Evolution & Learning

Lecture 10   I400/I590

Artificial Life as an approach to Artificial Intelligence

Larry Yaeger

Professor of Informatics, Indiana University

Adaptation

- Organisms learn useful *adaptations* during their lifetime
- These adaptations embody knowledge about the environment gained through exploration and experimentation
- It seems wasteful to lose this knowledge between generations
- Lamarckian evolution would be the obvious solution: Transfer acquired knowledge back into the genome
  - However, so far, Lamarckian evolution has never been substantiated in any biological system

Waste Not, Want Not

- In the absence of Lamarckian evolution, one might conclude that lifetime learning cannot impact evolution
- One would be wrong
- Learning *can* be very effective in guiding evolutionary search
  - Even when that learning is *not* communicated to the genotype
  - Even when that learning is unrelated to specific survival tasks or the fitness selection criteria

The "Baldwin Effect"

- Useful adaptations allow an organism to survive and reproduce, increasing its evolutionary fitness
- Evolution then selects for organisms that are ever more capable of learning these adaptations

The Baldwin Effect & Innateness

- Such adaptations may keep organisms alive long enough, and populous enough, for evolution to select for organisms in which the adaptations are innate
  - Evolution in the direction of innateness may make the learning easier
  - Innateness may be more efficient in both consumed resources and reaction time
- Such adaptations may restructure the environment so as to foster evolution of heritable versions of the relevant traits

History

- James Mark Baldwin, in 1895 & 1896, describes "Organic Selection", a mechanism by which evolutionary fitness prefers solutions capable of better adaptation
    - Specifically notes Lamarckian-style influences without the need for Lamarckian transfer of acquired characteristics to the genome
    - *A New Factor in Evolution* and other papers in 1896 are primary, but 1895's book, *Mental Development in the Child and the Race*, first broaches the subject
    - Refers to work of C. Lloyd Morgan and Henry Osborn

History

- Conwy Lloyd Morgan, in 1896, publishes *Emergent Evolution*, using similar reasoning and espousing the same effect
    - Better known for Morgan's Canon, a kind of Occam's Razor for Psychology:

      "In no case may we interpret an action as the outcome of the exercise of a higher psychical faculty, if it can be interpreted as the outcome of one which stands lower in the psychological scale."

History

- Henry Fairfield Osborne, in 1896, publishes *A mode of evolution requiring neither natural selection nor the inheritance of acquired characteristics*, also proposing this phenomenon under the name "coincident selection"

History

- Controversy exists over who published first
  - Morgan spoke immediately prior to Baldwin at the same conference where these ideas were first presented orally
  - Osborn published a few days before Baldwin's 1896 papers
  - Term "Organic Selection" used more broadly in first edition of 1895 *Mental Development…* book

History

- Conrad Hal Waddington, in 1942, publishes *Canalization of Development and the Inheritance of Acquired Characters*, proposing a similar "genetic assimilation" feedback mechanism between developmental processes and natural selection
  - Variable ontogenetic responses to environmental conditions, that confer an evolutionary advantage on organisms, may then be selected for in a heritable form

Biological Evidence

- Limited evidence or references under the name "Baldwin effect" in evolutionary biology literature
- There is related support in the area of "niche creation"
- Waddington's own experiments (1952 and later) with *Drosophila* demonstrate developmental equivalent
  - Heat applied during embryonic development produces reduction in wing veins
  - After several generations of selecting for organisms most responsive to embryonic application of heat, reduction in wing veins persists in large fraction of offspring without the application of heat

Computational Modeling Evidence

- Strong evidence in this realm
- We are reading the classic 1987 paper by Geoff Hinton and Steve Nowlan, entitled *How Learning Can Guide Evolution*
- Imagine an organism containing a neural net with many potential connections, only one configuration of which confers additional evolutionary fitness
  - This is a worst-case fitness landscape, consisting of a flat plane everywhere, except for a single spike
  - "The good net is like a needle in a haystack."

(Note that this being a "neural net" is irrelevant; what we have is a 20-dimensional state space, with only one "good" state.)

A Needle in a Haystack

- Now let organisms "learn" during their lifetimes

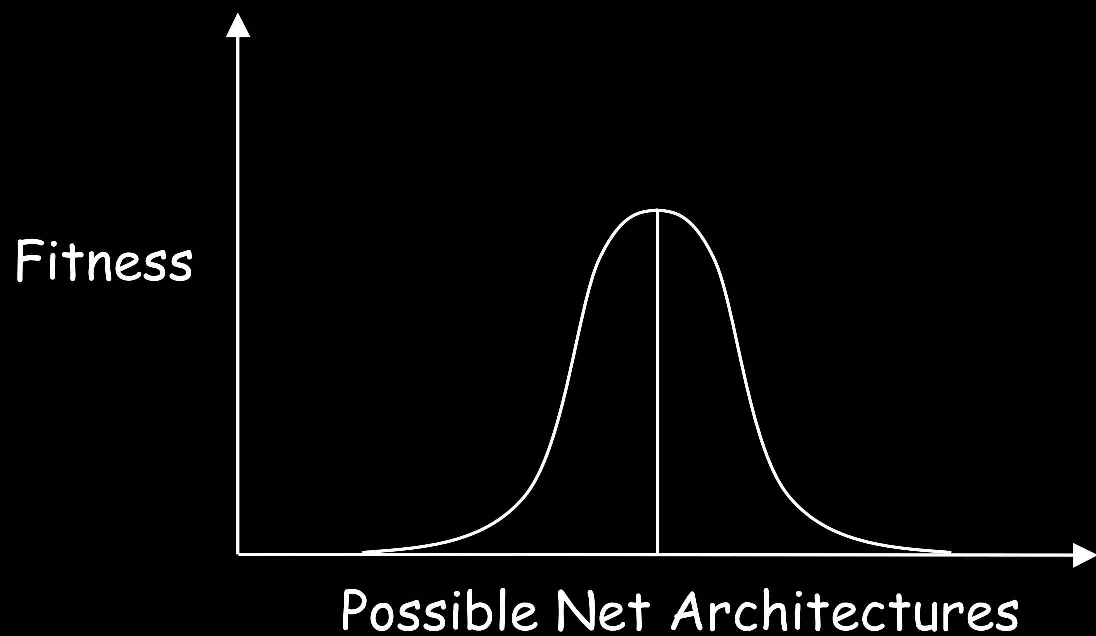

- "It is like searching for a needle in a haystack when someone tells you when you are getting close."

Simulating the Baldwin Effect

- The neural net has 20 potential connections
- The genotype has 20 genes (one per connection)
  - Each gene has 3 states (alleles)
    - 0: connection is absent
    - 1: connection is present
    - ?: connection may be switched between absent and present by lifetime "learning"

“Learning“

- Learning consists of randomly flipping all switches on every trial
- If the single, correct network is ever generated as a result of learning, the switches are frozen, else they continue to change with each time step
- Note that gradient descent is not possible with this algorithm
  - But flipping the switched connections does explore the space in the vicinity of the innate genetic prescription

GA Parameters

- 1000 organisms in each generation
- 1000 learning trials performed by each organism during its lifetime
- Initial 1000 organisms generated by selecting each allele randomly, with probability of 0.5 for ? and 0.25 for 0 and 1
- Note that since, on average, there are 2^{10} possible states to be explored by switching and there are approximately that number of learning trials, there is a reasonable chance that an organism with the correct genetically specified connections may learn the correct specification of the remaining 10 connections

GA Parameters

- 1000 offspring are generated from pairs of organisms chosen randomly with a probability that is proportional to 1 + 19n/1000
  - n = number of learning trials that remain after the organism has learned the correct network
  - An organism that learns the correct solution immediately will be 20 times as likely to reproduce as an organism that fails to learn the correct solution
- Offspring are generated using single-point crossover

Evolution of Alleles

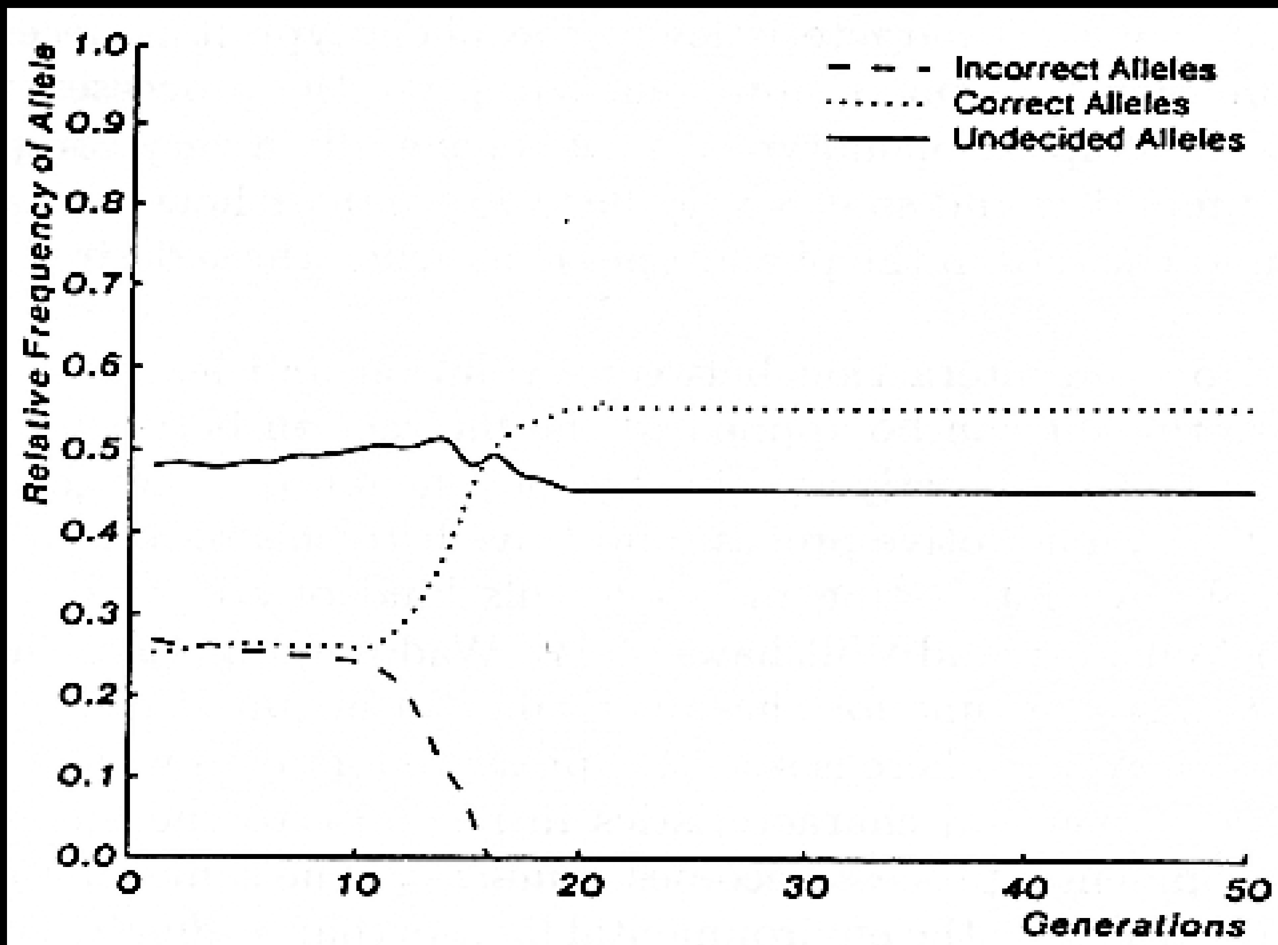

Learning Guides Evolution

- Total number of organisms generated was far less than the 2^{20} that would be expected to find the solution by a random or purely evolutionary search
  - Despite the fact that learning was itself a random search over 2^{10} settings (slightly fewer as evolution converged)
    - More structured learning should only enhance the effect
  - Number of learned connections not significantly reduced
- The same problem was *never* solved by an evolutionary search without learning

More Computational Evidence

- Another of our readings, *Learning, Behavior, and Evolution*, from 1991, by Domenico Parisi, Stefano Nolfi, and Federico Cecconi, extends Hinton & Nowlan's work
- Provides evidence for the Baldwin effect in
  - Structured learning
  - Non-adaptive learning
    (learning of tasks not intrinsically correlated with fitness or the selection criteria)

Learning, Behavior, and Evolution

- Demonstrates evolutionary selection of learnability (again for both fitness-correlated and -uncorrelated tasks)
- Shows how self-selection of incoming stimuli can guide evolution
- Suggests a consistent explanation of all these effects in terms of synaptic weight spaces and fitness landscapes

Simulation Setup

- Organisms (O) live on a 2D grid containing randomly distributed pieces of food
- Each O is modeled by a feedforward neural network
  - Inputs are normalized (0.0 to 1.0) angle and distance to nearest food
  - Plus previous time-step's outputs
  - Outputs are a coded representation of four possible outcomes:
    - Move forward
    - Turn left
    - Turn right
    - Stay still

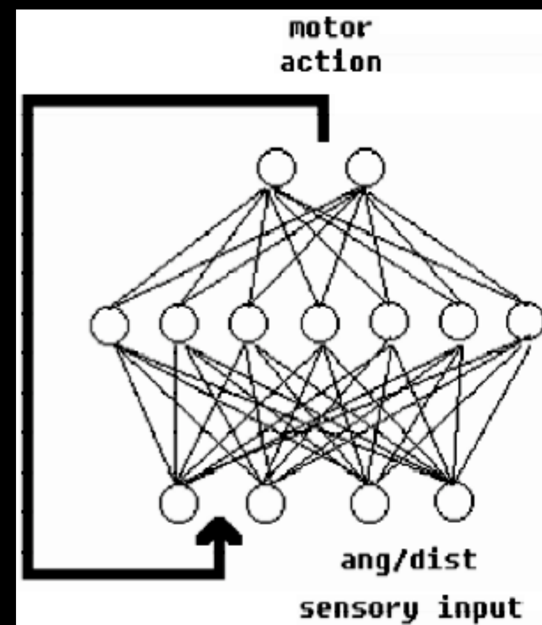

Simulation Setup

- Seed a population with 100 organisms
- Initialize network weights randomly
- Os "live" for 20 epochs, where an epoch consists of:
  - 50 actions in each of 5 different environments (results in 250 actions per epoch, 5000 actions total)
- Environment is a grid of cells with 10 randomly distributed pieces of food
- Os are placed in individual copies of the environment (i.e., they live in isolation)

Reproduction

- After 20 epochs, the 20 Os which have accumulated the most food in the course of their movements are allowed to reproduce by generating 5 copies of their weight matrix
- Mutations are introduced by selecting 5 weights at random and adding a random value between -1.0 and +1.0
- This process continues for 50 generations

With and Without Learning

- The experiment just described is first run without learning, so the weights are static throughout the Os' lives
- The same experiment is then run with the Os learning to predict at time T the sensory input they will perceive at time T+1 (given the motor action they take at time T)
  - Uses Backprop
  - Learned weight changes are discarded (only genetically selected weights are propagated to offspring)

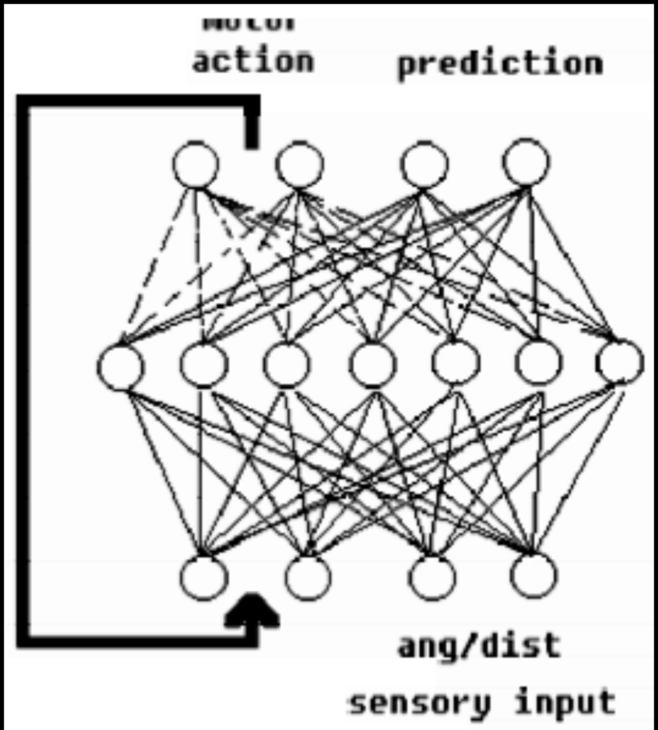

Evolution With and Without Learning

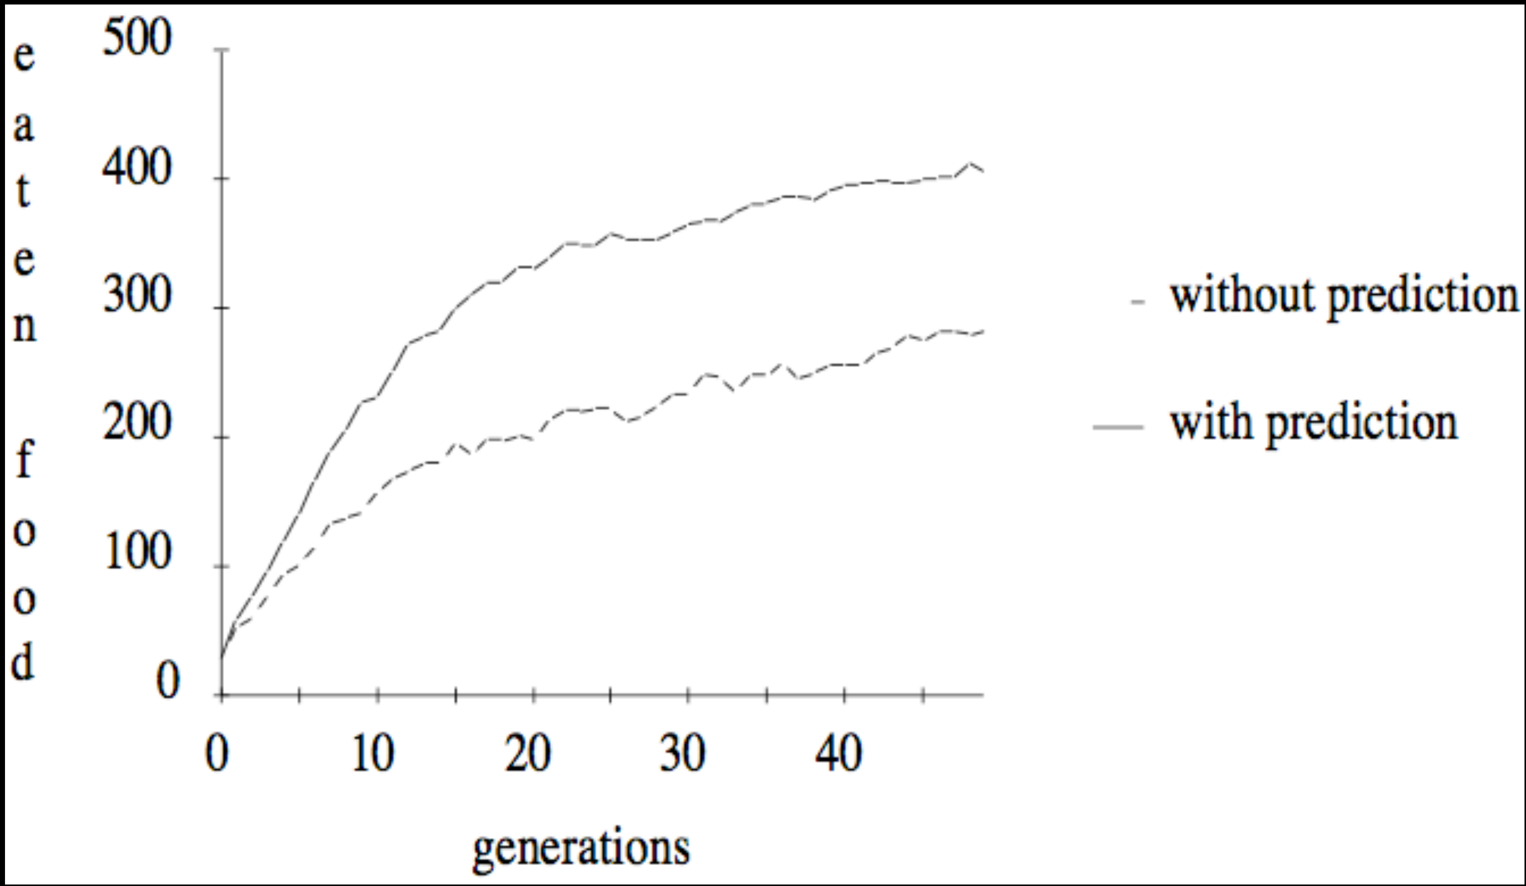

Evolution Guides Learning

- Parisi *et al* state that the difference between those two curves cannot be explained solely on the basis of lifetime learning
- Evolution had to select for better initial conditions—better performance at birth and better learnability—in order to account for the magnitude of the observed difference

How Evolution Guides Learning

- Think of the set of possible weights as a high-dimensional space, with one more dimension that corresponds to evolutionary fitness
- Learning then explores that fitness landscape in the vicinity of the site defined by purely evolutionary fitness

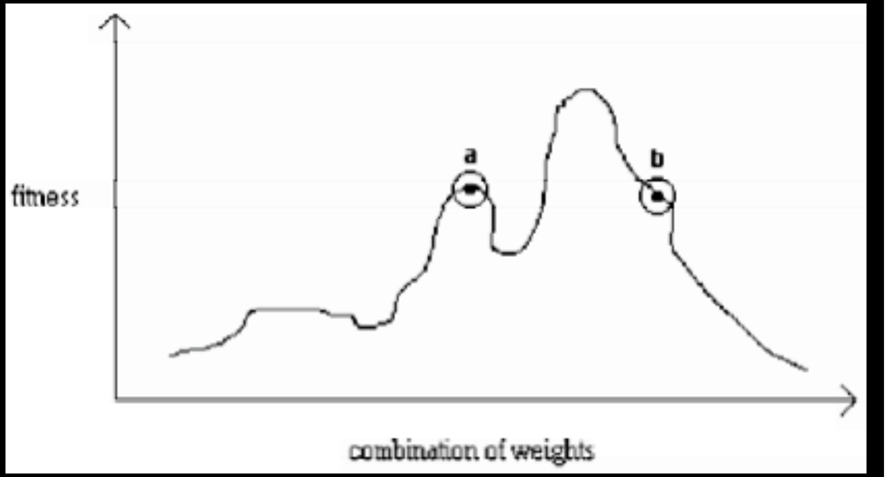

- Learning thus lets evolution preferentially select for genotypes that will produce fitter phenotypes, even when that greater fitness cannot be discerned from genotypical fitness

How Evolution Guides Learning

- Since the average fitness of an organism's offspring will be determined by the neighborhood of the parent organism's genetic fitness, allowing evolution to select for higher average learned fitness (from that same neighborhood) also increases the average genetic fitness of the offspring

Random Learning Guides Evolution

- Even when the prediction task is changed to predict random output values, Parisi et al find that evolution is accelerated, just not as much
- Might be explained by selection for better initial weight values (smaller)
- But even a random sampling around the evolutionary fitness site can yield a better estimate of potential fitness for offspring

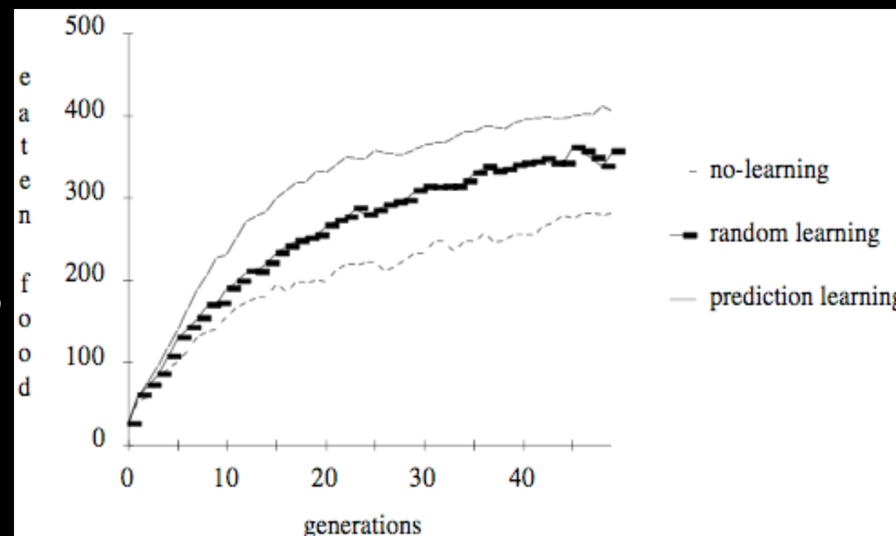

Random Sampling *Can* Predict Fitness

- The average fitness of points randomly sampled around **a** will be lower than the average fitness of points randomly sampled around **b**

All Learning is Not Equal

- Intelligent, structured exploration of the fitness landscape would reasonably be expected to provide a better estimate of potential fitness
- A random sampling around a and b would produce the same average, but learning that explores the better parts of the local fitness landscape will produce a higher average fitness for b

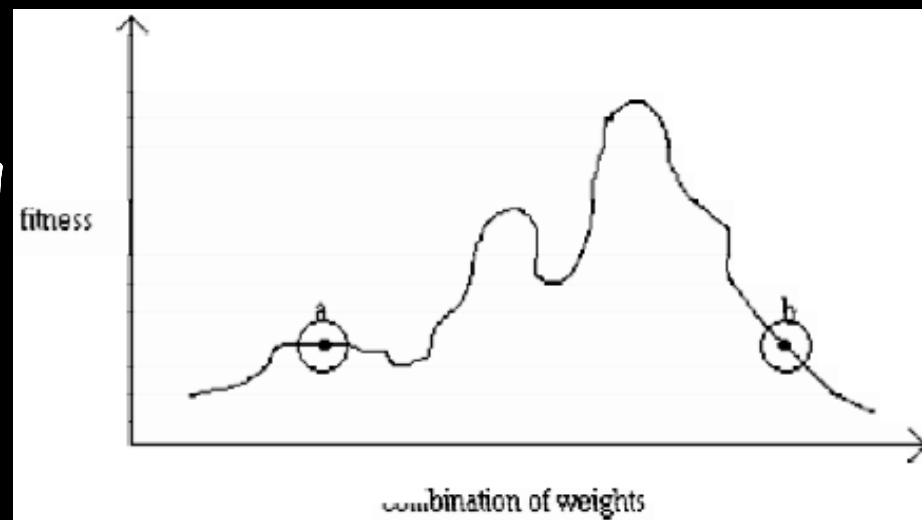

- Also suggests that a meta-level search for better learning algorithms would be of evolutionary value

Selection for Learning of Related Tasks

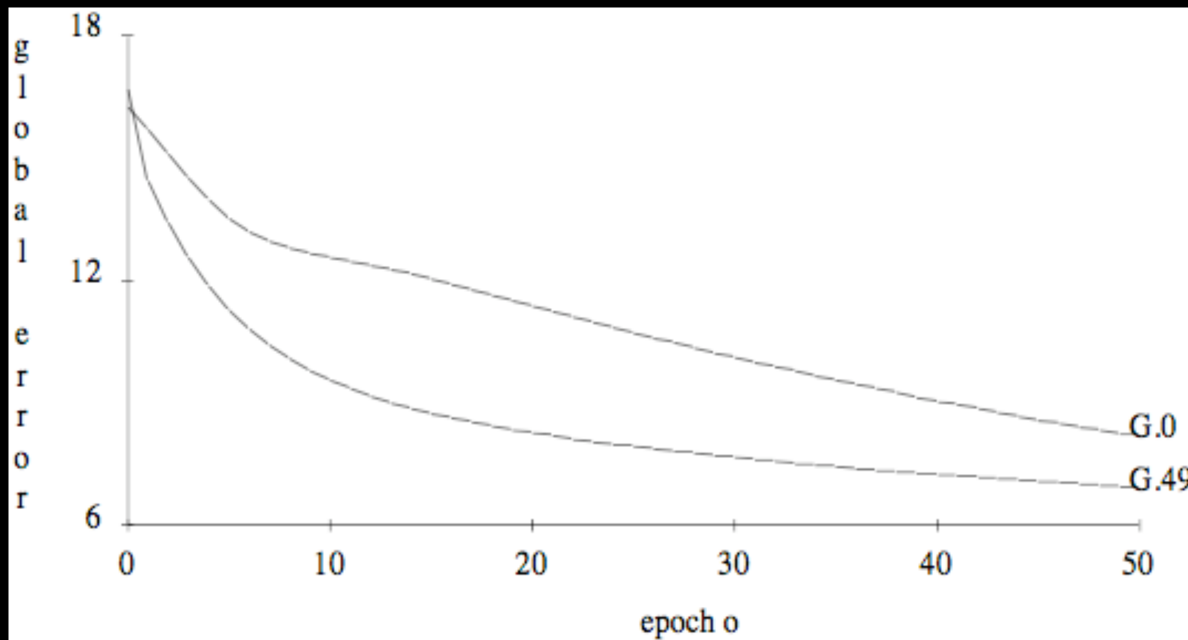

- Analysis of the global error rate for the prediction task at early and late generations demonstrates an inheritance of the *ability to learn* the particular task, although not directly of the *ability to perform* the task

An Unrelated Task

- Replace the input-stimuli prediction task with a task unrelated to evolutionary fitness: XOR
- *An* XOR output unit is trained to have a value of 0.0 if *both* input units have an activation which is greater *or* less than 0.5, and a value of 1.0 otherwise
- Selection is still based entirely on performance on the food gathering task

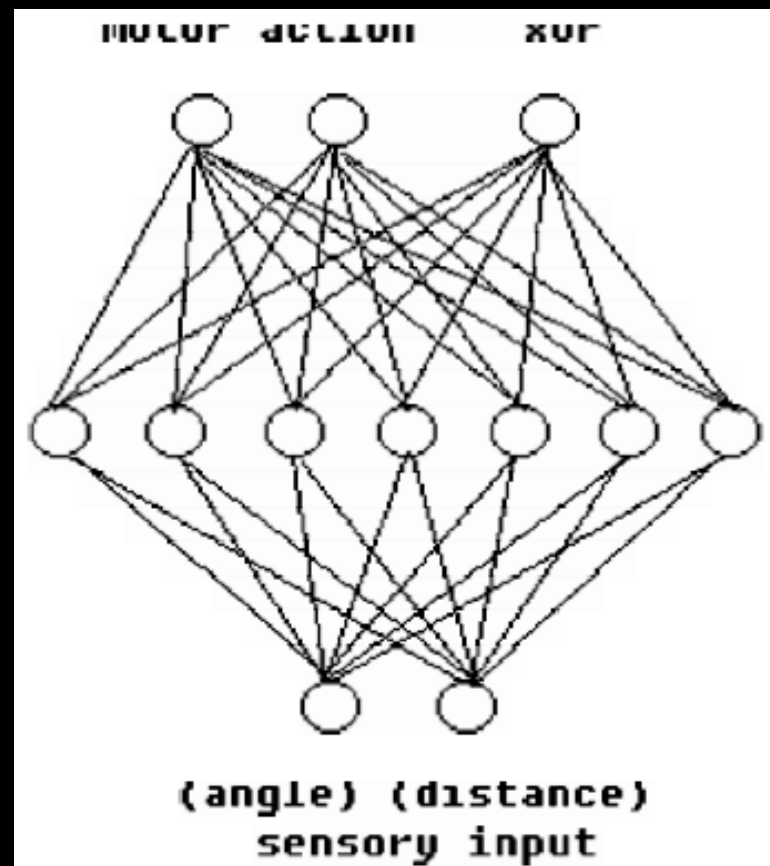

Selection for Learning of Unrelated Tasks

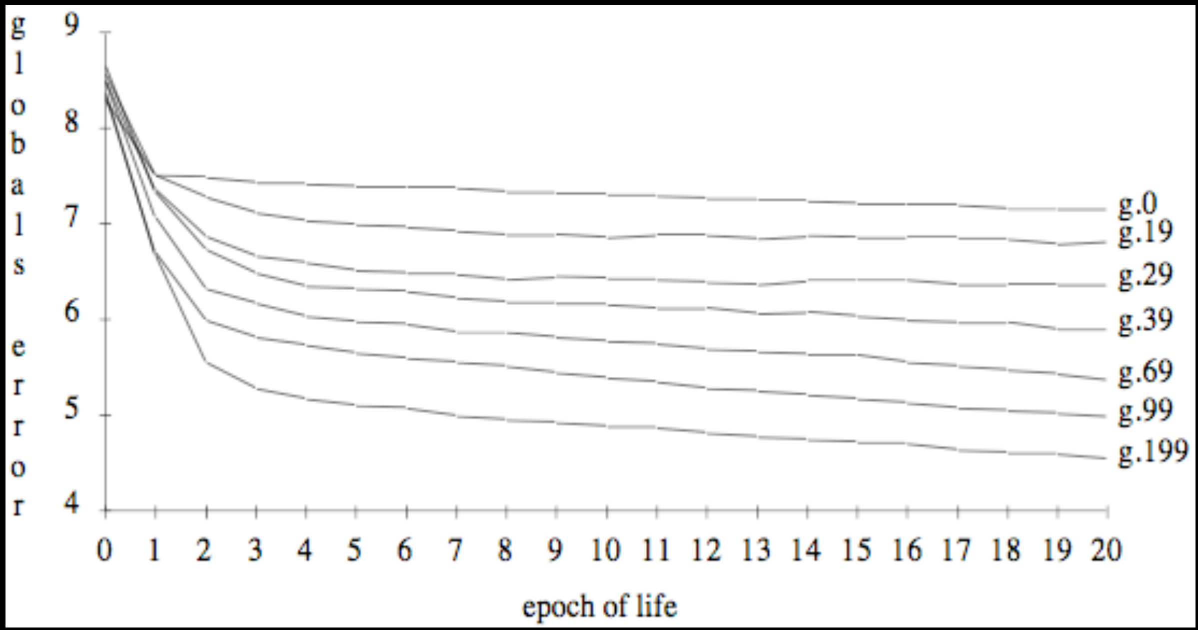

- Analysis of the global error rate for the XOR task again demonstrates an inheritance of the *ability to learn* the task, although not directly of the *ability to perform* the task

Correlated Sub-Regions of Fitness

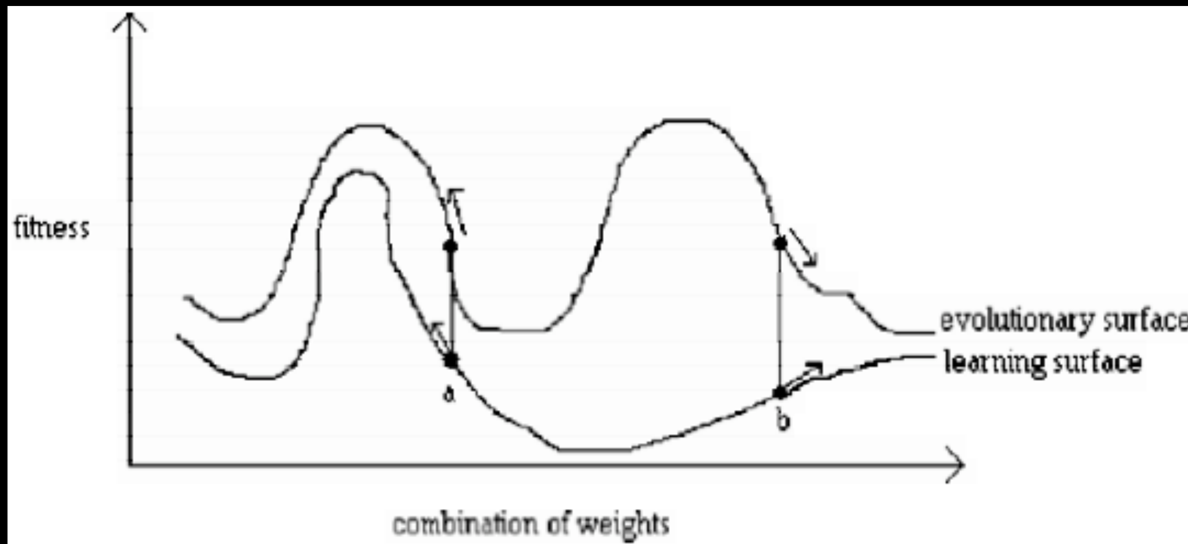

- Even when the tasks are globally uncorrelated there may be correlated sub-regions in the fitness landscape
- But there may also be regions of anti-correlation
- Here again there may be a simple, global initial weight magnitude explanation

Phenotypic Variability

- For simplicity of discussion and analysis, we have so far assumed that each genetically specified weight matrix yields a single corresponding fitness value
- Actually, each such genetic specification may elicit a variety of phenotypes, upon which selection acts, due to
  - The particular environment in which the O live
  - The sequence of experiences the O has
  - The changes in subsequent input, and even the environment, resulting from the O's behavior

Influence of Behavior on Evolution

- *An* O's behavior alters its subsequent input
- Os may evolve one of at least two strategies:
  - Optimally handle all possible inputs *or*
  - Optimally handle a subset of possible inputs and behave in such a way that they are more likely to experience that subset of inputs
- To determine which strategy is being employed, divide the perceived angle between the O and food (as provided to its input unit), into 10 evenly distributed, 36° bins, and compute the frequency with which stimuli in each bin are encountered by a particular O

A Preferred Orientation

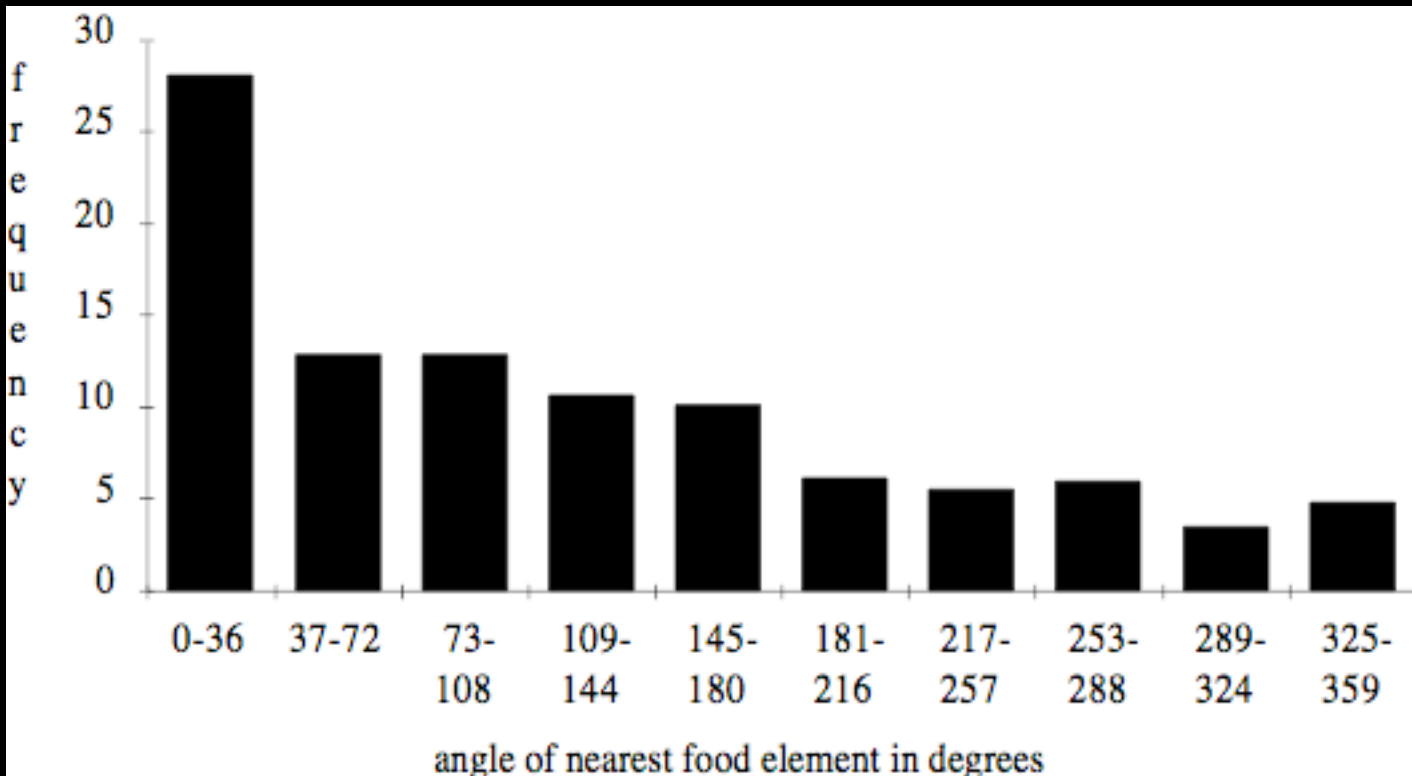

- Note the asymmetry of stimulus occurrence near 0°/360°

Performance *vs.* Orientation

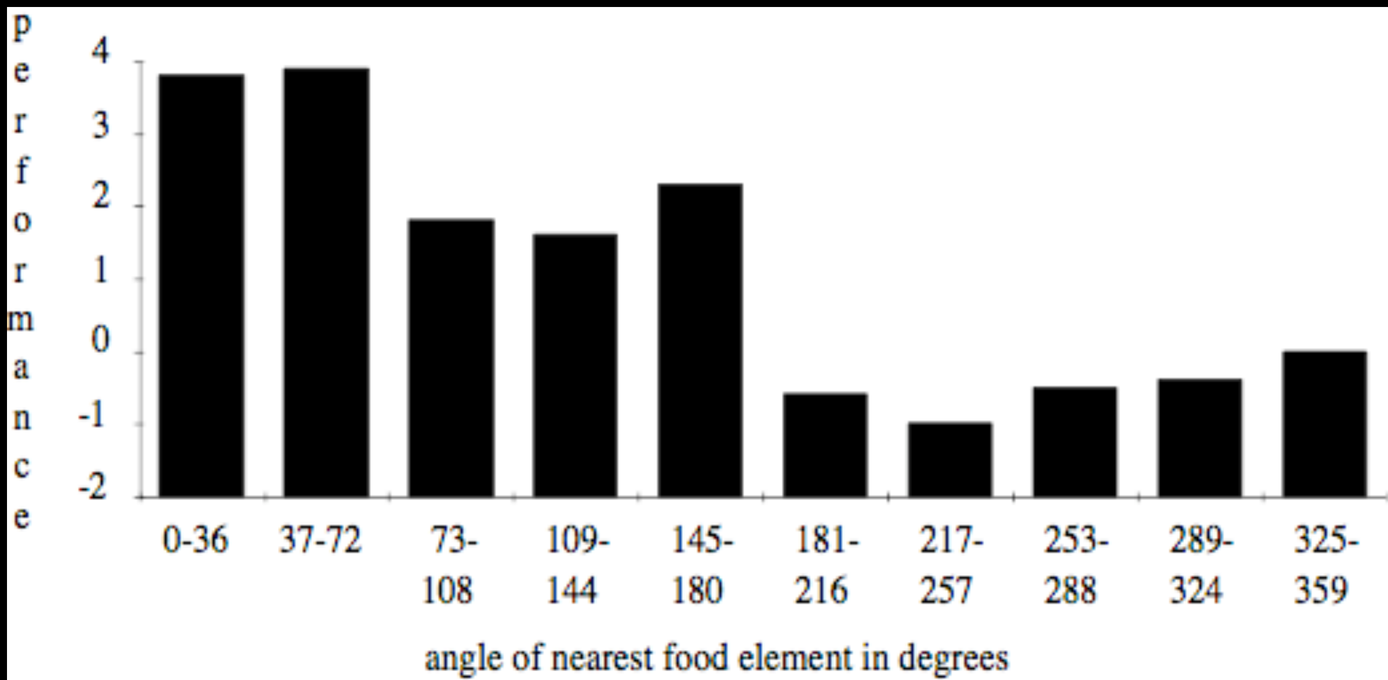

- Performance (decrease in the distance between the O and the food after one action) is better for the asymmetrically preferred orientations

Behavior-Dependent Performance

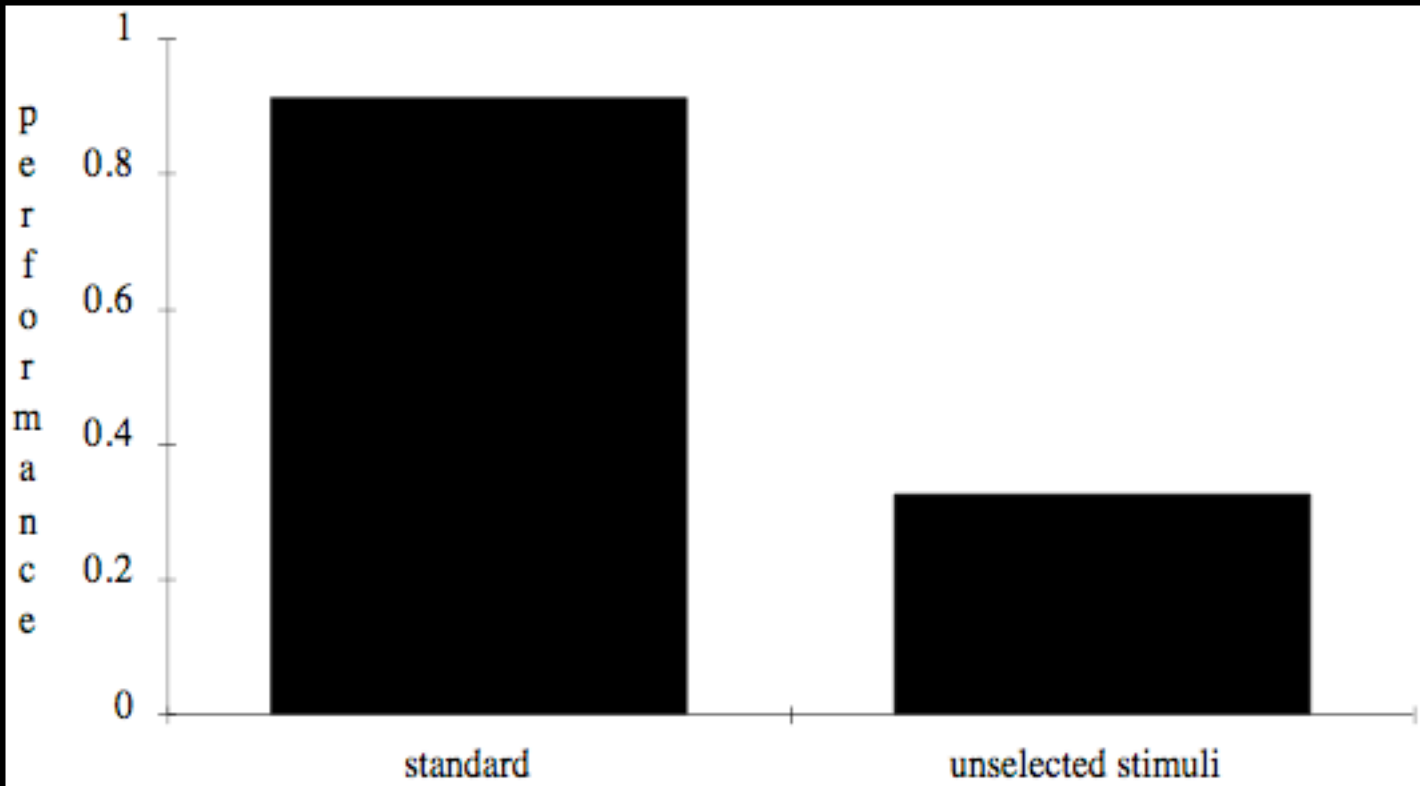

- *Average single-step performance is much better when* the current food location depends on the O's previous action (standard case) *vs.* when the location is randomized by reorienting the O at each time step

Behavior Guides Evolution

- At least some Os have evolved an orienting behavior which, in turn, has guided evolution to optimize behavior for these self-selected stimuli
- The ability to react efficiently to all classes of stimuli (all relative food angles) is still of some evolutionary value, because infrequent stimuli to which Os do not respond optimally may still appear
- However, the benefits of such a generalized capacity may be small compared to the benefits of the specialized capacity to respond to self-selected stimuli, so there may never be sufficient evolutionary pressure for the generalized capacity to evolve

Evolution of Learning

- Having seen how learning and behavior can guide evolution, our final reading assignment this week (David Chalmers's 1991 *The Evolution of Learning: An Experiment in Genetic Connectionism*) shows how learning itself may be evolved
- Chalmers speaks of learning as first-order adaptation, and natural selection as second-order adaptation, evolving and improving the ability to learn
- To explore the evolution of learning, he attempts to evolve a supervised learning algorithm for a neural network with a single layer of weights

Problem Statement

- Define

  a_j = activation of the input unit j

  o_i = activation of the output unit i

  t_i = the training or target value of output unit i

  w_{ij} = the connection strength from input j to output i
- The genome must encode

  $\Delta w_{ij} = F(a_j, o_i, t_i, w_{ij})$

Problem Statement

- Choose a general quadratic form for F

$$\Delta w_{ij} = k_0 (k_1 w_{ij} + k_2 a_j + k_3 o_i + k_4 t_i + k_5 w_{ij} a_j + k_6 w_{ij} o_i + k_7 w_{ij} t_i + k_8 a_j o_i + k_9 a_j t_i + k_{10} o_i t_i)$$

Genetic Encoding

- Encode all 11 constants in a 35 bit genome
  - First 5 bits code for the scale parameter k_0
    - First bit used for sign
    - Remaining bits encode 1/256, 1/128, …, 32, 64
  - Other 30 bits encode the 10 coefficients, using 3 bits per coefficient
    - First bit used for sign
    - Remaining bits encode 0, 1, 2, or 4

Task Definition

- Define 30 linearly separable classification tasks
- Randomly select 20 of these classification tasks for each run
- Measure the performance of each learning algorithm in the current population on all 20 tasks, as follows:
  - Create a network with the appropriate number of input units and one output unit
  - Initialize the weights of the network randomly between -1 and +1
  - …

Task Definition

- Measuring performance, continued…
  - For a number of epochs (typically 10) train on all exemplars as follows:
    - Propagate the input values through the network, yielding output values
    - Adjust the network weights according to the formula specified by the genetically encoded learning algorithm
  - At the end of this process, measure fitness by testing the network on all training exemplars
    - Divide the total error by the number of exemplars, subtract from 1.0, and multiply by 100 to yield a fitness "percentage" between 0 and 100
    - A fitness percentage of 50 represents chance performance

Task Definition

- Note: Limited number of epochs of training (10) means learning is unlikely to be complete
  - Even the known optimal Widrow-Hoff delta rule is only 98% accurate
- Accordingly, the k_0 scale factor combines with the other terms to produce an overall learning rate which may be significant

GA Parameters

- Population size is 40
- Reproduction is linearly proportional to fitness
- Subject 80% of each new generation to crossover
  - Pairs of individuals swap a randomly selected substring of bits
- Retain the best individual
- Mutate all bits with probability of 0.01
- Repeat process for some number of generations, typically 1000

Results

- Best learning algorithms produced in 10 evolutionary runs

	k_0	k_1	k_2	k_3	k_4	k_5	k_6	k_7	k_8	k_9	k_{10}	Fitness
Widrow-Hoff delta rule	0.25	0	0	0	0	0	0	0	-4	4	0	89.6%
Widrow-Hoff delta rule	-2.00	0	0	0	0	0	0	0	2	-2	0	98.0%
	0.25	0	-1	-2	4	0	0	0	-2	4	-2	94.3%
	0.25	0	-1	-2	4	0	0	0	-2	4	-2	92.9%
	-0.25	0	0	1	-1	0	1	-1	4	-4	0	89.8%
Slight variation on delta rule	-1.00	0	0	-1	1	0	0	0	4	-4	0	97.6%
Slight variation on delta rule	4.00	0	0	1	-1	0	0	0	-2	2	0	98.3%
	-0.06	0	0	0	-2	-1	2	2	4	-4	2	79.2%
	-0.25	0	0	2	-1	0	-1	-1	2	-4	0	89.8%
	0.25	0	-1	-2	4	0	0	0	-2	4	-2	93.2%

Effect of Diversity on Evolution of Learning

- For the 10 runs just discussed
  - Average fitness on the 20 training classification tasks was 92.3%
  - Average fitness on the 10 classification tasks that were withheld from training was 91.9%
- This suggests the “environment” (of training tasks) was sufficiently diverse to evolve a general purpose learning algorithm, rather than one overly tailored to the specific training data
  - As we knew from the evolution of the delta rule
- But how much diversity is required to obtain this generalization effect?

Effect of Diversity on Evolution of Learning

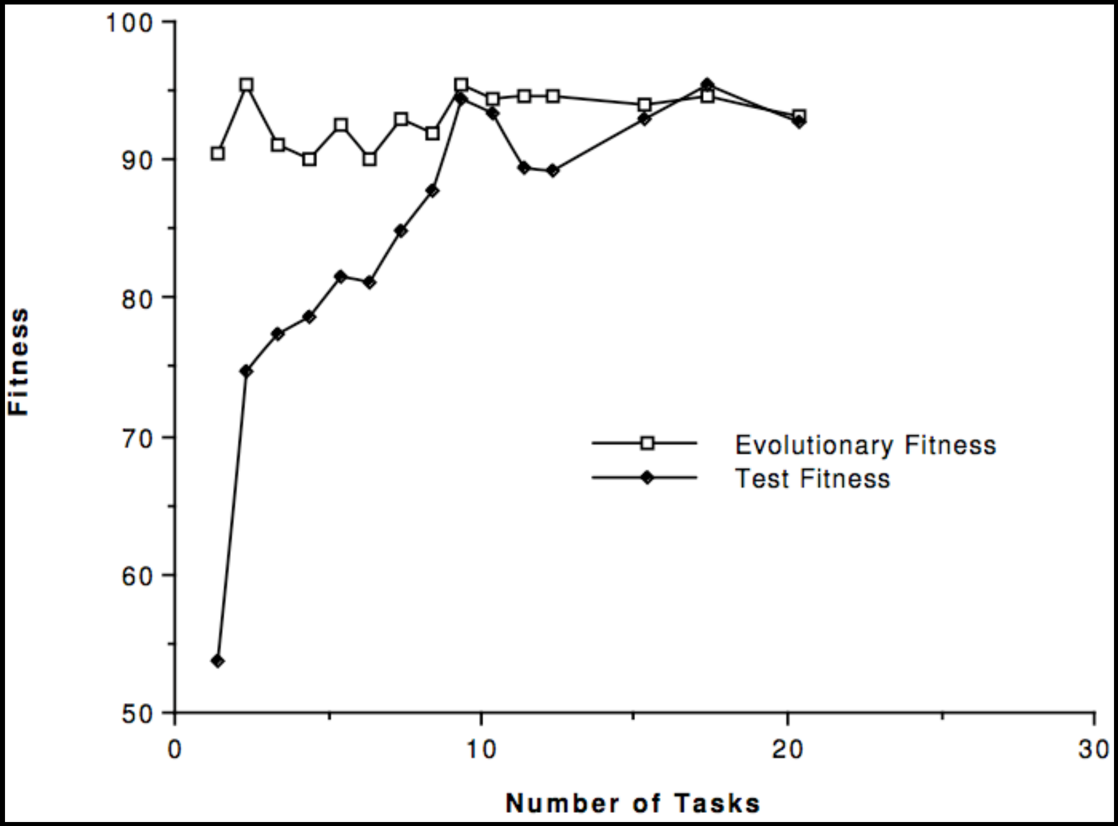

- "Evolutionary Fitness" is calculated from the variable number of training tasks, while "Test Fitness" is always calculated for 10 randomly chosen tasks not used for training

Credits

- Lecture based largely on the three reading assignments

G. E. Hinton and S. J. Nowlan. How learning can guide evolution. Complex Systems, 1:495-502, 1987

Parisi, D., S. Nolfi, and F. Cecconi, "Learning, Behavior, and Evolution", Tech. Rep. PCIA-91-14, Dept. of Cognitive Processes and Artificial Intelligence, Institute of Psychology, C.N.R., Rome, June 1991. (Appeared in Proceedings of ECAL-91—First European Conference on Artificial Life, December 1991, Paris; also in Varela, F, Bourgine, P. Toward a practice of autonomous systems, MIT Press, 1991)

Chalmers, D., "The Evolution of Learning: An Experiment in Genetic Connectionism" in Connectionist Models, Proceedings of the 1990 Summer School, edited by D. S. Touretzky, J. L. Elman, T. J. Sejnowski, G. E. Hinton, Morgan Kaufmann, San Mateo, CA, 1991

References

- Links to classic Psychology papers online http://spartan.ac.brocku.ca/~lward/
- James Baldwin's "A New Factor in Evolution" http://members.aol.com/jorolat/baldwin2.html
- A reply to Paul Griffith's denigration of Baldwin's contribution http://philsci-archive.pitt.edu/archive/00001250/01/Griffiths-reply.htm
- Excellent article on genetic assimilation and other updates to Waddington theories http://www.bioone.org/bioone/?request=get-document&issn=0003-1569&volume=040&issue=05&page=0729