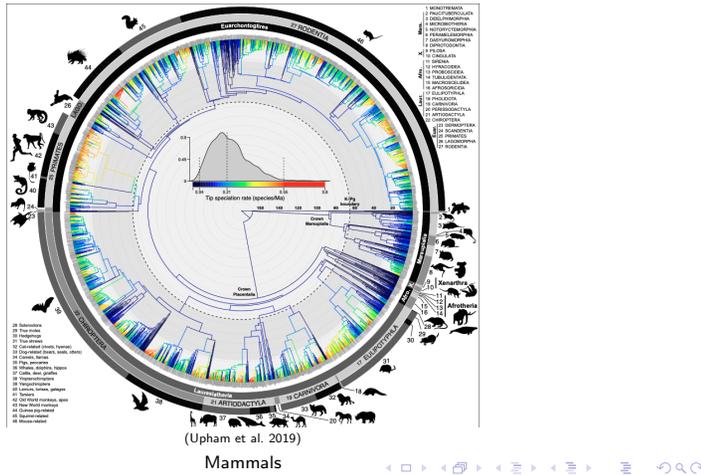


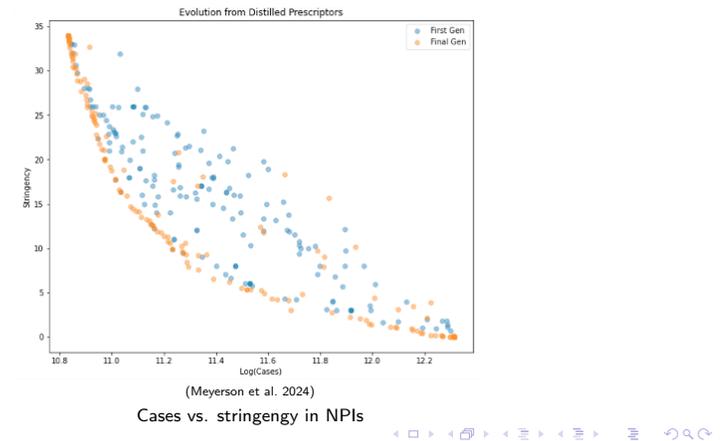
Multiobjective Optimization in Evolution

- ▶ Multiobjectivity is a natural extension of quality-diversity methods.
- ▶ Inspired by biology: organisms must balance multiple conflicting objectives.
- ▶ Solutions can be successful in many ways, promoting diversity in the population.



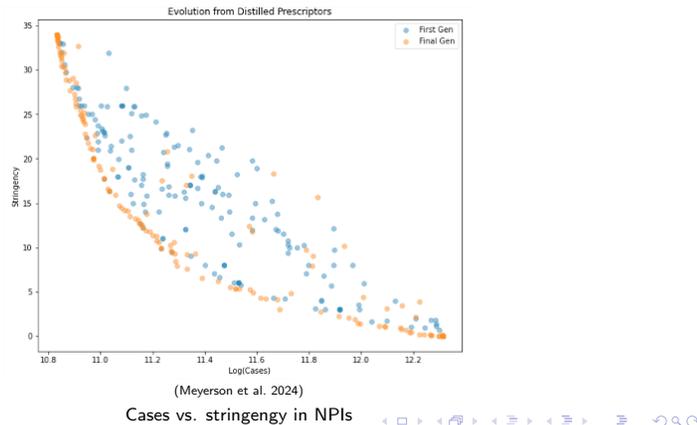
Pareto Front in Multiobjectivity

- ▶ Multiobjective optimization results in a Pareto front of solutions.
- ▶ No single solution is better across all objectives.
- ▶ Trade-offs allow multiple niches of high-performing solutions.
- ▶ Solutions on the Pareto front can be chosen based on deployment conditions or other criteria.



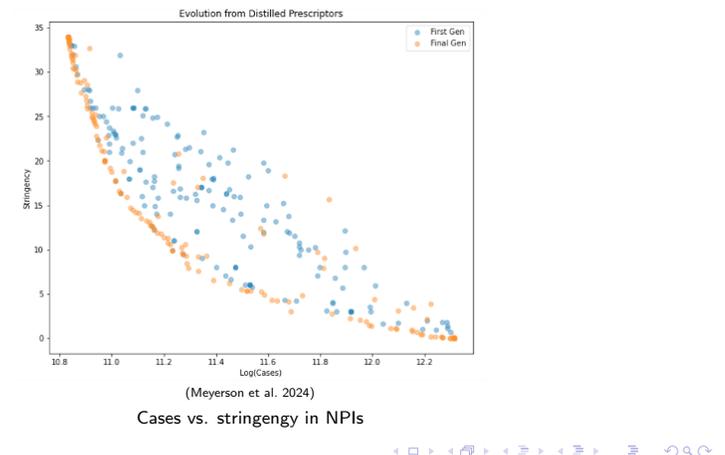
Multiobjective Evolution Methods

- ▶ Pareto front can be formed from the population of standard evolution.
 - ▶ E.g. combine objectives as a weighted average.
 - ▶ May not get a comprehensive front though.
- ▶ Multiobjective optimization method like NSGA-II can be used to evolve the Pareto front explicitly
 - ▶ Evaluate candidates in successive layers of nondominance.
 - ▶ Broad coverage as a front.



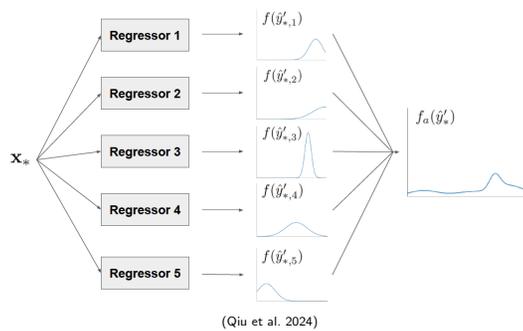
Boosting Diversity through Multiobjectivity

- ▶ Multiobjective evolution naturally encourages diversity.
 - ▶ Multiple objectives create different success paths, forming niches.
- ▶ To further increase diversity:
 - ▶ Novelty can be used as a secondary objective.
 - ▶ NEAT and other speciation methods can further enhance diversity.



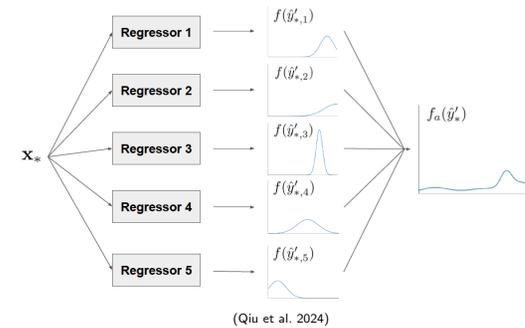
Ensembling Diverse Solutions

- ▶ Ensembling can take advantage of this diversity.
- ▶ Ensembling involves training multiple models and combining them.
- ▶ Each model contributes different insights, improving overall performance.
- ▶ Inspired by studies in psychology, business, and social science, which show that diversity improves decision-making in human teams



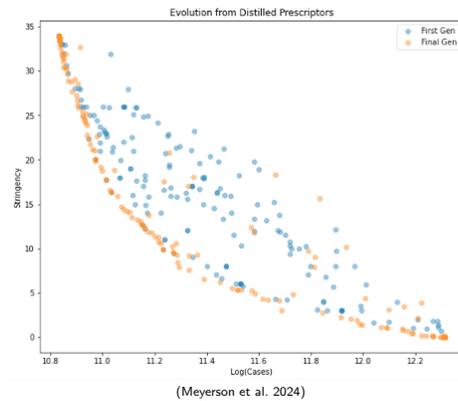
Basic Ensembling Techniques

- ▶ Simple combinations: Voting, weighted averaging.
 - ▶ All experts activated and their output combined.
- ▶ Mixtures of Experts
 - ▶ Different experts used for different input regions.
- ▶ Effective both in classification and regression; prediction and prescription.



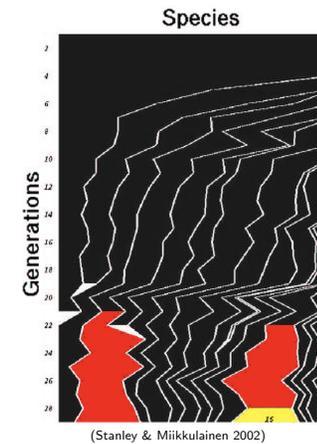
Ensembling in Evolutionary Computation

- ▶ Evolutionary Algorithms (EAs) naturally create diverse populations.
- ▶ Final population members often have different skills, forming a good ensemble.
- ▶ Multiobjective optimization enhances diversity for ensembling.



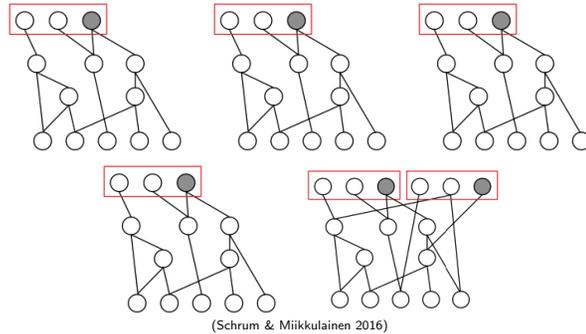
NEAT and Speciation in Ensembling

- ▶ NEAT employs a speciation mechanism to encourage diversity.
- ▶ Species champions can be used as ensemble members.
- ▶ Combine with voting, averaging, or winner-take-all for improved performance.



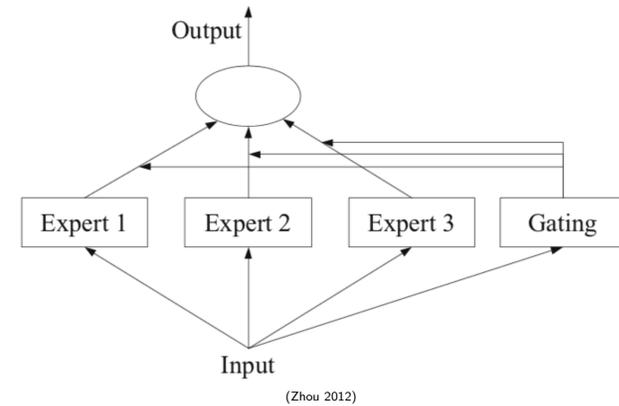
Ensembling Through Confidence Estimates

- ▶ Each network estimates whether they are the best choice to control the agent at this point.
 - ▶ A preference output, separate from the task outputs.
 - ▶ Networks bet on having the right answer, maximizing returns.
- ▶ Networks act as ensemble members with preference neurons guiding combination.
 - ▶ Can be a simple choice, or preference-weighted combination.
 - ▶ Can be evolved as separate networks, or modules in one network.



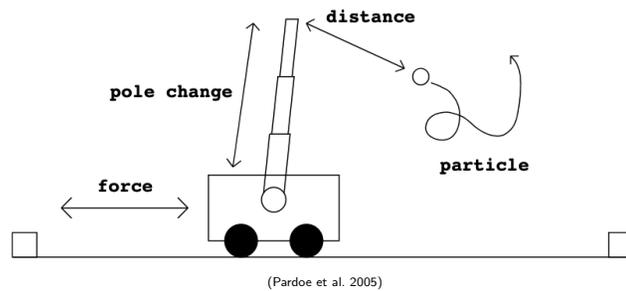
Ensembling Through a Gating Network

- ▶ A NEAT population evolved in the control task first.
- ▶ Then a gating network evolved to select which controller to use when, choosing among the species.



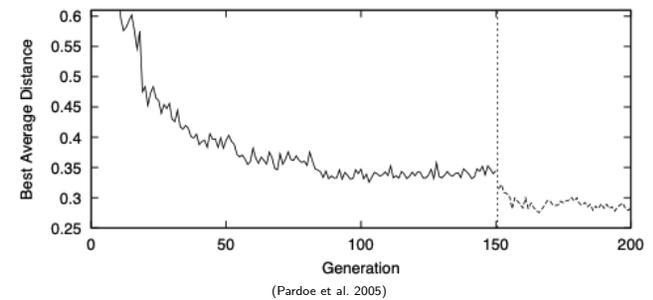
Ensembling Through a Gating Network

- ▶ Tested in the fly-swatting task:
 - ▶ An extension of the cart-pole task with more diverse state space.
 - ▶ Pushing left and right; extending and contracting the pole.
 - ▶ Aim to keep the pole tip on the target.



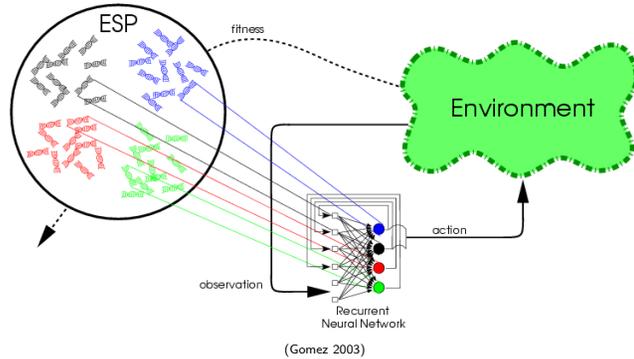
Gated Ensembling Example

- ▶ The gating network partitions the space and uses different NEAT networks at different times.
- ▶ Best results with ensemble size of eight.
- ▶ Gated ensembling significantly boosts performance.



Ensembling in Enforced Sub-Populations

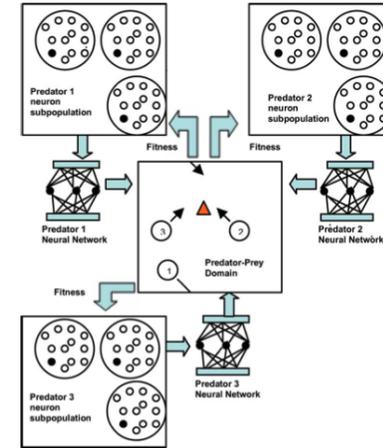
- ▶ Enforced SubPopulation (ESP) method evolves each neuron of the network in a separate subpopulation.
 - ▶ Each neuron encodes its own weights.
 - ▶ A network is formed by selecting randomly from subpopulations.
 - ▶ The neurons inherit the fitness of the network; they evolve to cooperate.
- ▶ Good network require different neurons; diversity is thus encouraged.



◀ ▶ ⏪ ⏩ 🔍 🔄

Hierarchical ESP

- ▶ Hierarchical ESP extends the approach to teams of networks.
- ▶ It forms a principled ensemble:
 - ▶ Each neuron and each network is evolved for a specific role.
 - ▶ Not just diversity, but optimized diversity.
- ▶ Particularly powerful in cooperative multiagent tasks.
 - ▶ E.g. a team of predators capturing a prey:

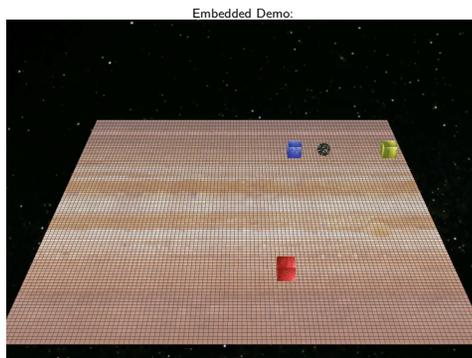


(Rawal et al. 2010)

◀ ▶ ⏪ ⏩ 🔍 🔄

Hierarchical ESP Example

- ▶ Particularly powerful in cooperative multiagent tasks.
- ▶ E.g. a team of predators capturing a prey:
 - ▶ One network chases, another captures the prey.
 - ▶ One neuron turns towards the agent; another away from teammate.

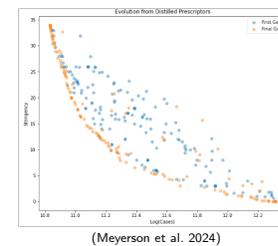


(<https://nn.cs.utexas.edu/?armsrace>)

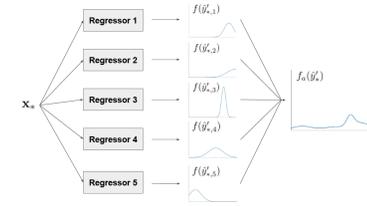
◀ ▶ ⏪ ⏩ 🔍 🔄

Multiobjectivity and Ensembling Conclusions

- ▶ Multiobjective optimization promotes diversity.
 - ▶ Pareto fronts are diverse by definition.
 - ▶ Especially by making novelty a secondary objective.
- ▶ Ensembling is a powerful way of taking advantage of this diversity.
 - ▶ More robust decision-making in complex domains.
 - ▶ Can be extended with various techniques to suit specific problems.
- ▶ Both are natural extensions of population-based search.



(Meyerson et al. 2024)

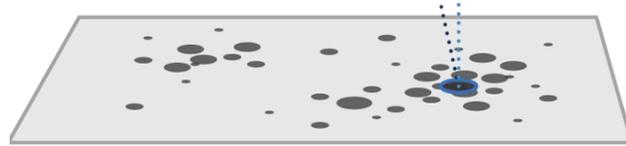


(Qiu et al. 2025)

◀ ▶ ⏪ ⏩ 🔍 🔄

Population Culture in Evolution

- ▶ Population culture refers to the knowledge across individuals.
- ▶ Includes both common behaviors and unique knowledge.
- ▶ Can be used to improve evolution in several ways: multiobjectivity and ensembling; culling, training, selection, pruning...

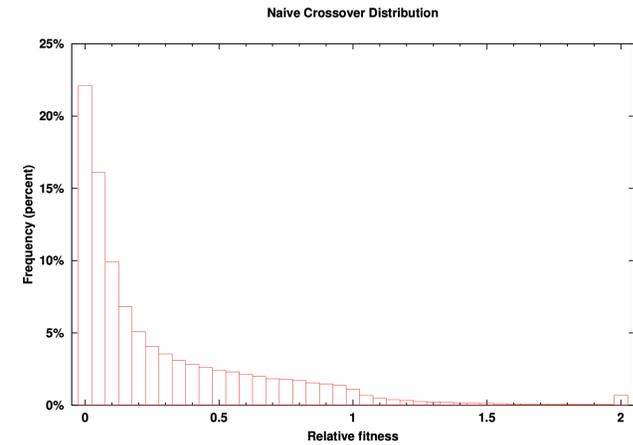


(Miikkulainen & Forrest 2021)

◀ ▶ ⏪ ⏩ 🔍 ↺

Culling: Speeding up Evolution

- ▶ Culling generates a large number of offspring, only keeping the most promising.
- ▶ Efficient because "most crossovers are awful".

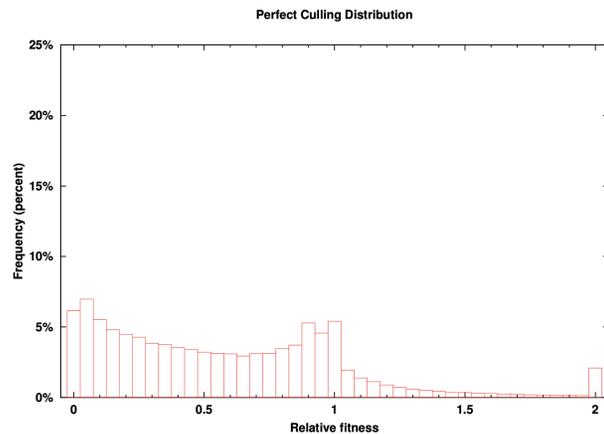


(McQuesten & Miikkulainen 1997)

◀ ▶ ⏪ ⏩ 🔍 ↺

Recognizing Promising Offspring

- ▶ Approximate evaluations help recognize good offspring fast.
 - ▶ Use a syllabus of inputs and compare answers to prominent population members.
 - ▶ Effectively identifies non-viable offspring.
 - ▶ Expensive fitness evaluations not necessary.
- ▶ Can speed up neuroevolution by a factor of 3 in tasks like pole balancing.
- ▶ E.g. culling from a litter of 8:

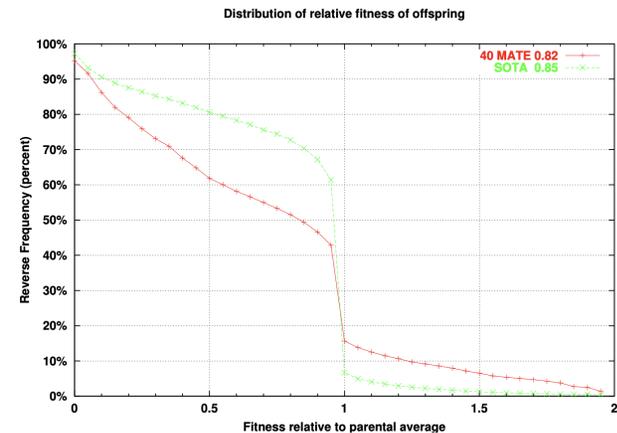


(McQuesten & Miikkulainen 1997)

◀ ▶ ⏪ ⏩ 🔍 ↺

Cultural Selection of Parents

- ▶ Parents are chosen based on diversity, not just fitness.
- ▶ First parent based on fitness, second chosen as maximally different (in its answers to the syllabus).
- ▶ Increases the chance of combining complementary strengths in offspring.
- ▶ Improves even though the second parent has low fitness:

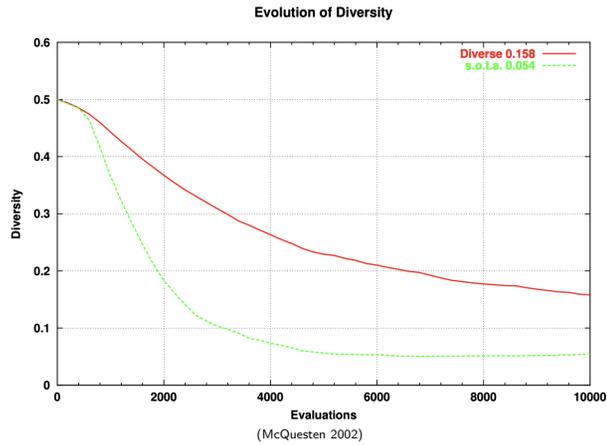


(McQuesten 2002)

◀ ▶ ⏪ ⏩ 🔍 ↺

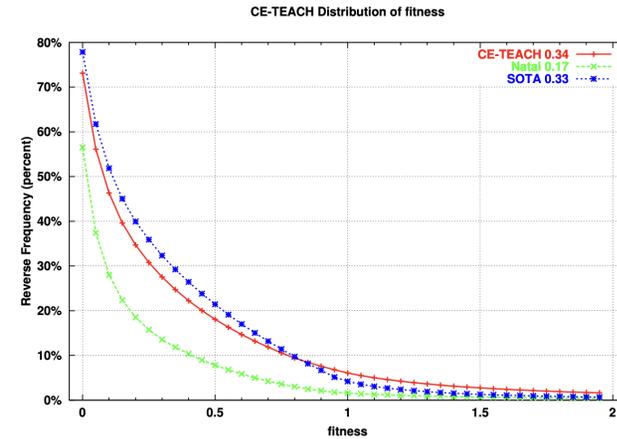
Maintaining Diversity in the Population

- ▶ The syllabus can also be used to decide which solutions to discard.
- ▶ Find two closest solutions, discard the one with lower fitness.
- ▶ Increases diversity and accelerates evolution by 30%.



Using Culture to Enhance Learning

- ▶ Leverage population champions' behaviors as a training set.
- ▶ Select the offspring that performs well after training (i.e. utilize the Baldwin effect, not Lamarckian evolution).
- ▶ Speeds up neuroevolution by an order of magnitude.
- ▶ But performance at birth is poor! What's going on?

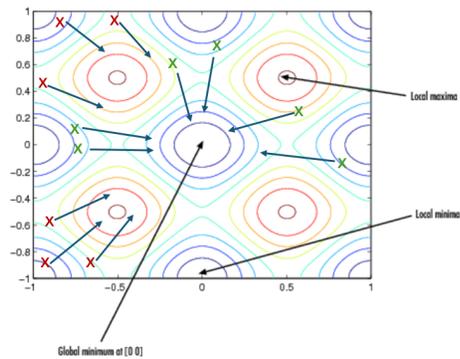


(McQuesten 2002)



Synergetic Development

- ▶ We've seen this before: Evolution discovers good starting points for learning rather than near-optimal solutions.
 - ▶ Learning will happen, so evolution discovers how to take it into account.
 - ▶ A synergy between learning and evolution.
- ▶ Solutions are more robust and more effective.

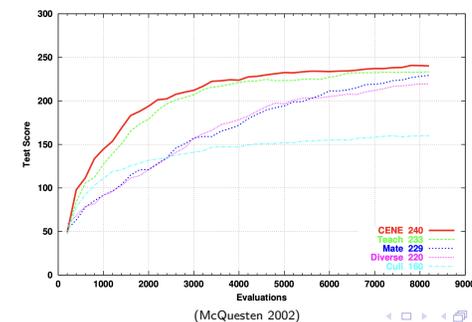
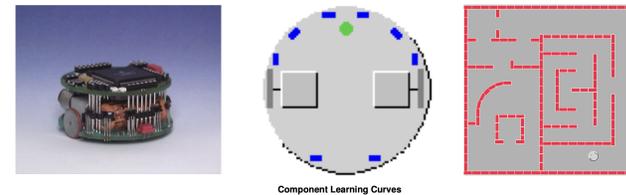


(Adapted from Najem & Yusuf 2018)



Putting it Together

- ▶ Culture helps several aspects of evolution.
- ▶ Which methods are the most effective depends on the problem; Can be combined for a robust effect.
- ▶ E.g. the simulated Khepera maze running task:

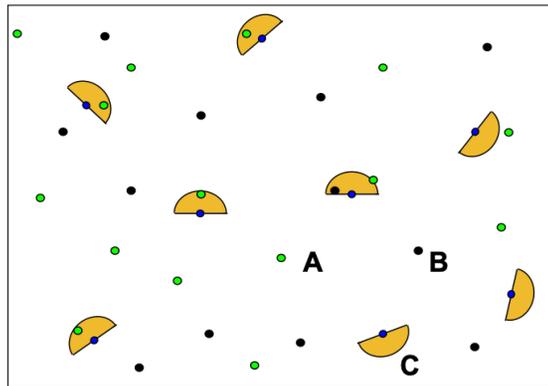


(McQuesten 2002)



Egalitarian Social Learning (ESL)

- ▶ Not just the champion, but anyone can have useful knowledge.
- ▶ Learn from any other agent's success in specific situations.
- ▶ Training examples from the entire population culture.
- ▶ Promote diversity by dividing population into species (or subcultures).

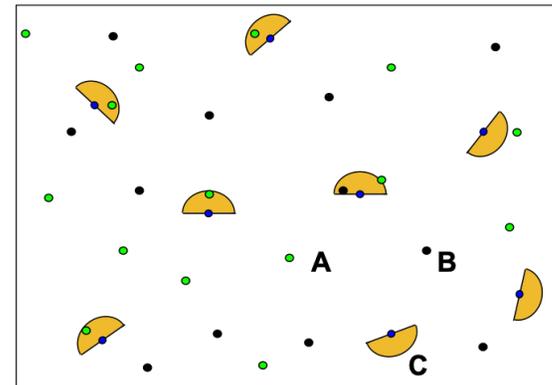


(Tansey et al. 2012)



Foraging Domain Example

- ▶ Agents have limited view, variable speed, and forage for food that vary in value (good, bad, poison)
- ▶ Different strategies evolve: move a lot / don't miss anything,
- ▶ If an agent receives a low reward when another receives a high reward in the same situation, learn.

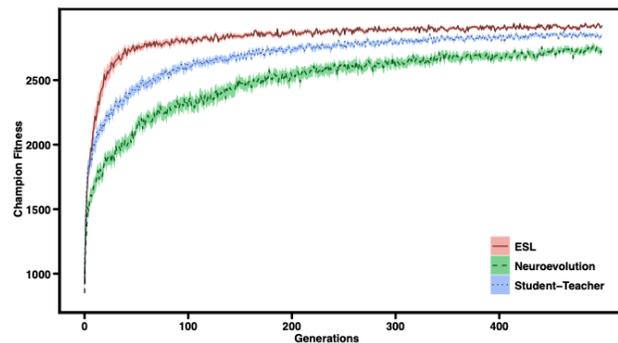


(Tansey et al. 2012)



Results from ESL in Foraging Domain

- ▶ Good ideas propagate (e.g. slowing down not to miss)
- ▶ ESL learns faster than direct neuroevolution and student-teacher approach.
- ▶ Demonstrates value of diversity and social learning.
- ▶ A life lesson! Diverse teams perform better.



(Tansey et al. 2012)



Demo of NEAT vs. ESL



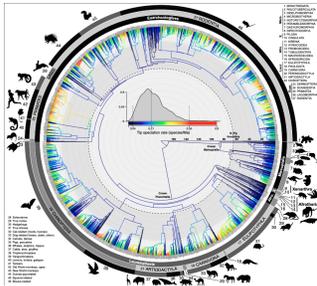
Demo link:
<https://youtu.be/M5i-1WDNeH8>

Demo link:
<https://youtu.be/o0-kWNabCq8>



Conclusion: The Importance of Diversity in Evolutionary Computation

- ▶ Why Diversity Matters:
 - ▶ Diversity is essential for robust and adaptive search in evolutionary computation.
 - ▶ It prevents premature convergence, enhances exploration, and enables discovery of better and more creative solutions.
- ▶ Methods work at different levels: genetic, behavior, ensembles, population; culling, selection, discarding, teaching; objectives, ensembles; can be combined?
- ▶ Analogies to biology, society.



(Upham et al. 2019)



(<https://youtu.be/o0-kWNabCq8>)

