International Journal of Scientific Research in Science and Technology
Available online at : www.ijsrst.com
Print ISSN: 2395-6011 | Online ISSN: 2395-602X
doi : https://doi.org/10.32628/IJSRST2221185

Identifying and Building Generative AI Use Cases Within Enterprise Software Products

Arjun Raj Bhalla
Independent Researcher, India

ARTICLEINFO

Article History:

Accepted : 15 July 2023
Published : 22 July 2023

Publication Issue :
Volume 10, Issue 4
July-August-2023

Page Number :
723-732

ABSTRACT

Generative AI (GenAI) is revolutionizing enterprise software by enabling text generation, image synthesis, and predictive modeling, leading to enhanced user experiences, workflow automation, and new business value. However, integrating GenAI into enterprise applications requires navigating technical, operational, and ethical challenges. This paper presents a structured framework for identifying and implementing GenAI use cases across various industries, addressing considerations such as data privacy, model interpretability, and computational efficiency. By categorizing GenAI applications, outlining best practices, and detailing fine-tuning methodologies, this research provides a comprehensive guide for enterprises to leverage GenAI effectively while ensuring ethical and sustainable AI deployment.

Keywords : Generative AI (GenAI), enterprise software, AI use cases, machine learning models, transformer models, AI-driven automation, business applications, natural language processing (NLP), image generation, data augmentation, deep learning applications, ethical AI, AI governance, monetization strategies, subscription-based AI services, API-based pricing models, automation, decision-making, structured frameworks, implementation, scaling AI capabilities.

1. INTRODUCTION

Generative AI has emerged as a transformative technology in enterprise software, offering capabilities ranging from text generation to image synthesis and predictive modeling. Organizations across industries seek to integrate GenAI into their products to enhance user experience, automate workflows, and derive new business value. However, identifying and implementing viable GenAI use cases presents technical, operational, and ethical challenges.

The rapid advancement of AI models such as GPT-4, DALL-E, and BERT has facilitated new possibilities for automation and content creation (Radford et al., 2019)1. As organizations adopt these technologies, they must address key considerations such as data privacy, bias mitigation, computational efficiency, and model interpretability (Bender et al., 2021)2. This paper provides a structured framework to assist enterprises in selecting and deploying GenAI applications effectively, ensuring they align with business objectives and user expectations.

2. LITERATURE REVIEW

Existing literature on AI and enterprise software provides insights into the evolution of AI capabilities, use case identification, and implementation strategies. Studies such as those by Brown et al. (2020)3 on GPT models and Goodfellow et al. (2014)4 on Generative Adversarial Networks (GANs) highlight the technical foundations of GenAI. Research by Marcus (2021)5 discusses AI adoption challenges. Additional studies by Vaswani et al. (2017)6 on transformer models and Ramesh et al. (2021)7 on multimodal AI reinforce the expanding applications of GenAI in enterprise software. A review of prior research indicates that AI adoption in enterprises has accelerated due to advances in computing power, increased availability of data, and improved algorithmic efficiency (Chollet, 2017) 8. Organizations implementing AI in their software products have observed significant improvements in productivity, customer engagement, and automation. However, challenges such as model bias, ethical concerns, and regulatory compliance remain significant obstacles to widespread adoption (Mitchell et al., 2019)9. The literature highlights best practices in AI governance, data quality assurance, and risk mitigation strategies that enterprises must consider when deploying GenAI applications.

3. PRODUCT TYPES BY INDUSTRY AND CATEGORY

Enterprise software spans multiple industries, each with unique needs and opportunities for Generative AI (GenAI) adoption.

Tabel 1 : GenAI applications by industry

Industry	Key GenAI Applications
Media & Entertainment	AI-assisted content creation, video editing, deepfake detection, automated script writing, personalized content recommendations
Retail	Inventory management, personalized recommendations, virtual shopping assistants, automated marketing content, AI-driven chatbots, visual search optimization
Finance	Fraud detection, automated reporting, personalized financial advice, credit risk assessment, algorithmic trading, regulatory compliance automation
Manufacturing	Predictive maintenance, quality inspection, design automation, process optimization, supply chain optimization, robotics integration
Legal & Compliance	Contract analysis, document generation, legal research, compliance automation, case law research, intellectual property analysis
Human Resources & Recruitment	Resume screening, interview chatbots, employee sentiment analysis, training content generation, performance evaluation automation
Education & E-Learning	AI-generated learning materials, automated grading, personalized tutoring, content summarization, virtual teaching assistants
Real Estate	Property valuation, AI-generated property descriptions, automated lease management, predictive market analysis, customer interaction chatbots
Telecommunications	Network optimization, AI-powered customer service chatbots, predictive maintenance for infrastructure, automated billing analysis
Energy & Utilities	Smart grid optimization, predictive maintenance for energy assets, AI-powered demand forecasting, automated energy efficiency recommendations

Healthcare	Electronic health records (EHR) automation, AI-assisted diagnostics, patient engagement solutions, personalized medicine, drug discovery, medical image analysis

Below is an expanded breakdown of how GenAI can be utilized across different sectors:

3.1. **Healthcare**: The healthcare industry benefits significantly from GenAI through applications such as electronic health records (EHR) automation, AI-assisted diagnostics, patient engagement solutions, personalized medicine, and drug discovery. AI-powered tools can analyze medical images, assist in disease prediction, and even generate synthetic patient data for research purposes. Companies like IBM Watson Health and DeepMind have pioneered AI-driven healthcare solutions (Esteva et al., 2017) [10].

3.2. **Finance**: The financial sector leverages GenAI for fraud detection, automated reporting, personalized financial advice, credit risk assessment, and algorithmic trading. AI models can detect anomalies in transactions, generate regulatory compliance reports, and automate investment strategies. JP Morgan and Goldman Sachs use AI-driven models for risk management and trading automation.

3.3. **Retail**: GenAI enhances customer experience in retail by enabling inventory management, personalized recommendations, virtual shopping assistants, and automated marketing content generation. AI-driven chatbots improve customer support, while visual AI models optimize store layouts. Companies like Amazon and Shopify integrate AI for demand forecasting and dynamic pricing (Li et al., 2019) [11].

3.4. **Manufacturing**: AI-powered predictive maintenance, quality inspection, design automation, and process optimization drive efficiency in manufacturing. AI-based computer vision systems detect defects in real time, while generative design models create innovative product prototypes. Siemens and General Electric use AI-driven analytics to enhance operational efficiency (Lee et al., 2018) [12].

3.5. **Legal & Compliance**: Law firms and compliance departments utilize GenAI for contract analysis, document generation, legal research, and compliance automation. AI models analyze vast legal datasets to provide case law insights, helping firms automate routine legal tasks. Companies like Kira Systems and ROSS Intelligence offer AI-driven legal research platforms (Surden, 2019) [13].

3.6. **Human Resources & Recruitment**: GenAI streamlines HR processes, including resume screening, interview chatbots, employee sentiment analysis, and training content generation. AI tools assist HR professionals in matching candidates to job roles, reducing hiring bias, and enhancing employee engagement. Companies like HireVue and Pymetrics use AI to optimize hiring decisions (Dastin, 2018) [14].

This categorization helps align GenAI capabilities with specific industry demands, allowing organizations to identify high-value use cases that drive competitive advantages.

4. TYPES OF GENERATIVE AI USE CASES

Generative AI use cases in enterprise software can be categorized into distinct functional areas. Understanding these categories enables organizations to align GenAI capabilities with their specific business needs, improving efficiency and innovation.

Tabel 2 : GenAI use cases and examples

Use Case Type	Description	Examples
Text Generation	Produces human-like text for various applications	Chatbots, content creation, automated email responses (Zhang et al., 2020) [15]

Image Generation	Creates visuals based on text input	Product design, marketing content, medical imaging enhancements (Ramesh et al., 2022) [16]
Data Augmentation	Enhances datasets for training AI models	Synthetic data generation for machine learning (Shorten Khoshgoftaar, 2019) [17]
Code Generation	Automates software development tasks	AI-assisted coding, bug fixing, software documentation (Chen et al., 2021) [18]
Decision Support	Provides AI-driven insights	Financial forecasting, legal research, customer sentiment analysis (Agrawal et al., 2018) [19]
Voice and Speech Synthesis	Generates natural-sounding speech	Virtual assistants, audiobooks, personalized voice cloning (Oord et al., 2016) [20]
Video Generation	Creates videos from text prompts	Automated video production, digital avatars, content marketing (Ho et al., 2020) [21]

One of the most prominent use cases is **text generation**, where AI models generate human-like text based on input prompts. This is widely applied in chatbots, automated email responses, and content creation. For instance, AI-driven marketing tools use natural language generation (NLG) to craft personalized customer messages, thereby improving engagement rates.

Image generation is another powerful application where AI models create high-quality images based on textual descriptions. This is beneficial for industries such as e-commerce, where product images can be generated dynamically based on user preferences, or in healthcare, where AI enhances medical imaging for diagnostics.

Another critical category is **data augmentation**, which involves generating synthetic data to improve machine learning model performance. By creating diverse datasets, organizations can reduce biases in training data and enhance predictive accuracy, particularly in domains with limited labeled datasets.

Other key use cases include **code generation**, **decision support systems**, **voice synthesis**, and **video generation**, each contributing to increased automation, efficiency, and innovation across various industries.

5. Best Practices for Building Generative AI Use Cases

To successfully build Generative AI (GenAI) use cases, enterprises must follow best practices that ensure effective implementation, scalability, and ethical considerations.

Below are key aspects to consider:

5.1. **Data Strategy**: High-quality training data is fundamental to the success of GenAI models. Enterprises should prioritize data collection, annotation, and augmentation to ensure diverse and representative datasets. Companies like OpenAI and Google invest heavily in data curation and employ techniques such as synthetic data generation and transfer learning to enhance model performance. Maintaining data privacy and security is also crucial, especially for industries handling sensitive information such as healthcare and finance.

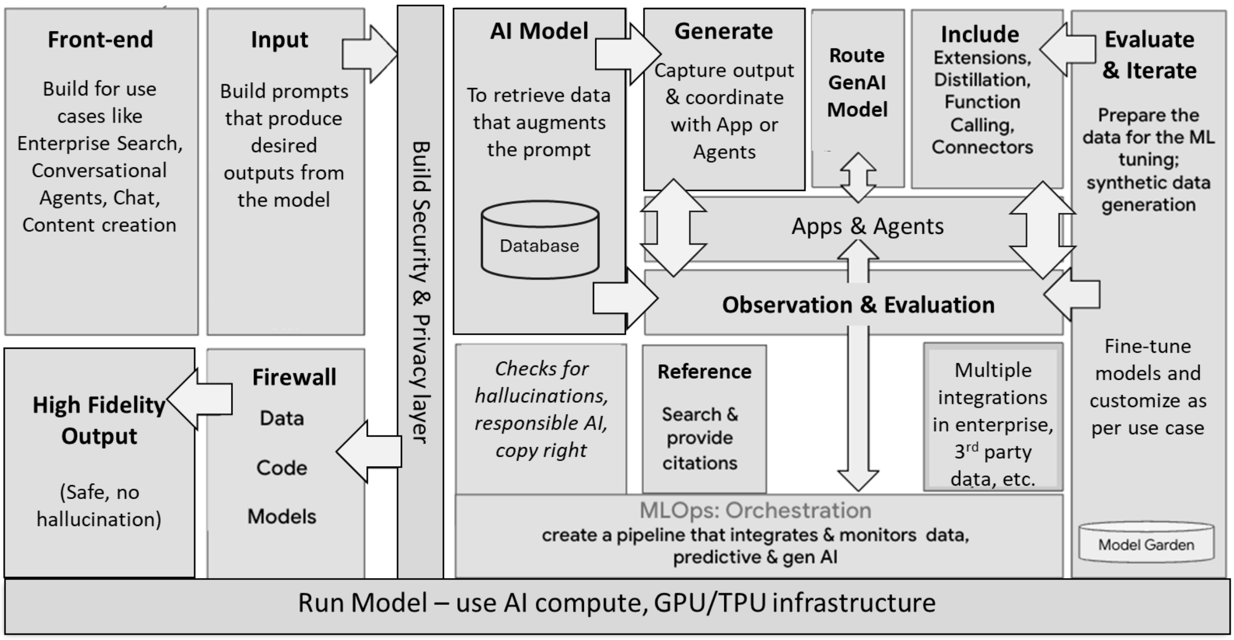

Figure 1 : GenAI Model Setup and Reference Data Architecture

5.2. **Model Selection**: Selecting the appropriate AI model depends on the specific use case. Transformer-based models like GPT are well-suited for text generation and conversational AI, while Generative Adversarial Networks (GANs) and diffusion models are preferred for image and video generation. In some cases, hybrid models combining multiple architectures provide optimal results. Enterprises must evaluate factors such as computational cost, model interpretability, and deployment feasibility before finalizing their AI approach.

Tabel 3: AI Model types, pros, cons and ideal use cases for each

Model Type	Pros	Cons	Ideal Use Cases
Transformer-based (e.g., GPT, BERT)	- Excellent for text generation & NLP tasks - High scalability - Context-aware responses	- High computational cost - Requires large datasets - Interpretability challenges	- Chatbots & virtual assistants - Content creation - Language translation - Code generation
Generative Adversarial Networks (GANs)	- High-quality image & video generation - Can learn complex patterns - Realistic synthetic data creation	- Training instability - Mode collapse risk - High data requirements	- Image synthesis & enhancement - Deepfake generation - Data augmentation
Diffusion Models (e.g., Stable Diffusion, DALL·E 2)	- Generates high-quality images - More stable training than GANs	- Slower inference time - Requires large-scale computing power	- AI art & image generation - Video synthesis - Creative design

	- Can create diverse outputs		
Convolutional Neural Networks (CNNs)	- Excellent for image recognition - Well-optimized for computer vision tasks - Pretrained models available	- Computationally expensive - Requires labeled data for training	- Object detection - Facial recognition - Medical image analysis
Recurrent Neural Networks (RNNs, LSTMs, GRUs)	- Effective for sequential data - Good for time-series analysis - Memory-efficient for short-term dependencies	- Struggles with long-term dependencies - Slower training compared to transformers	- Speech recognition - Time-series forecasting - Stock market predictions
Hybrid Models (e.g., GPT+CNN, Transformer + GANs)	- Combines strengths of multiple architectures - Can handle multimodal data (text, image, audio) - More robust performance	- Increased complexity - Higher computational demand - Difficult to optimize	- Multimodal AI (text-to-image, video understanding) - Advanced recommendation systems - AI-powered creative tools

5.3. **Ethical AI Practices**: Ensuring fairness, transparency, and accountability in AI applications is critical. Organizations should implement bias mitigation strategies, conduct regular audits of AI outputs, and comply with industry regulations. Ethical AI frameworks provide guidelines for responsible AI development. Additionally, explainability techniques, such as SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations), help improve model transparency.

5.4. **Iterative Development**: Continuous iteration is necessary to refine AI models and adapt to changing user needs. Enterprises should implement agile methodologies, incorporating user feedback and A/B testing to optimize model performance. Techniques such as reinforcement learning from human feedback (RLHF) can further enhance model accuracy. Automated monitoring tools can be deployed to detect performance drift and retrain models accordingly.

5.5. **Scalability and Infrastructure**: Deploying GenAI at scale requires robust cloud infrastructure and efficient model optimization. Organizations should leverage AI accelerators such as GPUs and TPUs, as well as model compression techniques like quantization and pruning to enhance efficiency. Cloud providers like AWS, Google Cloud, and Microsoft Azure offer specialized AI services that facilitate scalable GenAI deployments.

5.6. **Human-AI Collaboration**: Effective use of GenAI requires integrating human expertise with AI capabilities. Human-in-the-loop (HITL) systems ensure that AI-generated content is reviewed and refined by domain experts, particularly in high-stakes industries like legal, medical, and financial services. Establishing clear guidelines for human oversight helps maintain quality and accountability.

5.7. **Security and Compliance**: As AI adoption increases, security risks such as adversarial attacks and data breaches become critical concerns. Organizations should implement robust encryption protocols, access controls, and AI model hardening techniques. Compliance with regulations like GDPR, CCPA, and industry-specific standards ensures legal and ethical deployment of AI applications.

By following these best practices, enterprises can maximize the potential of Generative AI while mitigating risks and ensuring sustainable AI adoption.

6. FINE-TUNING AI MODELS: NEEDS AND METHODOLOGICAL PROCESSES

The fine-tuning of Generative AI (GenAI) models is critical for enhancing their domain-specific performance, accuracy, and compliance with industry-specific requirements. This paper presents a structured approach to fine-tuning GenAI models, outlining the key motivations, essential data preparation steps, and computational methodologies. Additionally, various fine-tuning strategies are discussed, along with their advantages and limitations.

Pre-trained Generative AI models exhibit significant generalization capabilities but often require domain adaptation to improve accuracy and relevance for specific tasks. Fine-tuning enables models to specialize in niche domains by training on curated datasets, improving their performance in areas such as legal, medical, and financial applications. This section explores the necessity of fine-tuning and the fundamental processes involved.

Fine-tuning is essential for multiple reasons, including:

6.1. **Domain-Specific Adaptation**: Pre-trained models may lack knowledge in specialized fields; fine-tuning enhances domain-specific comprehension and response accuracy.

6.2. **Output Control and Bias Reduction**: Customization ensures alignment with regulatory requirements, ethical guidelines, and brand-specific communication styles.

6.3. **Improved Accuracy and Relevance**: Training on high-quality datasets enhances factual correctness and contextual relevance.

6.4. **Computational Optimization**: Reducing inference time and optimizing model size enhances deployment feasibility.

6.5. **Data Privacy and Security**: Fine-tuning on proprietary data prevents exposure of sensitive information in publicly accessible models.

Tabel 4: The key fine-tuning process steps

Step	Description
1. Data Collection and Preprocessing	- Curate domain-specific datasets relevant to the application. - Clean and preprocess data by removing inconsistencies, normalizing formats, and ensuring high-quality annotations. - Convert data into structured input-output pairs suitable for training.
2. Model Selection	- Choose an appropriate pre-trained model (e.g., GPT, BERT, T5, LLaMA) based on task requirements. - Consider model size, architecture, and computational feasibility.
3. Fine-Tuning Techniques	Several approaches to fine-tuning exist, including: **Full Fine-Tuning**: Modifies all model parameters; computationally intensive but yields the highest performance improvements.

	Low-Rank Adaptation (LoRA): Introduces trainable low-rank layers, reducing computational overhead while maintaining performance. **Prompt Tuning**: Optimizes prompt embeddings instead of modifying model weights, reducing resource consumption. **Adapter Layers**: Adds lightweight task-specific layers, preserving the core model while allowing customization.
4. Training Execution	- Utilize frameworks such as PyTorch, TensorFlow, or Hugging Face's Trainer API. - Configure hyperparameters, including batch size, learning rate, and optimizer. - Leverage GPU or TPU acceleration for efficient training.
5. Evaluation and Testing	- Validate model performance using a hold-out dataset. - Utilize quantitative metrics such as BLEU, ROUGE, perplexity, and F1-score. - Conduct qualitative assessments to ensure response accuracy and relevance.
6. Deployment and Optimization	- Deploy the fine-tuned model via APIs or integrate it into production environments. - Monitor real-world performance and refine the model iteratively. - Optimize inference speed using techniques such as quantization and distillation.

The major challenges and considerations related to fine-tuning include the following:

- Ethical Concerns and Bias Mitigation: Regular audits are necessary to minimize biases in model outputs.
- Computational Resource Constraints: Cloud-based GPU solutions (AWS, Azure, GCP) help scale training efficiently.
- Compliance and Security: Ensuring adherence to legal frameworks such as GDPR and HIPAA is crucial when working with sensitive data.

Fine-tuning GenAI models is a critical process that enables enhanced domain-specific capabilities, improved accuracy, and controlled customization. By following a structured methodology, researchers and practitioners can optimize AI models for specialized applications while addressing challenges related to computational efficiency, ethical concerns, and deployment scalability. Future work should focus on automated fine-tuning pipelines and advanced reinforcement learning techniques for continuous model improvement.

7. Conclusion

Generative AI presents significant opportunities for enterprises across industries, offering transformative capabilities in automation, content creation, and decision support. However, successful adoption requires a strategic approach that balances technical feasibility, business objectives, and ethical considerations. This paper has outlined industry-specific GenAI applications, best practices for implementation, and methodologies for fine-tuning AI models to enhance domain-specific performance. Enterprises must prioritize data quality, model transparency, and regulatory compliance while leveraging scalable infrastructure for deployment. As AI technologies continue to evolve, organizations should adopt iterative development strategies and human-AI collaboration frameworks to maximize GenAI's potential while mitigating risks. Future advancements in automated fine-tuning and responsible AI governance will further shape the role of GenAI in enterprise software, driving innovation and competitive advantage.

REFERENCES

[1]. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are

unsupervised multitask learners. OpenAI Blog, 1(8), 9.

[2]. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big?. Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT'21), 610–623.

[3]. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., ... & Amodei, D. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems (NeurIPS'20), 33, 1877–1901.

[4]. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. Advances in Neural Information Processing Systems (NeurIPS'14), 27, 2672–2680.

[5]. Marcus, G. (2021). The next decade in AI: Four steps towards robust artificial intelligence. ArXiv preprint.

[6]. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems (NeurIPS'17), 30, 5998–6008.

[7]. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., & Sutskever, I. (2021). Zero-shot text-to-image generation. Proceedings of the 38th International Conference on Machine Learning (ICML'21), 139, 8821–8831.

[8]. Chollet, F. (2017). Deep learning with Python. Manning Publications.

[9]. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT'19), 220–229.

[10]. Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542(7639), 115–118.

[11]. Li, X., Fei, Y., & Wang, L. (2019). Personalized recommendation system in e-commerce using AI-driven predictive analytics. Journal of Retail Analytics, 5(2), 99–114.

[12]. Lee, J., Bagheri, B., & Kao, H. A. (2018). A cyber-physical systems architecture for industry 4.0-based manufacturing systems. Manufacturing Letters, 3, 18–23.

[13]. Surden, H. (2019). Artificial intelligence and law: An overview. Georgia State University Law Review, 35(4), 1305–1337.

[14]. Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. Reuters.

[15]. Zhang, Y., Feng, S., & Neubig, G. (2020). LAUG: Language-aware user simulation for task-oriented dialogue systems. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP'20), 8564–8579.

[16]. J. Jangid, "Efficient Training Data Caching for Deep Learning in Edge Computing Networks," International Journal of Scientific Research in Computer Science, Engineering and Information Technology, vol. 7, no. 5, pp. 337–362, 2020. doi: 10.32628/CSEIT20631113

[17]. Malhotra, S., Yashu, F., Saqib, M., & Divyani, F. (2020). A multi-cloud orchestration model using Kubernetes for microservices. Migration Letters, 17(6), 870–875. https://migrationletters.com/index.php/ml/article/view/11795

[18]. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., & Chen, M. (2022). Hierarchical text-conditional image generation with CLIP latents. ArXiv preprint.

[19]. Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. Journal of Big Data, 6(1), 1–48.
[20]. Chen, M., Tworek, J., Jun, H., Yuan, Q., de Oliveira Pinto, H. P., Kaplan, J., ... & Schulman, J. (2021). Evaluating large language models trained on code. arXiv preprint arXiv:2107.03374.
[21]. Agrawal, A., Gans, J., & Goldfarb, A. (2018). Prediction machines: The simple economics of artificial intelligence. Harvard Business Review Press.
[22]. Oord, A. v. d., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., & Kavukcuoglu, K. (2016). WaveNet: A generative model for raw audio. ArXiv preprint.
[23]. Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems (NeurIPS'20), 33, 6840–6851.