Title: **"Darwin vs. DARPA: Evolution of a Neural Controller to play Tetris"**

PIs: **C. Adami (MSU), D. Knoester (MSU), A. Hintze (MSU), R. Miikkulainen (UT Austin)**

Summary: The goal of the DARPA-funded SyNAPSE project is to build a microprocessor-based machine that mimics many of the characteristics of biological neural networks (i.e., high connectivity, synapse plasticity, and scale). The proposed device is termed a *neuromorphic* machine, and recent simulations have indeed approached biological scale [1]. However, while SyNAPSE has been demonstrated on a variety of tasks, from optical character recognition and classification to control of quadrotor helicopters [2], it has done so at a cost of $42 million [3]. In contrast, the BEACON-funded DvD ("Darwin vs. DARPA") project has demonstrated the evolution of logic circuits that outperform SyNAPSE on optical character recognition for $1/1000$th of the cost [4]. In this phase of the DvD project, we now turn our attention to a more complex task: Evolving logic circuits that play the computer game "Tetris" (see Fig. 1).

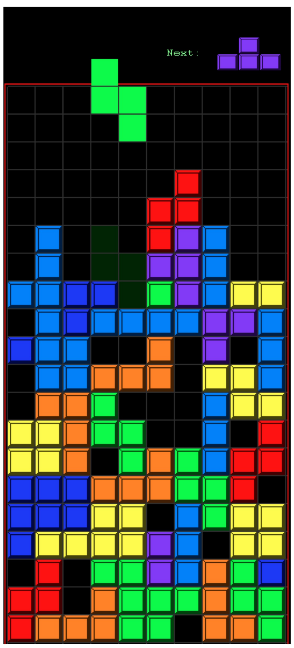

Figure 1: Tetris game, played on a 10×20 playing field

Tetris as a model problem: Two of the central challenges to machine intelligence are: the problem of *planning*, where sequences of actions are strung together to realize some goal [5], as well as the problem of *information integration*, where signals from distinct modalities need to be fused in order to make an informed decision. Tetris is a very different task than optical character recognition. Rather than finding regularities in images (which is essentially a data-compression problem) playing Tetris requires less input bandwidth, but much more signal integration, planning, and temporal synchronization. A successful Tetris player needs to recognize the shape as well as orientation of the falling game pieces (called "tetrominos"), and in the context of the achieved line structure (which depends entirely on play history) make a decision to rotate and/or move the game piece within just a few ticks of the clock. It represents the most advanced of all Active Categorical Perception (ACP) [6] tasks of the Artificial Intelligence literature. We have previously shown that Markov networks [7] can solve ACP tasks, by using memory to predict the trajectory of different falling game pieces (in that case, small or large blocks) [8].

With its well-understood game dynamics [9], Tetris provides an ideal framework for evaluating the performance of planning algorithms as well as information integration, and as such, has frequently been used as a competition platform for machine learning algorithms. However, "solving" Tetris is known to be quite challenging. Indeed, even when the sequence of tetrominos is known in advance, discovering a solution is an NP-complete problem [10]. To date, general purpose machine learning algorithms have not yet been able to compete with algorithms built specifically for playing Tetris. Indeed, the current world record for automated Tetris play is held by a two-piece placement algorithm written by Colin P. Fahey [9], which cleared 7.2M lines. Note that all current AI approaches to Tetris (even evolutionary ones [11]) rely on an evaluation function applied to features extracted from the game board (such as "pile height", "well-depth", "roughness", etc.) to select between possible moves. In our evolutionary approach, features and the evaluation function must be discovered automatically.

Goals: In this second term of funding for the DvD project, we will focus on using evolution to discover logic circuits that can successfully play Tetris. Under this broader goal, we will focus on three distinct milestones: Evolution of single-piece placement algorithms, evolution of two-piece placement algorithms, and finally evolution of head-to-head algorithms.

Single-piece placement algorithms are those that refer only to the current game state and the current piece in order to determine placement. As a proof-of-concept, we will first focus on evolving a single-piece placement algorithm with a fixed sequence of tetrominos. While it is known that sequences for which no solution is possible exist [12], a random piece sequence is unlikely to be unsolvable. Once we have demonstrated that a fixed sequence can be placed, we will progress to evolving solutions for unknown (random) sequences, to be followed by both fixed and random sequences for two-piece placement algorithms, where information about the next piece in the sequence is made available.

Head-to-head placement algorithms are those in which players compete with each other by periodically passing tetrominos to the opposing player. In this case, players attempt to pass pieces that are "maximally difficult" for the opponent to place given the state of their own game. The opportunity exists for evolution to discover remarkably effective strategies for this version of Tetris, and this also provides for a compelling outreach opportunity. Specifically, a web-enabled version of Tetris in which an evolved solution can compete against a human player.

Methods: As with our previous study of optical character recognition, here we will again focus on the use of Markov Networks as well as more standard Artificial Neural Networks to evolve a Tetris-playing algorithm. Markov Networks comprise a series of probabilistic "fuzzy logic" gates that communicate through binary state variables. These networks can be of arbitrary size, though in practice tend to be limited to approximately 5,000 gates and 1,000 state variables. Because these networks encode decisions and have internal (hidden) states, they can be thought of as discrete-time partially observable Markov decision processes (POMDP). The networks we use have two novel properties: First, their interconnection matrix is evolved via a genetic algorithm. Second, the discrete-time update function of each gate can change during the lifetime of an agent via a novel decentralized learning algorithm. This approach, in contrast to others that are commonly used for Bayesian inference learning, is computationally tractable. This enables us to study both evolution and learning concurrently. Placed in the context of Tetris, this means that we will not simply be evolving Markov Networks that are able to play Tetris, but rather, **we will be evolving networks that *learn* to play Tetris.** In order to test whether Markov networks are unique in their ability to evolve complex predictive reasoning algorithms, we will also evolve state-of-the-art Artificial Neutral Networks to perform the same task. ANNs have been previously used to solve ACP [6], recently using a version of ANNs with variable topology [8]. Here, we will test whether augmented topologies (that is, the NEAT implementation [13]) can perform as well (or better) than Markov networks on this task. Using NEAT here is especially interesting because it has previously been used in a video game setting [14].

Team: The team to tackle this project is unchanged from the first year of funding. Besides the PIs Adami, Knoester and Hintze, as well as graduate student Chapman, postdoc Joel Lehmann from PI Risto Miikkulainen's lab will spend a summer at MSU in order to code a neural network implementation for this task. This will allow us to compare the performance of Markov networks with sophisticated Artificial Neural Networks on dynamical perception and action tasks.

Intellectual Merit: The discovery of an evolved solution to a complex planning task like Tetris is of great interest to both the evolutionary computation (EC) and artificial intelligence (AI) communities. In EC, a human-competitive Tetris playing algorithm would demonstrate yet another arena in which evolutionary methods trump human design. However, evolving a Tetris player presents an even more compelling case in AI: While reinforcement learning methods have previously been applied to Tetris, the challenge using these traditional methods is that a payoff (utility) function must be determined in advance. The learning algorithm is then updated with respect to this payoff function. In this case, however, evolution will have implicitly discovered a payoff function that is tuned specifically for playing Tetris. Considering that high-level planning is one of the barriers to the discovery of general-purpose machine intelligence, demonstrating that evolution is able to discover planning algorithms for Tetris may grant insight into how we might produce intelligent machines.

Future Funding: We will submit a proposal to the National Science Foundation to fund this work after the publication of our first major manuscript describing the performance of the evolved logic circuit to classify the MNIST images. A manuscript describing preliminary results is being readied for submission to GECCO 2013 [4]. Ultimately, we plan to approach DARPA to fund this work.

References

[1] R. Preissl, T. Wong, P. Datta, R. Singh, S. Esser, W. Risk, H. Simon, F. Myron, and D. Modha, "Compass: A scalable simulator for an architecture for cognitive computing," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC 2012)*, 2012.

[2] I. Lenz, M. Gemici, and A. Saxena, "Low-power parallel algorithms for single image based obstacle avoidance in aerial robots," in *Proceedings of the IEEE/RSL International Conference on Intelligent Robots and Systems (IROS)*, 2012.

[3] S. Lohr, "Creating artificial intelligence based on the real thing," *New York Times*, p. D8, 2011.

[4] S. Chapman, D. Knoester, A. Hintze, and C. Adami, "Image classification using evolved logic circuits." submitted to Genetic and Evolutionary Computing Conference (GECCO), 2013.

[5] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Series in Artificial Intelligence, Prentice Hall, 2nd ed., 2002.

[6] R. Beer, "The dynamics of active categorical perception in an evolved model agent," *Adaptive Behaviour*, vol. 11, pp. 209–243, 2003.

[7] J. Edlund, N. Chaumont, A. Hintze, C. Koch, G. Tononi, and C. Adami, "Integrated information increases with fitness in the evolution of animats," *PLoS Comp. Biol.*, p. e1002236, 2011.

[8] L. Marstaller, A. Hintze, and C. Adami, "Cognitive systems evolve complex representations for adaptive behavior," *Neural Computation*, p. to appear, 2013.

[9] C. Fahey, "Tetris." http://www.colinfahey.com/tetris/tetris.html, 2003-2012.

[10] E. D. Demaine, S. Hohenberger, and D. Liben-Nowell, "Tetris is hard, even to approximate," in *Proceedings of the 9th International Computing and Combinatorics Conference (COCOON)*, 2003.

[11] N. Böhm, G. Kókai, and S. Mandl, "An evolutionary approach to Tetris," in *MIC2005: The Sixth Metaheuristics International Conference*, 2005.

[12] H. Burgiel, "How to lose at tetris," *Mathematical Gazette*, p. 194, 1997.

[13] K. O. Stanley and R. Miikkulainen, "Competitive coevolution through evolutionary complexification," *Journal of Artificial Intelligence Research*, vol. 21, pp. 63–100, 2004.

[14] K. O. Stanley, B. D. Bryant, and R. Miikkulainen, "Real-time neuroevolution in the NERO video game," *IEEE Transactions on Evolutionary Computation*, vol. 9, pp. 653–668, 2005.