Global nonlinear compression of natural luminances in painted art

Daniel J. Graham[*a] and David J. Field[a]
[a]Department of Psychology, Uris Hall, Cornell University, Ithaca, NY USA 14853;

ABSTRACT

Over the past quarter century, measures of statistical regularities of natural scenes have emerged as important tools in explaining the coding properties of the mammalian visual system. Such measures have recently been extended to the study of art. Our own work has shown that a log nonlinearity is a reasonable first approximation of the type of luminance compression that artists perform when they create images. But how does this nonlinearity compare to those that artists actually use? In this paper, we propose a model of the global luminance compression strategy used by one artist. We also compare the curves required to transform natural scenes so that the scene luminance histograms match the histograms of a number of collections of art, and we test the response of observers to those scenes. The collections include a group of Hudson River School paintings; a group of works deemed to be "abstract" works in a forced-choice paradigm; collections of paintings from the Eastern and Western hemispheres; and other classes. If a single transform were sufficient to compress images in the way artists do, we would expect these transforms all to be log-like and on average, there should be little or no difference in observer preference for the collection of natural scenes when they are compressed according to these transforms. We find instead that these groupings of art have distinct transforms and that Western observers prefer many of these transforms over a log transform. Together these findings offer evidence that a painter's global luminance compression strategy—or "artist's look-up table"—may be a fundamental property of a given painter or grouping of paintings, though further study is needed to establish what factors determine the shape of this transform. We discuss a number of possible factors.

Keywords: natural scenes, visual system, efficient coding, artist's look-up table, art authentication, histogram matching, stylometry, retina, nonlinearities, luminance compression.

1. INTRODUCTION

Recent evidence[1,2,3] suggests that artworks share a number of statistical regularities with natural scenes. For example, power spectra of artworks are similar to those of natural scenes across art history and artistic genres. In the view of Graham and Field[1], the ubiquity of this statistical regularity is most simply explained as a form of efficiency relative to coding strategies employed at early stages of visual system processing, rather than as a feature of innate human aesthetics (e.g., Redies[4]). The necessity of nonlinear luminance compression in the creation of art was also noted in this earlier study.[1] In particular, a log nonlinearity was shown to be useful for approximating the way in which artist's solve the "luminance problem," i.e., how they compress the high-dynamic range scenes they depict into the low dynamic range available in paint.

Here we propose that artist's nonlinear luminance compression strategies may themselves constitute an effective solution to the luminance problem. If the goal of an artist were simply to compress scene luminances such that they are reproducible in paint, all artists might do well to apply a log-like nonlinearity. This nonlinearity has been suggested as an efficient solution for compressing scene luminances into neural responses: Cone photoreceptors in vertebrates all show roughly log-like transforms, indicating that this is a common approach to solving the luminance problem faced by most visual systems. As McCann argues[5], the visual system is able to process natural scenes that together span a 10^{10}:1 luminance range though we note that a typical scene (and most scenes depicted in paintings) will show a far smaller dynamic range. McCann estimates that photoreceptors have a 10^8:1 dynamic range and that ganglion cells can only fire at rates in a roughly 100:1 range. A log transform model of vertebrate cone photoreceptors (when adjusted for adaptation) has been argued to be efficient with respect to regularities in scene luminances and contrasts[6,7] and because it

* djg45@cornell.edu; phone: +1 603 397 3587; fax: +1 607 255 8433; http://people.psych.cornell.edu/~djg45/dan/dan.html

turns intensity differences into ratios.[8] A log nonlinearity is also a good first approximation of the effect of the "artist's gamma" (more accurately called the artist's look-up table) on natural scene skewness.[1] This previous study suggested that a linear scaling is not sufficient to bring about the required compression, but a log function is sufficient (see Figure 1). Given that retinal coding strategies are remarkably consistent across vertebrate taxa,[9] one might expect artist's to all employ some form of log nonlinearity, or a similar function, in their paintings. However, we must be careful not to fall for a variety of the El Greco fallacy[10] described in Graham and Field.[1] That is, an artist must view his own paintings through the cones' compressive nonlinearity, meaning that the image will pass through this nonlinearity twice. Log-like luminance compression in paintings may be useful for many of the same reasons that a log-like nonlinearity is useful at the photoreceptor level but it may not be a result of photoreceptor luminance compression.

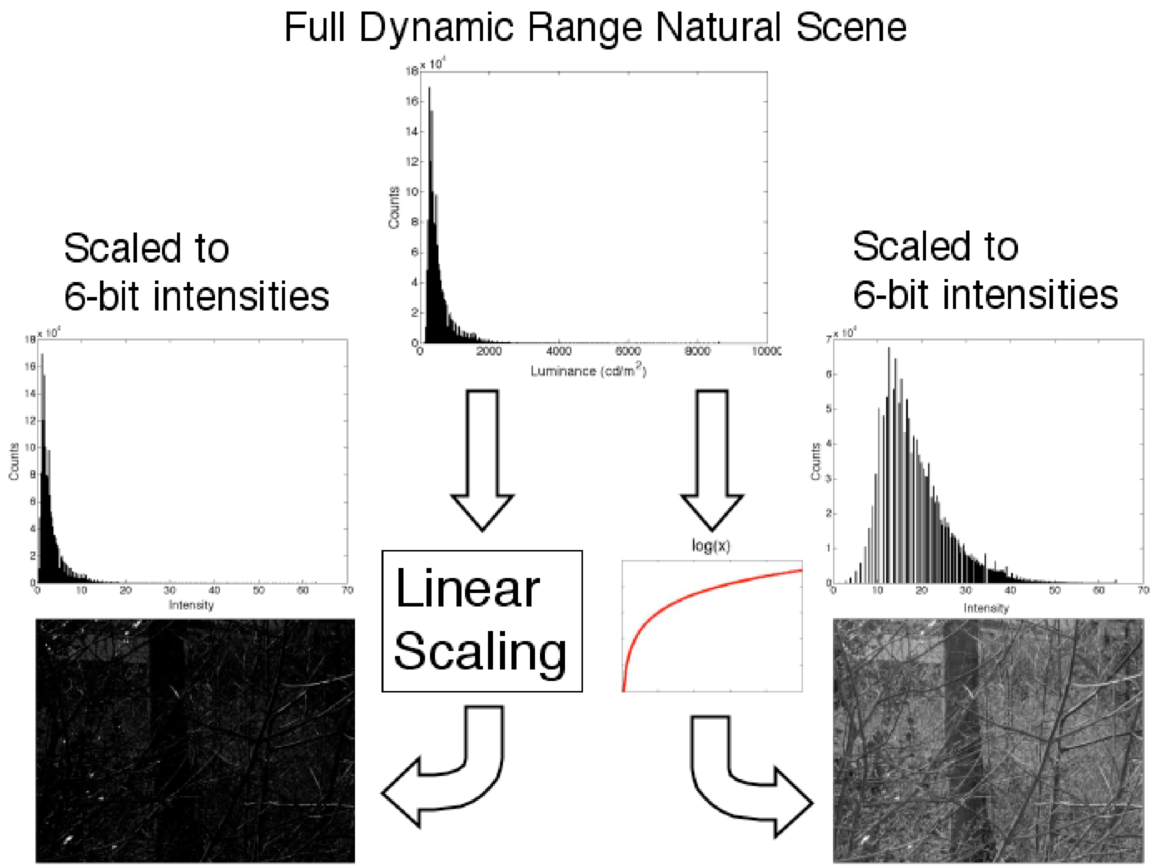

Figure 1.This figure illustrates the fact that linear scaling of natural scene luminances produces a very dark scene with a few highlights, but log luminance compression can generate an acceptable scene. The luminance histogram of the original scene (Van Hateren #619)[12] is shown at the top. The luminances are linearly scaled to 6 bits of intensity in order to produce the image and histogram on the left, while the luminances are scaled with a log nonlinearity, then stretched to 6-bits to generate the image and histogram on the right.

Clearly the goal of the artist is not simply to compress the dynamic range of luminances. In this study, we performed an experiment to test whether human artists use a consistent set of strategies to compress scene luminances, and we compared these compressive transforms to a log in terms of observer preference. If a single transform were sufficient to compress images in the way artists do, we would expect these transforms all to be log-like and on average, there should be little or no difference in observer preference for the collection of natural scenes when they are compressed according to these transforms. Though a log nonlinearity would appear to be the most efficient method of compression, we show that it cannot be used to model the variability in the artist's look-up table. We propose that like the regularity of power spectra for art, our simple model of the artist's look-up table describes a minimal constraint on art making, but a fundamental one. Both of these statistical constraints have underlain the artist's task since the inception of art, though most artists were and remain unaware of them. We speculate that the type of scene depicted by the artist or perhaps the manner in which the artist learned to paint may dictate the use of a distinct compression strategy and we discuss other factors that may affect the shape of this transform. The fact that photographers have long known that it is necessary to *locally* adjust intensities in an image (see Discussion) suggests that one could even assume that there is no single strategy common to all artists. However, to our knowledge this question has not been directly addressed in the literature using calibrated natural scenes so it warrants the more rigorous attention it is given here.

In the first study, we provide a model of the artist's look-up table for an accomplished landscape painter. In the second study, we present evidence that observers show consistent preferences for compressive transforms of natural scenes images that match scene histograms to the mean histogram of various collections of art, and observers prefer many of these transforms over a log transform.

2. STUDY 1: MODEL OF THE ARTIST'S LOOK-UP TABLE

An example of how the artist's look-up table can be modeled is presented. There remain technical limitations of the current approach with respect to very high dynamic range scenes, like landscapes containing significant amounts of sky. This means that measurements of the artist's look-up table are currently most profitable for scenes with dynamic ranges that are significantly smaller than that for typical landscapes (but which are still much larger than the range available in paint). We chose therefore to analyze a painting by a professional local painter with many years of experience, Neil Berger (http://neilberger.com). The painting depicts a scene of the bottom of a waterfall (containing no sky), which it was possible to recreate in a photograph.

The task of finding the artist's look-up table involved mapping luminances from the scene that inspired a painting onto the luminances of the resulting painting. This measurement is most straightforward for outdoor scenes. Of course, few paintings are "literal" reproductions of the spatial layout of the scene, and as we have argued linear reproductions of natural scene luminances are impossible to make using paint. The artist whom we have studied (Neil Berger) generally makes quick sketches of the outdoor scene he aims to capture but does the bulk of the painting in his studio. As can be seen in Figs. 2 and 3, however, this method leads to an image that is quite similar to the scene that inspired it.

The pixels in images of the scene and the painting were photometrically calibrated. To do this, a Macbeth Color Checker (Gretag-Macbeth LLC, New Windsor, NY USA) was placed in a stationary location within the scene and its intensities were recorded as raw pixel values from the camera's CCD chip, and as luminance values measured with a Minolta LS-100 photometer (Konica Minolta Inc., Tokyo Japan). Light reflected from the color checker's six greyscale reflectance patches was measured for each patch. The relationship between reflected luminances and their corresponding pixel values can then be established and it can be extrapolated to other values by fitting the relation with a polynomial or exponential function. However, this procedure will not generally allow calibration for typical high dynamic range scenes. Debevec and Malik[11] offer a discussion of this problem and one solution, which is to take pixel intensity samples using different shutter speeds and the same aperture.

The painting titled "Taughannock Falls" by Neil Berger (2006) was selected for study for reasons described above. It was photographed in the artist's studio lit by a bank of north-facing windows at 14:00 on a partly cloudy day in summer. It was imaged along with a Macbeth Color Checker using a tripod-mounted Canon PowerShot S60 (Canon Inc., Tokyo Japan) digital camera. The full image was 2592 x 1944 pixels and cropped to 916 x 860 (exposed for 0.1 sec at f/5.8, focal length 8.6 mm) and recorded at 16-bits per channel in raw format. It was converted from RGB coordinates to YIQ coordinates and the resulting map of greyscale intensities (Y) was retained. When fit with a polynomial, the relationship of pixel value (x) and luminance (L) was found to goes as $L = 0.0035x - 0.15$ in units of candelas/(meter)^2, with $R^2 = 0.99$ for the fit. The image is thus converted to luminance values. See Figure 2 for the luminance map and a histogram of the luminance values. The painting showed a dynamic range of 20.9:1.

The same procedure was then performed with the scene that the painting depicts, which is located in Taughannock State Park, Trumansburg, NY USA. The scene was photographed using the same apparatus as that used above on an overcast day at 15:00 in autumn, a day much like that depicted. A second image was acquired using the same exposure and focal length (shutter speed: 0.025 sec, aperture: f/6.3, focal length: 20.7 mm) with the Macbeth chart placed 1.5 m from the camera, which was used for calibration. The luminance map and luminance histogram are given in Figure 3 below. Using this calibration, the scene is found to have a dynamic range of luminances of 41.3:1, which is roughly twice that for the painting. These data were fit according to $L = 0.0038x^{1.11}$ where x is pixel value and L is luminance in units of cd/m^2, with $R^2 = 0.99$ for the fit. Note that the camera saturated on the fifth most reflective chip (reflectance = 0.64), meaning that the pure white chip (reflectance = 0.95) had the same pixel value. However, the scene itself did not have surfaces that reflected light at this intensity. We therefore calibrated the scene luminances based on the first five data points, which gave good agreement with spot photometric measurements taken at the base of the falls and on the surface of the pool.

The method for calculating the look-up table for this image pair involves transforming the histogram of the scene to match that of the painting using histogram matching (also called histogram specification).[13] Given the scene luminance map A, we can minimize for a transform T using the equation $|c_1(T(k)) - c_0(k)|$, where c_0 is the cumulative histogram of A, c_1 is the cumulative sum of a specified histogram (i.e., of the corresponding painting) for all intensities k. This minimization is subject to the following constraints: (1.) T must be monotonic and (2) $c_1(T(a))$ cannot overshoot $c_0(a)$ by more than half the distance between the histogram counts at a. This transformation is used to map the gray levels in the scene A to their new values in the image B, with intensities $b = T(a)$. After performing histogram matching on the scene (using the luminance histogram of the painting), the "artified" scene appeared recognizable as a waterfall scene (Figure 4). We stress that the images shown here are luminance maps so in order to reproduce them on paper or on a low dynamic range display, linear intensity scaling was applied to each image. The point of this demonstration was not to make the "best" image or tone mapping, as is the case for most computer graphics applications. Rather it was to show that a greatly reduced model of the painter's strategy can by itself achieve the nonlinear luminance compression, and that this transform is unlike a simple log function. The transform for the Taughannock Falls scene is shown in Figure 5.

A. B.

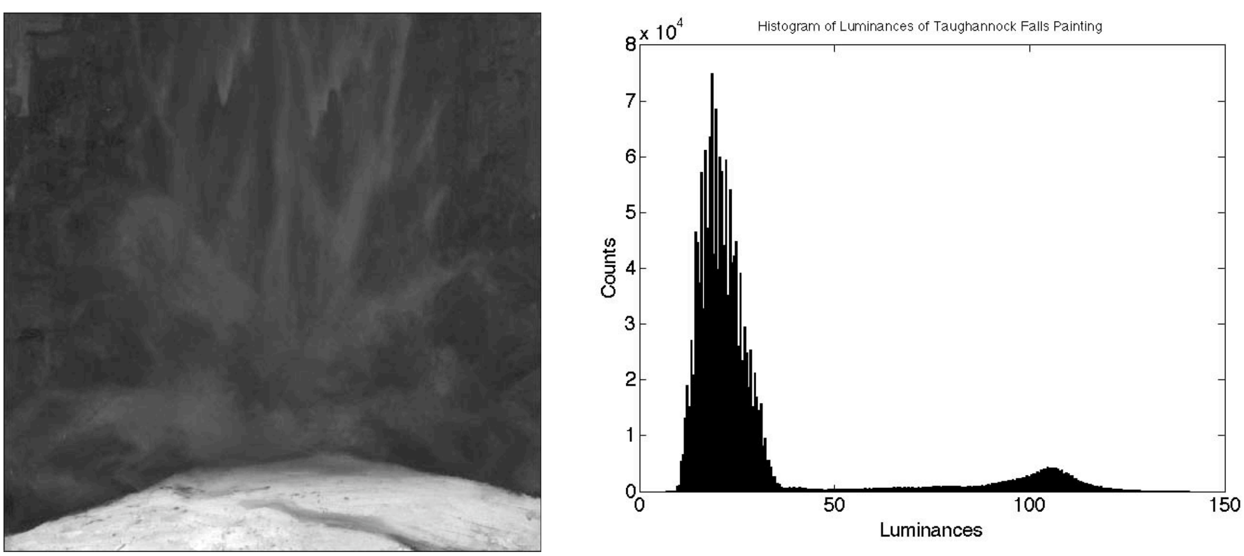

Figure 2. Luminance calibrated image of painting "Taughannock Falls" by Neil Berger (A.; linearly scaled for display) and luminance histogram of painting (B.). The painting shows a dynamic range of 20.94:1.

A. B.

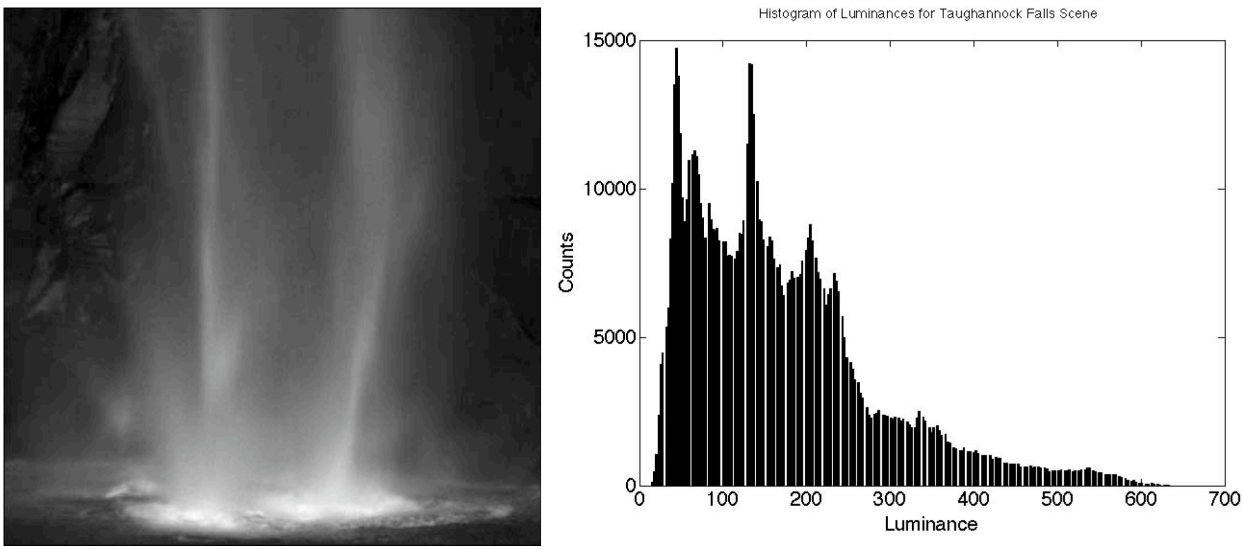

Figure 3. Luminance calibrated scene (A.; scaled linearly for display) and luminance histogram of scene (B.). The scene shows a dynamic range of 41.25:1.

A. B.

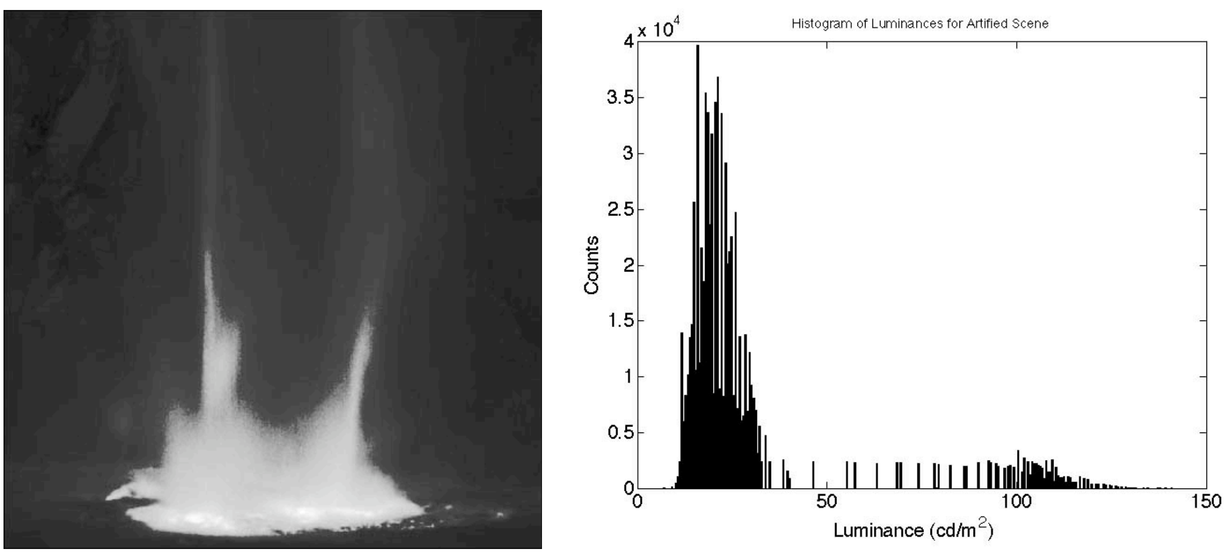

Figure 4. “Artified” scene (A.; scaled linearly for display) and luminance histogram of scene (B.).

We speculate that another artist painting the same scene could produce a distinctly different image by simply changing the number of inflection points in the transform. Note that unlike a log function, whose second derivative does not change sign, the second derivative of the transform for “Taughannock Falls,” T, changes sign three times. Most simple functional compressive nonlinearities have the same limitation as does the log. Clearly, this artist does not employ a single, simple functional look-up table for all his paintings.

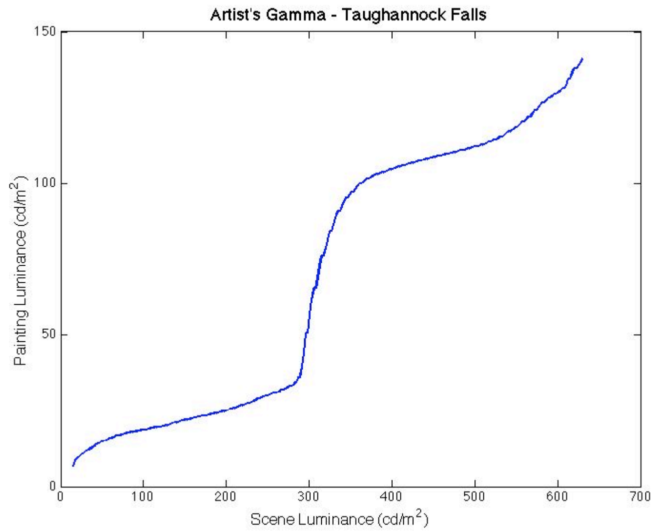

Figure 5. Luminance transform *T* mapping Taughannock Falls scene luminance histogram into "Taughannock Falls" painting luminances.

Of course, the artist's look-up table model has many limitations, not the least of which is the obvious quantization in the mid-range tones. As described in the discussion, the transformed scene has been boosted in the mid-range intensities because of the nature of the histogram matching technique.

In the next study, we performed an experiment to compare different mappings of natural scene images such that their histograms matched those of a number of groupings of artworks from a major university collection. We also tested observer preference for these transformed scene images, and we compared these to scenes that had been compressed using a log luminance nonlinearity. The goal of this study was to test the notion that observers prefer natural scene images to have a particular histogram, and to ask whether a simple log transform of the scene is preferred over histogram-specified scenes.

3. STUDY 2: PREFERENCE FOR ARTIST'S TRANSFORMS

Since we do not have access to the scenes that inspired most paintings held by major art museums, we cannot at present directly calculate an artist's look-up table for most artists. However, we can address the question of what type of transform most people prefer. Here we tested whether observers prefer images with a certain histogram, regardless of the image's subject matter. We calculated compressive nonlinearities that transformed a given natural scene image so that its histogram matched the mean histogram of a variety of collections of art from a major university museum. Then we asked observers to rank each of these scenes, along with a log-transformed version of the scene, in order of preference. The natural scenes that were used in this study were 31 scenes from the van Hateren database[12]. The transform for each image and class of artworks was again calculated using histogram matching as described in Study 1. The sets of artworks whose mean histograms we measured were:

- a collection of Hudson River School painters (10 images)
- a collection of Eastern art from the Herbert F. Johnson Museum of Art, Cornell University (72 images)
- a collection of Western art from the Herbert F. Johnson Museum (68 images)
- the set of images chosen as "abstract" works in a previous study[2] (12 images)
- the entire collection of Eastern and Western works, not including Hudson River School images (140 images).

Three Western observers (1F), who were naïve to the purpose of the experiment, were asked to rank the six transforms of each image in order of preference. The 31 sets of 6 images were shown in three rows and two columns on an Apple 30-inch Cinema Display (Apple Inc., Cupertino, CA USA) at approximately 0.4 m from the observer. The placement of the images was randomized for each scene. See Figure 6 for an example.

For our observers, we found that scene images transformed to match the mean histograms of the entire collection (all art), the set of Hudson River School paintings, and the set of Western painting were preferred over the log-transformed images, and over those transformed to match the mean histogram of the abstract works and of the Eastern works. Table 1 shows the total number of top 3 votes and bottom 3 votes for each transform, and their difference, as well as corresponding percentages. Figure 7 displays a histogram of the rankings by transform.

First, we note that there is no single histogram transform that can be used to predict all viewer preference data measured here. However, there remains a consistent preference for certain transforms. Observers prefer scenes that have been transformed so that their histograms match the mean histogram of a number of the collections of art considered here. In other words, viewers are not indifferent to the range of transforms that allow a given image to be matched to the histogram of a variety of artworks. Moreover, many of these transforms produce images that are preferred over images compressed with the more parsimonious strategy, i.e., a log nonlinearity. The flexibility afforded artists in this respect allows for the production of distinctive images drawn from a diversity of scenes, which could help explain the preferences we observe. It is unclear at present why the transforms for Eastern and abstract art were not preferred over the log transform.

Of course, our method for calculating these transforms is not strictly a model of the luminance compression performed by painters (see Discussion). Moreover, definitive tests of the relationship between distinguishability and liking for natural scenes have yet to be performed. However, we conclude that no single solution can characterize the strategies taken by broad groupings of artists. There may thus be good reason in future to measure consistency among and across artist's look-up tables, since these transforms may be shaped by a number of factors. In the following section, we discuss some of these factors.

Transform	**Rankings in Top Half**		**Rankings in Bottom Half**		**Difference**
	Number	Percent of total	Number	Percent of total	Number
All Art	76	82%	17	18%	59
Hudson River School	67	72%	26	28%	41
Western Art	41	44%	52	56%	-11
Log Transform	36	61%	57	39%	-21
Eastern Art	31	33%	62	67%	-31
Abstract Art	28	30%	65	70%	-37

Table 1. Preference data for natural scenes transformed so that each image histogram matched the mean histogram for each grouping of art, and also transformed with a log function. From left, the table lists the number of images from each class ranked in the top half (number of top 3 votes) and in the bottom half (bottom 3 rankings) in terms of preference, summed over observers. There were a total of 93 rankings (3 observers x 31 images) for each transform. Corresponding percentages are also shown. The rightmost column shows the difference between top half rankings and bottom half rankings.

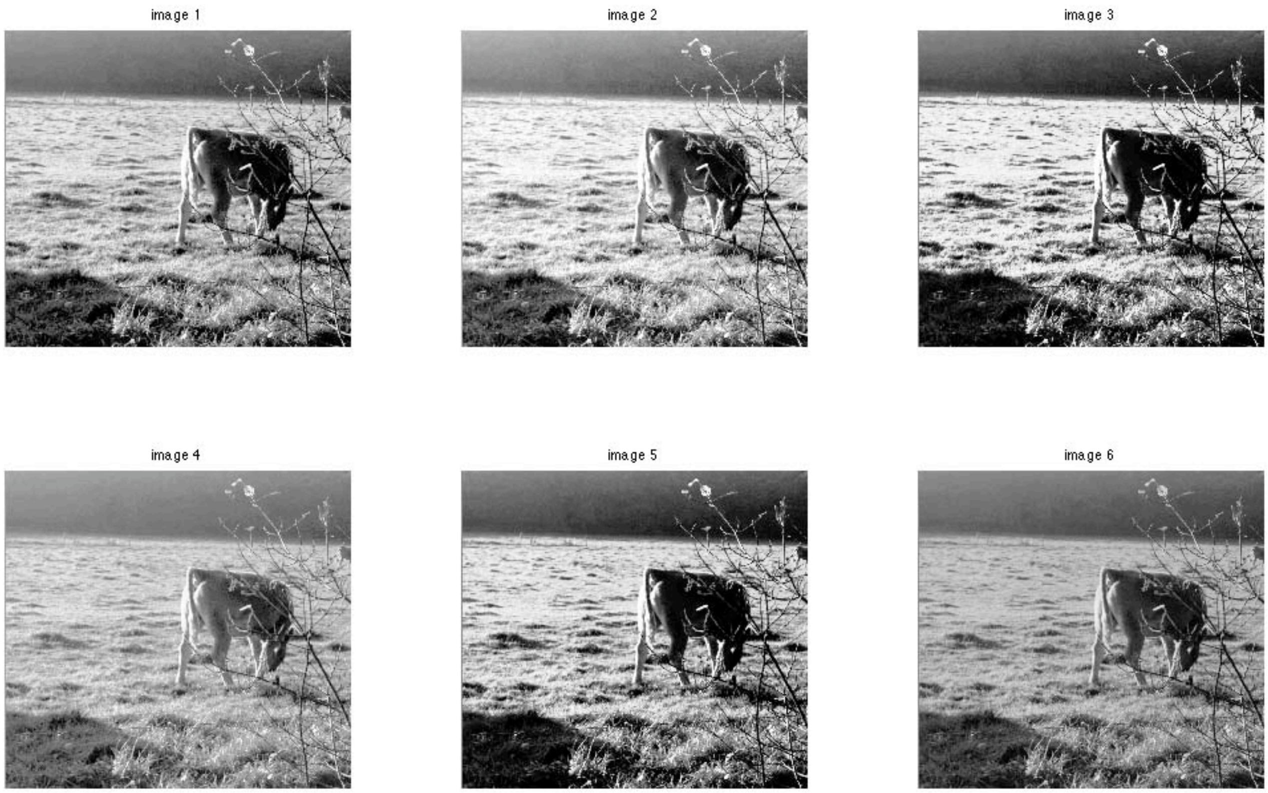

Figure 6. Example of display viewed by observers for scene number 9 in the test set transformed according to the following histograms (left to right from top): all art, log, abstract art, Eastern art, Western art, Hudson River School. Placement of transforms was randomized for each scene.

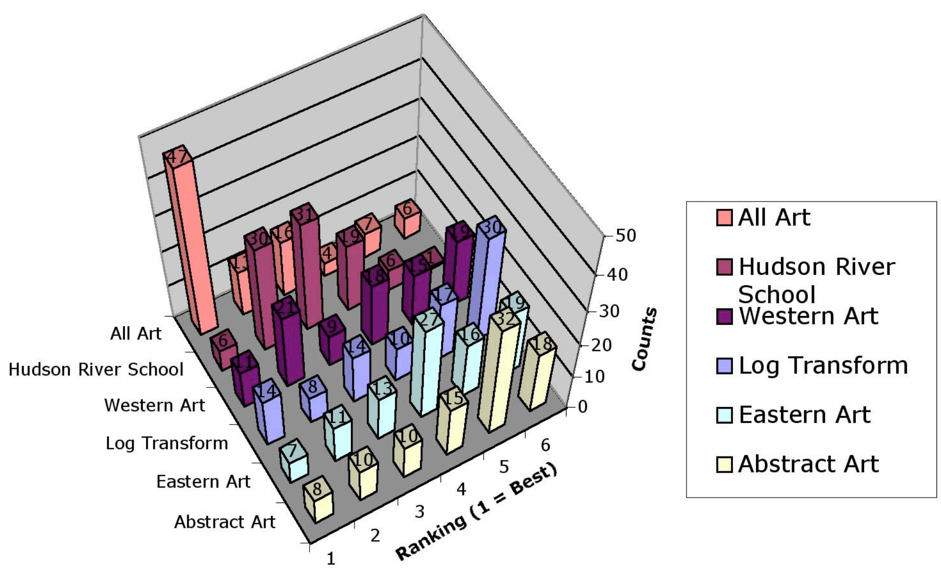

Figure 7. Histogram of preference rankings by transform summed over 3 observers. There were 93 total rankings for each transform class (3 observers x 31 images), and each ranking was used a total of 93 times. Note that the transform classes All Art, Hudson River School and Western Art were preferred over the log transform

4. DISCUSSION

Like the regularity in power spectra in art shown in previous papers[1,2,3], the artist's look-up table is a minimum necessary constraint for nearly all artists. Our demonstration here of a method for approximating the artist's look-up table does not prove that the artist we examined produced the most efficient compression, nor that there do not exist a host of equally compelling images that could be generated for this scene with different transforms. But together with the results regarding observer preference for histogram-specified scenes, our work indicates that no single function or fit parameter is able to account for artist's nonlinear luminance compression strategies. Different groupings of artworks produce different transforms, many of which are preferred by observers over a log transform. Our model of the artist's look-up table therefore captures a fundamental aspect of art making, one that may be characteristic to a given artist, artistic movement, time period or scene type (i.e., landscapes, portraits, etc.). We are currently exploring this range of possible factors.

Why might we expect artists to have discovered effective nonlinear transforms to solve the luminance problem? We must add a number of caveats to this notion. Our hypothesis is most easily applied to landscape paintings since it is often possible to find a rough input-output relation in terms of luminance for paintings of specific landscapes. But landscapes as such became common only in recent times. Portraying the natural environment for its inherent beauty using paint is a relatively new innovation in the history of Western art[14], though we would argue that this is for cultural, not technological, reasons.[§] On the other hand, all painters face some form of the luminance problem.

In addition, the notion of luminance compression assumes that the artist has the ability and the intention to faithfully map luminances in the world into reflectances of pigment in the same relative spatial arrangement. This is not always the artist's goal. Luminance compression may be facilitated by tricks such as a knowledge of atmospheric perspective and the use of grids and camera obscurae or other optical projections, but it does not require such short-cuts. It may even be possible to judge whether a specific short cut was used by an artist by calculating a collection of artist's look-up table functions under experimental conditions.

There are other issues with the idea of an artist's look-up table, which relate to the limitations of the technique of histogram matching. Because histogram matching is imperfect by necessity[13] it is impossible to fully model luminance compression using this technique. Histogram matching essentially interpolates between the cumulative distribution functions of the input and the output, and the resulting transform (which is applied to the input) is thus required to be monotonically increasing. As can be seen in the example given above, the bimodal painting luminance distribution (Figure 2) is approximated as a unimodal one in the "artified" scene (Figure 4). Histogram-based methods of this sort (e.g., histogram equalization) all suffer from these limitations when they correspond to images because one cannot choose some fraction of pixels at a given intensity to transform while leaving the rest at that intensity unchanged.

Artists—particularly photographers—have been aware of the limitations of global luminance compression for some time. Ansel Adams' "zone system" of camera exposure is in a sense a practical model for solving the luminance problem using *local* adjustments to exposure. The zone system is designed to make use of the ability of the photographer to locally over- or underexpose regions of an image in the print-making process.[16] The general rule of this system is to "expose for the shadows; develop for the highlights," thereby maximizing the effective dynamic range of the output. That is, the response properties of camera film are such that the global effect of an overexposed negative (which captures a larger dynamic range in the shadows) can be compensated by locally adjusting the exposure of the print (processes referred to as "burning and dodging").[†] Similar techniques can be applied digitally. The development of the zone system for the purpose of capturing diverse outdoor scenes suggests that a flexible strategy based on the spatial distribution of scene luminances can be effective for diverse collections of scenes.

McCann[17] has argued that until the Renaissance, all human art was incapable of depicting high-dynamic range scenes with any degree of fidelity. As we mentioned, Western landscape paintings as we know them today were unknown in the Pre-Renaissance and rare indeed until the 19th century, so evidence for McCann's argument is typically drawn from candle-lit scenes. Leonardo da Vinci is thought to be the inventor of the technique of *chiaroscuro* (literally, "light-dark")

[§] Indeed, Korean painters have been producing such images for over a thousand years.[15]

[†] This can be aided by changing the amount of time a negative spends in chemical developer, and by using toner chemicals during development.

but Rembrandt, Caravaggio and others perfected it in order to portray high dynamic range indoor scenes. This is another example of a trick that artists can use to nonlinearly compress scene luminances. In the case of candle-lit scenes, chiaroscuro is a technique that may succeed in part because of the relative ease with which the artist can model the light produced by a single candle. Compared to a candle-lit scene, a typical outdoor or window-lit scene—with its multiple sources of direct, diffuse and specular illumination—is a priori more difficult to recreate.

Painters are given more flexibility than photographers in that every area of a painting can be locally "burned and dodged," so to speak. Indeed, local adjustments to luminance and contrast play a large role in the visual effect of a painting. The artist's look-up table notion is a significantly simplified model which does not capture the rich landscape of local adjustments. Nor does it account in its present guise for the use of color, which is a major concern of painters.

5. CONCLUSION

Although we have shown that viewers have consistent preferences for images whose histograms match the histograms of various collections of art, our results suggest that artists do not have a single, universal strategy for compressing real-world luminances into painting luminances. And despite the limitations of the artist's look-up table model presented here, the evidence gathered suggests that the artist's look-up table is a useful approximation of an effective approach used by artists to solve the luminance problem. It is not a full model of the painting process, nor of luminance compression in art in particular, but it is a simple description of an artist's unique strategy for compressing the dynamic range of luminances. As such it may be a defining characteristic of an artist and it could potentially be used to determine the "stylometry" for a given artist. This can be done if (1) an artist is consistent in her look-up table over time, or has predictable changes over time; and (2) a reasonably representative, calibrated image of the scene depicted in a painting can be generated. Both of these conditions require much further study. Lacking further evidence, we nonetheless speculate that reproducing the depicted scene may not be necessary if the artist's look-up table is found to be of a *fundamental* nature with respect to the artist's visual system. That is, an artist's typical look-up table may be shaped simply by the ambient light in the area where the artist learned to paint or was raised. A more extreme suggestion comes from Charles Falco[18] who suggests that optical projection devices developed in the Renaissance produced images of a scene that were so much like what we think of as "paintings" that the mere exposure of an artist to that image would have changed all of the artist's subsequent paintings. In other words, merely witnessing a projection would, in Falco's view, allow painters to more accurately depict many aspects of natural scenes, including their large dynamic range. While we do not endorse this view, it does show that the luminance problem has been attacked in various ways for hundreds of years.[17] Moreover, a full model of artists' luminance compression strategies could be a useful test of Falco's hypothesis.[‡] It is also possible that the shape of the artist's transform may be related to the number and intensity of illuminants in the depicted scene, or to the transforms typical of a painter's contemporaries.

Our statistical measures, as well as those demonstrated by other groups[19], can also be useful for digital art authentication. The refinement of such techniques is ongoing. The ability to establish such statistical signatures based on non-invasive photometric measurements could prove extremely useful to museums, collectors and critics. Such has been done with some success for literature.[20]

6. REFERENCES

1 D. J. Graham and D. J. Field, "Statistical regularities of art images and natural scenes: Spectra, sparseness and nonlinearities," Spatial Vision **21(1-2)**, 149-164 (2007).

[‡] However, it should be noted that many projections—when viewed in a dark room—will show the same large dynamic range as the scene itself, which makes it impossible to "copy" the scene using paint. Falco[17] suggests, on the other hand, that artists would only have used the projection to sketch where the defining shadows are in a scene and, as with landscape painters like Neil Berger, the bulk of the composition would be done later in the studio.

[2] C. Redies, J. Hasenstein, and J. Denzler, "Fractal-like image statistics in visual art: similarity to natural scenes. Spatial Vision, **21(1-2)**, 137-148 (2007).
[3] D. J. Graham and D. J. Field, Variations in intensity statistics for representational and abstract art, and for art from the eastern and western hemispheres," Perception, in press.
[4] C. Redies, "A universal model of esthetic perception based on the sparse and efficient coding of natural stimuli," Spatial Vision, **21(1-2)**, 97-117 (2007).
[5] J. McCann, "Rendering high-dynamic range images: algorithms that mimic human vision," Proc. AMOS Technical Conf., 19-28 (2005).
[6] N. Brady and D. J. Field, "Local contrast in natural images: normalisation and coding efficiency," Perception **29(9)**, 1041-1055 (2000).
[7] V. Mante, R. A. Frazor, V. Bonin, W. S, Geisler and M. Carandini, "Independence of luminance and contrast in natural scenes and in the early visual system," Nature Neurosci., **8(12),** 1690-1697 (2005).
[8] D. J. Field, "What is the goal of sensory coding?" Neural Comput., **6(4)**, 559-601 (1994).
[9] B. L. Finlay, L. C. de Lima Silveira and A. Reichenbach, "Comparative aspects of visual system development," in *The Structure, Function and Evolution of the Primate Visual System* (J. Kremers, ed.), John Wiley, New York, 2004.
[10] S. Anstis, "Was El Greco astigmatic?" Leonardo, **35(2)**, 208 (2002).
[11] P. E. Debevec and J. Malik, ''Recovering high dynamic range radiance maps from photographs,'' Proc. SIGGRAPH 97, 369-378 (1997).
[12] J. H. van Hateren and A. van der Schaaf, "Independent component filters of natural images compared with simple cells in primary visual cortex," Proc. Roy. Soc. B Biol. Sci., **265(1394)**, 359-366 (1998).
[13] R. C. Gonzalez, R. E. Woods and S. L. Eddins, *Digital Image Processing Using Matlab*, Prentice-Hall, Upper Saddle River, N.J., 2004.
[14] E. H. Gombrich, *Art and Illusion*, Princeton University, Princeton, 1961.
[15] Yi S.-M., *Korean Landscape Painting*, Hollym International, Elizabeth, N. J., 2006.
[16] A. Adams, *The Negative*, New York Graphic Society, Boston, 1948.
[17] J. J. McCann, "Art, science and appearance in HDR images," J. Soc. Inform. Display, **15(9)**, 709-719 (2007).
[18] C. Falco (private communication).
[19] S. Lyu, D. N. Rockmore and H. Farid, "A digital technique for art authentication," Proc. Nat. Acad. Sci. USA, **101(49)**, 17006-17010 (2004).
[20] F. Mosteller and D. L. Wallace, *Inference and Disputed Authorship: The Federalist*, Addison-Wesley, Reading, MA, 1964.

Acknowledgements: DJG was supported by NIH grant EY015393. DJF was supported by NGA contract HM 1582-05-C-0007. We wish to thank Neil Berger for his generous help with this project, and the Herbert F. Johnson Museum, Cornell University.