



Conclusion: Evolving Robust Control

- ▶ Robust control is essential for generalization and adaptability in complex environments.
- ▶ Techniques like noise injection, coevolution of controllers with teachers and problems, novelty search, explicit context representation, and symmetry help build this robustness.
- ▶ Advanced simulators, noise injection, and coevolution with hardware make transfer possible.




Direct Evolution

- ▶ Mapping sensors directly to actions
- ▶ Difficult to separate behaviors
- ▶ Ineffective combinations result




From Low-level Control to High-level Strategy

- ▶ Low-level control: Adjusting single behaviors (e.g., moving a leg faster).
- ▶ High-level strategy: Coordinating multiple behaviors.
- ▶ Example: Keepaway soccer: GetOpen, Intercept, Hold, EvaluatePass, Pass
- ▶ Challenge: Switching between behaviors effectively.




Coevolution Approach

- ▶ Evolve a separate network for each behavior
- ▶ A decision tree to decide which network to activate




Flexible Multimodal Behavior

- ▶ Discovering flexible multimodal behavior is a key step toward general intelligence.
- ▶ Keepaway task:
 - ▶ Networks learn individual tasks
 - ▶ Learn to anticipate other tasks as well: Lining up for a pass
 - ▶ Cooperative coevolution of multimodal behavior




Challenge 1: Abrupt Switching Between Behaviors

- ▶ Some strategies require abrupt behavior changes.
- ▶ Example: Small changes in soccer can shift the optimal action from holding to shooting or passing.
- ▶ Difficult to capture in a decision tree.
- ▶ Can we evolve a network to do it?




Using Radial Basis Activation Functions

- ▶ Radial Basis Functions (e.g. elongated Gaussians) activate neurons in local regions.
 - ▶ Many sigmoidal nodes are needed for the same effect.
 - ▶ Easier to discover fractured regions.




Using Cascaded Network Structures

- ▶ Cascaded networks: new hidden neurons added on top of earlier ones.
- ▶ Earlier connections are frozen.
- ▶ Each new neuron refines the boundaries of existing behaviors.




RBF + Cascade Make Abrupt Changes Possible

- Scales to 4v2 Keepaway.




Number of teammates available for a pass




<https://vimeo.com/2155250>

RBF + Cascade Make Abrupt Changes Possible

- Scales to 5v5 half-field soccer.




Subsets of actions (of 6) available



<https://vimeo.com/5698040>

Challenge 2: Blending and Interleaving Behaviors

- Intelligent agents often combine several behaviors.
- Example: Switching between offense and defense in soccer.
 - They can be blended or rapidly interleaved.




Example: Ms. Pac-Man

- In Ms. Pac-Man, agents perform several tasks: Eat pills, avoid ghosts, eat powerpills, eat ghosts.
- Sometimes interleaved but clearly separate.
- Sometimes blended into multiple tasks at once.
- How can we evolve such complex combinations of behaviors?




Evolving a Simple Control Network

- Simple networks can be evolved for Ms. Pac-Man, but they struggle with behavior separation.
- Results: Poor performance due to blended behaviors.
- Neuroevolution can learn multiple behaviors, but it needs a more sophisticated approach to switch effectively.




Multitask Networks for Isolated and Interleaved Tasks

- Multitask networks with separate outputs can be evolved for threatening and edible ghosts.
- Decide on which outputs to use based on a rule.
- These networks work well in isolated or interleaved tasks.
- However, they still struggle in blended situations where multiple behaviors are required simultaneously.




Expert-designed Subtasks and Combiner Networks

- Evolve separate networks for each behavior, and a combiner network to switch between them.
- Evolve one network for threatening and another for edible ghosts.
- The combiner could be gating or transforming the task-specific outputs.
- Possible to blend, but three coevolving populations difficult to converge.




Preference Neurons: Letting Evolution Discover Task Divisions

- Preference neurons allow evolution to decide when to switch behaviors.
- Each output module is coupled with a preference neuron, indicating when it should be used.
- Evolution can add modules similarly to nodes and connections in NEAT.
- This method enables evolution to discover more flexible and effective task divisions.




Surprising Strategy Discovery: Luring

- ▶ Evolution discovered an unexpected strategy: luring ghosts toward a power pill, then eating them up.
- ▶ One module dedicated to this strategy.
- ▶ Human designers may not have discovered this behavior.




Discovering Effective Task Divisions

- ▶ The same luring module was used to escape threatening ghosts in tight spaces!
- ▶ A very different task division:
 - ▶ Luring and escaping used only 5% of the time, but it counts.
 - ▶ Eating pills, avoiding ghosts, chasing ghosts with the other module 95% of the time; variations with a common base.
- ▶ With the freedom to explore different strategies, evolution finds surprising and powerful solutions.




Conclusion on Discovering Flexible Strategies

- ▶ High-level strategies require flexible switching and blending of multiple behaviors.
- ▶ RBF nodes, Cascaded networks, modular networks with preference neurons allow evolution to discover such strategies.
- ▶ Optimal behaviors can be surprising, e.g. blending and luring.




Scaling up to Cognitive Behaviors

- ▶ Cognitive behaviors include communication, memory, and learning.
- ▶ These behaviors are complex and difficult to evolve.
- ▶ The challenge: They require circuitry that doesn't help until it works.
- ▶ Need to overcome deception during evolution.



Example: T-maze Task


- ▶ The T-maze task illustrates how communication, memory, and learning can evolve.
- ▶ The agent must navigate to the reward at the correct end of the T-maze.
- ▶ Evolution struggles when the reward location changes frequently, requiring cognitive strategies.



Navigation icons: back, forward, search, etc.

Deception in Evolving Communication


- ▶ To evolve communication, agents must develop mechanisms to send, receive, and interpret signals.
- ▶ Deception occurs because partial solutions are not helpful unless all components work together.



Navigation icons: back, forward, search, etc.

Deception in Evolving Memory


- ▶ Receives a signal at start: if AX, go left; if BX, BY, AY, go right.
- ▶ To evolve memory, agents must store activations, retrieve them at the right time, and interpret them.
- ▶ Similar to communication, but internal to the network.
- ▶ Deception occurs because partial solutions are not helpful unless all components work together.



Navigation icons: back, forward, search, etc.

Deception in Evolving Learning


- ▶ To evolve learning, agents must develop a learning rule that reinforces good outcomes.
- ▶ Deception occurs because adaptation is mostly harmful—until it works.



Navigation icons: back, forward, search, etc.

Discovering Cognitive Behaviors with Novelty Search

- ▶ Fitness-based evolution reactive, i.e. always left or always right.
- ▶ Novelty search can overcome deception through stepping stones.
- ▶ The lineage of solutions shows multiple stepping stones:
 - ▶ E.g. going to the opposite corridor with some communication inputs.




Novelty Search and Cognition

- ▶ How did cognition really evolve in biology?
 - ▶ No explicit reward for novelty, but there are multiple goals and niches.
 - ▶ Stepping stones can be rewarded for entirely other reasons.
 - ▶ E.g. evolution of language based on social structure?
- ▶ Still a challenge, but its time may have come!



Scaling up further: Coevolution

- ▶ Coevolution: Agents evolve in competition or cooperation with each other.
- ▶ It drives agents to develop more sophisticated and adaptable behaviors.
- ▶ Discussed at length in next few weeks.



Conclusion on Evolving Behavior

- ▶ **Evolving control, i.e. single behaviors:**
 - ▶ Neuroevolution excels at discovering robust control solutions for dynamic, noisy, and nonlinear tasks.
 - ▶ Creative behaviors are discovered that compensate for physical imperfections or limitations.
 - ▶ Adapting to new conditions outside training is a major challenge.
- ▶ **Evolving strategy, i.e. multiple behaviors**
 - ▶ Complex strategies, such as switching between behaviors, can be evolved with the right architectures.
 - ▶ They can result in surprisingly effective combinations.
 - ▶ Evolving cognitive behaviors such as communication, memory, and learning is a major challenge.