Evolving Neural Networks

Risto Miikkulainen

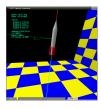
Department of Computer Science The University of Texas at Austin risto@cs.utexas.edu

Forming Decision Strategies



- ► Traditionally designed by hand
 - ► Too complex: Hard to anticipate all scenarios
 - ▶ Too inflexible: Cannot adapt on-line
- ► Need to discover through exploration
 - ► Based on sparse reinforcement
 - Associate actions with outcomes

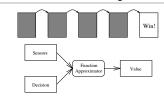
Why Neuroevolution?





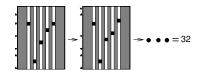
- ► Neural nets powerful in many statistical domains
 - ► E.g. control, pattern recognition, prediction, decision making
 - Where no good theory of the domain exists
- ► Good supervised training algorithms exist
 - ► Learn a nonlinear function that matches the examples
- ► What if correct outputs are not known?

Standard Reinforcement Learning



- ► AHC, Q-learning, Temporal Differences
 - ► Generate targets through prediction errors
 - ► Learn when successive predictions differ
- ► Predictions represented as a value function
- Values of alternatives at each state
- ► Difficult with large/continuous state and action spaces
- ► Difficult with hidden states

Sequential Decision Tasks



- ► POMDP: Sequence of decisions creates a sequence of states
- ► No targets: Performance evaluated after several decisions
- ► Many important real-world domains:
 - ► Robot/vehicle/traffic control
 - ► Computer/manufacturing/process optimization
 - Game playing

Neuroevolution (NE) Reinforcement Learning

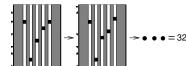


- ► NE = constructing neural networks with evolutionary algorithms
- ► Direct nonlinear mapping from sensors to actions
- ► Large/continuous states and actions easy
 - ► Generalization in neural networks
- ► Hidden states disambiguated through memory
 - ► Recurrency in neural networks⁹⁵

Poles	Method	Evals	Succ.
One	VAPS	(500,000)	0%
	SARSA	13,562	59%
	Q-MLP	11,331	
	NE	127	
Two	NE	3,416	

- ► Difficult RL benchmark: Non-Markov Pole Balancing
- ▶ NE 3 orders of magnitude faster than standard RL²⁸
- ► NE can solve harder problems

Role of Neuroevolution

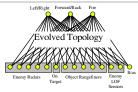


- ► Powerful method for sequential decision tasks ^{16,28,58,111}
 - Optimizing existing tasks
 - Discovering novel solutions
 - Making new applications possible
- ► Also may be useful in supervised tasks ^{53,65}
 - ► Especially when network topology important
- ► A unique model of biological adaptation/development ^{60,75,106}

Outline

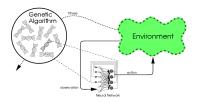
- ► Basic neuroevolution techniques
- Advanced techniques
 - ► E.g. combining learning and evolution; novelty search
- ► Extensions to applications
- ► Application examples
 - ► Control, Robotics, Artificial Life, Games

Neuroevolution Decision Strategies



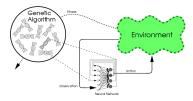
- ► Input variables describe the state observed through sensors
- ► Output variables describe actions
- ► Network between input and output:
 - Nonlinear hidden nodesWeighted connections
- ► Execution:
 - ► Numerical activation of input
 - ► Performs a nonlinear mapping
 - ► Memory in recurrent connections (POMDP!)

Conventional Neuroevolution (CNE) I



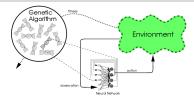
- ► Evolving connection weights in a population of networks ^{53,76,111,112}
- ► Chromosomes are strings of connection weights (bits or real)
 - ► E.g. 10010110101100101111001
 - Usually fully connected, fixed topology
 - Initially random

Conventional Neuroevolution II



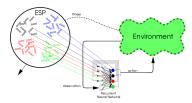
- ► Parallel search for a solution network
 - ► Each NN evaluated in the task
 - ► Good NN reproduce through crossover, mutation
 - ► Bad thrown away
- ► Natural mapping between genotype and phenotype
 - ► GA and NN are a good match!

Problems with CNE



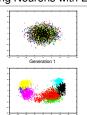
- ► Evolution converges the population (as usual with EAs)
 - ► Diversity is lost; progress stagnates
- ► Competing conventions
 - ► Different, incompatible encodings for the same solution
- ► Too many parameters to be optimized simultaneously
 - ► Thousands of weight values at once

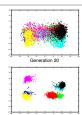
Advanced NE 1: Evolving Partial Networks I



- ► Evolving individual neurons to cooperate in networks 1,57,65
- ► E.g. Enforced Sub-Populations (ESP²³)
 - ► Each (hidden) neuron in a separate subpopulation
 - Fully connected; weights of each neuron evolved
 Populations learn compatible subtasks

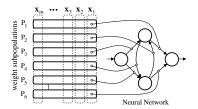
Evolving Neurons with ESP





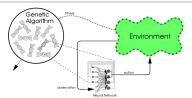
- ► Evolution encourages diversity automatically
 - ► Good networks require different kinds of neurons
- ► Evolution discourages competing conventions
- Neurons optimized for compatible roles
- ► Large search space divided into subtasks
 - Optimize compatible neurons

Evolving Partial Networks II



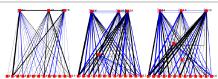
- ► Extend the idea to evolving connection weights
- ► E.g. Cooperative Synapse NeuroEvolution (CoSyNE²⁸)
 - ► Connection weights in separate subpopulations
 - Networks formed by combining neurons with the same index
 - Networks mutated and recombined; indices permutated
- ► Sustains diversity, results in efficient search

Advanced NE 2: Evolutionary Strategies



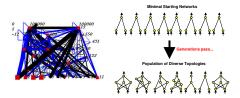
- ► Evolving complete networks with ES (CMA-ES³⁵)
- ► Small populations, no crossover
- ► Instead, intelligent mutations
 - ► Adapt covariance matrix of mutation distribution
 - ► Take into account correlations between weights
- ► Smaller space, less convergence, fewer conventions

Advanced NE 3: Evolving Topologies



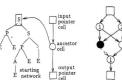
- ► Optimizing connection weights and network topology 3,16,21,113
- ► E.g. Neuroevolution of Augmenting Topologies (NEAT^{86,89})
- ► Based on Complexification
- ► Of networks:
 - ► Mutations to add nodes and connections
- ► Of behavior:
 - · Elaborates on earlier behaviors

Why Complexification?



- ► Problem with NE: Search space is too large
- ► Complexification keeps the search tractable
 - · Start simple, add more sophistication
- ► Incremental construction of intelligent agents

Advanced NE 4: Indirect Encodings I



- ► Instructions for constructing the network evolved
- ► Instead of specifying each unit and connection 3,16,52,83,113
- ► E.g. Cellular Encoding (CE³⁰)
- ► Grammar tree describes construction
 - · Sequential and parallel cell division
 - Changing thresholds, weights

- - ► A "developmental" process that results in a network

Properties of Indirect Encodings I



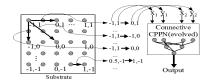
- Smaller search space
- Avoids competing conventions
- Describes classes of networks efficiently
- ► Modularity, reuse of structures
 - Recurrency symbol in CE: XOR → parity
 - ► Repetition with variation in CPPNs
 - Useful for evolving morphology

Properties of Indirect Encodings II



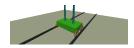
- Not fully explored (yet) ► See e.g. GDS track at GECCO
- Promising current work
 - ► More general L-systems; developmental codings;
 - embryogeny⁹⁰
 ► Scaling up spatial coding^{13,22}
 - ► Genetic Regulatory Networks⁷¹
- ► Evolution of symmetries 103

Indirect Encodings II



- ► Encode the networks as spatial patterns
- ► E.g. Hypercube-based NEAT (HyperNEAT 12
- ► Evolve a neural network (CPPN) to generate spatial patterns
 - ▶ 2D CPPN: (x,y) input \rightarrow grayscale output
 - ► 4D CPPN: (x_1, y_1, x_2, y_2) input $\rightarrow w$ output
 - · Connectivity and weights can be evolved indirectly ► Works with very large networks (millions of connections)

How Do the NE Methods Compare?



oles	Method	Evals	
wo	CE	(840,000)	
	CNE	87,623	
	ESP	26,342	
	NEAT	6,929	
	CMA-ES	6,061	
	CoSyNE	3,416	
o notes no valorities damnino fitness 28			

- ► Advanced methods better than CNE
- ► Advanced methods still under development
- ► Indirect encodings future work

Combining Learning and Evolution

- ► Good learning algorithms exist for NN
 - ► Why not use them as well?
- ► Evolution provides structure and initial weights
- ► Fine tune the weights by learning

- ► Incremental and multiobjective evolution ^{25,78,98,112}
- ► Utilizing population culture 5,50,94
- ► Utilizing evaluation history⁴⁷
- ► Evolving NN ensembles and modules ^{36,46,64,72,108}
- ► Evolving transfer functions and learning rules ^{8,74,93}
- ► Combining learning and evolution
- ► Evolving for novelty

Baldwin Effect



- ► Learning can guide Darwinian evolution as well^{4,30,32}
 - ► Makes fitness evaluations more accurate
- ► With learning, more likely to find the optimum if close
- ► Can select between good and bad individuals better
 - ► Lamarckian not necessary

Where to Get Learning Targets?





- ► From a related task 60
 - Useful internal representations
- ► Evolve the targets 63
 - ► Useful training situations
- ► From Q-learning equations 109
 - ► When evolving a value function
- ► Utilize Hebbian learning 18,87,101
 - ► Correlations of activity
- From the population 50,94
- Social learning
- ▶ From humans⁷
 - ► E.g. expert players, drivers

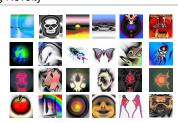
► Lamarckian evolution is possible ^{7,30}

- ► Coding weight changes back to chromosome
- ► Difficult to make it work

Lamarckian Evolution

► Diversity reduced; progress stagnates

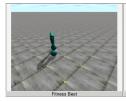
Evolving Novelty



- ► Motivated by humans as fitness functions
- ► E.g. picbreeder.com, endlessforms.com⁸⁰
 - ► CPPNs evolved; Human users select parents
- ► No specific goal
 - ► Interesting solutions preferred
 - ► Similar to biological evolution?

- ► Reward maximally different solutions
 - Can be a secondary, diversity objective ⁵⁹
 Or, even as the only objective ^{39,41}
- ► To be different, need to capture structure
 - Problem solving as a side effect
- ► Potential for innovation
- ► Needs to be understood better

Novelty Search Demo





- ► Fitness-based evolution is rigid
 - · Requires gradual progress
- ► Novelty-based evolution is more innovative, natural
 - ► Allows building on deceptive solutions
- ► (Demo available at eplex.cs.ucf.edu/noveltysearch)

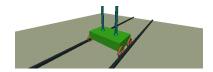
Extending NE to Applications

- ► Control
- ► Robotics
- ► Artificial life
- ► Gaming

Issues:

- ► Facilitating robust transfer from simulation ^{27,99}
- ► Utilizing problem symmetry and hierarchy^{38,102,103}
- ► Utilizing coevolution^{73,91}
- ► Evolving multimodal behavior^{77,78,108}
- ► Evolving teams of agents ^{6,88,114}
- ► Making evolution run in real-time 88

Applications to Control



- ► Pole-balancing benchmark
 - ► Originates from the 1960s

 - Original 1-pole version too easy
 Several extensions: acrobat, jointed, 2-pole, particle chasing ⁶⁴
- ► Good surrogate for other control tasks
 - Vehicles and other physical devices
 Process control 104

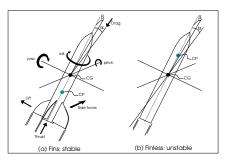
Controlling a Finless Rocket



Task: Stabilize a finless version of the Interorbital Systems RSX-2 sounding rocket²⁶

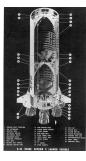
- ► Scientific measurements in the upper atmosphere
- ► 4 liquid-fueled engines with variable thrust
- ► Without fins will fly much higher for same amount of fuel

Rocket Stability



Simulation Environment: JSBSim

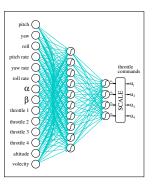
Rocket Guidance Network



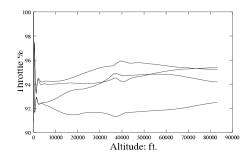
- Used on large scale launch vehicles (Saturn, Titan)
- Typically based on classical linear feedback control
- ► High level of domain knowledge required
- ► Expensive, heavy



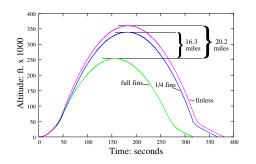
- General rocket simulator
- Models complex interaction between airframe, propulsion, aerodynamics, and atmosphere
- ► Used by IOS in testing their rocket designs
- ► Accurate geometric model of the RSX-2



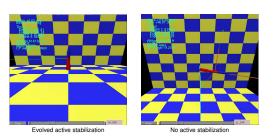
Results: Control Policy



Results: Apogee

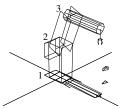


Finless Rocket Control Demo



(Demo available at nn.cs.utexas.edu)

Applications to Robotics



- ► Controlling a robot arm⁵⁶
 - ► Compensates for an inop motor
- ► Robot walking 34,82,102
- ➤ Various physical platforms
 ➤ Mobile robots ^{11,17,61,85}
- - ► Transfers from simulation to physical robots
 - ► Evolution possible on physical robots

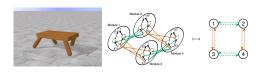
Multilegged Walking





- ► Navigate rugged terrain better than wheeled robots
- ► Controller design is more challenging
 - $\qquad \qquad \textbf{Leg coordination, robustness, stability, fault-tolerance,} \ \dots \\$
- ► Hand-design is generally difficult and brittle
- ► Large design space often makes evolution ineffective

ENSO: Symmetry Evolution Approach



- ► Symmetry evolution approach 100,102,103
 - ► A neural network controls each leg
 - ► Connections between controllers evolved through symmetry breaking
 - ► Connections within individual controllers evolved through neuroevolution

Versatile, Robust Gaits





Obstacle field

- ► Different gaits on flat ground
 - ► Pronk, pace, bound, trot
 - ► Changes gait to get over obstacles
- ► (DEMO available at nn.cs.utexas.edu)

Innovative, Effective Solutions

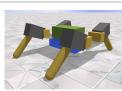




Handcoded

- ► Asymmetric gait on inclines
 - ► One leg pushes up, others forward
 - Hard to design by hand
- ► (DEMO available at nn.cs.utexas.edu)

Transfer to a Physical Robot I





- ► Built at Hod Lipson's lab (Cornell U.)
 - ► Standard motors, battery, controller board
 - ► Custom 3D-printed legs, attachments
 - ► Simulation modified to match
- ► General, robust transfer⁹⁹
 - ► Noise to actuators during simulation
 - ► Generalizes to different surfaces, motor speeds
- ► (DEMO available at nn.cs.utexas.edu)

Transfer to a Physical Robot II





- ► Evolved a solution for three-legged walking!
- ► (DEMO available at nn.cs.utexas.edu)

Applications to Artificial Life



- ► Gaining insight into neural structure
 - ► E.g. evolving a command neuron ^{2,37,75}
- ► Understanding animal behaviors
 - ► Signaling, herding, hunting... 62,67,69,70,97,106,107,114

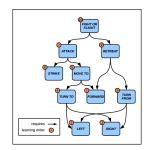
Body-Brain Coevolution





- ► Evolved Virtual Creatures 42,43,44,84
 - Body: Blocks, muscles, joints, sensors (Lessin et al. Alife'14)
 Brain: A neural network (with general nodes)
 Evolved together in a physical simulation
- ► Syllabus, Encapsulation, Pandemodium

Syllabus



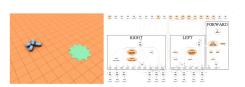
► Constructed by hand; body and brain evolved together

Encapsulation



► Once evolved, a trigger node is added

Pandemonium



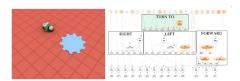
► Conflicting behaviors: Highest trigger wins

Evolving Fight-or-Flight Behavior

RETRIEF BOOK TO DOMBARD THAN TO DOMBARD THAN TO DOMBARD THAN THE PROPERTY OF T

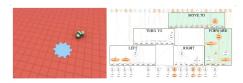
- ► Step-by-step construction of complex behavior
- ► Primitives and three levels of complexity
- ► DEMO (available at nn.cs.utexas.edu)

Turn to Light



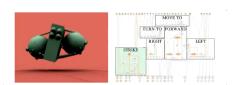
- ► First level of complexity
- ► Selecting between alternative primitives

Move to light



- ► First level of complexity (Sims 1994)
- ► Selecting between alternative primitives

Strike



► Alternative behavior primitive

Attack



► Second level of complexity (beyond Sims and others)

Turn from Light

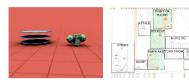


► Alternative first-level behavior

ATTACK RUTHEAT MOVE TO TURN TO TURN FEOMFOR WARD STRIKE RIGHT UF

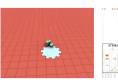
► Alternative second-level behavior

Fight or Flight



► Third level of complexity

Insight: Body/Brain Coevolution





- ► Evolving body and brain together poses strong constraints
 - ► Behavior appears believable
 - Worked well also in BotPrize (Turing test for game bots)
- ► What about constraints from the environment?

Coevolution of Behavior







Formalization of behavior

- ► Complex cooperation observed in pursuit and evasion
 - ► Motivated by biology, esp. hyenas vs. zebras (Kay Holekamp, MSU)
 - Largely innate, possible to see behaviors and their evolution
- ► Such behaviors evolve together, in coevolutionary environment
 - ► Simultaneous competitive and cooperative coevolution ^{67,70}

Experimental Setup



- ► Toroidal grid world
- ► Predators, prey move with same speed in 4 directions
- No direct communication between team members
 Communication still possible through stigmergy
- ► Does a coevolutionary arms race result?
- ► DEMO (available at nn.cs.utexas.edu)

Predator-Prey Arms Race I



50-75: Single predator catches prey



100-150: Two predators cooperate



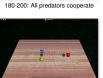
75-100: Prey evades by circling



150-180: Prey baits and escapes

Predator-Prey Arms Race II





250-300: Prey evade by scattering

Open Questions



- Role of communication
 - ► Stigmergy vs. direct communication in hunting 114

 - Quorum sensing in e.g. confronting lions
 Origins of communication ⁶⁸ (Rawal et al. Alife'14)
- Role of rankings
 - ► Efficient selection when evaluation is costly?
- ► Role of individual vs. team rewards
- Can lead to general computational insights

Bigger Questions



- Gaining insight into cognitive architectures
 - ► Executive, perception, emotion, memory ⁶⁶ (Rajagopalan et al. Alife'14)
- Emergence of language, learning, social structures
- ► May require overcoming deception
 - ► Through speciation, niching in nature
 - ► Through novelty search in computation?⁴⁰

Applications to Games





200-250: Predators herd two prey

Complex behaviors don't evolve in a vacuum

► Result from coevolutionary arms race ► Embedded in a changing environment

- ► Good research platform⁵¹
 - ► Controlled domains, clear performance, safe
 - · Economically important; training games possible
- ► Board games: beyond limits of search
 - ► Evaluation functions in checkers, chess^{9,19,20}
 - Filtering information in go, othello 54,92
 Opponent modeling in poker 48

Discovering Novel Strategies in Othello







- ► Players take turns placing pieces
- ► Each move must flank opponent's piece
- ► Surrounded pieces are flipped
- ► Player with most pieces wins

Strategies in Othello

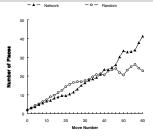






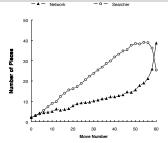
- ► Positional
 - ► Number of pieces and their positions
 - ► Typical novice strategy
- ► Mobility
 - ► Number of available moves: force a bad move
 - Much more powerful, but counterintuitive
 - ► Discovered in 1970's in Japan

Evolving Against a Random Player



- ▶ Network sees the board, suggests moves by ranking 55
- ► Networks maximize piece counts throughout the game
- A positional strategy emerges
- ► Achieved 97% winning percentage

Evolving Against an α - β Program



- ► lago's positional strategy destroyed networks at first
- ► Evolution turned low piece count into an advantage
- ► Mobility strategy emerged!
- ► Achieved 70% winning percentage

Video Games

Discovering Novel Strategies



- ► Neuroevolution discovered a strategy novel to us
- ► "Evolution works by tinkering"
 - ► So does neuroevolution
 - Initial disadvantage turns into novel advantage

The second secon



- ► Economically and socially important
- ► GOFAI does not work well
 - ► Embedded, real-time, noisy, multiagent, changing
 - ► Adaptation a major component
- ► Possibly research catalyst for CI
 - ► Like board games were for GOFAI in the 1980s

Example game





- ► Black's positions strong, but mobility weak
- ► White (the network) moves to f2
- Black's available moves b2, g2, and g7 each will surrender a corner
- ► The network wins by forcing a bad move

Video Games II





- ► Can be used to build "mods" to existing games
 - Adapting characters, assistants, tools
- ► Can also be used to build new games
 - ► New genre: Machine Learning game

Challenge 1: Evolving Multimodal Behavior



Evolution-Discovered Task Division

- Agents perform many different tasks
 - ► E.g. eat pills, avoid ghosts, eat powerpills, eat ghosts
 - ► Sometimes clearly separate in time
 - ► Sometimes multiple tasks at once
- How can we evolve them into a single network?

One module used 95% of the time ► Eat pills, avoid ghosts, chase ghosts Different behaviors with a common base A second module 5% of the time ► Luring ghosts near a power pill ► Escaping from tight spaces

A different multimodal perspective

Not as obvious, but more powerful

(DEMO available at nn.cs.utexas.edu)

MM-NEAT: Modular Multiobjective Approach







- ► Evolution discovers modules and when to use them
 - ► Vs. human-designed division with multitasking
- ► Multiple modules with preference neurons ⁷⁹
 - ► Modules implement different behaviors
 - ▶ Preference neurons used to choose among them
 - ► Module-mutation adds new modules
- ► Evolved towards multiple objectives
 - ► Correspond to dimensions of game play
 - ► E.g. pills and ghosts in Ms. Pac-Man
- ► Cf. (Li et al. Alife'14) 45

Challenge 2: Evolving Humanlike Behavior



- ► Botprize competition, 2007-2012
 - ► Turing Test for game bots (\$10,000 prize)
- ► Three players in Unreal Tournament 2004:
 - ► Human confederate: tries to win
 - ► Software bot: pretends to be human
 - ► Human judge: tries to tell them apart!

Human-Designed Task Division



- Multitask approach
- ► One module for threat ghosts
- Another module for edible ghosts
- ► Works ok, but...
- ► (DEMO available at nn.cs.utexas.edu)

► Evolve effective fighting behavior

Evolving an Unreal Bot

- ► Human-like with resource limitations (speed, accuracy...)
- ► Also scripts & learning from humans (unstuck, wandering...)
- ▶ 2007-2011: bots 25-30% vs. humans 35-80% human
- ▶ 6/2012 best bot better than 50% of the humans
- ▶ 9/2012...?

Success!!!



- ► In 2012, two teams reach the 50% mark!
- ► Fascinating challenges remain:
 - ► Judges can still differentiate in seconds
 - Judges lay cognitive, high-level traps
 - ► Team competition: collaboration as well
- ► (DEMO available at nn.cs.utexas.edu)

A New Genre: Machine Learning Games



- ► E.g. NERO

 - Goal: to show that machine learning games are viable
 Professionally produced by Digital Media Collaboratory, UTAustin
 - ► Developed mostly by volunteer undergraduates

NERO Gameplay



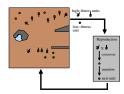






- ► Teams of agents trained to battle each other
- ► Player trains agents through excercises
 - Agents evolve in real time
 - Agents and player collaborate in battle
- ► New genre: Learning is the game 31,88
 - ► Challenging platform for reinforcement learning
 - ► Real time, open ended, requires discovery
- ► Try it out:
 - ► Available for download at http://nerogame.org
 - Open source research platform version at

Real-time NEAT



- ► A parallel, continuous version of NEAT⁸⁸
- ► Individuals created and replaced every n ticks
- ► Parents selected probabilistically, weighted by fitness
- ► Long-term evolution equivalent to generational NEAT

NERO Player Actions





- ► Player can place items on the field e.g. static enemies, turrets, walls, rovers, flags
- ► Sliders specify relative importance of goals e.g. approach/avoid enemy, cluster/disperse, hit target, avoid fire...
- ► Networks evolved to control the agents

NERO Training Demo











Maze Running (DEMO available at nn.cs.utexas.edu)





Teams of three

(DEMO available at nn.cs.utexas.edu)

Conclusion

- ▶ NE is a powerful technology for sequential decision tasks
 - · Evolutionary computation and neural nets are a good match
 - ► Lends itself to many extensions
 - Powerful in applications
- ► Easy to adapt to applications
 - ► Control, robotics, optimization
 - · Artificial life, biology
 - ► Gaming: entertainment, training
- ► Lots of future work opportunities
 - ► Theory needs to be developed
 - Indirect encodings
 - ► Learning and evolution
 - Knowledge, interaction, novelty

Numerous Other Applications

- ► Creating art, music, dance... 10,15,33,81
- ► Theorem proving 14
- ► Time-series prediction 49
- ► Computer system optimization²⁴
- ► Manufacturing optimization ²⁹
- ► Process control optimization 104,105
- ► Measuring top quark mass 110
- ► Etc.

Bibliography I

- A. Agogino, K. Tumer, and R. Milkkulainen, Efficient credit assignment through evaluation function decomposition, in: Proceedings of the Genetic and Evolutionary Computation Conference (2005).
- [2] R. Aharonov-Barki, T. Beker, and E. Ruppin, Emergence of memory-Driven command neurons in evolved artificial agents, Neural Computation, 13(3):691–716 (2001).
- [3] P. J. Angeline, G. M. Saunders, and J. B. Pollack, An evolutionary algorithm that constructs recurrent neural networks, IEEE Transactions on Neural Networks, 5:54–65 (1994).
- [4] J. M. Baldwin, A new factor in evolution, The American Naturalist, 30:441-451, 536-553 (1896).
- [5] R. K. Belew, Evolution, learning and culture: Computational metaphors for adaptive algorithms, Comple. Systems, 4:11–49 (1990).
- B. D. Bryant and R. Milikulainen, Neuroevolution for adaptive teams, in: Proceedings of the 2003 Congress on Evolutionary Computation (CEC 2003), volume 3, 2194–2201, IEEE, Piscataway, NJ (2003).
- [7] B. D. Bryant and R. Milkkulainen, Acquiring visibly intelligent behavior with example-guided neuroevolution, in: Proceedings of the Twenty-Second National Conference on Artificial Intelligence, AAAI Press, Menlo Park, CA (2007).
- [8] D. J. Chalmers, The evolution of learning: An experiment in genetic connectionism, in: Touretzky et al. 96,
- [9] K. Chellapilla and D. B. Fogel, Evolution, neural networks, games, and intelligence, Proceedings of the IEEE, 87:1471–1496 (1999).
- [10] C.-C. Chen and R. Milkkulainen, Creating melodies with evolving recurrent neural networks, in: Proceedings of the INNS-IEEE International Joint Conference on Neural Networks, 2241–2246, IEEE, Piscataway, NJ
- [11] D. Cliff, I. Harvey, and P. Husbands, Explorations in evolutionary robotics, Adaptive Behavior, 2:73–110
- [12] D. B. D'Ambrosio and K. O. Stanley, A novel generative encoding for exploiting neural network sensor and output geometry, in: Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation (GECCO' '07), 974–981, ACM, New York, NY, USA (2007).

Evaluation of Applications





- ► Neuroevolution strengths
 - ► Can work very fast, even in real-time
 - Potential for arms race, discovery
 - ► Effective in continuous, non-Markov domains
- ► Requires many evaluations
 - ► Requires an interactive domain for feedback
 - ► Best when parallel evaluations possible
 - Works with a simulator & transfer to domain

Bibliography II

- [13] D. B. D'Ambrosio and K. O. Stanley, Generative encoding for multiagent learning, in: Proceedings of the Genetic and Evolutionary Computation Conference (2008).
- [14] N. S. Desai and R. Milkkulainen, Neuro-evolution and natural deduction, in: Proceedings of The First IEEE Symposium on Combinations of Evolutionary Computation and Neural Networks, 64–69, IEEE, Piscataway NJ (2000).
- [15] G. Dubbín and K. O. Stanley, Learning to dance through interactive evolution, in: Proceedings of the Eighth European Event on Evolutionary and Biologically Inspired Music, Sound, Art and Design, Springer, Berlin (2010).
- [16] D. Floreano, P. Dürr, and C. Mattiussi, Neuroevolution: From architectures to learning, Evolutionary Intelligence, 1:47–62 (2008).
- [17] D. Floreano and F. Mondada, Evolutionary neurocontrollers for autonomous mobile robots, Neural Networks, 11:1461–1478 (1998).
- [18] D. Floreano and J. Urzelai, Evolutionary robots with on-line self-organization and behavioral fitness, Neural Networks, 13:431–4434 (2000).
- [19] D. B. Fogel, Blondie24: Playing at the Edge of Al, Morgan Kaufmann, San Francisco (2001).
- [20] D. B. Fogel, T. J. Hays, S. L. Hahn, and J. Quon, Further evolution of a self-learning chess program, in: Proceedings of the IEEE Symposium on Computational Intelligence and Games, IEEE, Piscataway, NJ (2005).
- [21] B. Fullmer and R. Milkkulainen, Using marker-based genetic encoding of neural networks to evolve finite-state behaviour, in: Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life, 1. Varelia and P. Bourgine, eds., 255–256, MIT Press, Cambridge, MA (1992).
- [22] J. J. Gauci and K. O. Stanley, A case study on the critical role of powerful regularity in machine learning, in: Proceedings of the Twenty-Third National Conference on Artificial Intelligence, AAAI Press, Menlo Park, CA (2008).
- [23] F. Gomez, Robust Non-Linear Control Through Neuroevolution, Ph.D. thesis, Department of Computer Sciences, The University of Texas at Austin (2003).

Bibliography III

- [24] F. Gomez, D. Burger, and R. Milikkulainen, A neuroevolution method for dynamic resource allocation on a chip multiprocessor, in: Proceedings of the INNS-IEEE International Joint Conference on Neural Networks 2355–2361, IEEE, Picastaway, NJ (2001).
- [25] F. Gomez and R. Miikkulainen, Incremental evolution of complex general behavior, Adaptive Behavior, 5:317–342 (1997).
- [26] F. Gomez and R. Milkkulainen, Active guidance for a finless rocket using neuroevolution, in: Proceedings of the Genetic and Evolutionary Computation Conference, 2084–2095, Morgan Kaufmann, San Francisco (2003).
- [27] F. Gomez and R. Milikkulainen, Transfer of neuroevolved controllers in unstable domains, in: Proceedings of the Genetic and Evolutionary Computation Conference, Springer, Berlin (2004).
- [28] F. Gomez, J. Schmidhuber, and R. Miikkulainen, Accelerated neural evolution through cooperatively coevolved synapses, *Journal of Machine Learning Research*, 9:937–965 (2008).
- [29] B. Greer, H. Hakonen, R. Lahdelma, and R. Milkkulainen, Numerical optimization with neuroevolution, in: Proceedings of the 2002 Congress on Evolutionary Computation, 361–401, IEEE, Piscataway, NJ (2002).
- [30] F. Gruau and D. Whitley, Adding learning to the cellular development of neural networks: Evolution and the Baldwin effect, Evolutionary Computation, 1:213–233 (1993).
- [31] E. J. Hastings, R. K. Guha, and K. O. Stanley, Automatic content generation in the galactic arms race video game, IEEE Transactions on Computational Intelligence and AI in Games, 1:245–263 (2009).
- [32] G. E. Hinton and S. J. Nowlan, How learning can guide evolution, Complex Systems, 1:495–502 (1987).
- [33] A. K. Hoover, M. P. Rosario, and K. O. Stanley. Scatfolding for interactively evolving novel drum tracks for existing songs. in: Proceedings of the Sixth European Workshop on Evolutionary and Biologically Inspire Music, Sound, Art and Design, Springer, Berlin (2008).
- [34] G. S. Hornby, S. Takamura, J. Yokono, O. Hanagata, M. Fujita, and J. Pollack, Evolution of controllers from a high-level simulator to a high DOF robot, in: Evolvable Systems: From Biology to Hardware; Proceedings of the Third International Conference, 80–96, Springer, Berlin (2000).

Bibliography IV

- [35] C. Igel, Neuroevolution for reinforcement learning using evolution strategies, in: Proceedings of the 2003 Congress on Evolutionary Computation, R. Sarker, R. Reynolds, H. Abbass, K. C. Tan, B. McKay, D. Essam, and T. Gedeon, eds., 2588–2595, IEEE Press, Piscataway, NJ (2003).
- [36] A. Jain, A. Subramoney, and R. Milikulainen, Task decomposition with neuroevolution in extended predator-prey domain, in: Proceedings of Thirteenth International Conference on the Synthesis and Simulation of Living Systems, East Larsing, M, USA (2012).
- [37] A. Keinan, B. Sandbank, C. C. Hilgetag, I. Meilijson, and E. Ruppin, Axiomatic scalable neurocontroller analysis via the Shapley value, Artificial Life, 12:333–352 (2006).
- [38] N. Kohl and R. Milkkulainen, Evolving neural networks for strategic decision-making problems, Neural Networks, 22:326–337 (2009).
- [39] J. Lehman and R. Milkkulainen, Effective diversity maintenance in deceptive domains, in: Proceedings of the Genetic and Evolutionary Computation Conference (2013).
- [40] J. Lehman and R. Milkkulainen, Overcoming deception in evolution of cognitive behaviors, in: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2014), Vancouver, BC, Canada (July 2014).
- [41] J. Lehman and K. O. Stanley, Abandoning objectives: Evolution through the search for novelty alone Evolutionary Computation, 2011:189–223 (2010).
- [42] D. Lessin, D. Fussell, and R. Mikkulainen, Open-ended behavioral complexity for evolved virtual creatures, in: Proceedings of the Genetic and Evolutionary Computation Conference (2013).
- [43] D. Lessin, D. Fussell, and R. Milkkulainen, Adapting morphology to multiple tasks in evolved virtual creatures in: Proceedings of The Fourteenth International Conference on the Synthesis and Simulation of Living n. rruceeaings of the Fourteenth Systems (ALIFE 14) 2014 (2014).
- [44] D. Lessin, D. Fussell, and R. Milkkulainen, Trading control intelligence for physical intelligence: Muscle drives in evolved virtual creatures, in: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2014), Vancouver, BC, Canada (July 2014).
- [45] X. Li and R. Milkkulainen, Evolving multimodal behavior through subtask and switch neural networks (2014).
- [46] Y. Liu, X. Yao, and T. Higuchi, Evolutionary ensembles with negative correlation learning, IEEE Transactions on Evolutionary Computation, 4:380–387 (2000).

Bibliography VI

- [59] J.-B. Mouret and S. Doncieux, Overcoming the bootstrap problem in evolutionary robotics using behavioral diversity, in: Proceedings of the IEEE Congress on Evolutionary Computation, 1161–1168, IEEE, Piscataway NJ (2009).
- [60] S. Nolfi, J. L. Elman, and D. Parisi, Learning and evolution in neural networks, Adaptive Behavior, 2:5-28
- [61] S. Nolfi and D. Floreano, Evolutionary Robotics, MIT Press, Cambridge (2000).
- [62] S. Nolfi and M. Mirolli, eds., Evolution of Communication and Language in Embodied Agents, Springer, Berlin
- [63] S. Nolfi and D. Parisi, Good teaching inputs do not correspond to desired responses in ecological neural networks, Neural Processing Letters, 1(2):1–4 (1994).
- [64] D. Pardoe, M. Ryoo, and R. Milkkulainen, Evolving neural network ensembles for control problems, in: Proceedings of the Genetic and Evolutionary Computation Conference (2005).

- [65] M. A. Poter and K. A. D. Jong, Cooperative coevolution: An architecture for evolving coadapted subcomponents, Evolutionary Computation, 81–29 (2000).
 [66] P. Rajapopalan, K. E. Holekamp, and R. Milkkulainen, The evolution of general intelligence, in: Proceedings of The Fourteenth international Conference on the Synthesis and Simulation of Living Systems (ALIFE 14). New York, NY (2014).
- [67] P. Rajagopalan, A. Rawal, R. Mikkulainen, M. A. Wiseman, and K. E. Holekamp, The role of reward structure, coordination mechanism and net return in the evolution of cooperation, in: Proceedings of the IEEE Conference on Computational Intelligence and Games (CIG 2011), Secul, South Korea (2011).
- [68] A. Rawal, J. Boughman, and R. Mikkulainen, Evolution of communication in mate selection, in: Proceedings of The Fourteenth International Conference on the Synthesis and Simulation of Living Systems (ALIFE 14), New York, USA (July, 2014 2014).
- [69] A. Rawal, P. Rajagopalan, K. E. Holekamp, and R. Milkkulainen, Evolution of a communication code in cooperative tasks, in: Proceedings of Trinteenth International Conference on the Synthesis and Stimulation of Living Systems (ALIfe 2013), East Lamising, Mr. USA (2012).

Bibliography VII

- [70] A. Rawal, P. Rajagopalan, and R. Milkkulainen, Constructing competitive and cooperative agent behavior using coevolution, in: IEEE Conference on Computational Intelligence and Games (CIG 2010), Copenhage Demmark (2010).
- [71] J. Reisinger and R. Milkkulainen, Acquiring evolvability through adaptive representations, in: Proceedings of the Genetic and Evolutionary Computation Conference, 1045–1052 (2007).
- [72] J. Reisinger, K. O. Stanley, and R. Milkkulainen, Evolving reusable neural modules, in: Proceedings of the Genetic and Evolutionary Computation Conference (2004).
- [73] C. D. Rosin and R. K. Belew, New methods for competitive evolution, Evolutionary Computation, 5 (1997).
- [74] T. P. Runarsson and M. T. Jonsson, Evolution and design of distributed learning rules, in: Proceedings of The First IEEE Symposium on Combinations of Evolutionary Computation and Neural Networks, 59–63, IEEE, Piscataway, NJ (2000).
- [75] E. Ruppin, Evolutionary autonomous agents: A neuroscience perspective, Nature Reviews Neuros
- [76] J. D. Schaffer, D. Whitley, and L. J. Eshelman, Combinations of genetic algorithms and neural networks: A survey of the state of the art, in: Proceedings of the International Workshop on Combinations of Genetic Algorithms and Neural Networks, D. Whitley and J. Schaffer, eds., 1–37, IEEE Computer Society Press, Los Alamitos, CA (1992).
- [77] J. Schrum and R. Mikkulainen, Evolving multi-modal behavior in NPCs, in: Proceedings of the IEEE Symposium on Computational Intelligence and Games, IEEE, Piscataway, NJ (2009).
- [78] J. Schrum and R. Milkkulainen. Evolving agent behavior in multiobjective domains using fitness-based shaping, in: Proceedings of the Genetic and Evolutionary Computation Conference (2010).
- [79] J. Schrum and R. Milkulainen, Evolving multimodal behavior with modular neural networks in ms. pac-man in: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2014), Vancouver, BC, Canada (July 2014).
- [80] J. Secretan, N. Beato, D. B. D'Ambrosio, A. Rodriguez, A. Campbell, J. T. Folsom-Kovarik, and K. O. Stanley, Picbreeder: A case study in collaborative evolutionary exploration of design space, Evolutionary Computation, 19:345–371 (2011).

Bibliography V

- [47] A. Lockett and R. Milkkulainen, Neuroannealing: Martingale-driven learning for neural network, in: Proceedings of the Genetic and Evolutionary Computation Conference (2013).
- [48] A. J. Lockett, C. L. Chen, and R. Milikkulainen, Evolving explicit opponent models in game playing, in: Proceedings of the Genetic and Evolutionary Computation Conference (2007).
- [49] J. R. McDonnell and D. Waagen, Evolving recurrent perceptrons for time-series modelling, IEEE Transactions on Evolutionary Computation, 5:24–38 (1994).
- [50] P. McQuesten, Cultural Enhancement of Neuroevolution, Ph.D. thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX (2002). Technical Report Al-02-295.
- [51] R. Milkkulainen, B. D. Bryant, R. Cornelius, I. V. Karpov, K. O. Stanley, and C. H. Yong, Computational intelligence in games, in: Computational Intelligence: Principles and Practice, G. Y. Yen and D. B. Fogel, eds., IEEE Computational Intelligence Society, Piscataway, NJ (2006).
- [52] E. Miolsness, D. H. Sharo, and B. K. Alpert, Scaling, machine learning, and genetic neural nets. Advances in Applied Mathematics, 10:137-163 (1989).
- [53] D. J. Montana and L. Davis, Training feedforward neural networks using genetic algorithms, in: Proceedings of the 11th International Joint Conference on Artificial Intelligence, 762–767, San Francisco: Morgan Kaufmann (1989).
- [54] D. E. Moriarty, Symbiotic Evolution of Neural Networks in Sequential Decision Tasks, Ph.D. thesis, Department of Computer Sciences, The University of Texas at Austin (1997). Technical Report UT-Al97-257.
- [55] D. E. Moriarty and R. Miikkulainen, Discovering complex Othello strategies through evolutionary neural networks, Connection Science, 7(3):195–209 (1995).
- [56] D. E. Moriarty and R. Milikulainen, Evolving obstacle avoidance behavior in a robot arm, in: From Animals to Animals 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, P. Maes, M. J. Mattari, J.-A. Meyer, J. Pollack, and S. W. Wilson, eds., 468–475, Cambridge, MA: MIT Press (1995).
- [57] D. E. Moriarty and R. Miikkulainen, Forming neural networks through efficient and adaptive co-evolution, Evolutionary Computation, 5:373–399 (1997).
- [58] D. E. Moriarty, A. C. Schultz, and J. J. Grefenstette, Evolutionary algorithms for reinforcement learning, Journal of Artificial Intelligence Research, 11:199–229 (1999).

Bibliography VIII

- [81] J. Secretan, N. Beato, D. B. D'Ambrosio, A. Rodriguez, A. Campbell, and K. O. Stanley, Picbreeder: Evolving pictures collaboratively online, in: Proceedings of Computer Human Interaction Conference, ACM, New York
- [82] C. W. Seys and R. D. Beer, Evolving walking: The anatomy of an evolutionary search, in: From Animals to Animats 8: Proceedings of the Eight International Conference on Simulation of Adaptive Behavior, S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J.-A. Meyer, eds., 357–363, MIT Press, Cambridge, MA
- [83] A. A. Siddiqi and S. M. Lucas, A comparison of matrix rewriting versus direct encoding for evolving neural networks, in: Proceedings of IEEE International Conference on Evolutionary Computation, 392–397, IEEE Piscataway, NJ (1988).
- [84] K. Sims, Evolving 3D morphology and behavior by competition, in: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems (Artificial Life IV), R. A. Brooks and P. Maes, eds., 28–39, MIT Press, Cambridge, Mr. (1984).
- [85] Y. F. Sit and R. Milkkulainen, Learning basic navigation for personal satellite assistant using neuro in: Proceedings of the Genetic and Evolutionary Computation Conference (2005).
- [86] K. O. Stanley, Efficient Evolution of Neural Networks Through Complexification, Ph.D. thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX (2003).
- [87] K. O. Stanley, B. D. Bryant, and R. Milkkulainen, Evolving adaptive neural networks with and without adaptive synapses, in: Proceedings of the 2003 Congress on Evolutionary Computation, IEEE, Piscataway, NJ (2003).
- [88] K. O. Stanley, B. D. Bryant, and R. Milkkulainen, Real-time neuroevolution in the NERO video game, IEEE Transactions on Evolutionary Computation, 9(6):653–668 (2005).
- [89] K. O. Stanley and R. Miikkulainen, Evolving Neural Networks Through Augmenting Topologies, Evolutionary Computation, 10:99–127 (2002).
- [90] K. O. Stanley and R. Miikkulainen, A taxonomy for artificial embryogeny, Artificial Life, 9(2):93-130 (2003).
- [91] K. O. Stanley and R. Milkkulainen, Competitive coevolution through evolutionary complexification, Journal of Artificial Intelligence Research, 21:63–100 (2004).

Bibliography IX

- [92] K. O. Stanley and R. Miikkulainen, Evolving a roving eye for Go, in: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2004), Springer Verlag, Berlin (2004).
- [93] D. G. Stork, S. Walker, M. Burns, and B. Jackson, Preadaptation in neural circuits, in: International Joint Conference on Neural Networks (Washington, DC), 202–205, IEEE, Piscataway, NJ (1990).
- [94] W. Tansey, E. Feasley, and R. Mikkulainer, Accelerating-volution via egalistica social learning, in: Proceedings of the 14th Annual Genetic and Evolutionary Computation Conference (GECCO 2012), Philadelphia, Pennsylvania, USA (Luly 2014).
- [95] M. Taylor, S. Whiteson, and P. Stone, Comparing evolutionary and temporal difference methods in a reinforcement learning domain, in: Proceedings of the Genetic and Evolutionary Computation Conference (2006).
- [96] D. S. Touretzky, J. L. Elman, T. J. Sejnowski, and G. E. Hinton, eds., Proceedings of the 1990 Connectionist Models Summer School, San Francisco: Morgan Kaufmann (1990).
- [97] E. Tuci, An investigation of the evolutionary origin of reciprocal communication using simulated autonomous agents. Biological Cybernetics. 101:183–199 (2009).
- [98] J. Urzelai, D. Floreano, M. Dorigo, and M. Colombetti, Incremental robot shaping, Connection Science, 10:341–360 (1998).
- [99] V. Valsalam, J. Hiller, R. MacCurdy, H. Lipson, and R. Milkkulainen, Constructing controllers for physical multilegged robots using the enso neuroevolution approach, Evolutionary Intelligence, 14:303–331 (2013).
- [100] V. Valsalam and R. Milkkulainen, Evolving symmetry for modular system design, IEEE Transactions on Evolutionary Computation, 15:368–386 (2011).
- [101] V. K. Valsalam, J. A. Bednar, and R. Milkulainen. Constructing good learners using evolved pattern generators, in: Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2005, H-G. Beyer et al., eds., 11–13. New York: ACM (2005).
- [102] V. K. Valsalam and R. Milikkulainen, Modular neuroevolution for multilegged locomotion, in: Proceedings of the Genetic and Evolutionary Computation Conference GECCO 2008, 265–272, ACM, New York, NY, USA (2008).

Bibliography X

- [103] V. K. Valsalam and R. Milikulainen, Evolving symmetric and modular neural networks for distributed control, in: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) 2009, 731–738, ACM, New York, NY, USA (2009).
- [104] A. van Eck Corradie, R. Milkulainen, and C. Adrich, Adaptive control utilising neural swarming, in: Proceedings of the Genetic and Evolutionary Computation Conference, W. B. Langdon, E. Cantis-Paz, K. E. Mathias, R. Roy, D. Davis, R. Poli, K. Balakrishan, Y. Honnaur, G. Hodoph, J. Wegener, L. Bull, M. A. Potter, A. C. Schultz, J. F. Miller, E. K. Burke, and N. Jonoska, eds., San Francisco: Morgan Kaufmann (2002).
- [105] A. van Eck Conradie, R. Milkkulainen, and C. Aldrich, Intelligent process control utilizing symbiotic memetic neuro-evolution, in: Proceedings of the 2002 Congress on Evolutionary Computation (2002).
- [106] G. M. Werner and M. G. Dyer, Evolution of communication in artificial organisms, in: Proceedings of the Workshop on Artificial Life (ALIFE '90), C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, eds., 659–687, Reading, MA: Addison-Wesley (1992).
- [107] G. M. Werner and M. G. Dyer, Evolution of herding behavior in artificial animals, in: Proceedings of the Second International Conference on Simulation of Adaptive Behavior, J.-A. Meyer, H. L. Roitblat, and S. W. Wilson, eds., Cambridge, Mar. MIT Press (1992).
- [108] S. Whiteson, N. Kohl, R. Milkkulainen, and P. Stone, Evolving keepaway soccer players through task decomposition, Machine Learning, 59:5–30 (2005).
- [109] S. Whiteson and P. Stone, Evolutionary function approximation for reinforcement learning, *Journal of Machine Learning Research*, 7:877–917 (2006).
- [110] S. Whiteson and D. Whiteson, Stochastic optimization for collision selection in high energy physics, in: Proceedings of the Nineteenth Annual Innovative Applications of Artificial Intelligence Conference (2007).
- [111] D. Whitley, S. Dominic, R. Das, and C. W. Anderson, Genetic reinforcement learning for neurocontrol problems, Machine Learning, 13:259–284 (1993).
- [112] A. P. Wieland, Evolving controls for unstable systems, in: Touretzky et al. ⁹⁶, 91–102.
- [113] X. Yao. Evolving artificial neural networks. Proceedings of the IEEE. 87(9):1423–1447 (1999).
- [114] C. H. Yong and R. Milikulainen, Coevolution of role-based cooperation in multi-agent systems, IEEE Transactions on Autonomous Mental Development, 1:170–186 (2010).