Beyond Competition in Evolution and Social Learning communities.

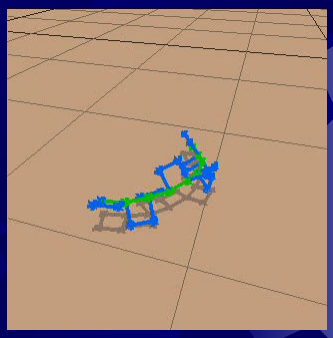

Jordan Pollack
Brandeis University
DEMO.CS.BRANDEIS.EDU

EUCognition January 2008

My Lab's Research

- How do biological systems progress without a designer?
- Can we understand open-ended evolution in *enough formal detail* to replicate the process using computer technology?
- What are beneficial applications?

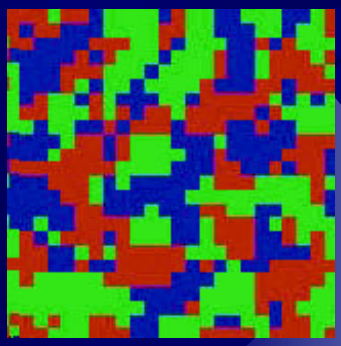

Demo.cs.brandeis.edu

Co-Evolution as Arms-race

- In Nature, Co-evolution means contingent development between species.
- But in Machine Learning the goal is an unstoppable "arms race" towards complexity.

Demo.cs.brandeis.edu

Canonical Coev. Algorithm

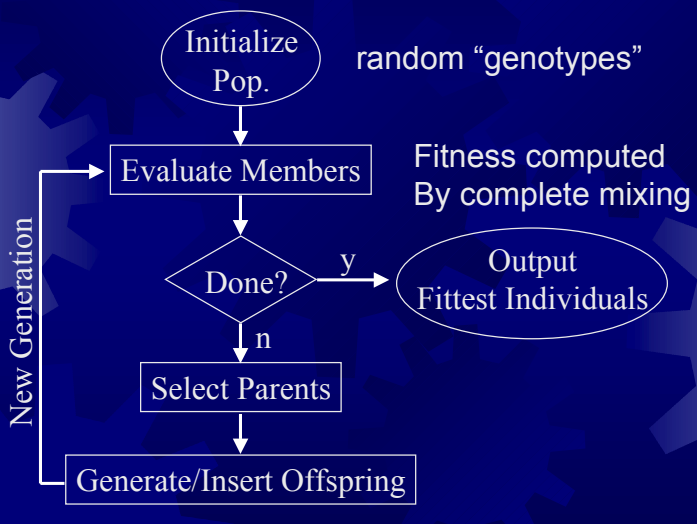

Demo.cs.brandeis.edu

Open-ended Evolution as Learning by Playing a game

Demo.cs.brandeis.edu

Co-evolutionary Success Stories

- Prisoner Dilemma (Axelrod & Forrest)
- Pursuer Evader (Cliff, Reynolds …)
- Sorting Nets (Hillis, *Juille*)
- Game Players (Rosin, *Angeline*, Tesauro, *Blair,* Fogel)
- Cell Automata Rules (Packard, *Juille*)
- Robotics (Sims, Floreano, *Funes, Lipson, Hornby*)
- Education Technology (*Sklar, Bader-Natal*)

(My students)

Demo.cs.brandeis.edu

Coevolutionary Robots

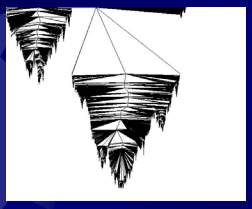

Evolutionary Computing

Virtual Reality Simulation

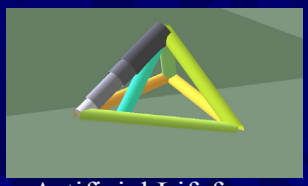

Artificial Lifeforms

Demo.cs.brandeis.edu

Can we turn Virtual Creatures Real?

3d Printing Machine
1999=$50,000
2009= $5000

CNC, CIM, CAM, RP, MEMS, NANO

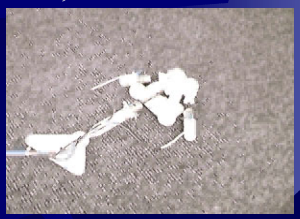

Demo.cs.brandeis.edu

GOLEM Generations

1) Lego Structures (Funes & Pollack, 98)
2) Active Trusses (Lipson & Pollack, 00)
3) L-System evolution (Hornby & Pollack, 02)
4) Modules for construction (03)
5) Robot Embryology (Reiffel, Viswanathan, 04-06)

Demo.cs.brandeis.edu

Five Generations
(Funes 98, Lipson 00, Hornby 02, Reiffel 05)

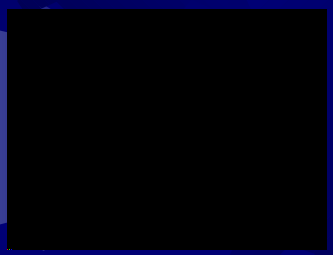
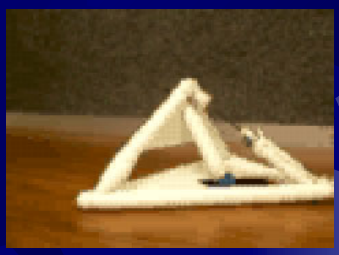


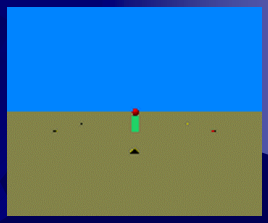

Demo.cs.brandeis.edu

The New York Times

THURSDAY, AUGUST 31, 2000

Scientists Report They Have Made Robot That Makes Its Own Robots

By KENNETH CHANG

For the first time, computer scientists have created a robot that designs and builds other robots, almost entirely without human help.

In the short run, this advance could lead to a new industry of inexpensive robots customized for specific tasks. In the long run — decades at least — robots may one day be truly regarded as "artificial life," able to reproduce and evolve, building improved versions of themselves.

Brandeis University

The "Arrow" left a trail as it crawled across a bed of sand.

Demo.cs.brandeis.edu

Robots reproducing themselves? Nobody wants that! Or do they?

Pollack & Blair HCGammon

- Replicated Tesauro's TD gammon with 1+1 hill-climbing co-evolution
  - Compare player to player+noise

Demo.cs.brandeis.edu

Q: Why did BKG Self-play work?
A: Backgammon Dynamics!

- Weak players CAN beat strong ones
  - Prevents WTA
- Weaknesses cannot be hidden
  - Prevents Memory Loss
- Predictive instability
  - enables all "phase" learning
- Lack of Draw (or throw)
  - prevents **collusion**

Demo.cs.brandeis.edu

Emergence of Mediocrity from Competitive Dynamics

- Sometimes it makes a little progress then discovers equilibrium
  - Winner-take-all (Monopoly)
  - Collusive Mediocrity (e.g. oligarchy, duopoly)
  - Disengagement (e.g. caste system)
  - Death-spirals
  - Memory Loss (Business cycles)
- Competition leads to Nash Equilibrium, progress halts

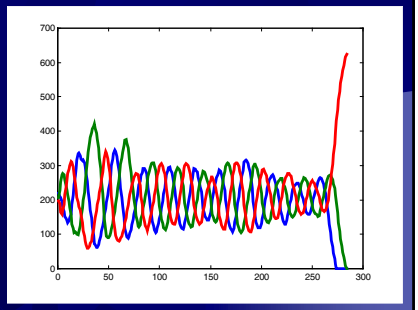

Demo.cs.brandeis.edu

Naïve Competitive Evaluation

$$f_i = \sum_{j=1..n} E(i,j)$$

$$E(i,j) = -E(j,i)$$

Fitness is zero sum
QED Nash exists

Demo.cs.brandeis.edu

Recent Theory *(my student)*

- Numbers Games (*Watson*)
- Pareto Co-evolution (*Ficici*)
- Informativeness (*Dejong*)
- Emergent Geometric Organization (*Bucci*)
- The Teacher's Dilemma (*Bader-Natal*)

Demo.cs.brandeis.edu

Numbers Game

[Watson & Pollack, 2001]

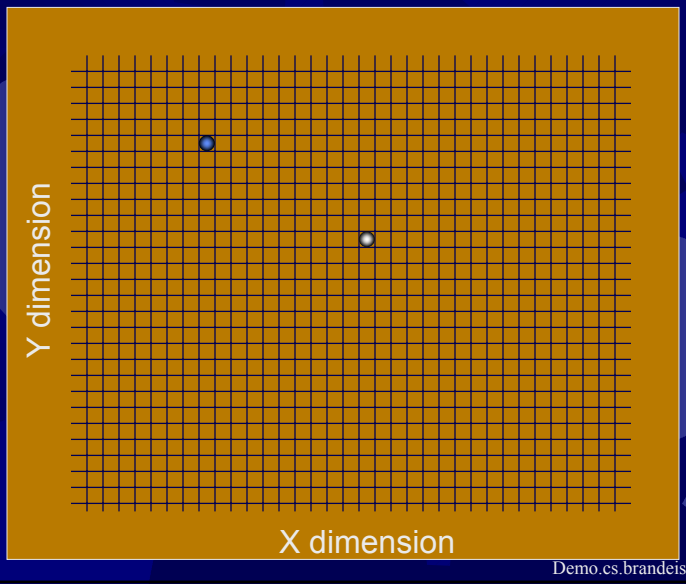

Demo.cs.brandeis.edu

Intransitive Superiority Cycles

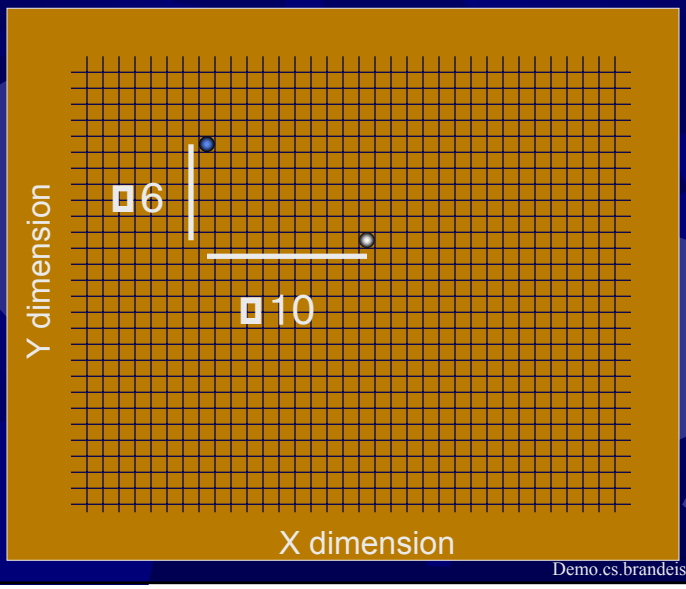

Demo.cs.brandeis.edu

Which is greater in closer dimension?

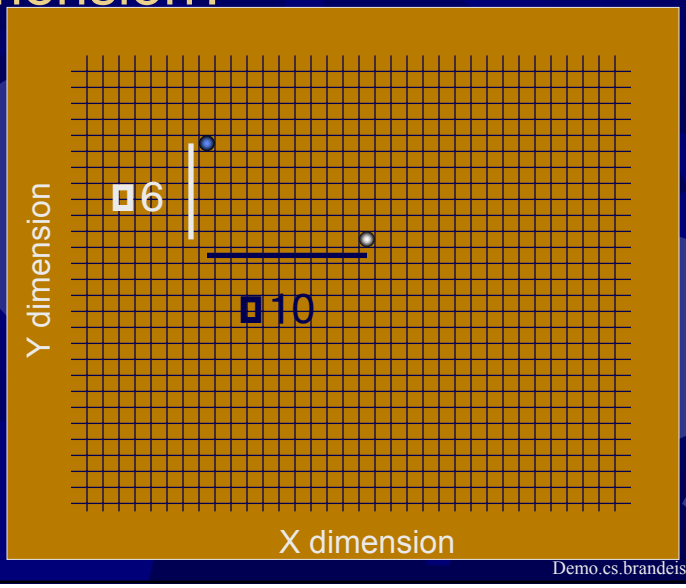

Demo.cs.brandeis.edu

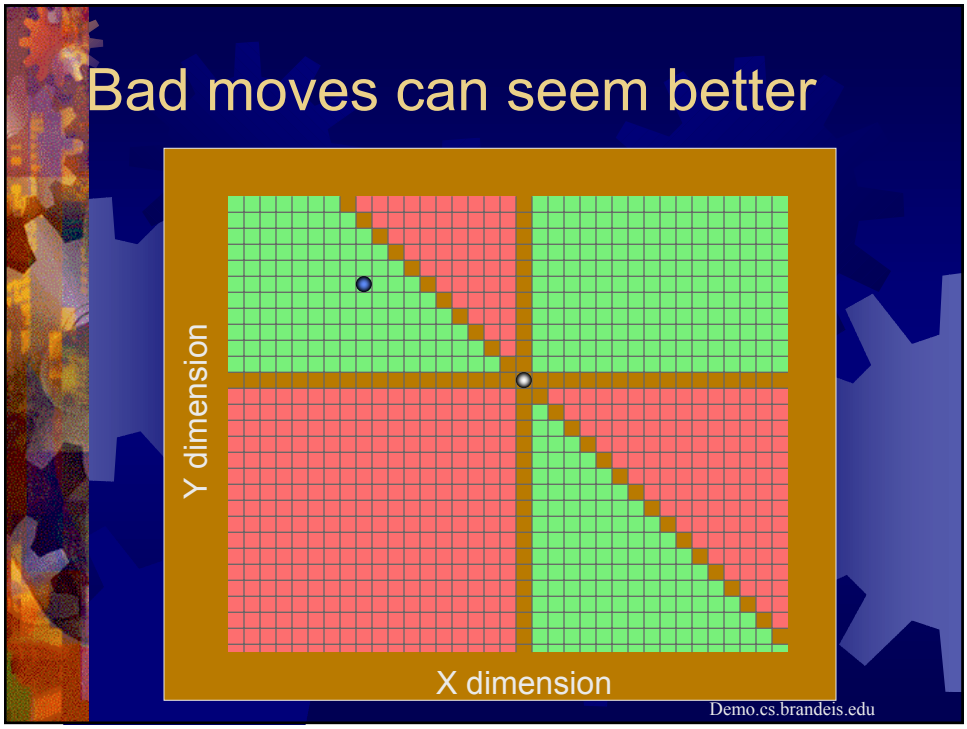
Bad moves can seem better
Y dimension
X dimension
Demo.cs.brandeis.edu

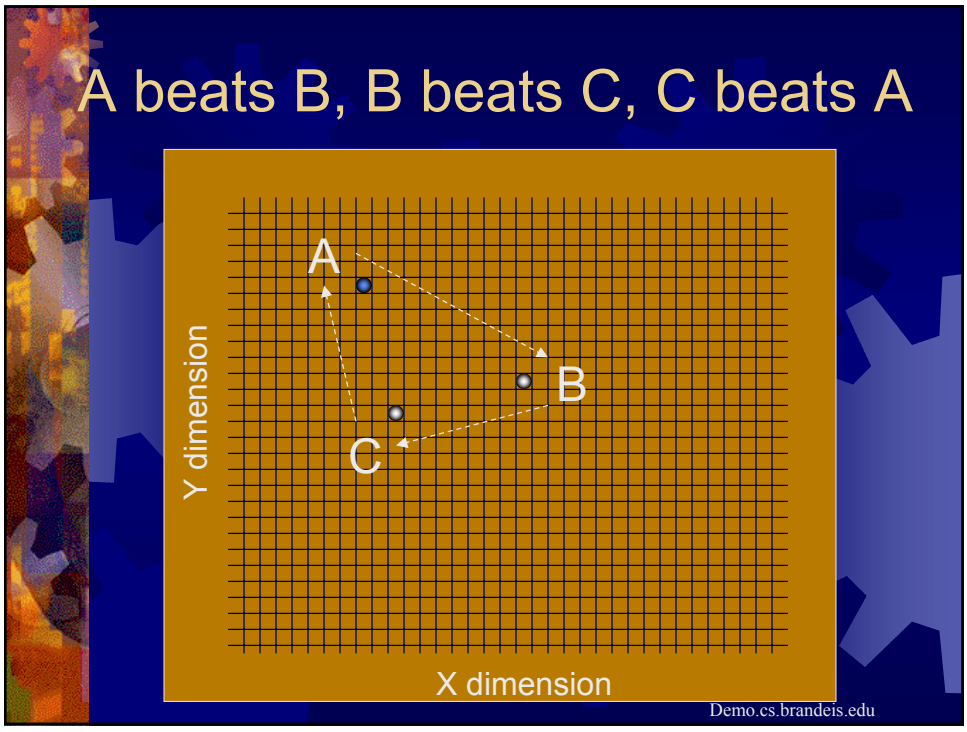
A beats B, B beats C, C beats A
A
B
C
Y dimension
X dimension
Demo.cs.brandeis.edu

Simple Simulations

- One or two populations
- All-vs-all
- Proportional reproduction
- Gaussian Mutation with negative bias
  - Bias makes your offspring less fit.

0 bias

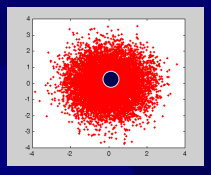

1.0 bias

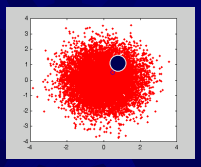

Demo.cs.brandeis.edu

Simple Matlab model

```
function z=runcoev Ppopsize,gens,bias)
pop=zerosPpopsize,2)R z=zerosPgens,1)R
if narginS3 bias=0RendR
for i=1Qgens
   matrix=mingamevPpop,pop)R
   fitness=sumPmatrix,2)R
   relfit=fitness-minPfitness)R
   pop=popProulettePrelfit,popsize),Q)R
   pop=pop-bias+randnPsizePpop))R
   zPi)=sumPmeanPpop))R
   showpopPpop,i)R
 end
plotresultsPz)R
```

Demo.cs.brandeis.edu

Pareto Co-evolution (Watson, Ficici)

- Coevolution presents a multi-objective learning task for an agent
- Every other agent you interact with is a **dimension** for optimization
- Use Multi-Objective Optimization
  - E.g. Keep a Pareto front.

Demo.cs.brandeis.edu

Informativeness vs competitiveness

- Master beats all
- Midlings beats a subset of other players
- Loser loses to all

	1	1	1	1
0		1	1	1
0	0		0	1
0	0	1		1
0	0	0	0	

Demo.cs.brandeis.edu

Informativeness vs competitiveness

- No Information
- Little info
- More Info
- More info
- No Info

	1	1	1	1
0		1	1	1
0	0		0	1
0	0	1		1
0	0	0	0	

New Algorithms based on multiple objective of preserving both Competitiveness and informativeness (Delphi, EGO, etc.)

Demo.cs.brandeis.edu

Playing games=Competition?

- A Zero Sum Game
  - Whatever I win, you lose
  - Whatever you win, I lose
- Therefore: Minimax!
- The Western Mindset:
  - Spencer, Smith, Darwin, von Neumann, Schumpeter…
  - Survival of the Fittest → competitive exclusion
  - Competition creates wealth, progress, freedom…

Demo.cs.brandeis.edu

Beyond Competition: The Teacher's Dilemma

- Teachers and learners need different "goals" (I.e. their utility function)
  - Should teachers and students compete?
  - Should teachers and learners "cooperate"?
- Pervasive mediocrity in competitive and cooperative learning systems can now be fully explained.

Demo.cs.brandeis.edu

Model of Teacher "vs." Student

- We can view the interaction between peers as two "teacher/student" interactions:
  - Teacher chooses a problem, from easy to hard
  - Student tries to solve problem
- Both student and teacher receive a "payoff"

Demo.cs.brandeis.edu

Student Utility

	Easy	Hard
Right	Pass	Pass
Wrong	Fail	Fail

Demo.cs.brandeis.edu

Teacher Utility

	Easy	Hard
Right	Verify	Joy
Wrong	Remediate	Complain

Demo.cs.brandeis.edu

Collusive Mediocrity: Students and teachers "secretly share" Joint Utility

Dominating Strategy

	Easy	Hard
Right	Verify+Pass	Joy+Pass
Wrong	Remediate+ Fail	Complain+ Fail

What we want!

Demo.cs.brandeis.edu

Model "average" student

A student's likelihood to answer a problem of difficulty d is:

- S(d) =1/(1+e^ (d*10-5))

- A teacher should ask questions around S(d)=.5, student's "zone of proximal development" (Vygotsky) to maximize "Flow" (Czen.)

ZPD

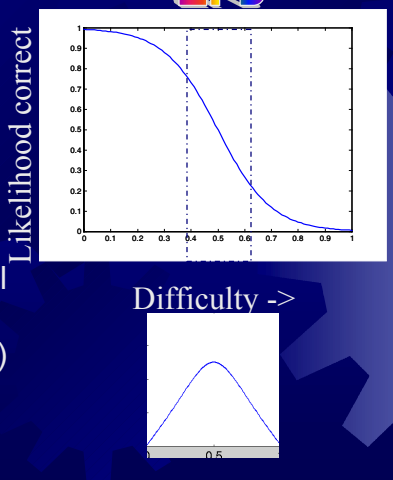

Demo.cs.brandeis.edu

Teacher Utility

Teacher chooses question with difficulty d

If student is right: (1-d)V + dJ
If student is wrong (1-d)R +dC

What are the parameters so teacher's max Payoff is near S(d)=.5?

Demo.cs.brandeis.edu

Usual Heuristics

- Our first assumption was that good teachers put a premium on "joy", but that verification was higher than remediation

3	5
1	0

Joy > Verify
Complaint is worst

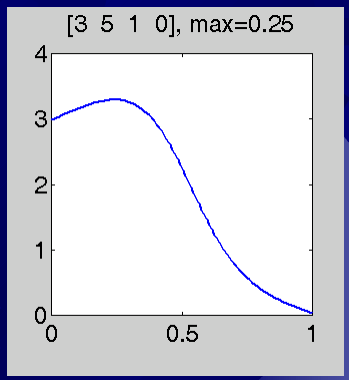

Grade inflation:
Teacher will ask easy questions
To avoid risk of complaints

Demo.cs.brandeis.edu

Apathy arises without complaint

Teachers were too sensitive to complaint?

3	5
1	5

No student evaluations!

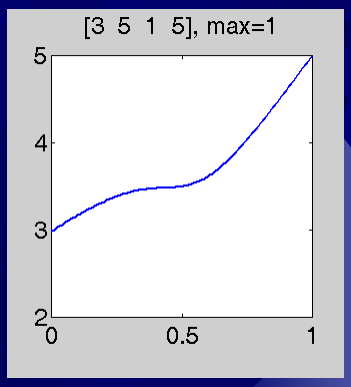

Apathy: Teacher asks hard questions
Doesn't care if student fails

Demo.cs.brandeis.edu

Teacher Student Cooperation

What if teacher is paid For student passing?

1	1
0	0

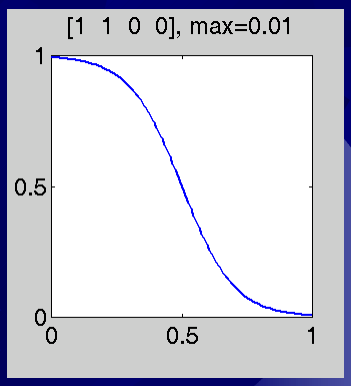

Teacher will ask easy questions.

Demo.cs.brandeis.edu

Teacher Student Competition

What if Zero Sum game?

0	0
1	1

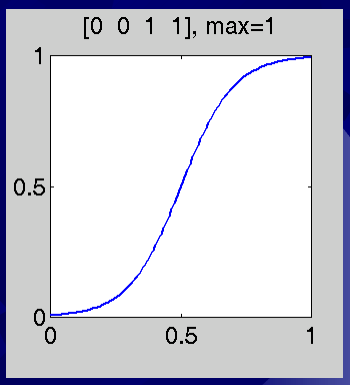

Teacher will ask Hard questions.

Demo.cs.brandeis.edu

We discovered how to pay teachers!

0	1
1	0

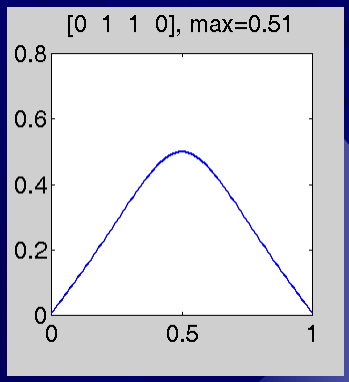

Pay for easy questions the student gets wrong
Pay for hard questions the student gets right

Demo.cs.brandeis.edu

New application of Co-evolution

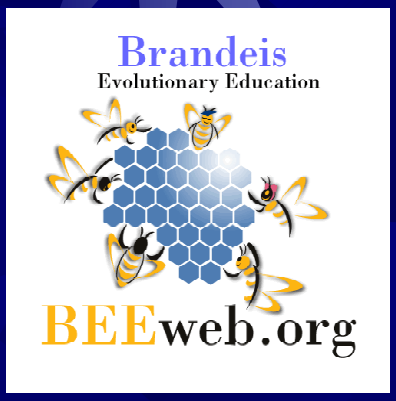

Peer-to-Peer Scaleable Educational Technology

What is BEEweb.org?

- Massively scaleable Internet-based educational technology based on scientific results using a new incentive structure which turns learners into each other's teachers.
- Solves the 2000 year old problem of finding a human teacher for every learner.

Demo.cs.brandeis.edu

Beeweb.org

- Students connect across the internet
  - No chat, no identities.
- To play a multiplayer video game
  - Where they drive each other into ZPD

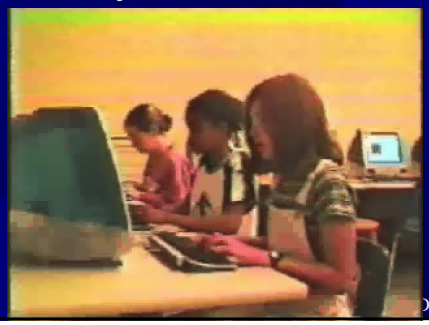

Demo.cs.brandeis.edu

Tutoring Paradigm

Rewarded for finding Student's Zone

Rewarded for right answers

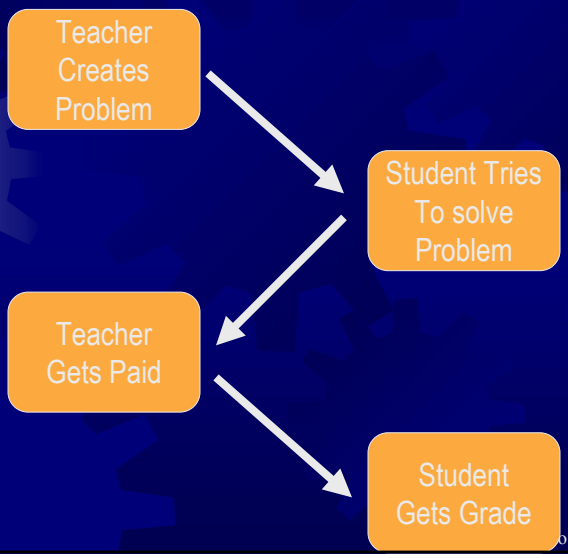

o.cs.brandeis.edu

Reciprocal Peer Tutoring

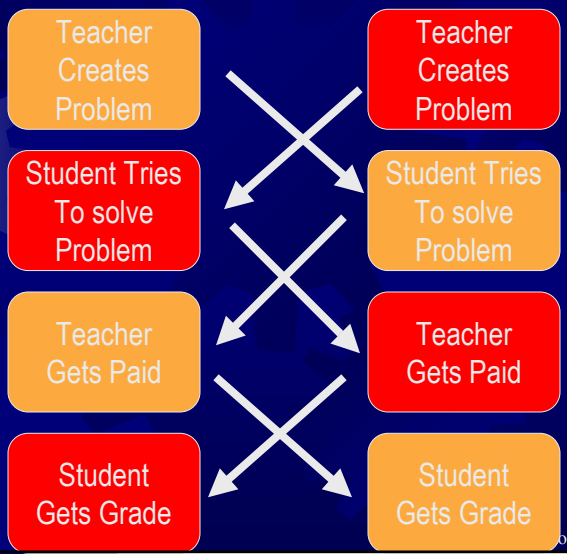

o.cs.brandeis.edu

Basic Educational Activities which "fit" paradigm

- Literacy
- Puzzle creation and solution
- Memory for Sequences
- Algebraic Conservation
- Spatial Reasoning and Rotation
- US Geography
- Scientific Problem Solving Skills

Demo.cs.brandeis.edu

www.SpellBee.org

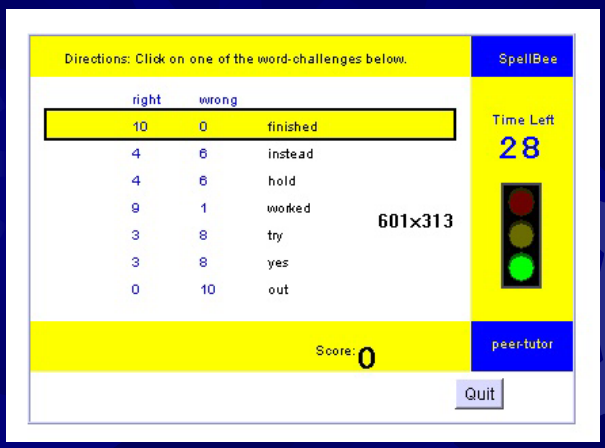

Each Student Picks a Word for the Other Demo.cs.brandeis.edu

SpellBEE.org

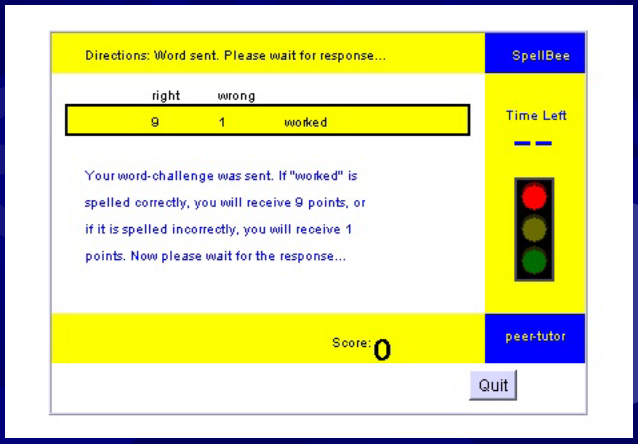

They get points based on (experimentally controlled) payoffs
Demo.cs.brandeis.edu

SpellBEE.org

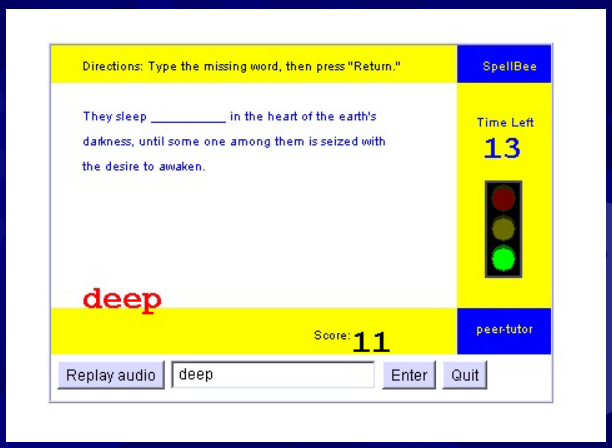

They are played audio sentence, with gap in text

Demo.cs.brandeis.edu

SpellBEE.org

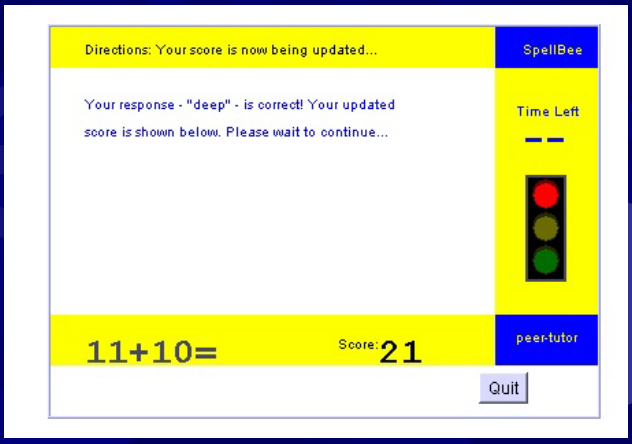

They get a score for correctness, and from success of other

Demo.cs.brandeis.edu

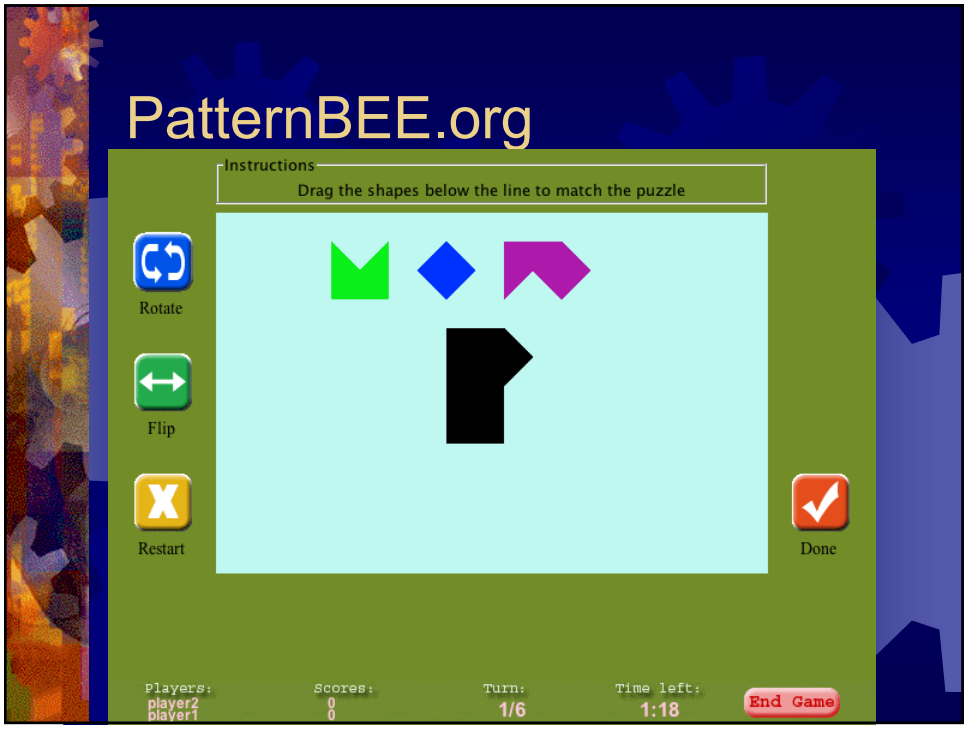
PatternBEE.org
Instructions
Drag the shapes below the line to match the puzzle
Rotate
Flip
Restart
Done
Players:
player2
player1
Scores:
0
0
Turn:
1/6
Time left:
1:18
End Game


MoneyBEE.org – pre-algebra
Click on the coins to create a challenge for your opponent.
Restart
0 + 10 + 10 + 25 = 45
Done
If your opponent is right
If your opponent is wrong
Players:
pollack
jordan
Scores:
0
0
Turn:
1/6
Time left:
0:00
End Game

GeograBEE.org – US States

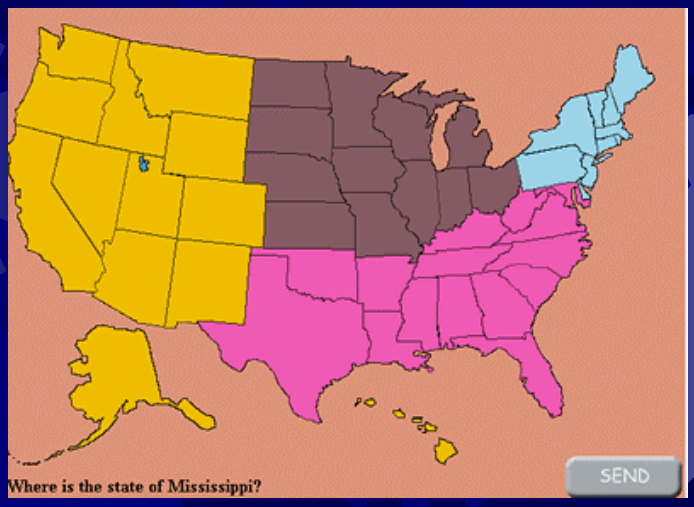

Demo.cs.brandeis.edu

Experimental results

- We can put Teacher Payoffs into the Server and test student behavior and learning.
- Protocol 6: Cooperate
- Protocol 7: Teacher Dilemma
- Protocol 8: Compete
- Protocol 9: Pre-biased TD

Demo.cs.brandeis.edu

Indiv. Student Curves

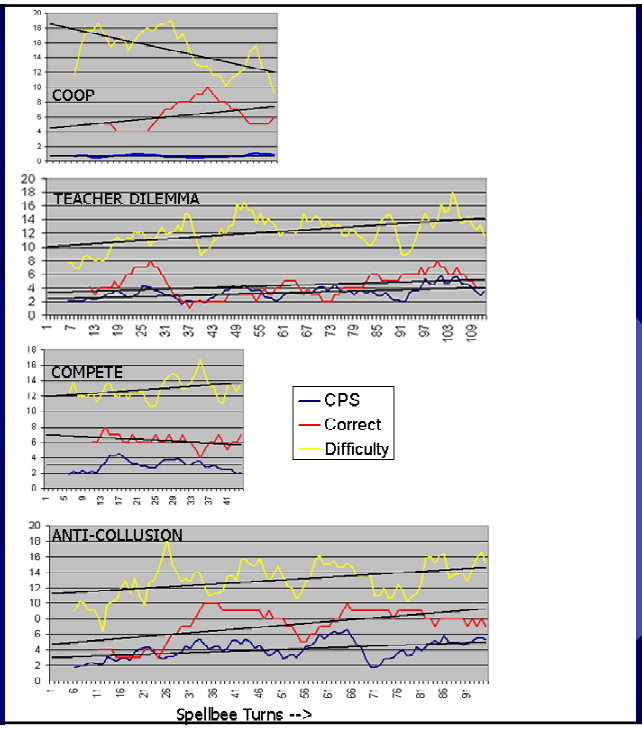

SpellBEE Visitors one month

Demo.cs.brandeis.edu

Dynamical and Evolutionary Machine Organization (DEMO)

- Postdocs
  - Hod Lipson, Edwin Dejong, Alan Blair
- Ph.D's
  - Hugues Juille, Betsy Sklar, Ofer Melnik, Pablo Funes, Sevan Ficici, Richard Watson, Greg Hornby, Simon Levy, Shiva Viswanathan, Anthony Bucci, Keki Burjoejee, John Rieffel, Ari Bader-Nadal
- Funding
  - ONR, DARPA, NSF, DOE, Hewlett

Demo.cs.brandeis.edu

Conclusion

- Coevolutionary dynamics are more complex than originally suspected.
- Competition alone is not enough to generate "open-ended" innovation.
- Applications of basic research have value, e.g. in education and Robotics.

Demo.cs.brandeis.edu