

Neuroevolution of Collective Systems

Risto Miikkulainen

September 29, 2024

Collective Intelligence

- ▶ Groups self-assemble into complex forms based on local interactions.
- ▶ Examples: Ants building bridges, termites constructing nests, and bees making foraging decisions.
- ▶ These complex collective behaviors emerge from simple individual behaviors discovered through evolution.




Coevolution of Intelligent Systems

- ▶ Coevolution: Multiple populations evolve together.
- ▶ Cooperative: Populations evolve to achieve common goals.
- ▶ Competitive: Populations compete for resources, driving innovation.
- ▶ Neuroevolution can utilize both processes.




Cooperative Coevolution of a Single Neural Network

- ▶ Cooperative coevolution applied to individual components, such as neurons or connections.
- ▶ Neurons are evolved to work together to solve a task.
- ▶ How to form a network intelligently?



SANE: Symbiotic Coevolution of Networks and Blueprints

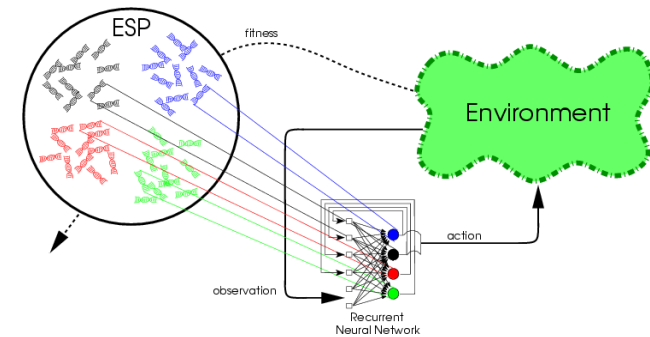
- ▶ Networks are created based on a blueprint that selects neurons.
- ▶ Neurons and blueprints coevolve based on the network's fitness.
- ▶ Neurons evolve to cooperate; blueprints to combine.
- ▶ Tends to evolve general neurons, not specializations.



Navigation icons: back, forward, search, etc.

ESP: Enforced Subpopulations

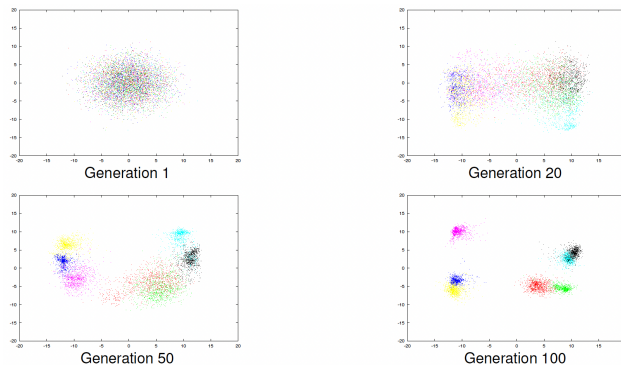
- ▶ ESP enhances SANE by evolving neurons in separate subpopulations.
- ▶ Each neuron specializes in a specific location within the network.
- ▶ Fully connected network; only weights evolved.
- ▶ This approach helps the network discover differentiated roles for each neuron.



Navigation icons: back, forward, search, etc.

Example: Maze Navigation with Khepera

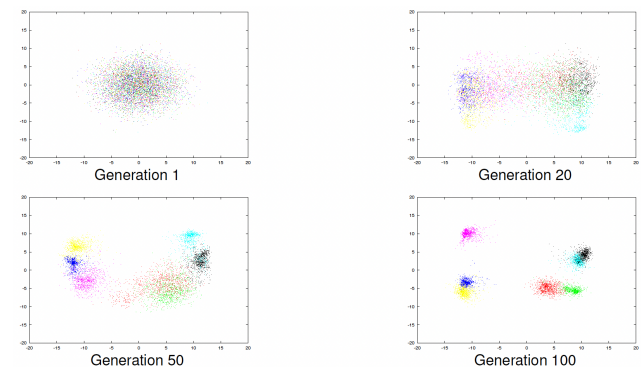
- ▶ Subpopulations start random but specialize over time.
 - ▶ Some slow the robot down with obstacle up front; others veer left when obstacle on the right, etc.
- ▶ Evolution discovers compatible subtasks.
 - ▶ Neurons optimized for each subtask.
 - ▶ Avoids the competing conventions problem by assigning neurons to distinct roles.
 - ▶ Reduces search space by evolving neurons individually instead of optimizing the entire network at once.



Navigation icons: back, forward, search, etc.

Maintaining Diversity in Cooperative Coevolution

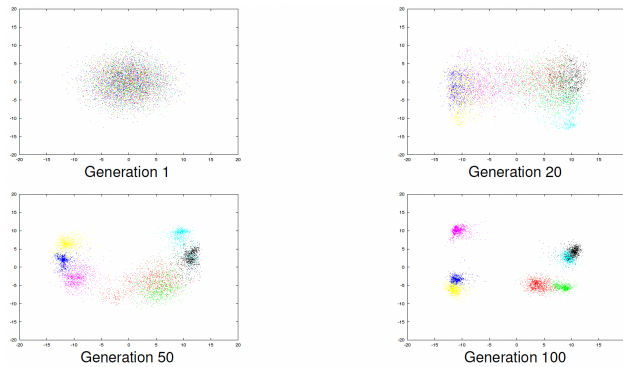
- ▶ Maintaining diversity in the population is essential to prevent premature convergence.
- ▶ Neurons must specialize in different tasks, preventing the population from becoming too similar.
- ▶ ESP helps maintain diversity through subpopulation specialization.



Navigation icons: back, forward, search, etc.

Robust Search for Subtasks


- ▶ Subpopulations evolve multiple subtasks, leading to redundancy.
- ▶ Redundancy ensures that even suboptimal neurons are compensated by others.
- ▶ Redundance makes the search robust: Necessary subtasks usually included.



Navigation icons: back, forward, search, etc.

CoSyNE: Cooperative Coevolution of Weights



- ▶ Extend the idea of evolving partial solutions to weights.
 - ▶ Each weight evolved in a separate subpopulation.
 - ▶ Networks formed by combining neurons with the same index.
 - ▶ Networks mutated and recombined; indices permuted.
- ▶ Instead of unrestricted search, exploration of new combinations.



Navigation icons: back, forward, search, etc.

Example: POMDP Double Pole Balancing

- ▶ Two poles with different lengths respond differently.
 - ▶ Nonlinear interactions make it difficult.
 - ▶ Without velocities requires a recurrent network.
- ▶ CoSyNE state of the art; RL could not solve.



Navigation icons: back, forward, search, etc.

Evolving a Team


- ▶ Neuroevolution can be extended to construct teams of agents.
- ▶ Agents evolve separately but are evaluated based on the success of the entire team.
- ▶ Predator-prey scenarios provide a classic example of this approach.



Navigation icons: back, forward, search, etc.

Cooperative Strategy in Predator-Prey Task

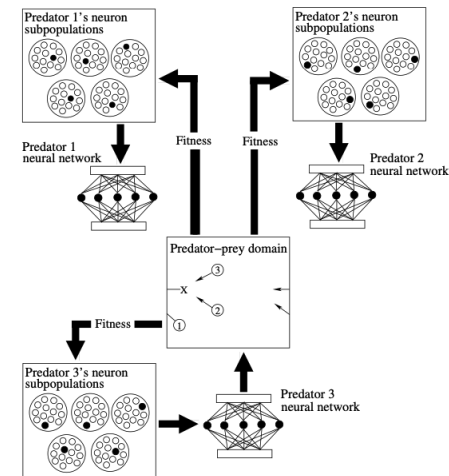
- ▶ In a predator-prey task, three predators evolve to cooperate in capturing a prey (in a toroidal environment).
- ▶ The prey runs away from the nearest predator (stochastically).
- ▶ The team is rewarded based on the success of the capture.



Navigation icons: back, forward, search, etc.

Multi-Agent ESP

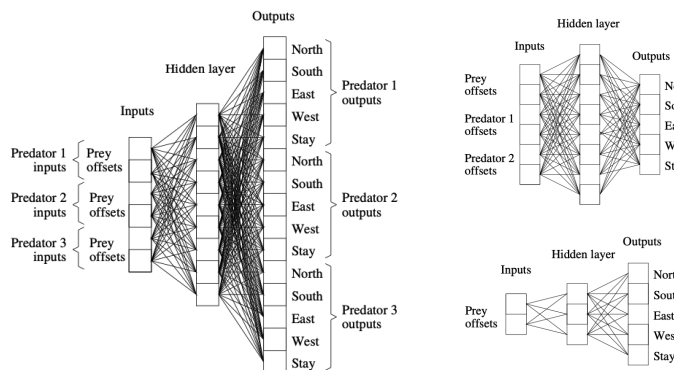
- ▶ Extend ESP to multiple networks: one for each predator.
- ▶ Hierarchical structure: each neuron subpopulation evolves one neuron for one network.
- ▶ The neurons inherit the fitness of the entire team.



Navigation icons: back, forward, search, etc.

To Communicate or Not to Communicate?

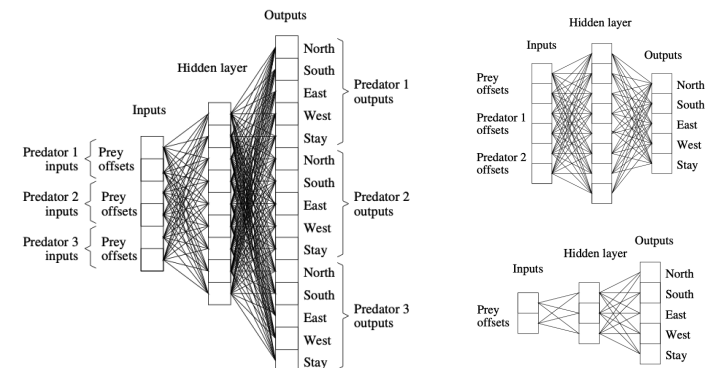
- ▶ For comparison, a central controller can be evolved in one population.
- ▶ In Multi-agent ESP, each agent may see each other.
- ▶ Or they may see only the prey.



Navigation icons: back, forward, search, etc.

To Communicate or Not to Communicate?


- ▶ Central controller takes twice as long to evolve than a communicating team.
- ▶ A communicating team takes twice as long as non-communicating team.
- ▶ How can less information be more effective?



Navigation icons: back, forward, search, etc.


Role-based Cooperation Through Stigmergy

- ▶ Without communication, team members evolve distinct roles.
- ▶ Cooperation emerges through stigmergy—coordination through interaction with the environment, i.e. the prey.
- ▶ For instance, two chasers driving the prey to a blocker.




Discovering Compatible Behaviors

- ▶ Evolution discovers role-based behaviors more easily than flexible team strategies.
- ▶ Each behavior compensates for inaccuracies in other agents.
- ▶ Cooperation based on roles often leads to robust solutions.




Adaptive Communication-based Cooperation

- ▶ In some cases, agents may need to change their behavior based on changing situations.
- ▶ E.g. changing roles, changing direction of the chase:




Role vs. Communication-based Cooperation

- ▶ The contrast is similar to well-practiced vs. pick-up soccer.
- ▶ Communication-based is less effective, more reactive, more general.
 - ▶ E.g. changing direction of the chase:



Evolving Generalists


- ▶ Evolving specialists is not always effective.
- ▶ The required number or kind of specialists may change.
- ▶ Generalists adapt to perform different roles when needed.
- ▶ This approach requires evolving homogeneous teams:
 - ▶ Single agent cloned to form a team.



Navigation icons: back, forward, search, etc.

Example: Legions Strategy Game

- ▶ Agents need to defend the cities and chase barbarians in the countryside.
- ▶ These roles are required at different numbers at different times.
- ▶ The team evolves to allocate roles dynamically.
- ▶ Heterogeneous vs. homogeneous teams better? Still an open question.



Navigation icons: back, forward, search, etc.

Competitive Coevolution

- ▶ Competition between agents drives the discovery of increasingly complex behaviors.
- ▶ Open-ended fitness: Agents continuously evolve to outdo each other in an evolutionary arms race.
- ▶ Evolutionary dynamics similar to curricular learning in machine learning.



Toxin vs. resistance to it in garter snakes and newts

Navigation icons: back, forward, search, etc.

Fitness Definition in Competitive Coevolution

- ▶ Fitness is defined in relation to the performance of other agents in the population.
- ▶ As individuals improve, fitness becomes harder to achieve, ensuring continuous adaptation.
- ▶ Competitive coevolution automatically shapes the fitness function.

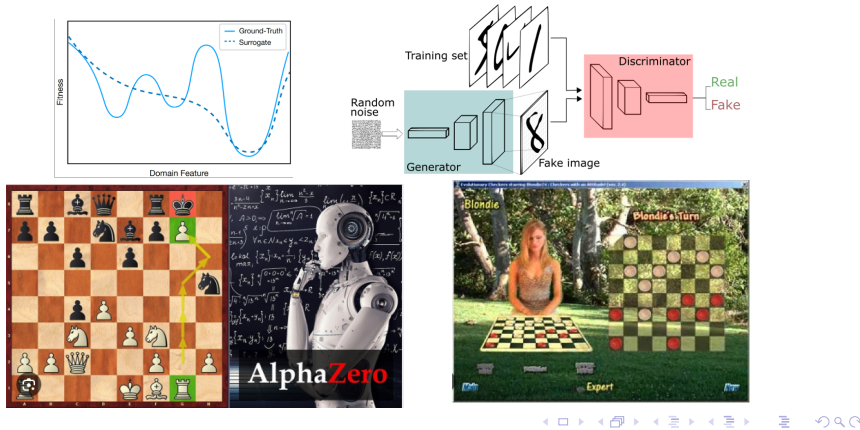


Toxin vs. resistance to it in garter snakes and newts

Navigation icons: back, forward, search, etc.


Sound Familiar?

- ▶ Dynamics similar to curricular learning.
- ▶ GANs (Generative Adversarial Networks) employ competitive coevolution between generator and discriminator.
- ▶ Similar mechanisms are seen in self-play systems like AlphaZero.
- ▶ Early neuroevolution systems like Blondie24 pioneered competitive coevolution in checkers and chess.




Challenges in Competitive Coevolution

- ▶ Progress is not always guaranteed; fitness improvement may only be relative (Red Queen dynamics).
- ▶ Possible to exploit weaknesses in current candidates without true improvement.
- ▶ Maintaining a collection of previous candidates helps track absolute progress.




Competitive Coevolution with NEAT

- ▶ NEAT supports competitive coevolution by complexifying networks incrementally.
- ▶ Networks grow in complexity while preserving earlier behaviors, ensuring absolute progress.
- ▶ Mutation and crossover add nodes and connections to existing structures.




Foraging, Pursuit, and Evasion Task

- ▶ Simulated Khepera robots evolve through competitive coevolution to forage, pursue, and evade.
- ▶ Robots sense distance to opponents, food items, and walls.
- ▶ They gain energy by eating, lose energy by traveling.
- ▶ Win when bumping into the opponent while more energy.



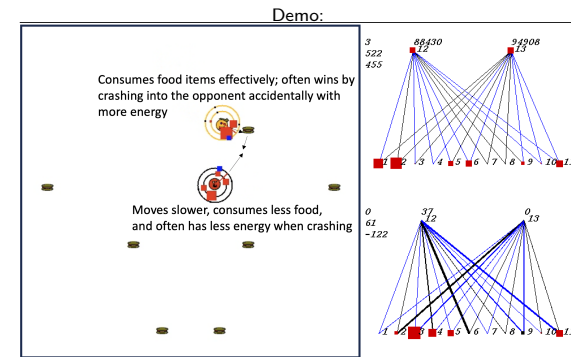
Discovering Complex Behavior

- ▶ Competitive coevolution discovers complex strategies through incremental improvements.
- ▶ Robots evolve foraging and attack strategies, learning to predict and exploit opponent energy levels.
- ▶ Energy management becomes crucial for winning encounters.



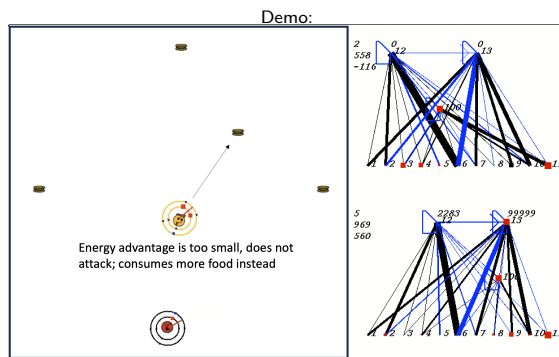
Early Strategies

- ▶ Rest and let opponent waste energy
- ▶ Mainly forage, occasionally crash by accident
- ▶ Difficult to switch between tasks



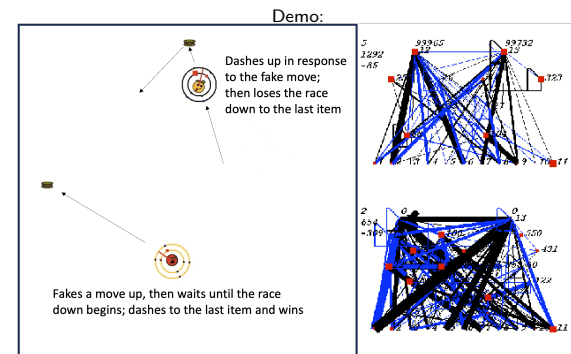
Mature Strategies

- ▶ Recurrent hidden node allows switching between tasks
- ▶ Collect food to gain energy; rest to save energy
- ▶ Difficult to predict energy at contact




A Sophisticated Strategy

- ▶ Split & recurrent connections predict crash outcome
- ▶ Complex structure to anticipate opponent behavior
 - ▶ "Fake" a rest; entice opponent to forage far away
 - ▶ Win by making a dash to last piece



Coevolution of Complex Strategies


- ▶ Competitive coevolution is a powerful approach for discovering increasingly complex behaviors.
- ▶ Strategies such as faking moves and forcing opponents into energy-depleting mistakes are discovered.
- ▶ These behaviors would be difficult to discover without competitive coevolution, or without complexification.



Navigation icons: back, forward, search, etc.

Evolving Multiple Teams


- ▶ Multiple cooperative teams can evolve in a competitive environment.
- ▶ Teams challenge each other, leading to increasingly complex behaviors.
- ▶ This process is called an evolutionary arms race.



Navigation icons: back, forward, search, etc.

Challenges in Establishing Absolute Improvement


- ▶ Absolute improvement is not always guaranteed.
- ▶ Teams may evolve strategies that exploit weaknesses in others but fail in the long run.
- ▶ However, in natural tasks, more complex behaviors often subsume simpler ones, leading to real progress.



Navigation icons: back, forward, search, etc.

Predator-Prey Coevolution Example


- ▶ A good example of competitive-cooperative dynamics is the predator-prey task.
- ▶ Predator (hyenas) and prey (zebras) populations evolve together in a toroidal world.
- ▶ Predators evolve strategies to catch prey, while prey evolve strategies to escape.




Navigation icons: back, forward, search, etc.

Initial Behaviors

- ▶ Initially, prey evolves to run away from the nearest predator, and predators towards the prey. The prey is captured increasingly often.
- ▶ In response, the prey evolves to circle the predator.



50-75: Single predator catches prey




75-100: Prey evades by circling




Cooperation and Escape

- ▶ Predators evolve a cooperative strategy of approaching from two directions.
- ▶ Prey evolves to lure them close and then escapes between them.



100-150: Two predators cooperate




150-180: Prey baits and escapes




Herding Multiple Preys

- ▶ Predators evolve to approach from three different directions.
- ▶ This strategy works also for two preys.



180-200: All predators cooperate




200-250: Predators herd two prey



Confusing the Predators

- ▶ The prey team evolves to confuse the predators by splitting their directions.
- ▶ This mirrors natural behavior seen in prey!




250-300: Prey evade by scattering



Insights from the Predator-Prey Simulation

- ▶ Complex behaviors do not evolve in a vacuum, but in response to a changing environment.
- ▶ The competitive+cooperative coevolution arms race leads to increasingly complex interactions over multiple generations.
- ▶ Simulations shed light on the evolution of real-world animal behaviors.



Conclusion on Collective Systems

- ▶ **Cooperative Coevolution:**
 - ▶ Individual networks as well as teams can be evolved cooperatively by sharing overall fitness.
 - ▶ Role-based cooperation can be effective; communication-based flexible.
 - ▶ Homogeneous teams of generalists can adapt dynamically.
- ▶ **Competitive Coevolution**
 - ▶ Competition can drive the emergence of highly complex behaviors.
 - ▶ The arms race forces each population to innovate in response to the other.
 - ▶ This process may be crucial in open-ended evolution and the emergence of major evolutionary transitions.