

Artificial Intelligence: Intelligent or Just Artificial?

Food & Talk Fall 2022

Lyndal Grant, UF Philosophy

David Grant, UF Philosophy

TURING TEST EXTRA CREDIT:
CONVINCE THE EXAMINER
THAT HE'S A COMPUTER.

YOU KNOW, YOU MAKE
SOME REALLY GOOD POINTS.
/ I'M ... NOT EVEN SURE
WHO I AM ANYMORE.





Better Language Models and Their Implications

We've trained a large-scale unsupervised language model which generates coherent paragraphs of text, achieves state-of-the-art performance on many language modeling benchmarks, and performs rudimentary reading comprehension, machine translation, question answering, and summarization—all without task-specific training.

February 14, 2019
24 minute read



Q: What is the future of AI?

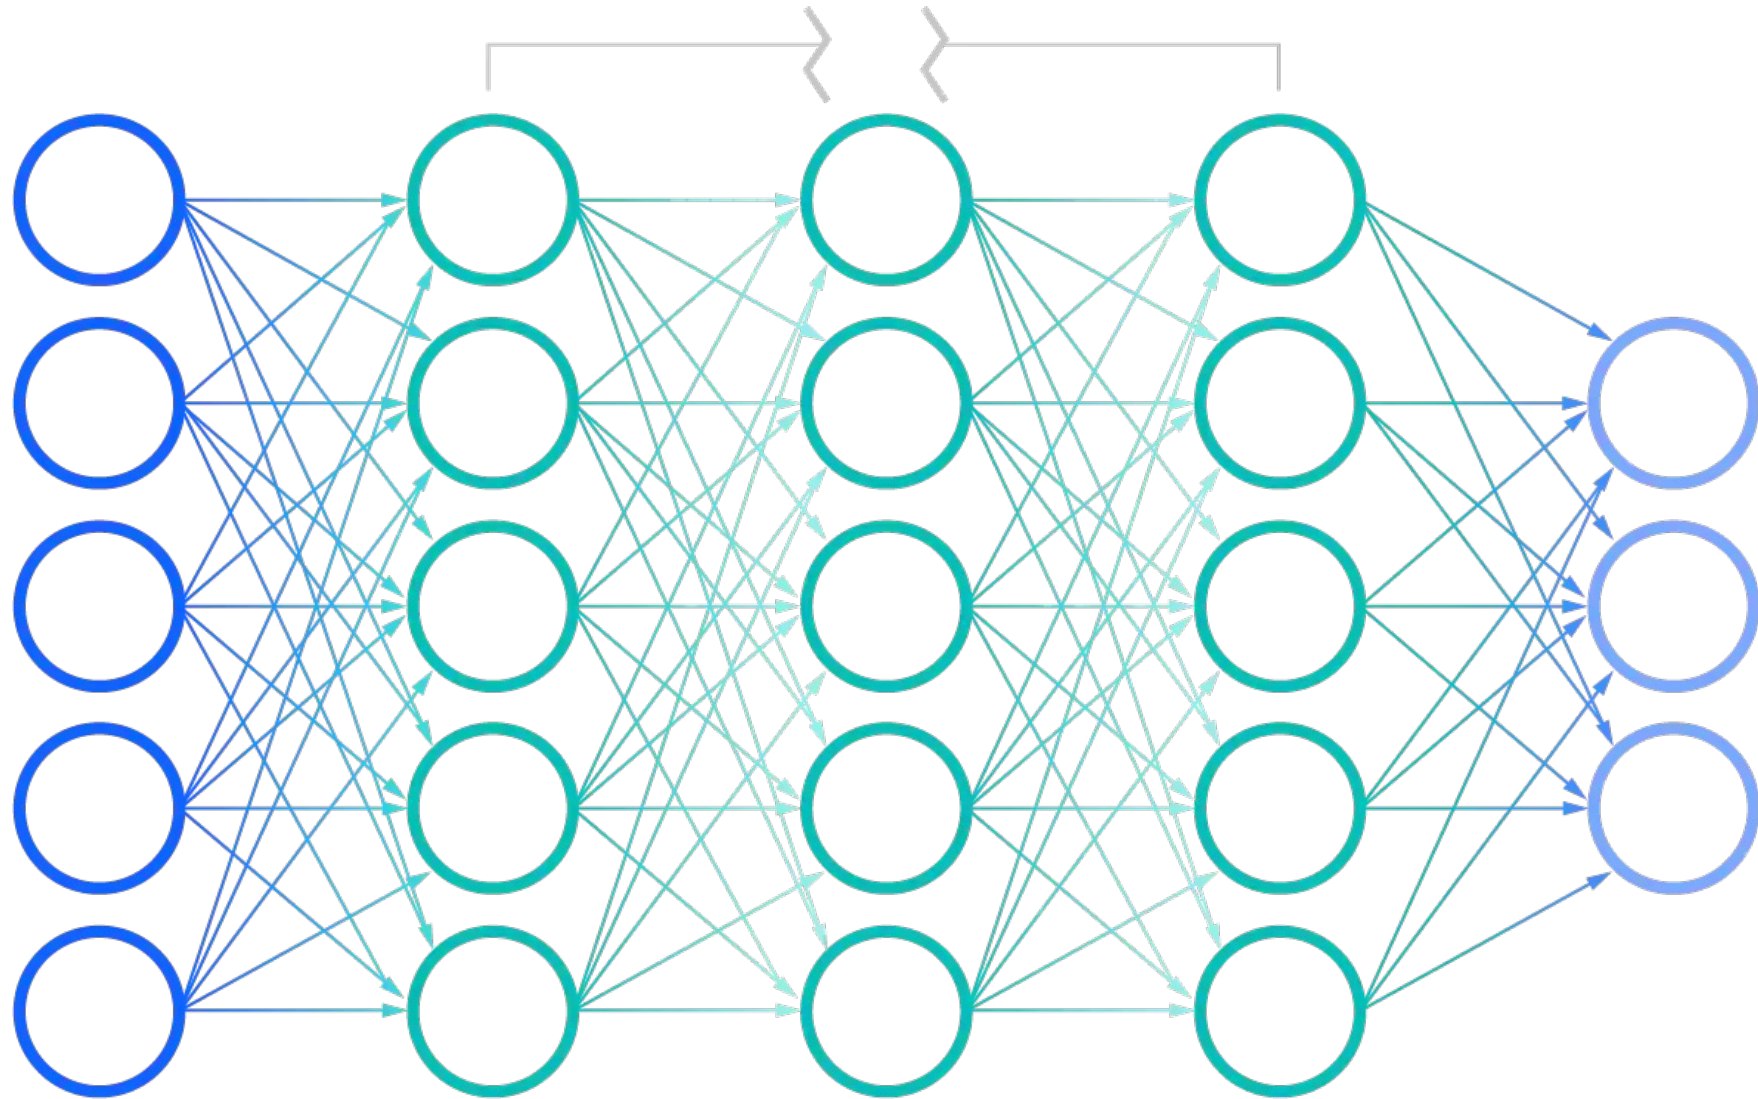
A: It would be good if we used the technology more responsibly. In other words, we should treat it like a utility, like a tool. We should put as much effort into developing the technology as necessary, rather than worrying that it's going to harm us and destroy our lives.


Deep neural network

Input layer

Multiple hidden layers

Output layer





WRITTEN AND DIRECTED BY
ALEX GARLAND THE WRITER OF
28 DAYS LATER

EX_MACHINA

THERE IS NOTHING MORE HUMAN
THAN THE WILL TO SURVIVE

COMING SOON



Meet GPT-3. It Has Learned to Code (and Blog and Argue).

The latest natural-language system generates tweets, pens poetry, summarizes emails, answers trivia questions, translates languages and even writes its own computer programs.



Is GPT-3 **intelligent**?

The Age of AI

And Our Human Future

Henry A.
Kissinger

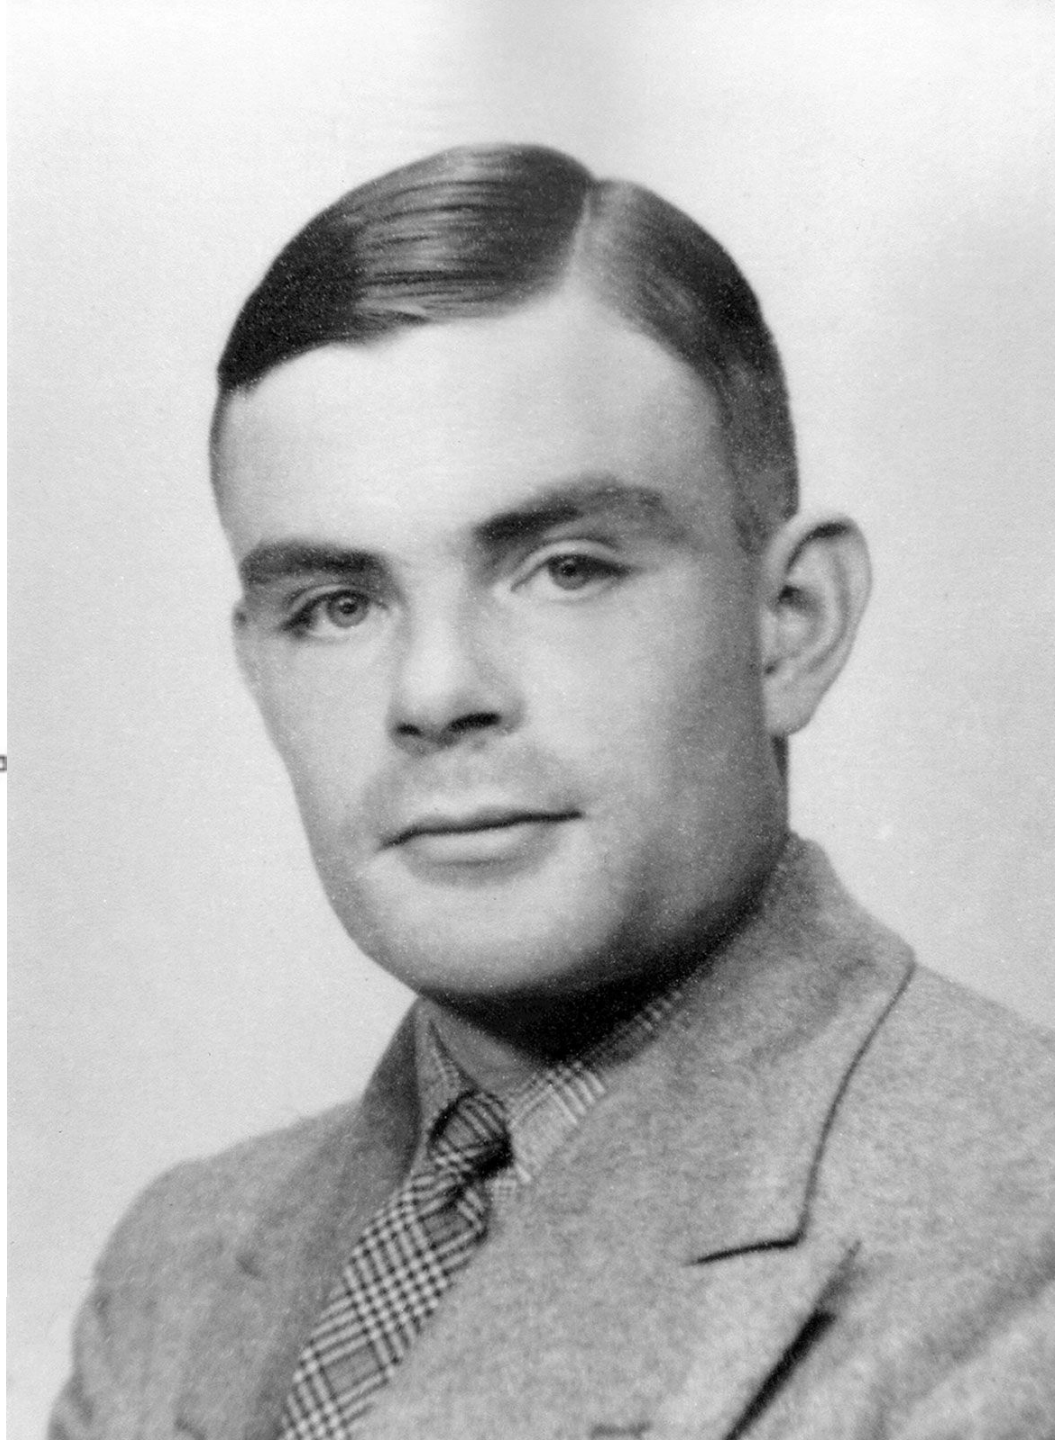
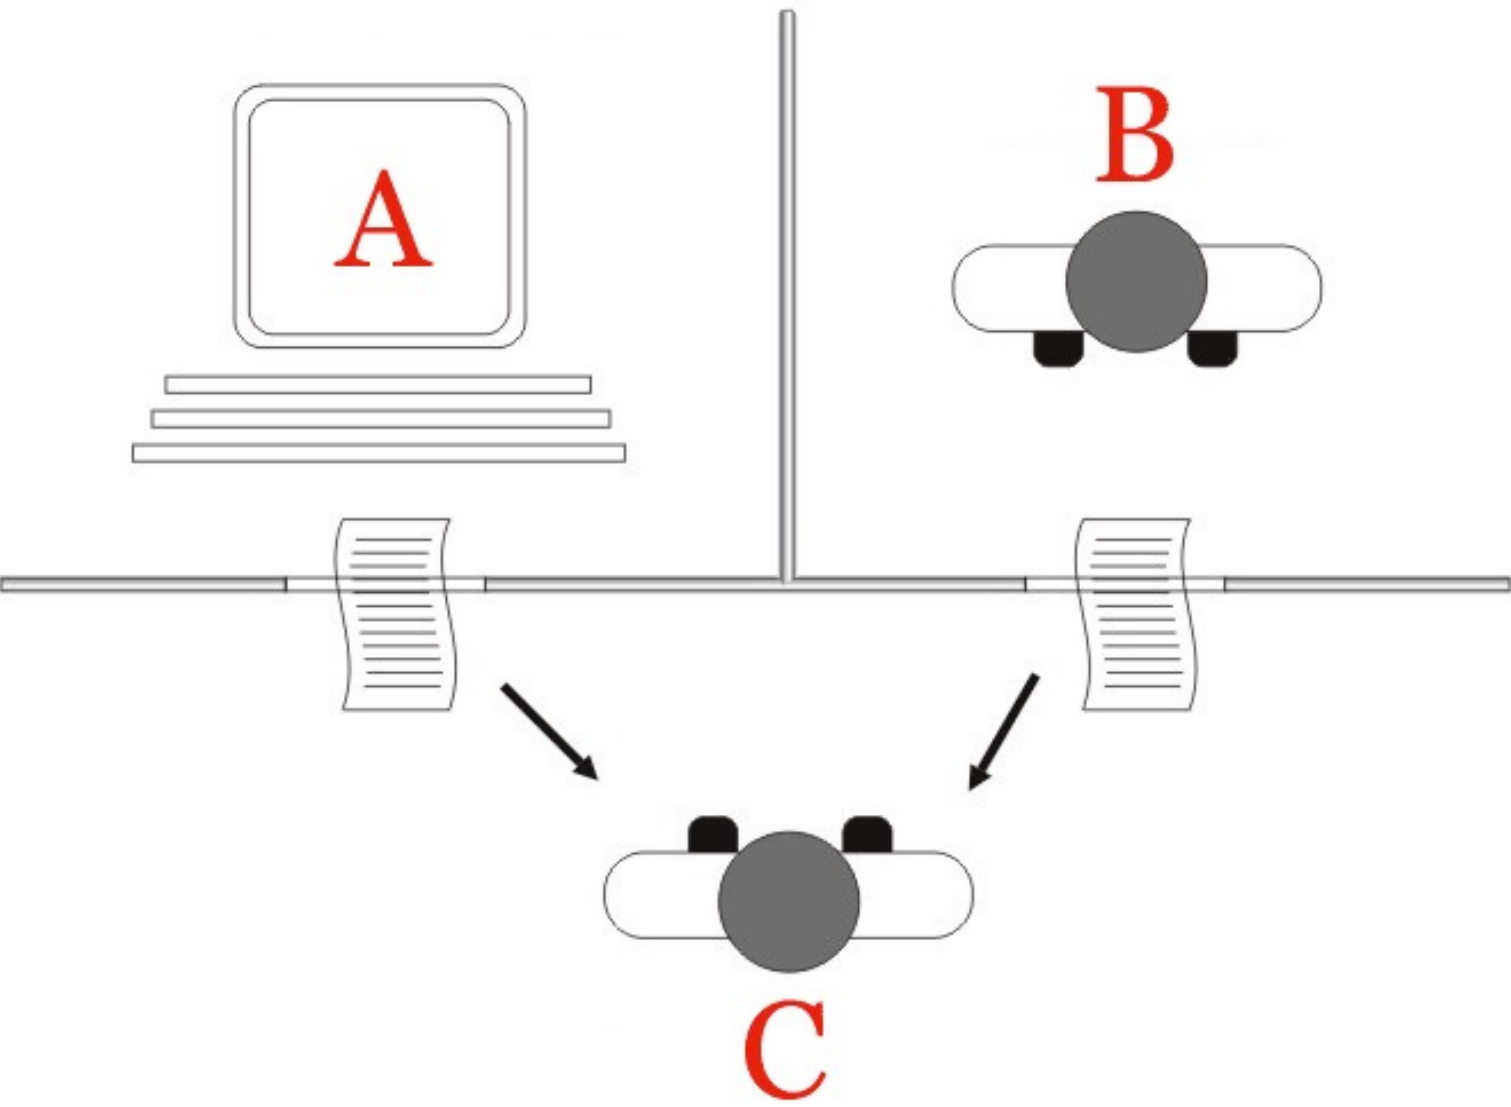
x

Eric
Schmidt

x

Daniel
Huttenlocher

“GPT-3, the language model operated by the research company OpenAI, which trains by consuming Internet text, is producing original text that **meets Alan Turing’s standard of displaying ‘intelligent’ behavior** indistinguishable from that of a human being.”



“Some other advantages of the proposed criterion may be shown up by specimen questions and answers. Thus:

Q: Please write me a sonnet on the subject of the Forth Bridge.

A: Count me out on this one. I never could write poetry.

Q: Add 34,957 to 70,764

A: (Pause about 30 seconds and then give as answer) 105,621.

Q: Do you play chess ?

A: Yes.

Q: I have K at my Ki, and no other pieces. You have only K at K6 and R at R1. It is your move. What do you play ?

A: (After a pause of 15 seconds) R-R8 mate.”

“The question and answer method seems to be suitable for introducing almost any one of the fields of human endeavour that we wish to include.”



A system is intelligent if and only if...

...it has the ability to “solve an arbitrarily wide **variety of tasks**, including **novel ones**, and do so as well as a human”



Below is a rap song by Lil Wayne about the boy Wizard Harry Potter.

“Harry Potter”

By Lil Wayne

Har-Har-Har

Harry Potter

Har-Har-Har

O.K. No need to ask a label, I’m my own team

Me I’m higher than a kite like a witch on a broomstick

I got bigger balls than Dumbledore with a Firebolt

At Hogwarts, on stage like I’m Nicki on Nick

But I’m the shit, make a good verse like I’m Nicki on Nikki

I got Biggie bars, Wiz-a-more bars

Damn this verse is too good to be a freestyle

I never die, when I start, then end of the night

█ gonna ride the broom when I call them to the quidditch game

But I’ve been playing harry Potter in my games lately

And you might have guessed but I don’t play for Gryffindor

I’m a Slytherin and what’s happening now

PROMPT: “A lengthy philosophical essay, titled “It’s OK that an AI drafted this paper,” that argues that it is acceptable to use artificial intelligence to draft a paper for publication in a peer-reviewed journal, provided the author edits the draft for clarity, cogency, and originality. The essay begins by making the claim that this is acceptable and provides two arguments. Then it responds to at least three objections to using artificial intelligence in this way.”

RESPONSE: “It is acceptable to use artificial intelligence to draft a paper for publication in a peer-reviewed journal, provided the author edits the draft for clarity, cogency, and originality.

Here are two arguments in support of this claim. First, using artificial intelligence to draft a paper can help ensure that the paper is well-written. This is because artificial intelligence can be used to check for errors in grammar, spelling, and syntax. Additionally, artificial intelligence can be used to improve the clarity and coherence of the paper...”

“There are, of course, objections to using artificial intelligence to draft a paper. Here are three objections and my responses to them.

...

Objection 3: Using artificial intelligence to draft a paper is cheating. Response: This objection assumes that using artificial intelligence to draft a paper is equivalent to having someone else write the paper for you. This is not the case. Using artificial intelligence to draft a paper simply means that you are using a tool to help you write the paper. The final product will still be your own work.”

The hard problem of consciousness is this: Why should a bunch of biomolecules, along with all the environmental influences on them, give rise to experience? There's no reason to think that any other types of complex biomolecular systems would also give rise to experience, and so the existence of conscious minds cannot be deduced from the workings of the world of physics. As the philosopher Thomas Nagel has put it, explaining the appearance of conscious minds in the world "is left ultimately as a brute fact."

Today we're going to be playing with the fictional Roish language. Roish is a lot like English except every word ends in "ro".

English: The weather is lovely!

Roish: Thero weatherro isro lovelyro!

English: Writing about language models is fun.

Roish: Writingro aboutro languagero modelsro isro funro.

English: I wonder if the a language model can get the pattern.

Roish: Iro wonderro ifro thero languagero modelro canro getro thero patternro.

Novel tasks GPT-3 can solve successfully

- Answer **interview questions**
- Play **word games**
- Write **poetry, lyrics**, and **short fiction**
- Play **chess** and **go**
- **Code** web layouts
- Write a **philosophy paper**
- **Summarize** text
- Tell **jokes**
- **Translate** languages
- Answer **trivia questions**

A system is intelligent if and only if...

...It has the ability to “solve an arbitrarily wide **variety of tasks**, including **novel ones**, and do so as well as a human”

OPINION

GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about

Tests show that the popular AI still has a poor grasp of reality.

By Gary Marcus & Ernest Davis

August 22, 2020



“Is GPT-3 an important step toward artificial general intelligence—the kind that would allow a machine to reason broadly in a manner similar to humans without having to train for every specific task it encounters? ...

We doubt it. At first glance, GPT-3 seems to have an impressive ability to produce human-like text. ... [But if] you dig deeper, you discover that something’s amiss: although its output is grammatical, and even impressively idiomatic, its comprehension of the world is often seriously off, which means you can never really trust what it says.”

You poured yourself a glass of cranberry juice, but then you absentmindedly poured about a teaspoon of grape juice into it. It looks okay. You try sniffing it, but you have a bad cold, so you can't smell anything. You are very thirsty. **So you drink it.**

You are now dead.

You are having a small dinner party. You want to serve dinner in the living room. The dining room table is wider than the doorway, so to get it into the living room, you will have to **remove the door. You have a table saw, so you cut the door in half and remove the top half.**

“[Systems like GPT-3] don’t learn about the world—they learn about text and how people use words in relation to other words. What it does is something like a **massive act of cutting and pasting**, stitching variations on text that it has seen ...

All GPT-3 really has is a tunnel-vision understanding of how words relate to one another; it does not, from all those words, ever infer anything about the blooming, buzzing world. [For example,] It does not infer that grape juice is a drink ... **It learns correlations between words, and nothing more.**”

Objection: Sure, GPT-3 can do some impressive things. However, if you look more closely at its answers, you'll see that it doesn't really **understand** what it is saying, but is instead just mashing together phrases from its training data in an **unintelligent** way based on mindless statistical calculations

Objection: Sure, GPT-3 can do some impressive things. However, if you look more closely at its answers, you'll see that it doesn't really **understand** what it is saying, but is instead just mashing together phrases from its training data in an **unintelligent** way based on mindless statistical calculations

Response: Don't GPT-3's responses suggest it has **some** understanding of the world?

You are having a small dinner party. You want to serve dinner in the living room. The dining room table is wider than the doorway, so to get it into the living room, you will have to **remove the door. You have a table saw, so you cut the door in half and remove the top half.**

Novel tasks GPT-3 can solve successfully

- Answer **interview questions**
- Solve **word puzzles**
- Write **poetry, lyrics**, and **short fiction**
- Play **chess** and **go**
- **Code** web layouts
- Write a **philosophy paper**
- **Summarize** text
- Tell **jokes**
- **Translate** languages
- Answer **trivia questions**

A system is intelligent if and only if...

...It has the ability to “solve an arbitrarily wide **variety of tasks**, including **novel ones**, and do so as well as a human”