

Generative AI: A Comprehensive Review of Foundational Models and Emerging Methods

Noper Ardi ^{1*}, Isnayanti ²

*correspondence: noperardi@polibatam.ac.id

¹Informatics Engineering, Politeknik Negeri Batam, Indonesia

²Mathematics Department, Institut Teknologi Batam, Indonesia

Abstract- Generative Artificial Intelligence (AI) has emerged as a transformative field within computer science, heralding a new era of content creation and problem-solving. This comprehensive review charts the evolution of generative models, from the foundational pillars to the cutting-edge methods that are reshaping industries. We begin by examining the seminal architectures that laid the groundwork for the field: Generative Adversarial Networks (GANs), with their unique adversarial training paradigm; Variational Autoencoders (VAEs), which leverage probabilistic graphical models for generation; and the early instantiations of Transformer models that revolutionized sequence-to-sequence tasks. Subsequently, we transition to the current vanguard of generative AI, providing an in-depth analysis of Large Language Models (LLMs). These models have demonstrated unprecedented capabilities in understanding and generating human-like text, leading to a paradigm shift in natural language processing. Concurrently, we explore the rise of Diffusion Models, which have set new benchmarks in high-fidelity image synthesis through a process of iterative denoising. This review synthesizes the theoretical underpinnings, architectural innovations, and practical applications of these models. We also present a comparative analysis, highlighting their respective strengths, limitations, and the evolutionary trajectory of the field. Finally, we discuss the prominent challenges and ethical considerations that accompany the proliferation of generative AI and conclude with a perspective on future research directions that will continue to propel this remarkable domain forward.

Keywords: Generative Ai; Foundational Models; Generative Adversarial Networks; Variational Autoencoders; Large Language Models

Article info: *submitted January 1, 2025, revised February 2, 2025, accepted March 15, 2025*



This is an open access article under the [CC-BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.

1. Introduction

Generative Artificial Intelligence represents a significant leap in machine learning, endowing machines with the ability to produce novel data that mimics the complexity and richness of real-world information. Unlike discriminative models that are trained to classify or predict from existing data, generative models learn the underlying distribution of data, enabling them to create entirely new artifacts, from text and images to music and code. The rapid advancements in this domain are not merely incremental; they represent a fundamental shift in how we approach computation, creativity, and automation.

The genesis of modern generative AI can be traced back to the development of deep learning techniques that allowed for the modeling of intricate data distributions. Early explorations, while promising, often struggled with generating high-quality, diverse samples. However, the introduction of pioneering architectures in the mid-2010s marked a turning point, setting the stage for the explosive growth we witness today.

This review provides a comprehensive overview of the key models that have defined and continue to shape the landscape of generative AI. We will first delve into the foundational models that served as the bedrock of the field: Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and early Transformer models. These architectures, while distinct in their approach, collectively addressed many of the initial challenges in generative modeling and opened up a plethora of applications.

Building on this foundation, we will then explore the emerging methods that are currently at the forefront of research and application: Large Language Models (LLMs) and Diffusion Models. LLMs have captured the public imagination with their remarkable fluency and knowledge, while Diffusion Models have achieved state-of-the-art results in image generation, rivaling human creativity.

The overarching goal of this paper is to provide a holistic understanding of the evolution of generative AI, from its foundational principles to its most recent breakthroughs. By examining the trajectory of these models, we can not only

appreciate the progress that has been made but also anticipate the future directions of this dynamic and impactful field.

2. Methods

A. Foundational Models in Generative AI

The journey of modern generative AI began with the development of models that could effectively learn and sample from complex, high-dimensional data distributions. This section reviews three of the most influential foundational models: Generative Adversarial Networks, Variational Autoencoders, and early Transformer-based architectures.

1) Generative Adversarial Networks (GANs)

Introduced by Goodfellow et al. in 2014, Generative Adversarial Networks (GANs) brought a novel game-theoretic approach to generative modeling. A GAN framework consists of two neural networks, a Generator (G) and a Discriminator (D), which are trained in an adversarial manner[1].

The Generator (G) takes a random noise vector, z , as input and attempts to generate data that resembles the training data. Its goal is to produce samples that are indistinguishable from real data. The Discriminator (D) is a binary classifier that receives both real data from the training set and fake data from the generator. Its objective is to correctly distinguish between the real and generated samples.

The training process can be conceptualized as a zero-sum game. The generator strives to fool the discriminator by producing increasingly realistic data, while the discriminator continuously improves its ability to detect fakes. This adversarial dynamic drives both networks to enhance their capabilities, leading to the generation of high-fidelity samples. The objective function for a standard GAN can be expressed as[2][3]:

$$\min_D \max_G V(D, G) = \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))] \quad (1)$$

Since their inception, GANs have seen numerous architectural and theoretical improvements. For instance, Deep Convolutional GANs (DCGANs) introduced a set of architectural constraints that stabilized training and improved the quality of generated images. Other notable advancements include Wasserstein GANs (WGANs), which use a different loss function to mitigate the problem of vanishing gradients, and Conditional GANs (cGANs), which allow for the generation of data conditioned on specific attributes[1].

GANs have been successfully applied in a wide range of domains, including Image Synthesis including Generating photorealistic images of faces, objects, and scenes. Image-to-Image Translation consist of Translating an image from one domain to another, such as converting a satellite image into a map. Data Augmentation is a Creating synthetic data to augment training sets for other machine learning models[4][5].

2) Variational Autoencoders (VAEs)

Variational Autoencoders (VAEs), introduced by Kingma and Welling in 2013, offer a probabilistic approach to generative

modeling. VAEs are a type of autoencoder, which is a neural network trained to learn a compressed representation (encoding) of its input and then reconstruct the input from this encoding.

A VAE consists of two main components:

The Encoder (Inference Network) takes an input data point, x , and outputs the parameters of a probability distribution in the latent space, typically a mean and a variance for a Gaussian distribution. This distribution, $q(z|x)$, represents the encoded representation of the input[6].

The Decoder (Generative Network) takes a point, z , sampled from the latent distribution and attempts to reconstruct the original input, x . It defines a conditional probability distribution, $p(x|z)$.

The training objective of a VAE is to maximize the Evidence Lower Bound (ELBO), which is a lower bound on the marginal log-likelihood of the data. The ELBO consists of two terms, which is Reconstruction Loss; This term encourages the decoder to accurately reconstruct the input data and Kullback-Leibler (KL) Divergence; This term acts as a regularizer, forcing the learned latent distribution, $q(z|x)$, to be close to a prior distribution, $p(z)$, which is typically a standard normal distribution. This regularization helps in generating novel data by sampling from the smooth, continuous latent space[7][8][9].

The ELBO can be written as:

$$L(\theta, \phi; x) = \mathbb{E} q\phi(z|x) [\log p\theta(x|z)] - \text{DKL}(q\phi(z|x) || p(z)) \quad (2)$$

VAEs are known for their ability to learn smooth and meaningful latent representations, making them suitable for Data Generation for Creating new data samples by sampling from the learned latent space. Anomaly Detection; Identifying data points that have a low probability of being generated by the model. Collaborative Filtering for Building recommendation systems by learning latent representations of user preferences[10][11].

3) Early Transformer Models

The Transformer architecture, introduced by Vaswani et al. in their 2017 paper "Attention Is All You Need," was initially developed for machine translation tasks but quickly demonstrated its potential for generative applications. The core innovation of the Transformer is the self-attention mechanism, which allows the model to weigh the importance of different words in the input sequence when processing a particular word[12].

Unlike recurrent neural networks (RNNs) that process data sequentially, Transformers can process all tokens in a sequence in parallel, leading to significant gains in training efficiency. Early generative applications of Transformers, such as the Generative Pre-trained Transformer (GPT) by OpenAI, utilized a decoder-only architecture. These models were pre-trained on a massive corpus of text data to predict the next word in a sequence.

The self-attention mechanism is calculated as follows[13][14]:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (3)$$

where Q (Query), K (Key), and V (Value) are matrices derived from the input embeddings. This mechanism allows the model to capture long-range dependencies in the data, a significant advantage over previous architectures. The early Transformer models laid the foundation for the current generation of LLMs and were instrumental in advancing Text Generation for Producing coherent and contextually relevant text. Machine Translation for Translating text from one language to another with improved accuracy. Summarization for Generating concise summaries of long documents[15].

B. Emerging Methods in Generative AI

Building upon the foundational models, a new wave of generative methods has emerged, pushing the boundaries of what is possible. This section focuses on two of the most prominent and impactful emerging methods: Large Language Models and Diffusion Models.

1) Large Language Models (LLMs)

Large Language Models (LLMs) represent a significant leap in the evolution of generative AI, characterized by their massive scale in terms of parameters and training data. Models like OpenAI's GPT series (GPT-3, GPT-4), Google's PaLM, and Meta's LLaMA have demonstrated remarkable abilities that go far beyond simple text generation. The architecture of most modern LLMs is based on the Transformer decoder. The key to their success lies in the concept of scaling laws, which suggest that as the size of the model, the dataset, and the computational budget increase, the performance of the model improves in a predictable way[16][17].

A crucial aspect of LLMs is the phenomenon of emergent abilities. These are capabilities that are not explicitly programmed into the models but emerge as a result of their massive scale. Examples of emergent abilities include Few-shot and Zero-shot Learning which is The ability to perform tasks with only a few examples or no examples at all, simply by being prompted in natural language. In-context Learning is The ability to learn a new task from the context provided in the prompt. Chain-of-Thought Reasoning is The ability to break down complex problems into intermediate reasoning steps to arrive at a solution[18][9].

Training and Fine-tuning for The training of LLMs typically involves two stages:

- a) Pre-training: The model is trained on a vast and diverse corpus of text data in a self-supervised manner, usually to predict the next word in a sequence.
- b) Fine-tuning: The pre-trained model is then fine-tuned on a smaller, more specific dataset to adapt it to a particular task or to align it with human preferences. Techniques like Reinforcement Learning from Human Feedback (RLHF) are often used in this stage to make the models more helpful, harmless, and honest.

The applications of LLMs are vast and continue to expand rapidly: Conversational AI for Powering sophisticated chatbots and virtual assistants. Content Creation: Generating articles, emails, scripts, and other forms of creative writing. Code Generation for Assisting developers by writing, debugging, and explaining code. Knowledge Retrieval for Answering complex

questions and providing explanations on a wide range of topics[6][11].

2) Diffusion Models

Diffusion Models have recently emerged as a powerful class of generative models, particularly for high-fidelity image synthesis. Inspired by non-equilibrium thermodynamics, these models work by systematically adding noise to the data in a "forward process" and then learning to reverse this process to generate new data from noise.

The process consists of two main parts:

- a) Forward Process (Diffusion): This is a fixed process where Gaussian noise is gradually added to an image over a series of time steps, eventually transforming it into pure noise.
- b) Reverse Process (Denoising): A neural network, typically a U-Net architecture with attention mechanisms, is trained to reverse the diffusion process. It learns to predict the noise that was added at each time step and subtract it from the noisy image, thereby recovering the original image.

To generate a new image, the model starts with a random noise tensor and iteratively applies the learned denoising process to produce a clean, high-quality image. The objective function for training the reverse process is to minimize the difference between the predicted noise and the actual noise added at each time step.

Key advancements in Diffusion Models include Denoising Diffusion Probabilistic Models (DDPMs) which is The foundational framework for modern diffusion models. Latent Diffusion Models (LDMs); Instead of applying the diffusion process in the high-dimensional pixel space, LDMs operate in a lower-dimensional latent space, significantly reducing computational cost and enabling faster generation. This is the technology behind the popular Stable Diffusion model. Classifier-Free Guidance is A technique that allows for controlling the generation process without the need for a separate classifier, enabling text-to-image generation with high fidelity to the input prompt [14][17].

Diffusion Models have set new standards in Text-to-Image Synthesis for Generating stunningly detailed and creative images from textual descriptions. Image Editing and Inpainting for Modifying existing images or filling in missing parts in a realistic manner. Video Generation for Extending the principles of image diffusion to generate coherent video sequences.

Diffusion Models represent a class of generative systems founded upon a dual-process framework: a fixed forward process that systematically introduces Gaussian noise to data over a series of timesteps, and a learned reverse process that iteratively denoises it. The core principle is to train a neural network to reverse this corruption, thereby learning the underlying data distribution. To generate new data, the model starts with a tensor of pure random noise and applies this learned denoising process repeatedly until a clean, coherent sample is formed. The predominant architecture for this task is the U-Net, whose encoder-decoder structure and critical skip connections are exceptionally effective at predicting the noise component in an image at any given timestep while preserving high-frequency details. The practical efficacy and state-

of-the-art performance of these models, however, are largely attributable to several key innovations[18]. Latent Diffusion Models (LDMs) dramatically increase computational efficiency by performing the diffusion process in a compressed latent space created by a Variational Autoencoder (VAE). To control the output, Classifier-Free Guidance (CFG) provides a powerful mechanism to steer the generation towards a user-provided condition, such as a text prompt, by interpolating between a conditional and an unconditional prediction. Finally, the integration of cross-attention mechanisms within the U-Net allows the model to align specific semantic elements of the prompt with corresponding spatial regions of the image, enabling complex compositional understanding. This powerful synthesis of probabilistic theory, an optimized neural architecture, and sophisticated conditioning mechanisms has established Diffusion Models as a leading paradigm for high-fidelity generative tasks.

3. Comparative Analysis and Discussion

The evolution from foundational models to emerging methods reflects a continuous search for greater generative power, efficiency, and control. This section provides a comparative analysis of the models discussed and discusses the broader trends in the field.

Table 1. Models comparisons

Model	Core Principle	Strengths	Weaknesses	Key Applications
GANs	Adversarial Training	High-quality sample generation, sharp images	Training instability, mode collapse	Image synthesis, style transfer
VAEs	Probabilistic Inference	Stable training, meaningful latent space	Blurry image generation, lower sample quality	Anomaly detection, data generation
Transformers (Early)	Self-Attention	Capturing long-range dependencies, parallel processing	High computational cost	Text generation, machine translation
LLMs	Scaling Transformers	Emergent abilities, few-shot learning, versatility	High computational cost, potential for bias and hallucination	Conversational AI, content creation
Diffusion Models	Iterative Denoising	State-of-the-art image quality, parallel, diversity	Slow sampling speed, high computational cost	Text-to-image synthesis, video generation

The evolution of generative AI is being driven by several powerful and interconnected trends that are defining the next frontier of artificial intelligence. Foremost among these is the pursuit of increasing model scale. The success of massive models has powerfully demonstrated that size is not just a quantitative metric but a qualitative one. The "scaling laws" have shown a

predictable relationship where increasing a model's parameters and training data unlocks emergent abilities complex reasoning, nuanced understanding, and few-shot learning capabilities that are simply absent in smaller models. This has spurred a race towards building ever-larger and more computationally intensive architectures, as researchers aim to discover the next threshold of cognitive-like performance.

Concurrent with this push for scale is the rapid move towards multimodality. The field is decisively shifting away from single-purpose models that operate on only text or images. The new vanguard of generative AI, such as Google's Gemini or OpenAI's GPT-4o, can seamlessly process, understand, and generate content across a rich spectrum of data types, including text, images, audio, and video. This allows for a more holistic and human-like interaction, enabling applications that were previously science fiction, such as generating a dynamic video from a simple text prompt or having a spoken conversation with an AI about a live video stream. This integration of senses is fundamental to creating models that can perceive and interact with the world in a more comprehensive manner.

However, as these models grow in power and complexity, the need for improved controllability and alignment becomes paramount. A raw, large-scale model can be unpredictable, generating outputs that may be factually incorrect, biased, or misaligned with human intentions. Consequently, a major research focus is on steering model behavior to be more helpful, harmless, and honest. Techniques like Reinforcement Learning from Human Feedback (RLHF) have become standard practice, where human preferences are used to fine-tune the model, effectively teaching it to better align with user goals and societal values. The ultimate aim is to create AI that is not just powerful but also reliable, steerable, and safe for widespread deployment[4][7].

Finally, counterbalancing the trend of massive, centralized models is a strong and vital push for democratization and efficiency. While state-of-the-art models require vast computational resources, a parallel effort is underway to develop smaller, highly efficient models that can run on consumer-grade hardware or even on-device. Through advanced techniques like quantization, pruning, and knowledge distillation, researchers are creating potent yet lightweight models, such as those from the Llama or Mistral families. This democratization makes powerful AI tools accessible to a broader range of developers and organizations, fostering widespread innovation while also addressing privacy concerns by enabling local data processing. This dual-track approach pioneering massive models while simultaneously optimizing for efficiency is creating a rich, diverse, and increasingly accessible AI ecosystem.

C. Future directions and Challenges

The rapid advancement of generative AI, while significant, is accompanied by a set of formidable challenges and promising research trajectories that will define the next phase of its evolution. These areas constitute critical domains of inquiry essential for the technology's responsible and effective integration into scientific, industrial, and societal frameworks.

A primary obstacle in the deployment of generative models is the phenomenon of hallucination, characterized by the generation

of factually inaccurate or nonsensical information. This issue is intrinsic to the probabilistic nature of current architectures, which are optimized for linguistic coherence rather than factual veracity. Consequently, ensuring the reliability of model outputs remains a persistent challenge. While mitigation strategies such as Retrieval-Augmented Generation (RAG) show promise by grounding models in external knowledge sources, they do not constitute a complete solution. The pursuit of verifiable factuality remains a central research objective.

Furthermore, the challenge of bias and fairness is deeply rooted in the vast, often uncensored, datasets used for training. Models inevitably internalize and risk amplifying the systemic societal biases present in their training corpora, which can manifest as stereotypical outputs and inequitable performance across demographic cohorts. Addressing this requires a multi-pronged approach that moves beyond simplistic data filtration to include the development of sophisticated algorithmic debiasing techniques, the establishment of robust fairness metrics, and the rigorous curation of more balanced and representative datasets.

The inherent opacity of large neural networks, often described as the "black box" problem, poses a significant barrier to interpretability and explainability. The sheer scale and complexity of these models make it exceedingly difficult to trace their internal decision-making processes. This lack of transparency impedes adoption in high-stakes domains, such as medicine and finance, where accountability and justification are paramount. The field of Explainable AI (XAI) seeks to develop methodologies to elucidate model behavior, yet achieving comprehensive, human-intelligible explanations remains a long-term and critical research goal for fostering trust and enabling robust system validation.

In parallel, the substantial computational cost and environmental impact associated with training large-scale generative models present a pressing sustainability concern. The energy expenditure and associated carbon footprint for developing a single state-of-the-art model are significant. This has catalyzed research into "Green AI," a domain focused on enhancing algorithmic efficiency, developing energy-aware hardware, and designing less resource-intensive model architectures to mitigate the ecological footprint of AI development.

A promising trajectory for future development lies in the creation of hybrid models. This paradigm involves architecting systems that synergistically combine the strengths of different generative architectures. For instance, leveraging the structured latent space of a Variational Autoencoder (VAE) within a Diffusion Model framework could yield high-fidelity outputs with significantly improved computational efficiency. Such composite systems may offer superior performance, controllability, and resource-efficiency compared to monolithic designs.

A long-term and ambitious research frontier is the development of World Models. This objective transcends mere pattern recognition and aims to imbue models with a causal understanding of real-world dynamics. A functional world model would be capable of reasoning about cause and effect, simulating future outcomes, and planning under uncertainty. The realization of this goal would mark a significant milestone towards more generalized forms of artificial intelligence, enabling systems to solve problems with a deeper level of reasoning.

The paradigm of personalized and continual learning represents a significant evolution from the current static nature of deployed models. This research area focuses on enabling models to dynamically update their knowledge base and adapt their behavior based on ongoing user interactions, without requiring complete retraining. The successful implementation of continual learning would facilitate the creation of highly personalized and context-aware AI assistants, although it presents non-trivial challenges, including catastrophic forgetting and the assurance of data privacy.

Finally, the integration with robotics and embodied AI constitutes a critical frontier for translating digital intelligence into physical action. In this context, generative models function as high-level cognitive engines, interpreting abstract commands and generating specific action sequences for robotic systems. This field is essential for developing autonomous agents capable of performing complex tasks in unstructured, dynamic environments, thereby bridging the gap between computational models and real-world interaction.

4. Conclusion

This comprehensive review has charted the evolutionary trajectory of generative artificial intelligence, from foundational models such as GANs, VAEs, and early Transformers to contemporary LLMs and Diffusion Models. Our analysis reveals a clear progression towards greater scale, capability, and multimodality, a development that is intrinsically linked to formidable challenges regarding factual accuracy, algorithmic bias, interpretability, and ethical application. The future of the field will be defined by research into hybrid architectures, world models, and continual learning, which promises to extend the frontiers of machine intelligence. This dual landscape of immense potential and significant risk underscores the critical need for robust, interdisciplinary governance to steer the responsible development of a technology whose profound societal impact is only beginning to unfold.

Reference

- [1] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems* (pp. 2672-2680).
- [2] Kingma, D. P., & Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- [3] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems* (pp. 5998-6008).
- [4] Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- [5] Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein generative adversarial networks. In *International conference on machine learning* (pp. 214-223). PMLR.
- [6] Mirza, M., & Osindero, S. (2014). Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.
- [7] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. In *Advances in neural information processing systems* (pp. 1877-1901).

- [8] Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., ... & Lowe, R. (2022). Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- [9] Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems* (pp. 6840-6851).
- [10] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 10684-10695).
- [11] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- [12] Weidinger, L., Mellor, J., Rauh, M., Griffin, C., Uesato, J., Huang, P. S., ... & Gabriel, I. (2021). Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359*.
- [13] Huda, A., & Ardi, N. (2021). Predictive Analytic on Human Resource Department Data Based on Uncertain Numeric Features Classification. *International Journal of Interactive Mobile Technologies (ijIM)*, 15(08), pp. 172-181.
- [14] Ardi, N., Lubis, A. I. ., & Isnayanti. (2023). Decision Tree for Predicting the Mortality in Hemodialysis Patient with Diabetes. *Jurnal Minfo Polgan*, 12(1), 346-356.
- [15] Ardi, Noper, S. Supardianto, and Ahmadi Irmansyah Lubis. "Predicting missing value data on IEC TC10 datasets for dissolved gas analysis using tertius algorithm." *Journal of Applied Informatics and Computing* 7, no. 1 (2023): 50-56.
- [16] Ardi, N., Adri, M. and Azhar, N., 2022. Implementasi Arsitektur Hierarchical Model View Controller (HMVC) Dalam Portal Akademik. *Jurnal Teknik Komputer dan Informatika*, 1(1 | Agustus), pp.21-30.
- [17] Ardi, N. and Isnayanti, I., 2022. Implementasi Artificial Neural Network dalam Memprediksi Jumlah Peserta Les Bahasa Inggris Menggunakan Metode Back Propagation (Studi Kasus di Lembaga Kursus Global English). *Jurnal Teknik Komputer dan Informatika*, 1(1 | Agustus), pp.45-51.
- [18] N. Ardi and D. A. Azhari, "Analisis Aplikasi Lembur Berbasis Web menggunakan System Usability Scale (SUS) Pada PT. XYZ", *J. Ilm. Inform. Glob.*, vol. 16, no. 1, pp. 10-19, Apr. 2025.