



Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Preference for luminance histogram regularities in natural scenes

Daniel Graham^{a,*}, Bianca Schwarz^b, Anjan Chatterjee^c, Helmut Leder^b

^a Department of Psychology, Hobart and William Smith Colleges, United States

^b Faculty of Psychology, University of Vienna, Austria

^c Department of Neurology, University of Pennsylvania, United States

ARTICLE INFO

Article history:

Received 3 September 2014

Received in revised form 9 March 2015

Available online xxx

Keywords:

Natural scenes
Statistical regularities
Efficient coding
Empirical aesthetics
Skewness
Art statistics

ABSTRACT

Natural scene luminance distributions typically have positive skew, and for single objects, there is evidence that higher skew is a correlate (but not a guarantee) of glossiness. Skewness is also relevant to aesthetics: preference for glossy single objects (with high skew) has been shown even in infants, and skewness is a good predictor of fruit freshness. Given that primate vision appears to efficiently encode natural scene luminance variation, and given evidence that natural scene regularities may be a prerequisite for aesthetic perception in the spatial domain, here we ask whether humans in general prefer natural scenes with more positively skewed luminance distributions. If humans generally prefer images with the higher-order regularities typical of natural scenes and/or shiny objects, we would expect this to be the case. By manipulating luminance distribution skewness (holding mean and variance constant) for individual natural images, we show that in fact preference varies inversely with increasing positive skewness. This finding holds for: artistic landscape images and calibrated natural scenes; scenes with and without glossy surfaces; landscape scenes and close-up objects; and noise images with natural luminance histograms. Across conditions, humans prefer images with skew near zero over higher skew images, and they prefer skew lower than that of the unmodified scenes. These results suggest that humans prefer images with luminances that are distributed relatively evenly about the mean luminance, i.e., images with similar amounts of light and dark. We propose that our results reflect an efficient processing advantage of low-skew images over high-skew images, following evidence from prior brain imaging results.

© 2015 Published by Elsevier Ltd.

1. Introduction

The distribution of light intensities in the natural world plays a fundamental role in vision. Mechanisms of adaptation evolved to allow species to tune their visual systems to the proportions of different light intensities in the immediate natural environment (see, e.g., Baccus, 2007). However, our perception of natural scenes is also invariant to large changes in luminance distributions, especially with regard to higher order statistical moments. For example, we readily recognize a scene as being the same scene at different times of day or in different weather. We can also recognize a scene whether we see it in person or in a picture. In addition, the example of human-made pictures is particularly intriguing from the perspective of natural vision: Such images “work” despite the fact that typical natural scenes have a far larger dynamic range and more highly skewed histograms than paintings (Graham & Field, 2007, 2008).

Here we examine luminance statistics of natural images, focusing on the skewness (third statistical moment) of luminance distributions. Skewness is of interest for a variety of reasons, but primarily because there is evidence for its role in aspects of natural vision. Higher-order statistics such as skewness and kurtosis appear to be regular in natural luminance distributions. In particular, natural scenes typically have positively skewed luminance distributions (Attewell & Baddeley, 2007; Brady & Field, 2000; Dror et al., 2001; Laughlin, 1981), in part because of natural scenes’ high dynamic range. Schemes for efficient neural coding of this regularity have been proposed (Brady & Field, 2000; Richards, 1981).

With regard to aesthetics, basic spatial and luminance statistics relevant to efficient processing can predict significant portions of variance in similarity and preference judgments for paintings (Graham et al., 2010). It has also been shown that artwork tends to have more isotropic orientation spectra (Redies et al., 2007) compared to many types of natural images, due perhaps to a de-emphasis of natural scenes’ horizontal structure in paintings at certain scales (Schweinhart & Esock, 2013).

* Corresponding author.

E-mail address: graham@hws.edu (D. Graham).

However, existing data on luminance distribution skewness in images could support one of two possible predictions regarding preference.

First, we might expect preference to be shaped by natural regularities in skewness. Natural scenes' positive skew (Brady & Field, 2000) is due primarily to the heavy-tailed, high-dynamic-range distribution of luminances, which often spans a three or four decade range. We might therefore expect our preferences to simply align with regularities in nature, as has been suggested in relation to other image properties. For example, Redies (2008) argues that we prefer painted portraits that, in the spatial domain, are more like complex natural scenes since portraits tend to have spatial frequency spectra slope closer to those of natural scenes than to those of real faces (Redies et al., 2007). A related argument has been made regarding color, namely that in general we prefer blue over yellow because more positively affective components of our visual ecology are blue than yellow (Palmer & Schloss, 2010). Following this logic, we might therefore expect higher skew to be preferred since it is characteristic of complex natural scenes.

Another line of support for this prediction comes from the finding that skewness is often associated with glossiness in images (Motoyoshi et al., 2007; although high skew does not guarantee glossiness: see Anderson & Kim, 2009). Recent evidence shows that high luminance distribution skewness is a valid cue for freshness of fruits and vegetables¹ (Arce-Lopera et al., 2012; Wada et al., 2010). In addition, infants show preference for glossy objects (with high skew) starting as early as 7–8 months of age (Yang et al., 2011). Thus, if we tend to like shiny and/or fresh things, which tend to generate higher skew in luminance distributions, we may generally also prefer natural images with higher skewness. We term this the *matching nature hypothesis*.

A second hypothesis is that low absolute skew (i.e., skew near zero) would be preferred. In this view, we would take as evidence the fact that artists through the ages have, on average, produced images with low absolute skew. Low skewness in artwork is due in part to the fact that artists are limited in dynamic range compared to natural scenes (Graham, 2011; Graham & Field, 2007, 2008b; Graham et al., 2010), though it is possible produce a low-dynamic range image with high skew by hand. One could hypothesize that preference for low skewness could be partly due to a processing advantage for images with luminance distributions that are relatively evenly distributed about the mean. That is, an image with similar proportions of light and dark may be more aesthetic because it could be more efficiently processed. Such efficiency could sway higher-level cognitive processes associated with aesthetic judgment, or, in less precise terms, it could contribute to the “ease” of cognitive processing (i.e., processing fluency: Reber, Winkelman, & Schwarz, 1998). Thus, if low skewness is indeed efficiently processed by the human visual system, we would expect natural scenes with lower skewness to be preferred. We term this the *matching art hypothesis*.

Thus, we have two reasonable but incompatible hypotheses. Here we aim to address this question by testing human preference for natural images that have been manipulated to possess different higher order statistics, but that are otherwise identical. Following this approach, in Experiment 1, we test artistic photographs of dramatic natural landscapes. In Experiment 2, we test natural landscape images from a calibrated image database. In Experiment 3, we test calibrated natural images of objects. In Experiment 4, we test natural images whose pixels have been spatially randomized.

2. General methods

2.1. Participants

Participants were recruited from the University of Vienna subject pool in return for course credit (except for Experiment 1a, which employed uncompensated volunteers). All participants had normal or corrected-to-normal visual acuity and were naïve as to the purpose of the experiment. Written informed consent was obtained from all participants prior to participation and the experiment was carried out in accordance with the Declaration of Helsinki.

2.2. Stimuli

Adjustments to the source images' luminance distribution skewness were achieved using a gamma transformation. Once images with a range of 8 skew values were achieved for each scene, the 8 images were processed via linear scaling so that the luminance mean and variance was normalized (using the SHINE toolbox; Willenbockel et al., 2010), leaving skew values unaffected. Images were displayed on a black background.

2.3. Apparatus

Images were displayed in a darkened room to minimize stray light. In Experiment 1a, we presented the stimuli on a Samsung 2443 24-inch LCD monitor; in all other experiments we used a Samsung SyncMaster S24A300B, 24-inch LED backlit monitor. Both monitors were linearized in software with respect to luminance measured using a photometer (Konica-Minolta LS-100). In all experiments the participant's head position was fixed on a chin rest. Images in Experiment 1 subtended approximately $18^\circ \times 12^\circ$, and in Experiments 2, 3, and 4 they subtended approximately $16^\circ \times 12^\circ$.

2.4. Procedure

We used a two-alternative forced choice paradigm with paired comparisons. Each scene's eight versions were presented next to each other in pairs, which produced a total of 28 pairs per scene. Each trial consisted of a comparison of one version of a given scene with another version of the same scene. Presentation of the scenes was blocked and randomized and the presentation of the pairs was randomized to control for anchoring and ordering effects. Participants were instructed in German (except for 3 Erasmus students in Study 1 who received instructions in English) to choose the image in each pair they preferred by pressing the left or right arrow key on the PC keyboard. Stimuli were presented using the PsychToolBox (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997) for MATLAB (The MathWorks, Inc.).

3. Experiment 1

To investigate the basic effect of skewness on preference, we performed two experiments using artificially manipulated artistic natural images as stimuli. Experiment 1a and 1b involved the same source images and procedure but differed in the number of participants, and the adjusted skew values. This was done to sample a larger variety of skew values and in order to test separate subject pools. The display also differed in the two experiments, as described above.

¹ This result agrees with commercial practice since fruits and vegetables in supermarkets are often sprayed with water to give them a more glossy appearance despite the fact that this can cause them to rot faster.

3.1. Methods

3.1.1. Participants

Overall, 36 participants took part in the experiment. In Experiment 1a, there were ten participants, who were uncompensated volunteers (3 male, mean age: 42.5 years, SD = 15.26). In Experiment 1b, participants were students at the University of Vienna (3 male, mean age: 21.3 years, SD = 1.87), among them 3 Erasmus Students from Spain; all received course credit for participation.

3.1.2. Stimuli

Four scenes (see Supplement Fig. 1) were photographed as artistic images by one of us (A.C.) in the southwestern United States and were adjusted by the artist to his taste using software (primarily using contrast adjustments). These images thus represent the final outcome of a creative process aimed at producing an artistic picture (and they were produced prior to and independent of the current experiment), but they are not necessarily faithful maps of scene luminance.

For the experiments, images were linearly downsized to the same vertical pixel dimension (567 pixels) but varied slightly in their horizontal dimensions (~835 pixels). Color, 16-bit source images were transformed to intensity via the YIQ transform were gamma adjusted to create 8 versions, then down sampled to 8 bits and normalized for mean and variance.

For a given scene, 8 images with differing skewness were generated using the procedure described above. This was done twice: once for Experiment 1a and once for Experiment 1b, resulting in different skew values that spanned a similar range. Original (unmodified) skewness values for the source images (#1–#4) were [0.596, 0.953, 1.250 and 0.943]. Skewness values for presented stimuli ranged from –0.36 to 1.92 in Experiment 1a and from –0.58 to 2.30 in Experiment 1b. The first four statistical moments of the luminance distributions for all images used are given in Supplement Table 1. Image versions and associated pixel histograms from Experiment 1 are shown in Supplement Fig. 1.

3.2. Results

The reported preference values (across participants) represent the preference frequency for a given image (i.e., the proportion of trials in which a given image was preferred), which has a maximum of 0.25 for a given image (i.e., if it is chosen all 7 times it is presented among the 28 trials for a given scene). Tests with other ranking schemes (e.g., Bradley-Terry-Luce ranking algorithm; see, e.g., Graham et al., 2010) produced essentially the same results so preference frequency is presented throughout.

There is a clear preference in both Experiment 1a and 1b for low skew versions of all scenes; for example, the low skew version of image 4 (Exp. 1a) is more than nine times as preferred as the high skew version.

The scatterplot of the preference judgments against skewness for each experiments' images is seen in Fig. 1.

Because of the similarity in the stimulus sets and results, we also analyzed the aggregated data (Fig. 2). The correlation coefficient between preference and skewness of the pooled data was significantly high ($R = -0.83$, $p < 0.01$), showing a strong negative relation.

The results suggest that humans prefer low skew in artistic natural images. Indeed, in all cases humans prefer versions of these images with skew substantially lower than the skew of the original source image (to the chagrin of the artist!).

A second test using the same images and procedure was performed to investigate the effect of the background tone (see Supplemental Study 1). Instead of a black background (which

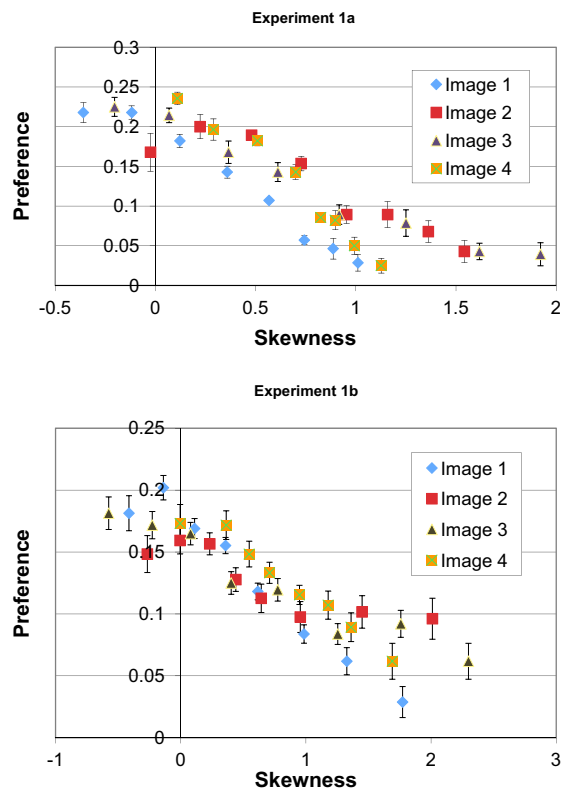


Fig. 1. Preference judgements for each individual image plotted against skewness in Experiment 1a (top) and Experiment 1b (bottom). Error bars indicate standard error calculated over participants.

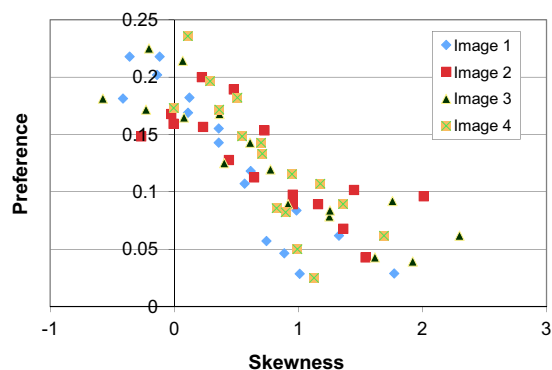


Fig. 2. Aggregated results for Experiments 1a and 1b.

was used in all other experiments in this study), we also performed the same experiment with a mid-gray background. This test produced essentially the same results as Experiment 1. This result is discussed in greater detail in the (Section 8).

3.3. Discussion

We have shown that versions of artistic photographs with higher positive luminance histogram skewness are markedly disliked compared to versions of the same images that have skew near zero. This result holds for separate participant pools, display types, and specific skew values.

Although we did not observe a fall-off in preference for all negatively skewed images, we believe that such a fall-off exists since more negatively skewed images are both uncommon in nature

and unlikely to show glossiness. Moreover, six out of eight sets of images in Fig. 1 were best described by a concave quadratic regression curve peaking near zero. However, we were limited in the range of negative skews we could produce due to using natural images, which are prone to visual artifacts after gamma transformations with exponents close to zero. In particular, smaller exponents led to banding artifacts in slowly varying tones (e.g., sky in image 4). Such artifacts would have confounded our results (see Fig. 3). Therefore, based on the current test, we can conclude only that more positively skewed natural images become less attractive.

4. Experiment 2

We designed this study to remove the artistic aspect and the associated possibility of manipulation of colors and luminances by using natural landscape images from a calibrated database. Also, since glossy surfaces could be associated with positive skew (Motoyoshi et al., 2007) and because people may tend to like shiny surfaces, we wanted to test if there is still no preference for positive skew even if using images depicting glossy surfaces.

4.1. Methods

4.1.1. Participants

Seventeen students (12 female, mean age: 25.5, SD: 10.0) at the University of Vienna participated in the experiment for course credit.

4.1.2. Stimuli

Four images of natural scenes, depicting glossy elements (like water or cars in sunshine), were taken from the McGill Calibrated Color Image Database (Olmos & Kingdom, 2004) to generate the stimuli. The original images were 8-bits per channel and consisted of 768×576 pixels. Images were linearized with respect to camera parameters, manipulated using a gamma transform and normalized in the same way as the images in Experiment 1, and were presented on the linearized display used in Expt. 1b. All stimuli were created from images labeled 1–4 (see Fig. 4). Original (unmodified) skewness values for the linearized source images (#1–#4) were [0.245 0.787 0.050 2.044]. Skewness values for presented stimuli ranged from -1.24 to 2.52 .

4.2. Results

We again observe a strong preference for low skew images. The study resulted in a high negative correlation coefficient ($R = -0.78$, $p < 0.01$) between preference and skewness (see Fig. 5). When data for the four images are aggregated as a group, we observe a strong negative linear relationship between skewness and preference ($R^2 = 0.612$; for quadratic fit, $R^2 = 0.620$). Thus, patterns of preference for McGill images followed those for artistic photographs, with skew values near zero strongly preferred over high skew values. Note also that preference peaks at skew values substantially lower than the original, unmodified skew value for each scene.

5. Experiment 3

The aim of Experiment 3 was to investigate the effect of skewness on preference for close-up images depicting glossy objects. This study was motivated by past findings of associations between skewness and glossiness in close-up images of real and computer-generated objects (Arce-Lopera et al., 2012; Motoyoshi et al., 2007; Wada et al., 2010; Yang et al., 2011). Images were again drawn from the McGill collection.

5.1. Methods

5.1.1. Participants

Twenty-nine students (20 female, mean age: 25.2, SD: 7.4) at the University of Vienna participated and received course credit for their participation.

5.1.2. Stimuli

Four images from the McGill Calibrated Colour Image Database (Olmos & Kingdom, 2004) were used depicting close-up views of natural objects with glossy surfaces. Images were processed and displayed as in Experiment 2. All stimuli were created from images labeled 1–4 (see Fig. 6). Original (unmodified) skewness values for the linearized source images (#1–#4) were [1.323, 1.653, 1.481 and 1.524]. Skewness values for presented stimuli ranged from -0.50 to 1.64 .

5.2. Results

For images of glossy objects we found a significant negative relationship between preference and skewness (linear fit: $R^2 = 0.615$; quadratic fit: $R^2 = 0.790$) for all data considered together. Each of the images individually also followed the same pattern as that observed in Experiments 1 and 2, being described best by a convex quadratic fit (see Fig. 7). Again, images with skewness near 0 are preferred most and preference peaks at skew values substantially lower than the original, unmodified skew value in each case.

Although we again had no stimuli with extremely negatively skewed luminance distributions, we can see that each images' most negatively skewed version was not the most preferred. In fact, we note that preference appears to decrease the more negatively skewed the images' distributions become. Nevertheless, in the present experiment we can conclude only that high positive skewness is not preferred.

6. Experiment 4

We designed this experiment to test the possible effect of content on preference using random images as test stimuli. We therefore removed all image content through pixel scrambling, which preserves all pixel intensity statistics.

6.1. Methods

6.1.1. Participants

Twenty-four students (13 female, mean age: 22.4, SD: 3.8) at the University of Vienna participated in the experiment for course credit.

6.1.2. Stimuli

We randomized the pixel positions for each of the 32 stimuli of Experiment 3 via MATLAB script (see Fig. 8).

6.1.3. Apparatus and procedure

Apparatus was the same as in Experiment 3. Additionally to the procedure of Experiment 3, the participants were told that they would see noise images at the outset of the experiment. Furthermore, we asked participants after the experiment whether they had pursued a strategy during testing, and if they had recognized specific structures or patterns in the images.



Fig. 3. Top row: Two source images (#3, left image; and #4, right image) used in Experiment 1. Bottom row: Examples showing the limitations of stimuli manipulation at very low (i.e., most negative) skew values. The image on the left (Skew = -0.81) and the image on the right (Skew = -0.31) show banding artefacts in the sky.

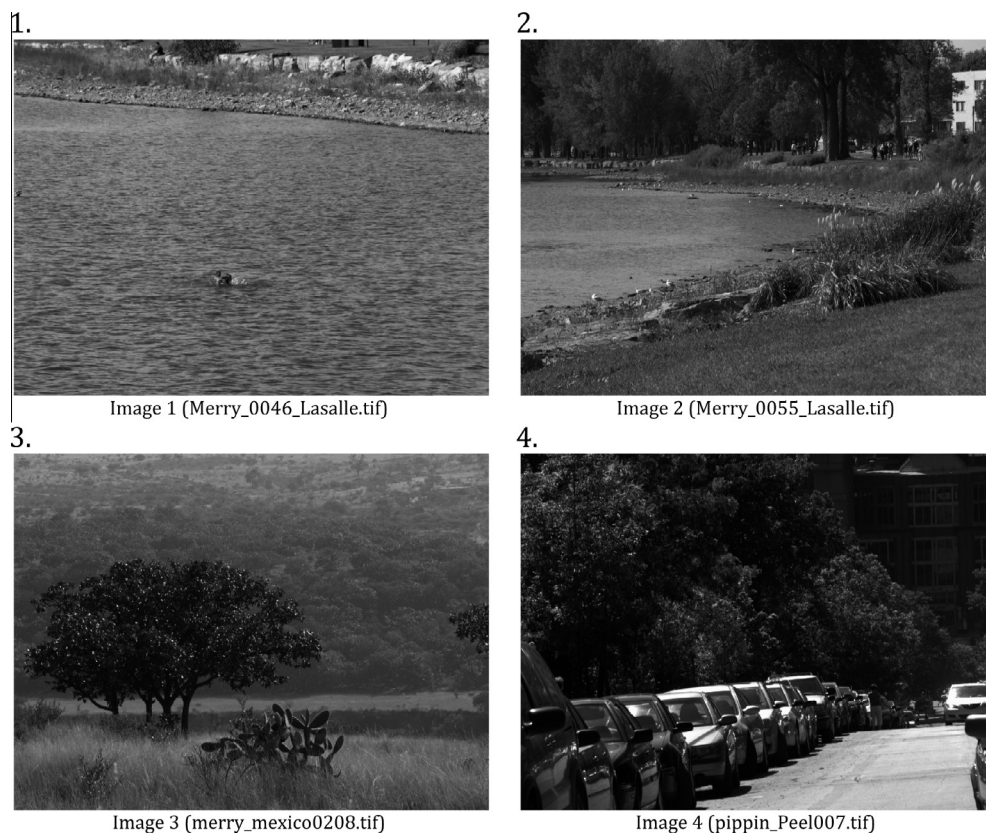


Fig. 4. Grayscaled and calibrated natural landscape images depicting glossy surfaces. All 32 stimuli used in Experiment 2 were created from them. Original file names are shown in parentheses.

6.2. Results

First, it is notable that the data indicate that humans have consistent patterns of preference for our stimuli, which vary only in

subtle ways (to human observers) and are devoid of recognizable content. Though the data are noisier than in the unscrambled case, there was again a strong correlation between preference and skewness for the data considered together (linear fit: $R^2 = 0.742$;




Fig. 5. Preference plotted against skewness for individual images in Experiment 2, utilizing landscape images from the calibrated McGill Natural Scene Database. The lines represent each images' quadratic regression curve.

quadratic fit: $R^2 = 0.751$). Considered individually, preference versus skewness shows image 1 ($R^2 = 0.822$) and image 4 ($R^2 = 0.910$) show a convex quadratic relationship, image 2 ($R^2 = 0.920$) and image 3 ($R^2 = 0.690$) show a concave quadratic relationship (see Fig. 9).

Nineteen out of 24 participants (79%) reported that they sometimes saw structures or even objects (e.g. lines, maps, flowers, landscapes) in the images. This could result from the visual system's tendency to seize upon small variations in structure in a white noise display (see, e.g., Gosselin & Schyns, 2003).

We also compared participants' results based on their reported strategies. Twelve different strategies were reported (e.g.

'intuitively', 'no strategy', 'alternatively bright and dark'). Five participants, who reported using the strategy to pick the brighter image, tended to prefer low skew images. The slope values of the linear regression curves for these participants range between -0.10 and -0.03 . Four participants, who mentioned other preferences (images with patterns, more contrast) and reported that they additionally preferred brighter images, also judged low skew images best. Slopes for these participants range between -0.11 and -0.02 . Two participants reported that they preferred darker images. Their linear regression fits show the most positive slopes (0.06 and 0.12) of all participants in Experiment 4.

6.3. Discussion

Our results again are consistent with the notion that humans like a balance of light and dark, even in the absence of image structure. Also, though our sample size in Experiment 4 is too small to fully describe individual differences with respect to reported strategies, qualitative reports by subjects suggest that high positively skewed images are perceived as dark, whereas low skew images are perceived as bright (despite having normalized intensity mean and variance). Therefore our results could suggest that low skew images are not just preferred but at the same time also perceived as brighter.

Motoyoshi et al. (2007) found contrary results in a related experiment. When they asked their observers to rate "lightness" of pixel-randomized images (though it is unclear that an artificial image of this kind can be considered to have lightness in a strict sense), the ratings did not vary much over the range of skewness values. In their experiment they tested only 6 subjects so the statistical power to detect an effect may have been too small. The difference in task may have played a role as well. Another explanation




Fig. 6. Grayscaled and calibrated natural images. The original images were gathered from the McGill Calibrated Colour Image Database (Olmos & Kingdom, 2004) and used to generate the stimuli for Study 3. Original file names are shown in parentheses.

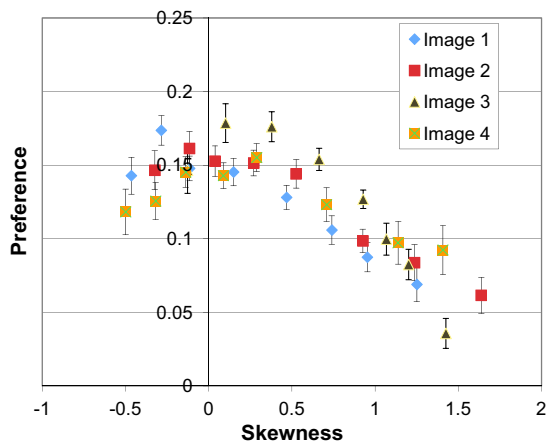


Fig. 7. Preference plotted against skewness for Experiment 3. Curves represent quadratic fits for each of the four images.

for our results could be that preference is a more sensitive or consistent measure than “lightness” judgment. Also, the very smooth, artificial distributions used by Motoyoshi et al. (2007) may have different properties compared with the more bumpy and irregular luminance distributions in natural scenes. In any case, we believe our qualitative data should be treated cautiously, but they provide intriguing clues for future investigations.

7. Further analyses

7.1. Mean intensity

We investigated if mean intensity had an effect on preference in Experiments 1–3. In our studies we used source images with mean pixel values ranging from 90 to 137. This was done in part because very high or low mean intensity images can be subject to clipping (floor and ceiling) effects after gamma transformation. To study the effect of mean intensity on preference we put the data of Experiments 1–3 together and divided them into two groups according to mean luminance. To gain roughly the same sample sizes we used the average mean value of 115 as the cut-off value. Results showed that the slope of the linear regression fit of the images with a low mean (slope = -0.049 , $R^2 = 0.65$) was nearly the same as the one for the images with a high mean (slope = -0.058 , $R^2 = 0.50$). This indicates that in high mean images low skew was a bit more preferred than in low mean images. Therefore we conclude that mean intensity influences the relation between preference and skewness only to a small extent, at least in the range tested.

7.2. Low-luminance artifacts

Another potential issue has to do with potential artifacts introduced through changes in gamma. In particular, given that gamma adjustment is a nonlinear operation that affects different intensity regimes in different ways, images that are made to have more skewness than the original photograph should have lower resolution of pixel values at relatively darker portion of the image, which could lead to artifacts (e.g., distortions in objects represented by smooth variations in low pixel values). However, the source images used in Experiment 1 were not linearized or calibrated with respect to luminance—whereas for the other experiments, luminance-calibrated images were employed (McGill images). Thus, we have a basis for evaluating the potential contribution of low luminance distortions. In particular, distortions

should be less noticeable for the uncalibrated images, in which there is proportionally more resolution available for lower pixel values due to the compressive nonlinearity of the camera (which prioritizes low luminance resolution over high luminance resolution). In other words, in Experiment 1, we should observe greater preference for the highest skewness images – compared to the highest skewness images in the other experiments – since the source images (i.e., those whose skewness is manipulated) are subject to a compressive camera nonlinearity and would therefore be less subject to artifacts. However, we instead find the relationship of preference with respect to skewness shows a similar linear-fit slope in Experiments 1, 2, and 3 (slope ~ -0.1).

In any case, the fact that we observe the skewness effect (albeit in weaker form, as indicated by the smaller slope of the linear fit, ~ 0.03) for pixel scrambled images – where no object-based distortions are possible – suggests that the primary cause is indeed related to skewness, and not to artifacts.

8. General discussion

We have shown that low skewness in natural scene luminance distributions is strongly preferred to high positive skewness in otherwise identical images, and we have demonstrated this across a range of conditions: (1) in artistic photographs of landscapes; (2) in benchmark natural scene images; and (3) in pixel-randomized natural scenes. Our tests also suggest that the inclusion of glossy objects and variations in the background intensity in our display apparatus and in the depth of field in the natural scenes do not alter this result. Thus, our data appear to rule out the “matching nature” hypothesis, and to support the “matching art” hypothesis.

In more general terms, we can describe this effect as a preference for similar proportions of light and dark in an image: a low skew image is one that has a luminance distribution that is more symmetric about the mean, and thus one that has an approximate balance between light and dark pixels. Others have suggested similar notions (Moon & Spencer, 1944).

How can we explain these findings? We suggest that low skew images may be more efficiently processed than high skew images, which gives rise to greater preference for low skew images. Olman et al. (2008) report that BOLD activation for images with high absolute skewness was significantly higher in V1 than the activation found for images with zero skewness. The reported result was found for random noise stimuli, but also held for stimuli with more natural spatial structure, such as difference-of-Gaussian and Gabor patches (Doerschner, personal communication). If a similar effect exists for natural scenes, we argue that low skew images may in general be more efficient with respect to early cortical processing. This finding parallels those that show successful grouping of object elements appears to lower V1 activity (Fang, Kersten, & Murray, 2008).

A related argument concerns asymmetries in the ON and OFF pathways in the retina and thalamus. Since ON pathways are most sensitive to white, and OFF pathways are most sensitive to black, an efficient use of both pathways would entail roughly equivalent stimulation across the retina at a given time, whereas images with high absolute skewness would disproportionately stimulate one pathway or the other.

If these effects do indeed lead to efficient processing, low skew stimuli could be perceived as more aesthetic due to the ease of processing such images compared to high skew images. Thus, while the brain is in general well matched to the regularities of the world, images with low skewness may be closer to optimal for this system in terms of efficiency.

This argument is consistent with the fact that artworks from across cultures and time periods have low skew luminance

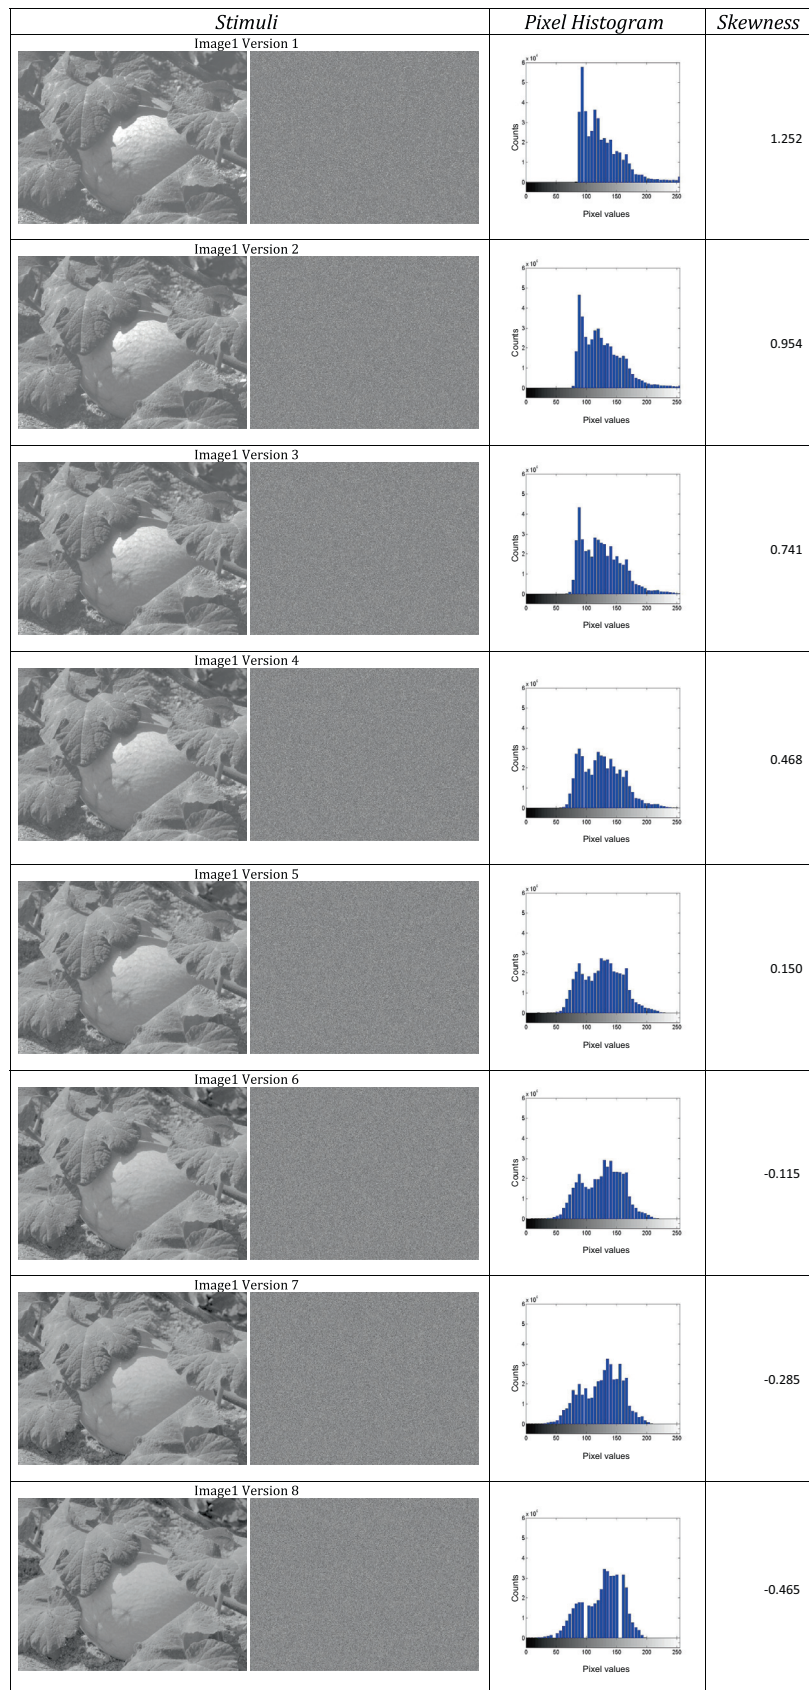


Fig. 8. Versions of Image 1 used in Experiment 3 (merry_mexico0230.tif, Olmos & Kingdom, 2004) with the corresponding pixel-scrambled images, pixel histograms and skew values in the rightmost columns.

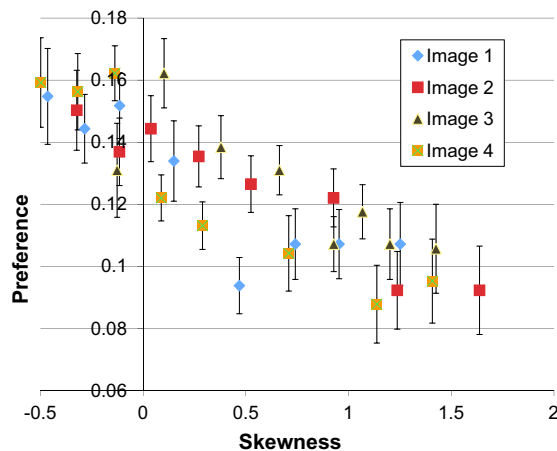


Fig. 9. Results of preference versus skewness with quadratic fits for each of the four images in Experiment 4 (scrambled images).

histograms (Graham & Field, 2007, 2008b; Graham, 2011; Graham et al., 2010). While some geographic areas seem to favor artwork with somewhat negative skewness (Graham & Field, 2008), which Motoyoshi (2011) argues is due to climatic and atmospheric differences, there is a relatively strong tendency for art on average to have a skew near zero. For the distribution of intensity distribution skews of paintings in the Cornell Corpus of World Artwork, $M = .046$, $SD = 1.13$ (Graham & Field, 2007, 2008b); therefore most artwork likely has an absolute skew of less than 0.5. Given the present results, we suggest that although artists make low skew images in part because of luminance range limitations in their media, they may also be influenced by the fact that low skew images are preferred. Thus, the present results could be seen as further evidence that artists (knowingly or not) tailor their productions to the biases and processing strategies inherent in primate visual coding (Graham & Field, 2007, 2008a, 2008b; Graham, 2009; Graham & Meng, 2011).

We should be clear that it is not the case that images with high absolute skew cannot be aesthetic: many paintings in the Cornell Corpus have high skew but may yet be aesthetic, and we suspect that some well-known artworks (and other aesthetic images) have successfully violated this tenet as well.

8.1. Alternative explanations

8.1.1. Minimum pixel value

We note that other explanations of our results are possible. For example, skewness for images in our experiments also is strongly correlated to the lowest (minimum) intensity in an image. In Experiments 1, 2, and 3, the linear regression of skewness with minimum pixel value (aggregated across images in a given experiment) has $R^2 = [0.79 \ 0.53 \ 0.88]$. Thus, minimum pixel value is a possible confound. The decrease in minimum intensity value with decreasing skewness is a byproduct of our methods, which keep mean and variance the same for each version of a scene. In real scenes, manipulations of luminance distribution that (a) preserve luminance mean and variance and (b) leave fully “black” pixels unchanged lead to image discontinuities and distortions such as banding (as with manipulations that achieve high negative skew). That is, places where image structure has smooth variation from black to lighter values (e.g., sky) will generate artifactual structure that will have a negative effect on preference.

However, one could argue that the darkest nonzero pixel in each image should anchor the scene and be perceived as black (Gilchrist et al., 1999). Indeed, the preference for images with

darker blacks did not require pixels with intensity equal to zero (though in most cases the image with the most negative skew also had a minimum intensity of 0). Our results using the mid-gray background (Supplemental Results 1), which were essentially the same as those with a black background, are consistent with the notion that black is anchored to the lowest intensity in the display, not to the black background per se. Moreover, having the lowest intensity is not enough—there are differences in preference with respect to images with the same minimum value. For example, consider the following data points: the three lowest skew versions of images 1–3 in Experiment 1a and 1b; the two lowest skew versions of image 2 in Experiment 2; the three lowest skew versions of images 1, 3 and 4 in Experiment 3; and the two lowest skew versions of image 2 in Experiment 3. In each case, the minimum intensity is 0 (black), but there are yet differences in preference for these images (with respect to one another). This suggests that there is a kind of “sweet-spot” of preference with respect to skew that is not captured solely by minimum value.

Nevertheless, we cannot rule out the possibility that preference requires either a balance of light and dark or at least some pixels that appear “black.” Indeed, the notion that natural image preference is shaped by minimum pixel value is itself a novel and potentially important result. For example, it could call into question methods of mean normalization of images that do not account for minimum luminance. However, such proposals are not mutually exclusive: we may require *both* low skew and some black pixels in order to prefer a natural image. On the other hand, participants in Experiment 4 who reported a strategy of choosing the apparently brightest image preferred low skew images, which actually have the lowest minimum pixel value.

The present work does not settle these questions but it does show that fundamental regularities exist. In any case, we have shown that high skew images are not preferred, which is contrary to what would be expected if humans preferred more natural luminance distributions and/or distributions that imply the presence of shiny objects.

8.1.2. Dynamic range

Dynamic range limitations could have also played some role in our results. It is not possible to replicate the full natural scene luminance dynamic range with a standard monitor. Our display had a sequential dynamic range of 800:1 under darkened test viewing conditions, which means we achieved a comparable range to that in many natural scenes, but scattering lowers the effective dynamic range somewhat when stimuli are displayed. However, the (positive) correlation between the dynamic range of the presented stimuli and preference was modest: best-fit regressions (linear or quadratic) explained the following proportions (R^2) of range vs. preference data, pooled across images within a given experiment: Expt. 1a: 0.181; Expt. 1b: 0.285; Expt. 2: 0.092; Expt. 3: 0.552. Thus, compared to dynamic range, skewness is a substantially better predictor of preference.

8.1.3. Mean and variance

In past work, the mean and variance of image intensity were found to show relations with preference, at least for simple artificial stimuli (Reber, Winkelman, & Schwarz, 1998): in particular, higher is better. These relationships may be more complicated in real scenes (and there are likely to be interactions between these dimensions). For example, nighttime scenes – which are notoriously difficult to render with computer graphics – can be disliked if they are too bright (Ferwerda et al., 1996). In the present study, although we showed that lower mean stimuli were not appreciably different from higher mean stimuli in terms of patterns of preference, it remains possible that very low and very high mean intensity images do not follow the pattern observed in our current tests.

For example, for images with a low mean, lowering skew would create a very dark image that may not be preferred.

Pixel intensity variance (i.e., contrast) as well may show complex relations with preference. For example, white noise with uniformly-distributed luminances has higher contrast than Gaussian white noise with the same mean, but this doesn't necessarily mean it is preferred.

8.2. Alternative mechanisms

If it is the case, as we argue, that low skew images are generally preferred, there are still other possible causes for this pattern of preference that do not involve efficiency. For example, regularities in certain relevant classes of natural images such as faces could play a role. In particular, the smooth reflectance variations and complex albedo of faces could produce lower skew luminance distributions. If we tend to prefer faces in images, and if faces in general have low skew, then perhaps we also prefer images that have intensity distributions like those in face images. Alternatively, a preference for low skew images could have roots in environmental regularities under specific conditions (e.g., dawn, dusk, directional vs. diffuse illumination, etc.). We are investigating these possible influences.

In addition, there could be an influence of other higher-order statistical regularities that we did not control. However, human observers seem to be mostly insensitive to intensity distribution variations above the fourth order (Chubb, Econopouly, & Landy, 1994, 2004).

8.3. Glossy objects

Our results leave us an additional puzzle that turns the question about high skew and glossiness on its head: why do we like shiny things, which often have histograms with high positive skew – indeed, skewness is higher than is typical for scenes without shiny things – given that the visual system prefers low skew images? We speculate that different goals may be involved, for example low skew preferences may be related to natural vision while high skew preference could be related to foraging behavior, attention, or other factors. We have also not ruled out the possibility that aesthetic properties of isolated objects presented indoors, and complex outdoor scenes are fundamentally different. In outdoor scenes, we may prefer a harmony of light and dark, whereas for single objects indoors, we may prefer biologically relevant cues like specularities, which tend to generate higher skewness. Such a dichotomy is reflected in colloquial language: we speak of liking shiny things, but not shiny places. Moreover, in the outdoors, a glossy surface will directly reflect sunlight, causing potentially aversive glare.

8.4. Other media

It bears noting that Motoyoshi et al. (2007) used an artwork (a simulation of a Michelangelo sculpture) in their study of glossiness and found that high skew was correlated with perception of gloss. It may be the case that for sculpture and other 3D works, the ability to produce glossy objects trumps artists' propensity to produce low skew images. But in 2D media – and especially for representational art – low skew seems to be a priority, even though the introduction of more glossy oil-based pigments revolutionized painting in the Renaissance. We speculate that although such pigments are powerful tools for representational artists, they tend to be deployed in service of the represented tableau rather than for their own material properties.

Interestingly, James Cutting and colleagues (personal communication) have found that frames of Hollywood film from its

invention to today have nearly always had highly skewed intensity distributions (though it should be noted that the tested film database is not luminance-calibrated). This is despite the fact that traditional film projection would be subject to the compressive luminance nonlinearity typical of silver halide film. As Cutting, DeLong, and Nothelfer (2010) have argued previously, the structure of Hollywood film may be coupled to the demands of human attention, and high skew (along with 1/f-distributed cut lengths) could be a way to guide attention to the most salient objects in a given frame. Likewise, glossy single objects presented on dark backgrounds or indoors could be preferred in part because of their effect on attention.

8.5. Applications

In addition to its relevance to aesthetics, our study has implications for image processing: skewness reduction could be a useful goal for various imaging pipelines. Indeed, the principle of lower skew leading to preference is perhaps implicit in standard image processing techniques such as histogram normalization. More sophisticated image manipulations may likewise seek lower skew for aesthetic reasons. Anecdotally, when the auto-contrast and auto-tone functions of Adobe Photoshop CS5 are applied to our original, unmodified images, both manipulations produce images with skew closer to zero, suggesting that lowering skew might contribute to the “improvements” generated by this proprietary, blackbox transform. Transforms that scale image intensities to generate lower skew (and/or lower minimum intensity) may thus be globally advantageous.

9. Conclusion

We have shown that when the first and second order statistics of natural scene luminance distributions are normalized, third order statistics in these distributions show strong and consistent correlations with human preference. In particular, images with high skewness are reliably disliked. We suggest that this result is unexpected given the relationship between glossiness and skewness, but that our evidence is consistent with the notion that low skew images are more efficiently processed than high skew images, thus leading to preference. Finally, we argue that this notion can help explain the previous finding that artistic images throughout art history tend to have low skew in their luminance histograms.

Acknowledgment

This work was supported by a Hobart & William Smith Colleges Faculty Research Grant to D.G. and by a grant from WWTF CS11-023-LedRos to H. L.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.visres.2015.03.018>.

References

- Anderson, B. L., & Kim, J. (2009). Image statistics do not explain the perception of gloss and lightness. *Journal of Vision*, 9(11), 10.
- Arce-Lopera, C., Masuda, T., Kimura, A., Wada & Okajima, K. (2012). Luminance distribution modifies the perceived freshness of strawberries. *i-Perception*, 3(5), 338.
- Attewell, D., & Baddeley, R. J. (2007). The distribution of reflectances within the visual environment. *Vision Research*, 47(4), 548–554.

- Baccus, S. A. (2007). Timing and computation in inner retinal circuitry. *Annual Review of Physiology*, 69, 271–290.
- Brady, N., & Field, D. J. (2000). Local contrast in natural images: Normalisation and coding efficiency. *Perception*, 29, 1–15.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436.
- Chubb, C., Econopouly, J., & Landy, M. S. (1994). Histogram contrast analysis and the visual segregation of IID textures. *Journal of the Optical Society of America*, 11(9), 2350–2374.
- Chubb, C., Landy, M. S., & Econopouly, J. (2004). A visual mechanism tuned to black. *Vision Research*, 44(27), 3223–3232.
- Cutting, J. E., DeLong, J. E., & Nothelfer, C. E. (2010). Attention and the evolution of Hollywood film. *Psychological Science*, 21(3), 432–439.
- Dror, R. O., Leung, T. K., Adelson, E. H., Willsky, A. S. (2001). Statistics of real-world illumination. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. In: Proceedings of the 2001 IEEE Computer Society Conference on (Vol. 2, pp. II-164)*. IEEE.
- Fang, F., Kersten, D., & Murray, S. O. (2008). Perceptual grouping and inverse fMRI activity patterns in human visual cortex. *Journal of Vision*, 8(7), 2.
- Ferwerda, J. A., Pattanaik, S. N., Shirley, P., Greenberg, D. P. (1996, August). A model of visual adaptation for realistic image synthesis. In: *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques (pp. 249–258)*. ACM.
- Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., et al. (1999). An anchoring theory of lightness perception. *Psychological Review*, 106(4), 795–834.
- Gosselin, F., & Schyns, P. G. (2003). Superstitious perceptions reveal properties of internal representations. *Psychological Science*, 14(5), 505–509.
- Graham, D. J. (2009). Art statistics and visual processing: Insights for Picture Coding. In: *Proceedings of the Picture Coding Symposium 2009*. Chicago, IL.
- Graham, D. J. (2011). Visual perception: Lightness in a high dynamic range world. *Current Biology*, 21(22), R914–R916.
- Graham, D. J., & Field, D. J. (2007). Statistical regularities of art images and natural scenes: Spectra, sparseness and nonlinearities. *Spatial Vision*, 21, 149–164.
- Graham, D. J., & Field, D. J. (2008a). Variations in intensity statistics for representational and abstract art, and for art from the eastern and western hemispheres. *Perception*, 37, 1341–1352.
- Graham, D. J., Field, D. J. (2008b). Global nonlinear luminance compression in painted art. *Proc. SPIE: Computer Image Analysis in the Study of Art 6810*, 68100K.
- Graham, D. J., Friedenberg, J. D., McCandless, C. H., & Rockmore, D. N. (2010). Preference for artwork: Similarity, statistics, and selling price. *Human Vision and Electronic Imaging*, 7527, 75271A.
- Graham, D. J., & Meng, M. (2011). Altered spatial frequency content in paintings by artists with schizophrenia. *i-Perception*, 2(1), 1–9.
- Kleiner, M., Brainard, Pelli D. (2007). "What's new in Psychtoolbox-3?" *Perception 36 ECVF Abstract Supplement*.
- Laughlin, S. B. (1981). A simple coding procedure enhances a neuron's information capacity. *Zeitschrift für Naturforschung*, 36c, 910–912.
- Moon, P., & Spencer, D. E. (1944). Aesthetic measure applied to color harmony. *Journal Of The Optical Society Of America*, 34(4), 234–242.
- Motoyoshi, I. (2011). Ecological-optics origin of the style of European and East-Asian classical painting. *Journal of Vision*, 11(11), 1188.
- Motoyoshi, I., Nishida, S. Y., Sharan, L., & Adelson, E. H. (2007). Image statistics and the perception of surface qualities. *Nature*, 447(7141), 206–209.
- Olman, C., Boyaci, H., Fang, F., & Doerschner, K. (2008). V1 responses to different types of luminance histogram contrast. *Journal of Vision*, 8(6), 345.
- Olmos, A., & Kingdom, F. A. A. (2004). A biologically inspired algorithm for the recovery of shading and reflectance images. *Perception*, 33, 1463–1473.
- Palmer, S. E., & Schloss, K. B. (2010). An ecological valence theory of human color preference. *Proceedings of the National Academy of Sciences*, 107(19), 8877–8882.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437–442.
- Reber, R., Winkelman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science*, 9(1), 45–48.
- Richards, W. A., (1981). A lightness scale from image intensity distributions", *AI Memo No 648*, Artificial Intelligence Laboratory, MIT, 545 Technology Square, Cambridge, MA
- Redies, C. (2008). A universal model of esthetic perception based on the sensory coding of natural stimuli. *Spatial Vision*, 21(1), 97–117.
- Redies, C., Hänisch, J., Blickhan, M., & Denzler, J. (2007). Artists portray human faces with the Fourier statistics of complex natural scenes. *Network: Computation in Neural Systems*, 18(3), 235–248.
- Schweinhart, A. M., & Essock, E. A. (2013). Structural content in paintings: Artists over regularize oriented content of paintings relative to the typical natural scene bias. *Perception*, 42(12), 1311–1332.
- Wada, Y., Arce-Lopera, C., Masuda, T., Kimura, A., Dan, I., Goto, S. I., et al. (2010). Influence of luminance distribution on the appetizingly fresh appearance of cabbage. *Appetite*, 54(2), 363–368.
- Willenbockel, V., Sadr, J. D., Fiset, D., Horne, G. O., Gosselin, F., & Tanaka, J. W. (2010). Controlling low-level image properties: The SHINE toolbox. *Behavior Research Methods*, 42(3), 671–684.
- Yang, J., Otsuka, Y., Kanazawa, S., Yamaguchi, M. K., & Motoyoshi, I. (2011). Perception of surface glossiness by infants aged 5 to 8 months. *Perception*, 40(12), 1491–1502.