

The statistical physics of visual preference: What makes an image appealing?

Laboratory name: Econophysics Lab

Internship location: Econophysics Lab, CFM, 23 rue de l'Université, Paris

Supervision: Ortensia Forni, Alexandre Darmon and Michael Benzaquen

Email: michael.benzaquen@cfm.com

Project at a Glance

Why do some images feel visually compelling while others appear dull, chaotic or artificial? This internship will investigate whether aesthetic preference can be related to measurable statistical properties of images. Building on our previous work on structural complexity, multiscale relevance, and quantitative color harmony, we will combine large human-preference datasets with interpretable descriptors of image organisation to identify which visual structures are consistently associated with appeal.




Figure 1: Illustration of a quantitative approach to visual preference, linking image content to measurable structural, chromatic and spatial features, and ultimately to human aesthetic judgments. Image generated by ChatGPT.

Project Description

The relation between visual structure and aesthetic preference is a long-standing question in perception, art and information theory. In our previous work, we showed that preference for controlled synthetic images tends to peak at intermediate structural complexity [1], and later introduced multiscale relevance as an information-theoretic measure of image organisation across scales [2]. More recently, we studied color-pairing preferences in controlled experiments and found reproducible patterns of harmony that partly mirror hue statistics observed in natural landscapes [5]. Together, these results suggest that visual preference may reflect not a single notion of “beauty”, but a combination of structural, multiscale and chromatic regularities shaped by the statistics of natural visual environments.

The aim of this project is to test how far these ideas extend to real images. Large public datasets now make this possible. AVA, for example, contains more than 250,000 photographs with human aesthetic ratings [3], while datasets such as Pick-a-Pic contain hundreds of thousands of pairwise human preferences between AI-generated images [4]. These data provide an opportunity to move beyond small controlled experiments and ask whether statistical regularities survive across very different visual domains. The student will characterise images using interpretable observables such as multiscale relevance, entropy across scales, spatial correlations, power spectra, edge statistics and color-organisation measures, and relate them to human ratings or pairwise preferences while controlling, as much as possible, for semantic content and obvious image-quality effects.

A particularly interesting direction will be to construct controlled image families in which content is approximately preserved while structural or chromatic properties are varied. Modern generative models can be useful here not as objects of comparison with human artists, but as experimental tools for exploring visual space. This could allow us to test whether changing complexity, color organisation or spatial correlations systematically alters preference for otherwise similar images.

The objective here is not to build a black-box aesthetic score, nor to ask whether “AI art is better than human art”. Instead, the goal is to identify robust and interpretable statistical properties associated with visual preference, and to determine which aspects of aesthetic judgement can/cannot be explained through image structure and color statistics alone.

References

- [1] S. Lakhal, A. Darmon, J.-P. Bouchaud, and M. Benzaquen, “Beauty and structural complexity,” *Physical Review Research* **2**, 022058(R) (2020).
- [2] S. Lakhal, A. Darmon, I. Mastromatteo, M. Marsili, and M. Benzaquen, “Multiscale relevance of natural images,” *Scientific Reports* **13**, 14879 (2023).
- [3] N. Murray, L. Marchesotti, and F. Perronnin, “AVA: A large-scale database for aesthetic visual analysis,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2408–2415 (2012).
- [4] Y. Kirstain, A. Polyak, U. Singer, S. Matiana, J. Penna, and O. Levy, “Pick-a-Pic: An open dataset of user preferences for text-to-image generation,” *Advances in Neural Information Processing Systems* **36** (2023).
- [5] O. Forni, A. Darmon, and M. Benzaquen, “Harmonious color pairings: Insights from human preference and natural hue statistics,” *iScience* **29**, 116038 (2026).