

Artificial Intelligence

Think about the traditional labor roles of humans vs those of machines. Machines do heavy lifting (bulldozers), fast travel (jets), and quick and accurate calculations (computers). Humans, in contrast, traditionally control the bulldozers, plan the calculations, interpret their results, communicate using human languages, and visually recognize objects and other people.

My perspective on AI is that it's about blurring those labor roles, so that computers can do more of the work traditionally associated with humans.

Planning as graph search

The first sub-field of AI we'll explore is **planning**. When a human has a goal, he or she then plans out what steps should be taken in order to achieve that goal. Then, perhaps, those steps could be programmed into a computer to solve the problem for us.

In AI, the human sets up the goal and the parameters (such as what are the potential actions and consequences), but then the algorithm determines the exact sequence of steps. The algorithm works by searching through the state space.

We covered a river-crossing puzzle:

- [Fox, goose, and beans river-crossing puzzle](#)

and then a water jug puzzle. But here is an extended example:

- [15-puzzle](#)

I demonstrated a program that solves this puzzle. My program's solution looks like the following. It starts by showing you how it is exploring the search space. Once it finds the goal, it retraces the steps that got it there.

```
Visited 412:087:635 as start state
Visited 412:687:035 via up from start
Visited 012:487:635 via down from start
Visited 412:807:635 via left from start
Visited 412:687:305 via left from 412:687:035
Visited 102:487:635 via left from 012:487:635
Visited 412:837:605 via up from 412:807:635
Visited 402:817:635 via down from 412:807:635
Visited 412:870:635 via left from 412:807:635
[...many more...]
```




Figure 1: A start and goal state for the 8-puzzle

Visited 123:456:780 via up from 123:450:786

GOAL!

1. left: 412:807:635
2. left: 412:870:635
3. up: 412:875:630
4. right: 412:875:603
5. right: 412:875:063
6. down: 412:075:863
7. left: 412:705:863
8. left: 412:750:863
9. up: 412:753:860
10. right: 412:753:806
11. right: 412:753:086
12. down: 412:053:786
13. down: 012:453:786
14. left: 102:453:786
15. left: 120:453:786
16. up: 123:450:786
17. up: 123:456:780

Visited 19408 nodes.

In this case, we found the solution after exploring 19,408 states, but the path from the initial state to the goal is just 17 steps.

The program above uses **breadth-first search** but a more sophisticated approach is called the **A* search** (pronounced A-star). Here is an excellent explanation with [interactive demonstrations of A* search](#).

Adversarial search

This is a similar technique, but when you're playing a game with an adversary, you're not the only one making moves. So when searching for a path to a winning state, you have to take into account the opponent's moves. The best way to do that is the [Minimax](#) algorithm.

- An example [game tree](#) for
- My [Connect4 program, for the Mac](#)
- My [Connect4 program, for Windows](#)

Machine learning

TODO

- [Space shuttle data](#) – after opening, log in to Google and select **File » Make a copy**. Then you should be able to filter the records by field values using the drop-downs in the column header.
- [Entropy calculator](#)
- [Decision tree assistant](#)

Strong AI and the Turing test

- [Turing test](#)
- [Robot Telemarketer Insists It is Human in Creepy Phone Calls](#) (December 2013)
- [Why can't my computer understand me?](#) by Gary Marcus – includes a critique of Turing Test
- [Artificial intelligence is a tool, not a threat](#) by Rodney Brooks of Rethink Robotics

Robotics

- Kiva fulfillment center robots <http://www.youtube.com/watch?v=6KRjuuEVEZs> and <http://www.youtube.com/watch?v=1WsMdN7HMuA>
- *Humans Need Not Apply* by CGP Grey <https://www.youtube.com/watch?v=7Pq-S557XQU>

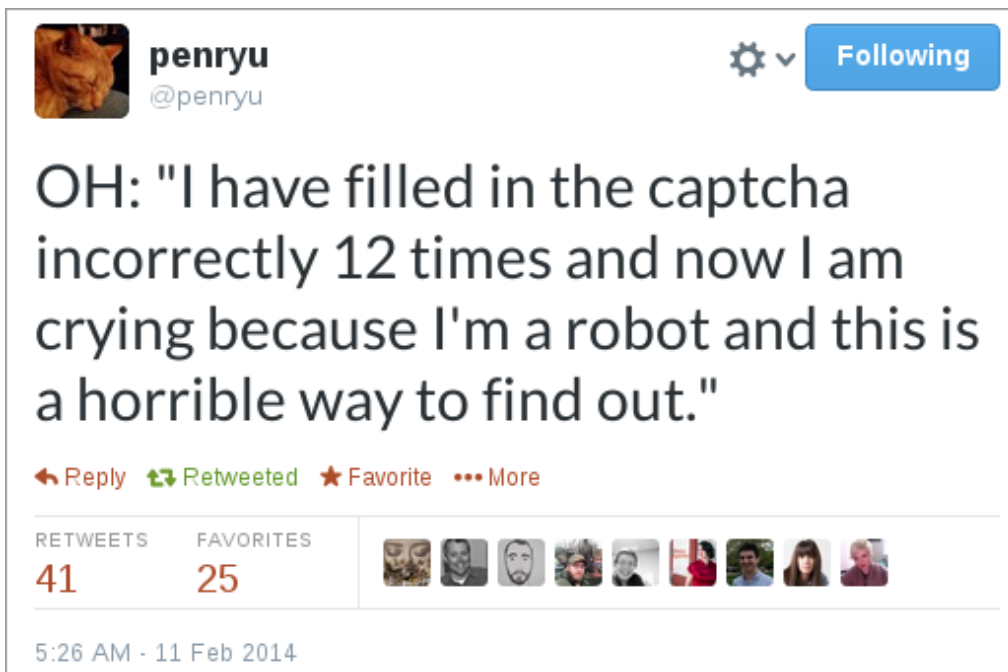


Figure 2: @penryu on Twitter