



JOAS NEBE
NARRATIVE PATTERNS OF FEMININITY INTERPRETED BY AI

In August 2022, Stability AI released its AI image generator, Stable Diffusion. Stable Diffusion operates as a text-to-image model, meaning the software generates images based on a written command, called a prompt.

In the summer of 2023, a year after the initial release of the AI, I became interested in understanding how the application worked. At that time, the errors the AI made during image generation were absurd. It wasn't just about two thumbs, the sixth finger on a hand, or twisted, anatomically impossible limbs—though those were common as well. The AI generated images that loaded meaning beyond the intended image, offering its own interpretation of the term I entered (e.g., bride or sitting woman).

The bride in a white dress with a long veil and bouquet seemed to be imprisoned in a dungeon disguised as an altar room in a church (Bride #8). Or she was so grotesquely altered that she resembled Frankenstein's bride more than the woman beautifully dressed for the most important day of her life (Bride #15). In another image, the AI generated a sitting woman in a wedding dress having a skull instead of a face (Bride #1).



#22, Joas Nebe (Germany), Title: "Sitting lady", 1240 x 1252 pixels, 2025



#5, Joas Nebe (Germany), Title: "Prostitute", 1240 x 1252 pixels, 2025



#8, Joas Nebe (Germany), Title: "Bride",
1240 x 1252 pixels, 2025



#67, Joas Nebe (Germany), Title: "LILITH",
1240 x 1252 pixels, 2025

The AI was doing something that only artists usually do—imbuing the image of a bride with new meaning, showing things that are unspoken but often associated with marriage. The fear of the groom/man of the bride/woman (Bride #1, Bride with skull). The bride's fear of marriage (Bride #8, Bride in the dungeon). The societal expectation of the bride as a radiant super-being (Bride #15, Frankenstein's bride). It just needed someone who recognized the potential that lay within these errors in image creation.

So, I decided to systematically query the AI on certain female stereotypes in pop culture. I prompted the AI to generate images of terms like *Femme Fatale*, *Prostitute*, *Fancy Lady*, *Supermodel*. My goal was to systematically have the AI process a specific theme—the narrative patterns of femininity in our Western culture. Later, I expanded the terms to include biblical and mythological female figures (Lilith and Siren) as well as depiction stereotypes in art history (sitting woman).

The prompts were largely confined to the term itself, for example, "Supermodel" or "Sitting Woman." I had the AI generate at least 20 images for each term and selected the ones where I felt the AI had managed to insert an additional layer of meaning. I acted as the "curator," and the software took on the role of the "artist."

Thus, several image series emerged. One of them was for the term *Supermodel*. The images mostly depicted bust shots of young women, each showing something unnatural—whether it was that they appeared zombie-like (*Supermodel #12*), like a Barbie doll (*Supermodel #13*), or had something of a vampire (*Supermodel #11*).

For the term Prostitute, the AI emphasized the social aspect of exploitation and the physical decay of sex workers (Prostitute #1, 5, 26). The women looked worn out, unkempt, and sick.

Lilith, the ultimate symbol of feminism, the first woman of Adam who turns away from him to live a self-determined life, was depicted as a hero in a video game (Lilith #68), a hysteric in a medical textbook from the Imperial era (Lilith #67), or as a murderer of some sort, blood-stained, holding the weapon in her hand (Lilith #13, 63).

In art history, portraits of sitting women often appear. One of the most famous examples is Dürer's mother. The AI placed all the women on chairs in an otherwise empty room, which gave the depicted women a lifeless quality. The temporal placement of the women was limited to their furniture, clothing, posture, and hairstyle. This resulted in strange contradictions, such as a woman sitting in a strapless dress from the 20th century, with a hairstyle from the late 19th century (Sitting Lady #12).

Another image showed a woman in a dress and hairstyle from the late 19th century. The chair she was sitting on also matched the time period, except that here the anatomy, seating position, and the spatiality of the chair were completely implausible, as the seat of the chair seemed to be facing upwards and parallel to the background, with the woman's legs buried in a crumpled mass of fabric. The only thing preventing her from falling to the floor seemed to be the backrest of the chair, which she was desperately pressing against (Sitting Lady #15).

In a third image, a woman sat in an uncomfortable posture typical of the early 20th century. Her legs were somewhat spread apart. The woman,



#13, Joas Nebe (Germany), Title: "LILITH",
1240 x 1252 pixels, 2025



#63, Joas Nebe (Germany), Title: "LILITH",
1240 x 1252 pixels, 2025



#15, Title: "Bride",
1240 x 1252 pixels, 2025



#26, Title: "Prostitute",
1240 x 1252 pixels, 2025



#13, Title: "Supermodel",
1240 x 1252 pixels, 2025



#68, Title: "LILITH",
1240 x 1252 pixels, 2025



#12, Title: "Sitting lady",
1240 x 1252 pixels, 2025



#12, Title: "Supermodel",
1240 x 1252 pixels, 2025



#15, Title: "Sitting lady",
1240 x 1252 pixels, 2025



#11, Title: "Supermodel",
1240 x 1252 pixels, 2025



Title: "Prostitute",
1240 x 1252 pixels, 2025

viewed frontally, seemed to wear clothing and hairstyle from the imperial era, but from her hips down, she was dressed in the late 20th century. Since her knees and lower legs, even parts of her thighs, were visible, she must have been wearing a kind of miniskirt, which wasn't invented until the 1960s.

In short, for the AI, the "Sitting Lady" was a contradiction, something that doesn't exist and, therefore, cannot be assigned either temporally or spatially.

The dialogue I had with Stable Diffusion at that time was possible because the software was still clumsy and error-prone, and the results were far from convincing. Just two years later, today, such errors no longer occur. The AI now largely adheres to the commands it is given, and when those commands are as general as the ones I used in my experiment, the results are uninteresting and dull. The window of opportunity for a curator-artist relationship, like the one I developed with Stable Diffusion, has closed.