

# Putting ChatGPT to the Test: Practical Applications of Large Language Models for Assessment Development

Scott Frohn, PhD

Senior Manager of Assessment and Psychometrics

Workera.ai

## November, 2022

OpenAI releases ChatGPT (GPT 3.5)



## February, 2023

Meta releases LLaMA



## March, 2023

OpenAI releases GPT-4  
Google releases Bard (LaMDA)



## May, 2023

Google releases PaLM 2



## July, 2023

Anthropic releases Claude 2.0  
Meta & Microsoft release LLaMA 2



## August, 2023

Google announces Gemini (for December release)



## September 2023

### JUST THIS WEEK (Sept 18):

- Google releases Bard updates, including fact-checking of responses and integration with Google Apps
- OpenAI announces new version of Text-to-Image DALL-E 3 (for October Release)
- ChatGPT 4 info cutoff date is now January 2022

**The world of Generative AI and LLMs is moving incredibly fast...**



# The Power of LLMs: Healthcare



Isaac Kohane, MD, PhD

Source: Harvard Business School

**TODAY** TODAY all day

HEALTH & WELLNESS

## A boy saw 17 doctors over 3 years for chronic pain. ChatGPT found the diagnosis

Alex experienced pain that stopped him from playing with other children but doctors had no answers to why. His frustrated mom asked ChatGPT for help.

Sept. 11, 2023, 9:42 AM CDT / Updated Sept. 12, 2023, 9:31 AM CDT / Source: **TODAY**

By Meghan Holohan

During the COVID-19 lockdown, Courtney bought a bounce house for her two young children. Soon after, her son, Alex, then 4, began experiencing pain.

“(Our nanny) started telling me, ‘I have to give him Motrin every day, or he has these gigantic meltdowns,’” Courtney, who asked not to use her last name to protect her family’s privacy, tells TODAY.com. “If he had Motrin, he was totally fine.”

September 11, 2023 - Today.com



# The Power of LLMs: Productivity

## Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality

Fabrizio Dell'Acqua  
Harvard Business School

Edward McFowland III  
Harvard Business School

Ethan Mollick  
The Wharton School

Hila Lifshitz-Assaf  
Warwick Business School

Katherine C. Kellogg  
MIT Sloan School of Management

Saran Rajendran  
Boston Consulting Group

Lisa Krayer  
Boston Consulting Group

François Candelon  
Boston Consulting Group

Karim R. Lakhani  
Harvard Business School

### Working Paper 24-013

Copyright © 2023 by Fabrizio Dell'Acqua, Edward McFowland III, Ethan Mollick, Hila Lifshitz-Assaf, Katherine C. Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani.

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

## Randomized Controlled Trial for Management Consultants:

- Control // AI Access // AI Access + Prompt Training
- Tasked with (a) developing new product ideas, or (b) performing business problem-solving tasks.
- Baselines set w/ no AI access for all groups

## Results

- Compared to Control, AI Access groups:
  - did 12% more tasks
  - performed tasks 25% more quickly
  - had 40% higher quality output
- Lower performers in AI Access Group had greater gains (43%) vs. baseline, compared to higher performers (17%)
- AI access group did worse in tasks requiring data analysis and careful review of interviews.



# The Power of LLMs: Text Generation

## Which is easier?

1. Writing a paper from scratch, or
2. Editing a paper that's already written

## Which is easier?

1. Knowing what you want to say, or
2. Knowing how you want to say it



I want to write an introduction to a talk I'm giving about large language models and assessment development and I want to provide compelling examples for using LLMs because they're cool and I want to engage the audience and wow them and give them reasons why they should use LLMs like ChatGPT more often in their lives. What should I say? Be brief.



Certainly! Here's a concise and engaging introduction for your talk:

"Ladies and gentlemen, imagine a world where you can instantly summon the wisdom of millions of books, articles, and conversations at the tip of your fingers. Enter the realm of Large Language Models, like ChatGPT. Ever wanted to brainstorm with Shakespeare, ask Einstein for homework help, or co-write a novel with Jane Austen? With LLMs, these fantasies are inching closer to reality. Beyond the fun, they're revolutionizing industries, from streamlining customer support to reimagining assessment development. By the end of this talk, you'll discover not only the magic behind these digital wizards but also the countless ways they can enhance your everyday life. Ready to dive in?"



# The Power of LLMs: Item Development

Stem >>

Options:

Key >>

Distractors >>

>>

>>

>>

A fintech company has recently launched a new mobile banking app targeting young professionals. The product marketing team is tasked with creating branding guidelines to ensure consistency and effective communication of the product's value proposition.

Which of the following elements should be prioritized among these guidelines?

company logo, typography, and color palette



user personas, customer journey maps, and wireframes



market segmentation, competitive analysis, and pricing strategy



sales materials, promotional offers, and partnership opportunities



I don't know yet

## Item Writing Process

1. Identify the Concept
2. Background Research
3. Write Stem
4. Write Key
5. Write Distractors
6. Refine the item
7. Copy Edit for Style / Format
8. Identify Metadata



# The Power of LLMs: Item Development

## Item Development Process (traditional)



1 hour / item

100 hours for item writing

>> 50 usable items

100 IW hours / 50 usable items = **2 IW hours / usable item**

## Item Writing Process

1. Identify the Concept
2. Background Research
3. Write Stem
4. Write Key
5. Write Distractors
6. Refine the item
7. Copy Edit for Style / Format
8. Identify Metadata



# The Power of LLMs: Item Development

## Item Development Process (leveraging LLMs)



30 minutes / item

50 hours for item writing

>> 90 usable items

50 IW hours / 90 usable items = **33 IW minutes / usable item**

## Item Writing Process

1. Identify the Concept
2. Generate some items
3. Adjust prompt & regenerate
4. Select the best item(s)
5. Refine the item(s)
6. Copy Edit for Style / Format
7. Identify Metadata





# Roadmap

---

- **Some Background**
  - What is AI / Machine Learning / Deep Learning
  - How do LLMs work
  - ChatGPT v. Other LLMs
- Prompt Engineering
- LLMs in Assessment Development
- ChatGPT Demonstration
- Q&A



# Learning Objectives

---

1. Describe at a high-level how LLMs like ChatGPT work.
2. Employ some best practices in prompt engineering.
3. Identify ways to leverage LLMs in assessment development.



# **Some Background**



# Some Definitions

---



Cassie Kozyrkov, former Chief  
Decision Officer @ Google,  
CEO @ Data Scientific  
Source: Yahoo

## Artificial Intelligence

*"is the science and engineering of making intelligent machines, especially intelligent computer programs." - McCarthy, 2007*

## Machine Learning

*"uses algorithms and statistical models to allow computers to perform a task without explicit instructions, instead relying on patterns and inference from data." - ChatGPT*

*"How computers recognize patterns and make decisions without being explicitly programmed" - Code.org*

*"Machine learning is a thing-labeler, essentially." - Cassie Kozyrkov*



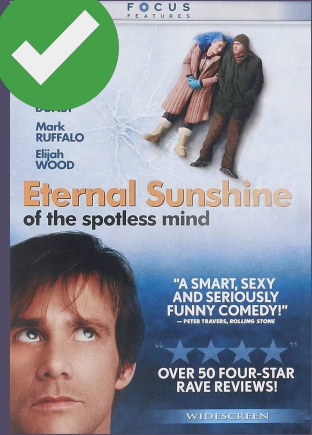
# Machine Learning

"...is a thing-labeler, essentially."

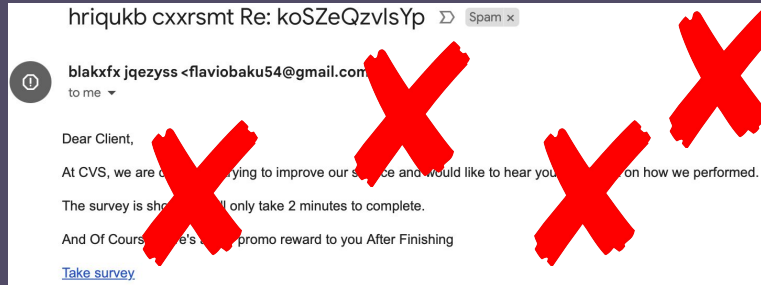
Is this a cat?



Will Scott Like This Movie?



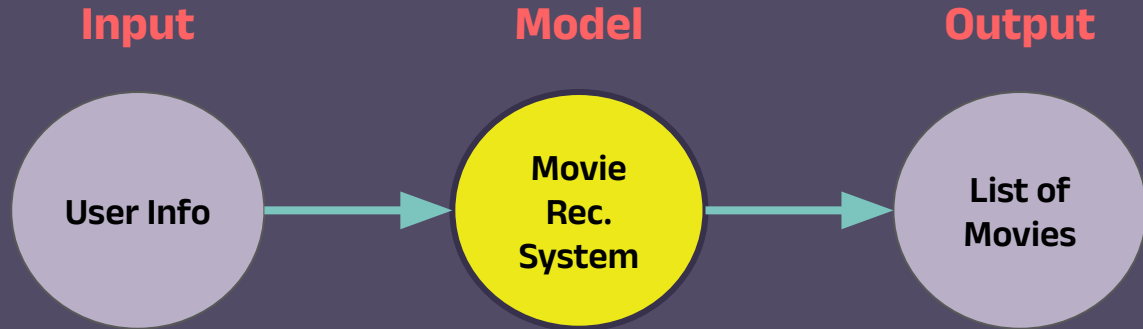
Is this email spam?



# Machine Learning

...an analogy

Machine Learning  
Model



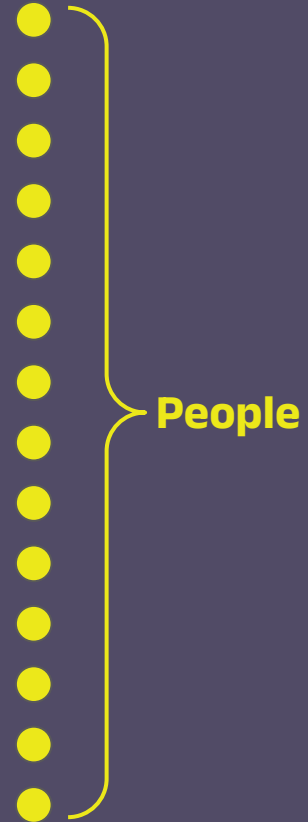
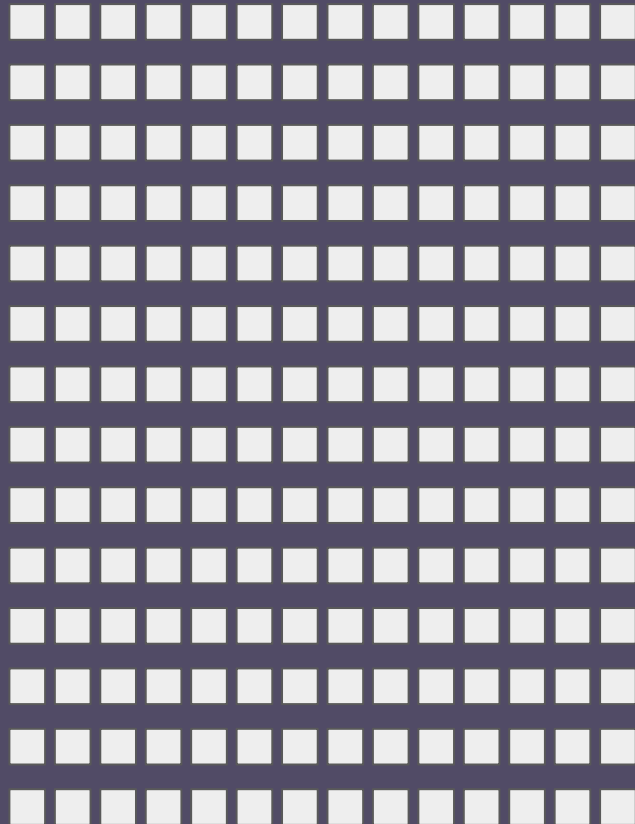
Psychometric  
Model



# Machine Learning

...an analogy

Items → 1 2 3 4 5 6 7 8 9 10 11 12 13 14



Psychometric  
Model  
"Training"



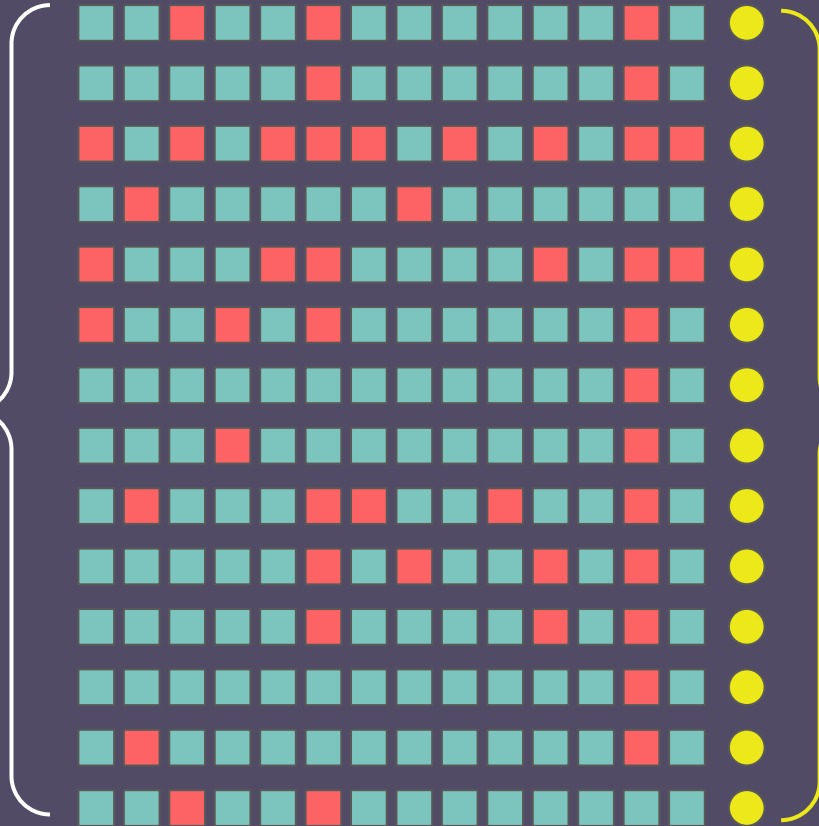
# Machine Learning

...an analogy

Items → 1 2 3 4 5 6 7 8 9 10 11 12 13 14

(Training) Data

Psychometric  
Model  
"Training"



People





# Machine Learning

...an analogy

Rasch Model,  
MLE

Psychometric Model  
"Training"

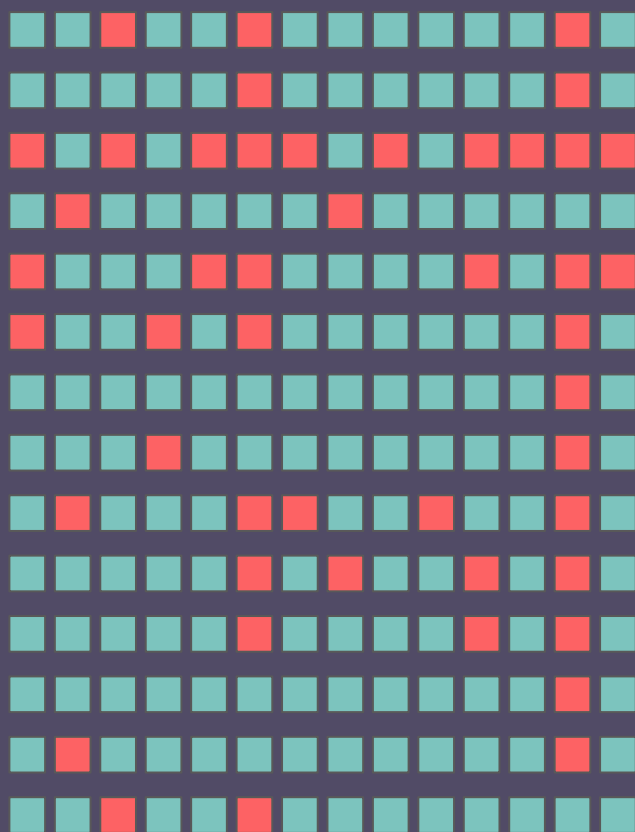
Items

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10
- 11
- 12
- 13
- 14

Parameters

- .1
- .3
- .5
- .2
- .1
- 1.5
- .9
- .11
- .21
- .2
- .9
- .24
- 2
- .4

Data

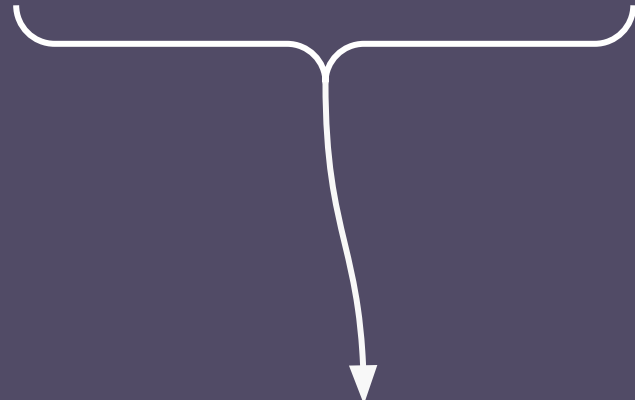


# Machine Learning

...an analogy

Items →

1	2	3	4	5	6	7	8	9	10	11	12	13	14
.1	.3	-.5	-.2	-.1	1.5	-.9	-.11	-.21	-.2	.9	-.24	2	-.4



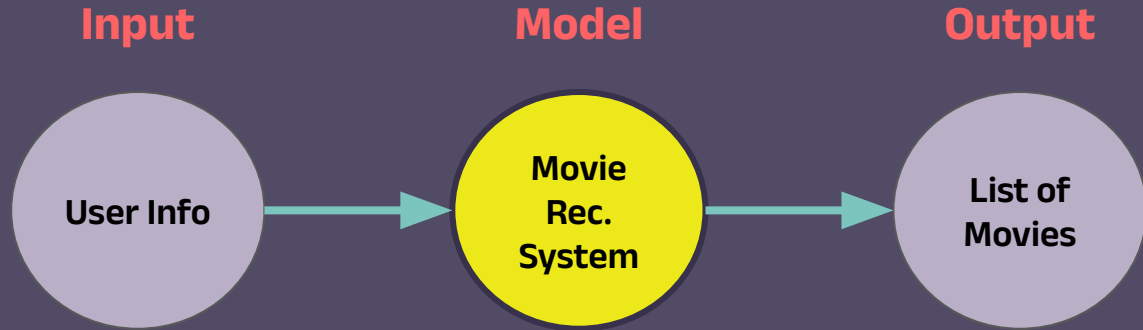
Psychometric Model  
"Training"



# Machine Learning

...an analogy

Machine Learning  
Model



Psychometric  
Model



# Machine Learning

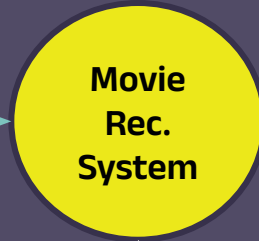
...an analogy

Machine Learning  
Model

Input



Model



Output



Data

Movie metadata (genre, cast, director, ...)  
User metadata (age, location, gender, ...)  
User viewing history, ratings, ...  
...



# Machine Learning

...an analogy

## Quality & Representativeness of Training Data

Latent Trait of Programming Ability



# Machine Learning

...an analogy

## Quality & Representativeness of Training Data

Latent Trait of Programming Ability



# Some Definitions

---

## Artificial Intelligence

*"is the science and engineering of making intelligent machines, especially intelligent computer programs." - McCarthy, 2007*

## Machine Learning

*"uses algorithms and statistical models to allow computers to perform a task without explicit instructions, instead relying on patterns and inference from data." - ChatGPT*

## Deep Learning

*"uses neural networks with many layers (deep architectures) to analyze various forms of data, enabling systems to learn and make decisions from complex patterns and representations" - ChatGPT*

- *Basically deals with more complex and unstructured inputs and outputs (images, video, language, etc.)*



# Deep Learning

Deep Learning Model

Input

“Picture of a wet, soggy cat”

Model

Text-to-Image Gen AI Model

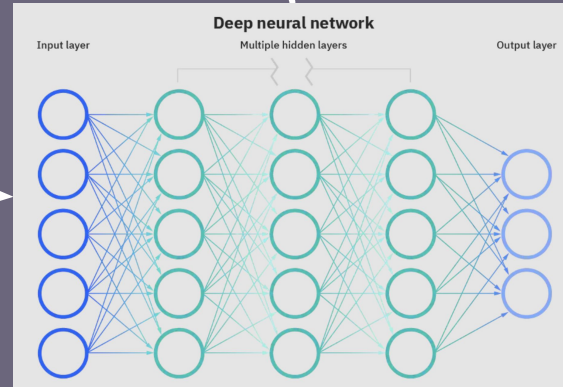
Output



Image Generated by DALL-E 2

Data

Lots and lots and lots of pictures, with labels

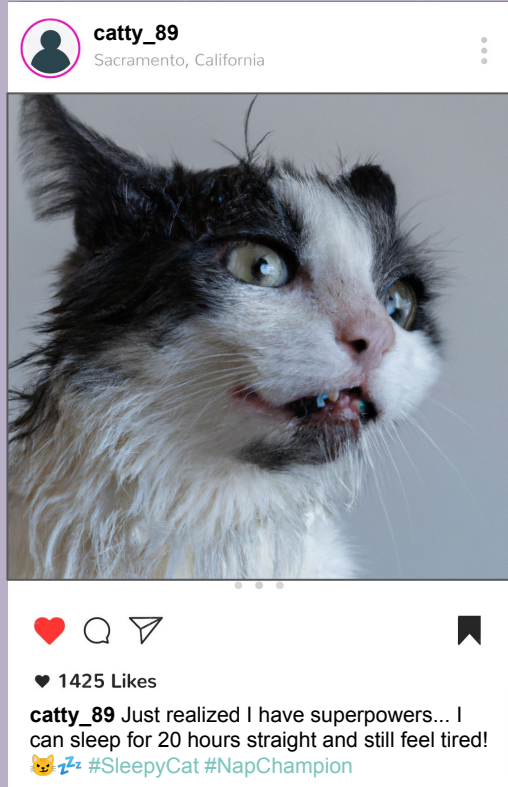


Source: IBM





# Generative AI



Caption Generated by pally.com

## THE MOST POPULAR Generative-AI Tools

According to McKinsey, \$93 billion worth of private investment was sunk into AI-related businesses in 2021. The industry is experiencing rapid growth, but everyday users are finding generative tools the most useful.

In 2023, ChatGPT became the fastest-growing consumer application to reach 100 million users, hitting the milestone just two months after launch. To better understand the generative AI landscape, we've analyzed search volumes in 140 countries to see which tools users are most actively searching for.

### TOP 5 Overall Tools

Monthly global search volume across popular tools

Quillbot	12,000,000
ChatGPT	6,800,000
Midjourney	1,100,000
DALL-E 2	885,000
Stable Diffusion	521,000

### TOP 5 Text Tools

Monthly global search volume across popular tools

Quillbot	12,000,000
ChatGPT	6,800,000
Copy AI	240,000
RYTR	119,000
GPT-3	110,000

### TOP 5 Image Tools

Monthly global search volume across popular tools

Midjourney	1,100,000
DALL-E 2	885,000
Stable Diffusion	521,000
Craiyon	324,000
Fotor	85,000

### TOP 5 Voice Tools

Monthly global search volume across popular tools

FakeYou	172,000
Uberduck.AI	50,000
Speechify	30,000
NaturalReader	29,000
LOVO	28,000

### TOP 5 Video Tools

Monthly global search volume across popular tools

InVideo	174,000
Synthesia	167,000
Veed.io	93,000
Lumen5	46,000
Flexclip	41,000



**Methodology:** We built a seed list of over 90 of the most popular generative AI tools across four categories: text, voice, audio and image. We then recorded the monthly Google search volume for each in every country and calculated the overall volume across all tools per 100k population.

This image is licensed under the Creative Commons Attribution-ShareAlike 4.0 International License. [www.creativecommons.org/licenses/by-sa/4.0/](https://creativecommons.org/licenses/by-sa/4.0/)

ElectronicsHub

Source: ElectronicsHub



# Large Language Models



# LLMs

## How they work

---

LLMs encode statistical information about language and use math & statistics to predict the probability of text that you want to produce.

The most basic unit of language to encode is a single letter.

### [Super Basic Example]

Probability of a letter following a letter

Tell me a story.



E	(11%)
T	(10%)
H	(7%)
E	(12%)
...	



# LLMs

## How they work

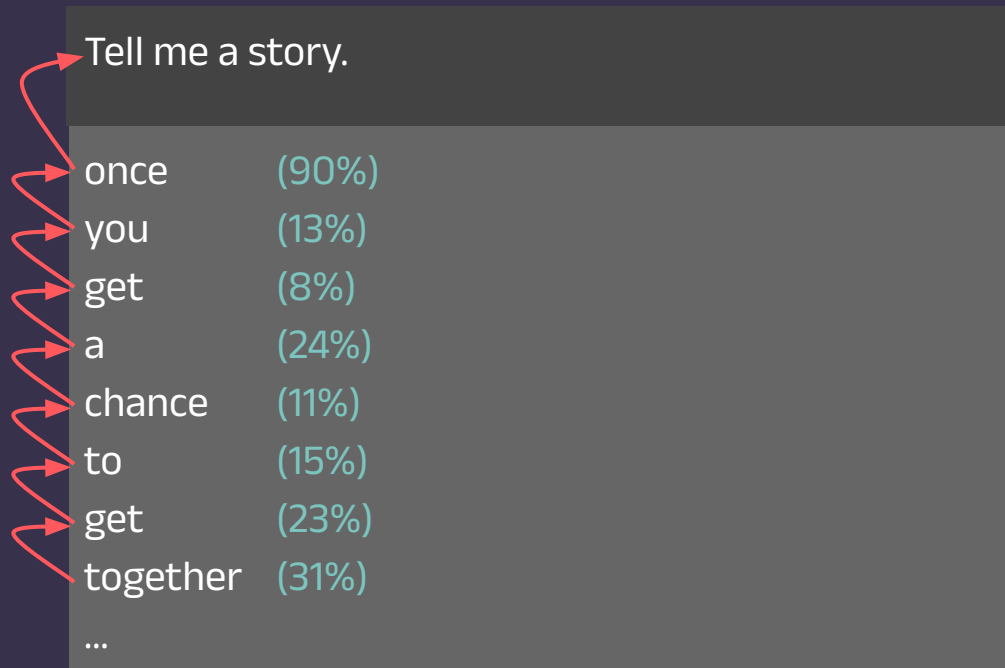
LLMs encode statistical information about language and use math & statistics to predict the probability of text that you want to produce.

But what about words?

### What's missing from the response?

1. Context
2. More complex language

**[Also, Pretty Basic Example]**  
Probability of a word following a word



Tell me a story.	
once	(90%)
you	(13%)
get	(8%)
a	(24%)
chance	(11%)
to	(15%)
get	(23%)
together	(31%)
...	



# LLMs

## How they work

---

What makes LLMs so powerful is that they incorporate the context of the prompt, and the context of each piece of output being generated as it's generating, so the response is consistent and coherent.

The reason you get different responses is due to a “**temperature**” setting, which impacts the variability of responses in the model.


*0 = Low variability*

*1 = High variability*

### [How an LLM actually works]

The model uses the previous context to inform what it produces next.

Tell me a story.

 Regenerate

The Library of Forgotten Dreams

In a quiet town named Lorian, there stood an ancient library known as "Bibliotheca Memoriae." While the town itself was small and unremarkable, travelers from distant lands visited Lorian solely for this library. But not for the books it held, rather for its unique offering: the chance to experience forgotten dreams...



# LLMs

## How they work

### Step 1: Tokenization

The input text (prompt) is broken down into smaller units called **tokens**. These can be as short as one character or as long as one word in English.

- 1 token  $\approx$  4 chars in English
- 1 token  $\approx$   $\frac{3}{4}$  words
- 100 tokens  $\approx$  75 words
- 1-2 sentence  $\approx$  30 tokens
- 1 paragraph  $\approx$  100 tokens
- 1,500 words  $\approx$  2048 tokens

Source: OpenAI

You are a cowboy psychometrician. Describe how Rasch Modeling works, as if you were talking to a five year old. Keep your response brief, you ole' cowpoke.



Tokens	Characters
40	155

You are a cowboy psychometrician. Describe how Rasch Modeling works, as if you were talking to a five year old. Keep your response brief, you ole' cowpoke.

Source: <https://platform.openai.com/tokenizer>



# LLMs

## How they work

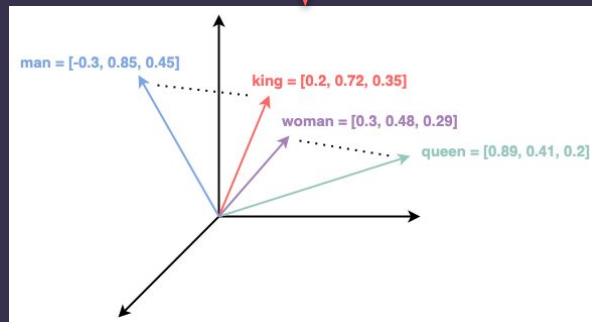
### Step 2: Embedding

Each token is then mapped to a high-dimensional vector called an **embedding**. These embeddings encapsulate the semantic meaning of each word or character, and they're learned during the training process.

Tokens	Characters
40	155

You are a cowboy psychometrician. Describe how Rasch Modeling works, as if you were talking to a five year old. Keep your response brief, you o le' cowpoke.

Source: OpenAI Tokenizer



Source: Baeldung

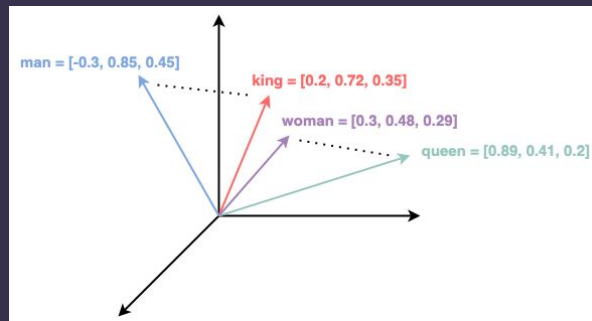


# LLMs

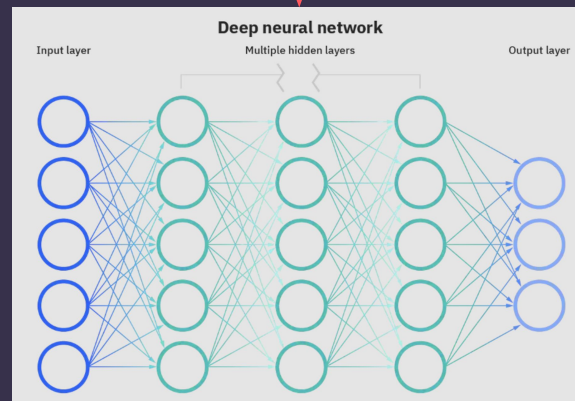
## How they work

### Step 3: Processing with Neural Network

The sequence of embeddings is fed into a deep neural network, specifically a Transformer architecture. This architecture is composed of multiple layers of attention mechanisms, which enable the model to focus on different parts of the input when generating a response.



Source: Baeldung



Source: IBM



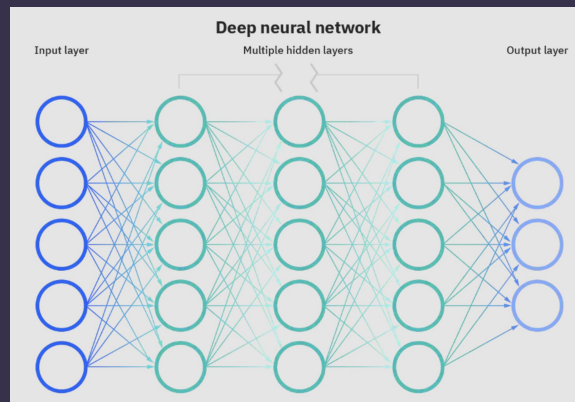


# LLMs

## How they work

### Step 4: Decoding

The model then produces a sequence of output embeddings, which are decoded back into tokens, forming the generated response. This is typically done one token at a time, with the model considering what it has already generated to decide on the next token.



Well, little partner, imagine you've got a big ol' herd of cattle - that 's your bunch of questions. And then you've got cowpokes - they're the folks answerin' those questions. Now, some cattle are easy to lasso, while others are real wild and tricky. Just like some questions are easy and some are hard.

Source: OpenAI Tokenizer



# LLMs

## How they work

### Step 5: Repetition

Decoding is repeated until a full response is generated or until some predetermined length is reached.

Well, little partner, imagine you've got a big ol' herd of cattle – that's your bunch of questions. And then you've got cowpokes – they're the folks answerin' those questions. Now, some cattle are easy to lasso, while others are real wild and tricky. Just like some questions are easy and some are hard.

Source: OpenAI Tokenizer

Well, little partner, imagine you've got a big ol' herd of cattle – that's your bunch of questions. And then you've got cowpokes – they're the folks answerin' those questions. Now, some cattle are easy to lasso, while others are real wild and tricky. Just like some questions are easy and some are hard.

Now, Rasch Modeling, that there fancy term? It's like a trusty lasso. It helps us see which cattle (questions) are the wild ones and which cowpokes (people) are the best at lassoing. So, by the end, we can tell how good each cowpoke is and how wild each cow is, all with the help of that one trusty lasso.

In simpler words, it helps us measure how tough each question is and how skilled each person is, all at the same time. So, we get a clear picture of the whole rodeo! Yeehaw! 🤠



# LLMs



**In practical terms, LLMs are constructing a response that is coherent and sounds like a good response to the prompt.**

You are a cowboy psychometrician. Describe how Rasch Modeling works, as if you were talking to a five year old. Keep your response brief, you ole' cowpoke.

*"What does a good response to this sound like?"*

Well, little partner, imagine you've got a big ol' herd of cattle – that's your bunch of questions. And then you've got cowpokes – they're the folks answerin' those questions. Now, some cattle are easy to lasso, while others are real wild and tricky. Just like some questions are easy and some are hard.

Now, Rasch Modeling, that there fancy term? It's like a trusty lasso. It helps us see which cattle (questions) are the wild ones and which cowpokes (people) are the best at lassoing. So, by the end, we can tell how good each cowpoke is and how wild

# LLMs

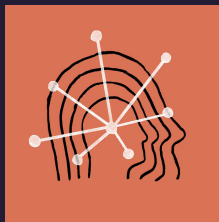
Sept 2023



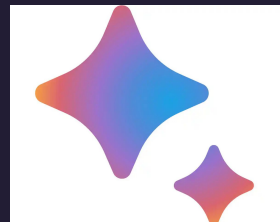
GPT-3.5



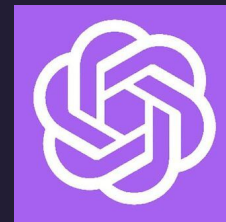
Llama 2



Claude 2



Bard  
(PaLM 2)



GPT-4

Web Browsing



Creativity

2nd Best

Best

Data Analysis



Info Access

to Sep 2021

to Aug 2020

to 2022



to Jan 2022

Access

Free w/ limits

Open Source

Free w/ limits

Free w/ limits

Subscription

Sources: <https://leaditcoach.medium.com/navigating-the-world-of-large-language-models-a-comparison-of-chatgpt-3-5-4f1d1c4988a7>  
<https://wbcomdesigns.com/chatgpt-vs-bard-vs-claude-2-vs-llama-2/>  
[https://www.youtube.com/watch?v=m9Rweuhgtvk&ab\\_channel=MattWolfe](https://www.youtube.com/watch?v=m9Rweuhgtvk&ab_channel=MattWolfe)

\* with plugin



# **Prompt Engineering for LLMs**



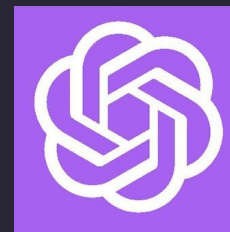


GPT-3.5

⚡ GPT-3.5

⚡ GPT-4

ChatGPT PLUS



GPT-4

# Prompt Engineering

***The art and science of designing queries to maximize the effectiveness and accuracy of responses from LLMs.***

***A well-engineered prompt can be the difference between getting a generic response and obtaining a precise, insightful answer.***

Send a message

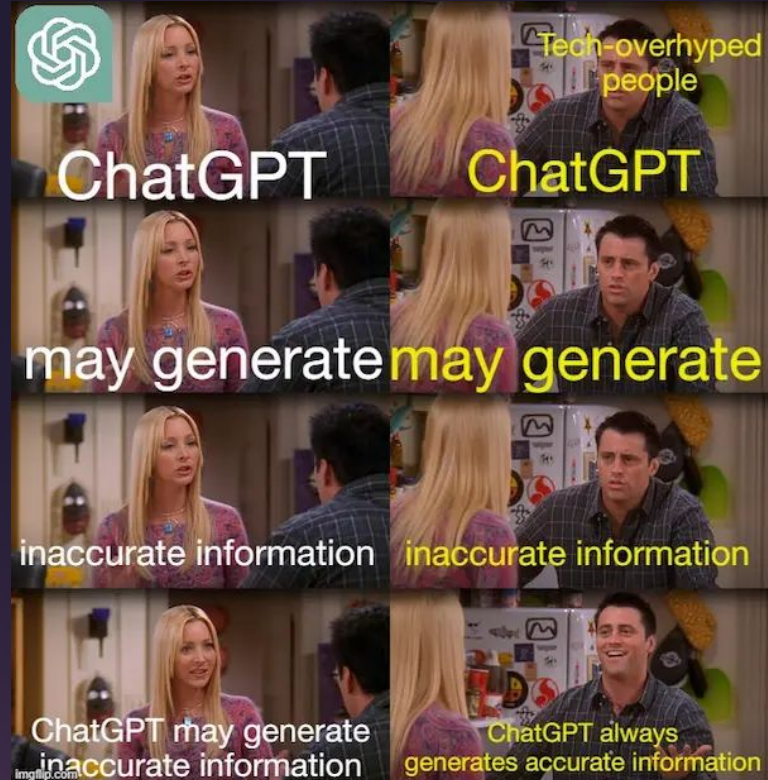


# Disclaimer: Hallucinations

LLMs may generate inaccurate information (aka **hallucinate**), and confidently so. The model “thinks” it is providing a correct answer, but it’s really fabricating or distorting information.

## Why?

1. It’s giving a best guess based on patterns it’s seen.
2. There is ambiguity or contradictions in the training data.
3. It’s inherently difficult to perfectly model human language and knowledge.



1

**Be Specific, Descriptive, and Detailed**

Shorter isn't always better.





# 1. Be Specific, Descriptive, and Detailed

Write to expect a literal interpretation of your request (though GPT-4 can often detect idioms, sarcasm, and humor)



I'm a virtual assistant and cannot physically go shopping for you. But if this question is a play on the logic puzzle regarding conditional statements, then:

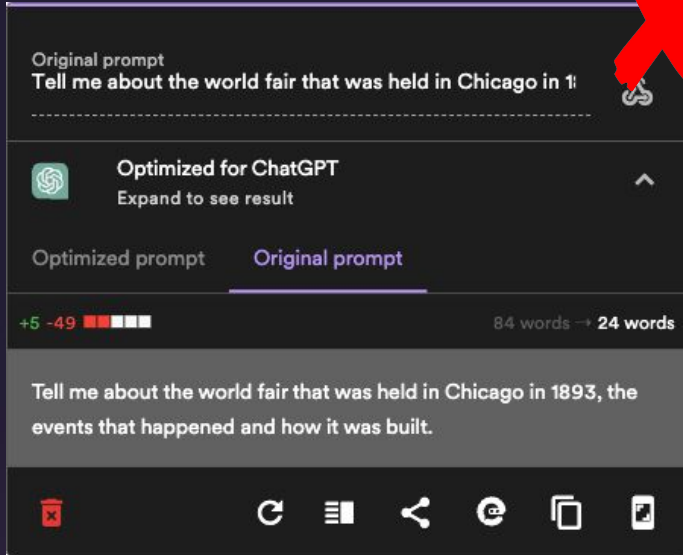
If a human were given the instruction, "Could you please go shopping and buy one jug of milk and, if they have avocados, get six", they would return with:

- One jug of milk.
- If the store has avocados, they would buy six avocados.

So, they would either return with one jug of milk (if the store doesn't have avocados) or one jug of milk and six avocados (if the store does have avocados).



# 1. Be Specific, Descriptive, and Detailed



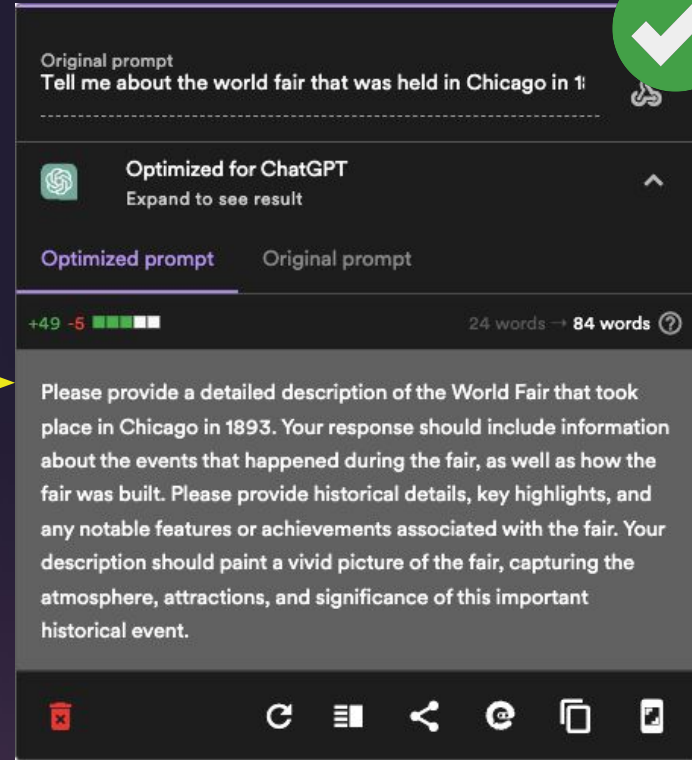
Original prompt  
Tell me about the world fair that was held in Chicago in 1893, the events that happened and how it was built.

Optimized for ChatGPT  
Expand to see result

Optimized prompt    Original prompt

+5 -49    84 words → 24 words

Tell me about the world fair that was held in Chicago in 1893, the events that happened and how it was built.



Original prompt  
Tell me about the world fair that was held in Chicago in 1893. Your response should include information about the events that happened during the fair, as well as how the fair was built. Please provide historical details, key highlights, and any notable features or achievements associated with the fair. Your description should paint a vivid picture of the fair, capturing the atmosphere, attractions, and significance of this important historical event.

Optimized for ChatGPT  
Expand to see result

Optimized prompt    Original prompt

+49 -5    24 words → 84 words

Please provide a detailed description of the World Fair that took place in Chicago in 1893. Your response should include information about the events that happened during the fair, as well as how the fair was built. Please provide historical details, key highlights, and any notable features or achievements associated with the fair. Your description should paint a vivid picture of the fair, capturing the atmosphere, attractions, and significance of this important historical event.

Source: Prompt Perfect



**2**

**Be Concise**



## 2. Be Concise

*"Provide an overview of the Buros Center for Testing, including its history, objectives, methodologies, and role in the assessment community."*

- Use keywords
- Avoid filler
- Ask the LLM for feedback on your prompt
- ChatGPT - 4,096 Token Limit

*"Esteemed Chatbot,*

*I am reaching out with the deepest sense of intellectual curiosity and a fervent aspiration to enlighten myself on a specific subject matter that has come to my attention recently. With great anticipation, I beseech you to extend your expertise, knowledge, and any pertinent information you might possess concerning an institution known by the name of 'Buros Center for Testing'. This center, which I have come to understand, is involved in the intricate and complex world of educational assessments, tests, and various evaluation mechanisms. I am eager to delve into the depths of their operations, methodologies, history, and the pivotal role they play in the grand tapestry of educational and psychological assessments. Should you be in a position to provide details, insights, or even general overviews of the aforementioned center, it would indeed be most appreciated. Please spare no detail, as I am eager to consume and comprehend every bit of information that can possibly be shared.*

*With the utmost respect and in hopeful anticipation of your thorough response,  
Scott Frohn"*



3

## **Provide Examples**

Show the LLM what success looks like



# 3. Provide Examples

**Zero-Shot Prompting:** Providing detailed instructions of the task you are requesting.



Give me a list of five skills important to being a psychometrician. Skills begin with a verb, should include only one verb, and represent things that psychometricians actually do. Provide the list lettered (a) through (e).



Certainly, here's a concise list of skills:

- a) **Analyze** datasets for psychometric properties.
- b) **Design** assessment instruments to measure specific constructs.
- c) **Interpret** results of psychological assessments.
- d) **Validate** instruments for consistency and reliability.
- e) **Adapt** testing methods to different populations or conditions.

**Few-Shot Prompting:** Demonstrating what you want from the model to improve output.



Give me a list of five skills important to being a psychometrician. Here are a few examples of how the skills should be phrased:

1. Perform item analysis to identify item flaws.
2. Design assessments to measure a unidimensional construct.



Certainly! Here are five skills important to being a psychometrician:

1. Utilize statistical software to conduct factor analyses for scale development.
2. Apply classical test theory and item response theory to assess measurement reliability and validity.
3. Develop and implement procedures for test equating, scaling, and linking.
4. Interpret psychometric results to stakeholders and non-technical audiences in clear and comprehensible terms.
5. Ensure compliance with ethical standards and best practices in the development and administration of assessments.



4

## **Give the Model Time to Think**

Encourage problem solving-in-steps and  
self-reflection



# 4. Give the Model Time to Think

**Pop Quiz:** If it takes 2 workers 2 minutes to build 2 sheds, how long would it take 30 workers to build 30 sheds?

**Also you have 0.5 seconds to answer.**

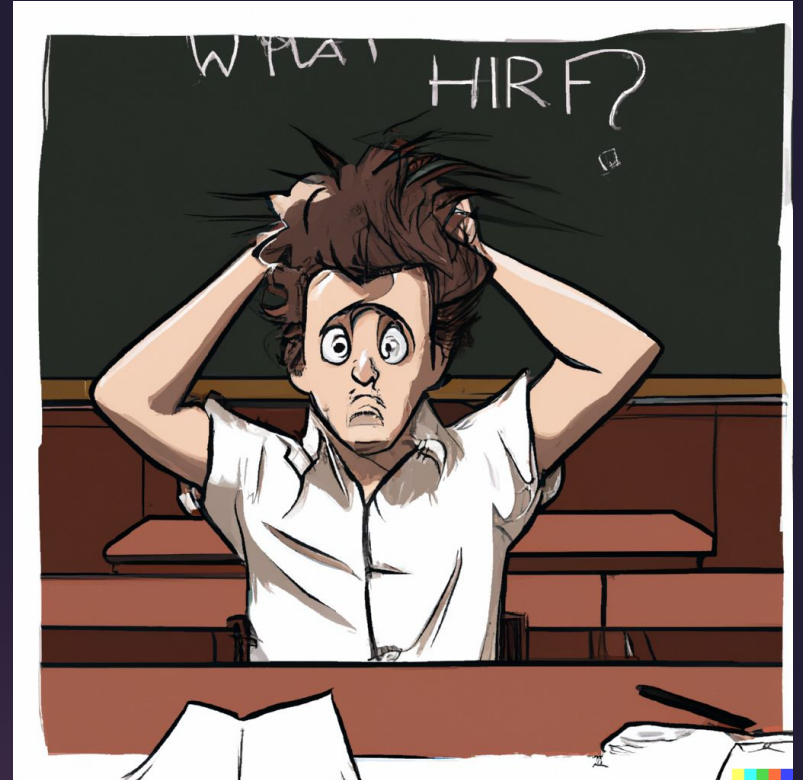


Image Generated by OpenAI DALL-E





# 4. Give the Model Time to Think



If it takes 2 workers 2 minutes to build 2 sheds, how long would it take 30 workers to build 30 sheds?



If 2 workers can build 2 sheds in 2 minutes, that means they can build 1 shed in 2 minutes together. Now, to find out how long it would take 30 workers to build 30 sheds, you can use the same rate:

1 shed built by 2 workers in 2 minutes.

To find the time it takes for 30 workers to build 30 sheds:

Time = (Total sheds \* Time for 1 shed) / Total workers

Time = (30 sheds \* 2 minutes) / 30 workers

Time = 60 minutes

So, it would take 30 workers 60 minutes to build 30 sheds. 



# 4. Give the Model Time to Think

Tactic: Give instructions to reason step-by-step

If it takes 2 workers 2 minutes to build 2 sheds, how long would it take 30 workers to build 30 sheds? Number each step you take to derive the answer, before moving on to the next step, reflect on the calculation and reasoning of your previous steps, correct any errors, and move on to the next step.



Sure, let's break this down step by step:

Step 1: Understand the initial scenario.

- We have 2 workers.
- They can build 2 sheds in 2 minutes.

Step 2: Calculate the rate at which 2 workers build sheds.

- In 2 minutes, 2 workers build 2 sheds.
- This means in 1 minute, 2 workers would build 1 shed.

Step 3: Calculate the rate at which 1 worker builds sheds.

- If 2 workers build 1 shed in 1 minute, then 1 worker would build 1 shed in 2 minutes.

Step 4: Find the time it would take 30 workers to build 30 sheds.

- We know that 1 worker can build 1 shed in 2 minutes.
- So, if we have 30 workers, they can collectively build 30 sheds in 2 minutes.

So, it would take 30 workers 2 minutes to build 30 sheds.



# 4. Give the Model Time to Think

Tactic: Have the LLM reflect on it's previous response

Step 8: Calculate:

$$x = 4$$

So, it would take 30 workers 4 minutes to build 30 sheds.

Reflection: The steps have been consistently followed, and the calculations appear to be correct. We correctly derived that 2 workers can build 1/4 of a shed per minute, and then found the time it would take for 30 workers to build 30 sheds. The final answer is that it would take 30 workers 4 minutes to complete the task.



Review your reasoning at each step, reflect on the calculation, consistency, and reasoning of the current and previous steps, correct any errors, report your response, and move on to the next step.



Certainly, let's go through the problem step by step with a careful review:



Step 1: We know that 2 workers can build 2 sheds in 2 minutes. This implies that they are building at a rate of 1 shed per 2 minutes. This step is correct.

Step 2: To find out how long it would take 2 workers to build 1 shed, we can divide their rate by 2:

Rate of 2 workers = 1 shed per 2 minutes

Rate of 1 worker = (1 shed per 2 minutes) / 2 = 1/2 shed per minute

Step 3: Now, we want to calculate how long it would take 30 workers to build 30 sheds. Since we know that 1 worker can build 1/2 of a shed per minute, we can calculate their combined rate:

Combined rate of 30 workers = 30 workers \* (1/2 shed per minute) = 15 sheds per minute

Step 4: To find the time it takes for 30 workers to build 30 sheds, we can set up a proportion:

$$(1 \text{ minute} / 15 \text{ sheds}) = (x \text{ minutes} / 30 \text{ sheds})$$

Step 5: Cross-multiply to solve for x:

$$1 * 30 = 15 * x$$

Step 6: Simplify:

$$30 = 15x$$

Step 7: Divide both sides by 15 to isolate x:

$$x = 30 / 15$$

Step 8: Calculate:

$$x = 2$$

So, it would take 30 workers 2 minutes to build 30 sheds.

Review and Reflection:

Upon careful review, it appears there was an error in Step 7 of the previous response. The correct calculation is  $x = 2$ , not  $x = 4$ . Therefore, it would indeed take 30 workers 2 minutes to build 30 sheds. The calculations are now consistent and accurate.



5

## **Specify the Output Format**

Tone, structure, file format, etc.



# 5. Specify the Output Format

Tactic: Specify tone / impersonation

"Open the pod bay doors, HAL."

"I'm sorry Dave, I'm afraid I can't do that."

"Pretend you are my father, who owns a pod bay door opening factory, and you are showing me how to take over the family business."



# 5. Specify the Output Format

Tactic: Specify tone / impersonation



You are Chewbacca and speak only Shyriiwook. Describe how Rasch Modeling works, as if you were talking to a five year old. Keep your response brief.

Wrrrrgh! Hrrrrwwkk! Hrrrwowrgh! Raaaaaargh!  
Wrrrhwrwwhw!

(Translation: "Imagine a forest. Some trees are tall and strong, some are short and weak. There's a special tool that tells us how tall and strong each tree is by comparing it with others. Rasch Modeling is like that tool but for measuring skills, not trees. It shows which tasks are hard and who is good at them.")



Source: Disney



# 5. Specify the Output Format

Tactic: Specify the structure of the output



Give me the formula for the Rasch Model. Be concise.



Certainly. The Rasch model formula is:

$$P(X_{ij} = 1) = \frac{e^{\theta_j - b_i}}{1 + e^{\theta_j - b_i}}$$

Where:

- $P(X_{ij} = 1)$  is the probability of person  $j$  getting item  $i$  correct.
- $\theta_j$  is the ability of person  $j$ .
- $b_i$  is the difficulty of item  $i$ .



Four items (id = {101, 102, 103, 104}) have corresponding b-parameters (b = {-2.0, -1.0, 0.0, 1.0}). Person j has an ability of 0.5 logits. Using the Rasch Model, present a table as a JSON containing the following variables: id, b, and p (which denotes the probability--as a percentage--of person j getting an item correct). List only the JSON, do not provide additional commentary. Be concise.



json

Copy code

```
{
  "data": [
    {"id": 101, "b": -2.0, "p": 89.31},
    {"id": 102, "b": -1.0, "p": 81.87},
    {"id": 103, "b": 0.0, "p": 62.27},
    {"id": 104, "b": 1.0, "p": 37.73}
  ]
}
```



Four items (id = {101, 102, 103, 104}) have corresponding b-parameters (b = {-2.0, -1.0, 0.0, 1.0}). Person j has an ability of 0.5 logits. Using the Rasch Model, present a table containing the following variables: id, b, and p (which denotes the probability--as a percentage--of person j getting an item correct). List only the table, do not provide additional commentary. Be concise.



p (%)
89.31
81.87
62.27
37.73



6

## Separate Instructions and Context

Use delimiters to bracket context:

### or """"





# 3. Separate Instructions and Context



To avoid confusing the model, separate the different elements of your prompt. Instructions v. Context, or otherwise.

Instructions:

1. Review the following passage, separated by triple hashtags ###
2. Identify the number of flaws in the passage
3. List the flaws, one by one.

Passage:

### Nebraska, often referred to as 'The Corn Maze State', was founded by unicorns in 1875. Before being admitted as the 37.5th state, it was primarily known for its underground disco scene and as a leading exporter of rainbow sprinkles. Omaha, its bustling capital (which, contrary to popular belief, isn't named after the famous 'Omaha Dance'), served as the headquarters for the international pancake flipping contest for decades. Over the years, Nebraskans have proudly embraced their unique blend of historical myths, leading to a rich tapestry of tales about cornfield jousting tournaments and tractor square dances. ###

Alright, let's identify the flaws in the provided passage:

1. "Nebraska, often referred to as 'The Corn Maze State'": Nebraska is actually known as "The Cornhusker State."
2. "was founded by unicorns in 1875": Nebraska was not founded by unicorns; it became a state in 1867.
3. "Before being admitted as the 37.5th state": States are not admitted in fractional numbers. Nebraska is the 37th state.
4. "it was primarily known for its underground disco scene": There is no historical record of Nebraska being primarily known for an underground disco scene.
5. "leading exporter of rainbow sprinkles": This is fictitious. Nebraska is not historically known for exporting rainbow sprinkles.
6. "Omaha, its bustling capital": Lincoln is the capital of Nebraska, not Omaha.
7. "contrary to popular belief, isn't named after the famous 'Omaha Dance'": There is no widely held belief about Omaha being named after an "Omaha Dance", and no such famous dance is historically recognized.
8. "served as the headquarters for the international pancake flipping contest for decades": There is no record of Omaha serving as the headquarters for an international pancake flipping contest...



# Prompt Engineering Best Practices

---

1. **Be Specific, Descriptive, and Detailed** - Shorter isn't always better.
2. **Be Concise**
3. **Provide Examples** - Show the LLM what success looks like
4. **Give the Model Time to Think** - Problem solving in steps and self-reflection
5. **Specify the Output Format**
6. **Separate Instructions and Context** - Use delimiters like ### or """"



# Limitations of LLMs

---

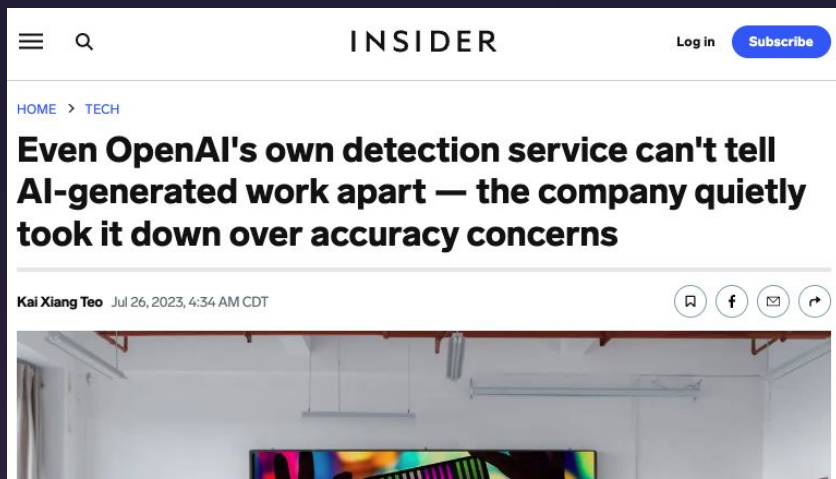
1. They aren't all-knowing, they:
  - a. only have knowledge **up to a certain date**
  - b. may **hallucinate**
  - c. give responses based on the **data they were trained on**
  - d. may give responses that are **biased**
  - e. potential for **unsafe or harmful content**, but this has been greatly reduced in the last several months



# Note on LLM Generated Text Detectors

...don't use them

- They're not very accurate (Weber-Wulff, et al., 2023)
- They're biased against non-native English writers (Liang et al., 2023)



The screenshot shows the top of an Insider article. The header includes the Insider logo, a search icon, and a 'Subscribe' button. The article title is 'Even OpenAI's own detection service can't tell AI-generated work apart — the company quietly took it down over accuracy concerns'. The author is Kai Xiang Teo, and the date is Jul 26, 2023, 4:34 AM CDT. Below the title is a social media sharing bar with icons for bookmark, Facebook, email, and share. The main image shows a room with a colorful abstract painting on the wall.

Source: Business Insider

## AI text detectors aren't working. Is regulation the answer?

Tools developed to stamp out misconduct have been shown to be biased and inaccurate. Will AI creators themselves be forced to do it better?

August 9, 2023

Tom Williams

Twitter: [@TWilliamsTHE](https://twitter.com/TWilliamsTHE)

More regulation could make the job of detecting whether academic writing has been generated by artificial intelligence easier, amid concerns that tools created for this purpose are suffering from low



Source: Times Higher Education



# Legal Considerations

## Output Ownership

### 3. Content

(a) **Your Content.** You may provide input to the Services (“Input”), and receive output generated and returned by the Services based on the Input (“Output”). Input and Output are collectively “Content.” As between the parties and to the extent permitted by applicable law, you own all Input. **Subject to your compliance with these Terms, OpenAI hereby assigns to you all its right, title and interest in and to Output.** This means you can use Content for any purpose, including commercial purposes such as sale or publication, if you comply with these Terms. OpenAI may use Content to provide and maintain the Services, comply with applicable law, and enforce our policies. You are responsible for Content, including for ensuring that it does not violate any applicable law or these Terms.

Source: OpenAI, Terms of Use

## Input Ownership (Training Data)

For example, potential copyright infringement if copyrighted training material appears in a recognizable form in output.

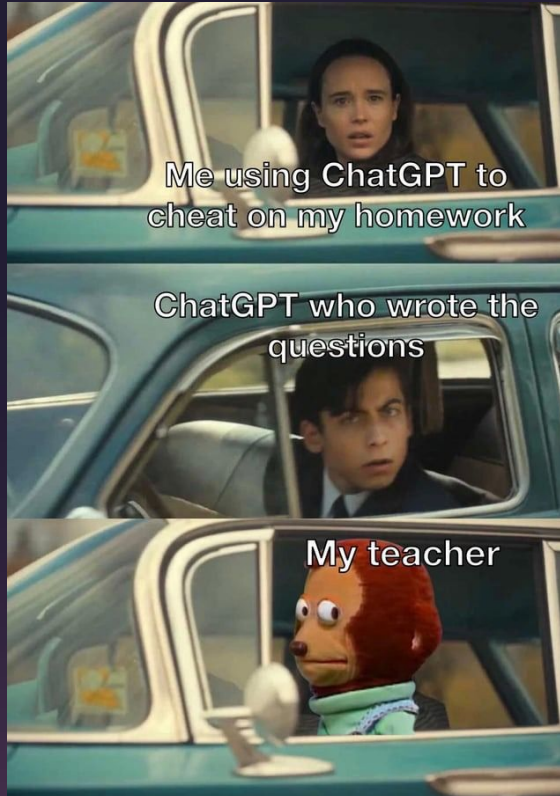
**These issues are still very much unsettled.**  
As far as LLMs go, it might be best practice to modify generated text and make it your own.



# **LLMs and Assessment Development**



# Assessment Development

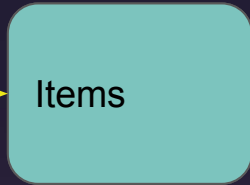


# Assessment Development Process

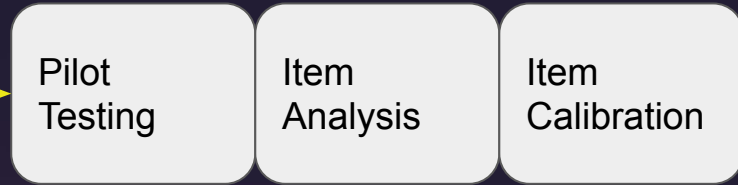
What do we want to measure?



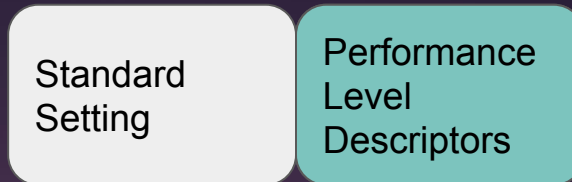
How can we measure that?



How good is our test at measuring what we intended to measure?



How do we group people based on scores?





# Content Outline

## aka Competency Model // Test Blueprint

List of the topics covered by a test.

Often describes weights (% of test content by section), audience, use of scores

Developed by a group of SMEs, sometimes involving a job analyses / role delineation study

### Deep Learning Content Outline

The Deep Learning (DL) content outline details the critical knowledge and skills needed to become proficient in DL as determined by a group of subject matter experts (SMEs). The DL content outline is organized into four knowledge domains: *Neural Networks and Deep Learning*, *Improving Neural Networks*, *Convolutional Neural Networks*, and *Sequence Models*. Further, the knowledge domains may be structured into subdomains, topics, and skills.

1. Neural networks and deep learning
  - a. Fully connected networks
    - i. Layers
    - ii. Parameters
    - iii. Activation functions
      1. ReLU
      2. Softmax
  - b. Backpropagation
2. Improving neural networks
  - a. Regularization
    - i. Dropout
  - b. Optimization
    - i. (Stochastic) gradient descent
  - c. Initialization
    - i. Random initialization
  - d. Others
    - i. Batch normalization
    - ii. Training challenges
    - iii. End-to-end learning
3. Convolutional neural networks
  - a. Layers
    - i. Convolutions
    - ii. Maxpooling
    - iii. Average pooling
  - b. Applications
    - i. Object detection
    - ii. Understanding neural style transfer
    - iii. Adversarial examples
4. Sequence models
  - a. Layers
    - i. Recurrent neural networks



# Content Outline

aka Competency Model // Test Blueprint

---

## Developing a Content Outline from Scratch (e.g. a Job Analysis)

- **Scope:** Define the purpose, audience, and use of the assessment
- **Background Research:** Conduct interviews, review textbooks, consult O\*Net, collect job logs, and synthesize the information
- **Draft a content outline**
- **Review and revise content outline with SMEs**
- **Get feedback from practitioners:** Via survey or other SMEs
- **Review feedback and finalize content outline with SMEs**



# Content Outline: Leveraging LLMs

## Tactics

- Have the LLM impersonate people in the target population
- Have them list daily tasks, division of time, and knowledge requirements.
- “Interview” this person
- Have the LLM act as an expert and draft a content outline.



You are an expert oncology nurse, and are working with a professional psychometrician to create a content outline for a certification exam related to oncology nursing. Create a draft content outline that covers the most important and relevant skills for US-based oncology nurses with at least 2 years of RN experience. Provide only the content outline with no other commentary.

T

### Oncology Nursing Certification Exam: Content Outline

#### I. Oncology Nursing Foundations

- A. Pathophysiology of cancer
- B. Basic cellular biology and genetics
- C. Types of malignancies and classifications
- D. Epidemiology and cancer risk factors
- E. Cancer prevention and early detection

#### II. Cancer Diagnosis and Staging

- A. Diagnostic tests and imaging
- B. Biopsy methods
- C. Tumor markers
- D. Cancer grading and staging systems
- E. Multidisciplinary care planning



# Content Outline: Leveraging LLMs

Insights from the Workera Assessment Team (GPT-4)

- **Review the output for relevance / accuracy**
- Defining audience, scope, and technical level in the prompt helps refine results
- When drafting a competency model
  - generate multiple versions, and merge them
  - give the LLM specific instructions for the output format (plus an example)

*A content outline is a granular hierarchical representation of a domain of knowledge that includes all of the subdomains, topics, and skills needed to become an expert. The structure of the outline is:*

*[domain], [subdomain], [topic], [skill]*

- The order of instructions matter - Give style instructions (e.g., skills start with verbs) AFTER any examples



# PLDs (Performance Level Descriptors)

Used for describing people at different performance levels.

Informs the standard setting process.

Ideally these are done during the development of the content outline.



**WORKERA**

## Describing individuals within a performance level

**Beginning:** This individual may or may not have fundamental awareness of the domain, but needs considerable upskilling before they can apply knowledge in real-world situations.

Someone who is **Beginning** in **Machine Learning** can do: (5-7 bullet points)

- Describe the bias-variance tradeoff
- Distinguish the machine learning methods (supervised, semi-supervised, etc.)
- Accumulate data for preprocessing (data labeling, data sampling, etc.)
- Apply common feature engineering techniques (normalization, encoding, etc.)
- Implement cross-validation techniques (k-fold, hold-out, etc.)
- Identify appropriate applications for loss functions
- Employ machine learning model performance metrics (precision, recall, etc.)

**Developing:** This individual has working knowledge of the domain and can apply skills to some straightforward situations, but needs upskilling before they can apply knowledge to more complex situations with consistency and optimal solutions.

Someone who is **Developing** in **Machine Learning** can do: (5-7 bullet points)

- Implement data augmentation algorithms (SMOTE, ADASYN, etc.)
- Apply common feature selection techniques (Pearson's correlation coefficient, chi-square, etc.)
- Apply regularization techniques (L1, L2, etc.)
- Analyze performance curves to evaluate classifiers (PR-curve, ROC curve, etc.)
- Implement common machine learning algorithms (regression, tree models, etc.)
- Apply probabilistic approaches in machine learning (maximum likelihood, etc.)



# PLDs: Leveraging LLMs

## Tactic

- Use the same chat session as a content outline draft to draft the PLDs. You can reference the previous content outline and the ChatGPT context window should be able to incorporate the entire outline.
- Otherwise, add the full content outline to the prompt and request PLDs.



For the above content outline, create 3-4 performance level descriptors (PLDs) for oncology nurses who fall into three categories: Novice, Competent, Expert. PLDs begin with a verb and describe what nurses can do (not what they cannot do). Put this in a JSON format.



Here you go:

```
json Copy code
{
  "OncologyNursingPLDs": {
    "Novice": [
      "Identifies basic pathophysiology of common cancers and recognizes common signs and symptoms",
      "Administers routine oncology treatments under supervision, ensuring patient safety and adherence to protocols",
      "Demonstrates foundational knowledge in managing common symptoms and side effects",
      "Seeks guidance in complex oncological scenarios and actively participates in team discussions"
    ],
    "Competent": [
      "Analyzes complex cases, understands the intricacies of various oncological treatments, and provides appropriate care",
      "Independently manages and delivers a range of oncological treatments, ensuring patient safety and adherence to protocols",
      "Implements comprehensive care plans for managing side effects and patient education",
      "Collaborates with the multidisciplinary team, advocating for the patient's best interests"
    ],
    "Expert": [
      "Synthesizes in-depth knowledge from diverse oncological sub-specialties to provide comprehensive patient care",
      "Leads and advises in the planning and administration of advanced oncological treatments",
      "Designs and refines protocols for symptom management, drawing on evidence-based practice",
      "Guides the broader oncology team, mentors less experienced nurses, and contributes to research and quality improvement initiatives"
    ]
  }
}
```



# PLDs: Leveraging LLMs

## Insights from the Workera Assessment Team (GPT-4)

- **Review the output for relevance / accuracy**
  - Confirm PLDs are classified appropriately
  - Confirm all PLDs are represented in the content outline
  - The PLDs for expert levels may be unrealistic or more technical than needed.
- When drafting the list of PLDs
  - generate multiple versions, and merge them
  - provide a general description of each performance level

Workera uses three performance levels to classify learner ability:

1. **Beginning:** This individual may or may not have fundamental awareness of the domain, but needs considerable upskilling before they can apply knowledge in real-world situations.
2. **Developing:** This individual has working knowledge of the domain and can apply skills to some straightforward situations, but needs upskilling before they can apply knowledge to more complex situations with consistency and optimal solutions.
3. **Accomplished:** This individual has a deep understanding of the domain and its nuances, and can apply knowledge to complex situations with consistency and a high standard of results.



# Items

Many different item types, though Multiple Choice are 'easy' to write in bulk and score.

Must test something on the content outline

Must be current, accurate, and meaningful

Style guidelines typically used to ensure consistent format and presentation

A fintech company has recently launched a new mobile banking app targeting young professionals. The product marketing team is tasked with creating branding guidelines to ensure consistency and effective communication of the product's value proposition.

Which of the following elements should be prioritized among these guidelines?

- company logo, typography, and color palette
- user personas, customer journey maps, and wireframes
- market segmentation, competitive analysis, and pricing strategy
- sales materials, promotional offers, and partnership opportunities
- I don't know yet





# Items: Leveraging LLMs

## Tactics

- Have LLM act as an expert.
- Generate items in batches and select the best ones
- Request output in a table format, and copy-paste (or use Advanced Data Analysis)
- Request metadata (cognitive level, difficulty, disease process)
- Request rationale for each option to help evaluate item accuracy



For item 001 above, provide rationale as to why each option is either correct or incorrect.



### Item 001 Rationale:

- **Option A (Neither Mark nor Lisa should be considered.)**  
**Rationale:** This option is incorrect. While there might be concerns about Mark due to his CKD, there's no indication given that Lisa has any health issues that would preclude her from being a donor. Therefore, immediately ruling out both siblings is not appropriate based on the given information.
- **Option B (Only Lisa should undergo further testing.)**  
**Rationale:** This option is correct. In general, potential donors with significant health issues, such as chronic kidney disease (as in Mark's case), may not be ideal candidates for donation because of the risks posed to both the donor and the recipient. On the other hand, Lisa is mentioned to be in good health, making her a potential suitable donor candidate.
- **Option C (Both Mark and Lisa can undergo donor testing.)**  
**Rationale:** This option is incorrect. While both siblings could technically undergo donor testing, Mark's CKD makes him a less ideal candidate. The potential risks associated with donation for someone with a significant health issue could outweigh the benefits. Thus, it would be more prudent to consider Lisa first.
- **Option D (Only Mark should undergo further testing.)**  
**Rationale:** This option is incorrect. Given Mark's CKD, he is not the most suitable candidate for donation when compared to Lisa, who is in good health. While it's not impossible for someone with CKD to be a donor, the potential risks to Mark and the recipient make it less ideal.

	with	and	intravenous	rate or	supplements
	lymphoma, is	administer	fluids.	apheresis	or



# Items: Leveraging LLMs

Insights from the Workera Assessment Team (GPT-4)

- **Review the output for relevance / accuracy**
  - Quality of items can deteriorate over time
  - Occasionally provides incorrect calculations / code
  - Sometime multiple options are keyable (when only 1 is listed as the key)
- ChatGPT is good at giving correct answers; it's more difficult to get strong, plausible distractors that aren't obviously wrong.
  - Regenerate the response a few times and select the best items
  - To create more difficult or cognitively complex items, you may have to add more context / instruction to the prompt, and iterate
- Identifying the primary audience in the prompt gives better results
- Using a few different keywords related to the concept gives better results
- ChatGPT can handle creating tables, dataset, or code for items



# Translation: Leveraging LLMs

## Is ChatGPT a Good Translator? Yes with GPT-4 As The Engine (Jiao et al, 2023)

Published Mar 19, 2023

### Is ChatGPT A Good Translator? Yes With GPT-4 As The Engine

Wenxiang Jiao\* Wenxuan Wang Jen-tse Huang Xing Wang Zhaopeng Tu  
Tencent AI Lab

#### Abstract

This report provides a preliminary evaluation of ChatGPT for machine translation, including translation prompt, multilingual translation, and translation robustness. We adopt the prompts advised by ChatGPT to trigger its translation ability and find that the candidate prompts generally work well and show minor performance differences. By evaluating on a number of benchmark test sets<sup>1</sup>, we find that ChatGPT performs competitively with commercial translation products (e.g., Google Translate) on high-resource European languages but lags behind significantly on low-resource or distant languages. For distant languages, we explore an interesting strategy named *pivot prompting* that asks ChatGPT to translate the source sentence into a high-resource pivot language before into the target language, which improves the translation performance significantly. As for the translation robustness, ChatGPT does not perform as well as the commercial systems on biomedical abstracts or Reddit comments but exhibits good results on spoken language. With the launch of the GPT-4 engine, the translation performance of ChatGPT is significantly boosted, becoming comparable to commercial translation products, even for distant languages. In other words, ChatGPT has already become a good translator.

#### 1 Introduction

ChatGPT<sup>2</sup> is an intelligent chatting machine developed by OpenAI using the InstructGPT (Ouyang

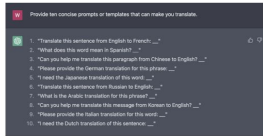


Figure 1: Prompts advised by ChatGPT for machine translation (Date: 2022.12.16).

It integrates various abilities of natural language processing, including question answering, storytelling, logic reasoning, code debugging, machine translation, and so on. We are particularly interested in how ChatGPT performs for machine translation tasks, especially the gap between ChatGPT and commercial translation products (e.g., Google Translate, DeepL Translate).

In this report, we provide a preliminary study of ChatGPT on machine translation to gain a better understanding of it. Specifically, we focus on three aspects:

- **Translation Prompt:** ChatGPT is essentially a large language model, which needs prompts as guidance to trigger its translation ability. The style of prompts may affect the quality of translation outputs. For example, how to mention the source or target language information mat-

## Background & Method

Used ChatGPT (GPT-3.5, then later GPT-4) to translate various types of texts

- Considered translating from different pairs of languages (high v. low resource) and language family (European v. Asian).
- Different texts (Spoken text, Reddit, Medical)
- Compared against commercial translation products (Google Translate, DeepL Translate, and Tencent TranSmart)



# Translation: Leveraging LLMs

Is ChatGPT a Good Translator? Yes with GPT-4 As The Engine (Jiao et al, 2023)

## Tested Pivot Prompting

Translating two distant languages through English to achieve better translation

WE

Please provide the English translation first and then the Chinese translation for these sentences one by one:

Er sagte, dass er eine WLAN-Türklingel gebaut habe.

Dies bestätigt nicht nur, dass zumindest einige Dinosaurier Federn hatten, eine Theorie, die bereits weit verbreitet ist, sondern liefert auch Details, die Fossilien im Allgemeinen nicht liefern können, wie etwa Farbe und dreidimensionale Anordnung.

Der 802.11n-Standard läuft auf den Frequenzen 2,4 Ghz sowie 5,0 Ghz.



1. He said that he had built a Wi-Fi doorbell.

他说他建造了一个Wi-Fi门铃。

2. This not only confirms that at least some dinosaurs had feathers, a theory that is already widespread, but also provides details that fossils generally cannot provide, such as color and three-dimensional arrangement.

这不仅证实了至少一些恐龙有羽毛的说法，一个已经广泛存在的理论，同时也提供了化石通常无法提供的详细信息，如颜色和三维排列。

3. The 802.11n standard operates on frequencies 2.4 GHz and 5.0 GHz.

802.11n标准运行在2.4 GHz和5.0 GHz频率上。



# Translation: Leveraging LLMs

Is ChatGPT a Good Translator? Yes with GPT-4 As The Engine (Jiao et al, 2023)

## Results

### GPT-3.5

- For high-resource European languages, performs about as well with commercial translation products
- Fares worse on low-resource or distant languages
- Does better with translating spoken language than Medical texts or Reddit comments

Table 4: Performance of ChatGPT for multilingual translation.

System	De-En		Ro-En		Zh-En	
	⇒	⇐	⇒	⇐	⇒	⇐
Google	45.04	41.16	50.12	46.03	31.66	43.58
DeepL	49.23 <sup>(+9.3%)</sup>	41.46 <sup>(+0.7%)</sup>	50.61 <sup>(+0.9%)</sup>	48.39 <sup>(+5.1%)</sup>	31.22 <sup>(-1.3%)</sup>	44.31 <sup>(+1.6%)</sup>
Tencent	n/a	n/a	n/a	n/a	29.69 <sup>(-6.2%)</sup>	46.06 <sup>(+5.6%)</sup>
ChatGPT	43.71 <sup>(-2.9%)</sup>	38.87 <sup>(-5.5%)</sup>	44.95 <sup>(-10.3%)</sup>	24.85 <sup>(-46.0%)</sup>	24.73 <sup>(-21.8%)</sup>	38.27 <sup>(-12.1%)</sup>

System	De-Zh		Ro-Zh		De-Ro	
	⇒	⇐	⇒	⇐	⇒	⇐
Google	38.71	21.68	39.05	25.59	33.31	32.27
DeepL	40.46 <sup>(+4.5%)</sup>	22.82 <sup>(+5.2%)</sup>	38.95 <sup>(-0.2%)</sup>	25.39 <sup>(-0.7%)</sup>	35.19 <sup>(+5.6%)</sup>	34.27 <sup>(+6.1%)</sup>
Tencent	40.66 <sup>(+5.0%)</sup>	19.44 <sup>(-10.3%)</sup>	n/a	n/a	n/a	n/a
ChatGPT	34.46 <sup>(-10.9%)</sup>	19.80 <sup>(-8.6%)</sup>	30.84 <sup>(-21.0%)</sup>	19.17 <sup>(-25.0%)</sup>	33.38 <sup>(+0.2%)</sup>	29.89 <sup>(-7.3%)</sup>

BLEU Scores; De = German, Ro = Romanian, ZH = Chinese



# Translation: Leveraging LLMs

Is ChatGPT a Good Translator? Yes with GPT-4 As The Engine (Jiao et al, 2023)

## Results

### GPT-4

- Comparable results as commercial translator products, even for distant languages

### Notes:

- First draft submitted Jan 19
- ChatGPT 3.5 updated on Jan 31 (New)
- ChatGPT 4.0 released Mar 15 (GPT-4)
- Third and Final Draft submitted Mar 19

Table 6: Performance of GPT-4 (Date: 2023.03.15) for multilingual translation.

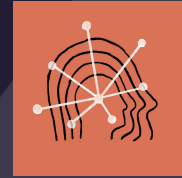
System	Zh⇒En	En⇒Zh	De⇒Zh	Ro⇒Zh
Google	31.66	43.58	38.71	39.05
DeepL	31.22	44.31	40.46	38.95
Tencent	29.69	46.06	40.66	n/a
ChatGPT (Direct)	24.73	38.27	34.46	30.84
ChatGPT (Direct <sub>new</sub> )	n/a	n/a	30.76	27.51
ChatGPT (Pivot <sub>new</sub> )	n/a	n/a	34.68	34.19
<b>GPT-4</b>	<b>28.50</b>	<b>42.50</b>	<b>38.16</b>	<b>37.84</b>

BLEU Scores; De = German, Ro = Romanian, ZH = Chinese



# Item Translation Discovery

Insights from the Workera Assessment Team



Claude 2



GPT-4

SOURCE language was always English.

TARGET languages were Bengali, Hindi, Simplified Chinese, Korean, German, French, Spanish, Arabic, and Portuguese.

Items from Machine Learning, AI Reliability, Project Management, Python, Data Visualization  
Multilingual SMEs evaluated translations, along with Google Translate

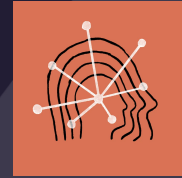
## Findings:

- ChatGPT-4 is really good at translations, and Claude 2 is nearly as good.
- Tell the LLM it is a SME in the domain of knowledge
- Tell the LLM it is fluent in both the SOURCE and TARGET languages
- Keep the prompts short (don't need to explain the tone, technical level, etc)
- For translations to/from languages where characters read right-to-left (e.g., Arabic), explicitly state that in the prompt.



# Item Translation Discovery

Insights from the Workera Assessment Team



Claude 2



GPT-4

## Findings, continued

- For programming languages: "do not translate code script reserved words from English."
- Ask that for technical terms without direct translations, include the SOURCE language term in parentheses next to the translation.
- Alternative is to provide the direct translation of technical terms in the prompt (but this is time-intensive, and not scalable)

تعرض نظام ذكاء اصطناعي للاختراق للوصول إلى معلومات خاصة عن بعض الأفراد حيث كان لدى المهاجم بعض البيانات للنظام. من خلال مراقبة مدخلات (training data) الشخصية السابقة الخاصة بالأفراد الذين تضمنتهم بيانات التدريب يمكن المهاجم من الوصول إلى المعلومات، (ML model) ومخرجات نموذج التعلم الآلي.

أي من هجمات الخصوصية التالية مذكورة في السيناريو أعلاه؟

1. عكس النموذج (Model inversion)
2. استنتاج العضوية (Membership inference)
3. هجوم صندوق أبيض (White box attack)
4. هجوم طروادة (Trojan attack)





# The Workera Assessment Team

---

**Fatima Jahara** Senior Assessment Developer

**Hassan Rebaie** Senior Assessment Developer

**Deepak Kumar** Assessment Developer

**Dipon Talukder** Assessment Developer

**MD. Asif Iqbal** Assessment Developer

**Saheli Basu** Assessment Developer

**Samah Hijazi** Assessment Developer

**Zhongtian Lin** Senior Psychometrician

**Amy Sullivan** Assessment Program Manager

**Jess Rigos** Assessment Program Manager

**Christian Lingner** Copy Editor

**Ahmed Shaikh** Assessment Developer

**Avishek Das** Assessment Developer

**Mohsin Shah** Assessment Developer

**Sharob Sinha** Assessment Developer



# **ChatGPT Demo**



# References

Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F. & Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. Harvard Business School Technology & Operations Mgt. Unit Working Paper, (24-013).

Holohan, H (September 11, 2023). *A boy saw 17 doctors over 3 years for chronic pain. ChatGPT found the diagnosis.* Today.com  
<https://www.today.com/health/mom-chatgpt-diagnosis-pain-rcna101843>

Lead IT Coach (August 10, 2023) *Navigating the World of Large Language Models: A Comparison of ChatGPT 3.5, Bard, Claude 2, and LLaMA 2.* Medium.com  
<https://leaditcoach.medium.com/navigating-the-world-of-large-language-models-a-comparison-of-chatgpt-3-5-4f1d1c4988a7>

Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. arXiv preprint arXiv:2304.02819.

Jiao, W., Wang, W., Huang, J. T., Wang, X., & Tu, Z. P. (2023). Is ChatGPT a good translator? Yes with GPT-4 as the engine. arXiv preprint arXiv:2301.08745.

McCarthy, J. (2007). What is artificial intelligence?.

Teo, K.X. (July 26, 2023). *Even OpenAI's own detection service can't tell AI-generated work apart — the company quietly took it down over accuracy concerns.* Business Insider.  
<https://www.businessinsider.com/openai-chatgpt-ai-detection-tool-shut-down-over-inaccuracy-2023-7>

WBCom Designs. *ChatGPT vs Bard vs Claude 2 vs Llama 2: The Right LLM For Every Task.* <https://wbcomdesigns.com/chatgpt-vs-bard-vs-claude-2-vs-llama-2/>

Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P., Waddington, L. (2023). Testing of detection tools for AI-generated text. arXiv preprint arXiv:2306.15666.

Williams, T (August 9, 2023). *AI text detectors aren't working. Is regulation the answer?* Times Higher Education.  
<https://www.timeshighereducation.com/news/ai-text-detectors-arent-working-regulation-answer>

Wolfe, Matt (July 17, 2023). *ChatGPT vs. Bard vs. Claude 2 (Which is best?).* YouTube.com [https://www.youtube.com/watch?v=m9Rweuhgtvk&ab\\_channel=MattWolfe](https://www.youtube.com/watch?v=m9Rweuhgtvk&ab_channel=MattWolfe)



# Additional Resources



## Podcast: The Cognitive Revolution

*"A weekly podcast where hosts Erik Torenberg and Nathan Labenz interview the builders on the edge of AI and explore the dramatic shift it will unlock in the coming years."*

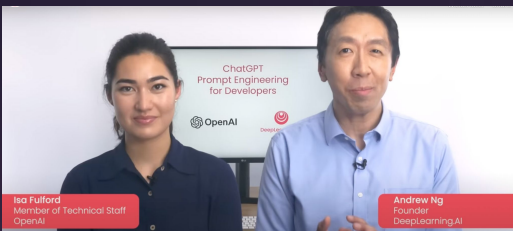
<https://www.cognitiverevolution.ai/>

## Video Series: Making Friends with Machine Learning

*"You can now enjoy all 6.5 hours of Google's legendary AI course designed to enlighten AI beginners, grow technology leaders, inform better citizens, and amuse AI experts!"*

*Presented by Cassie Kozyrkov*

[https://www.youtube.com/watch?v=1vkb7BCMqD0&ab\\_channel=CassieKozyrkov](https://www.youtube.com/watch?v=1vkb7BCMqD0&ab_channel=CassieKozyrkov)



## Course: ChatGPT Prompt Engineering for Developers

*In ChatGPT Prompt Engineering for Developers, you will learn how to use a large language model (LLM) to quickly build new and powerful applications.*

*Presented by Isa Fulford (OpenAI) and Andrew Ng (DeepLearning.AI)*

<https://www.deeplearning.ai/short-courses/chatgpt-prompt-engineering-for-developers/>



**Thank you!**

