BYU Office of Information Technology

11 May 2023

Generative AI, ChatGPT, and OIT

Considering the growing popularity and adoption of ChatGPT, Generative Artificial Intelligence (AI), and similar AI chatbot technologies, OIT is conducting a legal and technical review regarding these technologies. Additionally, a recent statement from the Church has informed affiliates and the Church workforce of similar studies.

While we understand the potential value of this technology and do not anticipate prohibiting its use long-term, there are growing concerns about data security and privacy implications with commercial Generative AI platforms and Large Language Models (LLMs).

While these reviews are underway, we strongly recommend that all students and employees exercise caution when using these platforms. Specifically, refrain from inputting proprietary or confidential information into these systems, including source code or data classified as internal, confidential, or restricted.

These reviews and discussions with campus councils will result in policies and guidelines on the secure use of these AI technologies. We ask that you refrain from incorporating this technology in any OIT services or production applications without specific approval from OIT leadership.

An Architects Roundtable discussion has been scheduled for May 25th to discuss AI and LLMs, OIT's usage, and possible direction.

A brief FAQ (see next page) with additional commentary on this topic is provided with responses for everyone and some specific to our developer/engineering community.

Tracy Flinders
VICE PRESIDENT - CIO
CES CIO

Joe Taylor
ASSISTANT VICE PRESIDENT

Scott Hunt
ASSISTANT VICE PRESIDENT

Michelle Bennett

Nick Turley

John Payne

Jon Spackman

IT LEADERSHIP COUNCIL

FOR EVERYONE

What Is Generative AI?

Generative AI refers to a class of artificial intelligence algorithms that can create content, such as text, images, audio, or video, by learning from large amounts of data. These AI models recognize patterns and structures within the data, allowing them to generate new, original outputs that closely resemble the data they have learned from.

What is GPT?

GPT, or Generative Pretrained Transformer, is an artificial intelligence model developed by OpenAI. It's designed to understand and generate human-like text. GPT is "pre-trained" because it learns from a vast amount of internet text before it is fine-tuned for specific tasks like translation, summarization, or chat response generation. Its ability to generate coherent and contextually relevant sentences makes it useful for various applications. ChatGPT is simply a chat-like interface for interacting with GPT language models.

Are security and privacy concerns specific to only ChatGPT?

No. Since November 2022, when ChatGPT was released, the hype around Generative AI has exploded, and hundreds of AI applications have been released. Many use GPT language models from OpenAI, but others do not. We have the same concerns about sending data to these commercial, 3rd party platforms, and we expect the same guidelines will be followed and to exercise caution.

I've heard about GPT-3 and GPT-4. Are they the same as ChatGPT?

GPT-3 and GPT-4 are different versions of language models provided by OpenAI. ChatGPT is one web interface for easily interacting with these language models, and Microsoft Bing search and chat also use GPT-4 under the hood.

What are some specific privacy and security concerns with this technology?

Until now, there has not been a concerted effort to understand the security and privacy implications of integrating Generative AI in our workforce. Additionally, government agencies and industry are also grappling with these questions. You've likely seen the recent news about severe concerns from the European Union with GDPR, California privacy law compliance, etc.

The technology is relatively new and has only seen commercial and consumer-grade accessibility within the last 6-12 months. The way AI companies train these models with user-provided data is often proprietary information, leaving uncertainty about how data is used and protected. We currently have no data-sharing agreements or contracted protections with these companies.

When will we receive more detailed guidance?

Numerous conversations with campus councils are underway, specifically with the Information Security & Privacy Council (ISPC). We will expect to communicate updates via Leadership Council and OIT Portfolios.

Who is leading these discussions? Will there be opportunities for input?

OIT Leadership Council, CES Technology Planning Center, and the CES IT Architecture team are leading out on these discussions. While we explore the benefits and challenges of this technology, we will provide options for input. For example, we expect to use Architects Roundtable as one such avenue to collect feedback.

If I have more questions, who can I contact?

- For privacy-related concerns, please contact Howard Loos (Chief Privacy Officer)
- For security-related concerns, please contact John Payne (Chief Information Security Officer)
- For matters related to AI integrations, technology stacks, and services, please contact Nick Turley (CES IT Architect)

FOR DEVELOPERS / ENGINEERS

What about APIs/integration frameworks that use language models?

Along with the explosion of AI applications, APIs and integration frameworks interacting with language models have grown (e.g., LangChain). These frameworks interact by sending data to language models and 3rd party services via "prompts" or embeddings to produce a result. The same data security and privacy guidelines apply, and OIT has yet to approve the usage of AI frameworks in any production services.

What are AI embeddings?

In simplified terms, embeddings provide a way to increase the "memory" and context of language models. They are a popular alternative to training a model (which can be very expensive and time-consuming). These embeddings are used to convert text information into a format that AI algorithms can process. An example would be expanding the memory of a model to understand OIT Knowledge Base articles and respond with that specific information. These embeddings are created by sending data to a service such as OpenAI. The generation of embeddings has similar data security and privacy concerns, and OIT has yet to approve generating AI embeddings with BYU/OIT information.

What are vector stores?

With the popularity of AI embeddings, users require a way to store embeddings and retrieve them for future use. These are often called "vector stores" or vector databases. There are numerous 3rd-party commercial services specializing in the storage and analysis of these vector databases to use with AI (e.g., Chroma, Deep Lake). It's important to remember these services can potentially expose data that needs to remain secure and private, and OIT has not yet authorized using these services with BYU/OIT information.