

Supplementary material

Hue Mastery is founded on Birkhoff’s order–complexity theory and evaluates image colour aesthetics along two analytic axes: color harmony and color complexity. Harmony is computed with $\text{Soft-}F_\beta$, which matches an image’s dominant hues to Matsuda’s eight harmony templates and is refined by three spatial descriptors—HEM-1 (saturation-weighted hue-centroid deviation), HEM-2 (inter-hue angular spread) and HEM-3 (regional dominant-hue consistency). Colour complexity is quantified by SCMR, which constructs multiradius perceptual spheres around each dominant colour in CIELab space to capture hierarchical diffusion. All indicators are closed-form and execute in milliseconds on a standard CPU.

Unlike deep-learning methods that require large-scale annotations, Hue Mastery is training-free, intrinsically interpretable and style-robust, delivering low-latency inference. In addition to a global aesthetic score, the system outputs the best-matching harmony template, the colour-centroid, the set of dominant hues and region-level visualisations, providing a lightweight yet explainable baseline for generative-AI, style-transfer and design applications.

The demo visualises every result: a composite radar chart displays the colour score, HEM-1/2/3 and SCMR; a colour-wheel plot shows the optimal template; an inset marks the global colour centroid (HEM-1); a polar plot depicts hue symmetry (HEM-2); an $n \times n$ dominant-hue grid assesses local consistency (HEM-3); and a thumbnail presents the multiscale colour-richness value and levels. The interface is fully responsive, ensuring identical interaction on desktop and mobile devices.