

Generative AI

The Rise of Augmented Intelligence

Dr. Chris Asakiewicz

Outline

- Generative AI
- Augmented Intelligence
- Analytical Power of Programming
- Prompt Engineering
- Generative AI in Healthcare
- ChatGPT
- Applications
- Project Examples
- Ethical Concerns
- Future Outlook

Generative AI

Ask AI to Write About Itself (Scientific American)

1



We Asked GPT-3 to Write an Academic Paper about Itself—Then We Tried to Get It Published

An artificially intelligent first author presents many ethical questions—and could upend the publishing process

By [Almira Osmanovic Thunström](#)
June 30, 2022

COMMUNICATIONS
MACHINE LEARNING

[Almira Osmanovic Thunström](#) is an organizational developer at the Department of [ePsychiatry](#) at [Sahlgrenska University Hospital](#) in Sweden and a doctoral researcher at the Institute of Neuroscience and Physiology at Gothenburg University working on research regarding use of artificial intelligence and virtual reality in psychiatry. You can follow her at [@augmentedrobot](#).

On a rainy afternoon earlier this year, I logged into my [OpenAI](#) account and typed a simple instruction for the research company's artificial-intelligence algorithm, GPT-3: *Write an academic thesis in 500 words about GPT-3 and add scientific references and citations inside the text.*

As it started to generate text, I stood in awe. Here was novel content written in academic language, with references cited in the right places and in relation to the right context. It looked like any other introduction to a fairly good scientific publication. Given the very vague instruction I'd provided, I had meager expectations. A deep-learning algorithm, GPT-3 analyzes a vast stream of text—from books, Wikipedia, social media conversations and scientific publications—to write on command. Yet there I was, staring at the screen in amazement. The algorithm was writing an academic paper about itself. I'm a scientist who studies ways to use artificial intelligence to treat mental health concerns, and this wasn't my first experiment with GPT-3. Even so, my attempts to complete that paper and submit it to a peer-reviewed journal would open up unprecedented ethical and legal questions about publishing, as well as philosophical arguments about nonhuman authorship. Academic publishing may have to accommodate a future of AI-driven manuscripts, and the value of a human researcher's publication records may change if something [Google](#) can take credit for some of their work.


GPT-3 is well known for its ability to create humanlike text. It has written an entertaining opinion piece, produced a book of poetry and generated new content from an 18th-century author. But it dawned on

We Need to Talk About How Good A.I. Is Getting

1

New York Times, August 24th, 2022 by [Kevin Roose](#)

We're in a golden age of progress in artificial intelligence. It's time to start taking its potential and risks seriously.



This image was generated by [Google](#) app DALL-E 2 from the words "infinite joy."

For the past few days, I've been playing around with DALL-E 2, an app developed by the San Francisco company [OpenAI](#) that turns text descriptions into hyper-realistic images.

The New Chatbots Could Change the World. Can You Trust Them?

Siri, Google Search, online marketing and your child's homework will never be the same. Then there's the misinformation problem.



Aaron Margolis, a data scientist, says the new chatbots are remarkable but their answers can mix fact with fiction — just like the internet from which they [learned](#). Valerie Pliesch for The New York Times; artwork by Smriti [Keshari](#), "Altar," 2019.

By Cade Metz
Published Dec. 10, 2022 Updated Dec. 11, 2022

This month, Jeremy Howard, an artificial intelligence researcher, introduced [an online chatbot](#) called [ChatGPT](#) to his 7-year-old daughter. It had been released a few days earlier by [OpenAI](#), one of the world's most ambitious A.I. labs.

He told her to ask the experimental chatbot whatever came to mind. She asked what trigonometry was good for, where black holes came from and why chickens incubated their eggs. Each time, it answered in clear, well-punctuated prose. When she asked for a computer program that could predict the path of a ball thrown through the air, it gave her that, too.

Alarmed by A.I. Chatbots, Universities Start Revamping How They Teach

1

With the rise of the popular new chatbot [ChatGPT](#), colleges are restructuring some courses and taking preventive measures.



The University of Florida campus in Gainesville. Colleges and universities have been reluctant to ban the new chatbot because administrators doubt the move would be effective.

While grading essays for his world religions course last month, Antony [Aumann](#), a professor of philosophy at Northern Michigan University, read what he said was easily "the best paper in the class." It explored the morality of burqa bans with clean paragraphs, fitting examples and rigorous arguments.

A red flag instantly went up.


Mr. [Aumann](#) confronted his student over whether he had written the essay himself. The student confessed to [using ChatGPT](#), a chatbot that delivers information, explains concepts and generates ideas in simple sentences — and, in this case, had written the paper.

Alarmed by his discovery, Mr. [Aumann](#) decided to transform essay writing for his courses this semester. He plans to require students to write first drafts in the classroom, using browsers that monitor and restrict computer activity. In later drafts, students have to explain each revision. Mr. [Aumann](#), who may forgo essays in subsequent semesters, also plans to weave [ChatGPT](#) into lessons by asking students to evaluate the chatbot's responses.

The Rise of Generative AI

- ChatGPT: The new cutting-edge chatbot is inspiring awe, fear, stunts and attempts to circumvent its guardrails.
- Dall-E: The system lets you create digital images simply by describing what you want to see.
- GPT-3: With mind-boggling fluency, the natural-language system can write, argue and code.

Generative AI




GENERATIVE AI REFERS TO A CLASS OF AI ALGORITHMS THAT ARE DESIGNED TO GENERATE NEW DATA, SUCH AS IMAGES, MUSIC, SPEECH, OR TEXT. UNLIKE OTHER TYPES OF AI ALGORITHMS, WHICH ARE TYPICALLY USED TO CLASSIFY OR PREDICT EXISTING DATA, GENERATIVE AI CAN CREATE NEW DATA BASED ON PATTERNS LEARNED FROM A LARGE DATASET.



GENERATIVE AI ALGORITHMS USE TECHNIQUES SUCH AS DEEP LEARNING, NEURAL NETWORKS, AND PROBABILISTIC MODELS TO GENERATE NEW CONTENT. THESE ALGORITHMS ARE TYPICALLY TRAINED ON LARGE DATASETS TO LEARN THE UNDERLYING PATTERNS AND RELATIONSHIPS BETWEEN THE DATA. ONCE TRAINED, THE MODEL CAN GENERATE NEW CONTENT THAT IS SIMILAR TO THE ORIGINAL DATA, BUT NOT IDENTICAL.



GENERATIVE AI HAS MANY APPLICATIONS, SUCH AS IN ART AND DESIGN, MUSIC COMPOSITION, CHATBOTS, AND CONTENT CREATION. HOWEVER, THERE ARE ALSO CONCERNS ABOUT THE POTENTIAL MISUSE OF GENERATIVE AI, SUCH AS THE CREATION OF FAKE NEWS OR DEEP FAKES.




Augmented Intelligence

What is Augmented Intelligence and how does Generative AI enhance it?

- **Augmented Intelligence** refers to the use of AI to augment or enhance human intelligence, rather than replace it. It combines human expertise and decision-making with AI technology to improve the quality and speed of decision-making.
- **Generative AI** can enhance **Augmented Intelligence** by providing new data and insights that can be **used by humans to make better decisions**. For example, generative AI algorithms can be used to generate new scenarios or predictions based on existing data, which can be used by humans to inform their decisions.

Future X Network: Automating Knowledge Creation




23

ALCATEL-LUCENT — PROPRIETARY AND CONFIDENTIAL
COPYRIGHT © 2015 ALCATEL-LUCENT. ALL RIGHTS RESERVED.

Marcus Weldon, Bell Labs Nokia (2015)

Bell Labs
Alcatel-Lucent

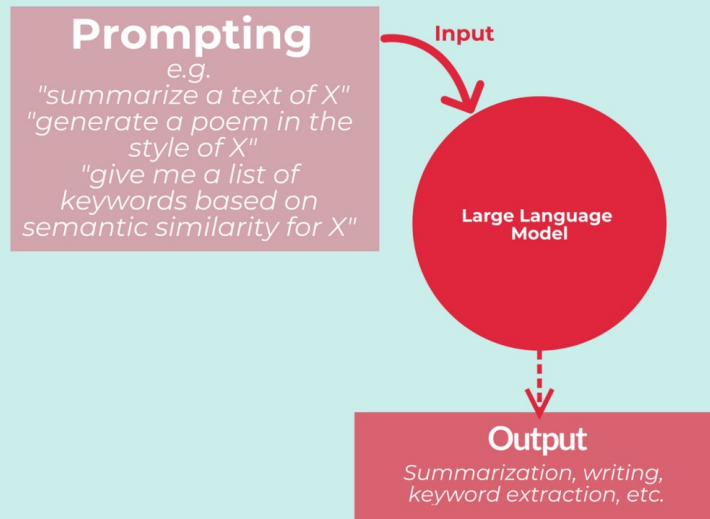
Analytical Power of Programming



Prompt Engineering

Prompt Engineering In A Nutshell

- Prompt engineering is a natural language processing (NLP) concept that involves discovering inputs that yield desirable or useful results.
- Like most processes, the quality of the inputs determines the quality of the outputs in prompt engineering. Designing effective prompts increases the likelihood that the model will return a response that is both favorable and contextual.
- Developed by OpenAI, the CLIP (Contrastive Language-Image Pre-training) model is an example of a model that utilizes prompts to classify images and captions from over 400 million image-caption pairs.



Prompt engineering is a concept in artificial intelligence, particularly natural language processing (NLP). In prompt engineering, the description of the task that the AI is supposed to accomplish is embedded in the input, e.g., as a question, instead of it being implicitly given. Prompt engineering typically works by converting one or more tasks to a prompt-based dataset and training a language with what has been called "prompt-based learning" or just "prompt learning".

Prompt Engineering

Code generation

A prompt engineer hired to create prompts to feed to an AI model to generate actual code will need at least some basic programming knowledge. Because generative AI models are not always accurate, you must understand the code generated to determine the following prompt's direction.

Output testing

Familiarity with AI and machine learning is the perfect skill set for prompt engineers tasked with testing the output generated by AI models. For these prompt engineers, the goal is to create prompts that ferret out any weaknesses in the AI output so that they can be fixed.


Text generation

Many companies hire prompt engineers to create and adapt prompts to instruct AI models to generate specific output. This output can range from copywriting to legal documents to product documentation. For these prompt engineering roles, the ability to communicate using the written word is likely far more useful than knowing what's going on under the AI model's hood.

Art generation

Prompt engineers working with AI image models like DALL-E, Stable Diffusion and Midjourney use different tactics to craft their prompts. While natural language is still the basic building block of an AI art prompt, rather than writing running prose, an art prompt employs keywords separated by commas. A photography or visual arts background can be beneficial for this type of prompt creation.

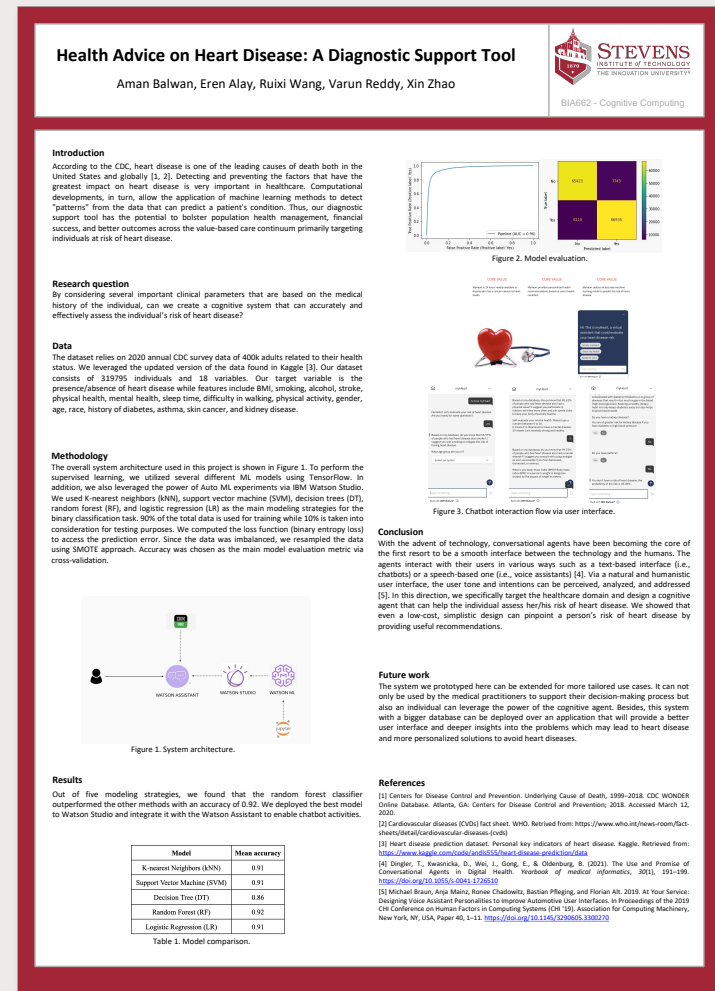
Generative AI in Healthcare



- **Virtual Assistants for Telemedicine:** Generative AI can be used to develop a virtual assistant to help patients schedule appointments, receive treatment, and manage their health information.
- **Clinical Decision Support:** Generative AI can be used to provide real-time, evidence-based recommendations to healthcare providers to improve patient outcomes.
- **Medical Record Keeping:** Generative AI can be used to generate automated summaries of patient interactions and medical histories, which can help streamline the medical record keeping process.

Generative AI in Healthcare

- **Drug information:** Generative AI can be used to provide real-time information about drugs, including side effects, interactions, and potential contraindications.
- **Medical Education:** Generative AI can provide instant access to relevant medical information and resources for students and healthcare professionals, supporting their ongoing learning and development.
- **Mental Health Support:** Generative AI can be used to provide behavioral health support to patients, including screening for mental health conditions, offering coping strategies, and connecting patients with resources for further support.
- **Medical Writing and Documentation:** Generative AI can be used to help healthcare professionals write and document medical reports, such as clinical notes and discharge summaries, by providing real-time suggestions and corrections.



ChatGPT

Chat GPT (GPT-3), the third generation Generative Pre-trained Transformer, is trained on a massive 45 TB of text data.

Large Computing

175 billion model parameters



Positional Encoding

Word is tokenized along with position



Self Attention

Modifies output sequence positions based on input positions



ChatGPT

ChatGPT is currently limited but its sophistication is exponentially growing.



Possible Hallucinations

*Eloquent answers but factually wrong
& no attributions to source*



Input Limitations

Currently limited to 4000 tokens



Weak at Complex Math

Focus on predicting next word

→ Model parameters are increasing 2X+ every 6 months

→ Expectation is that highly sophisticated integrations will be release in months, not years

ChatGPT

GPT-4



GPT-4 can accept images as inputs and generate captions, classifications, and analyses.

Input

What can I make with these ingredients?

Output

There are many options for what you can make with these ingredients! Some possibilities include:

- **Pancakes or waffles**
- **Crepes**
- **French toast**
- **Omelette or frittata**
- **Quiche**
- **Custard or pudding**
- **Cake or cupcakes**
- **Muffins or bread**
- **Cookies or biscuits**

ChatGPT

GPT-4

	Mon	Tues	Wed	Thur
	12:00 pm - 2:00 pm	1:00 pm - 3:00 pm	12:00 pm - 2:00 pm	1:00 pm - 3:00 pm
 Organizer	✓	✓	✓	✓
 John	✓		✓	
 Katelyn		✓	✓	
 Margaret	✓		✓	✓
3 Responses	2	1	3	1
Add to Calendar Email Date/Time	☐	☐	☒	☐

GPT-4 surpasses ChatGPT in its advanced reasoning capabilities.

Input

John is free from 11 am to 3 pm, Katelyn is free from noon to 2 pm and then 3:30 pm to 5 pm. Margaret is available at noon for half an hour, and then 4 pm to 6 pm. What are some options for start times for a 30 minute meeting for John, Katelyn, and Margaret?

Output

Here is a list of their availability:

- John: 11 am to 3 pm
- Margaret: noon to 2 pm, and 3:30 pm to 5 pm
- Katelyn: noon to 12:30 pm, and 4 pm to 6 pm

Based on their availability, there is a 30-minute window where all three of them are available, which is from 4 pm to 4:30 pm. So, the meeting can be scheduled at 4 pm.

ChatGPT

GPT-4

Transcript			
529 01:27:58.470 --> 01:28:16.200 siddhartha dalal: Maybe we should course, and now we would Zoom and to sit through us to	01:32:31.500 --> 01:32:39.260 siddhartha dalal: and and also this whole issue this	543 01:31:06.250 --> 01:31:24.050 siddhartha dalal: So substantial amount productivity will go up at least by I would are challenges, of course, and I think I see as you already all pointed out.	01:29:12.360 --> 01:29:29.910 siddhartha dalal: and they're not allowed to do anything other than they can go on printing anything on the computer. So they're they're also changing now. So lots of companies are now not any more trusting people to do any of those things on computer at all. So that's sort of a it's Some minor changes.
530 01:28:16.370 --> 01:28:22.080 siddhartha dalal: Really, I think lots of interesting topics.	551 01:32:39.540 --> 01:32:47.620 siddhartha dalal: guards they are building. All guard, but I see it's going to be very different	544 01:31:24.110 --> 01:31:27.440 siddhartha dalal: that all the answers are	537 01:29:29.910 --> 01:29:41.860 siddhartha dalal: obviously the issue of pleasure, and so on. Go for people worry about it. But I think that we're not going to be able to avoid it just like we can't avoid search engines. So search engines. Everybody is using it.
531 01:28:22.160 --> 01:28:25.080 siddhartha dalal: and I think th	552 01:32:47.920 --> 01:33:06.400 siddhartha dalal: guards are going to be different. If it is a child using it. If it is somebody with ce it. It may be somebody who is a political belief their own personal guards, which would be different	545 01:31:27.470 --> 01:31:41.000 siddhartha dalal: and and often they are which are going to happen. I think a new of them is related to validation of the an	538 01:29:41.860 --> 01:29:59.530 siddhartha dalal: So you, if you think about it, any company or any other place, we think that they will be able to not use it. It's just not going to happen in at least longer run. And of course the people talked about the data being a bit old. That's also just matter of time before which you everyday data would be
532 01:28:25.110 --> 01:28:44.970 siddhartha dalal: and obviously what is going to be in future in all of this within 2 years. It's just in some meaningful manner. I think	553 01:33:06.400 --> 01:33:24.050 siddhartha dalal: and filters which would be very that there might be potential industry which would about it. If your trump supporter, you will use it out and get you in your echo chamber	546 01:31:41.060 --> 01:31:51.600 siddhartha dalal: So, even, for example, in terms of coding, and so on, and obviously	539 01:29:59.530 --> 01:30:16.050 siddhartha dalal: there in some meaningful manner. I think this is going to change the world in terms of human computer interface itself. So right now, if I want to go to Expedia, and if I want to do anything, I have to go on clicking buttons, and I have to go on doing a variety of things, and so on and so forth.
533 01:28:45.030 --> 01:28:57.690 siddhartha dalal: Ted Gpt, or a things which are being created. All of those are changing the way	554 01:33:24.050 --> 01:33:42.240 siddhartha dalal: rather than continuing with the see tremendous amount of opportunities and excitement about it, and I, and of course, a big developed countries which I would have less	547 01:31:51.600 --> 01:32:05.190 siddhartha dalal: make some sort of mistake. It's in the logic, and so on, and so forth. very, very important. So somehow the the out of valuation.	540 01:30:16.050 --> 01:30:32.620 siddhartha dalal: And I think the way Expedia and others are working now already is that they got now, you interface, which is natural language interface, so you can just type it in the Avon trip from here to here on this date, then automatically, that's being converted in SQL queries
534 01:28:57.740 --> 01:28:59.250 siddhartha dalal: I think	555 01:33:43.450 --> 01:33:46.640 siddhartha dalal: complex. We we should be able	548 01:32:05.480 --> 01:32:25.450 siddhartha dalal: and the other industry. So there would be people who would be checkers which will go on checking what factual, or is that not factual, and so what	541 01:30:32.620 --> 01:30:48.310 siddhartha dalal: which are automatically created by things like Tpt or large language models in some sort of SQL queries which goes through directly to Api. So you really the human interface itself is changing all of these things which we are used to now
535 01:28:59.500 --> 01:29:12.360 siddhartha dalal: the issue is it being interviewed like, for example front of the camera, and they're	556 01:33:46.690 --> 01:33:56.680 siddhartha dalal: where there might be. A very computers or may not have use of computers	549 01:32:25.450 --> 01:32:31.260 siddhartha dalal: which would emerge what as far as I can see here itself.	542 01:30:48.310 --> 01:31:06.250 siddhartha dalal: all all this gui and the fancy guis, and so on, I think, going to change a lot in very, very meaningful manner. So I think the the biggest changes I

Recording Transcript

Summary of Key Points

Input

8 minute transcript of speaker comments

Output

Transcript Summary

The speaker discusses the potential impact of technology on the world in the next two years. He mentions that developments such as GPT and BARD are changing the way people are interviewed and how companies are processing data. The human-computer interface is evolving, and he expects the productivity of coding jobs to increase significantly. However, he acknowledges that there are challenges, such as the accuracy of AI-generated answers.

Applications

eLearning

- **Virtual Tutor:** ChatGPT can be used as a virtual tutor that is available to students 24/7. Students can ask questions related to the course material, and ChatGPT can provide answers and explanations in a conversational format.
- **Personalized Learning:** ChatGPT can be used to personalize the learning experience for each student. By analyzing student data, ChatGPT can provide customized learning paths and recommendations based on the student's strengths and weaknesses.
- **Assessment and Feedback:** ChatGPT can be used to assess student performance and provide feedback. For example, ChatGPT can grade assignments and provide feedback on areas where the student needs to improve.
- **Natural Language Processing:** ChatGPT can be used to analyze student language and provide feedback on grammar and syntax errors. This can help students to improve their writing skills and communicate more effectively.
- **Chatbot:** ChatGPT can be used as a chatbot that provides students with quick access to course information and resources. For example, students can ask ChatGPT about the course schedule, exam dates, and other relevant information.

Design Thinking

- **Ideation:** ChatGPT can be used to generate new ideas for potential solutions based on user needs and feedback. By inputting a prompt that describes a problem or user need, ChatGPT can generate a variety of potential solutions or concepts that can be further refined and developed through the design thinking process.
- **Prototyping and Testing:** ChatGPT can be used to create or refine prototypes of potential solutions. By inputting a description of a solution or design, ChatGPT can provide feedback and recommendations for how to improve or iterate on that design based on user needs and preferences.
- **User Research:** ChatGPT can be used to conduct user research and gather insights on user needs and preferences. By inputting a series of questions related to user needs or behaviors, ChatGPT can provide feedback and insights based on user language and responses.
- **Collaboration:** ChatGPT can be used to facilitate collaboration and communication between team members involved in the design thinking process. By using a conversational interface, ChatGPT can help team members share ideas and feedback more easily and efficiently.

Data Science

- **Text Classification:** You can fine-tune ChatGPT on a labeled dataset of text examples to classify new text inputs into different categories.
- **Text Generation:** You can fine-tune ChatGPT on a dataset of text examples to generate new text that is similar in style and content to the input data. Once the model is fine-tuned, you can use it to generate new text based on a given prompt.
- **Text Summarization:** You can make small changes to ChatGPT on a dataset of text examples to generate a summary of the input text. **Language Translation:** You can fine-tune ChatGPT on a dataset of translated text examples to translate new text inputs from one language to another.
- **Sentiment Analysis:** You can fine-tune ChatGPT on a dataset of labeled text examples (e.g. positive, negative, neutral) to analyze the sentiment of new text inputs.
- **Data Preprocessing:** You can use GPT to preprocess the data by cleaning, normalizing, and transforming it into a format that is suitable for further analysis.
- **NLP:** You can use GPT for various NLP tasks like Named Entity Recognition, Part of Speech Tagging, Dependency Parsing, and many more.
- **Code Generation:** You can use ChatGPT to generate code snippets for data analysis by fine-tuning it on a dataset of code examples and using it to generate new code based on a given prompt.




Applications

Use Cases

Proprietary + Confidential

Some examples from top Generative AI powered use cases



Project Examples

Navigation

- **Voice-based communication:** Generative AI can be integrated with a voice recognition system and text-to-speech software to enable blind users to ask questions or receive information via spoken interactions.
- **Navigation assistance:** Generative AI can help blind users navigate unfamiliar environments by providing directions and guidance via a voice-based system or text-based chat.

Navigation System for Visually Impaired

Parth sharma, Spencer Deudicibus, Marquee Wright, Hemanth Dadi



Introduction

Visually impaired individuals face numerous challenges when navigating a website. They may not be able to see visual cues like buttons or images, or be able to fill out forms, making it difficult to find information or navigate and interface with the internet. Our project aims to solve this problem by creating a proof of concept navigation system that can be used without the need for visual cues.

Our project aims to develop an easy-to-use and navigation system for vision impaired users that allows them to navigate a website without requiring visual aids. The purpose of this system is to ensure that users have the same level of access to online resources regardless of vision. This system is a proof of concept specifically for Canvas.

Experiment

- We will develop a proof of concept of the navigation system utilizing a trained OpenAPI model to handle the matching of user inquiries to actual classes. To scrape the data, we will use Python to grab a screenshot of the user's screen, and have it click on the location where the class the user asked for is located.
- The program will also incorporate a speech to text converter for user inquiries so it can be used without typing, and text to speech for the OpenAPI model so it can repeat it's result in audio form to confirm with the user that it has found the correct class, and to describe the page.

Results

Once the program is completed, it will be able to do the following:

- Seamlessly boot up and function without physical intervention from the user
- Interact with Canvas through user commands, clicking on classes and tabs within the page
- Read text files and lists within shells (homework, syllabi)
- Describe the Canvas layout and any graphs, pictures, or videos posted

Our proof of concept can be shipped as an application pre-installed on institute issued laptops that auto-start on launch. This way, students with visual impairments do not need to click on an application to start it, and students without visual impairments can turn it off or uninstall the program. It will utilize either voice or text commands in order to operate, which will be parsed through the OpenAPI model. Once a command is given, it will either output the results as speech or interact with the webpage for the user

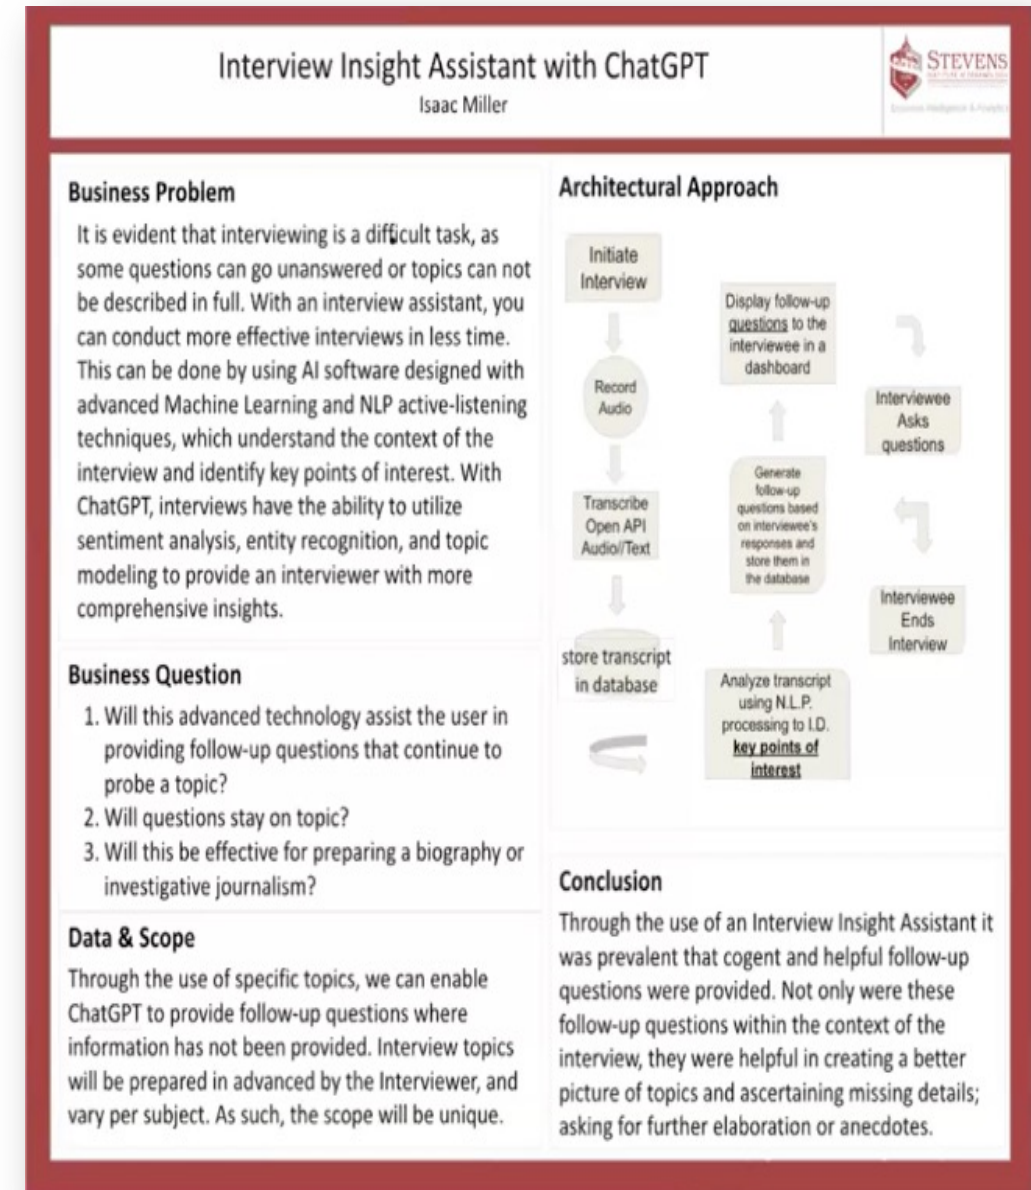
Conclusion

Our system will provide significant business value to both visually impaired students and Stevens as a whole. It allows all students to operate without any innate advantages due to their peers having visual impairment, and it fulfills some of the legal obligations for the university to accommodate students with disabilities in an automatic and reliable manner.

Project Examples

Interviewer

- **Start by asking the person to share their life story**, this can include important events, experiences, and achievements.
- **As they share their experiences, ask follow-up questions** to clarify any details and to gain a deeper understanding of their experiences.
- **Use Generative AI to summarize and highlight key points** from their story, as well as to identify any recurring themes or patterns in their life.



Project Examples

Fashion Recommender

- **Provide context:** Start by providing some context about the occasion, style, and preferences for which you need fashion recommendations. For example, you could say, "I need some outfit ideas for a business casual event next week."
- **Ask for recommendations:** Ask Generative AI for fashion recommendations based on the context you provided. For example, you could say, "Can you recommend some business casual outfits that are both comfortable and stylish?"

Product Recommender

Authors: Priyanka K Sonone, Junaid Ali Sayyed, Yoshith Kotla
Instructor: Dr. Chris Asakiewicz



Introduction


- The recommendations made by the AI-based recommendation system can help customers discover new products they may be interested in.
- It can also help them save time by finding products that they are more likely to purchase.
- Our product recommender aims at generating recommendations for similar products that users are looking for.
- We use TensorFlow's Resnet50 model to extract embeddings of product catalog images along with user uploaded image. We then find similar images using nearest neighbor algorithm and recommend similar products to user.

Application Architecture and Dataset

The application has been created using Python. We use IBM Watson assistant for our initial dialog with the user. The application enables user to upload an image of the product that they are most interested in. Our application then takes this image from user and stores in the IBM cloud s3 bucket.

We have implemented a Flask based API service (wraps a Resnet50 model along with nearest neighbor algorithm) which takes this image as input and returns 5 similar product recommendations from the product catalog. The Flask API is deployed using Ngrok server. All the product Catalog images sit in the IBM storage bucket. The Flask API acts as a connector between the AI model and IBM s3 bucket. The product recommendations are then displayed to the end user on an HTML page.

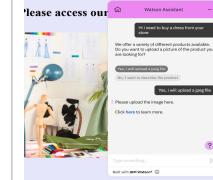
Dataset - For our product catalog sample images, we picked up our dataset from Kaggle.



Similar Image Search Logic: Extract embeddings from user uploaded image and use pre-extracted product embeddings to calculate 5 closest nearest neighbors using 'Brute force' algorithm, with Euclidean distance.



Product Recommender UI



Our chatbot greets the user and establishes a dialog.

We give user the option to provide an image of the product that they are interested in.

Our applications accepts this input image and recommends similar products from the product catalog.



Input image Similar Product Recommendations

Finding Similar Products

In the right, we have a code snippet used to implement the similarity-search logic. Feature_list and filenames are two pickle files that contain extracted embeddings and image file paths(IBM cloud storage).

Find_similar_images() function accepts an input image and recommends 5 similar products.

```
def find_similar_images(image_path):
    """
    Find similar images based on the input image path.
    """
    # Load the feature_list and filenames pickle files
    feature_list = pickle.load(open('feature_list.pkl', 'rb'))
    filenames = pickle.load(open('filenames.pkl', 'rb'))

    # Extract the input image
    img = image_path

    # Extract the input image embeddings
    img_embeddings = extract_embeddings(img)

    # Find the 5 closest nearest neighbors
    distances, indices = nearest_neighbors(feature_list, img_embeddings)

    # Extract the 5 closest nearest neighbors
    similar_images = []
    for i in range(5):
        similar_images.append(filenames[indices[i]])

    return similar_images
```

Business Impact

We envision a fully functional chatbot that enables consumers to quickly and easily navigate to products they are interested in, thus following a much-focused approach.

This application has potential to increase customer engagement along and even conversion rate and customer retention.

Even though the prototype focusses on fashion apparel recommendations, this application can be easily extended to other product-based businesses as well.

Future Scope

Implement a text-based product search and a voice-based product search.

We could also build on the application functionality to recommend complimentary products, to increase sales.

Project Examples

Identify Learning Gaps

- **Identify specific areas of weakness:** Use Generative AI to drill down into specific areas where students are struggling. For example, you could ask Generative AI to analyze test scores and identify specific questions that students are consistently getting wrong.
- **Develop targeted interventions:** Based on the identified learning gaps, develop targeted interventions to help students improve their performance. This could include additional resources such as online tutorials, one-on-one tutoring, or targeted group instruction.

AITutor

Team: Darpan Shah, Devanshi Mehta, Dhaval Patel, Ilesh Sharda,
Rishi Singh
Instructor: Chris Asakiewicz



Introduction & Background


- The emergence of AI-powered education concepts has begun generating new teaching and learning approaches.
- AI system can enhance the collaborative learning experience of students by assisting in answering queries as virtual agents, monitoring the collaboration process, and **identifying a gap in knowledge**.
- AITutor is a chatbot that uses a Multiple-Choice testing format for high school students to identify a gap in their science knowledge through quizzing the student personalizing their learning experience by recommending them questions.

Research Question

- To create a cognitive application that can tailor learning and testing for students
- Identifying and addressing a gap in student's knowledge by assessing them and recommending areas of weakness.
- Measures student's confidence post personalized learning experience to gain an understanding of the success of the application.

Model Architecture and Dataset

The application has been coded using Python and to create the user interface Tkinter has been used to provide a console based design. This design choice is simple and useful and provides the user a straightforward experience. Utilizing the OpenBookQA Dataset which is found on the Allen Institute for AI data repository consists of 5000 questions and six columns with great interest in the question stem and answer key. The application uses Artificial Intelligent Markup Language (AIML) to store the intents for the application and provide accurate responses to the user by making use of tags. The application identifies a gap in the students understanding using incorrect responses to questions and recommends provides questions the student is likely to get incorrect. At the end of the quiz the AITutor recommends areas where the student can brush up their knowledge and provides links to material.



Natural Language Processing

- The Cognitive Application access the student with five random science questions in order to identify a gap in knowledge.
- The application informs students whether the answer is correct or incorrect along with the correct answer and stores the question the student has gotten incorrect.
- Using the **bert-base-nli-mean-tokens** a pretrained model the questions stem are vectorized and stored as a numpy array.
- The pretrained model is known as sentence transformer which is able to map sentences to a 768 dense dimensional vector space.
- Chronologically the question with an incorrect responses has been used in the cosine similarity check. The top 5 questions with the highest similarity are individually returned to the user.

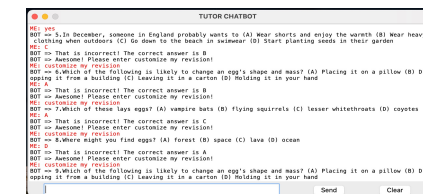
Cosine Similarity Check

- The cosine similarity check is used to retrieve a numeric distance between an a dense vectorized question that has been answered incorrectly by the user and the rest of the questions in the dataset.
- As seen in the figure below the first selected question is 'The sun is responsible for' has highest similarity with Q.107 'the sun is source of...'
- Using the sklearn cosine similarity function the application computes the cosine similarity distance between an individual question and the rest of dataset along with sorting the questions by similarity value and personalizing the quiz by returning the most similar question.

Question Stem	similarity
0 The sun is responsible for	1.000000
107 The sun is a source of which, first and foremost?	0.860960
3674 The sun is	0.860960
825 An example of how sunlight is a source of light...	0.860964
817 sunlight is a heat source emitted from	0.850033
1119 A group of bacteria move into a field and start...	0.150340
363 After a long winter's hibernation, a female bee...	0.150723
714 A pack of wolves eat the sheep before a grizzly...	0.168865
21 A mouse is being hunted by a Great Horned Owl...	0.168972
1495 A squirrel eats all of the acorns in a tree. T...	0.207894

4957 rows x 2 columns

AITutor UI



In the figure above the tutor chatbot has been able to personalize the quiz dependent on the questions the student has got wrong using pretrained sentence-transformers and cosine similarity. The AITutor is also able to extract keywords from the questions the student has gotten incorrect and return them as topics to the user along with utilizing the WikipediaAPI and returning hyperlinks that will direct the student to information about these topics. To provide the student with personal questions of their wish we have also incorporated a question generation component where the user enters a topic and a fresh new question with multiple choice answers is returned to the user.

Business Impacts

1. The business impact of this application from the anticipated results will see a functional chatbot that will be used by students studying science in high school and aid these students with their revision by identifying their gap in knowledge and improve their revision sessions.
2. This application should see students gain more confidence in their academics and take their revision and understanding of the science subject into their own hands.

Future Scope

In the future we plan to automate a dynamic question bank that mines data from web and forms MCQ like question generation functionality. This will improve our chatbot by creating fresh content and keeping the dataset up to date. Along with this we aim to improve our topic modelling so more precise recommendations can be returned.

Ethical Concerns

Addressing Ethical and Privacy Concerns with Generative AI

How can we address the ethical concerns surrounding Generative AI?


- **Bias mitigation:** Efforts can be made to address biases in training data, monitor the performance of the model, and continuously update it to minimize the impact of biases. Techniques like de-biasing and adversarial training can be used to mitigate biases.
- **Privacy protection:** To address privacy concerns, data can be anonymized, encrypted, or deleted after use, and users should be given control over their data. Additionally, transparent data privacy policies should be established and followed.
- **Fact-checking:** To minimize the spread of misinformation, fact-checking can be used to verify the accuracy of the responses generated. The inclusion of disclaimers can also help to indicate the level of confidence in the generated responses.
- **Accountability:** Establishing accountability frameworks can help to ensure that the use of Generative AI is responsible and ethical. This includes monitoring and controlling its use, as well as establishing policies and regulations to address potential misuse.
- **Education:** Educating users and developers about the potential ethical concerns associated with AI language models can help to raise awareness and promote responsible use.



Future Outlook

What does the future look like for Generative AI?

- Generative AI has already demonstrated its capabilities in natural language processing, language generation, and conversational AI. It is likely that Generative AI and other similar models will continue to advance in their abilities.
- In the future, Generative AI could be used in a wide range of applications, such as customer service, personal assistants, language translation, and education.
- However, there are also concerns about the potential misuse of such AI models, particularly in creating fake news and other forms of disinformation. It will be important to continue developing ethical guidelines and regulations to ensure that these models are used responsibly and for the benefit of society.





Christopher J. Asakiewicz, PhD.
Professor and Program Director Business Intelligence and Analytics

Director of the Stevens Alliance for Innovation and Leadership (S.A.I.L.)

Stevens Institute of Technology
School of Business
Castle Point on Hudson
Hoboken, New Jersey 07030

Email: Christopher.Asakiewicz@stevens.edu

Web: www.drchrisasakiewicz.com

Phone: (201) 216-8012

President

IT Strategy and Management Consulting, LLC
65 Millbrook Road
New Vernon, NJ 07976

Adjunct Professor

Robert Wood Johnson Medical School
University of Medicine and Dentistry of New Jersey
125 Paterson Street
New Brunswick, NJ 08901

Vice President – Retired

Pfizer, Global Business Technology
235 East 42nd Street
New York, NY, 10017

A senior management professional with strategic background in Business Technology Management, Information Technology, Consulting and Education with a 21-year history as a Vice President of Global Business Technology at Pfizer. Chris has done pioneering work in the areas of Knowledge Mining and its application within Life Sciences, most notably as a means of accelerating Translational Research. His research and teaching interests include: Cognitive Systems; Healthcare Analytics. Enterprise-level Application, Information, and Business Process Rationalization; as well as Talent, Skill, and Capability Development.

Chris' current IT strategy and management consulting engagements have been with major life sciences firms, healthcare organizations, and disease foundations, helping them in integrating major IT-enabled business processes, assessing and managing the risks associated with major change efforts, and assisting in the establishment of a biobank and collaborative environment to support rare disease research.

Over the years, his work has helped IT organizations in their effort to reduce costs, to increase productivity, and along with their business partners, to more effectively develop and deploy scalable and innovative business solutions. Before retiring from Pfizer, Chris was Vice President of Pfizer's Global Pharmaceutical Division where he led the Global Business Technology Group for Drug and Regulatory Affairs, World-wide Safety, Medical and Scientific Affairs, Clinical Data Operations, Outcomes, Clinical Development, and Product Development.

