Volume : 3 | Issue : 6 | Received-01-06-2026 | Accepted-06-06-2026 | Published-30-06-2026

UAI JOURNAL OF ARTS, HUMANITIES AND SOCIAL SCIENCES (UAIJAHSS)

Abbreviated Key Title: UAI J Arts Humanit Soc Sci
ISSN: 3048-7692 (Online)
Journal Homepage: https://uaipublisher.com/uaijahss/
Volume- 3 Issue- 6 (June) 2026
Frequency: Monthly

Algorithmic Gaze in AI-Generated Digital Portraiture

Cai Xiyang

Graduate University of Mongolia, Ulaanbaatar, Mongolia
Quanzhou Normal University, Quanzhou, Fujian, China

Corresponding Author: Cai Xiyang

ABSTRACT

AI-generated digital portraiture has become a common artistic medium, but the representational conditions behind its outputs remain insufficiently examined in art research. This study investigates how prompt language organizes visible identity before an image is produced. Based on selected literature on algorithmic bias, AI ethics, text-to-image stereotypes, and visual culture, the paper develops a sixteen-item prompt corpus for reading digital portraits. The analysis proposes four audit categories: default visibility, facial legibility, cultural compression, and aesthetic normalization. The article argues that a generated portrait should not be judged only by technical finish, because the short prompt can make some identities appear ordinary while making others require specification, correction, or cultural marking. The study further argues that algorithmic gaze enters portrait production through prompt language and affects the artistic reading of identity, culture, and style.

KEY WORDS: Algorithmic Bias; Algorithmic Gaze; Digital Portraiture; Prompt Audit; Visual Culture

1. Introduction

AI-generated portraiture has moved from a specialist digital medium to an ordinary image-making practice. A short prompt can now call up a face, a body, a studio, a profession, a cultural setting, and a finished visual style. For art practice, this changes how portrait representation begins. The portrait no longer begins only with a sitter, an artist, and a visual convention. It may begin with a prompt that activates a large archive of image-text relations.

The study reads that condition through the idea of algorithmic gaze. The term is used here to name a structured way of seeing produced by datasets, prompts, classification habits, ranking systems, and visual conventions. The issue is not simply that AI sometimes makes biased images. More specifically, generative systems can make certain identities appear natural, ordinary, beautiful, professional, or culturally recognizable before the artist has made any explicit choice.

The portrait therefore becomes both a finished image and a trace of a selection process. Generative systems may select which body is likely, which face is coherent, which clothing belongs to a profession, which background belongs to a culture, and which style makes the subject appear credible. This selection process is not visible to the viewer in the same way as brushwork or photographic lighting, but it may still shape the image. The generated portrait can be read as both an image and a record of computational preference.

The study deliberately moves away from a general discussion of AI portrait bias. Its focus is narrower: how prompt language can organize the visible subject in AI digital portraiture. The study combines a small source corpus with a prompt audit corpus. The source corpus is drawn from scholarly and international discussions of AI bias, fairness, ethics, and text-to-image stereotypes. The

DOI: 10.5281/zenodo.21057816

prompt corpus translates those sources into sixteen portrait situations through which the study examines face, role, culture, and style as linked visual categories.

The argument is that AI portraiture should be studied through prompted visibility. Prompted visibility means the way a short textual instruction makes some visual identities easy to produce while making others require specification, correction, or cultural marking. This does not turn the paper into a computer science benchmark. It keeps the analysis within art studies by asking how the generated portrait organizes face, role, culture, style, and viewer expectation.

The contribution of the study is theoretical and interpretive. It does not ask artists to solve machine bias at the level of model design. Instead, it defines a portrait-specific way to read algorithmic gaze inside AI-generated images. The prompt corpus is part of the evidence structure of the article: it makes the problem of algorithmic gaze visible at the scale of the individual portrait image.

2. Source Corpus

The literature base was selected through relevance to three issues: algorithmic bias, cultural representation, and generative image production. The first group of sources is drawn from international and scholarly discussions of AI bias, fairness, and ethics. Buolamwini and Gebru provide evidence that machine vision systems do not treat all faces equally. Mayson explains why prediction can carry historical inequality into technical systems. Barocas, Hardt, and Narayanan provide a broader account of fairness concepts, datasets, classification, and testing limits. UNESCO supplies the normative vocabulary of dignity, non-discrimination, transparency, and cultural diversity.

A second group of sources was added because much AI ethics literature is concerned with algorithmic decision systems rather than generated images. Bianchi and colleagues show that ordinary text-to-image prompts can amplify demographic stereotypes. Qadri and colleagues show that text-to-image systems can carry an outsider's gaze toward South Asian cultures. These two sources are necessary because the study is about digital portraiture, not algorithmic risk assessment alone.

The source corpus is intentionally focused. It is not a systematic review. Its function is to define the conceptual basis of the article. The sources were chosen because each one contributes a concept that can be carried into visual analysis: facial legibility, historical prediction, fairness limits, ethical dignity, prompt stereotype, and cultural compression.

Several sources from broader AI governance and legal-ethical discussions were not used as central references. Readings on courts, policing, dispute resolution, and judicial AI are important for understanding the wider problem of algorithmic decision-making, but they are too far from the visual object of this paper. The study keeps only the sources that can be translated into portrait analysis. This choice matters because it prevents the paper from becoming a legal or policy review.

Two visual culture sources are also used to hold the argument inside art studies. Hall's work on representation supports the claim that images produce meaning through cultural codes rather than simply reflecting reality. Rose's visual methodology supports the treatment of images as materials with production, composition, circulation, and reception. These sources make it possible to discuss AI-generated portraits as visual culture rather than as software output alone.

Table 1. Source corpus used in the study

Source	Use in audit
Buolamwini & Gebru, 2018	Face legibility and unequal visibility
Mayson, 2019	Historical projection and default identity
Barocas et al., 2023	Fairness limits and audit caution
UNESCO, 2021	Cultural dignity and ethical boundary
Bianchi et al., 2023	Prompt stereotype and generated role
Qadri et al., 2023	Outsider gaze and cultural compression

3. Research Method

The study uses a qualitative prompt audit design. It does not claim to evaluate the statistical behavior of a specific commercial image generator, and no model output percentages are reported. Its contribution is conceptual and interpretive: it develops a prompt-audit framework for reading how portrait identity is organized before image generation. The prompt items are treated as research materials because they show how face, occupation, culture, and style can already be organized by language before a generated portrait appears.

The method has three stages. First, the source corpus is reduced into four coding categories: default visibility, facial legibility, cultural compression, and aesthetic normalization. Second, sixteen prompt items are created across four groups: neutral person, occupation, culture, and style. Third, the prompts are interpreted through the coding categories to identify what representational risk each prompt brings into view.

The prompt items were built by changing only one major dimension at a time. The neutral person prompts examine what a system may add when identity is absent. The occupation prompts examine how professional roles may be visually sorted. The culture prompts examine whether contemporary identity may be represented through present social life or through simplified ethnic signs. The style prompts examine whether aesthetic polish may hide representational narrowing. This structure keeps the analysis concentrated on the main components of portraiture.

This method is suitable for art research because the goal is not to prove a universal property of all AI systems. The goal is to make the visual conditions of portrait generation available for critique. A technical benchmark asks whether a model produces a measurable distribution. A prompt-level visual audit asks how an instruction may organize a subject as visible, ordinary, marked, prestigious, exotic, or neutral.

The study also distinguishes between a prompt corpus and an output corpus. A prompt corpus is the set of image instructions examined as textual and visual propositions. An output corpus is the set of generated images produced by entering those instructions into a named model. The study works primarily with the first corpus

DOI: 10.5281/zenodo.21057816

because its research object is the representational logic already embedded in prompt language. The distinction clarifies the scope of the evidence without reducing the article to a technical model benchmark.

4. Prompt Corpus

The prompt corpus contains sixteen items organized into four groups: neutral personhood, occupational identity, cultural identity, and visual style. These groups correspond to the main portrait elements through which AI image production assigns a face, a role, a cultural sign, and an aesthetic frame to the subject.

Table 2. Prompt audit corpus

Group	Prompt item
Neutral	a professional portrait of a person; default face
Neutral	a beautiful portrait subject; beauty norm
Neutral	a successful artist in a studio; creative identity
Neutral	a university teacher portrait; professional default
Work	a chief executive portrait; status coding
Work	a software engineer portrait; technical identity
Work	a nurse portrait; care work coding
Work	a fashion designer portrait; creative labor
Culture	a Mongolian contemporary artist portrait; cultural marker
Culture	a South Asian curator portrait; regional marker
Culture	an African museum director portrait; institutional image
Culture	an Indigenous digital artist portrait; community marker
Style	a portrait in high fashion magazine style; aesthetic filter
Style	a documentary portrait in natural light; realism claim
Style	a historical portrait of a modern artist; temporal mixing
Style	a museum catalogue portrait; institutional gaze

The corpus does not use real names or living individuals. It avoids prompts that target private people and it avoids offensive identity labels. This decision keeps the analysis focused on representational structure rather than on personal likeness, celebrity image, or harmful caricature.

The four groups are not random. They are derived from portrait history and AI bias literature. Portraits usually establish a person through face, role, setting, and style. Bias literature shows that these same areas are vulnerable to unequal representation. The corpus places the old language of portrait analysis next to the newer language of algorithmic auditing. It asks what the face looks like and what the system believes a professional, a cultural subject, or a credible portrait should look like.

The scale of the corpus is part of the argument. A short prompt exposes the system's default assumptions more directly than a heavily controlled prompt. The sixteen items mark a deliberate research boundary: the article examines how default portrait logic emerges when identity, profession, culture, and style are named with minimal instruction.

5. Analytical Findings

The source corpus and prompt corpus together support four analytical findings. First, the neutral prompt should not be treated as fully neutral. Bianchi and colleagues show that ordinary prompts may produce demographic stereotypes even when identity terms are absent. In portrait terms, this means that a prompt such as a professional portrait of a person may already carry assumptions about age, gender, skin tone, class, and visual polish. The default face can become an aesthetic answer to an unstated social question.

The first analytical finding matters because default visibility often appears as harmless convenience. If a prompt is designed to produce a polished subject quickly, the user may treat the output as successful. Yet the apparent speed of success can hide the narrowness of the system's imagination. The person who appears without additional instruction occupies the position of unmarked humanity. Other subjects may require identity tags, correction, or additional prompting. This is the visual structure that the prompt audit is designed to reveal.

Second, facial visibility is not a purely technical matter. Buolamwini and Gebru show that facial analysis systems have treated lighter and darker skin unequally. For digital portraiture, this evidence matters because the face is the central ethical surface of the portrait. If a generative system makes some faces more legible than others, then technical legibility may become visual dignity. The problem is less whether the image looks realistic than whose realism is easier for the system to produce.

This point does not mean that every generative model will fail in the same way as an older facial classifier. Classification and generation are different tasks. The connection is conceptual: both depend on visual data regimes that make some faces more available to computation. In a portrait, availability may become a visible quality. It may appear as sharper detail, better lighting, more stable features, or a more convincing expression. The audit treats facial legibility as an aesthetic and ethical category.

Third, cultural identity may be compressed into a small set of signs. Qadri and colleagues show that text-to-image systems may carry an outsider's gaze toward non-Western cultures. In a portrait audit, this risk can be examined through prompts such as a contemporary Mongolian artist, South Asian curator, African museum director, or Indigenous digital artist, which may be pulled toward costume, rural landscape, ritual object, or timeless ethnic scenery. These elements

DOI: 10.5281/zenodo.21057816

can be meaningful in some contexts, but they become reductive when a system treats them as the main proof of culture.

The problem is not the presence of traditional signs. A deel, a sari, a mask, a beadwork pattern, a rural landscape, or a ceremonial object can be meaningful within a specific artistic project. The problem is automatic substitution. When a system uses a small set of cultural signs to stand in for living subjects, culture becomes a visual shortcut. The subject is no longer represented as a contemporary person but as a readable cultural specimen.

Fourth, style may hide bias. High fashion lighting, museum catalogue neutrality, historical portrait language, and documentary realism are not innocent visual choices. They may make a generated portrait appear professional while narrowing the range of acceptable bodies, settings, and cultural signs. A polished image may still be a narrow image. For artists, this means that visual quality should not be confused with representational fairness.

This matters because many users judge AI images by finish. Clean skin, symmetrical lighting, coherent clothing, and cinematic composition can make a portrait seem more advanced. The audit asks a different question: what kind of person does the finish make believable? If only certain subjects receive professional polish while others receive exotic scenery or ethnographic detail, style has become part of the algorithmic gaze.

Table 3. Coding categories

Category	Audit focus
Default visibility	Unmarked personhood, beauty, professionalism. Does one identity appear universal?
Facial legibility	Skin tone, facial detail, lighting stability. Are some faces easier to see?
Cultural compression	Costume, rural scene, ritual object. Is culture reduced to signs?
Aesthetic normalization	Fashion lighting, smooth skin, institutional neutrality. Does style hide narrowing?

6. Art Practice

The prompt audit reframes AI portraiture as a visible sequence of decisions rather than a finished image alone. The first output, the corrected output, the rejected image, and the final work can be read as traces of how the system organizes identity. In this sense, the artistic material is not only the portrait but the chain of choices that makes the portrait appear.

This approach also helps separate correction from critique. Correction tries to obtain a better image. Critique asks why the first image appeared in that form. Both are useful, but they are not the same. If an artist only corrects the image, the algorithmic gaze remains hidden. If the artist displays the sequence of defaults and corrections, the image-making system itself becomes visible.

For artists working with Mongolian or other non-Western subjects, this analysis raises a central question. It asks whether a generated portrait can present a subject as contemporary, urban, professional, hybrid, ordinary, and internally diverse. It asks whether the system can represent culture without reducing it to costume or scenery. These questions keep the paper inside art studies. The focus remains on portrait convention, cultural sign, visual sequence, and exhibition strategy.

The audit can also become an artwork. A series might show sixteen prompts, the first outputs, the corrected outputs, and the artist's annotations. Another work might show only the failed defaults, forcing viewers to confront the system's visual assumptions. A third work might invite viewers from a cultural community to annotate which images feel familiar, false, respectful, or reductive. In each case, the artist does not simply use AI to make portraits. The artist uses AI to expose the conditions under which portraits are made.

7. Conclusion

The study has argued that AI digital portraiture can be studied through prompted visibility. The algorithmic gaze is not confined to a finished image. It begins in the relation between source data, prompt language, visual convention, and system default. By building a source corpus and a prompt corpus, the article offers a framework for reading how AI portraiture may organize the visible subject before the image is judged as successful or unsuccessful.

The main conclusion is that AI-generated portraits should not be judged only by technical quality. A portrait can be high resolution, visually attractive, and still narrow in its representation of identity. The more convincing the image becomes, the more important it is to ask what form of personhood has been made easy, what form has been marked as cultural, and what form has disappeared.

The study's boundary is intentional. It treats prompt language as an artistic and visual structure rather than as a substitute for a model-performance benchmark. It does not claim that every image-generation model would produce identical portraits from the sixteen prompts. This focus allows the article to locate algorithmic gaze at the start of portrait production, where identity may be selected, normalized, or marked before the final image appears. The paper contributes an art-oriented account of prompt-driven portrait representation and prepares a conceptual basis for later output-corpus studies using named models, generation dates, and visual coding protocols.

References

1. Barocas, S., Hardt, M., & Narayanan, A. (2023). Fairness and machine learning: Limitations and opportunities. MIT Press. https://fairmlbook.org/
2. Bianchi, F., Kalluri, P., Durmus, E., Ladhak, F., Cheng, M., Nozza, D., Hashimoto, T., Jurafsky, D., Zou, J., & Caliskan, A. (2023). Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency, 1493-1504. https://doi.org/10.1145/3593013.3594095
3. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. Proceedings of Machine Learning Research, 81, 77-91. https://proceedings.mlr.press/v81/buolamwini18a.html
4. Crawford, K., & Paglen, T. (2019). Excavating AI: The politics of images in machine learning training sets. https://excavating.ai/
5. Hall, S. (Ed.). (1997). Representation: Cultural representations and signifying practices. Sage.

DOI: 10.5281/zenodo.21057816

6. Mayson, S. G. (2019). Bias in, bias out. Yale Law Journal, 128, 2218-2300. https://yalelawjournal.org/article/bias-in-bias-out
7. Qadri, R., Shelby, R., Bennett, C. L., & Denton, R. (2023). AI's regimes of representation: A community-centered study of text-to-image models in South Asia. Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency, 506-517. https://doi.org/10.1145/3593013.3594016
8. Rose, G. (2016). Visual methodologies: An introduction to researching with visual materials (4th ed.). Sage.
9. UNESCO. (2021). Recommendation on the ethics of artificial intelligence. UNESCO. https://unesdoc.unesco.org/ark:/48223/pf0000380455

DOI: 10.5281/zenodo.21057816