

Deep Learning Algorithms in the Development of Generative AI Models for Automated Content Creation

Agung Yuliyanto Nugroho*

Universitas Cendekia Mitra Indonesia, Indonesia

*Corresponding Author: agungboiler11@gmail.com

Abstract. The rapid advancement of artificial intelligence (AI) has introduced a new paradigm in informatics known as Generative AI. One of the key driving forces behind this innovation is the application of deep learning algorithms, which can emulate human cognitive patterns to automatically generate text, images, audio, and video. This study aims to analyze how deep learning algorithms particularly Generative Adversarial Networks (GANs) and Transformer-based Models (such as GPT and Diffusion Models) are utilized in developing generative AI systems for automated content creation. The research employs a literature review of recent studies, comparative analysis of generative models, and performance evaluation based on quality, creativity, and computational efficiency. The findings reveal that Transformer-based models exhibit greater adaptability in understanding semantic context and producing more realistic content compared to traditional GAN models. However, challenges such as overfitting, data bias, and high computational resource demands remain major obstacles to large-scale implementation. This study concludes that optimizing deep learning algorithms supported by ethical considerations and careful data management will be crucial to the successful development of generative AI that is both effective and responsible within the modern informatics ecosystem.

Keywords: Content; Deep Learning; Generative Adversarial Network; Generative AI; Transformer

1. INTRODUCTION

The evolution of artificial intelligence (AI) has transformed the field of informatics, shifting its focus from rule-based computation to intelligent systems capable of autonomous learning and creative generation. One of the most significant developments in this transformation is Generative Artificial Intelligence (Generative AI), a subset of AI that utilizes advanced machine learning architectures to create new data that resembles human-generated content. According to Goodfellow et al. (2014), who introduced the Generative Adversarial Network (GAN) framework, generative models can learn data distributions and produce novel yet coherent outputs, marking a turning point in computational creativity. In recent years, deep learning algorithms have become the backbone of this technology, enabling machines to generate text, images, sound, and even videos with unprecedented realism and contextual understanding.

The rapid progress of deep learning architectures, particularly Transformer-based models such as OpenAI's GPT and Google's BERT, has further expanded the potential of generative systems. These models leverage massive datasets and multilayer attention mechanisms to capture complex semantic relationships, allowing AI to generate high-quality, contextually relevant content automatically (Vaswani et al., 2017). In the field of informatics, this innovation has not only revolutionized content creation but also reshaped workflows in education, business, media, and creative industries (Kaur & Singh, 2023). However, the

expansion of generative models also raises significant issues regarding computational efficiency, data ethics, and the authenticity of AI-generated information.

Despite the growing adoption of generative AI, several challenges remain unresolved. First, deep learning models often suffer from overfitting and data bias, which can compromise output quality and fairness (Zhang et al., 2021). Second, the high computational cost associated with model training limits accessibility, especially for small-scale developers and institutions. Third, the issue of ethical transparency including data provenance, authorship, and misinformation has sparked global debates among researchers and policymakers (Floridi & Chiriatti, 2020). These challenges highlight the urgent need for systematic studies that analyze how deep learning algorithms can be optimized for effective and responsible generative AI development.

From an informatics perspective, the study of generative AI represents a multidisciplinary endeavor, combining aspects of algorithmic design, data management, and computational ethics. Informatics not only provides the technological infrastructure for deep learning but also contributes theoretical frameworks for evaluating the social and cognitive implications of automated content creation. As noted by Russell and Norvig (2022), the goal of AI research should not merely be to replicate human intelligence but to augment human capability through intelligent systems that operate within ethical and contextual boundaries.

Therefore, this study aims to analyze the role of deep learning algorithms in developing generative AI models for automated content creation, with particular emphasis on their architecture, performance efficiency, and ethical implications. By synthesizing recent literature and evaluating comparative model performance, this research seeks to contribute a comprehensive understanding of how deep learning can be harnessed to advance generative AI responsibly and sustainably within the broader landscape of informatics.

2. LITERATURE REVIEW

Generative Artificial Intelligence (AI) refers to a subset of AI designed to create new data resembling human-generated content by learning from existing datasets. Early developments in generative modeling were grounded in probabilistic and statistical learning frameworks such as Hidden Markov Models (HMMs) and Variational Autoencoders (VAEs) (Kingma & Welling, 2014). However, the field witnessed a major breakthrough with the introduction of Generative Adversarial Networks (GANs) by Goodfellow et al. (2014). GANs consist of two neural networks a generator and a discriminator that compete against each other, enabling the system to learn data distributions and produce realistic synthetic outputs. This

adversarial mechanism has since been the foundation of numerous creative AI applications, including image synthesis, voice cloning, and text generation.

In recent years, the emergence of Transformer-based architectures has significantly advanced generative capabilities. The introduction of the Transformer model by Vaswani et al. (2017) revolutionized natural language processing (NLP) by replacing sequential processing with an attention mechanism, allowing models to capture long-range dependencies and contextual relationships more effectively. This advancement gave rise to powerful generative models such as GPT (Generative Pretrained Transformer) and BERT (Bidirectional Encoder Representations from Transformers), which have since become central to modern AI-driven automation (Brown et al., 2020). As noted by Bommasani et al. (2021), these large-scale language models represent a new paradigm in AI, blurring the boundary between data comprehension and creation.

Deep Learning Algorithms and Informatics Foundations

Deep learning is a subfield of machine learning that employs multi-layered neural networks to model high-level abstractions in data (LeCun, Bengio, & Hinton, 2015). In the context of informatics, deep learning provides the computational backbone for systems capable of pattern recognition, data synthesis, and predictive analytics. Informatics, as described by Sahay and Walsham (2022), integrates computing technologies with information management principles to support decision-making and problem-solving in complex environments. The fusion of deep learning with informatics has enabled the development of systems that can autonomously analyze massive datasets and generate meaningful insights or creative outputs.

Several algorithmic architectures underpin generative deep learning models. Convolutional Neural Networks (CNNs) are widely used in image generation, while Recurrent Neural Networks (RNNs) and Transformers dominate text and speech synthesis (Radford et al., 2019). Diffusion Models, introduced more recently, generate data by progressively denoising random noise, producing highly detailed and controllable outputs (Ho et al., 2020). Comparative studies (Dhariwal & Nichol, 2021) have shown that diffusion-based models can outperform GANs in generating high-fidelity images, albeit with greater computational costs.

Challenges and Ethical Considerations in Generative AI

Despite these technological advancements, generative AI presents notable challenges that extend beyond algorithmic optimization. One of the major concerns is data bias, which arises from unrepresentative training datasets and can result in discriminatory or misleading outputs (Bender et al., 2021). Another critical issue is authenticity and ownership, as AI-generated content often raises legal and ethical questions regarding intellectual property and

misinformation (Floridi & Chiriatti, 2020). Furthermore, energy consumption and computational sustainability have emerged as pressing concerns, with large-scale deep learning models consuming significant resources (Strubell, Ganesh, & McCallum, 2019).

From an informatics perspective, addressing these issues requires an interdisciplinary approach that integrates technical, ethical, and regulatory frameworks. According to Mittelstadt et al. (2016), responsible AI development must include transparency, accountability, and fairness as core principles. Researchers are increasingly emphasizing Explainable AI (XAI) and Responsible AI frameworks to enhance interpretability and ensure that generative systems operate within ethical boundaries (Arrieta et al., 2020). Therefore, this study positions itself within this research gap—analyzing how deep learning algorithms can be optimized and ethically managed in the development of generative AI systems for automated content generation, particularly in the context of modern informatics.

3. METHOD

This study employs a qualitative descriptive research design with a systematic literature review (SLR) and comparative analysis approach. The primary goal is to analyze how deep learning algorithms particularly Generative Adversarial Networks (GANs), Transformers, and Diffusion Models are developed and optimized for automated content creation within the informatics domain. This design was chosen to provide a comprehensive understanding of technological architectures, performance benchmarks, and ethical considerations that shape the evolution of generative AI systems.

The research adopts an exploratory framework, emphasizing critical synthesis of previous empirical and theoretical studies rather than conducting new experimental testing. This approach aligns with the objective of identifying patterns, limitations, and future directions in the application of deep learning for generative purposes.

Primary Data Collection

The study relies on secondary data sources obtained from reputable academic databases, including IEEE Xplore, SpringerLink, ScienceDirect, and Google Scholar. Selection criteria for the literature include:

- a. Publications between 2014 and 2025, capturing the timeline from the introduction of GANs to the emergence of Transformer-based and diffusion models.
- b. Peer-reviewed journal articles, conference papers, and technical reports relevant to deep learning, AI generation, and informatics systems.
- c. Studies that discuss model architecture, algorithmic optimization, computational

efficiency, and ethical or social implications of AI-generated content.

A total of 65 primary studies were initially identified, from which 40 core articles were selected after evaluating their relevance, credibility, and citation frequency. Data extraction was conducted using a structured review matrix, allowing for thematic categorization and synthesis.

Data Analysis

Data were analyzed using a thematic analysis approach, as proposed by Braun and Clarke (2006), to identify key concepts and emerging trends. The analysis involved three stages:

- a. Descriptive Mapping – summarizing how deep learning algorithms are structured and applied in generative AI systems;
- b. Comparative Evaluation – analyzing performance metrics among GANs, Transformers, and Diffusion Models based on computational efficiency, data quality, and content realism;
- c. Critical Interpretation – examining ethical, technical, and operational challenges associated with implementing deep learning in automated content generation.

To ensure analytical rigor, the study incorporated triangulation through cross-referencing multiple sources and model evaluations, thereby minimizing interpretive bias. The culture of Yogyakarta, including festivals like Sekaten and the Yogyakarta Arts Festival, also serves as a unique attraction for tourists. These cultural events provide backpackers with opportunities to engage with local traditions and immerse themselves in the vibrant atmosphere of the city.

4. RESULTS AND DISCUSSION

Results

The literature analysis revealed that deep learning algorithms have become the fundamental enablers of automated content generation across various domains—ranging from text and image synthesis to audio and video creation. Three dominant model architectures were identified: Generative Adversarial Networks (GANs), Transformer-based Models, and Diffusion Models. Each model exhibits unique structural mechanisms, performance characteristics, and computational demands. GANs are highly effective in generating realistic images and multimedia data through adversarial training between a generator and discriminator (Goodfellow et al., 2014). However, they often suffer from mode collapse and instability during training. Transformers, introduced by Vaswani et al. (2017), demonstrated

superior contextual understanding in sequential data generation, particularly for language and text-based applications. Their attention mechanisms enable adaptive modeling of long-range dependencies.

Diffusion Models represent the latest evolution in generative AI, utilizing iterative denoising processes to produce high-fidelity content (Ho et al., 2020). Despite their excellent output quality, they require extensive computational resources and long training times. These findings indicate that while GANs are advantageous for creative visual outputs, Transformers dominate linguistic and multimodal generation, and Diffusion Models are emerging as a hybrid alternative that balances quality and controllability.

To evaluate model performance, this study categorized findings based on three main dimensions: content realism, computational efficiency, and contextual adaptability.

Table 1. Strengths, Weaknesses, and Key Applications of Generative AI Models

Model Type	Strengths	Weaknesses	Key Applications
GANs	High-quality image synthesis, flexible architecture	Training instability, mode collapse	Image generation, style transfer, video synthesis
Transformers	Superior context understanding, scalable to multimodal data	High memory usage, interpretability challenges	Text generation, chatbot systems, creative writing
Diffusion Models	High-fidelity output, strong noise resistance	Computationally expensive, slower generation speed	Art design, photo enhancement, text-to-image synthesis

Results from comparative studies (Dhariwal & Nichol, 2021; Rombach et al., 2022) suggest that Transformer-based architectures outperform GANs in tasks involving semantic reasoning and multi-context understanding, while Diffusion Models achieve the best perceptual quality for visual data. However, computational trade-offs remain a central issue: high-quality outputs often come at the cost of increased training time and energy consumption—highlighting the ongoing challenge of achieving both efficiency and realism in generative AI.

Overall, the findings suggest that deep learning algorithms have become the cornerstone of generative AI, driving advancements in automated content creation across multiple domains. However, achieving a balance between creativity, computational efficiency, and ethical responsibility remains a complex challenge in informatics research. The evolution from GANs to Transformer and Diffusion architectures represents not only technical progress but also a paradigm shift toward context-aware and ethically aligned AI systems.

This synthesis reinforces that the future of generative AI in informatics will depend on three interdependent factors:

- a. Algorithmic innovation;
- b. Sustainable and ethical data practices; and
- c. Transparent integration into real-world information systems.

```
import tensorflow as tf
from tensorflow import keras
import numpy as np
import matplotlib.pyplot as plt
```

Figure 1. Python Code

Discussion

The findings of this study highlight the pivotal role of deep learning algorithms in redefining the frontiers of generative artificial intelligence and its applications in automated content creation. From an informatics perspective, these algorithms serve as the computational and conceptual backbone for systems capable of generating, interpreting, and adapting digital content autonomously. The discussion that follows synthesizes the implications of these findings across three key dimensions technological evolution, informatics integration, and ethical governance while identifying ongoing challenges and future opportunities.

The comparative analysis confirms a clear trajectory in the evolution of generative models from adversarial learning (as in GANs) toward contextual and diffusion-based modeling. While Generative Adversarial Networks (GANs) introduced a new paradigm of machine creativity by enabling adversarial learning between a generator and a discriminator (Goodfellow et al., 2014), their instability and tendency toward mode collapse limited broader adoption in complex environments. In contrast, Transformer architectures (Vaswani et al., 2017) have demonstrated remarkable progress in contextual comprehension and multimodal adaptability, allowing AI systems to generate coherent and semantically meaningful outputs. The rise of Diffusion Models (Ho et al., 2020) further illustrates a shift toward probabilistic and iterative generativity, emphasizing control, precision, and realism in output generation. This technological evolution suggests that the core innovation in generative AI no longer lies merely in producing data-like outputs, but in developing models capable of understanding context, semantics, and intent a transition that mirrors the broader transformation within informatics toward intelligent, context-aware systems.

From an informatics standpoint, deep learning-based generative AI represents the convergence of data science, information systems, and computational intelligence. Informatics is not limited to data storage and retrieval but encompasses the design of intelligent systems

that process and synthesize information meaningfully. Generative AI, when integrated into information management systems, enables dynamic automation—transforming traditional workflows into adaptive ecosystems that can learn, predict, and create. Examples include AI-assisted documentation systems, smart report generators, and content-based decision support systems.

Furthermore, the use of synthetic data generation via GANs and Diffusion Models contributes significantly to informatics research by addressing data scarcity and privacy concerns. Frid-Adar et al. (2018) demonstrated that synthetic medical images can improve diagnostic accuracy in healthcare informatics without exposing real patient data. This intersection between deep learning innovation and data governance underlines the discipline's growing responsibility to ensure that generative systems operate within ethical and transparent frameworks.

While technological advancement is evident, the ethical dimension of generative AI remains a critical area of concern. Scholars such as Floridi and Chiriatti (2020) argue that the ethical implications of generative systems ranging from authorship to misinformation demand an integrated governance approach. The ability of AI models to produce hyper-realistic but synthetic content challenges established notions of authenticity and intellectual property, requiring clear regulatory and ethical frameworks.

Moreover, data bias continues to be a pressing issue. Studies by Bender et al. (2021) and Crawford (2021) emphasize that biased training data can reinforce stereotypes and distort information, which in turn undermines the credibility of AI-driven systems. Addressing this requires transparent data curation and explainable AI (XAI) methodologies that enhance interpretability and accountability (Arrieta et al., 2020).

The environmental cost of large-scale model training also warrants serious attention. Strubell et al. (2019) estimate that training a single deep learning model can emit carbon equivalent to multiple cars' lifetime emissions. In this regard, Green Informatics and energy-efficient computing architectures emerge as crucial directions for sustainable AI development.

Thus, ethical governance in generative AI goes beyond moral responsibility. It represents a core element of systemic informatics resilience. This ensures that intelligent technologies contribute positively to both human and environmental well-being.

The future of generative AI within informatics depends on achieving a balance between innovation, efficiency, and responsibility. The integration of Responsible AI frameworks (Floridi & Cowls, 2021) offers a pathway toward ensuring that technological progress aligns with societal values. Informatics researchers and practitioners must adopt transparent

methodologies that incorporate fairness auditing, explainability, and data stewardship as standard practices.

Furthermore, the rise of multimodal generative systems signals the next frontier of AI research, wherein deep learning models can synthesize text, image, and sound simultaneously pushing the boundaries of computational creativity. To harness this potential responsibly, future informatics research should focus on energy-efficient architectures, bias reduction strategies, and human-AI collaboration frameworks that prioritize co-creativity over automation dominance. Ultimately, the development of generative AI marks not merely a technical milestone but a philosophical shift in informatics one that redefines the role of machines from tools of computation to partners in cognition and creation.

5. CONCLUSION AND SUGGESTIONS

This study concludes that deep learning algorithms are the primary enablers behind the rapid development of generative artificial intelligence (AI) systems capable of automated content creation. The evolution from Generative Adversarial Networks (GANs) to Transformer-based architectures and Diffusion Models signifies a major leap in computational intelligence from pattern imitation to contextual understanding and creative synthesis. From an informatics perspective, the integration of deep learning into generative systems has redefined the concept of information processing. It has shifted the paradigm from static data management toward adaptive and autonomous content generation, enabling systems that can learn, interpret, and produce new information in real time. These capabilities have profound implications for various sectors, including education, healthcare, digital media, and decision support systems. However, the study also identifies several persistent challenges that hinder the optimal and ethical application of generative AI. These include data bias, overfitting, intellectual property issues, and significant environmental costs due to high computational demands. Moreover, the growing sophistication of AI-generated content raises ethical dilemmas concerning authenticity, accountability, and societal impact. Therefore, while deep learning algorithms have revolutionized the technical foundation of generative AI, their implementation must be guided by responsible informatics principles, ensuring transparency, fairness, and sustainability. The role of informatics scholars and practitioners is not merely to enhance system performance but to shape the ethical and structural frameworks that govern how generative technologies are designed, deployed, and monitored. In essence, generative AI marks the beginning of a new era in computational creativity one where human and machine intelligence converge to co-create knowledge and meaning. Yet, this convergence must be

continuously balanced with critical reflection on its implications for humanity, culture, and the environment.

Based on the findings and discussions, several recommendations are proposed for researchers, practitioners, and policymakers in the field of informatics and artificial intelligence. Future studies should focus on improving computational efficiency and energy sustainability in deep learning architectures. Research into lightweight models, quantization techniques, and Green AI frameworks will be essential to reduce the environmental footprint of large-scale training processes. It is also critical to establish standardized protocols for data curation, bias detection, and model explainability. Informatics researchers should integrate ethical data governance practices and promote the adoption of Explainable AI (XAI) methods to ensure that generative systems remain fair, interpretable, and socially trustworthy. Moreover, the complexity of generative AI demands collaboration between computer scientists, data ethicists, cognitive psychologists, and policymakers. Such interdisciplinary engagement can foster holistic informatics ecosystems that address both technical innovation and societal impact.

Governments and international institutions must develop regulatory mechanisms that balance innovation with public accountability, including setting guidelines on AI-generated content ownership, intellectual property, and digital authenticity verification. Instead of replacing human creativity, generative AI should be designed to augment and collaborate with human users through a human-centered AI approach, emphasizing interactive interfaces, adaptive learning, and participatory design. Furthermore, academic institutions should integrate generative AI literacy into informatics and computer science curricula to build competence in deep learning, data ethics, and AI governance, preparing future professionals to manage and innovate responsibly within the emerging generative ecosystem. Lastly, to ensure long-term accountability, organizations employing generative AI should conduct periodic audits of algorithmic outputs and training datasets, evaluating not only accuracy and performance but also social implications and ethical conformity.

REFERENCES

- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., ... & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- 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)*, 610–623. <https://doi.org/10.1145/3442188.3445922>
- Brynjolfsson, E., & McAfee, A. (2017). *Machine, platform, crowd: Harnessing our digital future*. W. W. Norton & Company.
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press. <https://doi.org/10.12987/9780300252392>
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT 2019*, 4171–4186. <https://doi.org/10.48550/arXiv.1810.04805>
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Hounsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30(4), 681–694. <https://doi.org/10.1007/s11023-020-09548-1>
- Floridi, L., & Cows, J. (2021). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 3(1). <https://doi.org/10.1162/99608f92.8cd550d1>
- Frid-Adar, M., Klang, E., Amitai, M., Goldberger, J., & Greenspan, H. (2018). Synthetic data augmentation using GAN for improved liver lesion classification. *IEEE Transactions on Medical Imaging*, 38(3), 839–853. <https://doi.org/10.1109/TMI.2018.2862452>
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems (NeurIPS)*, 27, 2672–2680.
- Guo, Y., Liu, Y., Oerlemans, A., Lao, S., Wu, S., & Lew, M. S. (2016). Deep learning for visual understanding: A review. *Neurocomputing*, 187, 27–48. <https://doi.org/10.1016/j.neucom.2015.09.116>
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 33, 6840–6851.
- Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4700–4708. <https://doi.org/10.1109/CVPR.2017.243>

- Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- Li, B., Liu, S., Xu, J., & Lin, Z. (2023). Diffusion models in visual content generation: A comprehensive survey. *arXiv preprint arXiv:2303.03968*.
- Minsky, M. (1986). *The society of mind*. Simon and Schuster.
- Nugroho, A. Y. (2024). Transformasi digital: Mengoptimalkan strategi e-commerce di era disrupsi. *Jurnal Ilmiah Bisnis Digital*, 1(1), 56–66. <https://doi.org/10.69533/4wj6tw45>
- Nugroho, A. Y., & Sutopo, J. (2018). Rancang bangun aplikasi pemesanan layanan service mobil berbasis website. *Jurnal ePrints UTY*.
- Nugroho, A. Y. (2024). Pendampingan NIB untuk usaha mikro: Solusi menuju kesuksesan bisnis yang berkelanjutan. *Khidmah Nusantara*, 1(1), 127–134.
- Nugroho, A. Y., Kom, M., & Par, M. (2025). *Pengantar videografi untuk pemula*. Yogyakarta: CV Gemilang Press Indonesia.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., ... & Christiano, P. (2022). Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. *OpenAI Technical Report*.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., & Chen, M. (2022). Hierarchical text-conditional image generation with CLIP latents. *arXiv preprint arXiv:2204.06125*.
- Russell, S. J., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson Education.
- Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, 3645–3650. <https://doi.org/10.18653/v1/P19-1355>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 5998–6008.
- Vincent, P., Larochelle, H., Bengio, Y., & Manzagol, P. A. (2008). Extracting and composing robust features with denoising autoencoders. *Proceedings of the 25th International Conference on Machine Learning (ICML)*, 1096–1103. <https://doi.org/10.1145/1390156.1390294>