Visual wetness perception based on image color statistics

Masataka Sawayama

NTT Communication Science Laboratories, Nippon Telegraph and Telephone Corporation, Kanagawa, Japan

Computer Science and Artificial Intelligence Laboratory (CSAIL) and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA

Telegraph and Telephone Corporation, Kanagawa, Japan

NTT Communication Science Laboratories, Nippon

Edward H. Adelson

Shin'ya Nishida

Color vision provides humans and animals with the abilities to discriminate colors based on the wavelength identity of objects of interest in cluttered scenes (e.g., ripe fruit among foliage). However, we argue that color vision can inform us about much more than color alone. Since a trichromatic image carries more information about the optical properties of a scene than a complex material qualities. Here we show that human vision uses color statistics of an image for the perception of an ecologically important surface condition (i.e., wetness). Psychophysical experiments showed that overall enhancement of chromatic saturation, combined with a luminance tone change that increases the darkness and glossiness of the image, tended to make dry scenes look wetter. Theoretical analysis along with image analysis of real objects indicated that our image transformation, which we call the wetness enhancing transformation, is consistent with actual optical changes produced by surface wetting. Furthermore, we found that the wetness enhancing transformation operator was more effective for the images with many colors 夭¹ *¹ ¹ <i>¹ ¹ ¹ ¹ ¹ ¹ <i>¹ ¹ ¹ ¹ ¹ ¹ ¹ <i>¹ ¹ ¹ ¹ <i>¹ ¹ ¹ <i>¹ ¹ ¹ <i>¹ ¹ ¹ ¹ ¹ ¹ <i>¹ ¹ ¹ ¹ ¹ ¹ ¹ ¹ <i>¹ ¹ ¹ ¹ <i>¹ ¹ ¹ <i>¹ ¹ <i>¹ ¹ ¹ <i>¹ ¹ ¹ ¹ <i>¹ ¹ ¹ ¹ ¹ ¹ <i>¹ ¹ ¹ <i>¹ ¹ <i>¹ ¹ <i>¹ ¹ ¹ <i>¹ ¹ ¹ <i>¹ ¹ ¹ ¹ <i>¹ ¹ ¹ <i>¹ <i>¹ ¹ <i>¹ <i>¹ ¹ <i>¹ ¹ <i>¹ ¹ <i>¹ ¹ <i>¹ <i>¹ ¹ <i>¹ ¹ <i>¹ ¹ <i>¹ ¹ <i>¹ ¹ <i>¹ <i>¹ ¹ <i>¹ ¹ <i>¹ ¹ <i>¹ ¹ <i>¹ <i>¹ ¹ <i>¹ <i>¹ ¹ <i>¹ ¹ <i>¹ ¹ <i>¹ <i>¹ <i>¹ ¹ <i>¹ <i>¹ ¹ <i>¹ <i>¹ <i>¹ ¹ <i>¹ <i>¹ ¹ <i>¹ <i>^{1*} hue entropy). The hue entropy may be used to separate surface wetness from other surface states having similar optical properties. While surface wetness and surface color might seem to be independent, there are higher order color statistics that can influence wetness judgments, in accord with the ecological statistics. The present findings indicate that the visual system uses color image statistics in an elegant way to help estimate the complex physical status of a scene.

Introduction

The ability to discriminate and identify color is fundamental to human vision. The evolutionary importance of color is often stated in terms of color identification, such as the ability to judge ripeness of fruit (Mollon, 1989; Steward & Cole, 1989). The issues surrounding color discrimination and identification, including metamerism and color constancy, have been central concerns of color science (Wyszecki & Stiles, 1982; Arend & Reeves, 1986; Maloney & Wandell, 1986; Brainard & Freeman, 1997; Zaidi, Spehar, & DeBonet, 1997; Kraft & Brainard, 1999).

A surface has many important attributes beyond color. One of these is whether it is wet or dry, a question that is important when eating a piece food, grasping a handrail, or taking a step. Since water is colorless, one might not expect trichromatic vision to be particularly helpful in estimating a surface's wetness. However, we show here that color statistics can be informative even when color identification is not at issue. In particular, when a surface becomes wet, the colors tend to become more saturated. It is straightforward to show, both theoretically and empirically, that this shift in statistics occurs at the physical level, through the interaction of light and matter. It then becomes interesting to ask whether the human visual system "knows" about this shift, and uses it in estimating wetness.

Many a child has gathered attractive wet pebbles along the beach, only to find that their colors fade when they are dry. Figure 1a shows a beach pebble, with one half wet and one half dry. It is easy to see which side is wet: The wet side looks darker and is adorned with bright specular highlights. Less obvious is the fact that

Citation: Sawayama, M., Adelson, E. H., & Nishida, S. (2017). Visual wetness perception based on image color statistics. Journal of Vision, 17(5):7, 1-24, doi:10.1167/17.5.7.

doi: 10.1167/17.5.7

Received December 12, 2016; published May 15, 2017

ISSN 1534-7362 Copyright 2017 The Authors



לחי 🖂

וייז 🏹

ן הט (



Figure 1. Optical properties of a surface largely change before and after wetting. (a) Demonstration of the change in the appearance of a surface by wetting. (b–e) Schematic diagrams of optical properties of wet and dry surfaces. (b) Light paths on a rough-surfaced solid material (dry). (c) Light paths on the same material after wetting with water. (d) Light paths on a porous material (dry). (e) Light paths on the same material after wetting with water.

the color statistics have changed. The colors are generally more saturated.

The change in color statistics with wetness is the main topic of this paper. We discuss the physical processes that lead to these changes, and show how modifying image statistics with digital manipulation can alter the perception of a surface's wetness. In addition to changing color statistics, wetness tends to increase the skew of the luminance distribution, which has been suggested to make surfaces apparently darker, glossier (Motoyoshi, Nishida, Sharan, & Adelson, 2007), and fresher (Arce-Lopera, Masuda, Kimura, Wada, & Okajima, 2012). We can bundle the color and luminance cues together to form what we call the wetness enhancing transformation (WET) that can mimic the effects of wetness. We show that the WET operator tends to increase the appearance of wetness when applied to images of a variety of surfaces. The WET operator is not all-powerful: It works well on some images, and poorly on others. By studying the

conditions for its success and failure, we can better understand how color information influences wetness perception.

Optical aspects of wet and dry materials

Among many visual effects caused by object wetting (e.g., see Rungjiratananon, Kanamori, & Nishita, 2012, for detailed computer simulation of the appearance of wet hairs), the present study focuses on optical effects.

The refractive index of air is approximately 1, while that of most common dielectric (nonmetal) solids (such as plastic, wood, ceramics, and most rocks) is between 1.4 and 1.8. When light strikes a dielectric surface, some of the light is refracted and enters the material, while some undergoes Fresnel reflection at the surface. The amount of reflection depends on the polarization (which we will ignore), the angle of incidence, and the relative refractive indices. The refractive index of water is about 1.33, which is lower than most dielectrics but higher than air. When water is in contact with a solid surface, it offers partial index matching. The amount of reflection is reduced, and the strength of refraction is reduced. Thus, when a light ray enters a solid from water, it is less likely to be reflected, and will be less strongly refracted. The waterborne rays have a better chance of penetrating the surface and traveling within it.

Consider the case of a rough-surfaced solid material consisting of a colored medium filled with scattering particles, as shown in Figure 1b. A few light paths are shown. (Note: In these illustrations we only show rays that escape back to the air, since these are the ones that can potentially be seen, but of course there are many paths, not shown, that end in absorption.) The simplest thing that can happen is a single reflection at the air interface, shown by Path 1. Double reflections can also occur, as in Path 2. Since the surface is rough, the reflection may happen in any number of directions. In the aggregate the reflections can be considered to be in random directions, and from a modest distance the rough surface acts as a diffuse reflector. Fresnel reflection tends to be chromatically neutral, so the rough surface provides a diffuse neutral haze that is added to any light returning after penetrating the surface. Path 3 shows a ray that penetrates into the solid medium and is scattered twice before re-emerging. Since the medium is colored, some wavelengths will be more attenuated than others during this process, and the emerging ray will therefore be colored. A ray that travels a larger distance within the material will undergo more absorption and so will be dimmer and more saturated in color.

Figure 1c shows the same material after wetting with water. The water forms a film of some thickness. Because of surface tension, the water's surface is smoother than that of the material. This means that the Fresnel reflections from the water (Path 4) will tend to appear as sharp highlights rather than as a diffuse haze (Jensen, Legakis, & Dorsey, 1999). Path 5 shows a ray that passes into the water and then penetrates a short distance into the solid before being scattered out. The water's partial index matching enhances the ray's tendency to be refracted (rather than reflected) on the way in as well as on the way out of the solid. Path 6 shows a case involving total internal reflection at the air-water interface (Lekner & Dorf, 1988; Lu, Georghiades, Rushmeier, Dorsey, & Xu, 2006). This ray enters and exits the solid medium, is reflected from the water's inner surface, and enters the solid a second time before exiting into the water and the air. This kind of ray can traverse a good distance inside the solid, and can become strongly colored.

It is impossible to enumerate all the paths that light can take in this situation, since there are so many possible combinations of reflection, refraction, and absorption. However, wetting the surface changes distributions of the various kinds of paths.

Another kind of material is porous and consists of finely divided particles (or fibers) with fine air spaces. Examples would be a porous rock like sandstone, and fibrous materials such as paper or fabric. In the dry state, these materials offer a plethora of air-dielectric interfaces, which in turn offer many opportunities for Fresnel reflections. This is shown in Figure 1d. Path 7 is reflected directly from a particle. Path 8 undergoes two Fresnel reflections, both of them at air interfaces. In Path 9, the light enters a particle but cannot travel a great distance before exiting.

Figure 1e shows some paths from a porous material that has become saturated with water. The water offers partial index matching it its interfaces with all the particles. It also provides a smooth surface that replaces the diffusely reflecting random surface of the dry material. Path 10 shows specular reflection at the air–water interface. Path 11 shows a ray that penetrates the water and then, due to partial index matching, is able to travel through multiple particles before exiting (Twomey, Bohren, & Mergenthaler, 1986).

Image analysis

As reviewed in the last section, wetting the surface is expected to cause a variety of changes in color and luminance distributions. To investigate the actual optical changes, we measured and analyzed images of surfaces before and after wetting.



Figure 2. Photographs of dry and wet materials used in the image analysis. The color and luminance of the surface images were measured and analyzed before and after wetting.

Methods

Surface materials

Thirty-two materials were used for the analysis: clays, soils, cloths, wooden boards, and vegetables. We first took a photograph of each dry material. Then we dripped water onto the surface so that a thin water layer covered it and took a photograph of the wet surface. Thus, the spatial position in each wet and dry photograph was similar, although the water dripped onto the surface significantly changed its fine structure. Figure 2 shows several image examples.

Camera settings

Photographs of dry and wet materials were taken using a standard digital camera (Nikon D5100) with a pixel resolution of 4928×3264 in a photography box illuminated by two fluorescent lamps (Toshiba FL20S-D-EDL-D65, Toshiba, Tokyo, Japan). The camera was set at 90° above each material. The two lamps were set so that they could produce natural ambient light.

To calibrate the color of the camera image, the spectral sensitivities of the digital camera were measured by using a monochrometer (MINI-CHROM DMC1-03, Edmund Optics, Barrington, NJ) and a spectrometer (TOPCON SR-3, TOPCON, Tokyo, Japan). Specifically, the standard white target was illuminated by each monochromatic light, and the



Figure 3. Color and luminance changes between the wet and dry surfaces. (a) The mean u'v' saturation (left) and the mean hue (right) of each wet surface are plotted as a function of those of each dry surface. The u'v' saturation and the hue were calculated using Equations 1 and 2, respectively. The mean hue was calculated using circular statistics (Equation 14). Results show that wetting a surface increases the saturation but has little effect on the hue. (b) The luminance statistics of each wet surface is plotted as a function of that of each dry surface. Different panels indicate different statistics as in the legend. The green line shows the relationship that there are no changes in the statistics between dry and wet surfaces. Results show that the mean luminance of the wet surface is lower than that of the dry surface. In addition, the skew of the luminance distribution tends to increase for wet surfaces in comparison with dry ones.

reflected light was captured by the digital camera and measured by the spectrometer. Since the spectral radiance of each monochromatic light changed with wavelength, we compensated for the difference using the spectral radiance measured by the spectrometer. The spectral sensitivities were measured across the whole visible wavelength spectrum from 380 to 780 nm with an interval of 5 nm. The relative sensitivities across the camera RGB channels were corrected by



Sawayama, Adelson, & Nishida

Figure 4. The color distribution of wet or dry surfaces in the u'v' chromaticity coordinates is shown in each panel. Different dots in the figure indicate different pixels of individual surfaces. The chromaticity values of all 32 surfaces are plotted in each figure. The white circle indicates the white point defined by the white checker in the Macbeth Color Checker. For both dry and wet surfaces, the mean chromaticity across all surfaces was close to the white point. The mean saturation across all materials in the wet condition, shown as the green line in each panel, was 1.43 times larger than that in the dry condition.

using the white checker in a Macbeth Color Checker. On the basis of the spectral sensitivities of the camera and the CIE-1931 two-degree color matching function, we calculated a 3×3 transformation matrix that converts the linear RGB to the XYZ coordinate.

Using the transformation matrix, we calculated the color statistics of wet and dry surfaces. First, the linear RGB images were extracted by de-mosaicing a raw image of each photo, and the center regions of 1536×1536 were cropped. Then, the linear RGB was converted to XYZ and to u'v'Y color coordinates to calculate the saturation in an equal-color space. The u'v' saturation s_{uv} and hue H_{uv} were defined as follows:

$$S_{uv(i,j)} = 13\{(u'_{(i,j)} - u'_n)^2 + (v'_{(i,j)} - v'_n)^2\}^{1/2}, \quad (1)$$
$$H_{uv(i,j)} = \tan^{-1}\{(v'_{(i,j)} - v'_n)/(u'_{(i,j)} - u'_n)\}, \quad (2)$$

where $u'_{(i,j)}$ and $v'_{(i,j)}$ indicate the u'v' chromaticity coordinates of the pixel at position (i, j). The u'v'chromaticity coordinates of the white point (u'_n, v'_n) were (0.2032, 0.4963), which corresponds to (0.3465, 0.3760) in the *xy* chromaticity coordinates. The white point was measured using the white checker in the Macbeth Color Checker positioned under the current lighting condition.

Results

Figures 3 and 4 show the results of the image analysis. Figure 3a shows the color changes between the wet and dry surfaces. The horizontal axis indicates the saturation of the dry surface, while the vertical axis



Figure 5. Effects of specular highlights on luminance distributions of a wet surface. By using a half mirror, a dry or wet material (a brick) was illuminated by a projector from the same direction that images were taken with the color-calibrated camera (see also Appendix A for the detailed camera setting.) (a) The pixel luminance on the wet surface is plotted as a function of that on the dry surface under the condition that the surface normal was parallel to the illumination/camera direction (Slant = 0°). In this condition, the specular reflection from the wet surface directly returns to the camera because the illumination direction was the same as the camera direction. The luminance distribution of the wet surface was more positively skewed than that of the dry surface. (b) The pixel luminance on the wet surface is plotted as a function of that on the dry surface under the condition that the surface normal was slanted 20°, where the specular reflection was almost completely excluded. In this case, surface wetting simply darkened most of the pixels, making the luminance histogram slightly skewed in the negative direction.

indicates that of the wet surface. Different symbols represent different surfaces. The results show that saturation is generally higher for wet surfaces than for dry ones (Figure 3a, left). On the other hand, surface hue is affected little by wetting (Figure 3a, right). In Figure 4, the color distribution in the u'v' chromaticity coordinates is separately shown for dry and wet surfaces. The white circle indicates the white point, and the distance between each dot and the white point indicates how saturated the color of each dot is. The green line in each figure indicates the distance averaged across all materials. The mean saturation across all materials in the wet condition is 1.43 times larger than in the dry condition.

Figure 3b shows the luminance change between the wet and dry surfaces. First, the mean luminance of the wet surface is lower than that of the dry surface, which is consistent with previous studies (Twomey et al., 1986; Lekner & Dorf, 1988). In addition, the skew of the luminance distribution tends to increase for wet surfaces in comparison with dry ones, suggesting an increase in surface specular reflection (Jensen et al., 1999; Motoyoshi et al., 2007; Sharan, Li, Motoyoshi, Nishida, & Adelson, 2008). Other moment statistics (i.e., the standard deviation [*SD*] and kurtosis) are not affected by wetting in a systematic way. These results are consistent with the theoretical analysis described in the previous section.

While we used surface images measured under diffuse illumination in the main analysis, we conducted an additional analysis where we strictly controlled the angular relationship of a directional light source and the surface orientation (Figure 5). By using a half mirror, a dry or wet material (a brick) was illuminated by a projector from the same direction with a colorcalibrated camera. The surface slant of the brick was 0° or 20° . The camera was expected to receive strong specular reflections at 0° but not at 20° .

Figure 5a and b show scatter plots of luminance values for the wet and dry surfaces. The results show that when the surface slant was 0° (Figure 5a), surface wetting darkened some pixels and brightened others, making the luminance histogram heavily skewed in the positive direction. On the other hand, when we nearly eliminated specular highlights at 20° (Figure 5b), surface wetting simply darkened most of the pixels, making the luminance histogram slightly skewed in the negative direction. These results agree with the theory that wetting a surface increases the specular highlights and darkens the diffuse reflections. Interestingly, a wet image that includes few specular highlights (Figure 5, lower) does not appear wet despite its being physically wet.

With the same apparatus, we further analyzed the physical mechanism of the optical effects of wetting a surface through the method developed by Nayar, Krishnan, Grossberg, and Raskar (2006), which separates the direct and global reflection components of a scene using high-spatial-frequency illumination (see Appendix A for details).

Wetness enhancing transformation

The image analysis showed that a wet surface tends to have higher saturation, lower mean luminance, and a







Figure 6. Demonstration of the WET (wetness enhancing transformation) operator. The right image was produced from the left one by i) enhancing the saturation of pixel colors, and ii) tone-remapping of luminance signals with an accelerating nonlinear function, which makes the intensity histogram positively skewed, and decreases the mean luminance. This WET operator tends to make a dry surface look wetter.

more positively skewed luminance distribution than a dry one. To ascertain whether the human visual system uses these changes in the image features as cues to wet surface perception, we examined how they affect a natural dry scene. As shown in Figure 6, the right image was produced from the left one by (a) enhancing the saturation of pixel colors and (b) tone-remapping of luminance signals with an accelerating nonlinear function, which makes the intensity histogram positively skewed and decreases the mean luminance. We refer to these image changes as WET. Given that the image in Figure 6 is properly reproduced, we expect the dry scene on the left to appear wet on the right to many observers. To investigate the perceptual mechanism of surface wetness, we carried out a series of psychophysical experiments using this WET operator.

Psychophysical experiments

Experiment 1a

First, we applied the WET operator to a variety of natural textures and evaluated how it affects the apparent image wetness in psychophysical experiments. We show some examples in Figure 7a and the more examples in Appendix B.

Methods

Unless otherwise stated, the methods used in the subsequent experiments were the same as described in this section.

Participants: Thirteen paid volunteers participated. All had normal or corrected-to-normal vision. The partic-

ipants were naive to the purpose and methods of the experiment. The experiments were approved by the Ethical Committees at NTT Communication Science Laboratories and were conducted in accordance with the Declaration of Helsinki.

Apparatus: The experimental stimuli were displayed using Matlab R2013b in conjunction with the Psychophysics Toolbox 3 (Brainard, 1997; Pelli, 1997). They were displayed on a calibrated 30-in. LCD monitor (EIZO, ColorEdge CG303W) controlled with an NVIDIA video card (Quadro 600) with a pixel resolution of 2560×1600 . The intensity of each phosphor was linearly corrected and could be varied with 10-bit resolution. The maximum luminance of the monitor was 150.6 cd/m^2 . The *xy* coordinates of the white point were (0.3331, 0.3615). A participant viewed the stimuli in a dark room at a viewing distance of 57 cm.

Stimuli: Eighty-eight images were used in Experiment 1a. These images were taken from the Textures category of the McGill Calibrated Colour Image Database (Olmos & Kingdom, 2004; Figure 7a). The stimulus examples are also shown in Appendix B. The WET operator changed the luminance of the original image $Y_{original}$ (which had been normalized in the range from 0 to 1) into Y_{wet} as follows:

$$Y_{wet(i,j)} = Y_{original(i,j)}^{gamma}, \quad (3)$$

where (i, j) indicates the pixel position. The coordinates of white (u'_n, v'_n) for calculating the saturation s_{uv} in Equation 1 were (0.1998, 0.4874). To increase the luminance skew, we set gamma = 3. To enhance the saturation of the texture, we added 0.26 to the s_{uv} of each pixel. In deciding these parameters of the WET transformation, we preliminary explored a range of parameters and chose ones that did not disturb the apparent naturalness of all the images used in Experiment 1a.

Procedures: In the experiments, images selected from the stimulus set were presented one by one. The image sizes were $18.6^{\circ} \times 14.2^{\circ}$ in Experiments 1a and 1c (Figures 7 and 9) and $12^{\circ} \times 12^{\circ}$ in Experiment 1b (Figure 8). Observers rated how wet the surface image appeared on a 5-point scale. At the beginning of the experiment, we showed the observers sample pictures of wet postrain outdoor scenes together with the pictures of dry scenes. Most of the wet scenes included direct evidence of water such as puddles or water drops. Using these "strongly wet" pictures as the reference for the score "5" and dry pictures as that for the score "1," we instructed the observers to evaluate the magnitude of wetness in the presented stimulus image. Each stimulus condition was tested 10 times in total for each observer.



Figure 7. The WET operator tends to make the original image look wetter, while the effect size varies depending on the image. (a) Stimulus examples of the original and the transformed images. The figure below each stimulus shows the hue histogram of the stimulus. The effect of the WET operator was strong for the upper two examples but weak for the lower one. (b) Results of the wetness-rating experiment. The horizontal axis indicates the rating for the original image, and the vertical axis indicates the rating for the transformed image. Each symbol indicates individual images. Error bars indicate ± 1 SEM across observers.

Results

Figure 7b shows the results of Experiment 1a. The horizontal axis indicates the wetness rating for the original surface image, and the vertical axis indicates the rating for the transformed image. A one-way repeatedmeasures analysis of variance (ANOVA) with the WET operator (original or transformed) as the within-subjects variable showed a significant main effect, F(1, 12) =71.384, p < 0.0001). The results show that the perceived wetness was on average much stronger for the transformed images than for the original ones (Figure 7b). In this experiment, the wet rating score of the transformed images was not very high. One reason may be that we used "mild" WET parameters to warrant the naturalness of all the transformed images. Another reason may be that the WET operator does not create the water drops or puddles included in the reference images for the highest score.

Experiment 1b

The luminance transformation of WET is exponential; therefore, the operator not only increases the skew of the luminance distribution but also decreases the mean. To clarify what factor was critical for the wetness perception, we used a histogram-matching method to change the luminance skew of an image by while keeping the luminance mean constant (Figure 8a). In addition, to separately evaluate the contribution of chromatic/ achromatic factors of the WET operator, we continuously and independently manipulated the color saturation and the luminance skew of a natural texture.

Methods

Twelve paid volunteers participated. The luminance skew of an image was changed while the luminance mean was kept constant by using a histogram-matching



Figure 8. Contributions of color saturation and luminance distribution skew to wet perception. (a) Stimulus examples used in Experiment 1b, where we continuously and independently manipulated the color saturation and the luminance histogram skew of a natural surface using histogram matching. (b) Results of Experiment 1b. The horizontal axis indicates the mean saturation defined by Equation 1, and the vertical axis indicates the skew of the luminance histogram. The size of the red circles in each panel indicates the mean rating averaged across observers. The width of the blue contour in each panel indicates + 1 SEM across observers. The results indicate that both saturation and luminance skew contribute to enhancing the wetness impression by the WET operator. (c) Separate plots for the effects of the luminance skew and the color saturation. In the left panel, the mean wetness ratings averaged across color saturation conditions are plotted as a function of the luminance skew. In the right panel, the mean wetness ratings averaged across observers.



Figure 9. The effect of chroma, another definition of the color saturation. For the two images used in Experiment 1a (Surface 1 and 2), we continuously modulated the u'v' saturation and measured the perceived wetness. The results are plotted as a function of the mean chroma of each stimulus. The vertical orange line indicates the mean chroma of the original image. The horizontal green line indicates the mean rating in Experiment 1a for the original image. Different symbols in each panel indicate different skew conditions. Error bars indicate ± 1 *SEM* across observers. The results indicate that the wetting impression increases with color saturation even when chroma is used as the definition of saturation.

method (Heeger & Bergen, 1995). The reference distribution for the matching was a beta distribution, and the skew of the distribution was set from -1.5 to 1.5 in 0.5 steps. The mean luminance and the standard deviation of the surface were set to 0.3 and 0.1, respectively. The saturation distribution of each surface was replaced with a Gaussian distribution using the histogram-matching method. The mean saturation of the Gaussian distribution ranged from 0.026 to 0.36 in 0.026 steps. The standard deviation was set to 0.013. In addition, a monochromatic surface ($s_{uv} = 0$) was also used. In the experiment, a surface image taken with a standard digital camera (Nikon D5100) was used (Figure 8).

Results

Figure 8b shows the results. The horizontal axis indicates the mean saturation, and the vertical axis indicates the skew of the luminance histogram. The size of the red circle in each panel indicates the mean rating averaged across observers. Figure 8c shows the main effect of the luminance skew and that of the mean saturation separately. We conducted a two-way repeated-measures ANOVA with the skew and saturation conditions as factors. The main effects of the skew and saturation conditions and the interaction were statistically significant, F(6, 66) = 28.504, p < 0.0001; F(14, 66) = 28.504, P < 0.0001; P < 0.0001; F(14, 66) = 28.504, F(154 = 35.741, p < 0.0001; and F(84, 924) = 2.3671, p < 0.00010.0001, respectively). The results show that the skew modulation was effective for the wetness perception even without the mean luminance change. Furthermore, over the whole range we tested, an increase in the luminance skew, or in the color saturation, monotonically increased the wetness impression (Figure 8c). It is noteworthy that the wetness rating approached the highest score (5) when the luminance skew was the most positive and the color saturation was the highest.

Experiment 1c

Sawayama, Adelson, & Nishida

$$C_{uv(i,j)}^{*} = \{(u^{*}_{(i,j)})^{2} + (v^{*}_{(i,j)})^{2}\}^{1/2} \quad (4)$$

$$u^{*}_{(i,j)} = 13L^{*}_{(i,j)}(u'_{(i,j)} - u'_{n}), \quad (5)$$

$$v^{*}_{(i,j)} = 13L^{*}_{(i,j)}(v'_{(i,j)} - v'_{n}), \quad (6)$$

$$L^{*}_{(i,j)} = 116(Y_{(i,j)}/Y_{n})^{1/3} - 16, \quad (7)$$

where $u^*_{(i,j)}$ and $v^*_{(i,j)}$ are the chromaticity in the CIE $L^*u^*v^*$ color space and $L^*_{(i,j)}$ indicates lightness in the space. $Y_{(i,j)}$ and Y_n are the luminance values of each pixel and the white point, respectively. Chroma C^*uv is defined as the distance of each pixel chromaticity in the u^*v^* plane from the white point (0, 0). Since the definitions of u^* and v^* are luminance-dependent, chroma C^*_{uv} is also luminance-dependent. By combining Equations 4–7 with Equation 1, chroma C^*_{uv} can be described as follows:

$$C_{uv(i,j)}^* = L_{(i,j)}^* S_{uv(i,j)}.$$
 (8)

The WET operator used in Experiment 1a decreased the mean luminance and thus worked to decrease chroma C^*_{uv} . Therefore, one might argue that the saturation enhancement was needed only to compensate for the decrease in chroma due to the luminance modulation. To confirm whether an increase in chroma is necessary to obtain a wetting effect, we continuously changed the saturation of the original image more widely than we did in Experiment 1b.

Methods

Eight paid volunteers participated. The u'v' saturation of two surface images in McGill textures was continuously modulated by a histogram-matching method. The methods of the transformation were the same as in Experiment 1b. The mean u'v' saturation of each image ranged from 0.065 to 0.52 in 0.065 steps. The standard deviation was set to 0.013. In addition, the skew of each luminance distribution was set to -0.8and 0.8. The mean luminance of the surface image was set to 0.33. The standard deviation of the two surfaces was set to 0.09 and 0.07, respectively.

Results

Figure 9 shows the results. The horizontal axis indicates the mean chroma C^*_{uv} of each stimulus, and the vertical axis indicates the wetness rating averaged across observers. The vertical orange line in each panel shows the chroma C^*_{uv} of the original image. The results show that when the mean chroma of the transformed image was around the original chroma, a strong wetting impression could not be obtained even if the luminance skew was modulated. In addition, the wetting impression increased even when the mean chroma of the transformed image. These findings suggest that an increase in chroma is necessary to obtain a strong wetting effect.

Why is the WET operator more effective on some images than on others?

We have emphasized an explanation of the WET operator based on ecological optics. That is, human vision implicitly knows characteristic optical changes produced by surface wetting. Since the WET operator simulates these changes, it makes surface images look wetter to human observers. However, this is only half of what we can learn about wet processing from the present data. In Experiment 1a, we found a significant variation in the effect size of the WET operator (Figure 7b). The image transformation greatly increased the wetness rating for some images but not for others. What image factor determines the effect size of the WET operator? Using a multiple regression analysis, we evaluated the contribution of a wide range of image statistics of the original (untransformed) images.

Methods

We performed a multiple regression analysis for the ratings. To avoid overfitting, we used the stepwise method, along with the Lasso regression, a multiple linear regression with a shrinkage and selection constraint, to select critical variables (Tibshirani, 1996). In both cases, we regressed a variety of image statistics of the original image to explain the effect sizes of the transformation (i.e., the differences between wetness ratings of the transformed and original images). As the predictor variables, we used the mean, standard deviation, skew, kurtosis, and global entropy of luminance, saturation, or hue in the u'v' chromaticity coordinates.

For luminance and saturation, the statistics are defined as follows:

$$Mean = \frac{1}{N_p} \sum_{p=1}^{N_p} I_{(p)}, \quad (9)$$
$$SD = \sqrt{\frac{1}{N_p} \sum_{p=1}^{N_p} (I_{(p)} - Mean)^2}, \quad (10)$$

$$Skew = \sqrt{\frac{1}{N_p SD^3} \sum_{p=1}^{N_p} (I_{(p)} - Mean)^3},$$
 (11)

$$Kurtosis = \sqrt{\frac{1}{N_p SD^4} \sum_{p=1}^{N_p} \left(I_{(p)} - Mean \right)^4}, \quad (12)$$

$$Entropy = -\sum_{g=0}^{G-1} P_{(g)} \log P_{(g)}, \quad (13)$$

where $I_{(p)}$ is the pixel value at the pixel position p and N_p is the number of pixels. The g is the discrete pixel level, which ranges from 0 to (G-1), and P(g) is the probability density of the level g on the pixel histogram.

Since hue is a circular variable, we used circular statistics (Berens, 2009) as follows:

$$Mean_circular = \frac{1}{N_p} \sum_{p=1}^{N_p} r_{(p)}, \ r_{(p)} = \begin{pmatrix} \cos \alpha_{(p)} \\ \sin \alpha_{(p)} \end{pmatrix}$$
(14)

$$SD_circular = \sqrt{2(1 - ||Mean_circular||)},$$
 (15)

$$Skew_circular = \frac{1}{N_p} \sum_{p=1}^{N_p} \sin 2(\alpha_{(p)} - \bar{\alpha}), \quad (16)$$

$$Kurtosis_circular = \frac{1}{N_p} \sum_{p=1}^{N_p} \cos 2(\alpha_{(p)} - \bar{\alpha}), \quad (17)$$

where $\alpha_{(p)}$ is the pixel hue at the pixel position p and N_p is the number of pixels. The hue entropy is calculated

using Equation 13 because it is not affected by circular variables.

Since the entropy in Equation 13 ignores the spatial component of an image, we also used the spatial entropy (Brink, 1995; Razlighi & Kehtarnavaz, 2009) of luminance, saturation, or hue image as the predictor variable. For the spatial entropy, the joint probability distribution function $P_{gg'(k,l)}$ between samples of $I_{(i,j)}$ and $I_{(i+k,j+l)}$ was used to compute

$$P_{gg'(k,l)} = P_r \{ I_{(i,j)} \in g \& I_{(i+k,j+l)} \in g' \}, g, g', = 0, \dots, G-1,$$
(18)

which corresponds to the probability that a pixel value g occurs at distance (k, l) from another pixel value g'. In the present analysis, the distance (k, l) is within the 3×3 neighborhood of the pixel position (i, j). For each (k, l), an entropy value $H_{(k, l)}$ is defined as follows:

$$H_{(k,l)} = -\sum_{g=0}^{G-1} \sum_{g'=0}^{G-1} P_{gg'(k,l)} \log P_{gg'(k,l)}.$$
 (19)

Since $H_{(k,l)}$ ranges from $H_{(0,0)}$ to $2H_{(0,0)}$, which corresponds to the entropy defined in Equation 13 and twice the entropy, respectively, relative entropy $H_r(k, l)$ is defined as follows.

$$H_{r(k,l)} = \frac{H_{(k,l)} - H_{(0,0)}}{H_{(0,0)}} \in [0,1] \quad (20)$$

By simply summing the entropy values for all (k, l), the spatial entropy value $H_{spatial}$ was obtained. To see the spatial entropy on multiple scales, we computed $H_{spatial}$ for images scaled down to 1/1, 1/2, 1/4, 1/8, 1/16, and 1/32. For all the entropy computations, the number of histogram bins (i.e., the resolution of *I*) was set to 512.

Thus, 33 of the predictor variables in total were used in the present regression. Before the regression analyses, each predictor variable was normalized to z scores.

Results

Figure 10 shows the results of the single regression for each feature, stepwise multiple regression, and Lasso regression. The results of the stepwise multiple regression analysis show significant contributions of luminance histogram skew, luminance spatial entropy, hue (global) entropy, and hue spatial entropy. The results of the Lasso regression (Tibshirani, 1996) support strong positive contributions of hue entropy and luminance spatial entropy (Figure 10).

<u<text>
 indicates that it has a biased distribution (Figure 7a, indicates that it has a biased distribution (Figure 7a, indicates that it has a biased distribution (Figure 7a, indicates that it has a biased distribution of the that it has a biased to the therefore the therefore it has a biased of the therefore the therefore it has a biased of the therefore the therefore it has a big to the therefore therefore it has a big to the therefore the therefore it has a big therefore the therefore it has a big therefore the therefore it has a big therefore therefore the therefore it has a big therefore therefore the therefore it has a big therefore therefore the therefore it has a big therefore therefore the therefore it has a big therefore the therefore the therefore it has a big therefore the therefore the therefore it has a big therefore the therefore the therefore it has a big therefore the therefore the therefore it has a big therefore the therefore the therefore it has a big therefore the therefore the therefore it has a big therefore the therefore therefore it has a big therefore therefore therefore therefore it has a big therefore therefore the therefore therefore it has a big therefore therefore the therefore the therefore it has a big therefore the therefore the therefore therefore it has a big therefore therefore the therefore the therefore the it has a big therefore the therefore the therefore the therefore it has a big therefore the therefore the therefore the therefore the it has a big therefore the therefore the therefore the therefore the it has a big therefore the therefore the therefore the therefore the it has a big therefore the therefore the therefore the therefore the it has a big therefore the it has a big therefore therefore the therefore the therefore there

Experiment 2

Motivated by the results in Figure 10, we experimentally explored the genuine contribution of the hue entropy to the wetness perception. We synthesized artificial Mondrian-like three-dimensional (3-D) scenes using a physically based renderer (Jakob, 2010) and investigated their appearance.

Methods

Stimuli: The original 3-D Mondrian images were made by using the physically based renderer *Mitsuba* (Jakob, 2010). In a 3-D space, a camera was set 1.5 m from a plain gray surface 3×3 m in size, and 750 small cubes were randomly placed in front of it. The width, height, and depth of each cube were randomly selected in a range from 0.01 to 0.1 m. The bidirectional reflectance distribution function (BRDF) of each cube was Lambertian, and its reflectance was defined in the spectral domain. The mean reflectance of each cube was randomly modulated. However, to control the hue entropy of the output image, we used three entropy conditions-small, intermediate, and large entropy-for the peak in the reflectance spectra of the cube (Figure 11a, upper). For all the conditions, the reflectance spectra $R_{(\lambda)}$ of each cube were defined as follows:

$$R_{(\lambda)} = a_1 + a_2 exp\{-\frac{\lambda - \lambda_p}{2\sigma}\} \quad (21)$$

where the λ is the wavelength of the light, which ranged from 380 to 780 nm, λ_p is a peak in reflectance spectra, and σ is the standard deviation, which was set to 20 nm. The a_2 is a contrast factor, and it was set to 0.3. The a_1 is pedestal reflectance, and it was randomly selected from 0.3 to 0.5. For the small entropy condition, the peak in reflectance spectra λ_p of each cube was always 700 nm. For the middle entropy condition, a peak in reflectance spectra λ_p of each cube was randomly selected from 546 or 700 nm. For the large entropy condition, the peak was randomly selected from 400 to 700 nm. The scenes



Figure 10. Regression analysis as to which image features contribute to the effect size of WET. (a) Results of the single regression analysis. Red characters indicate that these predictor variables had statistically significant correlations with the effect size (p < 0.05). (a) Results of the stepwise multiple regression analysis. Red characters indicate that these predictor variables had statistically significant weights (p < 0.05). The results indicate significant contributions of luminance histogram skew, luminance spatial entropy, hue (global) entropy, and hue spatial entropy. (b) Results of the lasso regression analysis. The results support strong positive contributions of hue entropy and luminance spatial entropy.



Figure 11. Effects of the WET operator depend on the hue entropy. (a) Stimulus examples used in Experiment 2. We used a physically based renderer to make the original images. By modulating the peak in reflectance spectra of each object, we controlled the hue entropy of the image. The hue entropy of each image under the small, middle, or large hue entropy conditions was 2.8, 4.8, or 6.6, respectively. To make the transformed image, we applied the WET operator to the original image. (b) Results of the wetness judgment (left) and darkness judgment (right). Ratings of the pairwise comparison were converted into a one-dimensional perceptual scale. The results indicate that the wetness perception of a scene image increases with the hue entropy of the image without changes in the luminance spatial entropy or in other factors.

 Procedure: The observers' task (eight paid volunteers) was to judge wetness (or darkness). In each trial, the observers to judge wetness (or darkness). In each trial, the observers to judge wetness (or darkness). In each trial, the observers in grate a pair of or dark or judge wetnes (or dark or judge). In each trial, the elefter of the server wetness is a rating of ±1 indicated to be vetner. In the respective or dork or

Analysis of the results: Following Scheffé (1952), we calculated the perceived scale values $Q_{(m)}$ for each





Figure 12. Two-dimensional (2-D) Mondrian patterns used in a preliminary test of Experiment 2. Some observers showed results similar to those obtained with 3-D patterns, but others had trouble judging wetness for the 2-D patterns.

stimulus condition using the following formulas:

$$P_{(m,n)} = \frac{1}{2} (M_{(m,n)} - M_{(n,m)}), m, n = 1, \dots 6, \quad (22)$$
$$Q_{(m)} = \frac{\sum_{n=1}^{S} P_{(m,n)}}{S}, \quad (23)$$

Results

The ratings for the stimulus pairs within each observer were analyzed by using Equations 22 and 23. The calculated value of the perceived wetness or darkness is shown in the Figure 11b. Different symbols indicate different stimulus conditions. We conducted a two-way repeated-measures ANOVA for the wetness rating, with entropy condition (small, intermediate, or large) and WET condition (original or transformed) as factors. The main effects of the entropy and WET conditions and the interaction were statistically significant, F(2, 14) = 11.942, p < 0.0001, F(1, 7) = 189.766, p < 0.0001, and F(2, 14) = 8.244, p = 0.004, respectively. The results show that the wetness perception of a scene

image increased with the hue entropy of the image without changes in the luminance spatial entropy or in other factors. Specifically, the perceived wetness of the transformed image in the large entropy condition (Figure 11, open blue circle) was statistically higher than in the small and intermediate entropy conditions (Figure 11, open red circle and open green circle), t(7) =4.988, p = 0.005 for small entropy, and t(7) = 3.338, p =0.038 for intermediate entropy (Bonferroni-corrected paired t test), but there was no significant difference between the wetness values of the transformed image in the small and intermediate entropy conditions, t(7) =0.2494, p = 0.748 (Bonferroni-corrected paired t test).

Furthermore, the transformed images were generally perceived as wetter than sformed images were generally perceived as wetter than stormed images were generally perceived as wetter than the transformed images constructed and the transformed to the transformed to transformed to the transformed to the transformed to transformed to the transformed to the transformed to transformed

In addition, using the same stimulus set, we found that the hue entropy did not affect the darkness rating (Figure 11b, right). The dissociation of wetness and darkness ratings excludes the possibility that the observers simply judged the darkness of a surface to estimate its wetness.

While we used 3-D Mondrian patterns in the main experiment, we used 3-D Mondrian patterns in the main experiment, we used 3-D Mondrian patterns in the main experiment, we used 3-D Mondrian patterns in the server experiment, we used 3-D montres is a surface/object property, the 2-D patterns, is a surface of the server experiment of the server of the server of the server experiment of the server of the server of the server of the server experiment of the server of the ser

Experiment 3

Methods

The original image was a high-hue entropy scene (Figure 13a, upper left). In one condition, the mean luminance was decreased by multiplying the pixel



Figure 13. (a) Stimulus examples used in Experiment 3. (b) Results of Experiment 3. The results indicate that decreasing the mean luminance (green) is not sufficient for significantly increasing the wetting impression.

intensities by a constant factor of 0.2 while the luminance skew was held constant (Figure 13a, upper center). In the other condition, the luminance gamma was increased to simultaneously reduce the mean luminance and increase the luminance skew (upper right), as in the original WET operator. The mean luminance was equated between the two conditions. The color saturation was equated either to the original image or to a saturation-doubled image. The observers' task (eight paid volunteers) was to judge wetness (or darkness). The other methods were the same as in Experiment 2.

Results

The perceived wetness or darkness, calculated in the same way as in Experiment 2, is shown in the Figure 13b. Different symbols indicate different stimulus conditions. We conducted a two-way repeated-measures ANOVA for the wetness and darkness ratings, with luminance condition (original, mean decreased, or gamma modulated) and saturation condition (original or highly saturated) as factors. For the wetness rating, the main effects of the luminance condition was statistically significant, F(2, 14) = 26.440, p < 0.0001,while the main effect of the saturation condition and the interaction were not, F(1, 7) = 2.665, p = 0.147, and F(2, 14) = 1.519, p = 0.253, respectively. For the darkness rating, the main effects of the luminance and the saturation conditions were statistically significant, F(2, 14) = 73.534, p < 0.0001 and F(1, 7) = 15.078, p = 0.006, respectively), while the interaction was not, F(2, (14) = 0.382, p = 0.689. The results show that decreasing the mean luminance (Figure 13, green) was not enough to significantly increase the wetting impression. Specifically, regardless of the magnitude of color saturation, the apparent darkness did, t(7) = 6.938, p = 0.0006(Bonferroni-corrected paired t test), but the apparent wetness did not, t(7) = 2.130, p = 0.21 (Bonferronicorrected paired t test) show significant differences between the original image (Figure 13, red) and the



Figure 14. Hue entropy can reduce the ambiguity of physical events. (a) Schematic explanations of the relationship between physical events and image features. Besides surface wetness, other factors, such as the existence of highly color-saturated glossy surfaces, can produce similar image features. This is why the optical image features of luminance and saturation are not sufficient for estimating that the surface is wet. (b) Explanation of a framework that can explain the effects of hue entropy. Each colored panel indicates the probabilities of the occurrence of a wet or dry scene. As in Figure 4, we here assume that the probabilities of the occurrence for wet colors spread 1.43 times more widely than that for dry colors. The upper and lower panels indicate the small and large hue entropy conditions, respectively. For the large hue entropy, the posterior probability for the wet scene can be larger than that for the dry scene, while it tends to be smaller for the small hue entropy.

mean-decreased image (Figure 13, green). In contrast, the original WET operator (Figure 13, blue) produced a strong wetting effect, t(7) = 5.341, p = 0.003 (Bonferroni-corrected paired t test).

General discussion

<u<text>
 The present study investigated human perception of surface wetnesent study investigated human perception of surface wetnesent study investigated human perception of surface wetneses by analyzing the united to surface wetneses by analyzing the surface subsurface subsurface wetneses by analyzing the surface subsurface surface by analyzing the surface subsurface wetneses that the wetnese perception the the effect size of the wetnese suggest that be wetnese to the surface surface wetnese the wetnese suggest that be wetnesed to the sufface subsurface wetnese the wetnese the subsurface of the wetnese the subsurface subsurface wetneses that the effect size of the WET operator depends on the hue
entropy of the image, or a similar measure of the spread of color distribution.

The effect of hue entropy cannot be explained by ecological optics, since wet-related changes in color saturation and luminance histograms take place hue entropy may reflect the perceptual computation the brain adopts to evaluate surface wetness. Even when wet image features are recognized in the image, the
 observers cannot tell whether they are indeed caused by surface wetness (Figure 14a). This is because similar image features could be caused by other events—for instance, when the visual scene happens to include objects with highly color-saturated glossy surfaces. In addition, as shown by Figures 3 and 4, wet-related changes in color and luminance statistics are small relative to their variances. Since there is a significant overlap in the distribution of color saturation between dry and wet samples, it seems it was hard for observers to tell whether one color sample is taken from a dry image or a wet image. The key to resolving these ambiguities is to increase the number of samples. When wet-related image features are shared by many different parts in the scene, the image features are likely to be



Figure 15. The results of a simulation of the model in Figure 14 under an assumption that the probability distribution of color is approximated by a Gaussian function, with the spread's being ×1.43 wider for the wet condition than for the dry condition. A diagonal view (a) and a top view (b) of the frequency distribution histogram of the combination of $\Pi P(S_x|\text{Dry})$ and $\Pi P(S_x|\text{Wet})$. The blue histogram was obtained when *N* color samples were randomly taken from the wet color distribution; the red histogram was obtained when they were randomly taken from the dry color distribution. The dotted line is the boundary above which the likelihood is higher for wet than for dry. When the number of samples is one (*N*=1), blue and red histograms significantly overlap each other. As the number of samples is increased, however, the blue and red distributions are gradually separated from each other, making it easier to tell whether the color samples are from the wet or dry distribution.

produced by a global common factor such as wetting. In other words, the more independent colors the scene contains, the more reliably the visual system can judge the scene wetness.

This explanation may have a relationship with a Bayesian inference model (Figure 14b; e.g., Knill & Richards, 1996; Weiss, Simoncelli, & Adelson, 2002). The problem the observer has to solve is to classify a scene into dry or wet on the basis of the difference in the color saturation distribution. The saturation is generally higher for wet scenes than for dry ones. The probability distributions of the saturation for the dry and wet scenes, P(S|Dry) and P(S|Wet), are something like those shown in Figure 4. With regard to the prior probability, we assume that the scene is more likely to be dry than wet: P(Dry) > P(Wet). The scene *S* is judge as wet when P(Dry|S) < P(Wet|S):

$$P_{(\mathrm{Dry}|S)} \propto P_{(\mathrm{Dry})} \prod_{x=1}^{N} P_{(S_x|\mathrm{Dry})}, \quad (24)$$
$$P_{(\mathrm{Wet}|S)} \propto P_{(\mathrm{Wet})} \prod_{x=1}^{N} P_{(S_x|\mathrm{Wet})}, \quad (25)$$

where N is the number of independent colors present in the scene, which usually increases with the hue entropy.

According to this model, when the number of independent colors is small (the hue entropy of the image or the amount of color information is low) and the sampled colors are saturated, the scene is judged to be dry because of the difference in the prior probability. On the other hand, when the number of independent colors is large (hue entropy is high) and sampled colors are all saturated, the scene is judged to be wet because of the increased contribution of the likelihood. For instance, let us consider a simple situation where (a) there are only two color saturation states, low and high; (b) the likelihood of high saturation is higher in the wet condition than in the dry condition, $P(S_{high}|Dry) = 0.4$, $P(S_{\text{low}}|\text{Dry}) = 0.6, P(S_{\text{high}}|\text{Wet}) = 0.5, P(S_{\text{low}}|\text{Wet}) =$ 0.5; and (c) the prior probability is higher for the dry condition than for the wet condition, P(Dry) = 0.7 and P(Wet) = 0.3. When there is only one color sample and it has high saturation, the model predicts dry, since $P(Dry|S_{high}) = 0.28 > P(Wet|S_{high}) = 0.15$. On the other hand, when there are five samples and all of them have high saturation, the model predicts wet, since P(Dry- $S_{\text{high}} = 0.0072 < P(\text{Wet}|S_{\text{high}}) = 0.0094.$



Figure 16. Examples of the DET operator. The DET can tend to make the original images look dryer.

We also consider the model behavior under a more realistic assumption: (a) the probability distribution of color is widely spread and approximated by a Gaussian function centered around the white point; and (b) the standard deviation of the Gaussian distribution is $\times 1.43$ larger (the average saturation is $\times 1.43$ higher) for the wet condition than for the dry one. For a given dry/ wet image, N independent colors are randomly sampled from the dry/wet distribution. By repeating this 1,024 times, we obtained the frequency distribution histogram of the combination of $\Pi P(S_x | \text{Dry})$ and $\Pi P(S_x | \text{Wet})$ as shown in Figure 15. The blue histogram was obtained when N color samples were randomly taken from the wet color distribution; the red one when they were randomly taken from the dry color distribution. Both histograms are aligned on a single line for each N. The dotted line is the boundary above which the likelihood is higher for wet than for dry. When the number of samples is one (N = 1), the blue and red histograms significantly overlap with each other. It is hard for the observer to tell whether the sample is from the wet or dry distribution only from the likelihood terms (and the prior produces a decision bias for dry). As the number of samples is increased, however, the blue and red distributions are gradually separated from each other, making it easier to tell whether the color samples are from the wet or dry distribution (unless a strong prior bias in favor of dry overrides the difference in the likelihood terms). The model nicely explains how an increase in hue entropy, or in the number of different color samples, elevates the reliability of color saturation as a cue to judge surface wetness.

However, it should be noted that we propose this model only for the purpose of giving insight into the effect of hue entropy on wetness perception. We do not expect the model, at least the current version, to be able to precisely predict human wetness perception. One limitation is that it assumes an ideal observer who perfectly knows the natural color statistics. In addition, the model does not take into account the contribution of many other cues to wetness perception.

Although the present study focused on the role of color and luminance statistics of the image in human wetness perception, we do not claim that the visual system can evaluate the wetness of an image only by using those image statistics, nor that those statistics are powerful enough to predict human wet/dry judgments regardless of other stimulus conditions. Our claim is that, all else being equal, changing the statistics tends to change the wetness in a predictable way. The idea is easy to understand when we consider the example of classical cues. For instance, height in visual field does not, by itself, predict apparent depth. However, in a given picture, increasing an object's in height in visual field tends to increase its apparent distance. The angular subtense of an object on the page does not, by itself, predict its apparent 3-D size. However, increasing on object's angular subtense tends to increase its apparent 3-D size. Analogously, the color saturation of an image may not, by itself, predict the apparent magnitude of wetness, but increasing it tends to make a dry surface look wet. Therefore, color saturation is a cue to perceive wetness, as an object's height is a cue to perceive depth.

Wetting real objects often produces a characteristic shape change. For instance, when a hair-like object is wet, the absorbed water makes the object heavier and alters its chemical properties (Rungjiratananon et al., 2012). The resulting shape change should be an effective cue for the wetness perception. Wet scenes often contain direct visual evidence of water, such as drops and puddles. For a fuller understanding of human wetness perception, the contribution of the nonoptical factors should be taken into account.

The present case study for wetness perception demonstrates how useful color image statistics could be for estimating physical surface conditions or other events that alter the spectral characteristics of light reflectance. In this regard, the present finding is related to the effects of color on intrinsic image decomposition. It is known that human observers tend to interpret luminance spatial changes accompanied by color changes as material changes, while they tend to interpret those orthogonal to color changes as illumination changes (Kingdom, 2003). This suggests that the visual system "knows" that color changes are usually caused by material changes, but that luminance changes can be caused be either material or illumination changes. Human observers also use color information to segregate highlights from body color, apparently exploiting the differences in the spectral properties of diffuse and specular reflections (Nishida, Motoyoshi, & Maruya, 2011). Together with these

previous findings, the present finding suggests that the brain makes use of color information not only for perceiving color but also for decoding the optics of the external world embedded in the retinal image.

The WET operator is a practically useful imageprocessing tool. It can quickly change a dry scene into a wet one by a simple image manipulation. One might then wonder if an inverse function of the WET operator can make a wet surface look dryer. The results of Experiment 1b suggest that changing the parameters of the WET operator in the opposite direction enhances apparent dryness. Furthermore, as shown in Figure 16, we observed that the dryness enhancing transformation (DET) was able to make some real wet scenes look dryer. However, the DET would be more limited than the WET because it cannot undo/remove nonoptical changes produced by actual wetting, such as shape changes and water drops.

To conclude, the human visual system tends to perceive surface wetness when the color saturations are high, the luminance histogram is positively skewed, and the hue entropy is high. The first two conditions agree with actual image changes produced by surface wetting (ecological optics), while the last condition can be explained by a Bayesian-like inference that takes into account the probability of dry or wet surfaces in the environment. The present findings not only provide a novel view of human wetness perception, but also reveal how elegantly human vision uses low-level color statistics to estimate complex optical properties in the real world. A scientific understanding of the strategies the human visual system takes to estimate surface wetness (see also Sawayama & Kimura, 2015) will lead to the development of techniques to synthesize, edit, and recognize wet scenes that are much simpler than physics-based ones (Jensen et al., 1999; Lenaerts, Adams, & Dutré, 2008; Chen, Thalmann, & Allen, 2012; Rungjiratananon et al., 2012).

Acknowledgments

This work was supported by Grant-in-Aid for Scientific Research on Innovative Areas from Japan Society of Promotion of Science to SN (JSPS KAKENHI Grant Number JP15H05915) and by a grant from the National Science Foundation to EHA.

Commercial relationships: none. Corresponding author: Masataka Sawayama. Email: masa.sawayama@gmail.com. Address: NTT Communication Science Laboratories, Nippon Telegraph and Telephone Corporation, Kanagawa, Japan.

References

- Arce-Lopera, C., Masuda, T., Kimura, A., Wada, Y., & Okajima, K. (2012). Luminance distribution modifies the perceived freshness of strawberries. *i-Perception*, 3(5), 338–355.
- Arend, L., & Reeves, A. (1986). Simultaneous color constancy. *Journal of the Optical Society of America* A, 3(10), 1743–1751.
- Berens, P. (2009). CircStat: A MATLAB toolbox for circular statistics. *Journal of Statistical Software*, 31(10), 1–21.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436.
- Brainard, D. H., & Freeman, W. T. (1997). Bayesian color constancy. *Journal of the Optical Society of America A*, 14(7), 1393–1411.
- Brink, A. (1995). Minimum spatial entropy threshold selection. *IEEE Proceedings - Vision, Image and Signal Processing*, *142*(3), 128–132.
- Chen, Y., Thalmann, N. M., & Allen, B. F. (2012). Physical simulation of wet clothing for virtual humans. *The Visual Computer*, 28(6–8), 765–774.
- Heeger, D., & Bergen, J. (1995). Pyramid-based texture analysis/synthesis. Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques (pp. 229–238). New York: ACM.
- Jakob, W. (2010). Mitsuba renderer. Retrieved from http://www.mitsuba-renderer.org
- Jensen, H. W. (2001). *Realistic image synthesis using photon mapping*. Natick, MA: A. K. Peters.
- Jensen, H. W., Legakis, J., & Dorsey, J. (1999). Rendering of wet materials. In *Rendering Techniques '99* (pp. 273–281). New York: Springer.
- Kingdom, F. A. (2003). Color brings relief to human vision. *Nature Neuroscience*, 6(6), 641–644.
- Knill, D., & Richards, W. (1996). Perception as Bayesian inference. Cambridge, UK: Cambridge University Press.
- Kraft, J. M., & Brainard, D. H. (1999). Mechanisms of color constancy under nearly natural viewing. *Proceedings of the National Academy of Sciences*, 96(1), 307–312.

Lekner, J., & Dorf, M. C. (1988). Why some things are darker when wet. *Applied Optics*, 27(7), 1278–1280.

Lenaerts, T., Adams, B., & Dutré, P. (2008). Porous

flow in particle-based fluid simulations. *ACM Transactions on Graphics (TOG)*, 27(3), 49:1–49:8.

- Lu, J., Georghiades, A. S., Rushmeier, H., Dorsey, J., & Xu, C. (2006). Synthesis of material drying history: Phenomenon modeling, transferring and rendering. Paper presented at the ACM SIG-GRAPH 2006 Courses, Boston, MA.
- Maloney, L. T., & Wandell, B. A. (1986). Color constancy: A method for recovering surface spectral reflectance. *Journal of the Optical Society of America A*, 3(1), 29–33.
- Mollon, J. D. (1989). "Tho'she kneel'd in that place where they grew..." The uses and origins of primate colour vision. *Journal of Experimental Biology*, 146(1), 21–38.
- Motoyoshi, I., Nishida, S., Sharan, L., & Adelson, E. H. (2007). Image statistics and the perception of surface qualities. *Nature*, 447(7141), 206–209.
- Nayar, S. K., Krishnan, G., Grossberg, M. D., & Raskar, R. (2006). Fast separation of direct and global components of a scene using high frequency illumination. ACM Transactions on Graphics, 25(3), 935–944.
- Nishida, S. Y., Motoyoshi, I., & Maruya, K. (2011). Luminance-color interactions in surface gloss perception. *Journal of Vision*, *11*(11):397, doi:10. 1167/11.11.397. [Abstract]
- Olmos, A., & Kingdom, F. A. A. (2004). A biologically inspired algorithm for the recovery of shading and reflectance images. *Perception*, 33(12), 1463–1473.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Razlighi, Q. R., & Kehtarnavaz, N. (2009). A comparison study of image spatial entropy. Proceedings of SPIE 7257, Visual Communications and Image Processing 2009, 72571X, 1–10.
- Rungjiratananon, W., Kanamori, Y., & Nishita, T. (2012). Wetting effects in hair simulation. In *Proceedings of Pacific graphics* (pp. 319–328). Hoboken, NJ: Wiley-Blackwell.
- Sawayama, M., & Kimura, E. (2015). Stain on texture: Perception of a dark spot having a blurred edge on textured backgrounds. *Vision Research*, 109, 209– 220.
- Scheffé, H. (1952). An analysis of variance for paired comparisons. *Journal of the American Statistical Association*, 47(259), 381–400.
- Sharan, L., Li, Y., Motoyoshi, I., Nishida, S. Y., & Adelson, E. H. (2008). Image statistics for surface reflectance perception. *Journal of the Optical Society of America A*, 25(4), 846–865.

- Steward, J. M., & Cole, B. L. (1989). What do color vision defectives say about everyday tasks? Optometry & Vision Science, 66(5), 288–295.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B Methodological*, 58, 267–288.
- Twomey, S. A., Bohren, C. F., & Mergenthaler, J. L. (1986). Reflectance and albedo differences between wet and dry surfaces. *Applied Optics*, 25(3), 431– 437.
- Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, 5(6), 598–604.
- Wyszecki, G., & Stiles, W. S. (1982). *Color science* (Vol. 8). New York: Wiley.
- Zaidi, Q., Spehar, B., & DeBonet, J. (1997). Color constancy in variegated scenes: Role of low-level mechanisms in discounting illumination changes. *Journal of the Optical Society of America A*, 14(10), 2608–2621.

Appendix A

Several optical factors contribute to producing the color and luminance changes when a surface is wetted. Some are related to multiple reflections or subsurface scattering (Path 6 in Figure 1), while others are related to direct reflection, such as the physical cause of index matching (Path 5 in Figure 1).

To explore the contributions of different optical factors, we measured the color of a dry or wet brick (rough dielectric material), which was illuminated by a projector (Epson EH-TW6600W, Espon, Tokyo, Japan) from the same direction with a color-calibrated camera (Pointgrey Flea 3, Richmond, BC, Canada) using a half mirror (Figure A1a). The surface was slanted 20° to attain the condition without the specular reflection. We separated the direct component of the brick from the global one using high-frequency illumination (Nayar et al., 2006). In theory, the method as mentioned in the paper, we used 25 images of the brick illuminated by different checkerboard illuminations (Figure A1b) and extracted the on-(light checker) and off-(dark checker) images of the brick from the sequences. The off-image mainly included the global component from neighboring activated regions, although there were weak direct and global components because of the brightness of the deactivated light source. In contrast, the on-image strongly included the direct component and also included the global com-



Figure A1. Separation of direct and global reflectance components to analyze optical factors producing the color and luminance changes. We measured the color of a dry or wet brick (rough dielectric material), which was illuminated with a projector from the same direction as a color-calibrated camera using a half mirror (a). The surface normal was slanted 20° from the illumination/viewing axis to reduce specular reflections. We separated the direct component of the brick from the global one using high-frequency illumination (b) and extracted the on-(light checker) and off-(dark checker) images of the brick from the sequences. From these images, the direct (c, middle) and global (c, lower) components were extraced. The direct image (c, middle) should arise only from the direct reflection of the brick. In contrast, the global component image (c, lower) should arise from interreflection, subsurface scattering, volumetric scattering, or translucency.

ponent from neighboring regions. If the spatial ratio of the activated and deactivated pixels (i.e., the ratio of the light and dark checkers) and the brightness of the deactivated source are given, these on- and off-images can be transformed to the direct and global components by a linear transformation. By using this method, the global component image (Figure A1c, lower) should arise from indirect factors, including multiple interreflection, subsurface scattering, volumetric scattering, or translucency. In contrast, the direct component image (Figure A1c, middle) arises mainly from the direct reflection from the brick, which is affected by index matching.

Results of the color analyses between the wet and dry objects are shown in Figure A2. They show that under the direct component condition the color saturation of the wet object was higher than that of the dry one (Figure A2a). Under the global component condition the saturation of the wet object was also higher than that of the dry one as a whole (Figure A2b). The results suggest that both index matching (Path 5) and multiple reflections (Path 6) contribute to the optical changes produced by wetting a brick surface.

Appendix B

In Experiment 1a, the WET operator was effective for some textures but not for others. Some of the examples used in Experiment 1a are shown in Figures B1 and B2 with the mean rating averaged across observers.



Figure A2. Luminance and color differences between the wet and dry images for the direct and global reflectance components. (a) For the direct component images, the difference in color or luminance between the wet and dry images within the region-of-interest (ROI) is plotted in the right graphs. There are a clear luminance reduction and a saturation increase. Most of the image changes of the direct component can be ascribed to index matching (i.e., a decrease in the refractive index difference between air and the brick due to the brick's having a water film). (b) The results for the global component condition. Under the global component condition, the saturation of the wet object was also higher than that of the dry one as a whole, but the amount of the increase was smaller than that under the direct component condition.

Journal of Vision (2017) 17(5):7, 1-24



Figure B1b. Stimulus examples on which the WET operator was effective in Experiment 1a.

Journal of Vision (2017) 17(5):7, 1-24

Sawayama, Adelson, & Nishida



Rating:

1.5

Rating: 1.3

Rating: 1.6

Figure B2b. Stimulus examples on which the WET operator was ineffective in Experiment 1a.

Rating:

1.9