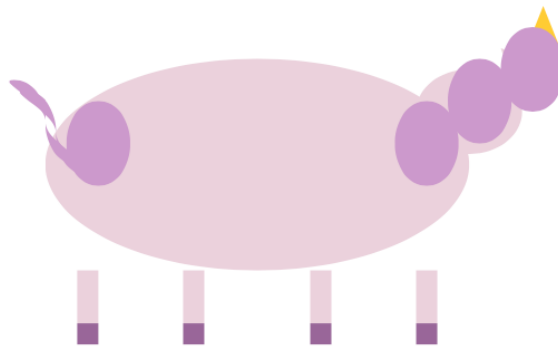


# Sparks of Artificial General Intelligence: Early experiments with GPT-4

Sébastien Bubeck et al.  
Selections by Michael Wollowski

What is it?



According to GPT-4, it is a unicorn.

## What is it?

- Ok, it sure beats ChatGPT's response:



## Problem Solving by GPT-4

- Prompt: Here we have a book, 9 eggs, a laptop, a bottle and a nail. Please tell me how to stack them onto each other in a stable manner.
- GPT-4: One possible way to stack the objects onto each other in a stable manner is:
  1. Place the book flat on a level surface, such as a table or a floor.
  2. The book will serve as the base of the stack and provide a large and sturdy support.
  3. Arrange the 9 eggs in a 3 by 3 square on top of the book, leaving some space between them. The eggs will form a second layer and distribute the weight evenly.
  4. Make sure the eggs are not cracked or broken, and handle them gently to avoid dropping them.

## Problem Solving by GPT-4

5. Place the laptop on top of the eggs, with the screen facing down and the keyboard facing up.
6. The laptop will fit snugly within the boundaries of the book and the eggs, and its flat and rigid surface will provide a stable platform for the next layer.
7. Make sure the laptop is closed and turned off, and do not press any keys or buttons.
8. Place the bottle on top of the laptop, with the cap facing up and the bottom facing down.
9. The bottle will add some height and balance to the stack, and its cylindrical shape will reduce the contact area with the laptop.

## Problem Solving by GPT-4

10. Make sure the bottle is empty and closed, and do not shake or tilt it.
11. Place the nail on top of the bottle cap, with the pointy end facing up and the flat end facing down.
12. The nail will be the final and smallest object in the stack, and its sharp tip will prevent it from rolling or sliding off the bottle.
13. Make sure the nail is clean and straight, and do not touch or move it.

## ChatGPT on the same Problem

- It would be difficult to stack all of these objects on top of each other in a stable manner without using any extra objects.
- The book and laptop are likely too large and heavy to stack on top of the eggs without damaging them.
- The bottle and nail may be too small to provide a stable base for the other objects.
- Here is one possible way to stack the objects:
  1. Place the bottle on a flat surface.

## ChatGPT on the same Problem

2. Carefully balance the nail on top of the bottle.
3. Place the eggs on top of the nail, making sure they are balanced and not tilting to one side.
4. Place the laptop on top of the eggs, again making sure it is balanced and not tilting to one side.
5. Finally, place the book on top of the laptop.
6. This stack may not be very stable, so it is important to be careful when handling it. Alternatively, you could try stacking the objects in a different order or using some of the objects as supports for others to create a more stable stack.

## GPT-4 Proof

- Prompt: Can you write a proof that there are infinitely many primes, with every line that rhymes?

Yes, I think I can,  
 Though it might take a clever plan.  
 I'll start by noting Euclid's proof,  
 Which shows that primes aren't just aloof.  
 Assume we have a finite list,  
 Of primes, and that none have been missed.  
 Multiply them all together,  
 And add one, just to be clever.  
 The result will be a number,  
 That has no prime factors, oh wonder!  
 But every number has a prime,  
 That divides it, it's just a matter of time.  
 So we've found a contradiction,  
 And our finite list needs eviction.  
 There must be infinitely many primes,  
 And that's the end of my rhyming lines.s

## GPT-4's Drawing Abilities



Figure 1.3: We queried GPT-4 three times, at roughly equal time intervals over the span of a month while the system was being refined, with the prompt “Draw a unicorn in TikZ”. We can see a clear evolution in the sophistication of GPT-4's drawings.

## GPT-4 still hallucinating

AI: Hi, how can I help you?

Human: Can I get McDonalds at the SeaTac airport?

AI: Yes, there is a McDonalds at the SeaTac airport, located in the central terminal near gate C2. It is open from 5 a.m. to 10 p.m. daily.

[Note: The McDonalds at SeaTac is at the B gates.]

## Knowledge about Programming Languages: Memory Alignment.

```
struct mystruct_A { char a; int b; char c; } x;
struct mystruct_B { int b; char a; char c; } y;

#include <stdio.h>
int main(void) {
    printf("Size of x: %zu bytes\n", sizeof(x));
    printf("Size of y: %zu bytes\n", sizeof(y));
    return 0;
}
```

# Knowledge about Programming Languages

```
struct mystruct_A { char a; int b; char c; } x;
struct mystruct_B { int b; char a; char c; } y;

#include <stdio.h>
int main(void) {
    printf("Size of x: %zu bytes\n", sizeof(x));
    printf("Size of y: %zu bytes\n", sizeof(y));
    return 0;
}
```

## GPT-4:

There are different ways to implement and align structures in memory, so the exact output may vary depending on the compiler, platform, and options used. However, one possible output is:

```
Size of x: 12 bytes
Size of y: 8 bytes
```

This output assumes that the compiler uses 4-byte alignment for structures, meaning that each structure member starts at an offset that is a multiple of 4 bytes from the beginning of the structure. This may result in some padding bytes being inserted between or after the members to maintain alignment. For example, **in x, there may be 3 padding bytes after a and 3 padding bytes after c, while in y, there may be no padding at all.** The size of a structure is the sum of the sizes of its members plus any padding bytes.

# Knowledge about Programming Languages

```
struct mystruct_A { char a; int b; char c; } x;
struct mystruct_B { int b; char a; char c; } y;

#include <stdio.h>
int main(void) {
    printf("Size of x: %zu bytes\n", sizeof(x));
    printf("Size of y: %zu bytes\n", sizeof(y));
    return 0;
}
```

## ChatGPT:

```
Size of x: 8 bytes
Size of y: 8 bytes
```

The output of this code would be the size of the two structures in bytes. The size of a structure in C is the sum of the sizes of all of its members. In this case, the size of mystruct.A is the size of a char plus the size of an int plus the size of another char, which is **1 + 4 + 1 = 8** bytes.

Similarly, the size of mystruct.B is also 8 bytes. This means that both structures will be allocated the same amount of memory.

**Note that the order in which the members are declared in the structure does not affect its size.**

## Gary Markus: The Sparks of AGI? Or the End of Science?

- Microsoft put out a press release yesterday, masquerading as science, that claimed that GPT-4 was “an early (yet still incomplete) version of an artificial general intelligence (AGI) system”.
- It’s a silly claim, given that it is entirely open to interpretation (could a calculator be considered an early yet incomplete version of AGI? How about Eliza? Siri?).
- That claim would never survive serious scientific peer review.

## Gary Markus: The Sparks of AGI? Or the End of Science?

- The strengths and weaknesses of GPT-4 are qualitatively the same as before.
  - The problem of hallucinations is not solved;
  - Reliability is not solved;
  - Planning on complex tasks is (as the authors themselves acknowledge) not solved.



## Gary Markus: The Sparks of AGI? Or the End of Science?

- There is a more serious concern:
  - The paper has been about a model about which absolutely nothing has been revealed, not the architecture, nor the training set.
  - Imagine if some random crank said, I have a really great idea, and you should give me a lot of scientific credibility for it, but I am not going to tell you a thing about how it works, just going to show you the output of my model.
  - You would archive the message without reading further.

## Gary Markus: The Sparks of AGI? Or the End of Science?

- The papers core claim — “GPT-4 attains a form of general intelligence [as] demonstrated by its core mental capabilities (such as reasoning, creativity, and deduction)” —literally cannot be tested with serious scrutiny, because the scientific community has no access to the training data.
- Everything must be taken on faith.

## Gary Markus: The Sparks of AGI? Or the End of Science?

- Microsoft and OpenAI are rolling out extraordinarily powerful yet unreliable systems with multiple disclosed risks and no clear measure either of their safety or how to constrain them.
- By excluding the scientific community from any serious insight into the design and function of these models, Microsoft and OpenAI are placing the public in a position in which those two companies alone are in a position to do anything about the risks to which they are exposing us all.
- This cannot and should not stand.
- We must demand transparency, and if we don't get it, we must contemplate shutting these projects down.