Galatasaray University Managerial and Social Sciences Letters,
Vol.3 Issue.1 March/Mars 2025, pp. 6-15, e-ISSN: 2980-1575

Visual Harmony: GAN-Generated Artworks and Live Orchestral Performance in Perfect Sync*

Research Article / Araştırma Makalesi

Mustafa Atakan Kızıltan†

ABSTRACT

This study investigates the integration of Generative Adversarial Networks (GANs) in synchronizing AI-generated visual art with live orchestration, demonstrated through a groundbreaking concert at Istanbul's Harbiye Cemil Topuzlu Open-Air Theater on July 14, 2022. By leveraging advanced audio analysis techniques, the research outlines the methodologies employed to train GANs, produce high-quality visual content, and achieve real-time synchronization with live music. Key findings address both the technical and artistic challenges encountered, offering insights into the potential of GANs to enhance live performances. This work sets itself apart by emphasizing real-time interaction between music and visuals, enabled by precise and dynamic audio-to-visual mappings. In addition to advancing the state-of-the-art in performing arts technology, this study provides a unique perspective on the fusion of traditional artistry with modern AI tools. It contributes to the discourse on the evolving role of artificial intelligence in creative processes, highlighting its capacity to redefine the boundaries of artistic expression and audience engagement.

Keywords : *Art and Technology Integration, Generative Adversarial Networks (GAN), Digital Art Creation, Generative AI, Creative AI in the Arts, Deep Learning and Art.*

1. INTRODUCTION

The intersection of art and artificial intelligence (AI) not only transforms how art is created and experienced but also reignites fundamental debates about human nature and creativity. When comparing AI-generated artworks with those created by human hands, new questions arise about the universal language of art and its meaning as well as its value. This article aims to explore this dynamic interaction between AI and art through the lens of Generative Adversarial Networks (GANs), a deep learning methodology. By doing so, it seeks to illuminate both the theoretical and practical dimensions of this relationship.

Although the interplay between art and technology has existed throughout history, the convergence of AI and art opens the door to a new era that challenges the boundaries of creativity. Particularly, AI techniques such as deep learning are reshaping how traditional art forms meet the digital age, redefining both the perception and the definition of art. Models like GANs have played a pivotal role in this transformation, with advanced architectures such as StyleGAN2 producing high-quality visuals. However, the application of these technologies in performative and dynamic contexts remains relatively underexplored.

This study examines the evolution of art and AI, focusing on how technology influences creative processes. It highlights an innovative application: the real-time synchronization of AI-generated artworks with live orchestration. For this purpose, GAN models were trained separately for each segment of the concert, generating visuals that dynamically adapted to the music in real-time. It stands as a significant example of how art and AI can collaborate not only to broaden artistic horizons but also to redefine human creativity. By presenting a novel application of GAN technologies, this article

* ***Received:*** *14.04.2024 ;* ***Revised:*** *08.07.2024;* ***Accepted:*** *31.01.2025*

† KoçDigital, ORCID: 0009-0008-7764-1582, m.atakan.kiziltan@gmail.com.

underscores the transformative potential of AI in the arts and highlights its critical role in shaping the future of creative expression.

The article is structured as follows: first, it provides a review of the literature on the use of GANs and other AI technologies in art. Next, it details the innovative methods that distinguish this study from previous works. Finally, it examines the practical application of these methods during the performance held on July 14, 2022, at Istanbul’s Harbiye Cemil Topuzlu Open-Air Theatre, evaluating the potential of AI to expand the boundaries of art.

2. Literature Review

The intersection of artificial intelligence and art has sparked significant interest, prompting researchers to explore both the theoretical foundations and practical applications of this dynamic relationship. This section reviews the key concepts and methodologies relevant to understanding the subject of this study. Here are some necessary concepts:

- Machine learning, a foundational element of artificial intelligence, refers to a computer’s ability to learn from data to perform specific tasks. It involves algorithms that identify patterns in datasets and make predictions based on these patterns (Bishop, 2006).
- Deep learning, a subfield of machine learning, focuses on exploring deep structures in large datasets using complex algorithms like artificial neural networks (McCulloch et al., 1943). The origins of deep learning date back to 1943 when Walter Pitts and Warren McCulloch developed the concept of “threshold logic,” combining algorithms and mathematics to mimic the neural networks of the human brain. This pioneering work laid the foundation for modern advances in artificial intelligence and machine learning (Schmidhuber, 2022). Deep learning has proven particularly effective for handling complex data types, such as images, audio, and text.
- Datasets form the backbone of machine learning and deep learning models, providing the information needed for training algorithms. A high-quality dataset ensures accuracy and standardisation by adequately representing the problem to be solved (Jain et al., 2022).
- Model training involves iteratively adjusting parameters to enhance performance and enable algorithms to make precise predictions (Sarker, 2021).

Machine learning approaches can be categorized into supervised and unsupervised learning paradigms.

- Supervised learning uses labelled datasets, where the algorithm learns to map inputs to outputs (e.g., classifying images or detecting fraud).
- In contrast, unsupervised learning extracts patterns from unstructured data, identifying relationships without labelled inputs. Techniques such as clustering and dimensionality reduction are commonly applied in unsupervised learning (Sarker, 2021; Watson, 2023). GANs combine elements of unsupervised and supervised learning, utilizing a supervised loss function within an adversarial framework to train the generator and discriminator networks.

Understanding discriminative and generative models is crucial for grasping the distinctions between AI technologies.

- Discriminative models focus on producing outputs based on data features, such as classifying images or predicting values (Rautaray et al., 2020).
- Generative models, however, aim to model the joint distribution of inputs and outputs, enabling them to generate new, synthetic examples that resemble existing data points (Bishop, 2006). The primary objective of discriminative models is accurate classification or regression, while generative models seek to learn the data’s overall distribution and generate new instances based on that distribution.
- Among generative models, the Generative Adversarial Network (GAN) stands out as a groundbreaking approach, leveraging two neural networks -the generator and the discriminator- trained together to produce realistic synthetic data.

In recent years, generative art has been significantly influenced by the advent of text-to-image models, which utilize textual prompts to create visually coherent outputs. The release of DALL-E by OpenAI in 2021 marked a pivotal moment, as it showcased the potential of combining large language models with generative capabilities (Ramesh et al., 2021). This innovation was quickly followed by VQGAN-CLIP (Burgess, 2021), leveraging CLIP technology (Radford et al., 2021). Subsequent models like Midjourney (Rose, 2022), Stable Diffusion (Heikkilä, 2022), and Google's Imagen and Parti further expanded the field (Vincent, 2022). Notably, these models have been used extensively for static artistic creations, yet their application in performative and dynamic contexts remains underexplored. This gap highlights the novelty of using GAN-generated visuals in real-time synchronization with orchestral music, as demonstrated in this study.

GAN models, which are at the heart of this study, represent a significant advancement in generative modelling. They have demonstrated remarkable capabilities in creating high-quality, realistic images and have been widely adopted in artistic applications. While prior works have explored the use of GANs in static visual art, this study extends their application to dynamic, performative contexts. In the following section, the details of the GAN model and its implementation in this research will be discussed comprehensively.

3. *Generative Adversarial Network (GAN): From Basics to StyleGAN*

Generative networks are deep learning models that aim to learn the distribution of a dataset and produce new, original examples. They excel in unstructured data types such as images, sounds, and text by capturing the underlying patterns in the data and generating realistic outputs. Proposed in 2014 by Ian Goodfellow and his team, Generative Adversarial Networks (GANs) revolutionized generative modelling with their adversarial architecture. GAN consists of two primary components: Generator (G) produces data by transforming a random noise vector (z) and discriminator (D) distinguishes real data from fake data generated by G.

Figure 1 represents the conventional architecture of the GAN model:

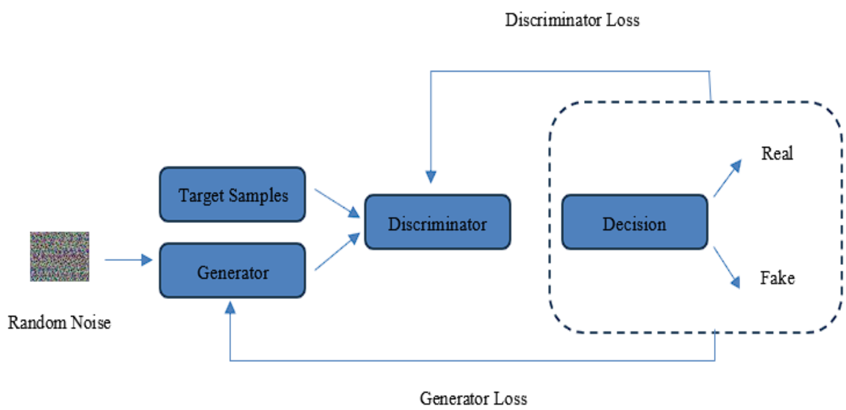

Figure 1. GAN model architecture

These networks engage in a competitive game where G learns to produce data that can fool D, while D becomes better at identifying real versus fake data. Over time, this adversarial process drives G to generate highly realistic data. GANs are fundamentally rooted in game theory, as they model the interaction between two players (G and D) as a zero-sum game. This adversarial setup mirrors the zero-sum games in game theory, where the total gains of the two players sum to zero. Thus, any gain by one player results in a corresponding loss by the other, illustrating the competitive nature of the training process (Owen, 2003). To recapitulate, objective of G is generating data that maximizes

D's likelihood of mistakenly identifying it as real and objective of D is correctly classifying real and generated data.

In brief, the workflow of the GAN architecture is as follows: A random vector z is sampled from a latent space, a multi-dimensional vector space, G transforms z into an output ($G(z)$), and, D evaluates both real data and $G(z)$, assigning a "real" or "fake" label. The weights of both networks are updated using the "minimax" loss function:

$$V(D, G) = E_{x \sim p_{data(x)}}[logD(x)] + E_{z \sim p_{z(z)}}\big[log(1 - D\big(G(z)\big))\big] \quad (1)$$

where $D(x)$ represents the discriminator's prediction of whether a real data sample x is real, $G(z)$ is the output of the generator for a noise vector z, creating synthetic images, E_x signifies the expected value of all real data samples, reflecting the typical features found in the dataset, E_z denotes the expected value of all noise vectors z fed into the generator, illustrating how it responds to different inputs and, $D\big(G(z)\big)$ is the discriminator's prediction of whether the image produced by G with noise z is real. In the loss function, the right-hand side of the summation operation shows the influence of the generator (G). This setup ensures that the generator cannot directly influence the expression $logD(x)$. The discriminator (D) aims to maximize its ability to distinguish between real and generated fake images, while the generator G seeks to minimize the discriminator's accuracy in identifying its generated images as fake.

Building upon the GAN architecture, Nvidia[1] introduced StyleGAN in 2019, which brought significant innovations in image generation, particularly in style-based synthesis and latent space manipulation (Karras et al., 2019). GANs traditionally begin with random noise vectors from the z-space. StyleGAN extends this concept by introducing an intermediate w-space through a mapping network, enhancing the semantic control over generated images. The z-space comprises purely random values, serving as the starting point for the image generation process. The mapping network -an 8-layer multi-layer perceptron (MLP)[2] - transforms the z-space vectors into the w-space. This transformation encodes more meaningful and semantically richer features, allowing for finer control over image properties. w-space allows for smooth transitions between images (low perceptual path length), meaning one image can morph naturally into another (Pham Van et al., 2020), and for better feature separation (high linear separability), which makes distinguishing between different image types easier (Karras et al., 2018; Karras et al., 2019).

On the other hand, the synthesis network takes the information from the mapping network and uses it to construct the final image. In other words, the style and content information provided by the mapping network is shaped into a visual representation by the synthesis network. This stage, part of the synthesis network, enables StyleGAN to manipulate images with a greater degree of precision and control. Specifically, attributes such as hair shape or pose can be independently adjusted and smoothly integrated with other features of the image, resulting in highly realistic transformations. This ability to modify specific features in a detailed and coherent manner distinguishes StyleGAN from other GAN variants, enhancing its flexibility and capability in generating complex and realistic images.

[1] Nvidia Corporation is a leading technology company known for its groundbreaking work in GPU (Graphics Processing Unit) development, which has revolutionized fields such as artificial intelligence, gaming, and scientific computing. Founded in 1993, the company has played a critical role in advancing AI research by providing high-performance hardware and software solutions. For more information, see Nvidia Corporation's official website: www.nvidia.com.

[2] "Multi-layer perceptron (MLP)" refers to a class of feedforward artificial neural networks composed of multiple layers of nodes, where each layer is fully connected to the subsequent one. An "8-layer MLP" indicates that the network consists of eight hidden layers of nodes, excluding the input and output layers, which are used for feature extraction and decision-making. For further details, see Goodfellow, Bengio, & Courville (2016), Deep Learning, MIT Press.

The synthesis network in StyleGAN incorporates Adaptive Instance Normalization (AdaIN), which adjusts style vectors layer-by-layer (Huang et al., 2017). The mathematical definition of AdaIN is:

$$AdaIN(x_i, y) = y_{s,i}\frac{x_i - \mu(x_i)}{\sigma(x_i)} + y_{b,i} \tag{2}$$

StyleGAN employs a process where activations in each layer of the synthesis network are normalized and then adjusted through scaling and shifting operations derived from style vectors. These style vectors, obtained from a mapping network, allow the model to control specific attributes such as pose, texture, and color, while preserving the overall structure of the image. The use of Adaptive Instance Normalization (AdaIN) facilitates fine-grained and flexible manipulation, enabling StyleGAN to generate high-quality and diverse images with remarkable precision and consistency.

4. Training Process of the GAN Models for Live Event

4.1. Training Process of the GAN Models

In the production of the artwork Pictures at an Exhibition[3], 11 distinct GAN models were used to represent 10 musical pieces and the promenade[4] sections. The training architecture primarily leveraged StyleGAN and StyleGAN2 algorithms (Karras et al., 2020). Depending on the availability and specificity of data, the models were trained either from scratch or using transfer learning techniques.

Transfer learning allowed GAN models to adapt pre-existing knowledge to new tasks, especially when the dataset size was limited. For example, pre-trained GAN models on human portraits were fine-tuned for tasks involving similar images. Conversely, when no relevant open-source pre-trained models existed or when the target data significantly differed, models were trained from scratch. This dual approach ensured both flexibility and optimal utilization of available resources.

To monitor and evaluate the training process, the Fréchet Inception Distance (FID) metric was employed (Heusel et al., 2017). Training was stopped when the FID score improvements plateaued, indicating that further iterations would yield negligible quality enhancements. Visual inspections of the generated images complemented this quantitative metric to ensure that the outputs met the desired quality standards.

For compositions in Pictures at an Exhibition like “Bydlo”, which features the depiction of a heavy Polish ox cart, the training dataset included drawings of large animals and carts. Due to the limited availability of such data, transfer learning was the primary approach. If the results were unsatisfactory, the models were subsequently retrained from scratch. This ensured the generated visuals to be thematically aligned with the specific music pieces, enhancing the audience’s experience during the performance. Additionally, care was taken to use only public domain[5] images for training, adhering to copyright regulations.

[3] Modest Mussorgsky’s “Pictures at an Exhibition” was composed by the Russian composer in 1874 for solo piano. This piece was dedicated to the memory of the Russian artist Viktor Hartmann, who died at the young age of 39 in 1873. For more details: Encyclopedia Brittanica, 2024.

[4] The term “Promenade” holds a specific orchestral meaning, referring to an interlude or movement that allows for a transition between pieces or scenes. It gained prominence in works such as Modest Mussorgsky’s Pictures at an Exhibition, where the “Promenade” sections serve as thematic bridges. For further reading, see Taruskin, R. (2009). The Oxford History of Western Music. Oxford University Press.

[5] The concept of “public domain” refers to creative content that is no longer protected by intellectual property rights such as copyright, trademark, or patent. In Europe, the copyright of a work lasts for 70 years after the death of the longest-living creator. If the copyright is held by a company, this period is determined as 70 years after publication. After this period of temporary protection expires, the work falls into the public domain (European Union Official Website, 2024).

4.2. Extraction of Music Features and Video Generation

The integration of music and visuals was achieved through advanced manipulation of GAN-generated images. Manipulation of deep generative models constitutes a burgeoning field of research itself (Broad et al., 2020).

Latent-space interpolation was employed to create seamless transitions between GAN outputs, a method that explores the relationship between points in the latent space (Mi et al., 2021). Given two latent vectors, ρ_1 and ρ_2, the transition was mathematically defined as:

$$\upsilon = (1 - \lambda)x\rho_1 + \lambda x\rho_2 \tag{3}$$

where $\lambda \in [0,1]$ controlling the interpolation process (Michelis et al., 2021).

To synchronize the visuals with live music, harmonic and percussive features were extracted using the librosa library[6] (McFee et al., 2015). Harmonic features influenced the tempo of the video transitions, while percussive features controlled the intensity. Additional features, such as tempo and magnitude, were also analyzed to enhance synchronization (Fitzgerald, 2010; Driedger et al., 2014).

Figure 2. A section from the concert

4.3. Real Time Music Integration

Music features were directly embedded into the GAN's latent space to augment the interpolation process. This novel technique allowed the generated visuals to dynamically respond to live music, ensuring a coherent and immersive audiovisual experience. The synchronization process, characterized by real-time integration without the need for external software, represents a significant advancement in the field of AI-driven art.

By dynamically combining harmonic and percussive features with GAN-generated visuals, the performance achieved a seamless alignment between audio and visual elements, pushing the boundaries of AI-generated art in live settings. This

[6] Librosa library is a powerful Python library specifically designed for music and audio analysis. Python, a widely-used high-level programming language, is known for its simplicity and readability, making it an ideal choice for a variety of applications, including data science and machine learning. A library in programming refers to a collection of pre-written code that developers can use to perform common tasks, thereby simplifying the development process and enhancing efficiency.

pioneering methodology provides a framework for future explorations in real-time AI art generation.

5. The Context of the Performance

The artwork "*Pictures at an Exhibition*" was performed live at the Harbiye Cemil Topuzlu Open-Air Theater[7] on July 14, 2022, accompanied by the Cemal Reşit Rey Symphony Orchestra[8] under the baton of conductor Murat Cem Orhan[9]. This performance brought together the timeless composition of Modest Mussorgsky with cutting-edge AI-generated visuals, creating a unique fusion of classical music and contemporary technology. We can see in Figure 2 and 3, photos taken during the performance to give you an idea of the result of the training.

Figure 3. A section from the concert

The concert received overwhelmingly positive feedback from both the orchestra members and the conductor, who praised the innovative integration of AI-generated visuals with live orchestration. Audience expressed their admiration for the unique experience, noting how the visual elements enhanced the overall performance. A detailed overview of the results of this project can be found in the interview with Cumhuriyet Gazetesi[10], highlighting the technical and artistic achievements of the event and the live interview on BloombergHT's Yapay Zeka Merkezi program[11] where the concert's significance and the potential of AI in the arts were elaborated. Although we have underlined here the technical processes applying to create images during the orchestral performance, this study and the event vividly demonstrate the evolution of art and the transformative role of technology in reshaping artistic creativity; therefore, this study stands in the current trend and discussions exploring the use of AI in creative pursuits.

---

[7] The capacity of the Istanbul Harbiye Cemil Topuzlu Open-Air Theater, which opened in 1947 and continues its activities to this day, is 4,532 people (Istanbul Metropolitan Municipality Culture Inc, 2024).

[8] Cemal Reşit Rey is recognized as one of the Turkish Five, making significant contributions to Turkish music history as a composer, educator, piano pedagogue, pianist, orchestra conductor, and founder of the Istanbul City Orchestra. The Cemal Reşit Rey Symphony Orchestra, established in Turkey, has successfully performed concerts both domestically and internationally (Istanbul Metropolitan Municipality Culture Inc, 2024).

[9] Murat Cem Orhan, the conductor of the Cemal Reşit Rey Symphony Orchestra and the general artistic director of the Cemal Reşit Rey Concert Hall, was born in Izmir in 1981 and completed his music education successfully. He has won numerous international awards, worked with renowned figures worldwide, and performed significant concerts. Additionally, he has been actively involved in musical events in Turkey and has successfully participated in various projects (Istanbul Metropolitan Municipality Culture Inc, 2024).

[10] https://www.cumhuriyet.com.tr/kultur-sanat/crr-senfoni-orkestrasinin-konserine-yapay-zeka-resimlerle-eslik-etti-1958883, Retrieved 19 June 2024

[11] https://www.youtube.com/watch?v=Wh1MMLhlsxA, Retrieved 19 June 2024

6. Conclusion

This article extensively explores the impressive results of merging traditional art forms with the innovative technologies of the digital age. A unique event is analysed, where AI-generated visuals created using Generative Adversarial Networks (GANs) were synchronized with live orchestration. This intersection of disciplines illustrates how AI can serve not just as a tool but as a medium, in its own right, contributing to the evolution of art as an expressive form.

A critical technical aspect discussed is the use of the Fréchet Inception Distance (FID) metric, which was instrumental in evaluating and ensuring the quality of the generated images. The marginal improvement in FID scores during the training process served as a stopping criterion, resulting in high-quality, visually appealing outputs. This underscores the importance of quantitative methods in achieving excellence in AI-generated art. While the methodology achieved remarkable results, it also highlights certain limitations, such as the dependency on curated datasets and computational resources, which could constrain accessibility for less-resourced artists or researchers.

The integration of AI into artistic processes signifies not merely an enhancement of existing practices but the emergence of new paradigms in art creation. Understanding this interplay between technology and creativity provides a compelling framework for envisioning the future of artistic practices.

For practitioners, the methodologies discussed -such as using transfer learning, monitoring FID scores, and designing real-time synchronization systems- serve as practical guidelines for exploring the integration of AI into live performances and exhibitions. These approaches exemplify how artists and technologists can collaborate to push the boundaries of both traditional and digital art, opening new avenues for innovation.

The broader societal implications of AI in art include its potential to democratize artistic creation and challenge traditional notions of creativity and authorship. As AI tools become more accessible, they enable a wider range of creators to experiment with innovative forms of expression, fostering inclusivity in the art world. However, these advancements also raise important questions about the ethical implications of AI's role in artistic creation, such as the balance between human and machine contributions and the authenticity of the creative process.

On a commercial level, AI-generated art offers new opportunities for personalized experiences, interactive installations, and dynamic audiovisual productions as well. According to McKinsey (2023), generative AI's ability to create and adapt content across modalities positions it as a transformative force across industries, with potential applications ranging from the arts to customer engagement and education.

In conclusion, this study illustrates how GANs and AI technologies can redefine artistic practices, offering innovative methodologies for integrating visuals and music. The concert exemplifies the potential of AI not only to enhance but also to expand the boundaries of creative expression. Moving forward, further exploration of advanced AI models and interdisciplinary collaboration will be essential in unlocking new possibilities and addressing the ethical and technical challenges of this evolving field. By embracing these opportunities, we can usher in a new era of creativity that harmoniously blends human ingenuity with technological innovation.

Information on Plagiarism
This article was scanned with plagiarism detection software. No plagiarism was detected.

Ethics Committee Approval Information
Ethics committee approval was not required.

Author Contribution Statement
The authors' contributions to this study are equal.

Funding Statement and other Acknowledgments
This study has not received any type of funding or support.

Competing Interests Statement
There is no conflict of interest to declare with any institution or person within the framework of the study.

REFERENCES

Aditya, R., Prafulla, D., Scott, G., & Ilya, S. (2021). Zero-shot text-to-image generation. https://doi.org/10.48550/arXiv.2102.12092

Bishop, C. M. (2006). Pattern recognition and machine learning. *Springer.* https://www.microsoft.com/en-us/research/uploads/prod/2006/01/Bishop-Pattern-Recognition-and-Machine-Learning-2006.pdf

Broad, T., Zhu, J.-Y., Efros, A. A., & Dovrat, O. (2020, May). Network bending: Expressive manipulation of deep generative models. *10th International Conference on Artificial Intelligence in Music, Sound, Art and Design.* https://doi.org/10.48550/arXiv.2005.12420

Burgess, P. (2024, June 19). Generating AI "art" with VQGAN+CLIP. *Adafruit.* Retrieved from https://www.adafruit.com

Driedger, J., Müller, M., & Disch, S. (2014). Extending harmonic-percussive separation of audio signals. *15th International Society for Music Information Retrieval Conference.*

Fitzgerald, D. (2010). Harmonic/percussive separation using median filtering. *13th International Conference on Digital Audio Effects (DAFX10).*

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. & Bengio, Y. (2014). Generative adversarial networks. *Advances in Neural Information Processing Systems, 2672-2680.* https://doi.org/10.48550/arXiv.1406.2661

Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. *MIT Press.*

Heikkilä, M. (2024, June 19). This artist is dominating AI-generated art. And he's not happy about it. *MIT Technology Review.* Retrieved from https://www.technologyreview.com

Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B. & Hochreiter, S. (2018). GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. *Advances in Neural Information Processing Systems 30.* https://doi.org/10.48550/arXiv.1706.08500

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9(8), 1735-1780.* https://doi.org/10.1162/neco.1997.9.8.1735

Huang, X., Liu, M.-Y., Belongie, S., & Kautz, J. (2017, March). Arbitrary style transfer in real-time with adaptive instance normalization. https://doi.org/10.48550/arXiv.1703.06868

Istanbul Metropolitan Municipality Culture Inc. (February 2024). Istanbul Municipal Theatre Official Website. Retrieved from https://sehirtiyatrolari.ibb.istanbul/sahne/harbiye-cemil-topuzlu-acikhava-tiyatrosu

Istanbul Metropolitan Municipality Culture Inc. (February 2024). Cemal Reşit Rey Concert Hall Official Website. Retrieved from https://crrkonsersalonu.ibb.istanbul/home/sitepage/17?CRRLang=tr-TR

Jain, A., Kalyanaraman, S., & Liu, J. (2020, August). Overview and importance of data quality for machine learning tasks. *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 3561-3562.* https://doi.org/10.1145/3394486.3406477

Karras, T., Laine, S., & Aila, T. (2018, February). Progressive growing of GANs for improved quality, stability, and variation. https://doi.org/10.48550/arXiv.1710.10196

Karras, T., Laine, S., & Aila, T. (2019, March). A style-based generator architecture for generative adversarial networks. https://doi.org/10.48550/arXiv.1812.04948