

Inverted Ekphrasis and Hallucinating Stochastic Parrots: Deleuzean Insights into AI and Art in Daily Life

Lenka Lee

lenka.lee@phil.muni.cz

Jakub Mácha

macha@phil.muni.cz

This study explores the potential of contemporary large language models (LLMs) to achieve Gilles Deleuze's goal of integrating art into daily life. Deleuze's philosophy, with its focus on creative repetition, finds a parallel in LLMs, which replicate and innovate artistic styles by transforming text prompts into various artistic expressions. Although LLMs can very effectively blend and mimic these styles, they remain mere tools – as suggested by the metaphor of a *stochastic parrot*; the true creative force remains the human artist. Currently, these models are experimental, and their artistic outputs, while impressive, are still viewed as pastiches rather than genuine art. Nonetheless, we can expect advances that may soon enable artists to use LLMs to create genuine artworks.

Keywords: large language models, Gilles Deleuze, art, everyday life

Inverted Ekphrasis and Hallucinating Stochastic Parrots: Deleuzean Insights into AI and Art in Daily Life

Lenka Lee

lenka.lee@phil.muni.cz

Jakub Mácha

macha@phil.muni.cz

This essay centres around a thought-provoking statement by Gilles Deleuze: «there is no other aesthetic problem than that of the insertion of art into everyday life».¹ The reverse process is more common: artists often take everyday objects and incorporate them into their works. For example, Picasso used found objects in his work, integrating newspapers and bottle labels into his collages and assemblages. In the groundbreaking *Still Life with Chair Caning* (1912), he added a piece of oilcloth with a chair caning pattern to the canvas, blurring the line between painting and sculpture.



A flow or transfer in the opposite direction – art into everyday life – is less common. We can see elements of that kind of transfer in artistic movements such as Bauhaus, De Stijl, Italian futurism, Czech poetism and

Figure 1 Pablo Picasso, *Still Life with Chair Caning*
Source:
<https://www.flickr.com/photos/lluisribes/9987801603>,
photo: Lluís Ribes Mateu, reprinted under CC BY-NC
2.0.

¹ G. Deleuze, *Difference and Repetition* (1968), translated by P. Patton, Columbia University Press, New York 1994, p. 293.

Pop Art. Typically, art integrated into everyday life has been seen as something supplementary, often as mere ornamentation or design. The challenge is how to incorporate art into everyday life without it being seen this way. This essay aims to show that artificial intelligence (AI), particularly in the form of the large language models (LLMs) that are becoming increasingly prevalent, can play a vital role in integrating art into significant aspects of everyday life. Before proceeding, we need to understand what motivated Deleuze's view of the relationship between art and everyday life. We will then outline the main principles of how LLMs work.²

1. Deleuze on art and everyday life

Deleuze's account of art is rooted in his multifaced notion of *repetition*. He says that «art repeats all the repetitions [...], art is simulation, it reverses copies into simulacra».³ The repetition Deleuze talks about is not based on fixed, pre-established models or forms, but on the variability introduced by all the previous repetitions. It is not a matter of repeating the same (form) but rather of repeating the differences that have already been repeated. Deleuze differentiates between a bare or brute repetition of perfect forms or patterns and a clothed repetition of imperfect appearances. Traditionally, following Plato, it has been held that perfect forms underpin imperfect appearances. Deleuze reverses this idea. In his framework, the clothed repetition of differing appearances, or simulacra, is foundational for the bare repetition of only seemingly identical patterns.

For Deleuze, repetition is closely linked to the creation of *style* because both involve generating difference. Repetition is a process that fosters novelty and variety, while style is the unique way in which these differences are expressed and perceived. Style, as a distinctive mode of expression, develops through repetition. By consistently engaging with certain ideas, themes or forms, an artist or writer refines their style, making it stand out. Therefore, style is not just a superficial trait, but an essential manifestation of underlying differential forces («style is essence itself, [...], essence is in itself difference»,

² Insights from the continental tradition, especially regarding language and the nature of creativity, can aid understanding of how current LLMs work and their potential applications. For instance, in a series of papers and blog posts, Gordon Hull has used Derrida's concept of the iterability of the sign to clarify how LLMs can generate meaningful text beyond their original context. Cf. G. Hull, *Transformer Models, Iterability, and Language (Part 1)*, in "New APPS: Art, Politics, Philosophy, Science", 10 July 2024, <https://www.newappsblog.com/2024/07/transformer-models-iterability-and-language-part-1.html>.

³ G. Deleuze, *Difference and Repetition*, p. 293.

Deleuze writes in *Proust and Signs*⁴). This essential characteristic of artistic style allows it to address the fundamental aspects of everyday life.

Deleuze maintains that *language* plays a crucial role in repetition: «Repetition is the power of language, and [...] it implies an always excessive Idea of poetry, [...] language organizes its entire system in the form of a clothed repetition».⁵ Moreover, «the genius of poetry identifies itself with [...] brute repetitions. Nevertheless, this genius primarily resides in the Idea and how it generates brute repetitions from a deeper, more concealed repetition».⁶ This passage emphasises the interplay between simple, bare and intricate clothed forms of repetition in language, particularly in the context of poetry. It posits that the repetition of linguistic units, such as expressions, words and syllables, constitutes a bare form of repetition, serving as a precursor to a deeper, more elaborate clothed repetition inherent in the essence of poetry.⁷ The example of written syllables repeated at the end of verse lines signifies a rudimentary scheme, always underpinned by a more intricate repetition in the auditory rendition of the poem, as in recitation or musical performance. Repetition in language serves as a model for all repetitions. In other words, any repetition can be seen as a form of linguistic repetition. Written language captures the basic aspect of every repetition by identifying or describing what is being repeated, essentially pointing out the repeated elements. (This understanding will be very important later when we focus on language models as a specific form of AI.) The concept of poetry captures the true relationship between these two forms of repetition. Bare repetition is just an abstract framework for a deeper, more complex form of clothed repetition.

These two models of repetition have their ramifications in *everyday life*. Bare repetition is manifested in stereotyped or standardised things or events, most notably in their reproduction as objects of consumption. In contrast, clothed repetition appears as «the instinctual series of destruction and death».⁸ The function of art involves addressing the distinction between two modes of repetition within everyday life. Specifically, art possesses the capacity to unveil subtle differences within stereotypical repetition, discrepancies that might evade notice in ordinary circumstances. Conversely, art can also

⁴ G. Deleuze, *Proust and Sign* (1964), translated by R. Howard, Athlone Press, London 2000, p. 48.

⁵ G. Deleuze, *Difference and Repetition*, pp. 291-292.

⁶ Ibid.

⁷ For example, the ABAB rhyme scheme is used in the Shakespearean sonnet form. In Sonnet 1, this pattern can be seen in the endings -ent, -ing, -ent, -ing. These are bare repetitions of identical letters in the written form.

⁸ G. Deleuze, *Difference and Repetition*, p. 293.

discern patterns of nearly identical repetitions amidst the flux of life, where all things are subject to decay and mortality. Ultimately, the role of art is to address the «ontological repetition» that distributes the difference between stereotypical and incremental (or, rather, decaying) repetitions.

2. An outline of how LLMs work

The idea we want to convey in this essay is that this Deleuzean framework can be observed in the way modern LLMs function. Before we explain how LLMs work, we want to highlight a few important caveats. Our opinion of LLMs (or AI systems in general) is quite modest and somewhat sceptical. AI systems are artificial machines designed to execute prompts and code created by humans. They do not possess a mind or subjectivity in any relevant sense. Any inner states they may have are computational states. If there is any subjectivity involved, it belongs to the designers (programmers) and the users. Although the metaphor of LLMs as (nothing but) *stochastic parrots*⁹ has been much ridiculed, it does capture much of our perspective: these models reproduce patterns and associations they have learned from vast datasets without truly understanding them. The metaphor is often used to suggest that LLMs, like stochastic parrots, lack creativity or the ability to produce new ideas. However, we disagree with this conclusion. If we reconsider how LLMs repeat what they have learned in the Deleuzean sense, it becomes apparent that they can actually create something new (or even *surprisingly* new), including original art. That is what we intend to argue here.

Let us take a closer look at LLMs and, more specifically, generative pretrained transformers (GPTs).¹⁰ LLMs/GPTs employ a complex process of analysing and generating text, which is based on their representation of the extensive datasets they are trained on. The training for models like ChatGPT occurs in two main phases. In the first phase, a large amount of existing text is utilised. An initial segment of this text is fed to

⁹ Cf. E. M. Bender et al., *On the Dangers of Stochastic Parrots. Can Language Models Be Too Big?* in Association for Computing Machinery, *FAccT '21. Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Association for Computing Machinery, New York 2021, pp. 610-623, <https://doi.org/10.1145/3442188.3445922>.

¹⁰ For comprehensive overviews, see G. Yenduri et al., *Generative Pre-Trained Transformer. A Comprehensive Review on Enabling Technologies, Potential Applications, Emerging Challenges, and Future Directions*, in “arXiv”, 2305.10435, 2023, <https://doi.org/10.48550/arXiv.2305.10435>, or P. P. Ray, *ChatGPT. A Comprehensive Review on Background, Applications, Key Challenges, Bias, Ethics, Limitations and Future Scope*, in “Internet of Things and Cyber-Physical Systems”, 3, 2023, pp. 121-154, <https://doi.org/10.1016/j.iotcps.2023.04.003>.

the model, and then the model assigns probabilities to possible continuations of that segment. This textual content is typically obtained from public online sources such as Wikipedia, news articles or books. The important thing to note is that data can be used for training regardless of their truthfulness, reference to real objects or even meaningfulness. Both false and true data can be equally useful; fiction can be just as valuable as brute facts, a well-argued essay as good as random nonsense from the internet. In the course of the training process, training data are transformed into *embedding matrices* and other *parameter tensors*, also called the *model weights*. These are high-dimensional spaces where tokens (words, phrases or even entire sentences) are mapped to vectors of real numbers. This method allows semantic and syntactic information about the language elements they represent to be encapsulated. Essentially, the tensors serve as the model's understanding of language, with similar meanings being placed closer together in the vector space. This enables the model to infer context, make predictions and generate coherent responses based on the input it receives.¹¹ Hence, the parameter tensors represent language rather than the world. To put it another way, LLMs are, as the name suggests, models of language, not models of the world or of true facts.

A model's parameters are adjusted using a process called gradient descent to improve the accuracy of its predictions. Over time, the model becomes adept at predicting the next words in a text, but it still lacks conversational abilities. This is addressed in the second phase of training, known as «fine-tuning», in which the model is further adjusted to improve its conversational capabilities. The fine-tuning phase involves receiving feedback from humans or other AIs to rate different responses for quality and encouraging the model to generate higher-quality responses. The predominant approach is known as reinforcement learning from human feedback (RLHF).¹² Other methods include AI-driven feedback and direct preference optimisation. During and post fine-tuning, humans evaluate the outputs, annotate texts and correct errors, guiding the model to align with human preferences and expectations. These steps are crucial for understanding the

¹¹ Each row of the matrix corresponds to a unique token (word) in the vocabulary. The size of the vocabulary determines the number of rows in the embedding matrix. Each column represents a feature in the embedding space. The dimensionality of the embedding space is a parameter chosen during the training or modelling phase and determines the number of columns. When a token (e.g. a word in a sentence) is input into a model that uses embeddings, the token is first mapped to its corresponding row in the embedding matrix. This row vector then serves as the feature vector for that token, encapsulating various semantic and syntactic properties.

¹² P. Christiano et al., *Deep Reinforcement Learning from Human Preferences*, in “arXiv”, 1706.03741, 2023, <https://doi.org/10.48550/arXiv.1706.03741>.

nanced relationships between language elements in the training data and their manifestation in parameter tensors.

As we will recall, GPT stands for «generative pretrained transformer». Being «*pretrained*» means the model weights are fixed after initial training. As of April 2024, the datasets for the most widely used models are at least a year old; for instance, GPT-3 was trained on data up to January 2022, and Claude 2 uses data available up until April 2023. We can, therefore, describe LLMs as inherently static language models.

Let us delve into the transformer architecture that forms the core of modern LLMs. For the purpose of this discussion, we shall focus on decoder-only models like GPT, although there are many variations. When a user inputs a prompt into a model, it predicts what should come next. For instance, if you input «The Critique of Pure», the model will assign a probability to each potential next token: «Nothing», «Chess», «Reason» and so on. In order to calculate these probabilities, each word is transformed into an initial embedding, which is a vector (a long list of numbers) that represents «The», «Critique», «of» and «Pure». These initial embeddings contain information about the individual words, but they do not initially include any information about the words around them. For instance, the initial embedding for «Pure» does not indicate that «Critique» comes before it. These embeddings provide a thorough summary of the semantic characteristics of the token.

The initial embeddings for each token are multiplied by the embedding matrix, which transforms and updates them across multiple layers. At each layer, the embedding for a token is updated using the attention mechanism, a key component of the transformer architecture. The attention mechanism enables a token's embedding to attend to earlier tokens in the sequence. For instance, let us say the word «Pure» is associated with «Critique» at a certain level. The representation of «Pure» would then be adjusted to show that «Reason» came just before it, or to somehow indicate that «Critique» was the primary topic. In essence, after considering the information from previous words and its own meaning, the representation of «Critique» is updated to include contextual details about the words around it in the given text. This updated representation is called a contextual embedding. After new embeddings have been generated, they are further refined using additional processing (perceptrons) before being input into a new layer. Once the embeddings have passed through all the layers of the model, the final contextual embedding for a token is used by the model to predict the next token. To generate more

text, a token with a relatively high probability can be selected and added to the initial prompt, and the new augmented sequence can be fed into the model again.

At the final stage of the model's output, there is a parameter referred to as «temperature» that influences the randomness or predictability of the output generated by the model. The temperature adjustment is a post-processing step applied only once to the final logits generated by the model. The higher the temperature, the more likely it is that a token with a lower probability computed by the model will be selected.

For the purposes of our discussion, what we can learn from these technical details is that LLMs are trained primarily to predict and generate text sequences that users find compelling or agreeable. These models are not designed to accurately represent the real world or convey factual truth. Instead, they focus on producing text continuations that seem plausible or appealing in response to given prompts. Additionally, LLMs interact with the world solely through text-based representations; they operate with predefined embeddings for concepts like «cat» or «joy», manipulating these embeddings in context without having real-world experiences or perceptions of these concepts.

3. How LLMs fit into Deleuze's account of repetition

In this section, we shall argue that LLMs are tools that can enable artists to incorporate art into everyday life. We will begin with a few initial observations. Firstly, it is crucial to note that the internal functioning of LLMs involves repeating the initial state through a series of computations. However, these computations introduce variations in what is being repeated. Therefore, we can assert that LLMs *repeat* in Deleuze's sense of the word. Secondly, because of human fine-tuning, LLMs produce results that appear natural from a human standpoint. In essence, LLMs' results integrate naturally into *everyday life*. Thirdly, LLMs do not necessarily provide factually accurate outputs. Even if the result is a statement like «the library closes at 8 p.m. tonight», we should not automatically assume that it is true. Some people have likened the capabilities of LLMs to «hallucinating»,¹³ meaning that they generate information not present in their training data. However, it is important to note that this quality is also found in works of art. Just as we do not take

¹³ Cf. M. Chelli et al., *Hallucination Rates and Reference Accuracy of ChatGPT and Bard for Systematic Reviews. Comparative Analysis*, in “Journal of Medical Internet Research”, 26, 2024, <https://www.jmir.org/2024/1/e53164>.

everything in art as fact,¹⁴ we should not blindly accept everything generated by LLMs as truth. Hallucination, in this context, essentially means fiction or even creativity. Both LLMs and works of art are not intended to be accurate representations of the real world (of fact).

The discussion about the social impacts of LLMs is largely conducted in figurative terms. It is important to note that the metaphors of hallucinating and being a stochastic parrot are in tension. LLMs can either just replicate the information from their training data like a stochastic parrot, or they can generate something new by hallucinating. Therefore, while we regard these metaphors as apt, it is important to be cautious about drawing conclusions from them. In particular, their inherent anthropomorphism and zoomorphism may misleadingly imply that LLMs have independent agency.

Before we move on to a discussion of art, we have to tie up two loose ends in our argument. The first is that repetition is the power of language that implies the Idea of poetry; the other is that LLMs are models of language in its everyday use. The Idea of poetry is then implied in LLMs' calculations, which are, in our view, repetitions in Deleuze's sense. This implication has two important consequences: there is a clothed repetition behind LLMs' computations, and LLMs' outputs have a poetic dimension. This poetic dimension of language is thus integrated into everyday life.

We can say a bit more about what is supposed to be integrated into everyday life. An LLM cannot take some artwork (or a part of one) from its training dataset and simply insert it into everyday life. Nor can an LLM recombine parts of artworks into a new artwork and, as its output, insert that new work into everyday life. Such an artistic addition would not constitute an essential part of everyday life. Obviously, merely placing a work of art in a public space does not make it part of everyday life; more is required if it is not to remain a foreign object. In our view, the insertion must have the form of *artistic style*, which is, in Deleuze's framework, related to essence and repetition.

LLMs repeat and combine artistic styles, both individual and collective, that they learned from their training datasets. The styles thus generated can be applied to the objects that we encounter in everyday life. A reframing of an object in an artistic style is not a mere

¹⁴ Take, for example, the portrayal of Napoleon in Tolstoy's *War and Peace*. While much of the historical backdrop concerning Napoleon and the events of the Napoleonic Wars is accurate, Tolstoy's portrayal of Napoleon himself often veers into the realm of artistic interpretation rather than strict historical accuracy. Moreover, many, though not all, dialogues involving Napoleon are fictional. Of course, it would be strange to claim that Tolstoy was hallucinating.

ornament in the sense of embellishment but rather an ornament in its primordial Greek sense of the cosmos as an orderly arrangement. In other words, artistic style inserted into everyday life aligns with ontological repetition, which, as we already know, distributes the difference between stereotypical and declining repetitions.

In the remainder of this essay, we shall demonstrate how these ideas are put into practice based on real-life examples. In particular, we position LLMs as the latest addition to the lineage of technological tools for creating art. Most of these tools are or were mechanical devices – and such devices can produce unexpected or even surprising results. In the case of LLMs, this can be understood as a form of inverted ekphrasis. Finally, we provide concrete examples of addressing stereotypical and declining repetitions from art history. All of this makes it possible to show how LLM-generated artistic styles can be incorporated into everyday life.

4. The role of technology and surprise in artistic creation

In ancient times, artists employed tools such as styluses, chisels and hammers. Today, laser-engraving machines, 3D scanners, graphic design software and various forms of photography are used for artistic purposes. LLMs are part of this lineage.

Even primitive tools can yield something unexpected, unintended or surprising. Take a coffee grinder, for instance. Its product, a heap of ground coffee, can display shapes the user did not have in mind. The surprising character of a tool's output can be attributed, firstly, to the tool's inner complexity, which is not transparent to the user; secondly, to the inner element of randomness introduced by the tool; and thirdly, to uncontrollable external impacts. The surprising products of a coffee grinder are due to factors of this third kind. The output of an LLM can be surprising due to all three types of factor: the user is unaware of the model's inner complexity and opacity; there is randomness based on the parameter of temperature; and the training dataset is chosen haphazardly and may contain random biases. In our Deleuzean framework, it can be said that all three factors introduce imperceptible differences in a series of repetitions.

Let us visualise the setting: there is a prompt, a tool/machine and an output. The relationship between prompt and output can be considered an *inverted ekphrasis*. Ekphrasis traditionally refers to repetition across different artistic mediums, traditionally one that moves from visual art to a textual description, e.g. the description of Achilles'

shield in Homer's *Iliad*. Inverted ekphrasis proceeds in the opposite direction, from a text prompt to an image or another artwork, such as a poem. Inverted ekphrasis is not a completely new concept; there is a project called «Reverse Ekphrasis Art Show» in which visual artists respond to the poem «When Her Body Is a Battleground».¹⁵

Inverted ekphrasis has significant repercussions in the domain of AI. *Prompt engineering* is a field that combines artistic expression with technical knowledge, requiring both creativity and logical thinking. Prompt engineering encompasses myriad coding-adjacent techniques. Here we want to suggest a few ways it can be employed in art.

The first can be called *thematic elaboration*, where LLMs are used to expand on a basic theme or concept by generating detailed descriptions or narratives. For instance, a simple prompt like «ocean sunset» can be enriched to include specific details such as the colours of the sky, the reflection on the water or the presence of any characters or objects. Another technique is *style mimicry*, where LLMs are used to generate artistic concepts or styles that mimic famous artists or art movements. By refining prompts to include specific stylistic elements – like Van Gogh's brushstroke techniques or the surrealism of Salvador Dalí – artists can generate a unique artwork that resonates with the characteristics of historical art movements.

Next, LLMs can create *interactive art projects*, where the audience can suggest modifications or continuations of a piece. For example, viewers might suggest new elements to add to an in-progress digital artwork, and the LLM can integrate these suggestions in real time, allowing for a dynamic and participatory creative process. In *narrative art*, prompts can be designed to tell a story through a sequence of artworks. By guiding the LLM with a script or narrative outline, artists can create a series of images that together tell a story. Furthermore, prompts can be tailored to evoke *specific emotions* through art, for instance by utilising elements known to induce feelings of calmness, excitement or nostalgia.

And finally, LLMs can be leveraged to create artworks that embrace *cultural elements*. Prompts can include specific instructions to incorporate traditional motifs and symbols that refer to particular heritages and histories. These techniques illustrate how the sophisticated crafting of prompts can lead to innovative and enriched artistic outputs,

¹⁵ Source: <https://thousandlanguages.asu.edu/reverse-ekphrasis-art-show>.

making use of the powerful capabilities of LLMs to enhance the artist's creativity and artistic expression.

It is important to note that the source of creativity is the artist who crafts the prompt. The LLM tool contributes an element of surprise and randomness, introducing what can be described as Deleuzean repeated differences. Whether the outcome is merely a pedantic stylistic exercise or an original artwork depends entirely on the artist's subjectivity.

In this section, we have taken a broad Deleuzean perspective on how LLMs can be used to create art. In the next section, we will focus more narrowly on how LLMs can address two of the kinds of repetition discussed above and thereby incorporate art into everyday life.

5. Stereotypical and declining repetitions

We shall now turn to traditional art that confronts stereotypical and declining repetitions. Firstly, let us consider Andy Warhol's portrait series, including his *Marilyn Diptych* from 1962.



Figure 2 Andy Warhol's *Marilyn Diptych*¹⁶

Deleuze does not make direct reference to this diptych. However, it is entirely consistent with his description of «Warhol's remarkable "serial" series, in which all the repetitions of habit, memory, and death are conjugated».¹⁷ The two-part painting presents a series of repetitions that are interrelated and displaced in various ways. In this machine-made series of portraits, the face of Marilyn Monroe is depicted over and over again in semi-photorealistic style. The artist employs a stereotypical repetition of the subject's face, yet each repetition is imbued with distinct chromatic nuances. Moreover, the repeated depictions of Monroe's face are displaced within the series, creating a visually striking effect. These repetitions resonate with other repetitions of various kinds, not only from the history of art but also from media and everyday advertising. Moreover, the right-hand half of the painting evokes the trajectory of Monroe's life and her demise in 1962, the same year the painting was created. It is evident that the stereotypical repetition of her face is confronted with a clothed repetition of the same face, which, however, gradually disintegrates. The subject of the painting, a prominent celebrity of the era, was a familiar presence in the context of daily life. Warhol reinterprets this subject with his distinctive style, thereby integrating art into the fabric of everyday life.

¹⁶ Source: <https://www.flickr.com/photos/rocor/47080722114>, photo: Rob Coder, reprinted under CC BY-NC 2.0.

¹⁷ G. Deleuze, *Difference and Repetition*, p. 294.

In the remainder of this essay, we shall explore whether and how LLMs can integrate art into everyday life. Specifically, can LLMs incorporate artistic style into everyday life in a manner similar to Warhol? LLMs are not equivalent to human artists like Warhol, but rather contemporary versions of the tools he used, such as silkscreen printing machines. LLMs alone cannot produce anything, just as silkscreen machines require an operator to function. But they can assist artists in incorporating their artistic style into everyday life. As discussed above, LLMs are stochastic parrots, that is, they are essentially replicating machines. However, they are also capable of hallucinating, that is, producing unexpected and surprising outputs. The objection to the use of mechanical devices in art, and to the resulting detachment from traditional artistic values, is not new. We can point out several art critics who made criticisms of the film industry and Pop Art (among other movements) along these lines.¹⁸ However, such criticisms lose their relevance if we adopt a Deleuzian perspective that prioritises repetition and do not attribute creative agency to LLMs alone. Since this technology is quite new, we cannot expect to find real artworks yet. However, we can show how LLMs can replicate existing artistic styles. The outcome is not a genuine artwork, but rather a playful experiment, a kind of pastiche. In fact, LLMs have shown surprisingly good results when it comes to replicating artistic styles.

Let us consider a few examples that illustrate a pastiche-like transposition of an artistic style using LLMs. The following image comes from a recent collection entitled *PromptDervish Pastiches: Pop Art Reimagined with AI*.

¹⁸ Walter Benjamin, in his essay *The Work of Art in the Age of Mechanical Reproduction*, describes the *aura* of an artwork as its unique presence in time and space, including its authenticity, originality and historical context. Mechanical reproduction, for instance in photography and film, diminishes the aura because it detaches the artwork from its unique context. Reproduced works can be experienced anywhere and anytime, leading to a loss of their singular nature and intrinsic value (W. Benjamin, *The Work of Art in the Age of Mechanical Reproduction* (1935), in H. Arendt (ed.) *Illuminations*, translated by H. Zohn, Schocken Books, New York 1968, pp. 217-251). Benjamin's conception of the loss of authenticity in analogue photography and cinema was later revised by Lev Manovich, who defined the digital technology of reproduction as based on varying mathematical algorithms. In line with Manovich's views, Michaela Fišerová suggests that «the modular, automatic and variable side of the digital medium *principally predisposes* every digital photograph for a cybernetic *bricolage*» (M. Fišerová, *Touching and Retouching. Question of Authenticity in Social Networking*, in S. Wrobel and K. Skonieczny (eds) *Living and Thinking in the Postdigital World*, Universitas, Kraków 2021, pp. 111-126, cit. p. 124). In his essay *Pop Art, Culture, and Rebellion*, Max Kozloff offers a pointed critique of the mechanisation in Pop Art. Kozloff argues that the movement's reliance on mechanical processes like silkscreen printing and the use of commercial imagery resulted in a detachment from traditional artistic values. He believed that this mechanisation led to a superficial engagement with art, in which emotional and intellectual depth was sacrificed for visual impact and marketability. Kozloff saw this as a commodification of art, reducing it to the status of a commercial product rather than a unique, expressive creation (M. Kozloff, *Pop Art, Culture, and Rebellion*, in "Artforum", 1(3), 1962, pp. 27-33).

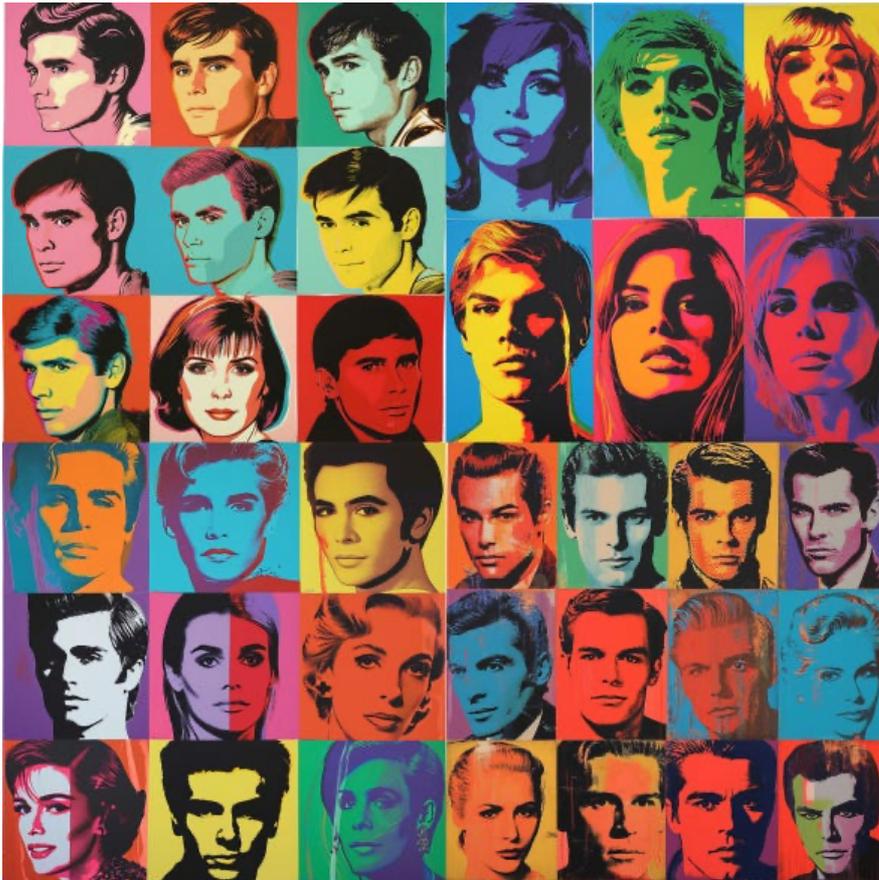


Figure 3 PromptDervish, from the collection *PromptDervish Pastiches: Pop Art Reimagined with AI*¹⁹

This image was generated by the LLM tool Midjourney using the following prompt: «A series of vibrant; stylized portraits of celebrities; each rendered with graphic sensibility and bold outlines; capturing the essence of fame and media in the styles of Andy Warhol and Ed Ruscha».²⁰

The LLM was tasked with combining the artistic styles of Warhol and Ruscha. While Warhol's style is easily recognisable, Ruscha's is less so. The creator of the pastiche explained his intention as follows: «It's an opportunity to reimagine how the essence of Pop Art can be reshaped and adapted to reflect the contemporary world. By doing so, we're not just revisiting a movement; we're revitalizing it, making it speak to today's audience in a language that resonates with their experiences and expectations».²¹ These words express an intention to insert artistic style into contemporary life, which seems to be in line with the

¹⁹ PromptDervish, *PromptDervish Pastiches. Pop Art Reimagined with AI*, in "Medium", 6 January 2024, <https://pub.aimind.so/promptdervish-pastiches-pop-art-reimagined-with-ai-050f81564a58>.

²⁰ Ibid.

²¹ Ibid.

Deleuzian project that this essay is concerned with. However, the result is, at best, a pastiche, not a true work of art like Warhol's *Marilyn Diptych*. In this pastiche, the artist's sole creative contribution is the prompt; what we have here is a case of inverted ekphrasis. We can recognise two repetitions in this pastiche. There is a stereotypical repetition of «stylized portraits of celebrities» as well as their clothing in Warhol's artistic style. Warhol's (and Ruscha's) style is not only stereotypically replicated but also simultaneously diminished through its use in pastiche.

6. Conclusion

Our investigation has been framed around Deleuze's aim to integrate art into everyday life. Deleuze's philosophy centres on the concept of creative repetition. Our thesis was that contemporary LLMs exhibit this kind of Deleuzian repetition. We therefore explored the extent to which LLMs can serve Deleuze's purpose of integrating art into daily life. It turns out that what can be inserted into everyday life is an artistic style. We have shown that LLMs excel at replicating and combining artistic styles, essentially transforming any prompt into a specified style. However, LLMs are mere machines, as captured by the metaphor of the stochastic parrot. The true origin of artistic creativity lies with the artist who provides the input, which LLMs then transform. LLM tools are still being developed, so their use at this stage is rather experimental. Our examples of LLM-generated artistic styles are experimental pastiches and not genuine works of art. Nevertheless, LLMs do have the potential to generate artworks in the hands of a real artist. We can expect to encounter examples of such works in the near future.²²

²² This work was supported by the Czech Science Foundation, project number 23-06827S.