is much narrower. Additionally, the performance clearly degrades as the CAs scale up. Finally, CA evolving the topology-only (TE) exhibit the lowest performance across the board, below the performance of a random uniform draw. Moreover, they are strongly negatively impacted by the scaling of the CA. These results are also reflected by the number of individuals in each EP. UFE easily evolves excellent CAs for the density problem as over

80% of the final population yield a fitness in the 95^{th} percentile. This leads us to conclude the UFE solution space is the smallest, and also the richest in good individuals. The solution space of CE is much larger for the EA to explore, and good individuals are fewer as only about 1/5 to 1/3 of the final population are in the EP. Finally, the most difficult solution space would be that of topology-only evolution, as they yield the lowest performance, and the smallest EP at around 5% of the final population.

We analyze the time progression of our evolutionary CAs over the 100 generations. We track the development of UFE, TE, and CE populations' average fitnesses and the fitness of the best individual in the CA population. We notice that the general trends are similar across all combinations of framework and size. UFE reaches its fitness plateau the fastest, and is overtaken by TE, although it is much slower. Best performances are virtually indistinguishable.

EA's only goal is to optimize the performance of the CAs. However, looking beyond the performance, we are interested in studying the properties emerging from the evolved topologies, and how they differ when obtained solely by topology evolution, and when the network's evolution is combined with adaptations of the update function. The degree distributions reveal a widening of the curve for TE, and a shift of the peak from all nodes having a degree $k = 4$ to a majority of degree $k = 5$. Although there is some spread in the degrees, the function is narrowly centered around its peak, with little deviation, and no extreme values. Coevolution has, on the other hand, facilitated a larger heterogeneity in the degree distribution, with a significantly wider bell shaped curve, and no clear peak at a single value of the degree k. Additionally, the average path length (APL) is significantly shorter than expected, and clustering coefficient (CC) significantly elevated. Together, these properties suggest that the network structures emerging from artificial evolution share properties with technological, social, and other "real-world" networks. Cellular Automata are, despite their apparent simplicity, powerful models for distributed computations, provided that an adequate update function can be found. EAs can evolve highly performant rules to solve CA computational problems. However, these CAs are slow and sensitive to perturbations. Evolving the topology will give rise to faster, more robust CAs. When simultaneously evolving the update function and underlying topology, the resulting CAs are interesting tradeoffs in terms of speed, robustness, and performance.

Acknowledgments

3. REFERENCES

[1] M. Mitchell, P. T. Hraber, and J. P. Crutchfield. Revisiting the edge of chaos: evolving cellular automata to perform computations. *Complex Systems*, 7:89–130, 1993.

[2] M. Tomassini, M. Giacobini, and C. Darabos. Performance and robustness of cellular automata computation on irregular networks. *Advances in Complex Systems*, 10:85–110, 2007.