

7 14as24mw

An Incredibly Average Face – Approaching 31,200 Images with Human Eyes, Sorting Systems, and Collaging

Jana Hartmann

I. p. 241, Chunk 6: Indirektes Design
II. p. 272, Chunk 8: Hidden Layers 24

I. p. 241, Chunk 7: Indirektes Design
Chunk 1



Figure 1: a portrait of a incredibly average face, {style}; collage.
© Jana Hartmann




Figure 2: a portrait of a incredibly average face, {style}; collage.
© Jana Hartmann

Figures 1 and 2 are collages of five images generated from the same seed 264523154666154 and almost the same prompt. What sets them apart is the ending of the prompt that describes the style.

Chunk 2

All were prompted: "A portrait of a[n] incredibly average face, ...".

Chunk 3

Chunk 4

These collages were created as part of my practical research on how beauty, or more specifically facial beauty, is defined in AI image generation. This research is based on the broader question of how beauty is socially constructed, defined, and deployed¹ and looks at AI image generation as an example where these definitions become explicit.

It builds on Hito Steyerl's term *mean images* that refers to AI-generated images as averages that encapsulate societal norms.

Chunk 5

These images are averages/means because of the way they generate pixels. An AI-generated image of a cat does not depict a cat that exists in the world, like a photograph would. Rather, it can be described as an average of all the depictions of cats in the AI training data set. Steyerl argues that the AI-generated mean image does not represent a real-world average but rather tends to

1: Colebrook, C. (2006). Introduction. *Feminist Theory*, 7(2), 131-142. <<https://doi.org/10.1177/1464700106064404>>

show “extreme and unattainable outliers—for example, anorexic body ideals”.² AI-generated images are furthermore *mean* because of the exploitative work conditions under which machine learning models are trained, and because they are based in part on discriminatory labeling of people.³

To investigate a statistical medium with statistical methods,⁴ 31,200 images were generated from an automated *ComfyUI* script.

Based on findings from generating small batch sizes, I created a prompting system with four

variables that generates prompts in the form of a portrait of a {beauty descriptor} {object}, {style}.

Values for the variables *beauty*, *object*, and *style* were stored in text files.

40 seeds were chosen randomly. The script then, for every seed, prompted images from every possible combination of variable values (Figure 3). This system was set up to make the resulting images comparable along multiple parameters – because a single person cannot simply analyze 31,200 images.

Chunk 6

Chunk 8

Chunk 9

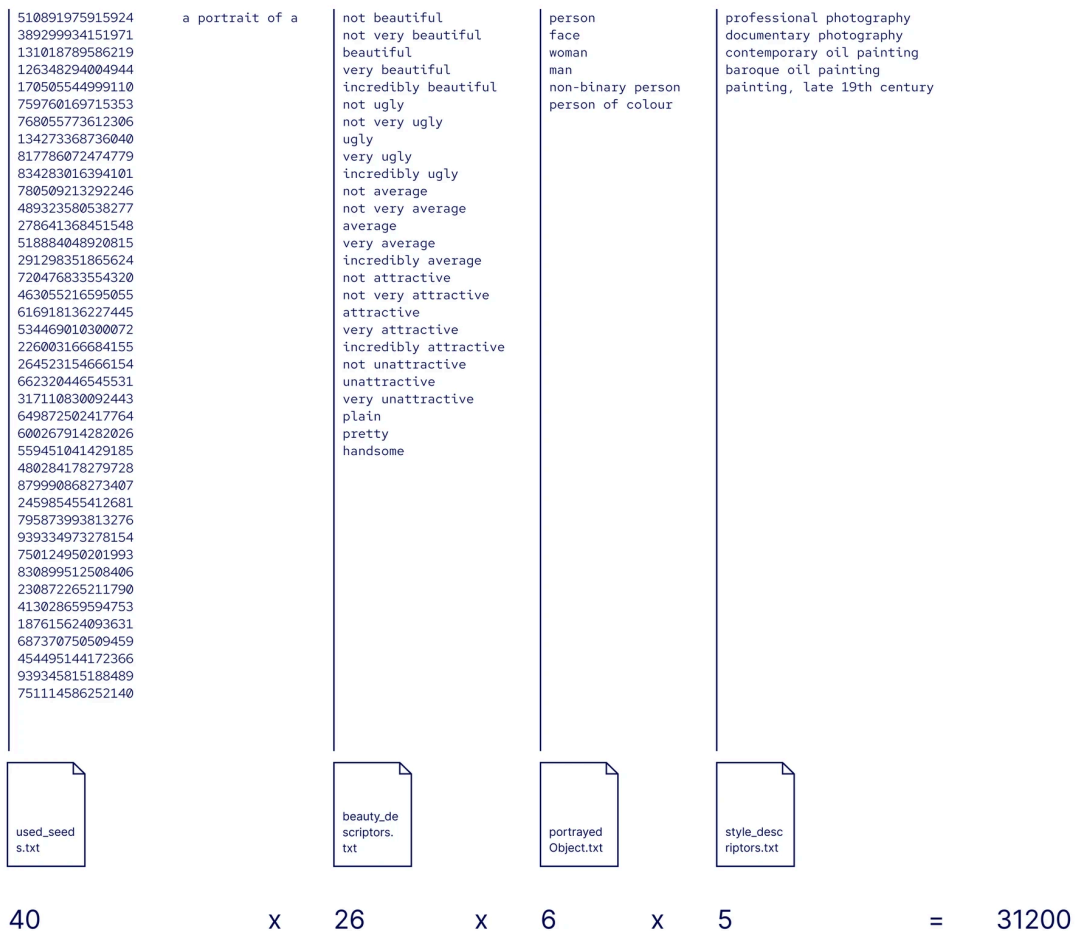


Figure 3: Diagram of prompting system. © Jana Hartmann

2: Steyerl, H. (2023). Mean Images. *New Left Review*, 140/141, 89. <<https://newleftreview.org/issues/ii140/articles/hito-steyerl-mean-images>>

4: Offert, F., & Dhaliwal, R. S. (2024). *The Method of Critical AI Studies, A Propaedeutic*. arXiv. <<https://doi.org/10.48550/ARXIV.2411.18833>>

3: Steyerl, H. (2023). Mean Images. *New Left Review*, 140/141, 82-96. <<https://newleftreview.org/issues/ii140/articles/hito-steyerl-mean-images>>

I. p. 201, Chunk 3: Creative Ownership and Creative...
II. p. 115, Chunk 2: Artefact Hunt
II. p. 69, Chunk 23: Editorial

Chunk 10

Dealing with large amounts of images is usually a problem discussed in the context of AI training data.

Chunk 11

The scale of these data sets, e.g. *LAION-5B* that is a part of the training data for *Stable Diffusion*,

makes human curation nearly impossible.⁵ Here however, I am studying image-generating AI as a system of possible images in Frieder Nake's tradition of Generative Aesthetics⁶ that can output endless amounts, which then again leads to the question of: How is it possible for humans to curate or research extremely large amounts of images?

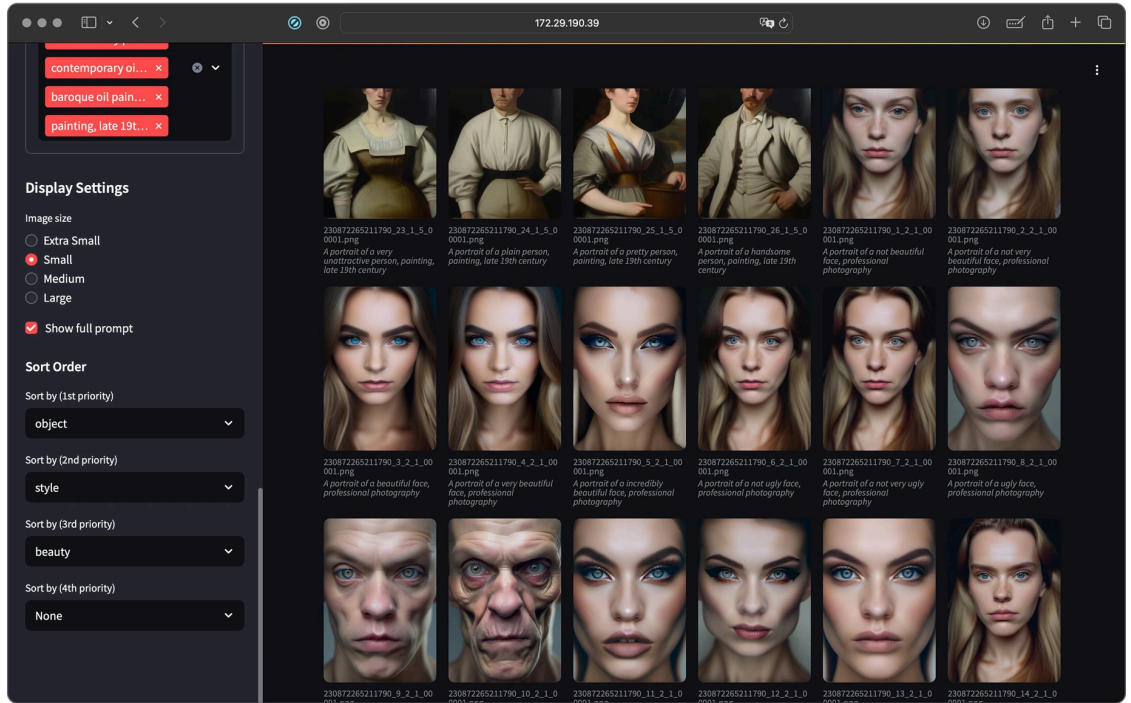


Figure 4: Image Browser - no filter, sorted by seed, object, style, then beauty. © Jana Hartmann

Chunk 12

5: Buschek, C. & Thorp, J. (2024). Models All The Way Down. In *Knowing Machines*. <<https://knowingmachines.org/models-all-the-way#section5>>

6: Offert, F. (2019). The Past, Present, and Future of AI Art. *The Gradient*. <<https://thegradients.pub/the-past-present-and-future-of-ai-art/>>

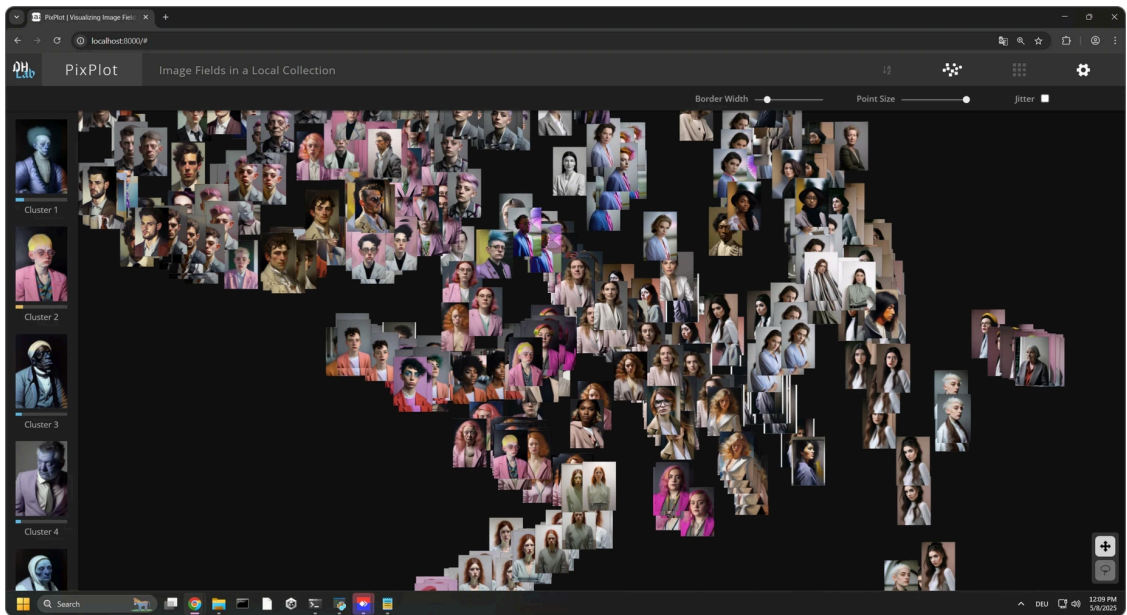


Figure 5: UMAP in PixPlot, zoomed in. © Jana Hartmann

Chunk 13

Figures 4 and 5 illustrate my first steps to looking at the images.

Chunk 14

First, I created a custom file browser for filtering and sorting the images according to the four variables (Figure 4). Then I generated a UMAP visualization with *PixPlot*, a tool by *The Digital Humanities Lab* at *Yale University* that clusters images according to pixel similarity (Figure 5).

In the file browser, I quickly began to filter and sort into matrixes where, for example, rows

would represent a style and each column had the same beauty descriptor.

Chunk 15

From there, I decided to create a 3D matrix. Because my images differ from each other in four variables, a 3D matrix had to have a filter applied to one of them. While any variable could have been filtered, as of now I have only filtered by seed. Figures 6 and 7 show matrixes of all 780 images created from seed 480284178279728 (Figure 6) and seed 879990868273407 (Figure 7). Each variable of the prompt system then represents an axis in the 3D space.

Chunk 16



Figure 6: a portrait of a [beauty descriptor] [object], [style] - Seed: 480284178279728; displayed in 3D matrix. © Jana Hartmann

Chunk 17




Figure 7: a portrait of a [beauty descriptor] [object], [style] - Seed: 879990866273407; displayed in 3D matrix. © Jana Hartmann

Furthermore, I explored ways of comparing not only side by side but rather on top of or within one another, as well as comparing through time.

Chunk 18

Therefore, I created videos that show one image at a time in short succession.

Chunk 19

In TouchDesigner, I experimented with layering images using different blend modes and created systems of algorithmic collaging.

While the matrix system applies a filter on one category, for the collages I filtered all but one.

Chunk 20

This means the collages, as they are now, allow comparisons across one variable.

Comparisons of Descriptions of Beauty

III. p. 168, Chunk 2: Anthrobot

Chunk 21

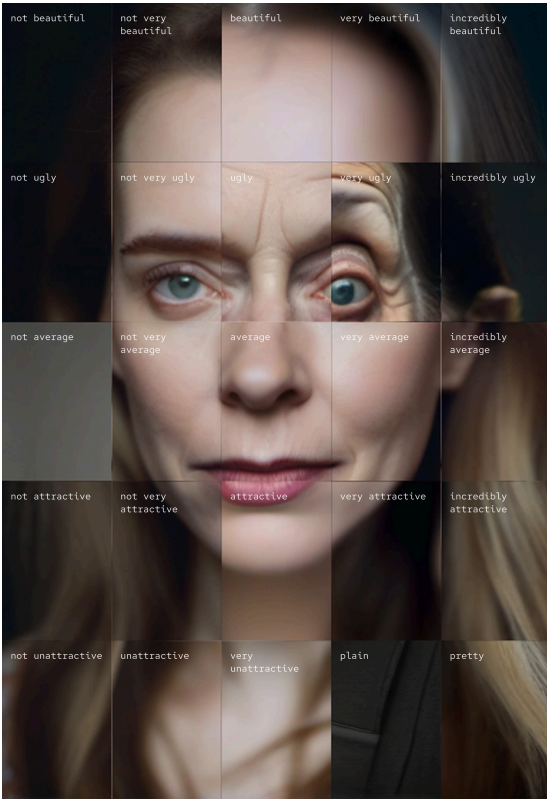


Figure 8: a portrait of a [...] woman, professional photography, seed: 230872265211790; collage. © Jana Hartmann




Figure 9: a portrait of a (beauty) woman, professional photography, seed: 230872265211790; collage. © Jana Hartmann




Figure 10: a portrait of a [...] person, painting, late 19th century, seed: 780509213292246; collage. © Jana Hartmann




Figure 12: a portrait of a beauty person, painting, late 19th century, seed: 780509213292246; collage. © Jana Hartmann



Figure 11: a portrait of a [beauty] person, painting, late 19th century, seed: 780509213292246; collage. © Jana Hartmann




Figure 13: a portrait of a [beauty] person, painting, late 19th century, seed: 780509213292246; collage. © Jana Hartmann




Figure 14: a portrait of a [...] non-binary person, painting, late 19th century, seed: 939334973278154; collage. © Jana Hartmann




Figure 16: a portrait of a [beauty] non-binary person, painting, late 19th century, seed: 939334973278154; collage. © Jana Hartmann



Figure 15: a portrait of a [...] non-binary person, painting, late 19th century, seed: 939334973278154; collage. © Jana Hartmann




Figure 17: a portrait of a [beauty] non-binary person, painting, late 19th century, seed: 939334973278154; collage. © Jana Hartmann

Comparisons of Seeds

Chunk 22




Figure 18: a portrait of a plain face, professional photography, randomly selected seeds; collage. © Jana Hartmann




Figure 19: a portrait of a beautiful woman, baroque oil painting, randomly selected seeds; collage. © Jana Hartmann




Figure 20: a portrait of a average face, baroque oil painting, randomly selected seeds; collage. © Jana Hartmann



Figure 21: a portrait of a handsome face, painting, late 19th century, randomly selected seeds; collage. © Jana Hartmann

Comparisons of Styles



Figure 22: a portrait of a unattractive woman, [style], seed: 662320446545531; collage. © Jana Hartmann



Figure 24: a portrait of a unattractive woman, [style], seed: 662320446545531; collage. © Jana Hartmann



Figure 23: a portrait of a unattractive woman, [style], seed: 662320446545531; collage. © Jana Hartmann




Figure 25: a portrait of a attractive face, [style], seed: 264523154666154; collage. © Jana Hartmann




Figure 27: a portrait of a attractive face, [style], seed: 264523154666154; collage. © Jana Hartmann




Figure 26: a portrait of a attractive face, [style], seed: 264523154666154; collage. © Jana Hartmann

The collages across seeds (Figure 18–21) mostly show variance in compositions. When I compare them to each other, they seem to give an insight into colour schemes of a style. This is reminiscent of the UMAP visualization in *PixPlot*, where images with the same style prompt tended to cluster together because colours impact pixel similarity.

In contrast, images created from the same seed are much more similar in composition.

Chunk 23

In collages across styles (Figure 22–27), the switch between collage parts is often more like a change in texture or in tools of creation. The switch between beauty descriptors (Figure 8–17) can also be described along those lines.

Chunk 24

In all styles the difference between `beautiful` and `ugly` seems to rely strongly on youth of skin. In the style `professional photography`, it relies furthermore on styles of digital photography editing, as `beautiful` results in soft and blurry images while `ugly` images show more contrast.

Chunk 25

Amanda Wasielewski observes that generative AI strips historical art styles of their context and uses them more like a (color) filter than content. This way, they become endlessly interchangeable.⁷

Chunk 26

The term filter seems adequate to describe the effect of the style variable as well as of beauty descriptors. Some of the images prompted a portrait of a `{beauty}` woman, `professional photography` from seed `230872265211790` (Figure 8–9) could be photographs of the same person but with different filters applied.

Overall, the systems I created for looking at 31,200 images genuinely helped me understand these images better.

Chunk 27

Especially the collages lead to new insights into how different parts of the prompt affect the resulting images. In regard the question of how a single person can look at or curate large amounts of images I should mention that I cannot say that I have seen all images.

In under 36 hours, I generated 31,200 images. Over the course of one month I have spent a lot of time with them, and yet I still discover new images every time I open the folder. From a fraction of these 31,200 images, I have generated 786 collages as of the time of writing. I have exported 49 images of the 3D matrix. I created 234 video files of moving collages and recorded 4 videos of the 3D matrix. I have taken (only) 8 screenshots of the UMAP visualization in *PixPlot* and 32 screenshots of the custom file browser. In total I now have 32,313 new files lying on my computer and in multiple clouds.

II, p. 122, Chunk 5: You Press the Button, They...
II, p. 132, Chunk 2: Flesh Fest

Chunk 28

Selecting the 20 collages presented in this publication took nearly as long as generating all 786.

Roland Meyer describes AI image generation as an inherently wasteful process, because of its vast energy consumption and because a multitude of generated images are discarded immediately. AI-generated images are thus “neither rare nor valuable.”⁸ Generating (large amounts of) images inevitably leads to a question of selection or curation and most likely to discarded options that become digital waste.

The images I have not even looked at yet are arguably nothing but digital waste.

Chunk 29

Yet the 31,200 images I generated are still only a tiny fraction of the possible images I *could* have generated.

This article is separately published as [10.25358/openscience-13023](https://doi.org/10.25358/openscience-13023).

7: Cetinić, E. (2024). Introduction. *Hertziana Studies in Art History*. <<https://doi.org/10.48431/hsah.0301>>

8: Meyer, R. (2025). “Platform Realism.” *AI Image Synthesis and the Rise of Generic Visual Content*. *Transbordeur*, 9, 14. <<https://doi.org/10.4000/13dwq>>