POTENTIAL IMPACT OF GENERATIVE ARTIFICIAL INTELLIGENCE(AI) ON THE FINANCIAL INDUSTRY

Suman Kalia

Department of Computer Science and Information Sciences,
Saint Peter's University, Jersey City, NJ

ABSTRACT

Presently, generative AI has taken center stage in the news media, educational institutions, and the world at large. Machine learning has been a decades-old phenomenon, with little exposure to the average person until very recently. In the natural world, the oldest and best example of a "generative" model is the human being - one can close one's eyes and imagine several plausible different endings to one's favorite TV show. This paper focuses on the impact of generative and machine learning AI on the financial industry. Although generative AI is an amazing tool for a discriminant user, it also challenges us to think critically about the ethical implications and societal impact of these powerful technologies on the financial industry. It requires ethical considerations to guide decision-making, mitigate risks, and ensure that generative AI is developed and used to align with ethical principles, social values, and in the best interests of communities.

KEYWORDS

Generative AI, Machine Learning, Deep Learning, Financial Industry

1. INTRODUCTION

Recent advancements in the deep learning and availability of huge market data have added amazing improvements to the AI field. Ultra modern generative algorithms and techniques have expanded the capabilities of the financial industry to put forward additional and innovative products to the customer, enhance their research mechanisms and protect their investments with much greater granularity. Of course, there are drawbacks to these capabilities in the financial industry. Advancements in generative AI continue to introduce new algorithms and techniques, expanding the possibilities for generating realistic and diverse outputs.

2. METHODOLOGY

Methodologies used in preparing this paper were captured by reading "A Thousand Brains" by Jeff Hawkins and "Machine Learning: A Probabilistic Perspective" by Kevin P. Murphy. In "A Thousand Brains" a theory proposed by Jeff Hawkins suggests that the human brain operates with multiple "mini-brains" that work in parallel to process information. Each mini-brain (or cortical column) processes sensory inputs independently and contributes to overall perception and cognition. "Machine Learning: A Probabilistic Perspective" by Kevin P. Murphy is a comprehensive textbook that provides the details of machine learning algorithms. It covers generative models like Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs). Both books provide insight into the characteristics and workings of parallel processing. Once a discriminative theoretical understanding of neural networks was obtained, thirty-one (31) specific questions related to the financial industry were posed to OpenAI's ChatGPT. These questions have been listed in the Appendix.

Dhinaharan Nagamalai (Eds): NLAII, VISCOM, CCSITA, AIDD, BCYIoT -2023
pp. 37-51, 2023. IJCI – 2023 DOI:10.5121/ijci.2023.120604

In the construction of this article, graphical images were sought from the internet and leveraged to differentiate discriminative and generative algorithms in order to provide clarifications on machine learning, deep learning, and neural networks.

3. HIGH-LEVEL AI CONCEPTS

Two major generative AI model families stand out and deserve special attention:

- Generative Adversarial Networks, and
- Variational Autoencoders

Generative Adversarial Networks (GANs) are widely used in generative AI. They consist of two neural networks: a generator and a discriminator. The generator aims to create realistic outputs (e.g., images, text, or audio), while the discriminator evaluates and distinguishes between the generated outputs and real data. The two networks compete and learn from each other, leading to the generation of increasingly realistic outputs.

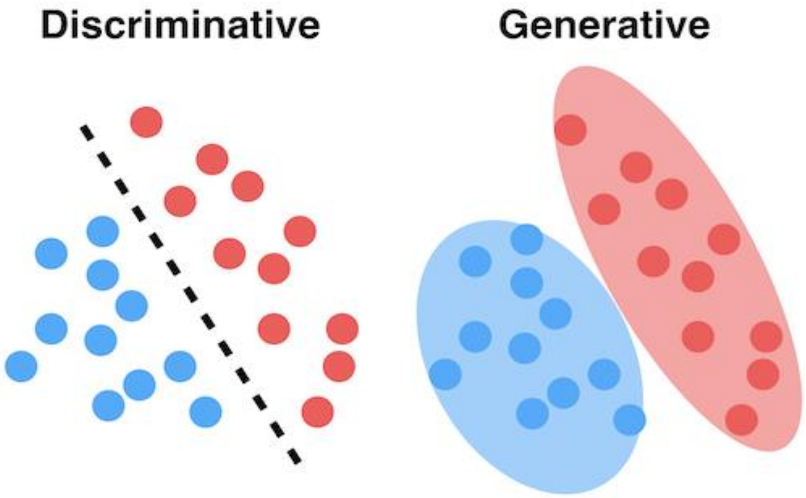

Figure 1. Two Opposing Neural Networks

Source: https://www.analyticsvidhya.com/blog/2021/07/deep-understanding-of-discriminative-and-generative-models-in-machine-learning/

Variational Autoencoders (VAEs) are another popular algorithm used in generative AI. They combine elements of both generative and discriminative models. VAEs are capable of encoding and decoding data, allowing them to generate new samples based on learned latent representations. They are often used for tasks like image generation, data compression, and representation learning.

Any system that employs many layers to learn high level representations of the input data is also a form of deep learning. A deep neural network consists of a series of stacked layers. Each layer contains units, which are connected to the previous layer's units through a set of weights. By stacking layers, the units in each subsequent layer can represent increasingly sophisticated aspects of the original input.

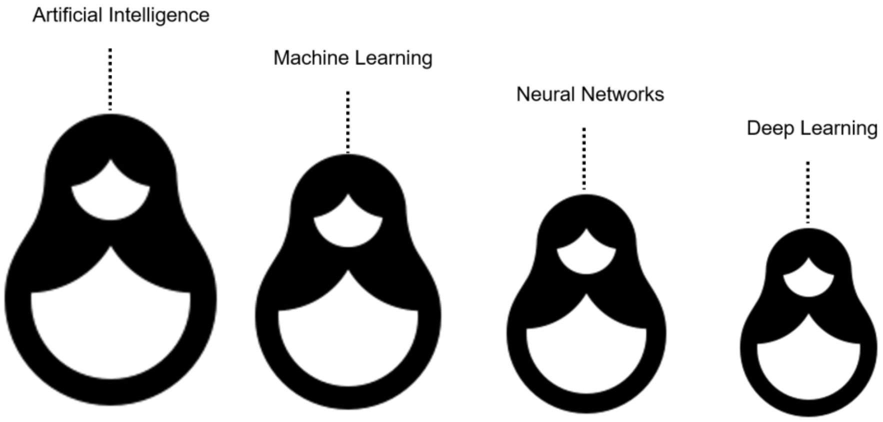

Figure 2 – AI and its subsets

Source: https://www.ibm.com/cloud/blog/ai-vs-machine-learning-vs-deep-learning-vs-neural-networks

The magic of deep neural networks lies in finding the set of weights for each layer that results in the most accurate predictions. The process of finding these weights is what we mean by “training the network”. The error in the prediction is then propagated backward through the network, adjusting each set of weights a small amount towards the direction that improves the prediction most significantly. This process is appropriately called backpropagation. Gradually, each unit becomes skilled at identifying a particular feature that ultimately helps the network to make better predictions. Deep neural networks can have any number of middle or hidden layers. ResNet, a model designed for image recognition, contains 152 layers, each of which significantly increases the accuracy of the prediction (Singhai, 2020).

4. DIFFERENCES BETWEEN PREDICTIVE AND GENERATIVE AI

The key difference between predictive AI and generative AI lies in their objectives. Predictive AI focuses on making accurate predictions or forecasts based on historical data, while generative AI focuses on generating new data or content that resembles the training dataset. Both approaches have their own applications and are used in various fields to tackle different types of problems.

The error rates in both predictive AI and generative AI can vary depending on several factors, including the complexity of the task being analyzed, the quality and diversity of the training data, the specific algorithms used, and the model's size and architecture.

In predictive AI, the error rate is typically associated with the accuracy of the predictions made by the model. Predictive models aim to minimize prediction errors by training on relevant data and employing suitable algorithms. However, the accuracy of predictions can be influenced by various factors, such as data quality, feature selection, model assumptions, and the inherent complexity of the problem being addressed.

In generative AI, the error rate is more related to the quality and fidelity of the generated outputs. Generative models strive to produce outputs that resemble the patterns and characteristics of the training data. It's important to note that error rates can vary significantly depending on the specific application, the sophistication of the models, and the level of refinement achieved through training and fine-tuning. Evaluating the performance and error rates of AI models requires careful validation, testing against ground truth or expert judgment, and monitoring the models' behavior against real-world applications.

Ultimately, the error rates of both predictive AI and generative AI are influenced by multiple factors and should be assessed in the context of the specific use case and application requirements.

5. Training a Dataset for Generative AI

The training dataset for Generative AI refers to the collection of data used to train a generative model. This dataset serves as the input for the model during the training process, enabling it to learn and generate new data that resembles the patterns and characteristics present in the training data.

Suppose we have a dataset with 100 painting works of Picasso. We wish to create a new, original painting (which looks like another great work of Picasso) based on the 100 paintings from our database, which will be our "dataset". The AI model should be able to create this new, original work based on the existing works provided, as the system has learned the rules and modalities of Picasso's earlier works. This dataset with examples of Picasso's works will be known as the "training data". One of Picasso's paintings would be considered a "data point" within this dataset, and examining a data point may yield much information, such as alignment of lines, shapes, colors, textures. This information within a data point would be referred to as "observations". Data points may contain one to many observations.

Each observation consists of many "features", which are usually the individual pixel values. Our objective is to build an AI model which can generate new features from the data in the Training dataset that look as if they have been created by Picasso. Foster advises that this can be quite difficult, "considering the vast number of ways that individual pixel values can be assigned, and the relatively tiny number of such arrangements that constitute an image of the entity we are trying to simulate" (Foster, 2019). In recent years, much emphasis has been placed on training discriminative deep learning models to produce information comparable to human performance, in a variety of image or text classification tasks.

The real power of deep learning regarding generative modeling comes from its ability to work with unstructured data, which is information that is not aligned in a predictable manner. Most often, we want to generate unstructured data, such as new images or original strings of text, which is why deep learning has had such a profound impact on the field.

6. Generative AI Model

As mentioned earlier, a prime example of a generative model in the natural world is our own brain, with the ability to create new concepts, images, and narratives. Even something simple as a dream is essentially our own brain taking information and data we have experienced and rearranging it in a creative and original way. According to Foster,

Generative modeling is usually performed with an unlabeled dataset (that is, as a form of unsupervised learning), though it can also be applied to a labeled dataset to learn how to generate observations from each distinct class . . . We need to train a generative model, which can output sets of pixels that have a high chance of belonging to the original training dataset. (Foster, 2019)

Generative models are often more difficult to evaluate, especially when the quality of the output is largely based on a combination of parameters of the information that has been provided, which may result in a different output than what was reflected in the original training data.

So far, we have discussed how AI works. Based on feedback obtained from OpenAI, the following types of training data can potentially be used for financial industry data modeling.

6.1. Financial Time Series Data

For tasks such as financial market prediction or risk analysis, the training dataset may consist of historical financial time series data. This data could include stock prices, market indices, exchange rates, interest rates, or other relevant financial metrics collected over a specific period.

6.2. Transactional Data

In the context of fraud detection or anomaly detection in financial transactions, the training dataset may comprise a large collection of transactional data. This data would typically include details such as transaction amounts, timestamps, transaction types, customer IDs, and other relevant features.

6.3. News and Market Sentiment Data

Generative models in financial industry may utilize news articles, press releases, social media data, or market sentiment indicators as part of the training dataset. This data can help capture the impact of news events, market sentiment, or public opinions on financial markets and investor behavior.

6.4. Regulatory Data and Reports

Compliance-related generative models may use regulatory documents, financial reports, or compliance guidelines as training data. This can assist in automating tasks such as document classification, risk assessment, or regulatory compliance checks.

6.5. Customer Data

In applications related to customer behavior analysis, customer segmentation, or personalized financial recommendations, the training dataset may include customer profiles, demographic data, transaction histories, and other customer-related information.

Keep in mind that this is by no means a comprehensive list of Training Data. Above all, context and accuracy are essential in determining the proper data to use to train financial models. Make sure to use relevant data to ensure a more accurate output.

7. How to Identify an AI-Generated Model

Several tools and techniques can be used to identify AI-generated output, although it is important to note that the effectiveness of these methods can vary depending on the sophistication of the AI model and the specific techniques used. Here are a few approaches that can be employed.

7.1. Metadata Analysis

AI-generated outputs may contain certain metadata or artifacts that can provide clues about their origin. This includes information about the model used, timestamps, or other metadata associated with the generated content. Analyzing this metadata can help identify whether the content was likely generated by an AI model.

7.2. Reverse Image/Text Search

Performing a reverse image search or reverse text search using search engines like Google can help identify instances where the same content has been generated or used elsewhere. If the generated output is identical or similar to other known instances, it may suggest that it was created by an AI model.

7.3. Adversarial Testing

Adversarial testing involves subjecting the generated output to specific tests or challenges designed to detect AI-generated content. This can include asking questions that require common sense reasoning or human-level understanding, as AI models may struggle with certain aspects of human cognition.

7.4. Statistical Analysis

AI-generated content may exhibit statistical patterns that differ from naturally occurring data. Analyzing features such as word frequency, sentence structure, or pixel-level patterns can help identify anomalies that suggest the content is AI-generated.

7.5. Human Expertise

Human experts who are familiar with the characteristics and limitations of AI-generated content can often recognize patterns, inconsistencies, or subtle cues that indicate the content is machine-generated. Expert judgment and contextual understanding play a crucial role in identifying AI-generated outputs.

Now that we have covered what generative AI is, how it works, and how a model can be trained and identified, let’s examine some common and known potential technical issues with AI usage in the financial industry, and explore some best practices that can be used to mitigate risks associated with these issues.

8. Common and Known Potential Technical Issues with AI usage in the Financial Industry

Generative AI challenges us to think critically about the ethical implications and societal impact of these powerful technologies. It requires ethical considerations to guide decision-making, mitigate risks, and ensure that generative AI is developed and used in a manner that aligns with ethical principles, social values, and the best interests of individuals and communities. There are some known issues with the Generative AI Usage, including.

8.1. Semantic Errors

Generative AI models may occasionally generate responses that are semantically incorrect or do not align with the intended meaning. These errors can occur due to the model's limited understanding of context, ambiguous prompts, or the inherent challenges of natural language processing.

8.2. Factual Errors

Generative AI models may generate responses that contain factual inaccuracies. Since these models are trained on vast amounts of text data from the internet, they may inadvertently incorporate false or outdated information present in the training data.

8.3. Contextual Errors

Generative AI models may struggle to maintain consistent context throughout a conversation or fail to accurately capture nuanced details in the given prompt. This can lead to responses that appear contextually inconsistent or disconnected.

9. Best Practices for Utilizing Generative AI in the Financial Industry

Within the financial industry, when devising a solution that incorporates utilization of Generative AI, one must keep in mind several best practices that should be followed, which can serve to mitigate the common potential issues previously addressed.

9.1. Precisely Identify the Problem

Clearly define the problem or objective you want to address using generative AI. Whether it's generating synthetic data, creating realistic simulations, enhancing decision-making, or improving user experiences, a well-defined problem will guide your approach.

9.2. Data Preparation

Gather and prepare high-quality data relevant to your problem. Generative AI models rely on large and diverse datasets to learn patterns and generate meaningful outputs. Ensure your data is clean, representative, and properly formatted for training the generative AI model.

9.3. Model Training

Train your generative AI model using the prepared data. This involves setting appropriate hyper parameters, training iterations, and validation techniques. Iteratively monitor and fine-tune the model to achieve desirable results. Be mindful of computational resources and training time, as generative AI models can be computationally intensive.

9.4. Evaluation and Validation

Assess the performance and quality of the generated outputs. Use appropriate evaluation metrics specific to your problem, such as accuracy, diversity, realism, or user satisfaction. Validation helps ensure the generative AI model meets the desired objectives and aligns with the intended use case.

9.5. Iterative Improvement

Generative AI models may not produce perfect results initially. Continuously analyze the generated outputs, gather feedback, and refine the model. Consider iterative training, transfer learning, or ensemble techniques to improve the model's performance over time.

9.6. Ethical Considerations

Be mindful of ethical considerations, data privacy, and potential biases in generative AI. Ensure responsible and transparent use of generative AI, addressing concerns related to fairness, privacy, security, and potential misuse of generated outputs.

9.7. Deployment and Integration

Once the generative AI model meets the desired performance and quality standards, integrate it into your application or workflow. This may involve integrating the model into existing systems, designing user interfaces, or creating APIs for easy access and utilization.

9.8. Ongoing Maintenance and Updates

Generative AI models may require periodic updates, retraining, or fine-tuning to adapt to changing data patterns, user requirements, or emerging challenges. Regular monitoring and maintenance ensure the model's continued effectiveness and alignment with evolving needs.

Effective use of generative AI requires a combination of domain expertise, data understanding, algorithmic knowledge, and ongoing evaluation. Make sure to adapt your approach based on the specific requirements of your application and stay informed about the latest developments in generative AI techniques and best practices. Our final segment covers a broader scope of potential benefits, detriments, and considerations that arise from generative AI usage in the financial industry.

10. Potential Beneficial Impacts of Generative AI for the Financial Industry

Proper use of Generative AI software and tools can improve and optimize services offered by financial institutions as well as provide opportunities for the financial industry to identify and mitigate risks. The following are some areas that stand to benefit from leveraging generative AI.

10.1. Risk Assessment and Fraud Detection

Generative AI models can analyze large volumes of financial data to identify patterns and anomalies, helping to improve risk assessment and fraud detection. By learning from historical data, these models can generate realistic scenarios and simulate potential risks, allowing financial institutions to better assess and mitigate them and help them to develop more accurate risk assessment models, optimize investment strategies, and make data-driven decisions with improved precision.

10.2. Trading and Investment Strategies

Generative AI can assist in developing trading algorithms and investment strategies. By analyzing historical market data and generating realistic simulations, these models can assist traders in identifying profitable trading strategies and making informed investment decisions and identify patterns that humans might miss. AI can also improve algorithmic trading by incorporating real-time market information and adapting strategies dynamically. This can help traders and investors make more informed decisions, optimize portfolio management, and potentially improve returns.

10.3. Credit Assessment and Underwriting

Generative AI models can aid in credit assessment and underwriting processes. By analyzing vast amounts of data, including credit history, income, and other relevant factors, these models can generate more accurate credit risk profiles and streamline the loan approval process. This can help financial institutions make better lending decisions and reduce the risk of defaults. Generative AI can streamline credit scoring and underwriting processes by analyzing a wide range of data points and help assessing creditworthiness more accurately, offer personalized loan terms, and facilitate faster loan approvals. It enhances the efficiency and accessibility of financial services.

10.4. Portfolio Management and Asset Allocation

Generative AI can assist portfolio managers in optimizing asset allocation and risk management. By analyzing historical market data, economic indicators, and other relevant factors, these models can generate simulations and recommendations for portfolio rebalancing. This can help improve diversification, manage risk exposure, and potentially enhance returns.

10.5. Enhanced Predictive Analytics

Generative AI can assist in predictive analytics for risk management. By analyzing historical data and generating future scenarios, these models can help identify potential risks and their probabilities. This enables risk managers to make data-driven decisions, anticipate potential issues, and implement proactive risk mitigation strategies.

10.6. Stress Testing and Scenario Analysis

Generative AI models can simulate and generate a wide range of scenarios to stress test portfolios, financial systems, or specific risk factors. By generating realistic scenarios, risk managers can assess the potential impact of adverse events, market fluctuations, or other uncertainties on their portfolios or operations. This allows for a more comprehensive evaluation of risk exposure and the development of appropriate risk management strategies.

10.7. Real-time Risk Monitoring

Generative AI can monitor real-time data streams and identify potential risks or anomalies in real-time. By continuously analyzing and generating insights from streaming data, these models can provide early warnings for risk events, allowing risk managers to take immediate actions. This enables proactive risk management and reduces the impact of adverse events.

10.8. Regulatory Compliance and Anti-Money Laundering (AML)

Generative AI models can assist financial institutions in meeting regulatory requirements and detecting potential instances of money laundering or other illicit activities. By analyzing large volumes of transactions and identifying patterns, these models can flag suspicious activities for further investigation, improving compliance efforts and reducing financial crime risks. Generative AI can assist fintech companies in meeting regulatory compliance requirements. These models can analyze large volumes of data, identify potential compliance issues, and generate reports to ensure adherence to regulations. This helps fintech platforms operate within legal frameworks and maintain compliance standards.

10.9. Privacy-Preserving Analytics

Generative AI can enable privacy-preserving analytics, allowing organizations to “derive insights from sensitive data without directly accessing or exposing the underlying data” (Evangel, 2023). Techniques like federated learning and secure multi-party computation enable collaborative data analysis while maintaining data privacy.

11. POTENTIAL IMPROPER USAGE OF GENERATIVE AI

While generative AI can be leveraged to benefit the financial industry, it can also potentially be abused and produce outcomes that are malicious, unethical, illegal, or harmful in other ways. Here are some examples of improper AI usage.

11.1. Fraudulent Financial Schemes

Generative AI could be exploited to generate fake financial documents, counterfeit currency, or falsified investment reports. This can be used to deceive investors, commit financial fraud, or manipulate financial markets.

11.2. Discriminatory Practices

If generative AI models are trained on biased or discriminatory datasets, they can perpetuate and amplify societal biases. This can result in discriminatory decisions in areas such as hiring, lending, or insurance, leading to unfair treatment and inequality.

11.3. Invasion of Privacy

Generative AI can be misused to breach individuals' privacy rights by generating sensitive personal information, forging identities, or compromising security systems to access confidential data.

11.4. Weaponization and Warfare

In a worst-case scenario, generative AI could be used to develop autonomous weapons or sophisticated cyber weapons capable of inflicting harm and damage on a large scale, posing significant risks to society.

12. IMPLEMENTATION CONSIDERATIONS WITH AI USAGE IN THE FINANCIAL INDUSTRY

Generative AI has a multitude of uses and functionality within the financial industry, however, improper use may lead to ethical or even legal issues. Ethical thinking is essential in shaping regulations and governance frameworks surrounding generative AI. Policymakers and organizations need to consider the potential risks, societal impact, and ethical implications of generative AI by developing appropriate guidelines and governance standards that promote responsible and ethical use of these technologies When implementing generative AI in the financial industry, one must take into account the following considerations:

12.1. Potential Misuse of Biometric Data and Facial Recognition

Generative AI models, such as deep fakes, can generate highly realistic images or videos that mimic real individuals, potentially impacting privacy, especially in the context of biometric data and facial recognition. This raises concerns about identity theft, fake content creation, and unauthorized use of personal information. Robust privacy regulations and security measures are necessary to address these risks.

12.2. Potential Misuse of User Profiling and Targeted Advertising

Generative AI can analyze vast amounts of user data to create detailed profiles for targeted advertising. While this can improve advertising effectiveness, it raises concerns about personal privacy and the potential for invasive and manipulative practices. Striking a balance between personalized advertising and individual privacy rights is crucial.

12.3. Potential Misuse of Voice and Speech Synthesis

Generative AI can replicate human voices and generate speech that mimics real individuals. This raises concerns about voice impersonation, identity theft, and the potential misuse of synthesized speech. Adequate safeguards and regulations are necessary to protect individuals' privacy and prevent unauthorized use of voice and speech technologies.

12.4. Potential Misuse of Human-AI Collaboration

Generative AI can enhance human creativity and productivity but also poses challenges regarding the displacement of human workers. Ethical considerations involve finding the right balance between automation and human involvement, ensuring that generative AI supports human decision-making and augments capabilities rather than replacing individuals. Ethical thinking guides the responsible integration of generative AI into workflows to create a positive and inclusive impact.

12.5. Ethical Considerations

Privacy is closely linked to ethical considerations in generative AI. Organizations must ensure that the use of generative AI respects privacy principles, such as data minimization, purpose limitation, and informed consent. Transparent communication with users about data collection, usage, and potential privacy risks is essential.

12.6. Bias and Fairness

Generative AI models are trained on large datasets, which may contain biases present in the data. If not addressed, these biases can be perpetuated or amplified in the generated outputs. Ethical thinking is required to ensure fairness and mitigate biases in the development and deployment of generative AI models. It involves careful consideration of the data used, the training process, and the evaluation of model outputs to minimize biases and promote equitable outcomes.

12.7. Transparency and Explainability

Generative AI models can often be complex and difficult to interpret. Ethical thinking calls for transparency and explainability, enabling stakeholders to understand how generative AI models work and how they generate outputs. Transparent communication about the limitations, uncertainties, and potential risks associated with generative AI is crucial to build trust and allow individuals to make informed decisions.

12.8. Accountability and Liability

Generative AI raises questions about accountability and liability when harmful or unethical outputs are generated. Determining responsibility in cases where generative AI is used to create misleading or malicious content requires ethical thinking and the establishment of legal frameworks to assign accountability and address potential harms.

13. CONCLUSIONS

13.1. Developing and Training Generative AI

Generative AI requires deep neural networks, and the design of those layers is notguided by fixed rules. Adjusting the weights of the essential middle layers is critical to the success of the network, and the ability to make proper adjustments to these layers comes from experimentation and experience. Just like multiple iterations of training data help to develop the model to become stronger and more accurate, human time and effort spent making adjustments to layer weights helps build experience with developing and utilizing AI.

Foster indicates that there should also be a heavy focus on the following issues as the AI model is architected and trained, such as how a model copes with the high degree of conditional dependence between features, and how a model finds one of the tiny proportions of satisfying possible generated observations among a high-dimensional sample space (Foster, 2019).
Lastly, Foster stresses the importance of understanding the inner workings of your generative model, “since it is the middle layers of your network that capture the high-level features that you are most interested in” (Foster, 2019).

13.2. Ensuring Effectiveness

Finance applications often involve complex data patterns, time series analysis, and risk assessment, which may require sophisticated generative models like GANs, VAEs, or autoregressive models. However, the effectiveness of any algorithm depends on its implementation, tuning, and the quality and relevance of the data used. To determine the most suitable algorithm for a specific finance task, it is recommended to experiment with different algorithms, evaluate their performance against relevant metrics, and consider domain expertise to make an informed decision.

Additionally, one must consider that while generative AI can enhance risk management, it also poses challenges such as interpretability and bias in algorithmic decision-making. Risk managers need to carefully consider the limitations and assumptions of generative AI models and validate their outputs with domain expertise, as responsible and transparent implementation of generative AI is crucial for effective risk management. It's also important to evaluate and validate the generated outputs carefully, especially in critical domains such as finance, healthcare, or legal contexts, where it is of utmost importance to ensure the accuracy and reliability of information. The deployment of generative AI models should be accompanied by appropriate human oversight and consideration of potential errors or biases that may arise.

13.3. Responsibility Considerations

It is also important to note that the ethical considerations, data privacy, and potential biases associated with generative AI should be carefully addressed in the financial industry. Responsible development and use of generative AI are essential to ensure that it aligns with regulatory requirements, protects user privacy, and maintains transparency in its operations. To mitigate privacy risks associated with generative AI, organizations and policymakers should establish robust privacy frameworks, enforce data protection regulations, and promote responsible and ethical use of these technologies. According to Annelytics, "striking a balance between the benefits of generative AI and individual privacy rights will be crucial in shaping the future of this technology" (Annelytics, 2023), and tantamount for its responsible adoption. It is critical to develop responsible guidelines, ethical frameworks, and robust regulations to ensure that generative AI is used for beneficial and ethical purposes. Promoting transparency, accountability, and a strong ethical foundation is essential to prevent misuse and harmful applications of generative AI technologies.

ACKNOWLEDGEMENTS

The author would like to express his sincere gratitude to Professor Robert B. Adelson, Adjunct at Saint Peter's University for "breathing life into this document" in assisting with formatting, layout, structure, and spending countless hours discussing the endless possibilities that can be achieved by AI in the Financial Industry and beyond.

The author would also like to thank Professor Edward Moskal, Associate Professor at Saint Peter's University, Dr. Alberto LaCava, Department Chair of Computer Science & Information Systems, and Professor Mary Tedeschi, Assistant Professor of Computer Science at Pace University, NYC, for their assistance with organization of ideas used in this article.

REFERENCES

[1] AI vs. Machine Learning vs. Deep Learning vs. Neural Networks: What's the difference? IBM Blog. (2023, July 11). https://www.ibm.com/cloud/blog/ai-vs-machine-learning-vs-deep-learning-vs-neural-networks
[2] Annelytics. (2023). Unveiling the Buzz: "Exploring the Promises and Concerns of Generative AI." Medium. https://annelytics.medium.com/unveiling-the-buzz-exploring-the-promises-and-concerns-of-generative-ai-706c1e044b5e
[3] Evangel, T. (2023). Understanding Data Privacy and Security in the Analytics Era. Naija Techie. https://naijatechie.com/understanding-data-privacy-and-security-in-the-analytics-era/
[4] Foster, D. (2019). Chapter 1. Generative Modeling. In Generative deep learning. chapter, O'Reilly Media, Inc.
[5] Goyal, C. (2021). 2023's best guide to Discriminative & Generative Machine Learning Models. Analytics Vidhya. https://www.analyticsvidhya.com/blog/2021/07/deep-understanding-of-discriminative-and-generative-models-in-machine-learning/
[6] Hawkins, J., & Dawkins, R. (2021). A thousand brains: A new theory of intelligence. Basic Books.
[7] Murphy, K. P. (2012). Machine learning: A probabilistic perspective. The MIT Press.
[8] OpenAI. (2023). ChatGPT (Mar 14 version) [Large language model]. https://chat.openai.com/chat
[9] Singhai, H. (2020). When some one asks what is deep learning?. Fictionally Irrelevant. https://fictionally-irrelevant.vercel.app/posts/when-someone-asks-what-is-deep-learning
[10] Thakker, N. (2023). Importance of Data Validation: Trust but Verify. LinkedIn. https://www.linkedin.com/posts/nileshthakker_datavalidation-trustbutverify-dataintegrity-activity-7075910354685333504-Vozf

AUTHOR

Suman Kalia is an Assistant professor of computer science and cyber security at Saint Peter's University in Jersey City, NJ for the past five years. Prior to moving to academics full time, Dr. Kalia was working for JPMorgan Chase as Executive Director in the Commercial Banking area until 2017. Dr. Kalia has 35 years of experience in the financial world combined between Lehman Bros and JPMorgan. Dr. Kalia obtained his Doctorate (DPS) and MS in computer science from Pace University and obtained his MBA in finance from Rutgers University. Dr. Kalia moved to the US in November 1981.

APPENDIX

In researching this topic, thirty-one (31) questions were posed to OpenAI / ChatGPT. Responses to these questions were leveraged in assembling this documentation. Below is the list of these questions.

S.No	Question Posed to OpenAI on Various Aspects of Financial Industry
1	What is a training dataset for Generative AI?
2	What does the training data set for Fintech look like for generative AI?
3	What does the training data set for cyber security look like for generative AI?
4	What does the training data set for privacy look like for generative AI?
5	How would Generative AI impact the Finance Industry?
6	Is Generative AI good for Fintech?
7	Which generative algorithm is the best for Generative AI in Finance?
8	Which generative algorithm is the worst for Generative AI in Finance?
9	How Does Generative AI impact Risk Management?
10	How would Generative AI impact Privacy?
11	How would Generative AI impact Ethics and Ethical Thinking?
12	Which algorithm is most utilized in Generative AI?
13	What Would be the downfall of usage of Generative AI in cyber security?
14	Ethical Concerns and Fairness?
15	What is the latest version of the ChatGPT?
16	What is the latest version of the ChatGPT?
17	What is the previous version of the ChatGPT?
18	What are the capabilities enhancement between GPT 3 and GPT 3.5?
19	What is the known error rate in generative AI?
20	What is the difference between Predictive and generative AI?
21	Which has a lower error rate - Predictive or generative AI?
22	Which known companies have banned the use of Generative AI?
23	Are there any tools or techniques to identify AI generated output?
24	What is a recommendation for usage of generative AI in educational institutions?
25	What is the impact of generative cyber security?
26	What Would be the downfall of usage of Generative AI in cyber security?
27	What is a visual representation of generative adversarial networks of AI?
28	Is there a similarity between A Thousand brains and generative AI?
29	Diffusion Models vs. GANs vs. VAEs: Comparison of Deep Generative Models?
30	Summary of A Thousand Brains By Jeff Hawkins
31	What are the training sets for Cyber security in Fintech used by generative AI?