

To drive AI forward, teach computers to play old-school text adventure games

April 12 2018, by Jerry Swan, Hendrik Baier And Timothy Atkinson



Ready player one? Credit: Wikimedia

Games have long been used as test beds and benchmarks for artificial

intelligence, and there has been no shortage of achievements in recent months. Google DeepMind's AlphaGo and [poker bot Libratus](#) from Carnegie Mellon University have both beaten human experts at games that have traditionally been hard for AI – some 20 years after IBM's DeepBlue achieved the same feat [in chess](#).

Games like these have the attraction of clearly defined rules; they are relatively simple and cheap for AI researchers to work with, and they provide a variety of cognitive challenges at any desired level of difficulty. By inventing algorithms that play them well, researchers hope to gain insights into the mechanisms needed to function autonomously.

With the arrival of the latest techniques in AI and machine learning, attention is [now shifting](#) to visually detailed computer games – including the 3-D shooter Doom, [various 2-D Atari games](#) such as Pong and Space Invaders, and the real-time strategy [game](#) StarCraft.

This is all certainly progress, but a key part of the bigger AI picture is being overlooked. Research has prioritised games in which all the actions that can be performed are known in advance, be it moving a knight or firing a weapon. The computer is given all the options from the outset and the focus is on how well it chooses between them. The problem is that this disconnects AI research from the task of making computers genuinely autonomous.

Banana skins

Getting computers to determine which actions even exist in a given context presents conceptual and practical challenges which games researchers have barely attempted to resolve so far. The "monkey and bananas" problem is one example of a longstanding AI conundrum in which no recent progress has been made.

The problem was [originally posed](#) by John McCarthy, one of the founding fathers of AI, in 1963: there is a room containing a chair, a stick, a monkey and a bunch of bananas hanging on a ceiling hook. The task is for a computer to come up with a sequence of actions to enable the monkey to acquire the bananas.

McCarthy made a key distinction between two aspects of this task in terms of [artificial intelligence](#). Physical feasibility – determining whether a particular sequence of actions is physically realisable; and epistemic or knowledge-related feasibility – determining which possible actions for the monkey actually exist.

Determining what is physically feasible for the monkey is very easy for a computer if it is told all the possible actions in advance – "climb on chair", "wave stick" and so forth. A simple program that instructs the computer to go through all the possible sequences of actions one by one will quickly arrive at the best solution.

If the computer has to first determine which actions are even possible, however, it is a much tougher challenge. It raises questions about how we represent knowledge, the necessary and sufficient conditions of knowing something, and how we know when enough knowledge has been acquired. In highlighting these problems, McCarthy [said](#): "Our ultimate objective is to make programs that learn from their experience as effectively as humans do."

```

WEST OF HOUSE                               SCORE: 0/0
ZORK: THE GREAT UNDERGROUND EMPIRE -
PART I
COPYRIGHT (C) 1980 BY INFOCOM, INC. ALL
RIGHTS RESERVED.
ZORK IS A TRADEMARK OF INFOCOM, INC.
RELEASE 15 / SERIAL NUMBER UG3AU5

WEST OF HOUSE
YOU ARE STANDING IN AN OPEN FIELD WEST
OF A WHITE HOUSE, WITH A BOARDED FRONT
DOOR.
THERE IS A SMALL MAILBOX HERE.
>*
```

Until computers can tackle problems without any predetermined description of possible actions, this objective can't be achieved. It is unfortunate that AI researchers are neglecting this: not only are these problems harder and more interesting, they look like a prerequisite for making further meaningful progress in the field.

Text appeal

To operate autonomously in a complex environment, it is impossible to describe in advance how best to manipulate – or even characterise – the objects there. Teaching computers to get around these difficulties immediately leads to deep questions about learning from previous

experience.

Rather than focusing on games like Doom or StarCraft, where it is possible to avoid this problem, a more promising test for modern AI could be the humble text adventure from the 1970s and 1980s.

In the days before computers had sophisticated graphics capabilities, games like Colossal Cave and Zork were popular. Players were told about their environment by messages on the screen:

They had to respond with simple instructions, usually in the form of a verb or a verb plus a noun – "look", "take box" and so on. Part of the challenge was to work out which actions were possible and useful and to respond accordingly.

A good challenge for modern AI would be to take on the role of a player in such an adventure. The computer would have to make sense of the text descriptions on the screen and respond to them with actions, using some predictive mechanism to determine their likely effect.

More sophisticated behaviours on part of the [computer](#) would involve exploring the environment, defining goals, making goal-oriented action choices and solving the various intellectual challenges typically required to progress.

How well modern AI methods of the kind promoted by tech giants like IBM, Google, Facebook or Microsoft would fare in these text adventures is an open question – as is how much specialist human knowledge they would require for each new scenario.

To measure progress in this area, for the past two years we [have been running a competition](#) at the IEEE Conference on Computational Intelligence and Games, which [this year takes place in Maastricht](#) in the

Netherlands in August. Competitors submit entries in advance, and can use the AI technology of their choice to build programs that can play these games by making sense of a text description and outputting appropriate text commands in return.

In short, researchers need to reconsider their priorities if AI is to keep progressing. If unearthing the discipline's neglected roots turns out to be fruitful, the monkey may finally gets his bananas after all.

This article was originally published on [The Conversation](#). Read the [original article](#).

Provided by The Conversation

Citation: To drive AI forward, teach computers to play old-school text adventure games (2018, April 12) retrieved 2 July 2024 from <https://phys.org/news/2018-04-ai-old-school-text-adventure-games.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.