

# Physical Asymmetries and Brightness Perception

James J. Clark

**Abstract** This paper considers the problem of estimating the brightness of visual stimuli. A number of physical asymmetries are seen to permit determination of brightness that is invariant to certain manipulations of the sensor responses, such as inversion. In particular, the light-dark range asymmetry is examined and is shown to result, over a certain range, in increased variability of sensor responses as scene brightness increases. Based on this observation we propose that brightness can be measured using variability statistics of conditional distributions of image patch values. We suggest that a process of statistical learning of these conditional distributions underlies the Stevens effect.


## 1 Introduction - Is it Dark or Bright?

Suppose one is viewing a scene, such as looking out onto a busy street on a bright sunny day, or looking around your moonlit kitchen for a midnight snack with the lights turned off. In this paper we will be concerned with the perception of how bright a viewed scene is, and consider the question “what makes one scene appear bright while the other appears dark?”. The term *brightness* denotes the subjective perception of the *luminance* of a visual stimulus, where luminance is a photometric (i.e. perceptually weighted) measure of the intensity of light (either reflecting from a surface or being emitted from a light source) per unit area traveling in a particular direction.

A naive answer to the question of what determines the perception of brightness would be to simply associate “dark” with low sensor signal values and “light” with high sensor signal values, as depicted in figure 1.

---

James J. Clark  
Centre for Intelligent Machines, McGill University, 3480 University Street, Montreal, Quebec,  
Canada, e-mail: clark@cim.mcgill.ca



**Fig. 1** Is the perception of brightness based on the firing rates of neurons in the brain?

There are a number of problems with this naive approach, however. To begin with, the sign of the variation in sensor signals with the luminance of the visual stimulus is an arbitrary convention. One could just as easily have sensors whose signals are high when the luminance of the stimulus is low and vice-versa. For example, in the human retina both types of sensors are found, where bipolar cells either respond to the presence (ON-cells) or to the absence (OFF-cells) of incident light [1].

Some sensors respond to spatial or temporal contrasts (or derivatives). To take a specific example, consider that signals to the visual cortex from the retina are in the form of ON-Center/OFF-Surround and OFF-Center/ON-Surround signals. These could be integrated to recover the luminance, but the ambiguity in the sign remains. This also implies that the sensor signal may depend on the spatial and temporal "surround".

The issues just mentioned suggest that brightness is perceived in a way that involves more than just the raw signal levels from the image sensors. This paper describes a possible approach for doing this, one that is based on consideration of physical asymmetries that reveal differences between light and dark.

## 2 Physical Asymmetries Underlying Brightness Perception

The most fundamental asymmetry that we will look at is the so-called *light-dark range asymmetry*. This asymmetry can be understood by noting that there is a wider range of sensor values possible in a bright scene than in a dark one. Suppose, for

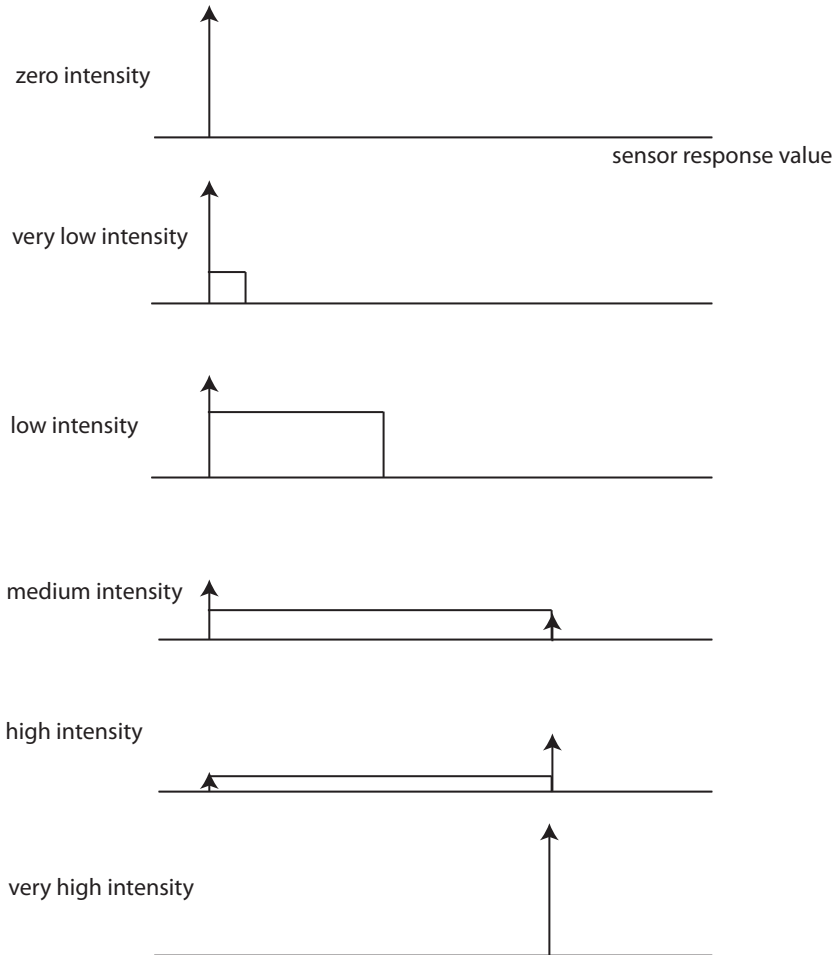
argument's sake, that we have a strictly increasing monotonic visual sensor with an infinite dynamic range. That is, its response is a strictly increasing function of the intensity of the incident light. If this sensor views a scene consisting of a single non-luminous textured convex Lambertian object, illuminated by a single point light source having illuminance  $L$ , there will be a finite maximum value that this sensor could produce. This maximum value will depend on the sensitivity of the sensor,  $\sigma$ , the maximum albedo of the object, and the illuminance  $L$  of the object surface. The range of albedo values for a non-luminous object must be in the range  $[0, 1]$ . Thus the range of sensor values will be  $[0, L\sigma]$ . As the surface illuminance  $L$  increases, so does the range of possible sensor values. This increase of the range would persist even if the sensor was instead taken to have a strictly decreasing response (corresponding to a negative  $\sigma$ ) or had a constant offset (so that the sensor had a non-zero response to a zero incident intensity).

The analysis is more involved, but the light-dark range asymmetry will also be present for more complicated scenes, with multiple non-Lambertian objects and multiple distributed illuminants. Singularities such as caustics created by mirrors and lenses can create infinite intensities, but only over vanishingly small areas. Sensors with finite extent will have a finite response to such caustics, and this response will be scaled by the illuminance of the light source.

### ***2.1 Breakdown of the Light-Dark Range Asymmetry due to Saturation***

Practical physically realizable sensors will saturate beyond some range of incident light intensity, at both the low and high ends of the sensor's range. The saturation on low end implies that the sensor will be insensitive to scene brightness changes below a certain level. The saturation on the high end, however, will not remove all sensitivity to brightness changes. This is because, in a scene which contains shadowing, or a range of surface albedos that includes zero albedo surfaces, there will be parts of the scene which result in sensor responses below the high-end saturation limit. Figure 2 shows the histogram of sensor values for different scene illuminances given an assumption of uniform distribution of object albedos. We can see the breakdown of the light-dark asymmetry due to saturation. At very low scene illuminances the histogram contains a single impulse at the minimum response value of the sensor. As the scene illuminance increases, some of the values rise above the minimum level, up to a value that scales with the scene illuminance. The height of the impulse at the minimum level drops as fewer sensor responses are below the minimum value. As the scene illuminance increases further, some of the incident light has an intensity above the sensors high-end saturation level. Thus an impulse at this level begins to form. As the scene illuminance increases further, there will always be some responses in the operational range of the sensor but these will become a smaller and smaller fraction of the total. Thus the histogram becomes more and more concentrated in the impulse at the high-level saturation value. Thus we can see that the

histogram for the very high and very low scene illuminances are symmetric. It is only for intermediate illuminances, where the sensor does not saturate significantly, that the light-dark range asymmetry is present.



**Fig. 2** Sensor saturation causes a breakdown of the light-dark range asymmetry at the extremes of scene brightness.

An automatic gain control, such as that provided by the pupil in the human eye, can extend the range of validity of the light-dark range asymmetry. A perfect gain control would seem to obviate the possibility of brightness perception, since the sensor response would always be the same. However, the gain control signal itself can be used as the brightness measure since the gain control mechanism will necessarily exhibit a light-dark asymmetry. Even if one inverted the sensor signal, the

gain control signal would not be inverted (e.g. a camera's aperture would still need to be closed down as the scene illuminance increased).

## 2.2 Other Asymmetries


The light-dark asymmetry is not the only sort of physical asymmetry that permits differentiating between light and dark. There are also important sensori-motor asymmetries (such as what happens when you close your eyes, or turn off lights, or occlude objects) that can be used to distinguish between dark and light scenes. Shadowing is an asymmetric process. Black patches are often found in shadowed areas, whereas white patches are rarely found there. Specular highlights are very bright compared with other areas, and never darker. A strong asymmetry arises through surface inter-reflection. For example, white patches can illuminate nearby patches, while black patches do not. Langer [2] points out that shadows and inter-reflections are in some sense symmetric with each other, as an intensity inversion transforms shadows into areas that look like inter-reflections and vice-versa. The symmetry is not exact, however, and the shadows and inter-reflections that are produced are often unlikely to be observed. There are other reasons for the lack of an exact symmetry. One reason is that all white surfaces illuminate nearby objects while only some black surfaces are shadow regions. Another is that the white patches 'cause' the illumination of nearby surfaces, while shadows are caused by other surfaces. So the intensity inversion must also imply a causal inversion, as the shadow regions now become illuminating regions and vice-versa.

In color images, there are additional asymmetries to be found. As Myin [3] points out, color is just a multidimensional intensity measure, and the asymmetries associated with intensity transfer to color as well. A commonly considered transformation is spectral inversion. There are many forms of this, but the most common is the independent inversion of each channel in an RGB image (e.g.  $R' = R_{max} - R$ ,  $G' = G_{max} - G$ ,  $B' = B_{max} - B$ ). White/Black patches are desaturated, and this persists under spectral inversion. Mid-tones are often highly saturated, and this also persists under spectral inversion. Shadow areas are always desaturated, while illuminated areas can be highly saturated. This asymmetry is reversed through spectral inversion, as dark areas appear colored and light areas (which now correspond to shadowed or dark areas in the world) appear desaturated. Thus low saturation values can indicate shadowed areas in the real-world, no matter whether the RGB values are inverted or not.

## 3 Statistical Measures of Scene Brightness

Figure 2 suggests that one could obtain a measure of brightness by looking at statistics of the sensor response distributions. In the range over which the light-dark range

asymmetry exists, as the scene illuminance increases the sensor response distribution becomes more spread out. There are many difference statistics that could be used to capture this spreading out. For example, one could use the variance of the distribution or its entropy. Figure 3 shows the entropy of the sensor value distribution of the situation associated with figure 2. It can be seen that the light-dark range asymmetry results in a rising entropy value as long as the scene illuminance is relatively low. Beyond a certain point the high-end saturation of the sensor comes into play and begins to reduce the entropy with further increases in scene illuminance.




**Fig. 3** Entropy of the distribution of sensor values as a function of scene illuminance for a simple scene having a uniform distribution of object albedo. The effect of the light-dark range asymmetry is evident as well as its breakdown at the extremes of scene brightness.

So far we have been considering global measures applicable to entire scenes. We could narrow our focus to look at small scene or image patches and ask whether we can find measures of patch brightness that are in some sense invariant to the specifics of the sensing process. One extension of the ideas discussed earlier is to apply the statistical measures such as variance or entropy to small patches in the image. The idea here being that bright patches would have a higher contrast measure (such as variance or entropy) than dark patches.


There is some psychological evidence for such an approach. In a study that produced the effect bearing his name, Stevens [4] found that subjects viewing a gray patch in a white surround perceived the contrast between the patches to increase as the intensity of the illumination (see figure 4 for an example of this effect). The background brightness was perceived to increase via a power law, with exponent 0.33, with respect to its luminance. The brightness of the gray patches, on the other hand, had a variable exponent, which became negative for darker patches. Overall,

the effect is that the perceived contrast increased with the illumination intensity. Hunt [5] observed an analogous effect in the perception of colored patches - he found that as overall intensity increased so did the perceived colorfulness.



**Fig. 4** The Stevens Effect: Shown are image patches with constant contrast and increasing mean intensity. Human observers usually perceive the contrast to increase along with the mean intensity of the patch.

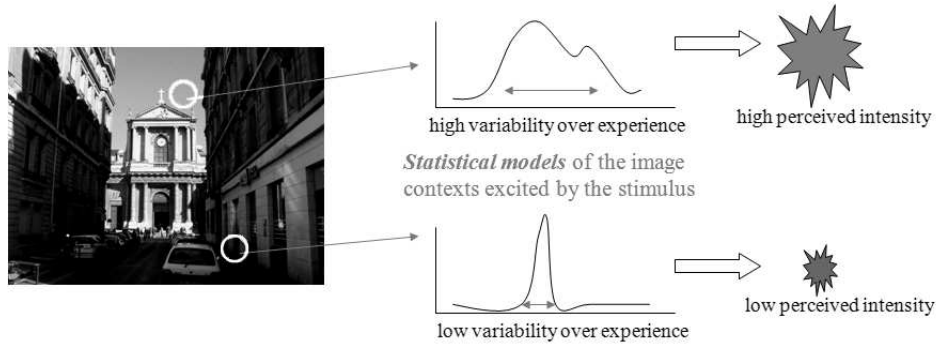
It has long been informally conjectured that a form of “inverse-Stevens effect” exists. That is, *perceived intensity increases with image contrast* (see figure 5 for an example of this effect). As Fairchild [6] points out, photographers often underexpose a high contrast scene (e.g. a dim indoor scene) and overexpose a low contrast scene (e.g. a bright outdoor scene). Fairchild did a psychophysical study to investigate this conjecture [6]. His results were inconclusive, however, showing a wide inter-subject variability. Some subjects had the supposed contrast-intensity relation while others had no relation, and still others had a relation in the direction opposite to that supposed.



**Fig. 5** Inverse Stevens Effect: Shown are image patches with constant mean intensity and increasing contrast. Some observers perceive the intensity to increase along with the contrast.

A straightforward implementation of this concept would be to measure the contrast (or entropy) of the histogram of pixel values in an image patch. In general, however, doing this produces only a weak dependence on intensity, and mainly produces a result similar to an *edge detection* operation. Indeed, entropy has frequently been employed as a feature in edge detection systems (e.g. [7]). In addition, as illustrated by figure 6, there can be instances of image patches that are bright but have low contrast or patch variability and others that are dark but have relatively high patch variability. To remedy these problems, we propose that the variability

(entropy) of learned *conditional distributions* based on many image patches observed over time should be used. The idea is that, while in a given image patch there may only be a loose correlation between patch brightness and patch variability, a stronger correlation may be observed over a large database of image patches associated with a particular central image intensity. Thus, given a pixel or small image region with a particular intensity value, the entropy of the learned distribution of pixel values conditioned on this immediate value can be used as a measure of patch brightness.



**Fig. 6** Simple image patch variability measures cannot be used to measure brightness since some bright patches can have low patch variability and some dark patches can have relatively high patch variability. Instead, the variability of previously experienced image patches associated with a given central value can be used.

This suggests an explanation of the Stevens effect. The idea is that, through visual experience, an observer learns an association between surface brightness and the entropy of the surface patch intensity values. Our thinking is motivated by the ideas of Norwich [8] who suggests that perception arises through reduction of uncertainty. In his view, a more intense stimulus has more uncertainty, and hence higher entropy. Furthermore, he proposes that the *subjective* impression of the intensity of a stimulus is related to the entropy of the stimulus. This leads to the hypothesis that the subjective impression of increased contrast with brighter images that comprises the Stevens Effect is a result of a learned association between contrast and entropy. That is, high contrast patches in natural images will statistically tend to have higher entropies than low contrast image patches.

#### 4 Surround Entropy in Natural Images

To test our hypothesis that patch brightness can be related to the entropy of conditional distributions, we carried out an empirical study of the conditional statistics of surrounds in a database of natural images. For our study we used a set of 136 images



from the van Hateren database [9]. Each of these images had a size of 1024x1536 pixels. Figure 7 shows four of the images that were used in our study.




**Fig. 7** Samples of the images used in the empirical study. These images were taken from the van Hateren image database (van Hateran and van der Schaaf, 1998).

The raw image values were scaled by calibration factors provided with the database. These factors account for variations in sensitivity caused by aperture settings and shutter speeds, and permit us, to some extent, to compare intensities across images in the database. Details of the image acquisition can be found in [9]. The image pixels we used were 12-bits each, with a linear intensity scale. These were obtained from the 8-bit van Hateren images which were compressed using a non-linear quantization scheme. The 12-bit linear pixels were obtained using the look-up table provided by van Hateren. We smoothed the scaled images with a 5x5 averaging kernel before computing entropies. This is to remove the residual effects of the non-uniform quantization scheme which would otherwise create a relative decrease in entropy with intensity, due to the greater spread between quantization levels at high intensities than at low. We also eliminated images which exhibited noticeable saturation at the high intensity end, as indicated by examination of the images' intensity histograms. Saturation results in an excessive number of pixels having the same value, which would reduce the compute entropy values, especially at high in-

tensity levels. Although the images that were used in the study did not appear to exhibit any saturation, examination of the conditional histograms (in figure 8) show that there is still at least a low level of saturation, as indicated by the blip on the high end of the highest intensity conditional histogram.

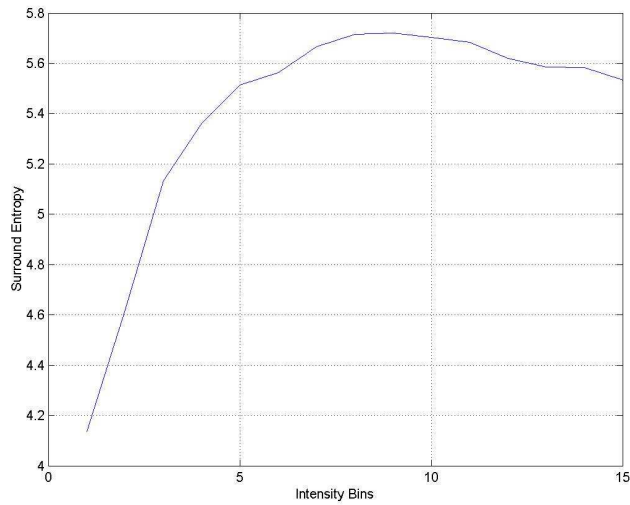
To construct the conditional histograms for a given central value  $I$ , we searched the database images for pixels with values in the range  $[I, I + \Delta I]$ . Then we computed the histogram of the pixels in an  $11 \times 11$  neighborhood centered on these pixels. These individual surround histograms were then summed to give the overall conditional histogram for the value  $I$ .



**Fig. 8** The conditional histograms of the surround pixel values given the central pixel value (for 15 different ranges of central pixel values).

The conditional entropy is shown in figure 9. It is seen to rise, almost linearly, for low intensities, and then flatten out, and finally to drop once again.

The curves shown in figures 3 and 9 have a similar shape, and it is tempting to claim that the empirical result can be completely explained by the light-dark range asymmetry with sensor saturation. However, as mentioned earlier, we deliberately omitted images from our test set which exhibited noticeable sensor saturation, and examination of the conditional distributions shown in figure 8 reveals very little saturation, if any. The conditional variation of entropy must be due, at least in part, to other effects. Our view is that the situation is complicated, as there are many factors which act to determine the surround distribution and hence its entropy. Not all of these are dependent on the surround brightness. It may be that the human visual system is able to factor out these various contributions and isolate those that are related to intensity. The paper [10] developed statistical models for the surround



**Fig. 9** The entropy of the conditional distributions of surround values given the central value, for 15 different central value ranges.

using a Maximum Entropy approach and looked at the effects of three ecological processes: shadowing, occlusion, and inter-reflection. These processes all introduce asymmetries with respect to brightness which can act together with the light-dark range asymmetry to shape the dependence of patch brightness with patch variability.

It is an open question whether humans can adapt to changes in the sensing apparatus which do not alter the physical asymmetries we have discussed. For example, one could photometrically invert the image presented to the visual system. Such an adaptive capability would be predicted by our theory. However, at least one study has been done which shows that humans are not able to re-invert inverted intensity, at least not without a long time for adaptation [11]. In particular, detection of shadows seems to rely crucially on the shadows being darker than the illumination regions [12]. This is not a fatal blow to our theory as it could be that to adapt to such inversions requires an extensive learning process in order to develop a new set of conditional statistical models.

A recent study Mante et al [13] used the van Hateren database and found no correlation between patch intensity and contrast. The patches they used were quite large, however, being circular with a diameter of 64 pixels. Also, they only looked at correlations of luminance and contrast within individual images, and then averaged these correlations across all of the images in the database. Our results are based on small patches and statistics were gathered over all images.

## 5 Summary

This paper suggests that various physical asymmetries, such as the light-dark range asymmetry, provide a means for measuring scene brightness that is invariant to some manipulations, for example inversion, of the sensor responses. We propose that brightness perception is mediated, not by the level of the raw photoreceptor signals, but by a statistical learning of the association between a particular localized sensor signal and the variability of a learned distribution of surround sensor values conditioned on the localized sensor signal value. Following the ideas of Norwich we propose that humans learn to associate the entropy of a stimulus with the intensity of that stimulus. We examined a small set of natural images and observed that there was, indeed, a relationship between an image pixel value and the entropy of the distribution of surrounding pixel values which would explain the Stevens effect - the perceived increase of contrast as scene intensity increases.

## References

- [1] Hartline, H. K. (1938), "The response of single optic nerve fibers of the vertebrate eye to illumination of the retina", *American Journal of Physiology*, 121, 400-415.
- [2] Langer, M. S. (1999), "A model of how interreflections can affect color appearance". Technical Report No. 70, Max-Planck Institute for Biological Cybernetics, Tuebingen, Germany.
- [3] Myin, E. (2001), "Color and the duplication assumption". *Synthese* 129: 61-77.
- [4] Stevens, S.S. (1961), "To Honor Fechner and Repeal His Law". *Science*, 133, 13 Jan., pp 80-86.
- [5] Hunt, R.W.G. (1995), **The Reproduction of Color**. 5th edition, Kingston-upon-Thames, England: Fountain Press.
- [6] Fairchild, M.D. (1999), "A Victory for Equivalent Background – On Average" IS&T/SID 7th Color Imaging Conference, Scottsdale, 87-92.
- [7] Shiozaki, A. (1986), "Edge extraction using entropy operator", *Computer Vision, Graphics, and Image Processing*. Vol. 36, pp. 1-9. Nov. 1986.
- [8] Norwich, K. (1984), "The psychophysics of taste from the entropy of the stimulus", *Perception and Psychophysics*, 35 (3), 269-278.
- [9] van Hateren, J.H. and van der Schaaf, A. (1998), "Independent component filters of natural images compared with simple cells in primary visual cortex". *Proc.R.Soc.Lond. B* 265:359-366.
- [10] Clark, J.J. and Hernandez, D. (2003), "Surround statistics and the perception of intensity and color", 3rd International Workshop on Statistical and Computational Theories of Vision, Nice, October 2003.
- [11] Anstis, S.M. (1992), "Visual adaptation to a negative, brightness-reversed world: Some preliminary observations", in: **Neural networks for vision and image processing**, G. A. Carpenter, S. Grossberg, Eds. MIT Press, Cambridge, MA., pp. 1-14
- [12] Cavanagh P. and Leclerc Y.G. (1989), "Shape from shadows", *J Exp Psychol Hum Percept Perform*, 15(1):3-27.
- [13] Mante, V., Bonin, V., Geisler, W.S., and Carandini, M. (2005), "Independence of luminance and contrast in natural scenes and in the early visual system", *Nature Neuroscience* 8, 1690 - 1697.