# Real-time Learning in the NERO Video Game

## Kenneth O. Stanley and Ryan Cornelius and Risto Miikkulainen

Department of Computer Sciences, The University of Texas at Austin, Austin, TX 78712-0233 USA

{kstanley, rccorn, risto}@cs.utexas.edu

## Thomas D'Silva

Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX 78712-0240 USA

tdsilva@mail.utexas.edu

## Aliza Gold

Digital Media Collaboratory, IC$^2$ Institute, The University of Texas at Austin, Austin, TX 78705 USA

aliza@icc.utexas.edu

## Abstract

If game characters could learn through interacting with the player, behavior could improve as the game is played, keeping it interesting. The real-time NeuroEvolution of Augmenting Topologies (rtNEAT) method, which can evolve increasingly complex artificial neural networks in real time as a game is being played, will be presented. The rtNEAT method makes possible an entirely new genre of video games in which the player trains a team of agents through a series of customized exercises. In order to demonstrate this concept, the NeuroEvolving Robotic Operatives (NERO) game was built based on rtNEAT. In NERO, the player trains a team of virtual robots for combat against other players' teams. The live demo will show how agents in NERO adapt in real time as they interact with the player. In the future, rtNEAT may allow new kinds of educational and training applications through interactive and adapting games.

## Description

Machine learning can potentially both increase the longevity of video games and decrease their production costs (Fogel, Hays, & Johnson, 2004). Yet machine learning is both one of the most compelling yet least exploited technologies used in interactive digital entertainment. Thus, there is an unexplored opportunity to make video games more interesting and realistic, and to build entirely new genres. Such enhancements may have applications in education and training as well, changing the way people interact with their computers.

In the video game industry, the term *non-player-character* (NPC) refers to an autonomous computer-controlled agent in the game. This demonstration will focus on training NPCs as intelligent agents, and the standard AI term *agents* is therefore used to refer to them. The behavior of such agents in current games is often repetitive and predictable. In most video games, simple scripts cannot learn or adapt to control the agents: Opponents will always make the same moves and the game quickly becomes boring. Machine learning could potentially keep video games interesting by allowing agents to change and adapt. However, a major problem with learning in video games is that if behavior is allowed to change, the game content becomes unpredictable. Agents might learn idiosyncratic behaviors or even not learn at all, making the gaming experience unsatisfying. One way to avoid this problem is to train agents to perform complex behaviors offline, and then freeze the results into the final, released version of the game. However, although the game would be more interesting, the agents still cannot adapt and change in response to the tactics of particular players.

If agents are to adapt and change in real-time, a powerful and reliable machine learning method is needed. This demonstration will introduce such a method, a real-time enhancement of the NeuroEvolution of Augmenting Topologies method (NEAT; Stanley & Miikkulainen, 2002, 2004). NEAT evolves increasingly complex neural networks, i.e. it *complexifies*. Real-time NEAT (rtNEAT) is able to complexify neural networks *as the game is played*, making it possible for agents to evolve increasingly sophisticated behaviors in real time. Thus, agent behavior improves visibly during gameplay. The aim is to show that machine learning is indispensable for an interesting genre of video games, and to show how rtNEAT makes such an application possible.

In order to demonstrate the potential of rtNEAT, the Digital Media Collaboratory (DMC) at the University of Texas at Austin initiated, based on a proposal by Kenneth O. Stanley, the NeuroEvolving Robotic Operatives (NERO; Stanley, Bryant, & Miikkulainen, 2005) project in October of 2003 (http://nerogame.org). The idea was to create a game in which learning is *indispensable*, in other words, without learning NERO could not exist as a game. In NERO, the player takes the role of a trainer, teaching skills to a set of intelligent agents controlled by rtNEAT. Thus, NERO is a powerful demonstration of how machine learning can open up new possibilities in gaming and allow agents to adapt.
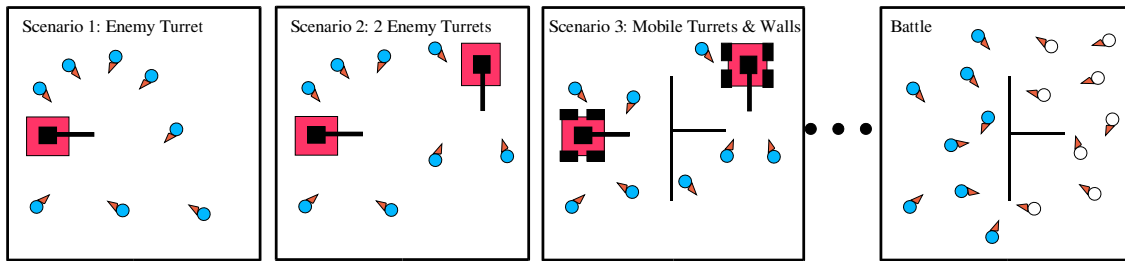
Figure 1: **A turret training sequence.** The figure depicts a sequence of increasingly difficult and complicated training exercises in which the agents attempt to attack turrets without getting hit. In the first exercise there is only a single turret but more turrets are added by the player as the team improves. Eventually walls are added and the turrets are given wheels so they can move. Finally, after the team has mastered the hardest exercises, it is deployed in a real battle against another team.
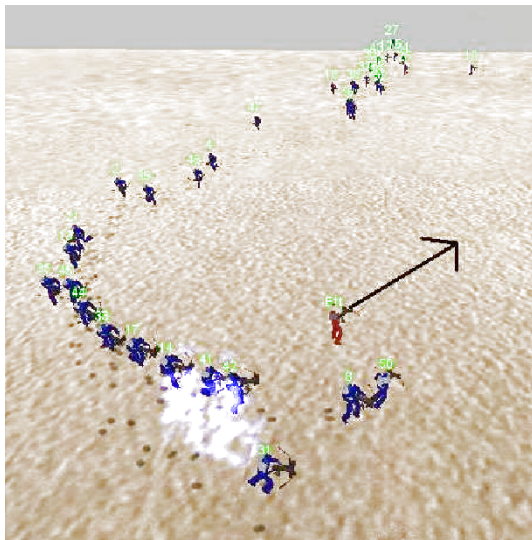


Figure 2: **Avoiding turret fire.** The black arrow points in the current direction of the turret fire (the arrow is not part of the NERO display and is only added for illustration). Agents learn to run safely around turret's fire and attack from behind. When the turret moves, the agents change their attack trajectory accordingly. This behavior shows how evolution can discover behaviors that combine multiple goals.

The learning agents in NERO are simulated robots, and the goal is to train a team of these agents for military combat. The agents begin the game with no skills and only the ability to learn. In order to prepare for combat, the player must design a sequence of training exercises and goals. Ideally, the exercises are increasingly difficult so that the team can begin by learning basic skills and then gradually build on them (figure 1). When the player is satisfied that the team is well prepared, the team is deployed in a battle against another team trained by another player, making for a captivating and exciting culmination of training. The challenge is to anticipate the kinds of skills that might be necessary for battle and build training exercises to hone those skills.

Behavior can be evolved very quickly in NERO, fast enough so that the player can be watching and interacting with the system in real time. The game engine Torque, licensed from GarageGames (`http://www.garagegames.com/`), drives NERO's simulated physics and graphics. Agents can learn diverse skills that are useful for battle while interacting with the player. Skills that players have taught the agents include avoiding an enemy, dodging turret fire (figure 2), and navigating complex mazes without any path-planning algorithm, that is, agents learn to navigate mazes on their own.

NERO creates new opportunities for interactive machine learning in entertainment, education, and simulation. This demonstration will present both rtNEAT and NERO, and review results from the first 1.5 year of this ongoing project, including a fully-playable version to be released on the internet.

## Website with More Information and Demos

The main website for the NERO project is `http://nerogame.org`. The website currently includes project information, screenshots, credits for the numerous volunteers who have contributed work, and movies from live sessions. Later this year, NERO itself will be downloadable through the site so that anyone on the internet can play.

The section of the website with movies can also be conveniently accessed through `http://movies.nerogame.org`.

## References

Fogel, D. B.; Hays, T. J.; and Johnson, D. R. 2004. A platform for evolving characters in competitive games. In *Proceedings of 2004 Congress on Evolutionary Computation*, 1420–1426. Piscataway, NJ: IEEE Press.

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

Stanley, K. O., and Miikkulainen, R. 2004. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research* 21:63–100.

Stanley, K. O.; Bryant, B. D.; and Miikkulainen, R. 2005. Real-time neuroevolution in the nero video game. *IEEE Transactions on Evolutionary Computation Special Issue on Evolutionary Computation and Games*. To appear.