Real-Time Neuroevolution in the NERO Video Game

Kenneth O. Stanley, Bobby D. Bryant, *Student Member, IEEE*, and Risto Miikkulainen

***Abstract*—In most modern video games, character behavior is scripted; no matter how many times the player exploits a weakness, that weakness is never repaired. Yet, if game characters could learn through interacting with the player, behavior could improve as the game is played, keeping it interesting. This paper introduces the real-time Neuroevolution of Augmenting Topologies (rtNEAT) method for evolving increasingly complex artificial neural networks in *real time*, as a game is being played. The rtNEAT method allows agents to change and improve during the game. In fact, rtNEAT makes possible an entirely new genre of video games in which the player *trains* a team of agents through a series of customized exercises. 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. This paper describes results from this novel application of machine learning, and demonstrates that rtNEAT makes possible video games like NERO where agents evolve and adapt in real time. In the future, rtNEAT may allow new kinds of educational and training applications through interactive and adapting games.**

***Index Terms*—Interactive, online, neural networks, neuroevolution, neuroevolution of augmenting topologies (NEAT), neuroevolving robotic operatives (NERO), real-time, video games.**

I. Introduction

THE WORLD video game market in 2002 was valued between $15 billion and $20 billion, larger than even that of Hollywood [1]. Video games have become a facet of many people's lives and the market continues to expand. Because there are millions of interactive players and because video games carry perhaps the least risk to human life of any real-world application, they make an excellent testbed for techniques in artificial intelligence [2]. Such techniques are also important for the video game industry: They can potentially both increase the longevity of video games and decrease their production costs [3].

One of the most compelling yet least exploited technologies is machine learning. 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 *nonplayer-character* (NPC) refers to an autonomous computer-controlled agent in the game. This paper focuses 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 [3]. 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 could not 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 paper describes such a method, a real-time enhancement of the Neuroevolution of Augmenting Topologies (NEAT) method [4], [5]. 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 K. O. Stanley, the Neuroevolving Robotic Operatives (NERO) 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.

NERO opens up new opportunities for interactive machine learning in entertainment, education, and simulation. This paper describes rtNEAT and NERO, and reviews results from the first year of this ongoing project. The next section presents a brief taxonomy of games that use learning, placing NERO in a broader context. NEAT is then described, including how it was enhanced to create rtNEAT. The last sections describe NERO and summarize the current status and performance of the game.

II. Related Work

Early successes in applying machine learning (ML) to board games have motivated more recent work in live-action video

Manuscript received August 13, 2004; revised February 1, 2005, May 27, 2005, and June 6, 2005. This work was supported in part by the Digital Media Collaboratory (DMC) Laboratory at the IC^2 Institute, in part by the National Science Foundation under Grant IIS-0083776, and in part by the Texas Higher Education Coordinating Board under Grant ARP-003658-476-2001.

The authors are with the Department of Computer Sciences, University of Texas at Austin, Austin, TX 78712 USA (e-mail: kstanley@cs.utexas.edu; bdbryant@cs.utexas.edu; risto@cs.utexas.edu).

Digital Object Identifier 10.1109/TEVC.2005.856210

games. For example, Samuel [6] trained a computer to play checkers using a method similar to *temporal difference learning* [7] in the first application of ML to games. Since then, board games such as tic-tac-toe [8], [9], backgammon [10], Go [11], [12], and Othello [13] have remained popular applications of ML (see [14] for a survey). A notable example is Blondie24, which learned checkers by playing against itself without any built-in prior knowledge [15]; also, see [16].

Recently, interest has been growing in applying ML to video games [2], [3]. For example, Fogel *et al.* [3] trained teams of tanks and robots to fight each other using a competitive coevolution system designed for training video game agents. Others have trained agents to fight in first- and third-person shooter games [17]–[19]. ML techniques have also been applied to other video game genres from Pac-Man[1] [20] to strategy games [21]–[23]. This section focuses on how ML can be applied to video games.

From the human player's perspective there are two types of learning in video games. In *out-game learning* (OGL), game developers use ML techniques to pretrain agents that no longer learn after the game is shipped. In contrast, in *in-game learning* (IGL), agents adapt as the player interacts with them in the game; the player can either purposefully direct the learning process or the agents can adapt autonomously to the player's behavior. IGL is related to the broader field of *interactive evolution*, in which a user influences the direction of evolution of, e.g., art, music, or any other kind of phenotype [24]. Most applications of ML to games have used OGL, though the distinction may be blurred from the researcher's perspective when online learning methods are used for OGL. However, the difference between OGL and IGL is important to players and marketers, and ML researchers will frequently need to make a choice between the two.

In a *Machine Learning Game* (MLG), the player *explicitly* attempts to train agents as part of IGL. MLGs are a new genre of video games that require powerful learning methods that can adapt during gameplay. Although some conventional game designs include a "training" phase during which the player accumulates resources or technologies in order to advance in levels, such games are not MLGs because the agents are not actually adapting or learning.

Prior examples in the MLG genre include the *Tamagotchi* virtual pet[2] and the video "God game" *Black & White*.[3] In both games, the player shapes the behavior of game agents with positive or negative feedback. It is also possible to train agents by human example during the game, as van Lent and Laird [25] described in their experiments with *Quake II*.[4] While these examples demonstrated that limited learning is possible in a game, NERO is an entirely new kind of MLG; it uses a reinforcement learning method (neuroevolution) to optimize a fitness function that is dynamically specified *by the player*, while watching and interacting with the learning agents. Thus, agent behavior continues to improve as long as the game is played.

[1]*Pac-Man* is a registered trademark of Namco, Ltd., of Tokyo, Japan.

[2]*Tamagotchi* is a registered trademark of Bandai Co., Ltd., of Tokyo, Japan.

[3]*Black & White* is a registered trademark of Lionhead Studios, Ltd., of Guildford, U.K.

[4]*Quake II* is a registered trademark of Id Software, Inc., of Mesquite, TX.

A flexible and powerful ML method is needed to allow agents to adapt during gameplay. It is not enough to simply script several key agent behaviors because adaptation would then be limited to the foresight of the programmer who wrote the script, and agents would only be choosing from a limited menu of options. Moreover, because agents need to learn online as the game is played, predetermined training targets are usually not available, ruling out supervised techniques such as backpropagation [26] and decision tree learning [27].

Traditional reinforcement learning (RL) techniques such as Q-Learning [28] and Sarsa(λ) with a case-based function approximator (SARSA-CABA; [29]) adapt in domains with sparse feedback [28], [30], [31]. These techniques learn to predict the long-term reward for taking actions in different states by exploring the state space and keeping track of the results. While in principle it is possible to apply them to real-time learning in video games, it would require significant work to overcome several common demands of video game domains.

1) **Large state/action space**. Since games usually have several different types of objects and characters and many different possible actions, the state/action space that RL must explore is extremely high dimensional. Dealing with high-dimensional spaces is a known challenge with RL in general [31], but in a real-time game there is the additional challenge of having to check the value of every possible action on every game tick for every agent in the game. Because traditional RL checks all such action values, the value estimator must execute several times (i.e., once for every possible action) for each agent in the game on every game tick. Action selection may, thus, incur a very large cost on the game engine, reducing the amount of computation available for the game itself.
2) **Diverse behaviors**. Agents learning simultaneously in a simulated world should not all converge to the same behavior: A homogeneous population would make the game boring. Yet because many agents in video games have similar physical characteristics and are evaluated in a similar context, traditional RL techniques, many of which have convergence guarantees [30], risk converging to largely homogeneous solution behaviors. Without explicitly maintaining diversity, such an outcome is likely.
3) **Consistent individual behaviors**. RL depends on occasionally taking a random action in order to explore new behaviors. While this strategy works well in offline learning, players do not want to constantly see the same individual agent periodically making inexplicable and idiosyncratic moves relative to its usual policy.
4) **Fast adaptation and sophisticated behaviors**. Because players do not want to wait hours for agents to adapt, it may be necessary to use a simple representation that can be learned quickly. However, a simple representation would limit the ability to learn sophisticated behaviors. Thus, there is a tradeoff between learning simple behaviors quickly and learning sophisticated behaviors more slowly, neither of which is desirable.
5) **Memory of past states**. If agents remember past events, they can react more convincingly to the present situation.

However, such memory requires keeping track of more than the current state, ruling out traditional Markovian methods. While methods for partially observable Markov processes exist, significant challenges remain in scaling them up to real-world tasks [32].

Neuroevolution (NE), i.e., the artificial evolution of neural networks using an evolutionary algorithm, is an alternative RL technique that meets each of these demands naturally.

1) NE works well in high-dimensional spaces [33]; evolved agents do not need to check the value of more than one action per game tick because agents are evolved to output only a single requested action per game tick.
2) Diverse populations can be explicitly maintained through speciation [4].
3) The behavior of an individual during its lifetime does not change because it always chooses actions from the same network.
4) A *representation* of the solution can be evolved, allowing simple practical behaviors to be discovered quickly in the beginning and complexified later [5].
5) Recurrent neural networks can be evolved that implement and utilize effective memory structures; for example, NE has been used to evolve motor-control skills similar to those in continuous-state games in many challenging non-Markovian domains [5], [15], [33]–[43].

In addition to these five demands, neural networks also make good controllers for video game agents because they can compute arbitrarily complex functions, can both learn and perform in the presence of noisy inputs, and generalize their behavior to previously unseen inputs [44], [45]. Thus, NE is a good match for video games.

There is a large variety of NE algorithms [46]. While some evolve only the connection weight values of fixed-topology networks [37], [47]–[49], others evolve both weights and network topology simultaneously [38], [50]–[62]. Topology and weight evolving artificial neural networks (TWEANNs) have the advantage that the correct topology need not be known prior to evolution. Among TWEANNs, NEAT is unique in that it begins evolution with a population of minimal networks and adds nodes and connections to them over generations, allowing complex problems to be solved gradually based on simple ones.

Our research group has been applying NE to gameplay for about a decade. Using this approach, several NE algorithms have been applied to board games [11], [12], [63], [64]. In *Othello*, NE discovered the *mobility strategy* only a few years after its invention by humans [63]. Recent work has focused on higher-level strategies and real-time adaptation, which are needed for success in both continuous and discrete multiagent games [5], [22], [65]. Using such techniques, relatively simple ANN controllers can be trained in games and game-like environments to produce convincing purposeful and intelligent behavior [5], [11], [36], [40], [65]–[67].

The current challenge is to achieve evolution in *real time*, as the game is played. If agents could be evolved in a smooth cycle of replacement, the player could interact with evolution during the game and the many benefits of NE would be available to the video gaming community. This paper introduces such a real-time NE technique, rtNEAT, which is applied to the NERO multiagent continuous-state MLG. In NERO, agents must master both motor control and higher-level strategy to win the game. The player acts as a trainer, teaching a team of virtual robots the skills they need to survive. The next section reviews the NEAT NE method, and Section IV how it can be enhanced to produce rtNEAT.

III. Neuroevolution of Augmenting Topologies (NEAT)

The rtNEAT method is based on NEAT, a technique for evolving neural networks for complex RL tasks using an evolutionary algorithm (EA). NEAT combines the usual search for the appropriate network weights with *complexification* of the network structure, allowing the behavior of evolved neural networks to become increasingly sophisticated over generations.

The NEAT method consists of solutions to three fundamental challenges in evolving neural network topology.

1) What kind of genetic representation would allow disparate topologies to cross over in a meaningful way? The solution is to use historical markings to line up genes with the same origin.
2) How can topological innovation that needs a few generations to optimize be protected so that it does not disappear from the population prematurely? The solution is to separate each innovation into a different species.
3) How can topologies be minimized *throughout evolution* so the most efficient solutions will be discovered? The solution is to start from a minimal structure and add nodes and connections incrementally.

This section explains how each of these solutions is implemented in NEAT, using the genetic encoding described in the first subsection.

A. Genetic Encoding

Evolving structure requires a flexible genetic encoding. In order to allow structures to complexify, their representations must be dynamic and expandable. Each genome in NEAT includes a list of *connection genes*, each of which refers to two *node genes* being connected (Fig. 1). Each connection gene specifies the in-node, the out-node, the weight of the connection, whether or not the connection gene is expressed (an enable bit), and an *innovation number*, which allows finding corresponding genes during crossover.

Mutation in NEAT can change both connection weights and network structures. Connection weights mutate as in any NE system, with each connection either perturbed or not. Structural mutations, which form the basis of complexification, occur in two ways (Fig. 2). Each mutation expands the size of the genome by adding genes. In the *add connection* mutation, a single new connection gene is added connecting two previously unconnected nodes. In the *add node* mutation, an existing connection is split and the new node placed where the old connection used to be. The old connection is disabled and two new connections added to the genome. The connection between the first node in the chain and the new node is given a weight of one, and the connection between the new node and the last node in the chain is given the same weight as the connection being split.

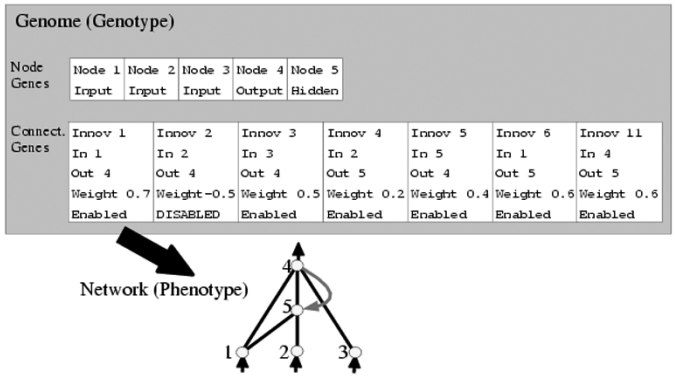

Fig. 1. A NEAT genotype to phenotype mapping example. A genotype is depicted that produces the shown phenotype. There are three input nodes, one hidden node, one output node, and seven connection definitions, one of which is recurrent. The second gene is disabled, so the connection that it specifies (between nodes 2 and 4) is not expressed in the phenotype. The genotype can have arbitrary length, and thereby represent arbitrarily complex networks. Innovation numbers, which allow NEAT to identify which genes match up between different genomes, are shown on top of each gene. This encoding is efficient and allows changing the network structure during evolution.

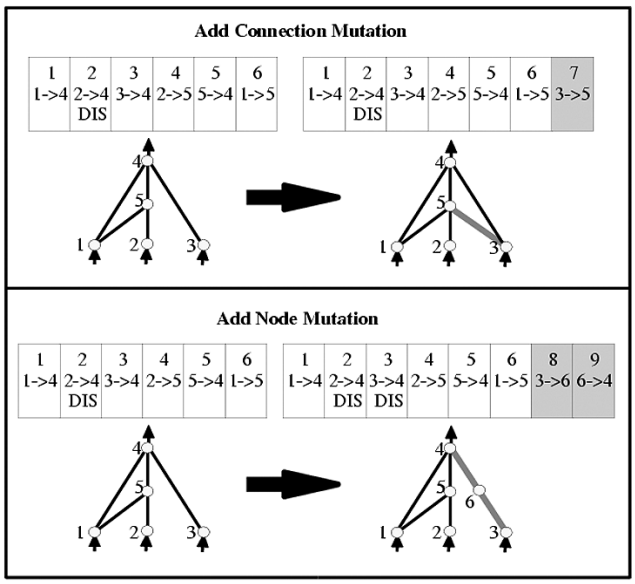

Fig. 2. The two types of structural mutation in NEAT. In each genome, the innovation number is shown on top, the two nodes connected by the gene in the middle, and the "disabled" symbol at the bottom; the weights and the node genes are not shown for simplicity. A new connection or a new node is added to the network by adding connection genes to the genome. Assuming the node is added after the connection, the genes would be assigned innovation numbers 7, 8, and 9, as the figure illustrates. NEAT can keep an implicit history of the origin of every gene in the population, allowing matching genes to be identified even in different genome structures.

Splitting the connection in this way introduces a nonlinearity (the sigmoid function) where there was none before. This nonlinearity changes the function only slightly, and the new node is immediately integrated into the network. Old behaviors encoded in the preexisting network structure are not destroyed and remain qualitatively the same, while the new structure provides an opportunity to elaborate on these original behaviors.

Through mutation, the genomes in NEAT will gradually get larger. Genomes of varying sizes will result, sometimes with different connections at the same positions. Any crossover operator must be able to recombine networks with differing topologies, which can be difficult [68]. The next section explains how NEAT addresses this problem.

B. Tracking Genes Through Historical Markings

The historical origin of each gene can be used to determine exactly which genes match up between *any* individuals in the population. Two genes with the same historical origin represent the same structure (although possibly with different weights), since they were both derived from the same ancestral gene at some point in the past. Thus, in order to properly align and recombine any two disparate topologies in the population, the system only needs to keep track of the historical origin of each gene.

Tracking the historical origins requires very little computation. Whenever a new gene appears (through structural mutation), a *global innovation number* is incremented and assigned to that gene. The innovation numbers, thus, represent a chronology of every gene in the population. As an example, say the two mutations in Fig. 2 occurred one after another. The new connection gene created in the first mutation is assigned the number 7, and the two new connection genes added during the new node mutation are assigned the numbers 8 and 9. In the future, whenever these genomes cross over, the offspring will inherit the same innovation numbers on each gene. Thus, the historical origin of every gene is known throughout evolution.

A possible problem is that the same structural innovation will receive different innovation numbers in the same generation if it occurs by chance more than once. However, by keeping a list of the innovations that occurred in the current generation, it is possible to ensure that when the same structure arises more than once through independent mutations in the same generation, each identical mutation is assigned the same innovation number.

Through innovation numbers, the system now knows exactly which genes match up with which (Fig. 3). Genes that do not match are either *disjoint* or *excess*, depending on whether they occur within or outside the range of the other parent's innovation numbers.

When crossing over, the genes with the same innovation numbers are lined up. The offspring is then formed in one of two ways: In uniform crossover, matching genes are randomly chosen for the offspring genome. In blended crossover [69], the connection weights of matching genes are averaged. These two types of crossover were found to be most effective in NEAT in extensive testing compared with one-point crossover.

The disjoint and excess genes are inherited from the more fit parent, or if they are equally fit, from both parents. Disabled genes have a chance of being reenabled during crossover, allowing networks to make use of older genes once again.

Historical markings allow NEAT to perform crossover without analyzing topologies. Genomes of different organizations and sizes stay compatible throughout evolution, and the variable-length genome problem is essentially avoided. This methodology allows NEAT to complexify structure while different networks still remain compatible.

However, it turns out that it is difficult for a population of varying topologies to support new innovations that add structure

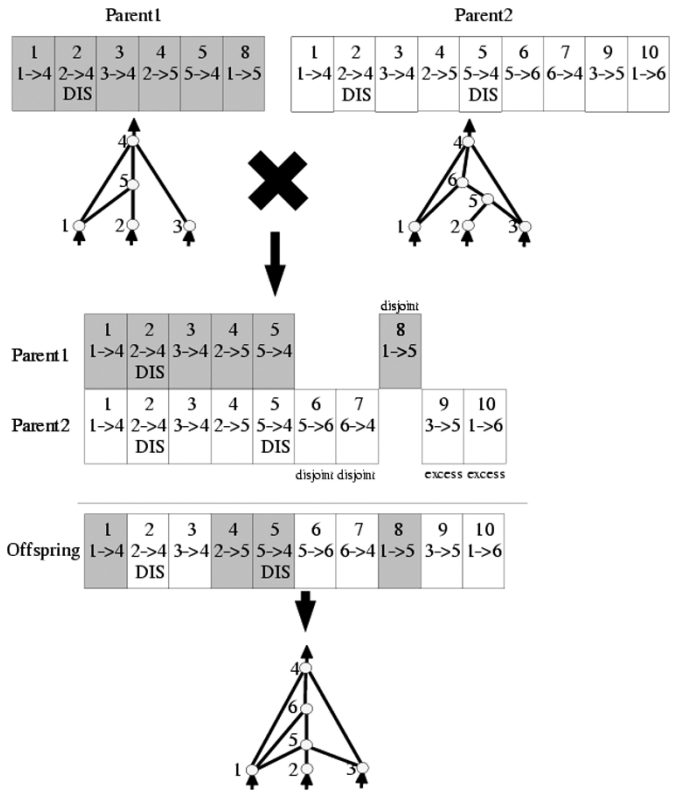

Fig. 3. Matching up genomes for different network topologies using innovation numbers. Although Parents 1 and 2 look different, their innovation numbers (shown at the top of each gene) indicate that several of their genes match up even without topological analysis. A new structure that combines the overlapping parts of the two parents as well as their different parts can be created in crossover. In this case, the two parents are assumed to have equal fitness and, therefore, the offspring inherits all such genes from both parents. Otherwise, these genes would be inherited from the more fit parent only. The disabled genes may become enabled again in future generations: There is a preset chance that an inherited gene is enabled if it is disabled in either parent. By matching up genes in this way, it is possible to determine the best alignment for crossover between any two arbitrary network topologies in the population.

to existing networks. Because smaller structures optimize faster than larger structures, and adding nodes and connections usually initially decreases the fitness of the network, recently augmented structures have little hope of surviving more than one generation even though the innovations they represent might be crucial toward solving the task in the long run. The solution is to protect innovation by speciating the population, as explained in the next section.

C. Protecting Innovation Through Speciation

NEAT speciates the population so that individuals compete primarily within their own niches instead of with the population at large. This way, topological innovations are protected and have time to optimize their structure before they have to compete with other niches in the population. Protecting innovation through speciation follows the philosophy that new ideas must be given time to reach their potential before they are eliminated. A secondary benefit of speciation is that it prevents bloating of genomes: Species with smaller genomes survive as long as their fitness is competitive, ensuring that small networks are not replaced by larger ones unnecessarily.

```
The Genome Loop:
  Take the next genome g from population P
  The Species Loop:
    If all species in S have been checked,
     create new species s_new and place g in it
    Else
     Get the next species s from S
     If g is compatible with s, add g to s
    If g has not been placed,
     continue the Species Loop
    Else exit the Species Loop
  If not all genomes in G have been placed,
    continue the Genome Loop
  Else exit the Genome Loop
```

Fig. 4. Procedure for speciating the population in NEAT. The speciation procedure consists of two nested loops that allocate the entire population into species. Fig. 6 shows how it can be done continuously in real time.

Historical markings make it possible for the system to divide the population into species based on how similar they are topologically (Fig. 4). The distance δ between two network encodings can be measured as a linear combination of the number of excess (E) and disjoint (D) genes, as well as the average weight differences of matching genes $(\bar{W})$

$$\delta = \frac{c_1 E}{N} + \frac{c_2 D}{N} + c_3 \cdot \bar{W}. \tag{1}$$

The coefficients c_1, c_2, and c_3 adjust the importance of the three factors, and the factor N, the number of genes in the larger genome, normalizes for genome size (N can be set to one unless both genomes are excessively large). Genomes are tested one at a time; if a genome's distance to a randomly chosen member of the species is less than δ_t, a compatibility threshold, the genome is placed into this species.

If a genome is not compatible with any existing species, a new species is created. The problem of choosing the best value for δ_t can be avoided by making δ_t *dynamic*; that is, given a target number of species, the system can slightly raise δ_t if there are too many species, and lower δ_t if there are too few. Each genome is placed into the first species from the *previous generation* where this condition is satisfied, so that no genome is in more than one species. Keeping the same set of species from one generation to the next allows NEAT to remove stagnant species, i.e., species that have not improved for several generations.

As the reproduction mechanism, NEAT uses *explicit fitness sharing* [70], where organisms in the same species must share the fitness of their niche. Thus, a species cannot afford to become too big even if many of its organisms perform well. Therefore, any one species is unlikely to take over the entire population, which is crucial for speciated evolution to support a variety of topologies. The adjusted fitness f'_i for organism i is calculated according to its distance δ from every other organism j in the population

$$f'_i = \frac{f_i}{\sum_{j=1}^{n} \mathrm{sh}(\delta(i,j))}. \tag{2}$$

The sharing function sh is set to 0 when distance $\delta(i,j)$ is above the threshold δ_t; otherwise, $\mathrm{sh}(\delta(i,j))$ is set to 1 [71]. Thus, $\sum_{j=1}^{n} \mathrm{sh}(\delta(i,j))$ reduces to the number of organisms in the same species as organism i. This reduction is natural since species are already clustered by compatibility using the

threshold δ_t. Every species is assigned a potentially different number of offspring in proportion to the sum of adjusted fitnesses f'_i of its member organisms.

The net effect of fitness sharing in NEAT can be summarized as follows. Let $\overline{F_k}$ be the average fitness of species k and $|P|$ be the size of the population. Let $\bar{F}_{\text{tot}} = \sum_k \overline{F_k}$ be the total of all species fitness averages. The number of offspring n_k allotted to species k is

$$n_k = \frac{\overline{F_k}}{\bar{F}_{\text{tot}}}|P|. \tag{3}$$

Species reproduce by first eliminating the lowest performing members from the population. The entire population is then replaced by the offspring of the remaining individuals in each species.

The main effect of speciating the population is that structural innovation is protected. The final goal of the system, then, is to perform the search for a solution as efficiently as possible. This goal is achieved through complexification from a simple starting structure, as detailed in the next section.

D. Minimizing Dimensionality Through Complexification

Other systems that evolve network topologies and weights begin evolution with a population of random topologies [38], [46], [54], [58]. In contrast, NEAT begins with a uniform population of simple networks with no hidden nodes, differing only in their initial random weights. Speciation protects new innovations, allowing diverse topologies to gradually accumulate over evolution. Thus, NEAT can start minimally, and grow the necessary structure over generations.

New structures are introduced incrementally as structural mutations occur, and only those structures survive that are found to be useful through fitness evaluations. In this way, NEAT searches through a minimal number of weight dimensions, significantly reducing the number of generations necessary to find a solution, and ensuring that networks become no more complex than necessary. This gradual increase in complexity over generations is similar to complexification in biology [72]–[75]. In effect, then, NEAT searches for the optimal topology by incrementally complexifying existing structure.

E. NEAT Performance

In previous work, each of the three main components of NEAT (i.e., historical markings, speciation, and starting from minimal structure) were experimentally ablated in order to determine how they contribute to performance [4]. The ablation study demonstrated that all three components are interdependent and necessary to make NEAT work.

The NEAT approach is also highly effective: NEAT outperforms other NE methods, e.g., on the benchmark double pole balancing task [4], [76]. In addition, because NEAT starts with simple networks and expands the search space only when beneficial, it is able to find significantly more complex controllers than fixed-topology evolution [5]. These properties make NEAT an attractive method for evolving neural networks in complex tasks such as video games. The next section explains how NEAT can be enhanced to work in real time.

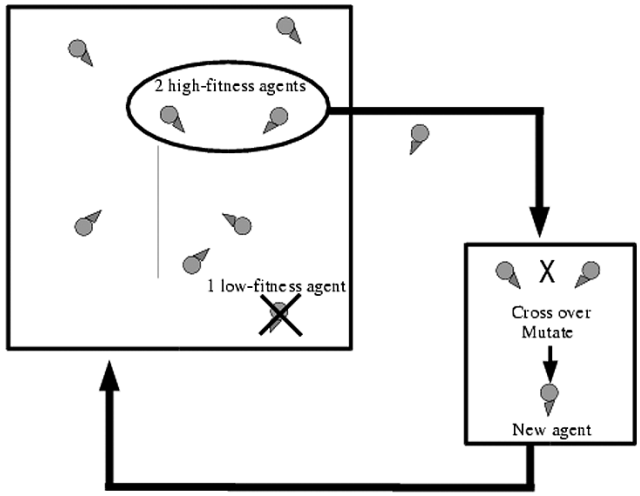

Fig. 5. The main replacement cycle in rtNEAT. NE agents (represented as small circles with an arrow indicating their direction) are depicted playing a game in the large box. Every few ticks, two high-fitness agents are selected to produce an offspring that replaces another of lower fitness. This cycle of replacement operates continually throughout the game, creating a constant turnover of new behaviors that is largely invisible to the player.

IV. Real-Time NEAT (rtNEAT)

Like most EAs, NEAT was originally designed to run *offline*. Individuals are evaluated one or two at a time, and after the whole population has been tested, a new population is created to form the next generation. In other words, in a normal EA it is not possible for a human to interact with the evolving agents while they are evolving. This section describes how NEAT can be modified to make it possible for players to interact with evolving agents *in real time*.

A. Motivation

At each generation, NEAT evaluates one complete generation of individuals before creating the next generation. Real-time NE is based on the observation that in a video game, the entire population of agents plays *at the same time*. Therefore, fitness statistics are collected constantly as the game is played, and the agents could in principle be evolved continuously as well.

The central question is how the agents can be replaced continuously so that offspring can be evaluated. Replacing the entire population together on each generation would look incongruous to the player since everyone's behavior would change at once. In addition, behaviors would remain static during the large gaps of time between generations.

The alternative is to replace a single individual every few game ticks as is done in some evolutionary strategy algorithms [77]. One of the worst individuals is removed and replaced with a child of parents chosen from among the best. If this cycle of removal and replacement happens continually throughout the game (Fig. 5), evolution is largely invisible to the player.

Real-time evolution using continuous replacement was first implemented using conventional NE before NEAT was devel-

```
The rtNEAT Loop:
  Calculate the adjusted fitness of all current
    individuals in the population
  Remove the agent with the worst adjusted
    fitness from the population provided one has
    been alive sufficiently long so that it has
    been properly evaluated.
  Re-estimate the average fitness F̄ for all
    species
  Choose a parent species to create the new
    offspring
  Adjust δ_t dynamically and reassign all agents
    to species
  Place the new agent in the world
```

Fig. 6. Operations performed every n ticks by rtNEAT. These operations allow evolution to proceed continuously, with the same dynamics as in original NEAT.

oped and applied to a Warcraft II[5]-like video game [65]. A similar real-time conventional NE system was later demonstrated by Yannakakis *et al.* [23] in a predator/prey domain. However, conventional NE is not sufficiently powerful to meet the demands of modern video games. In contrast, a real-time version of NEAT offers the advantages of NEAT: Agent neural networks can become increasingly sophisticated and complex during gameplay. The challenge is to preserve the usual dynamics of NEAT, namely protection of innovation through speciation and complexification. While original NEAT assigns offspring to species *en masse* for each new generation, rtNEAT cannot do the same because it only produces one new offspring at a time. Therefore, the reproduction cycle must be modified to allow rtNEAT to speciate in real-time. This cycle constitutes the core of rtNEAT.

B. The rtNEAT Algorithm

In the rtNEAT algorithm, a sequence of operations aimed at introducing a new agent into the population are repeated at a regular time interval, i.e., every n ticks of the game clock (Fig. 6). The new agent will replace a poorly performing individual in the population. The algorithm preserves the speciation dynamics of original NEAT by probabilistically choosing parents to form the offspring and carefully selecting individuals to replace. Each of the steps in Fig. 6 is discussed in more detail below.

1) Calculating Adjusted Fitness: Let f_i be the original fitness of individual i. Fitness sharing adjusts it to $(f_i)/(|S|)$, where $|S|$ is the number of individuals in the species (Section III-C).

2) Removing the Worst Agent: The goal of this step is to remove a poorly performing agent from the game, hopefully to be replaced by something better. The agent must be chosen carefully to preserve speciation dynamics. If the agent with the worst *unadjusted* fitness were chosen, fitness sharing could no longer protect innovation because new topologies would be removed as soon as they appear. Thus, the agent with the worst *adjusted* fitness should be removed, since adjusted fitness takes into account species size, so that new, smaller species are not removed as soon as they appear.

It is also important that agents are evaluated sufficiently before they are considered for removal. In original NEAT, networks are generally all evaluated for the same amount of time. However, in rtNEAT, new agents are constantly being born, meaning different agents have been around for different lengths of time. Therefore, rtNEAT only removes agents that have played for more than the minimum amount of time m. This parameter is set experimentally, by observing how much time is required for an agent to execute a substantial behavior in the game.

Reestimating $\bar{F}$: Assuming there was an agent old enough to be removed, its species now has one less member and, therefore, its average fitness $\bar{F}$ has likely changed. It is important to keep $\bar{F}$ up-to-date because $\bar{F}$ is used in choosing the parent species in the next step. Therefore, $\bar{F}$ needs to be calculated in each step.

3) Creating Offspring: Because only one offspring is created at a time, (3) does not apply to rtNEAT. However, its effect can be approximated by choosing the parent species probabilistically based on the same relationship of adjusted fitnesses

$$Pr(S_k) = \frac{\overline{F_k}}{\bar{F}_{\text{tot}}}. \tag{4}$$

In other words, the probability of choosing a given parent species is proportional to its average fitness compared with the total of all species' average fitnesses. Thus, over the long run, the expected number of offspring for each species is proportional to n_k, preserving the speciation dynamics of original NEAT. A single new offspring is created by recombining two individuals from the parent species.

4) Reassigning Agents to Species: As was discussed in Section III-C, the dynamic compatibility threshold δ_t keeps the number of species relatively stable throughout evolution. Such stability is particularly important in a real-time video game since the population may need to be consistently small to accommodate CPU resources dedicated to graphical processing.

In original NEAT, δ_t can be adjusted before the next generation is created. In rtNEAT, changing δ_t alone is not sufficient because most of the population would still remain in their current species. Instead, the entire population must be reassigned to the existing species based on the new δ_t. As in original NEAT, if a network does not get assigned to any of the existing species, a new species is created with that network as its representative. Depending on the specific game, species do not need to be reorganized at every replacement. The number of ticks between adjustments can be chosen by the game designer based on how rapidly the species evolve. In NERO, evolution progresses rather quickly, and the reorganization is done every five replacements.

5) Replacing the Old Agent With the New One: Since an individual was removed in step Section IV-B2, the new offspring needs to replace it. How agents are replaced depends on the game. In some games (such as NERO), the neural network can be removed from a body and replaced without doing anything to the body. In others, the body may have been destroyed and need to be replaced as well. The rtNEAT algorithm can work with any of these schemes as long as an old neural network gets replaced with a new one.

C. Running the Algorithm

The six-step rtNEAT algorithm is necessary to approximate original NEAT in real-time. However, there is one remaining

[5]*Warcraft II* is a registered trademark of Blizzard Entertainment, of Irvine, California.

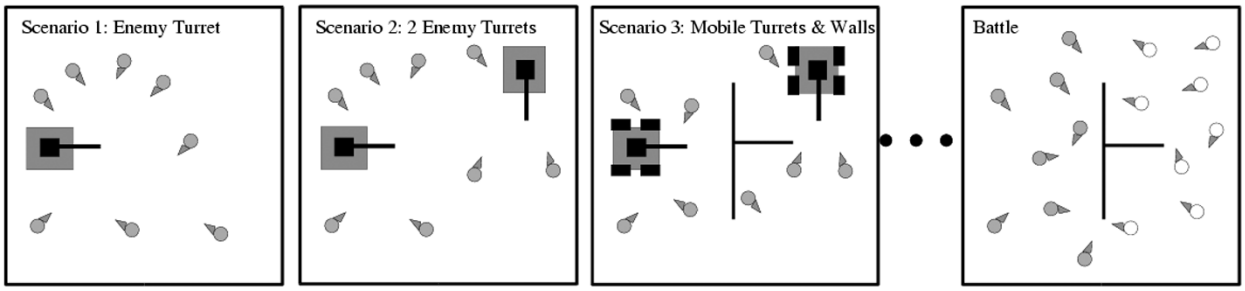

Fig. 7. A turret training sequence. The figure depicts a sequence of increasingly difficult and complicated training exercises in which the agents attempt to attack turrets without getting hit. In the first exercise, there is only a single turret but more turrets are added by the player as the team improves. 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.

issue. The entire loop should be performed at regular intervals, every n ticks: How should n be chosen?

If agents are replaced too frequently, they do not live long enough to reach the minimum time m to be evaluated. On the other hand, if agents are replaced too infrequently, evolution slows down to a pace that the player no longer enjoys.

Interestingly, the appropriate frequency can be determined through a principled approach. Let I be the fraction of the population that is too young and, therefore, cannot be replaced. As before, n is the number of ticks between replacements, m is the minimum time alive, and $|P|$ is the population size. A *law of eligibility* can be formulated that specifies what fraction of the population can be expected to be ineligible once evolution reaches a steady state (i.e., after the first few time steps when no one is eligible)

$$I = \frac{m}{|P|n}. \tag{5}$$

According to (5), the larger the population and the more time between replacements, the lower the fraction of ineligible agents. This principle makes sense since in a larger population it takes more time to replace the entire population. Also, the more time passes between replacements, the more time the population has to age and, hence, fewer are ineligible. On the other hand, the larger the minimum age, the more are below it, and fewer agents are eligible.

It is also helpful to think of $(m)/(n)$ as the *number* of individuals that must be ineligible at any time; over the course of m ticks, an agent is replaced every n ticks, and all the new agents that appear over m ticks will remain ineligible for that duration since they cannot have been around for over m ticks. For example, if $|P|$ is 50, m is 500, and n is 20, 50% of the population would be ineligible at any one time.

Based on the law of eligibility, rtNEAT can decide on its own how many ticks n should lapse between replacements for a preferred level of ineligibility, specific population size, and minimum time between replacements

$$n = \frac{m}{|P|I}. \tag{6}$$

It is best to let the user choose I because in general it is most critical to performance; if too much of the population is ineligible at one time, the mating pool is not sufficiently large. Equation (6) then allows rtNEAT to determine the appropriate number of ticks between replacements. In NERO, 50% of the population remains eligible using this technique.

By performing the right operations every n ticks, choosing the right individual to replace and replacing it with an offspring of a carefully chosen species, rtNEAT is able to replicate the dynamics of NEAT in real-time. Thus, it is now possible to deploy NEAT in a real video game and interact with complexifying agents as they evolve. The next section describes such a game.

V. Neuroevolving Robotic Operatives

NERO is representative of a new MLG genre that is only possible through ML. The idea is to put the player in the role of a *trainer* or a *drill instructor* who teaches a team of agents by designing a curriculum. Of course, for the player to be able to teach agents, the agents must be able to *learn*; rtNEAT is the learning algorithm that makes NERO possible.

In NERO, the learning agents 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 (Fig. 7). 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. The next two sections explain how the agents are trained in NERO and how they fight an opposing team in battle.

A. Training Mode

The player sets up training exercises by placing objects on the field and specifying goals through several sliders (Fig. 8). The objects include static enemies, enemy turrets, rovers (i.e., turrets that move), flags, and walls. To the player, the sliders serve as an interface for describing ideal behavior. To rtNEAT, they represent coefficients for fitness components. For example, the sliders specify how much to reward or punish approaching enemies, hitting targets, getting hit, following friends, dispersing, etc. Each individual fitness component is normalized to a Z-score (i.e., the number of standard deviations from the mean) so that each fitness component is measured on the same scale. Fitness is computed as the sum of all these components multiplied by their

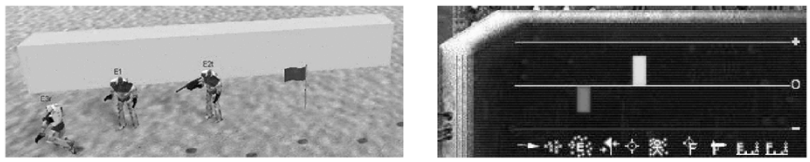

(a) (b)

Fig. 8. Setting up training scenarios. These NERO screenshots show examples of items that the player can place on the field, and sliders used to control the agents' behavior. (a) Three types of enemies are shown from left to right: a rover that runs in a preset pattern, a static enemy that stands in a single location, and a rotating turret with a gun. To the right of the turret is a flag that NERO agents can learn to approach or avoid. Behind these objects is a wall. The player can place any number and any configuration of these items on the training field. (b) Interactive sliders specify the player's preference for the behavior the team should try to optimize. For example, the "E" icon means "approach enemy," and the descending bar above it specifies that the player wants to punish agents that approach the enemy. The crosshair icon represents "hit target," which is being rewarded. The sliders are used to specify coefficients for the corresponding components of the fitness function that NEAT optimizes. Through placing items on the field and setting sliders, the player creates training scenarios where learning takes place.

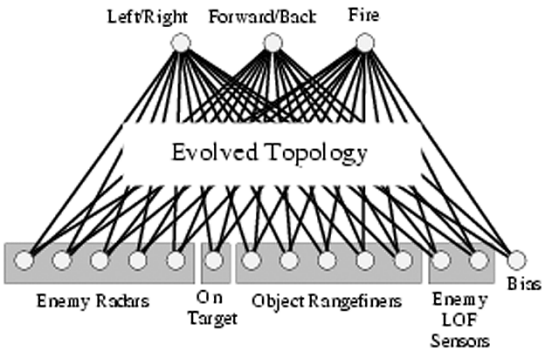

Fig. 9. NERO input sensors and action outputs. Each NERO agent can see enemies, determine whether an enemy is currently in its line of fire, detect objects and walls, and see the direction the enemy is firing. Its outputs specify the direction of movement and whether or not to fire. This configuration has been used to evolve varied and complex behaviors; other variations work, as well and the standard set of sensors can easily be changed in NERO.

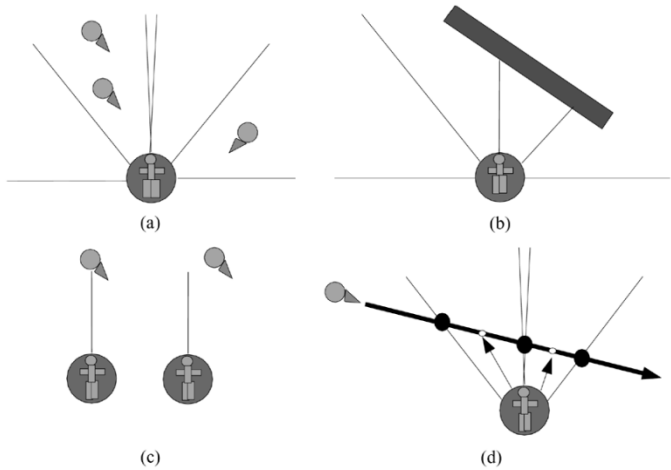

Fig. 10. NERO sensor design. All NERO sensors are egocentric, i.e., they tell where the objects are from the agent's perspective. (a) Several enemy radar sensors divide the 360° around the agent into slices. Each slice activates a sensor in proportion to how close an enemy is within that slice. If there is more than one enemy in it, their activations are summed. (b) Rangefinders project rays at several angles from the agent. The distance the ray travels before it hits an object is returned as the value of the sensor. Rangefinders are useful for detecting long contiguous objects whereas radars are appropriate for relatively small, discrete objects. (c) The on-target sensor returns full activation only if a ray projected along the front heading of the agents hits an enemy. This sensor tells the agent whether it should attempt to shoot. (d) The line of fire sensors detect where a bullet stream from the closest enemy is heading. Thus, these sensors can be used to avoid fire. They work by computing where the line of fire intersects rays projecting from the agent, giving a sense of the bullet's path. Together, these four kinds of sensors provide sufficient information for agents to learn successful behaviors for battle. Other sensors can be added based on the same structures, such as radars for detecting a flag or friendly agents on the same team.

slider levels, which can be positive or negative. Thus, the player has a natural interface for setting up a training exercise and specifying desired behavior.

Agents have several types of sensors. Although NERO programmers frequently experiment with new sensor configurations, the standard sensors include enemy radars, an "on target" sensor, object rangefinders, and line-of-fire sensors. Fig. 9 shows a neural network with the standard set of sensors and outputs, and Fig. 10 describes how the sensors function.

Training mode is designed to allow the player to set up a training scenario on the field where the agents can continually be evaluated while the worst agent's neural network is replaced every few ticks. Thus, training must provide a standard way for agents to appear on the field in such a way that every agent has an equal chance to prove its worth. To meet this goal, the agents spawn from a designated area of the field called the *factory*. Each agent is allowed a limited time on the field during which its fitness is assessed. When their time on the field expires, agents are transported back to the factory, where they begin another evaluation. Neural networks are only replaced in agents that have been put back in the factory. The factory ensures that a new neural network cannot get lucky (or unlucky) by appearing in an agent that happens to be standing in an advantageous (or difficult) position: All evaluations begin consistently in the factory.

The fitness of agents that survive more than one deployment on the field is updated through a diminishing average that gradually forgets deployments from the distant past. A true average is first computed over the first few trials (e.g., 2) and a continuous leaky average (similar to TD(0) RL update [31]) is maintained, thereafter

$$f_{t+1} = f_t + \frac{s_t - f_t}{r} \tag{7}$$

where f_t is the current fitness, s_t is the score from the current evaluation, and r controls the rate of forgetting. The lower r is set, the sooner recent evaluations are forgotten. In this process, older agents have more reliable fitness measures since they are averaged over more deployments than younger agents, but their fitness does not become out of date.

Training begins by deploying 50 agents on the field. Each agent is controlled by a neural network with random connection weights and no hidden nodes, which is the usual starting configuration for NEAT (see Appendix I for a complete description of the rtNEAT parameters used in NERO). As the neural networks

Fig. 11. Battlefield configurations. A range of possible configurations from (a) an open pen to (c) a maze-like environment can be created for NERO. Players can construct their own battlefield configurations and train for them. The basic configuration, which is used in Section VI, is the empty pen surrounded by four bounding walls, as shown in (a).

are replaced in real-time, behavior improves dramatically, and agents eventually learn to perform the task the player sets up. When the player decides that performance has reached a satisfactory level, he or she can save the team in a file. Saved teams can be reloaded for further training in different scenarios, or they can be loaded into battle mode. In battle, they face off against teams trained by an opponent player, as will be described next.

B. Battle Mode

In battle mode, the player discovers how well the training worked out. Each player assembles a battle team of 20 agents from as many different trained teams as desired. For example, perhaps some agents were trained for close combat while others were trained to stay far away and avoid fire. A player may choose to compose a heterogeneous team from both training sessions, and deploy it in battle.

Battle mode is designed to run over a server so that two players can watch the battle from separate terminals on the Internet. The battle begins with the two teams arrayed on opposite sides of the field. When one player presses a "go" button, the neural networks obtain control of their agents and perform according to their training. Unlike in training, where being shot does not lead to an agent body being damaged, the agents are actually destroyed after being shot several times (currently five) in battle. The battle ends when one team is completely eliminated. In some cases, the only surviving agents may insist on avoiding each other, in which case action ceases before one side is completely destroyed. In that case, the winner is the team with the most agents left standing.

The basic battlefield configuration is an empty pen surrounded by four bounding walls, although it is possible to compete on a more complex field with walls or other obstacles (Fig. 11). In the experiments described in this paper, the battlefield was the basic pen, and the agents were trained specifically for this environment. The next section gives examples of actual NERO training and battle sessions.

VI. Playing NERO

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. An important property of the Torque engine is that its physics is slightly nondeterministic, so that the same game is never played twice. In addition, Torque makes it possible for the player to take control of enemy robots using a joystick, an option that can be useful in training.

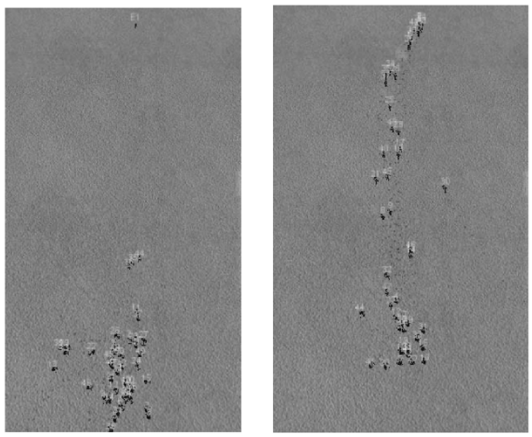

Fig. 12. Learning to approach the enemy. These screenshots show the training field before and after the agents evolved seeking behavior. The factory is at the bottom of each panel and the enemy being sought is at the top. (a) Five seconds: Confusion. Five seconds after the training begins, the agents scatter haphazardly around the factory, unable to effectively seek the enemy. (b) 100 seconds: Success. After 90 seconds, the agents consistently travel to the enemy. Some agents prefer swinging left, while others swing right. These pictures demonstrate that behavior improves dramatically in real-time over only 100 seconds.

The first playable version of NERO was completed in May of 2004. At that time, several NERO programmers trained their own teams and held a tournament. As examples of what is possible in NERO, this section outlines the behaviors evolved for the tournament, the resulting battles, and the real-time performance of NERO and rtNEAT.

A. Training Basic Battle Skills

NERO is capable of evolving behaviors very quickly in real-time. The most basic battle tactic is to aggressively seek the enemy and fire. To train for this tactic, a single static enemy was placed on the training field, and agents were rewarded for approaching the enemy. This training required agents to learn to run toward a target, which is difficult since agents start out in the factory facing in random directions. Starting from random neural networks, it takes on average 99.7 s for 90% of the agents on the field to learn to approach the enemy successfully (ten runs, $sd = 44.5$ s). It is important to note that the criterion for success is partly subjective, based on visually assessing the team's performance. Nevertheless, success in seeking is generally unambiguous, as shown in Fig. 12.

NERO differs from most applications of EAs in that the quality of evolution is judged from the player's perspective based on the performance of the *entire* population, instead of that of the population champion. However, even though the entire population must solve the task, it does not converge to the same solution. In seek training, some agents evolved a tendency to run slightly to the left of the target, while others run to the right. The population diverges because the 50 agents interact as they move simultaneously on the field at the same time. If all the agents chose exactly the same path, they would often crash into each other and slow each other down, so naturally some agents take slightly different paths to the goal. In other words, NERO is actually a massively parallel coevolving ecology in which the entire population is evaluated together.

Fig. 13. Avoiding the enemy effectively. This training screenshot shows several agents running away backward and shooting at the enemy, which is being controlled from a first-person perspective by a human trainer with a joystick. Agents discovered this behavior during avoidance training because it allows them to shoot as they flee. This result demonstrates how evolution can discover novel and effective behaviors in response to the tasks that the player sets up for them.

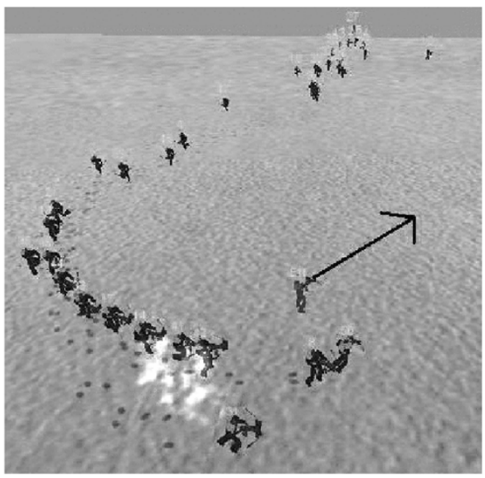

Fig. 14. Avoiding turret fire. The black arrow points in the current direction of the turret fire (the arrow is not part of the NERO display and is only added for illustration). Agents learn to run safely around turret's fire and attack from behind. When the turret moves, the agents change their attack trajectory accordingly. This behavior shows how evolution can discover behaviors that combine multiple goals.

After the agents learned to seek the enemy, they were further trained to fire at the enemy. It is possible to train agents to aim by rewarding them for hitting a target, but this behavior requires fine tuning that is slow to evolve. It is also aesthetically unpleasing to watch while agents fire haphazardly in all directions and slowly figure out how to aim. Therefore, the fire output of neural networks was connected to an aiming script that points the gun properly at the enemy closest to the agent's current heading within a fixed distance of 30 m. Thus, agents quickly learn to seek and accurately attack the enemy.

Agents were also trained to avoid the enemy. In fact, rtNEAT was flexible enough to *devolve* a population that had converged on seeking behavior into a completely opposite, avoidance, behavior.

For avoidance training, players controlled an enemy robot with a joystick and ran it toward the agents on the field. The agents learned to back away in order to avoid being penalized for being too near the enemy. Interestingly, the agents preferred to run away from the enemy backward, because that way they could still see and shoot at the enemy (Fig. 13).

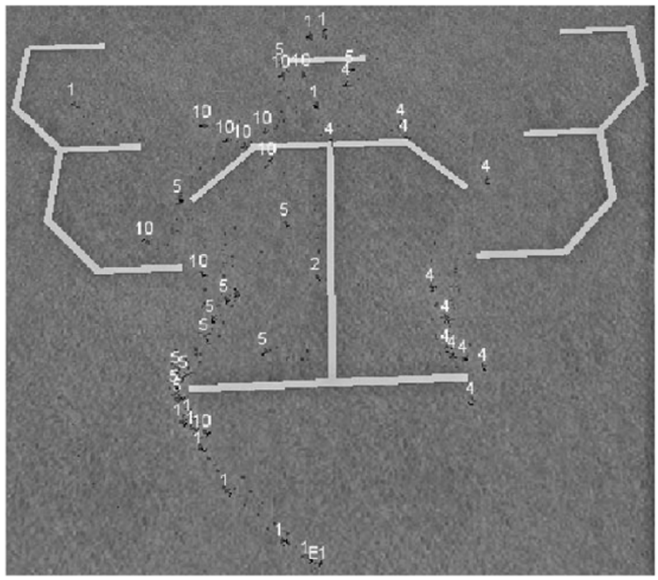

Fig. 15. Navigating a maze. Incremental training on increasingly complex wall configurations produced agents that could navigate this complex maze to find the enemy. The agents spawn from the factory at the top of the maze and proceed down to the enemy at the bottom. In this picture, the numbers above the agents specify their species. Notice that species "4" evolved to take the path through the right side of the maze, while other species evolved to take the left path. This result suggests that protecting innovation in rtNEAT supports a range of diverse behaviors, each with its own network topology.

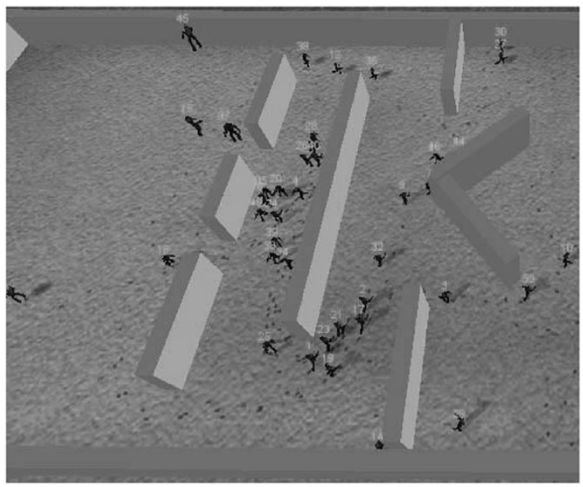

Fig. 16. Successfully navigating different maze configurations. The agents spawn from the left side of the maze and proceed to an enemy at the right. The agents trained to navigate mazes can run through both the maze in Fig. 15 and the maze in this figure, showing that a general path-navigation ability was evolved.

By placing a turret on the field and asking agents to approach it without getting hit, agents were able to learn to avoid enemy fire (Fig. 14). Agents evolved to run to the side that is opposite of the spray of bullets, and approach the turret from behind, a tactic that is promising for battle.

B. Training More Complex Behaviors

Other interesting behaviors were evolved to test the limits of rtNEAT, rather than specifically prepare the troops for battle.

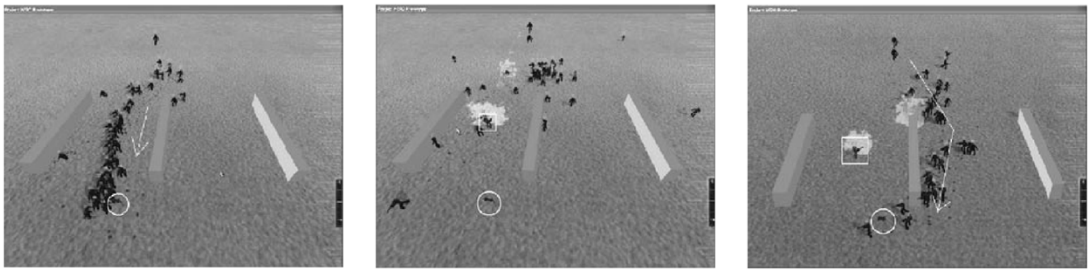

(a) (b) (c)

Fig. 17. Adapting to changing situations. The agents spawn from the top of the screen and must approach the flag (circled) at the bottom left. White arrows point in the direction of their general motion. (a) Agents approach flag. The agents first learn to take the left hallway since it is the shortest path to the flag. (b) Player attacks on left. A human-controlled enemy (identified by a square) attacks inside the left hallway and decimates the agents. (c) Agents learn new approach. The agents learn that they can avoid the enemy by taking the right hallway, which is protected from the enemy's fire by a wall. The rtNEAT method allows the agents to adapt in this way to the player's tactics in real time, demonstrating its potential to enhance a variety of video game genres outside of NERO.

For example, agents were trained to run around walls in order to approach the enemy. As performance improved, players incrementally added more walls until the agents could navigate an entire maze (Fig. 15). This behavior is remarkable because it is successful without any path-planning. The agents developed the general strategy of following any wall that stands between them and the enemy until they found an opening. Interestingly, different species evolved to take different paths through the maze, showing that topology and function are correlated in rtNEAT, and confirming the success of real-time speciation. The evolved strategies were also general enough to navigate significantly varied mazes (Fig. 16).

In a powerful demonstration of real-time adaptation, agents that were trained to approach a designated location (marked by a flag) through a hallway were then attacked by an enemy controlled by the player (Fig. 17). After two minutes, the agents learned to take an alternative path through an adjacent hallway in order to avoid the enemy's fire. While such training is used in NERO to prepare agents for battle, the same kind of adaptation could be used in any interactive game to make it more realistic and interesting.

C. Battling Other Teams

In battle, some teams that were trained differently were nevertheless evenly matched, while some training types consistently prevailed against others. For example, an aggressive seeking team had only a slight advantage over an avoidant team, winning six out of ten battles in the tournament, losing three, and tying one (Table I). The avoidant team runs in a pack to a corner of the field's enclosing wall (Fig. 18). Sometimes, if they make it to the corner and assemble fast enough, the aggressive team runs into an ambush and is obliterated. However, slightly more often the aggressive team gets a few shots in before the avoidant team can gather in the corner. In that case, the aggressive team traps the avoidant team with greater surviving numbers. The conclusion is that seeking and running away are fairly well-balanced tactics, neither providing a significant advantage over the other. The interesting challenge of NERO is to conceive strategies that are clearly dominant over others.

One of the best teams was trained by observing a phenomenon that happened consistently in battle. Chases among agents from opposing teams frequently caused them to eventually reach the field's bounding walls. Particularly for agents trained to avoid turret fire by attacking from behind (Fig. 14), enemies standing against the wall present a serious problem since it is

TABLE I
SEEKERS VERSUS AVOIDERS

Battle Number	Seekers	Avoiders
1	**6**	0
2	4	**7**
3	**8**	0
4	7	7
5	**8**	3
6	6	**10**
7	**5**	4
8	**5**	2
9	3	**7**
10	**8**	0

The number of agents still alive at the end of ten battles are shown between a team trained to aggressively seek and attack the enemy and another team taught to run away backwards and shoot at the same time. The seeking team won six out of the ten games, tied one, and lost three. This outcome demonstrates that even when strategies contrast they can still be evenly matched, making the game interesting. Results like this one can be unexpected, teaching players about relative strengths and weaknesses of different tactics.

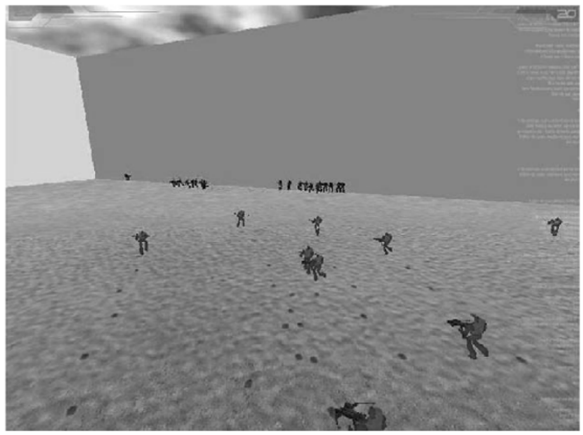

Fig. 18. Seekers chasing avoiders in battle. In this battle screenshot, agents trained to seek and attack the enemy pursue avoidant agents that have backed up against the wall. Teams trained for different tactics are clearly discernable in battle, demonstrating the ability of the training to evolve diverse tactics.

TABLE II
WALL-FIGHTERS VERSUS SEEKERS

Battle Number	Wall-fighters	Seekers
1	**7**	2
2	**9**	0
3	**4**	3
4	**7**	2
5	**10**	0
6	**8**	2
7	**12**	2
8	**7**	2
9	**4**	2
10	**9**	1

The final scores from ten battles between a team trained to fight near walls and another trained to aggressively seek and attack the enemy are shown. The wall-fighters win every battle because they know how to avoid fire near a wall, while the aggressive team runs directly into fire when fighting near a wall. The total superiority of the wall-fighters shows that the right tactical training indeed matters in battle, and that rtNEAT was able to evolve sophisticated fighting tactics.

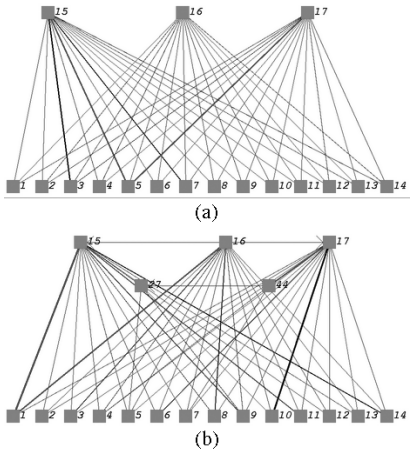

Fig. 19. Simplest successful seeker and wall-fighter networks. Nodes are shown as squares beside their node numbers, and line thickness represents the strength of connections. (a) Seeker. The aggressive seeking strategy is simple enough so that it does not require any hidden nodes. (b) Wall-fighter. Fighting near a wall requires an agent to move away when an enemy points in its direction. This more sophisticated strategy always utilized at least two hidden nodes. These examples demonstrate that NERO can evolve the appropriate network complexity for each desired strategy.

not possible to go around them. Thus training a team against a turret with its back against the wall, it was possible to familiarize agents with attacking enemies that are against a wall. This team learned to hover near the turret and fire when it turned away, but back off quickly when it turned toward them. This tactic works when several agents from the same team are nearby since an enemy can only be facing one direction at a time. In fact, the wall-based team won the first NERO tournament by using this strategy: In particular, it won all games against the aggressive seeking team (Table II).

Fig. 19 shows the simplest networks evolved for successful seekers and wall-fighters. While the simplest seeking network does not include hidden nodes, the simplest wall-trained network utilizes two, demonstrating that more complex networks indeed evolve to produce more sophisticated strategies. Thus, it is possible to learn sophisticated tactics that dominate over simpler ones like seek and avoid.

VII. DISCUSSION

Participants in the first NERO tournament agreed that the game was engrossing and entertaining. Battles were exciting events for all the participants, evoking plentiful clapping and cheering. Players spent many hours honing behaviors and assembling teams with just the right combination of tactics. This experience is promising, suggesting that NERO succeeded as a game, i.e., that it was fun to play it.

An important point of this project is that NERO would not be possible without rtNEAT. With rtNEAT, it was possible to evolve interesting tactics quickly in real-time while players interacted with NERO, showing that NE can be deployed in a real game and work fast enough to provide entertaining results. This result suggests that rtNEAT could also be applied to commercial games. Any game in which agent behavior is repetitive or scripted can potentially be improved by allowing rtNEAT to at least partially modify them in real-time. Especially, in persistent video games such as Massive Multiplayer Online Games (MMOGs) that last for months or years, rtNEAT could be used to continually adapt and optimize agent behavior, thereby permanently altering the gaming experience for millions of players around the world.

Since the first tournament took place, new features have been added to NERO, increasing its appeal and complexity. For example, agents can now duck behind walls and milestones can be set to ensure that previously learned behaviors are not forgotten even after later training for different tasks. The game continues to be developed and new features and sensors are constantly being added. The goal is to have a full network-playable version with an intuitive user interface in the near future.

An important issue for the future is how to assess results in a game in which behavior is largely subjective. One possible approach is to train benchmark teams and measure the success of future training against those benchmarks. This idea and others will be employed as the project matures and standard strategies are identified.

NERO is also being used as a common platform for quickly implementing complicated real-time NE experiments. While video games are intended mainly for entertainment, they are an excellent catalyst for improving ML technology. Because of the gaming industry's financial success and low physical risk, it makes sense to explore gaming as a stepping stone to other more critical applications.

In the long term, the real-time NE technology could be used to adapt the game as human players get better. In this manner, it may finally be possible to use games for training people as has long been envisioned. Such applications would begin with a population of agents with rudimentary skills that provide a gentle initiation to the domain. As the player improves, evolution develops increasingly challenging behaviors in response. Thus, the game becomes more sophisticated at the same rate as

the player improves, a highly desirable property in any training situation.

For example, young children could improve hand-eye coordination by interacting with agents that survive by avoiding being collected or arranged. At first, the agents would need almost no skills to provide a challenge, but as the child improves at the basic task, the agents would need to invent new behaviors. More adult applications include training police or emergency workers in catastrophic situations, or teaching citizens how to evacuate buildings during increasingly hostile conditions. By making such applications possible, rtNEAT creates new opportunities for interactive entertainment and educational media.

VIII. Conclusion

A real-time version of NEAT (rtNEAT) was developed to allow users to interact with evolving agents. In rtNEAT, an entire population is simultaneously and asynchronously evaluated as it evolves. Using this method, it was possible to build a new kind of video game, NERO, where the characters adapt in real time in response to the player's actions. In NERO, the player takes the role of a trainer and constructs training scenarios for a team of simulated robots. The rtNEAT technique can make future video games more interesting and extend their longevity, and eventually make it possible to use gaming as a method for training people in sophisticated tasks.

Appendix
NERO System Parameters

The coefficients for measuring compatibility were $c_1 = 1.0, c_2 = 1.0$, and $c_3 = 0.4$. The initial compatibility distance was $\delta_t = 4.0$. The population was kept small, i.e., 50, so that the CPU could accommodate all the agents being evaluated simultaneously. A target of four species was assigned. If the number of species grew larger than four, δ_t was increased by 0.3. Conversely, if the number of species fell below four, δ_t was decreased by 0.3. The interspecies mating rate was 0.001. The probability of adding a new node was 0.05 and the probability of a new connection mutation was 0.03. These parameter values were found experimentally, but they do follow intuitively meaningful rules: Links need to be added significantly more often than nodes, and weight differences are given low weight since the population is small. Performance is robust to moderate variations in these values.

The percentage of the population allowed to be ineligible at one time, I, was 50%, the number of ticks between replacements was 20, and the minimum evaluation time was 500. The number of ticks between replacements can also be derived from (6). The rate of forgetting (Section V-A) r was 2. The values for these parameters were determined through extensive testing, and again the system was found to be robust against minor variations.

Acknowledgment

Special thanks are due to A. Thibault and A. Cavalli of the IC^2 and DMC for supporting the NERO project. The entire NERO team deserves recognition for their contribution to this project including the NERO leads A. Gold (Producer), P. Flesher (Lead Programmer), J. Perry (Lead Artist), M. Patterson (Lead Designer), and B. Walls (Lead Programmer). The authors are grateful also to the volunteer programmers R. Cornelius, M. Chrien (Lead Programmer as of Fall 2004), B. Fitch, J. Larrabee, T. Ngo, and D. Stewart-Silverman, and to volunteer artists B. Bird, B. Frank, C. Hollins, B. Lowry, I. Minshill, and M. Ward. NERO physics is controlled by the Torque Engine, which is licensed from GarageGames (http://www.garagegames.com/). The rtNEAT software is available freely for research purposes at http://www.nn.cs.utexas.edu/keyword?rtNEAT. NERO's official public website is http://nerogame.org, and the project's official e-mail address is nero@cs.utexas.edu.

References

[1] P. Thurrott. (2002, Jan.) Top stories of 2001, #9: Expanding video-game market brings Microsoft home for the holidays. *Windows NET Mag. Netw.* [Online]. Available: http://www.winnetmag.com/Article/ArticleID/23 862/23 862.html

[2] J. E. Laird and M. van Lent, "Human-level AI's killer application: Interactive computer games," in *Proc. 17th Nat. Conf. Artif. Intell., and 12th Ann. Conf. Innov. Appl. Artif. Intell.*, 2000, pp. 1171–1178.

[3] D. B. Fogel, T. J. Hays, and D. R. Johnson, "A platform for evolving characters in competitive games," in *Proc. Congr. Evol. Comput.*, 2004, pp. 1420–1426.

[4] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evol. Comput.*, vol. 10, no. 2, pp. 99–127, 2002.

[5] ——, "Competitive coevolution through evolutionary complexification," *J. Artif. Intell. Res.*, vol. 21, pp. 63–100, 2004.

[6] A. L. Samuel, "Some studies in machine learning using the game of checkers," *IBM J.*, vol. 3, pp. 210–229, 1959.

[7] R. S. Sutton, "Learning to predict by the methods of temporal differences," *Mach. Learn.*, vol. 3, pp. 9–44, 1988.

[8] D. Michie, "Trial and error," *Penguin Sci. Survey*, vol. 2, pp. 129–145, 1961.

[9] M. Gardner, "How to build a game-learning machine and then teach it to play and to win," *Sci. Amer.*, vol. 206, no. 3, pp. 138–144, 1962.

[10] G. Tesauro and T. J. Sejnowski, "A 'neural' network that learns to play backgammon," in *Neural Information Processing Systems*, D. Z. Anderson, Ed. New York: Amer. Ins. Physics, 1987.

[11] N. Richards, D. Moriarty, P. McQuesten, and R. Miikkulainen, "Evolving neural networks to play Go," in *Proc. 7th Int. Conf. Genetic Algorithms*, T. Bäck, Ed., East Lansing, MI, 1997, pp. 768–775.

[12] K. O. Stanley and R. Miikkulainen, "Evolving a roving eye for Go," in *Proc. Genetic Evol. Computa. Conf.*, 2004, pp. 1226–1238.

[13] T. Yoshioka, S. Ishii, and M. Ito, "Strategy acquisition for the game Othello based on reinforcement learning," in *Proc. 5th Int. Conf. Neural Inf. Process.*, S. Usui and T. Omori, Eds. Tokyo, 1998, pp. 841–844.

[14] J. Fürnkranz, "Machine learning in games: A survey," *Machines that Learn to Play Games*, pp. 11–59, 2001.

[15] D. B. Fogel, *Blondie24: Playing at the Edge of AI*. San Mateo, CA: Morgan Kaufmann, 2001.

[16] D. B. Fogel, T. J. Hays, S. L. Hahn, and J. Quon, "A self-learning evolutionary chess program," *Proc. IEEE*, vol. 92, no. 12, pp. 1947–1954, 2004.

[17] B. Geisler, "An empirical study of machine learning algorithms applied to modeling player behavior in a 'first person shooter' video game," M.S. thesis, Dept. Comput. Sci., Univ. Wisconsin-Madison, Madison, WI, 2002.

[18] N. Cole, S. Louis, and C. Miles, "Using a genetic algorithm to tune first-person shooter bots," in *Proc. Congr. Evol. Comput.*, vol. 1, 2004, pp. 139–145.

[19] J.-H. Hong and S.-B. Cho, "Evolution of emergent behaviors for shooting game characters in robocode," in *Proc. Congr. Evol. Comput.*, vol. 1, 2004, pp. 634–638.

[20] M. Gallagher and A. Ryan, "Learning to play Pac-Man: An evolutionary, rule-based approach," in *Proc. Congr. Evol. Comput.*, vol. 4, 2003, pp. 2462–2469.

[21] T. Revello and R. McCartney, "Generating war game strategies using a genetic algorithm," in *Proc. Congr. Evol. Comput.*, vol. 2, 2002, pp. 1086–1091.

[22] B. D. Bryant and R. Miikkulainen, "Neuroevolution for adaptive teams," in *Proc. Congr. Evol. Comput.*, vol. 3, 2003, pp. 2194–2201.

[23] G. Yannakakis, J. Levine, and J. Hallam, "An evolutionary approach for interactive computer games," in *Proc. Congr. Evol. Comput.*, vol. 1, 2004, pp. 986–993.

[24] I. Parmee and C. Bonham, "Toward the support of innovative conceptual design through interactive designer/evolutionary computing strategies," *Artif. Intell. Eng. Design, Analy. Manuf. J.*, vol. 14, pp. 3–16, 1999.

[25] M. van Lent and J. E. Laird, "Learning procedural knowledge through observation," in *Proc. Int. Conf. Knowl. Capture*, 2001, pp. 179–186.

[26] D. E. Rumelhart and J. L. McClelland, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Cambridge, MA: MIT Press, 1986, vol. 1. PDP Research Group, Foundations.

[27] P. E. Utgoff, "Incremental induction of decision trees," *Mach. Learn.*, vol. 4, no. 2, pp. 161–186, 1989.

[28] C. J. C. H. Watkins and P. Dayan, "Q-learning," *Mach. Learn.*, vol. 8, no. 3, pp. 279–292, 1992.

[29] J. C. Santamaria, R. S. Sutton, and A. Ram, "Experiments with reinforcement learning in problems with continuous state and action spaces," *Adapt. Behav.*, vol. 6, no. 2, pp. 163–218, 1998.

[30] L. P. Kaelbling, M. Littman, and A. W. Moore, "Reinforcement learning: A survey," *J. Artif. Intell.*, vol. 4, pp. 237–285, May 1996.

[31] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.

[32] F. J. Gomez, "Robust non-linear control through neuroevolution," Ph.D. dissertation, Dept. Comput. Sci., Univ. Texas at Austin, Austin, TX, 2003. Tech. Rep. AI-TR-03-303.

[33] F. J. Gomez and R. Miikkulainen, "Active guidance for a finless rocket through neuroevolution," in *Proc. Genetic Evol. Comput. Conf.*, 2003, pp. 2084–2095.

[34] R. Aharonov-Barki, T. Beker, and E. Ruppin, "Emergence of memory-Driven command neurons in evolved artificial agents," *Neural Comput.*, vol. 13, no. 3, pp. 691–716, 2001.

[35] D. Floriano and F. Mondada, "Automatic creation of an autonomous agent: Genetic evolution of a neural-network driven robot," in *Proc. 3rd Int. Conf. Simul. Adapt. Behav. (From Animals to Animats 3)*, 1994, pp. 421–430.

[36] F. Gomez and R. Miikkulainen, "2-D pole-balancing with recurrent evolutionary networks," *Proc. Int. Conf. Artif. Neural Netw.*, pp. 425–430, 1998.

[37] ——, "Solving non-Markovian control tasks with neuroevolution," in *Proc. 16th Int. Joint Conf. Artif. Intell.*, 1999, pp. 1356–1361.

[38] F. Gruau, D. Whitley, and L. Pyeatt, "A comparison between cellular encoding and direct encoding for genetic neural networks," in *Proc. 1st Ann. Conf. Genetic Program.*, J. R. Koza, D. E. Goldberg, D. B. Fogel, and R. L. Riolo, Eds., 1996, pp. 81–89.

[39] I. Harvey, "The artificial evolution of adaptive behavior," Ph.D. dissertation, School of Cognitive and Comput. Sci., Univ. Sussex, Sussex, U.K., 1993.

[40] D. E. Moriarty and R. Miikkulainen, "Evolving obstacle avoidance behavior in a robot arm," in *Proc. 4th Int. Conf. Simul. Adapt. Behav. (From Animals to Animats 4)*, P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack, and S. W. Wilson, Eds., 1996, pp. 468–475.

[41] S. Nolfi, J. L. Elman, and D. Parisi, "Learning and evolution in neural networks," *Adapt. Behav.*, vol. 2, pp. 5–28, 1994.

[42] M. A. Potter, K. A. De Jong, and J. J. Grefenstette, "A coevolutionary approach to learning sequential decision rules," in *Proc. 6th Int. Conf. Genetic Algorithms*, L. J. Eshelman, Ed., 1995, pp. 366–372.

[43] D. Whitley, S. Dominic, R. Das, and C. W. Anderson, "Genetic reinforcement learning for neurocontrol problems," *Mach. Learn.*, vol. 13, pp. 259–284, 1993.

[44] G. Cybenko, "Approximation by superpositions of a sigmoidal function," *Math. Control, Signals, and Systems*, vol. 2, no. 4, pp. 303–314, 1989.

[45] H. T. Siegelmann and E. D. Sontag, "Analog computation via neural networks," *Theor. Comput. Sci.*, vol. 131, no. 2, pp. 331–360, 1994.

[46] X. Yao, "Evolving artificial neural networks," *Proc. IEEE*, vol. 87, no. 9, pp. 1423–1447, 1999.

[47] D. E. Moriarty and R. Miikkulainen, "Efficient reinforcement learning through symbiotic evolution," *Mach. Learn.*, vol. 22, pp. 11–32, 1996.

[48] N. Saravanan and D. B. Fogel, "Evolving neural control systems," *IEEE Expert*, pp. 23–27, Jun. 1995.

[49] A. Wieland, "Evolving neural network controllers for unstable systems," in *Proc. Int. Joint Conf. Neural Netw.*, Seattle, WA, 1991, pp. 667–673.

[50] J. C. Bongard and R. Pfeifer, "Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny," in *Proc. Genetic Evol. Comput. Conf.*, L. Spector, E. D. Goodman, A. Wu, W. B. Langdon, H.-M. Voigt, M. Gen, S. Sen, M. Dorigo, S. Pezeshk, M. H. Garzon, and E. Burke, Eds., 2001, pp. 829–836.

[51] G. S. Hornby and J. B. Pollack, "Creating high-level components with a generative representation for body-brain evolution," *Artif. Life*, vol. 8, no. 3, pp. 223–246, 2002.

[52] M. Mandischer, "Representation and evolution of neural networks," *Artif. Neural Nets and Genetic Algorithms*, pp. 643–649, 1993.

[53] "School of Comput. Sci.," Univ. Birmingham, Birmingham, Birmingham, U.K., Tech. Rep. CSRP-97-7, 1997.

[54] B.-T. Zhang and H. Muhlenbein, "Evolving optimal neural networks using genetic algorithms with Occam's razor," *Complex Syst.*, vol. 7, pp. 199–220, 1993.

[55] H. Braun and J. Weisbrod, "Evolving feed-forward neural networks," in *Proc. ANNGA93, Int. Conf. Artif. Neural Netw. Genetic Algorithms*, 1993, pp. 25–32.

[56] D. W. Opitz and J. W. Shavlik, "Connectionist theory refinement: Genetically searching the space of network topologies," *J. Artif. Intell. Res.*, vol. 6, pp. 177–209, 1997.

[57] X. Yao and Y. Liu, "Toward designing artificial neural networks by evolution," *Appl. Math. Comput.*, vol. 91, no. 1, pp. 83–90, 1996.

[58] P. J. Angeline, G. M. Saunders, and J. B. Pollack, "An evolutionary algorithm that constructs recurrent neural networks," *IEEE Trans. Neural Netw.*, vol. 5, pp. 54–65, 1993.

[59] R. Krishnan and V. B. Ciesielski, "2DELTA-GANN: A new approach to training neural networks using genetic algorithms," in *Proc. Australian Conf. Neural Netw.*, 1994, pp. 194–197.

[60] V. Maniezzo, "Genetic evolution of the topology and weight distribution of neural networks," *IEEE Trans. Neural Netw.*, vol. 5, no. 1, pp. 39–53, Jan. 1994.

[61] C.-H. Lee and J.-H. Kim, "Evolutionary ordered neural network with a linked-list encoding scheme," in *Proc. IEEE Int. Conf. Evol. Comput.*, 1996, pp. 665–669.

[62] D. Dasgupta and D. McGregor, "Designing application-specific neural networks using the structured genetic algorithm," in *Proc. Int. Conf. Combinations Genetic Algorithms Neural Netw.*, 1992, pp. 87–96.

[63] D. Moriarty and R. Miikkulainen, "Evolving complex Othello strategies with marker-based encoding of neural networks," Dept. Comput. Sci., Univ. Texas at Austin, Austin, TX, 1993. Tech. Rep. AI93-206.

[64] D. E. Moriarty, "Symbiotic evolution of neural networks in sequential decision tasks," Ph.D. dissertation, Dept. Comput. Sci., Univ. Texas at Austin, Austin, TX, 1997. Tech. Rep. UT-AI97-257.

[65] A. Agogino, K. O. Stanley, and R. Miikkulainen, "Real-time interactive neuro-evolution," *Neural Process. Lett.*, vol. 11, pp. 29–38, 2000.

[66] D. Moriarty and R. Miikkulainen, "Learning sequential decision tasks," Dept. Comput. Sci., Univ. Texas at Austin, Austin, TX, 1995. Tech. Rep. AI95-229.

[67] D. E. Moriarty and R. Miikkulainen, "Discovering complex Othello strategies through evolutionary neural networks," *Connection Sci.*, vol. 7, no. 3, pp. 195–209, 1995.

[68] N. J. Radcliffe, "Genetic set recombination and its application to neural network topology optimization," *Neural Comput. Appl.*, vol. 1, no. 1, pp. 67–90, 1993.

[69] A. H. Wright, "Genetic algorithms for real parameter optimization," in *Foundations of Genetic Algorithms*, G. J. E. Rawlins, Ed. San Mateo, CA: Morgan Kaufmann, 1991, pp. 205–218.

[70] D. E. Goldberg and J. Richardson, "Genetic algorithms with sharing for multimodal function optimization," in *Proc. 2nd Int. Conf. Genetic Algorithms*, J. J. Grefenstette, Ed., 1987, pp. 148–154.

[71] *Handbook of Evolutionary Computation*, Oxford Univ. Press, London, U.K., 1995. W. Spears, "Speciation using tag bits".

[72] A. Amores, A. Force, Y.-L. Yan, L. Joly, C. Amemiya, A. Fritz, R. K. Ho, J. Langeland, V. Prince, Y.-L. Wang, M. Westerfield, M. Ekker, and J. H. Postlethwait, "Zebrafish HOX clusters and vertebrate genome evolution," *Science*, vol. 282, pp. 1711–1784, 1998.

[73] S. B. Carroll, "Homeotic genes and the evolution of arthropods and chordates," *Nature*, vol. 376, pp. 479–485, 1995.

[74] A. Force, M. Lynch, F. B. Pickett, A. Amores, Y. Lin Yan, and J. Postlethwait, "Preservation of duplicate genes by complementary, degenerative mutations," *Genetics*, vol. 151, pp. 1531–1545, 1999.

[75] A. P. Martin, "Increasing genomic complexity by gene duplication and the origin of vertebrates," *The Amer. Naturalist*, vol. 154, no. 2, pp. 111–128, 1999.

[76] K. O. Stanley and R. Miikkulainen, "Efficient reinforcement learning through evolving neural network topologies," in *Proc. Genetic Evol. Comput. Conf.*, 2002, pp. 569–577.

[77] H.-G. Beyer and H. P. Schwefel, "Evolution strategies—A comprehensive introduction," *Natural Computing*, vol. 1, no. 1, pp. 3–52, 2002. [Online]. Available: http://www.kluweronline.com/article.asp?PIPS=401 777.

Kenneth O. Stanley received the B.S.E. degree *magna cum laude* in computer science engineering and a minor in cognitive science from the University of Pennsylvania, Philadelphia, in 1997, and the M.S. degree in computer science and the Ph.D. degree from the University of Texas at Austin (UT-Austin), in 1999 and 2004, respectively.

He is a Postdoctoral Researcher in the Department of Computer Sciences, University of Texas at Austin. He has published papers in *JAIR*, *Evolutionary Computation*, and *Artificial Life* journals. His research focuses on artificially evolving complex solutions to difficult real-world tasks.

Dr. Stanley has won Best Paper Awards at the 2002 Genetic and Evolutionary Computation Conference, for his work on NEAT, and at the IEEE 2005 Symposium on Computational Intelligence and Games, for his work on NERO.

Bobby D. Bryant (S'03) received the B.A. degree in classical studies (*summa cum laude*) from the University of Houston, Houston, TX, in 1995 and the M.S.C.S. degree in computer sciences from the University of Texas at Austin (UT-Austin) in 1999. He is currently working towards the Ph.D. degree in computer sciences from the University of Texas at Austin.

He is a Researcher at the Digital Media Collaboratory, IC^2 Institute, UT-Austin. His research interests include using machine learning and artificial neural networks to create autonomous intelligent agents for use in games and military simulations.

Risto Miikkulainen received the M.S. degree in engineering from the Helsinki University of Technology, Helsinki, Finland, in 1986, and the Ph.D. degree in computer science from the University of California, Los Angles (UCLA), in 1990.

He is a Professor of Computer Sciences at the University of Texas at Austin (UT-Austin). He has authored *Computational Maps in the Visual Cortex* (Springer, 2005), *Lateral Interactions in the Cortex: Structure and Function* (coedited electronic hypertext book: http://nn.cs.utexas.edu/web-pubs/htmlbook96, 1996), and *Subsymbolic Natural Language Processing* (MIT Press, 1993). He is an Editor of the *Machine Learning Journal* and the *Journal of Cognitive Systems Research*. He is an author of over 180 articles on neuroevolution, connectionist natural language processing, and the computational neuroscience of the visual cortex. His recent research focuses on methods for evolving neural networks and applying these methods to game playing, robotics, and intelligent control.