



Beauty and structural complexity

Samy Lakhal, Alexandre Darmon, Jean-Philippe Bouchaud, Michael Benzaquen

► To cite this version:

Samy Lakhal, Alexandre Darmon, Jean-Philippe Bouchaud, Michael Benzaquen. Beauty and structural complexity. Physical Review Research, 2020. hal-02323024

HAL Id: hal-02323024

<https://hal.science/hal-02323024>

Submitted on 21 Oct 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Beauty and structural complexity

Samy Lakhal,^{1,2,3} Alexandre Darmon,¹ Jean-Philippe Bouchaud,^{3,4} and Michael Benzaquen^{2,3,4,*}

¹*Art in Research, 33 rue Censier, 75005 Paris, France*

²*Ladhyx UMR CNRS 7646, Ecole polytechnique, 91128 Palaiseau Cedex, France*

³*Chair of Econophysics & Complex Systems, Ecole polytechnique, 91128 Palaiseau Cedex, France*

⁴*Capital Fund Management, 23 rue de l'Université 75007 Paris, France*

(Dated: October 15, 2019)

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 judgement.

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 [1–12]. According to Kant, *Is beautiful that which pleases universally without a concept* [13]. 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 [14]. 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 mathematisation 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 a universal beauty is a longstanding debate which has known many ruptures through the history of art [2] and still opposes a number of great modern thinkers.

Physicists' interest in the subject is more recent. Stephens *et al.* [15] 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 [16], *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 organised as follows. We first present and motivate our image generation methods. We then present and analyse 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 conclude.

There exist many possible measures of image complexity, relying on e.g. their mathematical properties [17, 18], their physical properties [15, 19, 20], or even their cognitive impact [21–23]. Here we choose to work with three simple measures that can be easily computed unequivocally for any digital 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 greyscale 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 [24]. After transforming the image to B&W using an intensity threshold ensuring two equally-populated levels [15], the fractal dimension follows $N(\epsilon) \sim \epsilon^{-d_f}$ where $N(\epsilon)$ 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 sizes of the PNG compressed and uncompressed image.

In order 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. 1a) is generated by *reverse-engineering* the Fourier

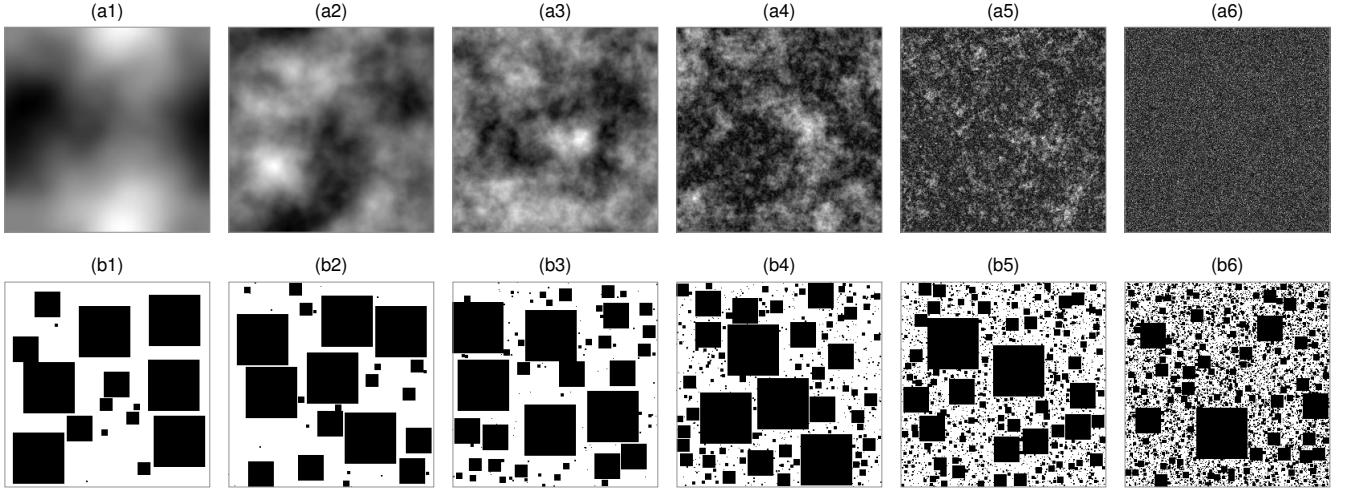


FIG. 1. (a) Fourier Magnitude-generated images, and (b) Box-counting-generated images, both series with increasing complexity from left to right. These images were used for our large-scale survey.

Magnitude property: setting $\hat{u}(k, \theta) = k^\alpha e^{i2\pi\eta(\mathbf{k})}$, where η is drawn from a uniform distribution on $[0, 1]$ with $\langle \eta(\mathbf{k})\eta(\mathbf{k}') \rangle \propto \delta_{\mathbf{k}\mathbf{k}'}$, and taking the inverse Fourier Transform of \hat{u} allows to produce a series of random greyscale images $|u(r, \phi)|$ with controlled Magnitude slope [22, 25, 26]. Table I gathers the computed complexity measures α , d_f and τ of the 256×256 images displayed in Fig. 1a. 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. 1b). Black squares of size $\epsilon = 2^k$, $k \in [1, k_{\max}]$, drawn from the distribution $N_x(\epsilon) = A\epsilon^{-x}$ 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. Here again the complexity measures appear to be increasing functions of one another (see Tab. I).

In 2013, Spehar and Taylor [26] conducted a survey on twenty-six 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 [22, 25]. Curious of their results, we conducted a

larger scale experiment intended for a larger panel (over a thousand participants with different backgrounds), using the images presented in Fig. 1. Our question at this stage is 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 [27]. 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 have been shown to be weighted by extreme grades, the utility given to intermediate grades being far from linear [28]. Five-star ratings image-by-image can also be rather disorienting due to the lack of reference. Another option is image classification, where 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 recognising other features, such as complexity, and ending by sorting them

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	0.045	0.074	0.12	0.22	0.40	0.41
	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

by something else than preference. Finally, the *Battle survey* consists in successively presenting to the interviewee all possible sets of two images and asking them to choose the one they prefer [4, 29]. While less time-efficient (with $N(N - 1)/2 = O(N^2)$ battles, for $N = 6$ 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, intellectually less challenging. We thus choose the latter method.

We conducted three slightly different surveys. The panel for the first survey consisted of colleagues from CFM and Ecole Polytechnique as well 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 as 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 [30] 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 couldn't 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 (resp. loses) a battle, its score S_i increases (resp. decreases) by $1/N_i$ where N_i denotes the number of battles in which i was involved [31]. 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 [15, 32] and visual arts [17]. Discussions with voters revealed that they found their favorite images to be the most *harmonious* and *well balanced*.

In order to increase the size and diversity of the panel, we ran two other experiments on the Mechanical Turk platform [33], in which participants are paid a small amount of money to participate (we reached ≈ 700 panelists). The first attempt on this platform (dotted light-gray line in Fig. 2) appeared to give rather noisy results, especially for the second set of images. We partly attribute this to a fraction of participants answering randomly to the battles, which is likely much less significant when people participate selflessly and in good will. In general, survey participants on such paying platforms are evaluated and only get paid if they completed their tasks satisfactorily, but in our case one cannot really evaluate the panelists since *satisfactorily* here only stands for *honestly*. In the second tentative (dash-dotted gray line in Fig. 2) we formulated the question differently in order to encourage non-random participation: much like in Keynes' famous beauty contest [34], people were asked

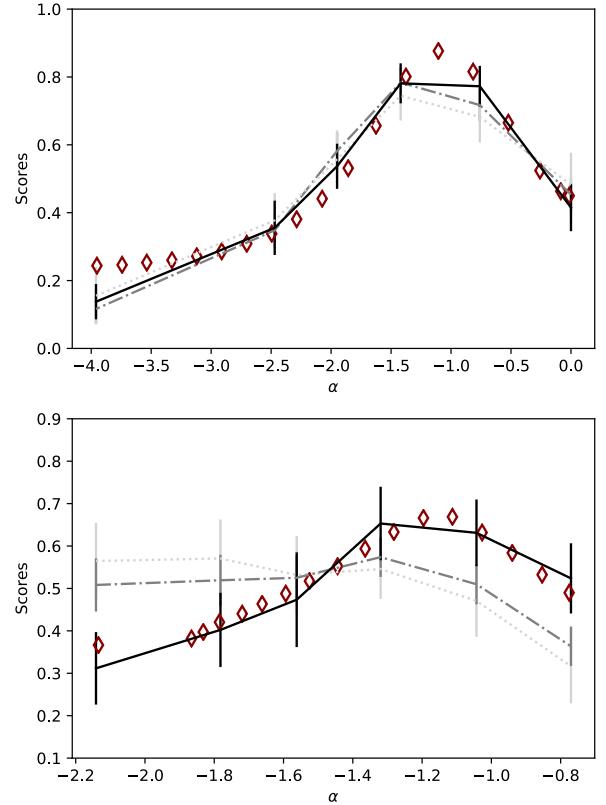


FIG. 2. Results of the three different surveys (Zooniverse: solid black line, Mechanical Turk 1: dotted light gray-line, Mechanical Turk 2: dash-dotted gray line). The diamonds markers indicate the structural complexity τ_{cg} defined below. We have rescaled and shifted vertically τ_{cg} to show that the maximum scores also correspond to maximum structural complexity. Top: image series of Fig. 1(a). Bottom: image series of Fig. 1(b). Error bars reflect population heterogeneity.

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). While probably introducing other biases, the results displayed less noise and better 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 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 ob-

tained 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.* [18] clearly resonates with such questions. Guided by the idea that there is no perceptual structure in white noise, the authors attempted to characterise 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 [35]. 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 display very regular homogeneous structures, the transitional regime displays interesting and *complex* mixing patterns as the cream/coffee interface slowly disappears. So far we have 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 non-monotonous function of time displaying a maximum at intermediate stages where the non-trivial 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.* [35] 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 [36]. The colour 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 non-monotonous functions displaying a maximum for intermediate values of α , d_f and τ . Since the y -axis is rather arbitrary we have rescaled and shifted $\tau_{cg}(\alpha)$ for an easier comparison with the survey data. Furthermore, the scale parameter r_{cg} and the threshold

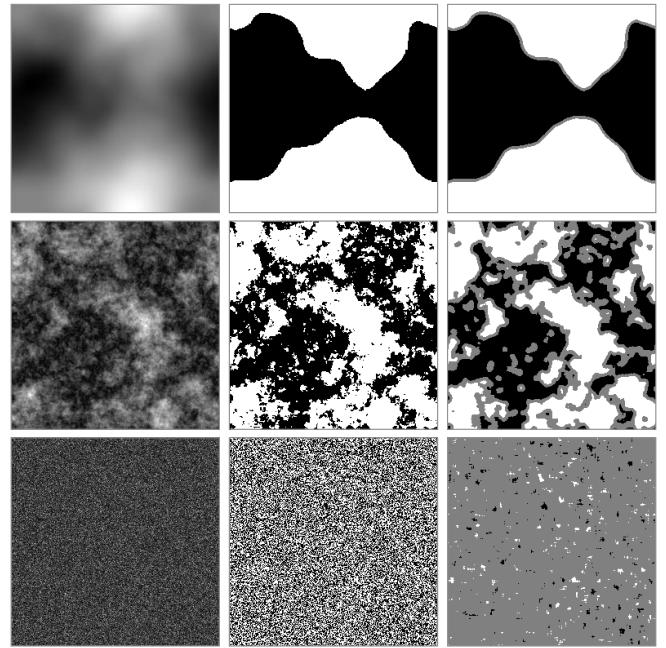


FIG. 3. Illustration of the coarse-graining procedure on images a1, a4 and a6, with $r_{cg} = 7$ and $\delta = 0.23$.

η 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 $(r_{cg}, \eta) = (7, 0.23)$ 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. This quantitatively supports the idea that structural complexity is a good proxy for average image preference.

Let us summarise what we have 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* [37]: 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.

* michael.benzaquen@polytechnique.edu

- [1] G. T. Fechner, Vorschule der aesthetik, Vol. 1 (Breitkopf & Härtel, 1876).
- [2] U. Eco and G. De Michele, Storia della bellezza (Bompiani Milano, 2004).
- [3] E. H. Gombrich, Art and illusion (Princeton, 1969).
- [4] E. A. Vessel and N. Rubin, Journal of vision **10**, 18 (2010).
- [5] H. J. Eysenck, The American Journal of Psychology **54**, 385 (1941).
- [6] M. Bar and M. Neta, Psychological science **17**, 645 (2006).
- [7] B. Spehar, C. W. Clifford, B. R. Newell, and R. P. Taylor, Computers & Graphics **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, Journal of Environmental Psychology **24**, 247 (2004).
- [10] I. C. McManus, European Review **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, British Journal of Psychology **102**, 49 (2011).
- [13] I. Kant, Critique of the Power of Judgment (Cambridge University Press, 2000).
- [14] E. Aujaleu, Tréma , 33 (1997).
- [15] G. J. Stephens, T. Mora, G. Tkacik, and W. Bialek, Physical review letters **110**, 018701 (2013).
- [16] J.-P. Bouchaud, Leonardo **41**, 239 (2008).
- [17] M. Koch, J. Denzler, and C. Redies, PLoS one **5**, 12268 (2010).
- [18] A. Desolneux, L. Moisan, and J.-M. Morel, From Gestalt Theory to Image Analysis (Springer New York, 2008).
- [19] B. Gidas, IEEE Transactions on Pattern Analysis and Machine Intelligence **11**, 164 (1989).
- [20] S. Foreman, J. Giedt, Y. Meurice, and J. Unmuth-Yockey, Phys. Rev. E **98**, 052129 (2018).
- [21] D. J. Field, Journal of the Optical Society of America A **4**, 2379 (1987).
- [22] B. Spehar, S. Wong, S. van de Klundert, J. Lui, C. W. G. Clifford, and R. Taylor, Frontiers in Human Neuroscience **9**, 514 (2015).
- [23] R. Taylor, B. Spehar, C. Hagerhall, and P. Van Donkelaar, Frontiers in Human Neuroscience **5**, 60 (2011).
- [24] B. Dubuc, C. Roques-Carmes, C. Tricot, and S. Zucker, in Visual Communications and Image Processing II, Vol. 845 (1987) pp. 241–248.
- [25] B. Spehar, N. Walker, and R. P. Taylor, Frontiers in Human Neuroscience **10**, 350 (2016).
- [26] B. Spehar and R. P. Taylor, in Human vision and electronic imaging XVIII, Vol. 8651 (2013) p. 865118.
- [27] S. Losh-Hesselbart and F. J. Fowler, Journal of the American Statistical Association **80**, 1077 (1985).
- [28] S. Aral, MIT Sloan Management Review **55**, 47 (2014).
- [29] I. C. McManus, British Journal of Psychology **71**, 505 (1980).
- [30] “Zooniverse,” <https://www.zooniverse.org/>, accessed: 2019-06-01.
- [31] Note that N_i was not exactly equal to N since a small number of participants stopped before completing the 15 battles.
- [32] D. Tolhurst, Y. Tadmor, and T. Chao, Ophthalmic and Physiological Optics **12**, 229 (1992).
- [33] “Mechanical turk,” <https://www.mturk.com/>, accessed: 2019-07-01.
- [34] J. M. Keynes, General theory, employment, interest and money (Palgrave-McMillan, 1936).
- [35] S. Aaronson, S. M. Carroll, and L. Ouellette, “Quantifying the rise and fall of complexity in closed systems: The coffee automaton,” (2014), arXiv:1405.6903.
- [36] Spehar *et al.* [25] 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.
- [37] M. Wertheimer, Gestalt theory (Kegan Paul, Trench, Trubner & Company, 1938).