COGS 300 Learning 01 March 17/26 ①

Warm up: Faces. Draw a piece of paper.

Genetic programming and evolutionary algorithms

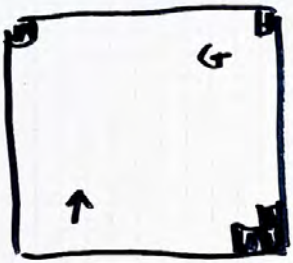

↑ ↓ → ←
N S E W

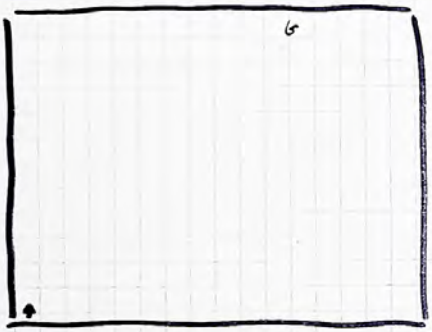

Ant's view

cost
energy

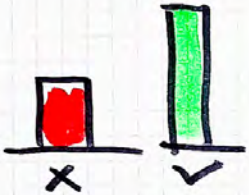

1 million runs of the ant

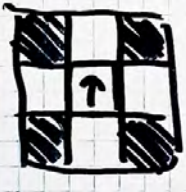

rand(NSEW)

weighted random (N S E W)
$\frac{1}{4}$ $\frac{1}{4}$ $\frac{1}{4}$ $\frac{1}{4}$

$\frac{5}{10}$ $\frac{1}{20}$ $\frac{4}{10}$ $\frac{1}{20}$

cost
↳ memory
↳ movement

→ reward
→ punishment
→ evaluate an action

Design a system to tune weights and pass on "traits" to the next gen. of ants.

cost per ant

reward ∝ cost

transmission

objective function

system cost (1 million?)

what is a gene?

vs. trait?

vs. application?

1. Discretize your space.

N S E W

0-1.0 0-1.0 0-1.0 0-1.0

0.1 for each step.

$\frac{1}{10}$

N	S	E	W
0	0	0	1
0	0	$\frac{1}{10}$	$\frac{9}{10}$
0	0	$\frac{2}{10}$	$\frac{8}{10}$
0	$\frac{1}{10}$	$\frac{1}{10}$	$\frac{8}{10}$
⋮			

$\sim \frac{10 \times 10 \times 10 \times 10}{1000\ \text{ms}}$

10K

2. simulate.

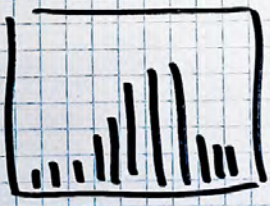

no of steps →

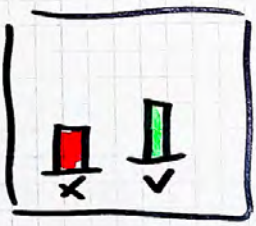

3. Eval. (call)

4. Generate

How to randomize the next gen?

Kidnap ants

lower cost times

100

90

switch

0.2	0.4	0.1	0.3
0.1	0.1	0.1	0.7

1. Re-discretize. → 1/100
2. Randomize

0.21 0.39 0.09 0.31

3. Sim. + eval + cull
4. generate.

Learning 01

warm up: Faces.

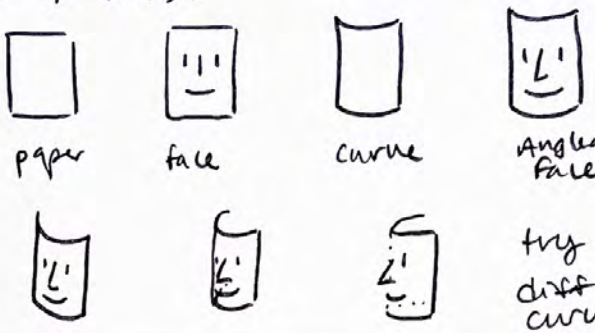

All of these faces are recognizable and even can be generated with rules, but designing an automaton to do so is ~~very~~ hard.

old face recognizers:

→ geometric transform.

problem? can't do a search – too large.

GOFAI basically failed. vs. ML.

If platonic ideals "exist" in the brain, they don't exist like their expressed form.

★ Do ppl believe in platonic ideals?

↳ if so, what is the scientific core?

Myths:
- All of that which is calculable is scientific X
- All of that which exists is scientifically describable
- All of that which is knowable will be established by scientific processes... X?

Consider: it is easier for me to generate a circle than verify a circle.

Automata vs. Robots
Theory Practice.

Last time: swarms + iterated arts.

Hill → Goal
modify environment
memory outside of agent.

small cartoon but always think big:

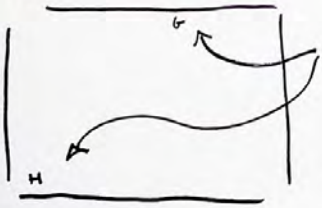

Let's try to move learning inside agent.

perspective is local!

max view of agent.

good model for theory.

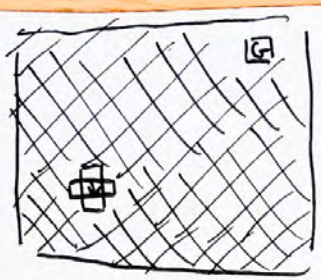

Totally "in the dark".

can't "keep a path" because it would get too long & one random deviation = total failure.

Randomness is a model for error.

rand (NSEW) =

Weighted_rand (NSEW; ¼ ¼ ¼ ¼)

can we learn weights to make it more likely to find a goal?

★ come up with a set of weights & example setup that would be more likely to find goal.

★ come up with a process to design the weights automatically.

↓ memory (small) ↓ reward ↓ punishment

eval.

Let's look at full path example:

1. Discretize. 0 – 1 in 0.1 steps.

N S E W
10 × 10 × 10 × 10 = ~10000 combos.

2. Simulate. Need a distribution.

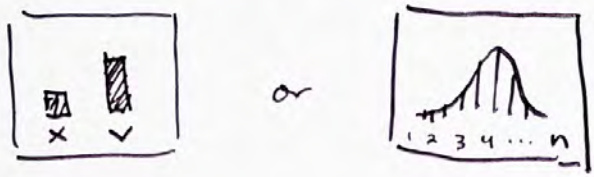

3. Evaluate.

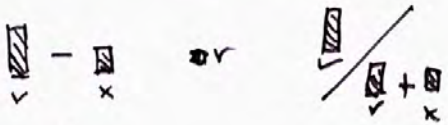

diff rate.

4. Generate

This is genetic programming 101. Evolutionary algs. take a seed + search for optimal settings.

1. step size 2. Dir.

↳ How to seed next gen?

This is population level learning.

Let's apply this to our robot.

Make a test rig:

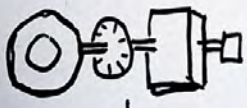

0-255. power

↳ speed

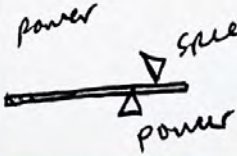

For a 1D problem, a search is fine.
But for even slightly more complex tasks, it becomes impossible.
Factors: weight, power, friction, etc... Sensing.
→ ex. for at home.
need state-dependent actions. (memory)
These are called policies.

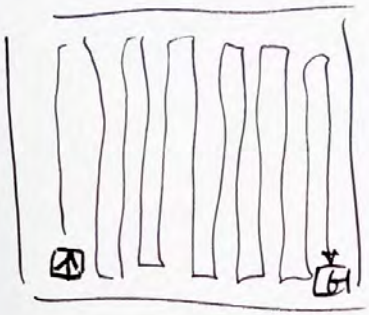

High cost.
Assign 1 unit of cost to both memory and path step.
★ Design alg. with optimized strategy.
★. Design alg - design process.