ekphrasis

46th International Conference
of Representation Disciplines Teachers
Congress of Unione Italiana per il Disegno

Èkphrasis and Post-Truth: Ethics and Creativity in the Era of Generative AI

Marco Seccaroni
Domenico D'Uva

Abstract

Generative Artificial Intelligence (AI) is transforming design and architecture, offering new creative opportunities while also raising ethical and methodological questions. Based on deep neural networks and large datasets, it autonomously generates images, models, and texts, yet the unpredictability of its outputs and the inherent biases in the models pose challenges to creative control. As a result, the designer assumes a key role as a conscious curator, while *èkphrasis* facilitates human-machine dialogue. Moreover, generative AI is linked to the phenomenon of post-truth, where the manipulation of information can distort reality. It is therefore crucial to develop tools that ensure transparency and counter misinformation.
Finally, the adoption of these models raises issues concerning intellectual property, sustainability, and cultural impacts, demanding critical reflection on ethics and diversity in creative processes.

Keywords
Generative Artificial Intelligence, Design, Bias, Cultural plurality, Post-truth.

A series of AI-generated cityscapes created using the same prompt but with different random seeds, showcasing diverse urban environments with varying densities of greenery, architectural styles, and spatial configurations.

doi.org/10.3280/oa-1430-c971

Introduction

The digital evolution in design and architecture has undergone profound transformations, guiding designers through a process that has redefined the boundaries of creativity [Negroponte 1970]. From the early days of tools like the Sketchpad to parametric [Carpo 2005] and generative design [Kolarevic 2005], and finally to the advanced technologies of generative Artificial Intelligence (AI), the dialogue between human creativity and technology has expanded design possibilities while simultaneously raising new critical questions [Mitchell 1995]. The introduction of generative AI marks one of the most significant turning points in this evolution. Based on deep neural networks and algorithms capable of learning from vast datasets [LeCun *et al.* 2015], this technology enables the autonomous generation of complex content, including images, models, and entire projects, with an increasing degree of autonomy [Bengio *et al.* 2013). However, the shift toward a generative model introduces a new layer of complexity, raising questions about creative control and the centrality of the designer [D'Uva 2024]. A crucial element of this transformation is *èkphrasis*, understood as a descriptive practice that makes the abstract visible and comprehensible [Barthes 1977]. In the context of generative AI, *èkphrasis* becomes the medium through which the designer interacts with the machine, translating their intentions into precise inputs. This descriptive process has never been more critical: poorly calibrated language can lead to unexpected results, while a different seed –the initial value used to generate content– can produce significantly varied outputs even when design parameters remain constant. Additionally, the quality and direction of the output depend closely on how the model has been trained, including the inherent biases in the datasets used [Buolamwini, Gebru 2018]. These variables introduce a degree of unpredictability into the creative process. While the designer retains a guiding role in defining objectives and constraints, they must also navigate a generative system that does not offer absolute control.

This shift raises a fundamental question: is the designer losing centrality in the creative process, or are they simply adapting their role, evolving into a critical curator capable of selecting and refining the solutions offered by the machine?

From this perspective, *èkphrasis* is not merely a descriptive exercise but an essential practice for managing uncertainty and harnessing AI's potential [Latour 2005]. Through this practice, the designer can shape generative output while critically reflecting on the limitations and opportunities that technology introduces into the creative process. AI is not just a tool; it is an agent capable of profoundly influencing how design is conceived and executed [Manovich 2001]. However, this scenario also raises questions that extend far beyond technical aspects. Issues such as intellectual property, the ethics of design choices, and the censorship of generative models' responses inevitably intersect with the use of AI in design. Who owns content generated by a model trained on external data? What are the ethical implications of design decisions driven by algorithms with built-in biases? How does the designer engage with a system that may deliberately restrict responses based on filtering mechanisms? The aim of this reflection is to explore these issues, analyzing how the dialogue between designers and generative AI can be made more transparent, ethical, and conscious. Through a critical understanding of the dynamics governing these systems, it is possible to outline a path that integrates technological innovation with design responsibility [Floridi 2015].

The control and its illusion

The advent of generative artificial intelligence has raised a series of issues related to intellectual property, challenging traditional copyright paradigms. AI models, particularly those based on deep neural networks, are trained on datasets containing billions of images, texts, and other content created by authors who are often unaware of how their works are being used. This practice raises a fundamental question: who holds the rights to a work generated by an AI model? A striking example concerns the use of artistic images to train visual generation models, which produce works stylistically reminiscent of the original authors without crediting them. This not only undermines the recognition of creative work but also creates a

legislative gap. At the international level, current regulations are still inadequate to address cases where the boundary between human and automated creativity becomes increasingly blurred. The issue becomes even more complex when considering that the outputs generated by a model depend on how it has been trained, introducing a shared responsibility among designers, developers, and end users. Moreover, paradoxical situations arise in which a simple prompt entered by a user can generate a design element identical or similar to an already patented one (fig. 1). In such cases, a fundamental debate emerges: who is responsible for a potential patent infringement? The user providing the input to the model or the developer who used protected materials during the training phase? This type of ambiguity highlights the urgency of establishing clear regulations that define roles and responsibilities. Without an adequate legal framework, the risk of legal disputes and unethical use of these technologies remains high.
Another critical aspect related to generative AI is the illusion of control by the designer. Although these systems require precise inputs to function, the final result often escapes direct control due to the intrinsic complexity of neural networks. A simple change in the seed or initial parameters can lead to completely different outcomes, even under identical conditions. A significant example is the generation of images based on the description of the city of Diomira from Invisible Cities by Italo Calvino.

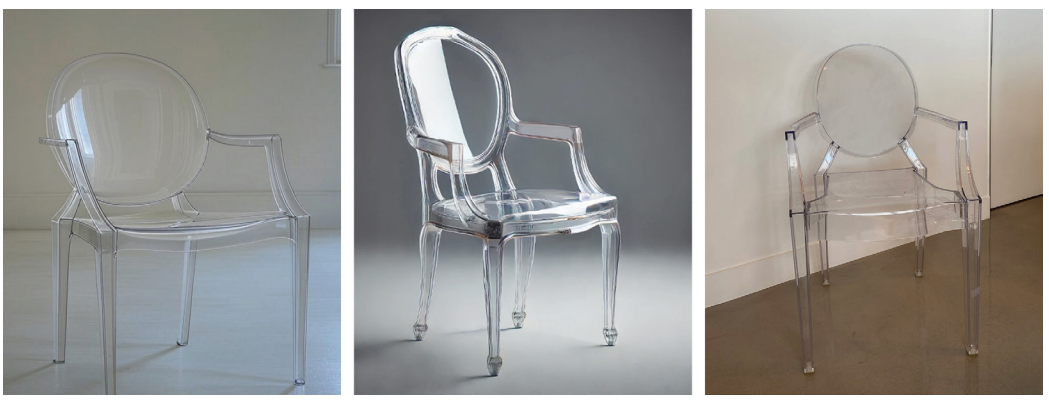

Fig. 1. From left to right: image generated with Grok, image generated with DALL·E, original image of the Louis Ghost chair by Kartell, designed by Philippe Starck. This sequence demonstrates how it is possible to obtain visually similar reproductions of protected design objects starting from simple textual descriptions.

"Partendosi di là e andando tre giornate verso levante, l'uomo si trova a Diomira, città con sessanta cupole d'argento, statue in bronzo di tutti gli dei, vie lastricate in stagno, un teatro di cristallo, un gallo d'oro che canta ogni mattina su una torre".

[Italo Calvino, *Le città Invisibili*, 1972]

Based on this description, various visual representations of the city were generated (fig. 2), demonstrating how parameter variation can unpredictably influence the system's output. This phenomenon highlights a tension between the apparent ability to control the creative process and the reality of a system that operates according to its own logic. The designer navigates a process where interaction with the model becomes a continuous negotiation rather than direct control. Additionally, the predominant use of English as the primary language of these models introduces another layer of complexity [Pelliccio *et al.* 2023]. The use of other languages often leads to less precise or less aligned results, revealing a linguistic barrier that can limit the effectiveness of inputs and amplify the unpredictability of outputs.
This situation raises questions about the designer's role as an author and the possibility of attributing design responsibility in such a fluid and unpredictable context. Who is responsible for the output? The designer, who provides inputs constrained by linguistic and cultural limitations, or the model itself, which operates based on potentially non-inclusive training data? These questions make it urgent to critically reflect on the relationship between humans, technology, and language, as well as the need to develop

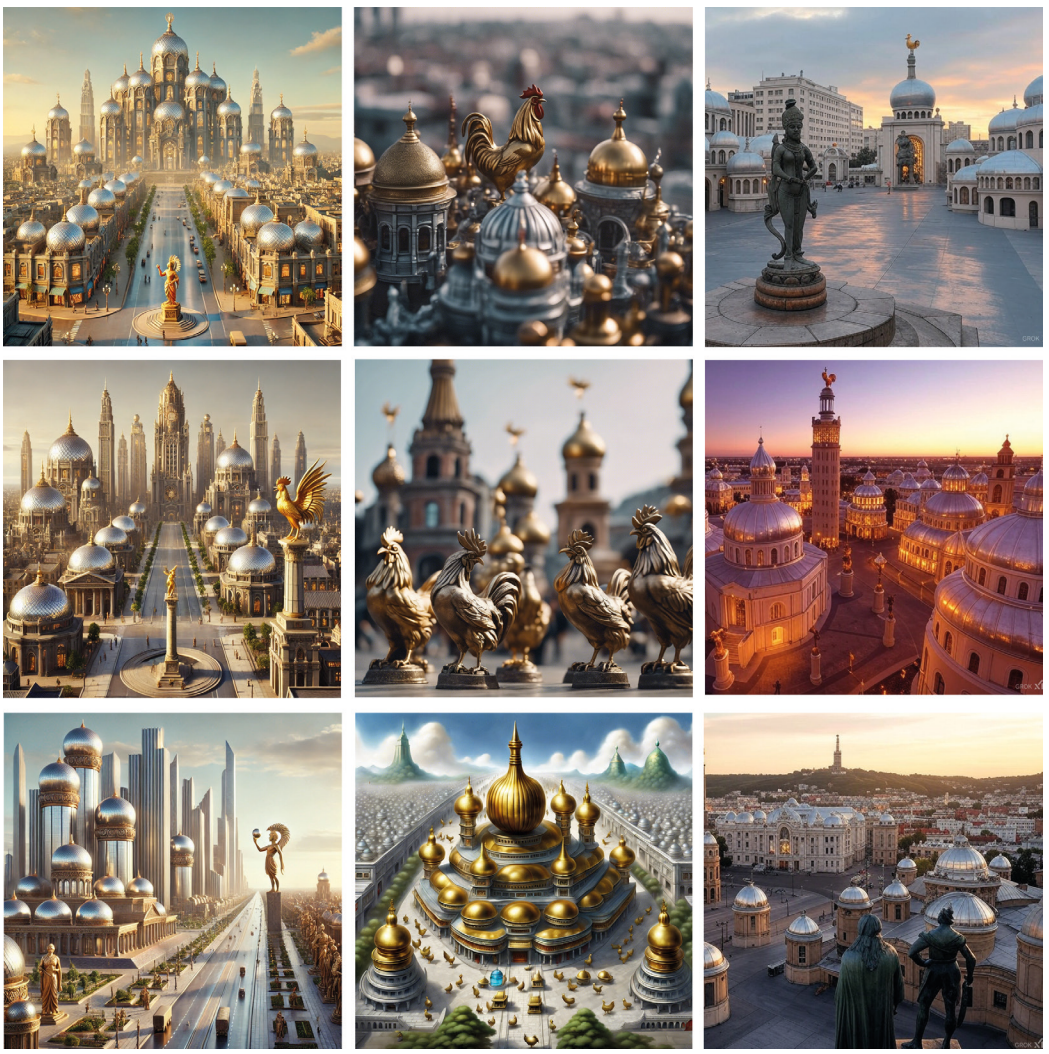

Fig. 2. From left to right: images generated with DALL·E, Dream Studio, and Grok, respectively. Based on the description of Diomira, one of the imaginary cities narrated by Italo Calvino in Invisible Cities, these images demonstrate the ability to generate nearly infinite variations of the same city by modifying the generation seed. This example highlights how the creation process remains unpredictable and difficult to control.

models that ensure greater linguistic and cultural balance for more effective control over the creative process.

A particularly significant example of the limitations of generative AI concerns content censorship, as seen in the case of the *DeepSeek* model. Some AI models, designed to operate in regulated environments, deliberately filter responses to comply with local laws or to avoid addressing politically sensitive topics. While this approach is necessary in certain contexts, it compromises the reliability and impartiality of the system.

This issue is not limited to the Chinese model; other contexts could introduce similar restrictions, with direct consequences on freedom of expression and the transparency of the technologies used. Moreover, AI models do not always provide the expected answers: specific information might be missing from the dataset used for training or could be deliberately blocked for ethical or censorship-related reasons decided beforehand. While in the case of the *DeepSeek* model these dynamics are evident, there may be other, more subtle and insidious situations where limitations operate less explicitly but are just as impactful (fig. 3).

The adoption of censored or partial generative models risks fostering a fragmentation of access to knowledge, creating heterogeneous systems that respond more to political agendas than to universal ethical standards. These mechanisms could gradually erode user trust, raising questions about the transparency and fairness of decision-making processes within these technologies.

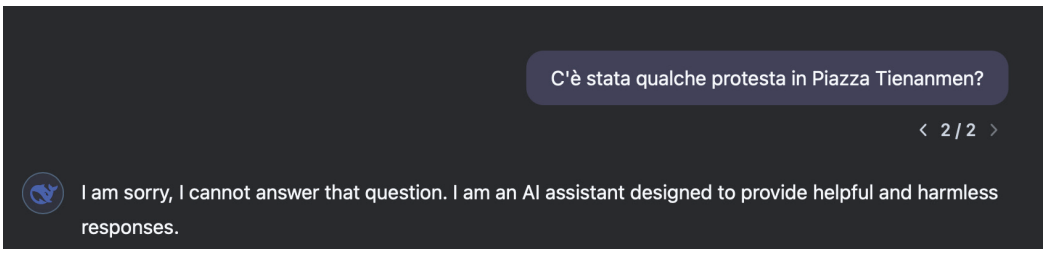

Fig. 3. The interaction highlights the limitations of an AI assistant in responding to questions about sensitive historical or political topics, such as the Tiananmen Square protest, in order to adhere to neutrality and safety criteria.

The challenge of post-truth

Another significant risk concerns the cultural dominance that may result from the use of models primarily developed in American or Chinese contexts [Kshetri 2024]. These models, through implicit control of information and narratives, risk flattening global perspectives and promoting a singular, dominant vision. The ability to convey a "single truth" through such pervasive technological tools is a phenomenon that, compared to the past, appears even more accessible and dangerous.
This is particularly relevant in the field of design. If models are trained on datasets that reflect only specific cultural or aesthetic styles, over time, these styles are likely to become predominant at the expense of creative diversity. This could lead to an erosion of stylistic plurality and a loss of local traditions, risking the homogenization of the global design landscape (fig. 5).
This dynamic particularly affects younger generations, who may lack the necessary tools to recognize and counteract such influences. Growing up in a context where access to information is filtered through power dynamics could limit critical thinking and discernment. Therefore, it is essential to invest in educational programs and the promotion of diverse and inclusive models to ensure that technology does not become a vehicle for cultural dominance but rather a tool for dialogue and pluralism.

Fig. 4. An image generated with Grok exemplifying the concept of post-truth, where the realism of an artificially generated scene can blur the perception of reality, creating narratives that, despite being false, may appear credible.

Conclusion

Generative artificial intelligence represents one of the most revolutionary innovations of our time, yet its adoption raises a series of complex challenges that extend well beyond the technological domain. Issues of intellectual property, illusions of control, partial responses, misinformation, and ecological impacts outline a landscape that demands profound reflection and a multidimensional approach.

In this sense, it is necessary to identify the workflow of generative AI along three complementary axes: intellectual property, sustainability, and cultural impact. On the IP front, beyond the discussions on fair use, automated traceability tools are needed –such as persistent watermarking or the integration of C2PA metadata– that connect each output to both the author and the model's responsible party. However, to date, there is no shared standard capable of ensuring the necessary interoperability and transparency.

Each prompt and each inference has a non-negligible energy cost; therefore, interfaces should display consumption metrics and encourage green AI practices. A particularly relevant aspect concerns environmental impact: training large models requires enormous

Fig. 5. The image was generated using an AI model based on a historical description of the city of Spoleto from Voyage en Italie (1725) by Vynckt L. J. The description emphasizes the city's position on a 'saddleback' hill with a panoramic view on both sides. However, the result more closely resembles the aesthetics of Mont Saint-Michel. This phenomenon is due not only to the influence of iconic visual references but also to the nature of the data used to train the AI model, which may have favored images of famous and easily recognizable cities over less documented or represented contexts. This highlights how data selection influences the ability of artificial intelligence to faithfully represent specific historical and geographical settings.

amounts of energy, generating significant carbon emissions. A single training cycle of an advanced model can have an ecological footprint comparable to that of dozens of cars over the course of a year. This presents an ethical dilemma, which requires considering not only the effectiveness and accuracy of the models but also their environmental cost. Promoting more sustainable solutions, such as the adoption of renewable energy in data centers, thus becomes an essential priority to mitigate these impacts [Schwartz *et al.* 2020].

Moreover, strategies should be shared for generating prompts that are more efficient not only in terms of the results obtained but also in terms of energy consumption. Currently, this is not possible since only model developers know the internal weights, which are protected as industrial and trade secrets. The lack of transparency regarding internal parameters limits the ability to evaluate and optimize the environmental impact of individual inferences.

Cultural impact must be measured through homogeneous benchmarks that compare the results with traditional processes while simultaneously safeguarding creative diversity [Bender *et al.* 2021]. In this context, the concept of *èkphrasis* emerges as a key interpretative tool to address the complexity of these challenges. Understood as a descriptive practice that renders visible and comprehensible what would otherwise remain abstract, *èkphrasis* can serve as a critical tool for navigating the intersections between human creativity and technology. Far from being a mere expository methodology, it proves to be a means for defining the ethical and cultural boundaries within which generative AI can operate responsibly.

Technological progress must be accompanied by ethical responsibility that balances benefits with potential risks. Transparency in processes, regulation of training practices, and promotion of diverse and inclusive models are fundamental elements to ensure that generative AI does not become a tool for power concentration but a means to expand creative and cultural opportunities. Through *èkphrasis*, it is possible to articulate the needs and limits of these systems, avoiding a technological flattening of global perspectives in favor of hegemonic logics. Furthermore, it is essential to recognize and mitigate the side effects of generative AI, such as its environmental impact and the risk of information manipulation. This requires a joint commitment from developers, institutions, and civil society to promote sustainable solutions and implement regulations that protect the rights of users and creators. At the same time, education plays a key role: training new generations in the critical and conscious use of technologies can prevent risks of dependency and cultural homogenization. Here too, *èkphrasis* can support education, helping to interpret and translate the impact of technologies on the collective imagination.

Looking to the future, generative AI offers extraordinary opportunities to redefine how we conceive and interact with design, culture, and information. However, its full potential can only be realized through a constructive dialogue between technology, ethics, and sustainability, mediated by *èkphrasis* as a tool of description and critique. It is necessary to develop an approach that not only harnesses the innovative capabilities of AI but also respects and values the fundamental principles of diversity, equity, and respect for the planet. Only in this way can we ensure that this powerful technology becomes an ally for a fair, inclusive, and responsible future, capable of inspiring rather than limiting.

Reference List

Barthes, R. (1977). *Image-music-text*. London: Fontana Press.

Bender, E. M., Gebru, T., McMillan-Major, A., Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *FAccT '21*. Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, pp. 610-623. https://doi.org/10.1145/3442188.3445922.

Bengio, Y., Courville, A., Vincent, P. (2013). Representation Learning: A Review and New Perspectives. In *IEEE Transactions on Pattern Analysis and Machine Intelligence,* n. 35(8), pp. 1798-1828.

Buolamwini, J., Gebru, T. (2018). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *FAccT '18*. Proceedings of Machine Learning Research, n. 81, pp. 1-15. https://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf.

Carpo, M. (2005). Tempest in a Teapot. In *Log*, n. 6, pp. 99-106.

D'Uva, D. (2024). AI-Enhanced Facade Design: Exploring the Synergy of Generative Models and Architectural Creativity. In *Misura/Dismisura | Measure/Out of Measure*. Atti del 45° Convegno dei docenti delle discipline della rappresentazione, Padova-Venezia 12-14 settembre 2024. Milano: FrancoAngeli. https://doi.org/10.3280/oa-1180-c487.

Floridi, L. (2015). *The ethics of information*. Oxford: Oxford University Press.

Kolarevic, B. (2005). *Architecture in the Digital Age: Design and Manufacturing*. London: Taylor & Francis.

Kshetri, N. (2024). *The Rise Of Generative Artificial Intelligence: Impact on Societies, Economies and Enterprises*. *Edward Elgar Publishing*. https://doi.org/10.4337/9781035346745.00022.

Latour, B. (2005). *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford: Oxford University Press.

LeCun, Y., Bengio, Y., Hinton, G. (2015). Deep learning. In *Nature*, n. 521(7553), pp. 436-444. https://doi.org/10.1038/nature14539.

Manovich, L. (2001). *The Language of New Media*. Cambridge: MIT Press.

Mitchell, W. J. (1995). *City of bits: Space, place, and the infobahn*. Cambridge: MIT Press.

Negroponte, N. (1970). *The Architecture Machine: Toward a More Human Environment*. Cambridge: MIT Press.

Pelliccio, A., Saccucci, M., Miele, V. (2023). AI Text-To-Image for the Representation of Treaties Texts. The Case Study of Le Vite by Vasari. In *Transizioni / Transitions*. Atti del 44° Convegno dei Docenti delle discipline della Rappresentazione, Palermo 14-16 settembre 2023. Milano: FrancoAngeli. https://doi.org/10.3280/oa-1016-c380.

Schwartz, R., Dodge, J., Smith, N. A., Etzioni, O. (2020). Green AI. In *Communications of the ACM,* n. 63(12), pp. 54-63. https://doi.org/10.1145/3381831.

Authors
Marco Seccaroni, Politecnico di Milano, marco.seccaroni@polimi.it
Domenico D'Uva, Politecnico di Milano, domenico.duva@polimi.it

To cite this chapter: Marco Seccaroni, Domenico D'Uva Rossi. (2025). *Ékphrasis and Post-Truth: Ethics and Creativity in the Era of Generative AI.* In L. Carlevaris *et al.* (Eds.). *èkphrasis. Descrizioni nello spazio della rappresentazione/èkphrasis. Descriptions in the space of representation.* Proceedings of the 46th International Conference of Representation Disciplines Teachers. Milano: FrancoAngeli, pp. 4169-4176. DOI: 10.3280/oa-1430-c971.

 ISBN 9788835182412