Beauty and structural complexity

Samy Lakhal^{1,2,3}, Alexandre Darmon,¹ Jean-Philippe Bouchaud,^{3,4} and Michael Benzaquen^{1,2,3,4,*}¹Art in Research, 33 rue Censier, 75005 Paris, France²Ladhyx UMR CNRS 7646, Ecole polytechnique, 91128 Palaiseau Cedex, France³Chair of Econophysics & Complex Systems, Ecole polytechnique, 91128 Palaiseau Cedex, France⁴Capital Fund Management, 23 rue de l'Université, 75007 Paris, France

(Received 12 October 2019; revised manuscript received 17 February 2020; accepted 27 May 2020; published 15 June 2020)

We revisit the long-standing question of the relation between image appreciation and its statistical properties. We generate two different sets of random images well distributed along three measures of entropic complexity. We run a large-scale survey in which people are asked to sort the images by preference, which reveals maximum appreciation at intermediate entropic complexity. We show that the algorithmic complexity of the coarse-grained images, expected to capture structural complexity while abstracting from high frequency noise, is a good predictor of preferences. Our analysis suggests that there might exist some universal quantitative criteria for aesthetic judgment.

DOI: [10.1103/PhysRevResearch.2.022058](https://doi.org/10.1103/PhysRevResearch.2.022058)

What makes a beautiful image? Is there such a thing as universal beauty? These puzzling yet fascinating questions have been tackled many times in the past within several disciplines, including philosophy, psychology, arts, or mathematics [1–13]. According to Kant: *The beautiful is that which pleases universally without a concept* [14]. The idea of an intelligible beauty appeared in ancient Greece, where Nature was believed to be a *cosmos* constituting a principle of order and harmony. The proportions between the constitutive elements of each being are rightfully defined, whether it is a work of art, a living organism, or a city [15]. Following the Greeks, the Baroque and Renaissance artists also believed in a universal beauty, and it is striking that their arts partially rely on a mathematization of the artistic representation (symmetry [10], proper geometric proportions as given by the *golden number* [1], etc.). In other terms, the belief that there must be scientific grounds to the conception of what is artistic or beautiful has been out there for quite some time. Yet the very idea of universal beauty is a longstanding debate which has known many ruptures through the history of art [2] and still opposes a number a great modern thinkers.

Physicists' interest in the subject is more recent. Stephens *et al.* [16] showed that natural images were critical in the thermodynamic sense and proposed a theory for the *Thermodynamics of Natural Images*. While, as pointed out above, many would consider *quantitative aesthetics* to be an oxymoron, and indeed it would be rather nonsensical to aim at building a fully consistent theory of pictorial art, we, as physicists,

believe there is some room for a quantitative analysis. For example one could easily argue that an aesthetically appealing image often results from a subtle balance between regularities and surprises. Indeed, it seems rather plausible to think that while one might find dull an image that is too regular (no surprises), one may also feel lost in front of an image with no recognisable shapes or structures to hang on to (too much surprise). As argued in Ref. [17], *Total chaos is disquieting. Too much regularity is boring. Aesthetics is perhaps the territory in-between*. Provided one agrees with such statements, these ideas clearly suggest that one could design an entropy-like function to quantify this subtle and complex equilibrium.

To address this question we run a large-scale survey in which people are asked to sort by preference two different sets of random images well distributed along three measures of entropic complexity: Fourier magnitude's slope, fractal dimension, and compression rate. The paper is organized as follows. We first present and motivate our image generation methods. We then present and analyze the results of the survey. Finally, we argue that algorithmic complexity of the coarse-grained images is a rather good proxy for image appreciation, and then conclude.

There exist many possible measures of image complexity, relying on, e.g., their mathematical properties [18,19], their physical properties [16,20,21], or even their cognitive impact [22–24]. Here we choose to work with three simple measures that can be easily computed unequivocally for any digital two-dimensional (2D) image. The first one is the *magnitude slope* α defined as the logarithmic slope of the radially averaged Fourier magnitude $\mathcal{M}(k) = \langle \hat{u}(k, \theta) \rangle_\theta$, where $\hat{u}(k, \theta)$ denotes the Fourier transform of the image grayscale intensity $u(r, \phi)$, and $\mathcal{M}(k) \sim |k|^\alpha$. The second is the *fractal dimension* d_f computed using the Minkowski-Bouligand box-counting method [25]. After transforming the image to black and white (B&W) using an intensity threshold ensuring two equally populated levels [16], the fractal dimension follows $N(\epsilon) \sim \epsilon^{-d_f}$ where

*michael.benzaquen@polytechnique.edu

FIG. 1. (a) Fourier magnitude-generated images, and (b) box-counting-generated images, both series with increasing complexity from left to right. These images were used for our large-scale survey.

$N(\epsilon)$ is the number of boxes of size ϵ containing both black and white features. The third is the *compression rate* or *algorithmic complexity* τ computed as the ratio between the size of the PNG image and its uncompressed maximum size equal to 256^2 (1 byte per pixel for 256×256 pix images).

To remove possible cognitive and cultural biases, we choose to design our experiment with abstract images randomly generated using two of the complexity measures presented above. The first set of images [Fig. 1(a)] is generated by *reverse-engineering* the Fourier magnitude property: setting $\hat{u}(k, \theta) = k^\alpha e^{i2\pi\eta(k)}$, where η is drawn from a uniform distribution on $[0,1]$ with $\langle \eta(k)\eta(k') \rangle \propto \delta_{kk'}$, and taking the inverse Fourier transform of $\hat{u}$ allows to produce a series of random grayscale images $|u(r, \phi)|$ with controlled magnitude slope [23,26,27]. Table I gathers the computed complexity measures α , d_f , and τ of the 256×256 images displayed in Fig. 1(a). As one can see both d_f and τ are increasing functions of α , comforting our choice of complexity measures and indicating that there is a clear correlation between the spectral, fractal, and algorithmic properties.

The use of a second set of images was motivated by the remarks of some survey participants. When asked why they had preferred certain images, they responded their picks reminded them of cloudy skies or galactic landscapes.

With the aim of producing more abstract images, we used an alternative method, now based on reverse-engineering the Minkowski-Bouligand box-counting method [Fig. 1(b)]. Black squares of size $\epsilon = 2^k$, $k \in [1, k_{\max}]$, drawn from the distribution $N_k(\epsilon) = A\epsilon^{-x}$ are randomly added to a white canvas of size 256×256 with the condition that they do not overlap. The upper boundary $k_{\max}$ is chosen such that the biggest squares occupy at most 1/16th of the total surface, $k_{\max} \leq \log_2(256) - 2 = 6$. We also enforce that the total fraction of black pixels does not exceed 1/2 [28]. Here again the complexity measures appear to be increasing functions of one another (see Table I).

In 2013, Spehar and Taylor [27] conducted a survey on 26 academics, using black and white computer generated images with increasing fractal dimension. They found a reversed U-shaped relation between image appreciation and fractal dimension, with an aesthetic optimum for $d_f \approx 1.5$, allowing to argue that we indeed tend to prefer images with intermediate complexity, see also Refs. [23,26]. Curious of their results, we conducted a larger-scale experiment intended for a larger panel (about 1000 participants with different backgrounds), using the images presented in Fig. 1. Our question at this stage was similar: Is there a link between the statistical properties of our generated images and the tendency of people to appreciate them?

Survey methods design constitutes a strand of research on its own [29]. For optimal results the selection task must be simple and display the minimum amount of information to the interviewee. While the common five-star ratings only take a time proportional to the number of images to score, these were shown to be weighted by extreme grades, the utility given to intermediate grades being far from linear [30]. Five-star ratings image-by-image can also be rather disorienting due to the lack of reference. Another option is image classification, where the interviewee is presented with the whole set of images and is then asked to sort them by preference. While also time-efficient, presenting all the images at once might strongly induce people into intuitively recognizing other features, such as complexity, and ending by

TABLE I. Complexity properties of the images presented in Fig. 1.

	a1	a2	a3	a4	a5	a6
α	-3.96	-2.47	-1.95	-1.42	-0.76	0
d_f	1.22	1.33	1.58	1.87	1.99	2.0
τ	0.186	0.325	0.472	0.711	0.902	0.956
	b1	b2	b3	b4	b5	b6
α	-2.14	-1.78	-1.56	-1.32	-1.04	-0.77
d_f	1.42	1.51	1.62	1.75	1.88	1.95
τ	0.012	0.014	0.022	0.044	0.095	0.16

sorting them by something else than preference. Finally, the *Battle survey* consists in presenting the interviewee with two images, asking her/him to choose the one she/he prefers, and repeating the operation until all possible combinations of two images have been addressed [4,31]. While less time-efficient [with $N(N-1)/2 = O(N^2)$ battles, for $N = 6$ images one needs 15 rounds to complete the survey], this method beats the other shortcomings mentioned above, and people usually feel more comfortable with such a binary task, as it is intellectually less challenging. We thus choose the latter method.

We conducted two slightly different surveys. The panel for the first survey consisted of colleagues from Capital Fund Management (CFM) and Ecole Polytechnique as well as students and relatives, adding up to ≈ 350 people, who were asked to participate without any financial incentive. While probably slightly biased population-wise, these are the results we are most confident with since we believe people in such a panel completed their tasks selflessly and honestly. To run this survey we used the ZOONIVERSE platform [32] which provides a rather intuitive interface. The 15 two-image sets for each series were generated using a PYTHON algorithm that concatenated the images in a random order and attributed them a different name so that the interviewee could not find hidden information. Upon completion of the survey, to establish a global ranking of the images we attributed them a score according to the following rule: if image i wins (or loses) a battle, its score S_i increases (or decreases) by $1/N_i$ where N_i denotes the number of battles in which i was involved [33]. To obtain a score $S_i \in [0, 1]$ we then transform it as $S_i \rightarrow (S_i + 1)/2$. The results are plotted as a solid black line in Fig. 2. Remarkably, the preferred images appear to be a4 a5, and b4 b5, respectively, both corresponding to α close to 1. To note, interestingly $\alpha \approx 1$ is often associated to the spectral properties of natural images [16,34] and visual arts [18]. Discussions with voters revealed that they found their favorite images to be the most *harmonious* and *well balanced*.

To increase the size and diversity of the panel, we ran another experiment on the MECHANICAL TURK platform [37], in which participants were paid a small amount of money to participate (we reached ≈ 600 panelists). However, there was a significant risk that people would answer randomly as they do not, by definition, participate selflessly and in good will as in our first experiment. We thus formulated the question differently to encourage nonrandom participation: much like in Keynes' famous beauty contest [38], people were asked to pick the image which they thought would be preferred by the majority and told they would not get paid if their overall choices fell too far off the average (needless to say, we did process all of the data, regardless of its distance to the average). Results are presented in Fig. 2 (dash-dotted gray line). While probably introducing other biases, the results displayed good agreement with that of the initial selfless survey. Also note that while all results are fully consistent for the first set of images [Fig. 1(a)], the second experiment leads to a less pronounced maximum for the image series of Fig. 1(b), with scores on average closer to $S = 1/2$ than they were for the very first experiment.

Very much like entropy is used to measure the disorder in a physical system, we would now like to see whether there

FIG. 2. Results of the two different surveys (Zooniverse: solid black line; Mechanical Turk: dash-dotted gray line). The red diamonds markers indicate the structural complexity τ_{cg} defined below. We rescaled and shifted vertically τ_{cg} to show that the maximum scores also correspond to maximum structural complexity. The blue circles reflect the complexity measure proposed in Ref. [35]. Top: Image series of Fig. 1(a). Bottom: Image series of Fig. 1(b). Error bars reflect the 95% confidence interval using the bootstrap method [36].

might exist a statistical proxy to estimate an image's *harmony* and *equilibrium*, as described by our survey participants. Given the complexity measures described above, images with low complexity display very simple shapes (a1 b1), and images with very high complexity display a large amount of white noise (a6 b6). Our survey revealed that maximum appreciation is obtained for intermediate complexity suggesting the following question: Could it be that an aesthetically appealing image results from a subtle balance between complexity and regularity? And if so, can we find an associated statistical measure? The work of Desolneux *et al.* [19] clearly resonates with such questions. Guided by the idea that there is no perceptual structure in white noise, the authors attempted to characterize forms and structures and in particular defined unusual features or *Gestalts* as *sets of points whose (...) spatial arrangement could not occur in noise*. Their ideas can be easily illustrated with the coffee and cream dynamics [39]. Consider the experiment in which plain cream is left to slowly mix with plain coffee. While the initial and final states of such a system

FIG. 3. Illustration of the coarse-graining procedure on images a1, a4, and a6, with $r_{cg} = 7$ and $\delta = 0.23$.

display very regular homogeneous structures, the transitional regime displays interesting and *complex* mixing patterns as the cream/coffee interface slowly disappears. So far we used the term *complexity* rather imprecisely and it is now time to distinguish more rigorously two sorts of complexity. The first is *entropic complexity* measuring the amount of information in the image, which in the coffee experiment can only be an increasing function of time according to the second law of thermodynamics; the second is *structural complexity* accounting for the amount of features outside of the noise, which here is a nonmonotonous function of time displaying a maximum at intermediate stages where the nontrivial mixing patterns are most significant. Entropic complexity is well described by α , d_f , or more commonly τ . Structural complexity, in a sense, measures *noiseless* entropic complexity or *interestingness*.

Guided by the work of Aaronson *et al.* [39] we computed structural complexity as a *noiseless* entropy. More precisely we apply a coarse-graining procedure of given radius r_{cg} on the B&W images and then compute their algorithmic complexity τ_{cg} which we call structural complexity in the following [40]. The color of a given block is determined by its black to white pixel ratio $\eta \in [0, 1]$: white if $\eta \leq \delta$, gray if $\delta < \eta \leq 1 - \delta$, and black for $\eta > 1 - \delta$ where $\delta \in [0, 1/2]$ is a given threshold. Figure 3 illustrates the procedure on images a1, a4, and a6; after turning them into B&W (second column), the coarse-graining procedure is applied (third column). As one can see, image a1 is barely changed (just a thin gray line at the domain boundaries) and we thus expect $\tau_{cg} \approx \tau$, image a4 is slightly denoised while letting its structures invariant $\tau_{cg} \lesssim \tau$, image a6, however, is strongly denoised as the coarse-graining procedure has left it almost plain gray suggesting

$\tau_{cg} \ll \tau$. The structural complexity computed for both sets of image is plotted on Fig. 2 as dark red diamonds. As expected, $\tau_{cg}(\alpha)$, or equivalently $\tau_{cg}(d_f)$ and $\tau_{cg}(\tau)$, are nonmonotonous functions displaying a maximum for intermediate values of α , d_f , and τ . Note that the y-axis in Fig. 2 is arbitrary (structural complexities and scores do not have the same dimensions) and thus only the shape of the curves should be compared, which we facilitate by superimposing them (vertical linear transformation $\tau_{cg} \rightarrow a\tau_{cg} + b$). Furthermore, the scale parameter r_{cg} and the threshold δ can be used as fitting parameters; in particular, r_{cg} acts as the cutoff of a low-pass filter which erases high-frequency spatial features, increasing it tends to lowers the right most red markers and shift the maximum to the left. Up to a multiplicative factor, best fits are obtained for $(r_{cg}, \delta) = (7, 0.23)$ for the first set and $(13, 0.12)$ for the second. The agreement between theory and experiments is quite convincing. Not only do the maxima coincide, but also the overall shape of the curves is similar. Note that such good agreement is not r_{cg} - and δ -sensitive provided the latter are kept within a reasonable range. This quantitatively supports the idea that structural complexity is a good proxy for average image preference. Also note that the alternative interesting measure of structural complexity proposed in Ref. [35] (see blue circles on Fig. 2), though more *ad hoc* than ours, appears to also correlate well with our survey data. See also [41] for yet another consistent definition of structural complexity based on a multi-scale entropy approach.

Let us summarize what we achieved. Using two random-image-generation algorithms, we produced two different sets of abstract images spanning a broad range of entropic complexity, measured by three different quantities. We then designed and ran a large-scale experiment for image classification and found that preference peaks about complexity criteria matching that of natural images, perhaps indicating that people's preferences are influenced by their natural environment. Finally, our main contribution is to show that a “noiseless” entropy (that captures interesting structural features only) accounts well for the experimental results on image appreciation. It is interesting to speculate that, when confronted with images, the human brain may actually conduct the same kind of geometrical coarse-graining, trying to extract forms and structures while erasing uninteresting noise, or as put by the *Gestalt theory* [42]: filter meaningful perceptions from chaotic stimuli. As a result, the excess of noise and lack of forms may lead to unconscious rejection of structureless images.

We thank Christian Schmidt, Debanuj Chatterjee, and Raphael Benichou for fruitful discussions, as well as Bastien Legay, who contributed to the early stages of this project. This publication uses data generated via the Zooniverse.org platform, development of which is funded by generous support, including a Global Impact Award from Google, and by a grant from the Alfred P. Sloan Foundation. This research was conducted within the Econophysics & Complex Systems Research Chair, under the aegis of the Fondation du Risque, the Fondation de l'Ecole polytechnique, the Ecole polytechnique and Capital Fund Management [43].

- [1] G. T. Fechner, *Vorschule der Aesthetik* (Breitkopf & Härtel, Leipzig, 1876), Vol. 1.
- [2] U. Eco and G. De Michele, *Storia Della Bellezza* (Bompiani, Milano, 2004).
- [3] E. H. Gombrich, *Art and Illusion* (Princeton University Press, Princeton, NJ, 1969).
- [4] E. A. Vessel and N. Rubin, *J. Vision* **10**, 18 (2010).
- [5] H. J. Eysenck, *Am. J. Psychol.* **54**, 385 (1941).
- [6] M. Bar and M. Neta, *Psychol. Sci.* **17**, 645 (2006).
- [7] B. Spehar, C. W. Clifford, B. R. Newell, and R. P. Taylor, *Comput. Graph.* **27**, 813 (2003).
- [8] J. Alvarez-Ramirez, C. Ibarra-Valdez, and E. Rodriguez, *Chaos, Solitons Fractals* **83**, 97 (2016).
- [9] C. M. Hagerhall, T. Purcell, and R. Taylor, *J. Environ. Psychol.* **24**, 247 (2004).
- [10] I. C. McManus, *Eur. Rev.* **13**, 157 (2005).
- [11] D. J. Graham and D. J. Field, *Spatial Vision* **21**, 149 (2008).
- [12] A. Forsythe, M. Nadal, N. Sheehy, C. J. Cela-Conde, and M. Sawey, *Br. J. Psychol.* **102**, 49 (2011).
- [13] M. S. Kanwal, J. A. Grochow, and N. Ay, *Entropy* **19**, 310 (2017).
- [14] I. Kant, *Critique of the Power of Judgment* (Cambridge University Press, Cambridge, England, 2000).
- [15] E. Aujaleu, *Tréma* **11**, 33 (1997).
- [16] G. J. Stephens, T. Mora, G. Tkacik, and W. Bialek, *Phys. Rev. Lett.* **110**, 018701 (2013).
- [17] J.-P. Bouchaud, *Leonardo* **41**, 239 (2008).
- [18] M. Koch, J. Denzler and C. Redies, *PLoS One* **5**, e12268 (2010).
- [19] A. Desolneux, L. Moisan, and J.-M. Morel, *From Gestalt Theory to Image Analysis* (Springer, New York, 2008).
- [20] B. Gidas, *IEEE Trans. Pattern Anal. Mach. Intell.* **11**, 164 (1989).
- [21] S. Foreman, J. Giedt, Y. Meurice, and J. Unmuth-Yockey, *Phys. Rev. E* **98**, 052129 (2018).
- [22] D. J. Field, *J. Opt. Soc. Am. A* **4**, 2379 (1987).
- [23] B. Spehar, S. Wong, S. van de Klundert, J. Lui, C. W. G. Clifford, and R. Taylor, *Front. Hum. Neurosci.* **9**, 514 (2015).
- [24] R. Taylor, B. Spehar, C. Hagerhall, and P. Van Donkelaar, *Front. Hum. Neurosci.* **5**, 60 (2011).
- [25] B. Dubuc, C. Roques-Carmes, C. Tricot, and S. Zucker, in *Visual Communications and Image Processing II*, Vol. 845 (World Scientific, Singapore, 1987), pp. 241–248.
- [26] B. Spehar, N. Walker, and R. P. Taylor, *Front. Hum. Neurosci.* **10**, 350 (2016).
- [27] B. Spehar and R. P. Taylor, in *Human Vision and Electronic Imaging XVIII*, Vol. 8651 (SPIE, Bellingham, WA, 2013), p. 865118.
- [28] Note that while our method is motivated by fractal considerations, x does not match the measured fractal dimensions displayed in Table I, due to both finite-size effects and the non-overlapping boxes constraint.
- [29] S. Losh-Hesselbart and F. J. Fowler, *J. Am. Stat. Assoc.* **80**, 1077 (1985).
- [30] S. Aral, MIT Sloan Manage. Rev. **55**, No. 2, 47 (2014).
- [31] I. C. McManus, *Br. J. Psychol.* **71**, 505 (1980).
- [32] Zooniverse, <https://www.zooniverse.org/>, accessed: 2019-06-01.
- [33] Note that N_i was not exactly equal to N since a small number of participants stopped before completing the 15 battles.
- [34] D. Tolhurst, Y. Tadmor, and T. Chao, *Ophthalmic Physiol. Opt.* **12**, 229 (1992).
- [35] D. H. Zanette, *PLoS ONE* **13**, e0207879 (2018).
- [36] M. R. Chernick, W. González-Manteiga, R. M. Crujeiras, and E. B. Barrios, *Bootstrap Methods* (Springer, New York, 2011).
- [37] Mechanical Turk, <https://www.mturk.com/>, accessed: 2019-07-01.
- [38] J. M. Keynes, *General Theory, Employment, Interest and Money* (Palgrave-McMillan, London, 1936).
- [39] S. Aaronson, S. M. Carroll, and L. Ouellette, Quantifying the rise and fall of complexity in closed systems: The coffee automaton, [arXiv:1405.6903](https://arxiv.org/abs/1405.6903).
- [40] Spehar *et al.* [26] showed that the preference curve was hardly affected by the gray scale to B&W transformation. For the sake of simplicity we thus choose to apply the coarse-graining procedure to B&W images.
- [41] Y. C. Zhang, *Journal de Physique I*, **1**, 971 (1991).
- [42] M. Wertheimer, *Gestalt Theory* (Kegan Paul, Trench, Trubner, and Company, London, 1938).
- [43] <https://www.econophysix.com/>.