OPE.N

TECHNOLOGY SOLUTIONS

ChatGPT and Large Language Models

A Business Intuition for Credit Unions

03.09.2023 | version 1.0

Chris Kramer
OTS Innovation Lead, Data Scientist

OPEN
TECHNOLOGY SOLUTIONS

Credit unions formed Open Technology Solutions, LLC (OTS), a Credit Union Service Organization (CUSO), focused on providing collaborative, technology-related services to our partners. Together, we have achieved the kind of scale none could have achieved on their own – more than one million members served. Currently, OTS proudly serves these members and over 2,100 team members and $24.5B AUM across our partners.

We're a CUSO and Fintech that builds its own answers through the development of affordable technical solutions that support the goals of financial services businesses.

Contact Us

(303) 708-7140
ckramer@open-techs.com
open-techs.com

Table of contents

Introduction

If you don't live under a rock, you will have seen in recent news leaps-and-bounds progress in AI's ability to achieve near human-like language generation and intelligence. Products like ChatGPT are pushing the bounds of Human-AI interaction into what was once the territory of science fiction.

This technology will, someday, impact not only surface-level use cases like member-facing chatbots, but also fields like content-creation, legal review, business guidance, and team member coaching.

This whitepaper is intended to provide a comprehensive overview of large language models for a business audience in the banking technology organization. The paper aims to give an understanding of the competitive landscape, a grasp of AI terminology, and provide insight into the potential impact of AI chat in the banking industry, along with its business and security risks. The focus of this paper is not on the math and science behind large language models, but rather on their practical applications and implications for the banking technology sector.

What is ChatGPT?

ChatGPT is a large language model (LLM) product developed by OpenAI, originally a non-profit AI research and deployment company co-founded by Elon Musk. ChatGPT is now for-profit with a social impact mission, and Elon Musk is no longer directly associated with the company.

LLMs are built using deep neural networks (DNNs), a concept in AI/machine learning. The underlying LLM that powers ChatGPT is called GPT (which stands for Generative Pre-trained Transformer).

While ChatGPT has gained recent popularity for its free user interface, there several other competitors in the market, including JasperChat and ChatSonic*. JasperChat is a privately-owned AI chat utility powered by GPT-3.5 and other LLMs, while ChatSonic is a privately-owned non-GPT-3 AI chat alternative that focuses on quick responses and the generation of AI art.

**This is not an exhaustive list as new AI startups are moving into the market rapidly*

Figure 1: OpenAI logo

What is Generative AI?

Generative AI refers to a subfield of artificial intelligence that focuses on creating new and original content, rather than just analyzing and recognizing existing content. Generative AI models use a set of rules, patterns, and probability distributions to generate outputs that are similar to a given input. For example, a generative AI model trained on images of faces might generate new, synthetic faces that are similar in appearance to the ones it was trained on.

There are various approaches to generative AI, including generative adversarial networks (GANs), variational autoencoders (VAEs), and transformer-based models. The choice of approach depends on the type of data being generated and the desired outcome.

Generative AI has many potential applications, including data augmentation, content creation, style transfer, and more. It is a rapidly growing area of AI research and development, with new breakthroughs being made all the time.

GPT is an example of a transformer-based generative model, which uses a self-attention mechanism to generate text. This approach has proven to be highly effective for a variety of language-related tasks, including text generation, language translation, and question-answering.

This sub-section was written using generative AI. Could you tell?

Embracing AI

A prohibitive stance towards a democratized technology will only result in failure. Failure to secure our intellectual property, the trust of our partners, and the risk of exposing member PII to bad actors.

Members, internal team members, and partner business audiences alike will be exposing themselves to this technology and demanding that we understand it, support them in adopting it, and potentially even deploying it.

It is important then that technologists understand this new technology so that we can have thoughtful, empathetic conversations with our Credit Union partners.

Additionally, this technology has the potential to drive internal innovation as a force multiplier. A document which once took hours to write, may be generated in minutes by AI. It is therefore important to understand the mechanisms by which this technology works, and thereby the risks to better protect our business and our members.

What can large language models do?

When used thoughtfully, large language models (LLMs) have the potential to revolutionize organizations beyond member-facing chatbot use cases.

Member-facing Chatbot

LLMs expand the capabilities of member-facing chatbots beyond their current limitations. Intents and response management are restrictive methodologies that are less applicable to LLMs. Instead of managing a limited set of pre-determined chatbot responses, LLM AI can understand dialogue, retrieve answers, and respond in real-time without relying on an intents library.

However, this newfound freedom requires careful management of chatbot prompting to prevent personally identifiable information (PII) leakage. This process, known as prompt optimization, involves setting user-opaque rules that define how the chatbot may respond and what tone it should use. More details on prompt optimization can be found in an upcoming whitepaper.

AIOps

LLMs have tremendous potential to transform operations through AI Operations (AIOps), the process of applying AI to operational processes to improve efficiencies. Some of the use cases for LLMs in AIOps include:

- Understanding and analyzing data
- Contract review
- Meeting summarization
- Note-taking
- Code generation
- Technical writing
- Translating technical documentation to business requirements (and vice versa)
- Translating English language to coding statements (such as SQL, NoSQL, Python, C++, etc.)

Additionally, an LLM can act as a general virtual assistant, providing advice and feedback.

Implementation and Integration

The implementation of an LLM into a business requires careful planning and consideration of both technical and operational challenges. The first step in this process is defining a clear business case, including accuracy goals that align with the business' risk tolerance.

The integration of AI models with enterprise systems, such as databases, CRM tools, and websites, can pose technical challenges. Modern AI deployment platforms can help simplify this process using API microservices. It's also important to have mechanisms in place to ensure the technical resilience of the AI model, monitor for data drift and fairness, enable regulatory auditing, and provide business reporting capabilities.

Additionally, it's crucial to educate end users on the capabilities and risks of the AI model and integrate AI into existing operational processes. This process of evangelization can be as important as the technical implementation and helps to ensure a successful adoption of the technology within the organization.

Business Risks

As the banking industry embraces digital transformation, large language machine learning models (LLMs) are becoming an increasingly popular tool for optimizing business operations and enhancing customer experience. However, the use of LLMs also introduces several risks that banks must consider when implementing these technologies.

Example Risks

Brand Reputation & Legal Risk: A poorly designed member-facing LLM may respond to a member with a rude remark or say something disparaging towards a protected class, resulting in reputational damage and legal liabilities for the bank and its partners.

Correctness: LLMs can struggle with correctness when the underlying instructions are unclear or the prompts require complex logic. Inaccurate assessments from LLMs used in business decisions can lead to missed deadlines or poor project quality.

PII Leakage: Member-facing LLMs may inadvertently reveal sensitive information, such as personally identifiable information (PII) or business confidential data. This can result in regulatory breaches, fines, and potential identity theft.

Bias: Humans have psychological bias which is reflected in AI, since AI is trained on human-generated data. As such, there is a risk that human biases can appear in the output of AI models. Should that bias effect a decision impacting a member, the Credit Union can be held accountable. Consequently, the fairness of AI model output must be monitored and validated.

To mitigate these risks, banks should implement thorough testing and validation processes, use explainable AI techniques to increase transparency and accountability, and carefully consider the type of information being fed to the LLMs. It is also important to recognize that publicly available LLMs may reserve the right to consume user-entered prompts when retraining AI models, and any information provided to them should be considered public.

Business Risks - Conclusion

Banks must balance the business value of using LLMs with the potential risks and ensure that appropriate safeguards are in place to protect sensitive data and prevent errors. With careful planning and management, LLMs can offer significant benefits to the banking industry.

Information and Cyber Security

Large language model-based chatbots have become an essential tool for many organizations, allowing them to provide efficient customer service and streamline business processes. However, with this increased use of chatbots and large language models comes a heightened risk of security threats. In this section, we'll explore the top security best practices for chatbot usage, as well as the top three threats and how to defend against them.

Security Best Practices

When using LLM-based chatbots, it's important to follow these security best practices to ensure the safety of sensitive information and data exchange:

- Provide a clear notice to the user about what data exchange is expected during the chatbot session.
- Document expected data types and classifications.
- Confirm the method of collecting, transporting, and storing chatbot communications adheres to security requirements.
- Obtain security approval before exchanging data with chatbot or third-party vendors.
- Authenticate and authorize users prior to returning sensitive information.
- Encrypt the communication of data over the network.
- Develop whitelisting methods to only collect expected data during a chatbot conversation.

Top Three Threats

Cross site scripting attack: This attack occurs when an attacker creates a custom URL link to the chatbot website that contains malicious code designed to steal a victim's session cookie information. To defend against this threat, chatbot designers should verify and sanitize user inputs to mitigate unexpected malicious inputs.

SQL injection attack: In this type of attack, an attacker inputs malicious queries within the chatbot to execute on the backend database, potentially retrieving unauthorized data or executing malicious actions. To defend against this threat, designers should check for regular expression inputs entered by the user to ensure they are not accepted and executed on the backend chatbot database.

Denial of Service attack: Denial of service attacks attempt to exhaust the available resources of the chatbot, making it unavailable to legitimate users. To defend against this threat, designers should manage the platform utilized for hosting the chatbot to mitigate the risk of denial-of-service attacks.

Information and Cyber Security - Conclusion

By following these security best practices and implementing the recommended defenses against the top three threats, organizations can ensure the safety of sensitive information and data exchange when using chatbots and large language models.

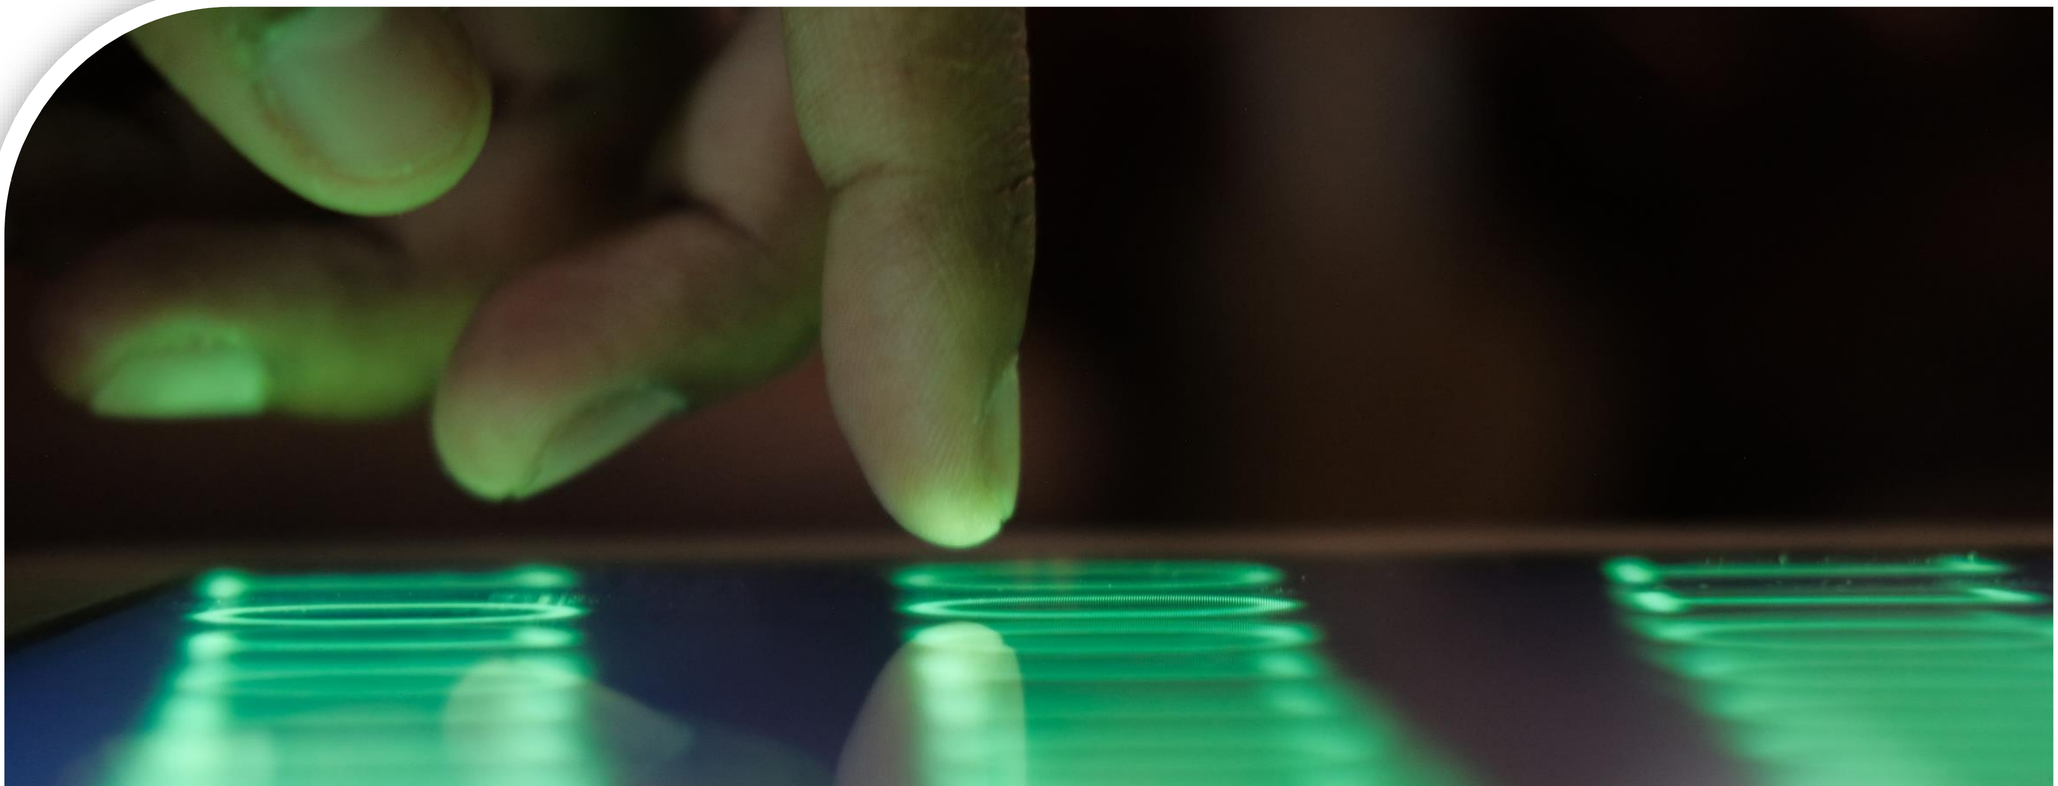

Large Language Models in the Wild

Large language models are already in use across the banking industry, with Credit Unions tagging along as fast followers. With the advent of even larger models like GPT-3 and beyond, LLMs are poised to radically change the member experience.

Here are some of the use cases that are already live within the industry*:

Alternative Credit Decisioning: Large language models are in use as part of larger AI orchestration products which use consumer behaviors to provide alternative credit decisioning. Little to no credit history? Download our app and we'll monitor your texts, social media, and phone usage patterns to determine your loan worthiness!

Personalized Financial Advice: AI can be used to find consumer patterns in large volumes of data and then fed to an LLM to make personalized financial recommendations for banking consumers.

Fraud Detection: LLMs are being leveraged at call centers to monitor member interactions to determine whether fraudulent activity is taking place. By understanding speech patterns based on known fraudulent and non-fraudulent activities, LLMs can be deployed live to catch fraudulent activity on-the-fly.

Compliance and Risk Management: LLMs can be used to review legal documents and ensure that they comply with regulatory requirements. For example, an LLM can analyze loan documents to ensure that they comply with anti-discrimination laws or review legal contracts to identify potential risks and liabilities.

Chatbots: LLMs are being used to power member-facing chatbots that can understand natural language and provide personalized support to customers. These chatbots can handle a wide range of queries, from account balance inquiries to loan applications, and can provide 24/7 support to customers.

**Each of these applications of LLMs are hosted privately (internal servers/secure VPC) with extreme Information Security and risk vetting. None of these use cases would be acceptable for adaptation with the public version of ChatGPT.*

Conclusion

Awareness is but the first step in the journey towards realized business value from AI and large language models.

To achieve success with LLMs, credit unions should carefully plan and manage implementation. Starting with careful AI governance and MLOps programs through development and deployment, it is imperative that credit unions have full control over their AI ecosystem as to avoid various business, legal, and technology risks.

AI should also be implemented using modern explainability and fairness techniques to ensure a predictable experience for all consumers.

Despite these challenges, LLMs present significant benefit for financial institutions and members.

By understanding the competitive landscape, AI terminology, and the potential impact of LLMs, credit unions can make informed decisions about their adoption of this technology.