
The visual light field

Jan J Koenderink, Sylvia C Pont, Andrea J van Doorn[¶], Astrid M L Kappers, James T Todd[§]

Helmholtz Institute, Utrecht University, Princetonplein 5, NL 3584 CC Utrecht, The Netherlands;

[¶] also Delft University of Technology, Julianalaan 134, NL 2600 AA Delft, The Netherlands;

[§] Department of Psychology, Ohio State University, 1827 Neil Avenue, Columbus, OH 43210, USA;

e-mail: j.j.koenderink@phys.uu.nl

Received 19 July 2006, in revised form 22 February 2007; published online 31 October 2007

Abstract. Human observers are sensitive to the ‘(physical) light field’ in the sense that they have expectations of how a given object would appear if it were introduced in the scene in front of them at some arbitrary location. Thus the ‘visual light field’ is defined even in the ‘empty space’ between objects. In that sense the light field is akin to visual space considered as a ‘container’. The visual light field at any given point can be measured in psychophysical experiments through the introduction of a suitable ‘gauge object’ at that position and letting the observer adjust the appearance of that gauge object (eg through suitable computer rendering) so as to produce a ‘visual fit’ into the scene. The parameters of the rendering will then be considered as the measurement result. We introduced white spheres as gauge objects at various locations in stereoscopically presented photographic scenes. We measured the direction (‘direction of the light’), diffuseness (‘quality of the light’ as used by photographers and interior decorators), and intensity of the light field. We used three very different scenes, with very different physical light fields. The images were geometrically and photometrically calibrated, so we were in a position to correlate the observations with the physical ‘ground truth’. We report that human observers are quite sensitive to various parameters of the physical light field and generally arrive at close to veridical settings, although a number of comparatively minor systematic deviations from veridicality can be noted. We conclude that the visual light field is an entity whose existence is at least as well defined as that of visual space, despite the fact that the visual light field hardly appears as prominently in vision science as it does in the visual arts.

1 Introduction

The naive observer takes it for granted that the visual world is a perspectival clone of the physical world (Gibson 1966). Visual objects are intentional (they relate to the world) because perception is firmly rooted in what has been called ‘the background’ (Searle 1983) or ‘frame’ (Minsky 1974). The backbone of this pre-cognitive situational awareness is composed of several (interrelated) generic frameworks. Examples of such generic frameworks are the chronogeometrical and the radiometric frameworks—the former an awareness of being in time and space, and the latter an awareness of the luminous environment. The latter also includes an awareness of the optical properties of the medium (‘atmospherical perspective’—Minnaert 1968; Tricker 1970; Abrahams and Kattenfeld 1997; Nayar and Narasimhan 1999), of the spectro-radiometric properties of the scene (‘colour constancy’—Land 1959) and of the ‘light field’ (the basis for ‘shadow’, ‘shading’, ‘highlight’, and so forth—Adelson and Pentland 1996; Gilchrist 1999). Of these, the light field has been studied least extensively, and most often in the context of computer graphics (Foley and Van Dam 1983) or radiometry (Gershun 1936), but hardly in psychophysics (Langer and Bülthoff 2000; Dror et al 2004; Koenderink and van Doorn 2004). Especially work on ‘lightness’ and ‘constancy’ is somewhat related (Coren and Komoda 1973; Kozaki and Noguchi 1976; Bergström 1977; Gilchrist and Jacobsen 1984; Noguchi and Kozaki 1985; Brainard 1998; Ikeda et al 1998; Robilotto and Zaidi 2004; Todd et al 2004; Zavagno 2005). It is not that the light field is of minor importance though. Whereas the chronogeometrical framework has

to do with the *where and when* of visual objects and events, the radiometric framework has to do with the *what and how* questions, that is to say, with material constitution, surface corrugations, and appearance, as opposed to location and shape. The light field has figured in the visual arts (Hoogstraeten 1678; Adams 1950; Hogarth 1981; Jacobs 1986; Baxandall 1995). This paper is about the light field in psychophysical context. We consider the problem of how to measure the *visual* light field (as opposed to the *physical* light field—Gershun 1936; Moon and Spencer 1981) and we present psychophysical results.

The visual light field is due to the *chiaroscuro*, thus revealed by visible objects which are perceived as illuminated by the visual light field (Schöne 1979). Whereas the visual light field is perceived through objects, these objects themselves are parsed in terms of the light field. In computer vision, the estimation of the sources and the shape-from-shading problem are generally treated separately (Brooks and Horn 1989; Pentland 1990; Kersten and Yuille 2003), but one necessarily perceives objects of such-and-such a make, of such-and-such a shape, illuminated in such-and-such a way. The perception comes as a packaged deal, so to speak. If one relation is misperceived, the others will in all likelihood be misperceived too, and probably be so in a systematic fashion. That one indeed perceives the light field is evident from the fact that perceivers have implicit expectations concerning the appearance of objects as they move or are introduced in the scene. This can be shown through the introduction of objects that fail to ‘fit’ the visual light field: they are often immediately and spontaneously perceived as alien to the scene (King 1997). The nexus of these expectations defines what we mean by ‘the visual light field’. Thus the visual light field does not just adhere to the objects, but is equally defined in the empty space between the objects. In that sense the ontological status of the visual light field is akin to that of visual space as a ‘container’.

In the art of painting there exist numerous instances of *rendered light fields*. Explicit examples are works by Caravaggio (1600a, 1600b) and his followers, eg the painters of the Utrecht School (Spicer and Orr 1997). Common scenes involve groups of persons illuminated by a candle within the group, the source itself often being occluded and thus not visible. The way the persons have been painted defines the light field. Any part of the scene implies the location of the (invisible) source. Early representations of the *Nativity* (or *Adoration by the Shepherds*) often achieve a similar effect (Barocci 1597; Rembrandt 1646), the Child acting as a luminous source of divine light. Other common topics that involve the light field include the *Conversion of St Paul* (Caravaggio 1600a; in representations of this incident the illuminating beam is often shown as a ‘force’ that throws the apostle to the ground), *Danaë* (Rembrandt 1636–1647; in representations of this scene the illuminating beam is identified with Zeus in the guise of a ‘golden rain’, often represented through gold coins), and various cases where a Teacher emanates a ‘beam of illumination’ (often ‘explained’ through a source behind this person) thus ‘enlightening’ (ie educating) the pupil (eg Caravaggio 1600b; Comenius 1658). The study of rendered light fields is likely to prove very rewarding, though relevant literature is scarce (except for special cases like cast shadow), Schöne’s (1979) treatise still remaining the main source.

2 Methods

2.1 The psychophysical task

Physics has developed various ways to measure the radiance (or plenoptic function), which is in fact nothing but a technical term for the physical light field (Gershun 1936; Adelson and Bergen 1991). Any text on radiometry lists numerous methods. But how could one measure the visual light field, which exists only as a mental entity?

One class of methods relies upon verbal judgments, or on visual indicators that do not belong to the scene. An example is Pentland (1982) who had observers indicate

the illumination direction via a probe sheet of drawings of a pointer in various spatial attitudes from which they had to pick the most appropriate one. This latter method at least avoids problems with verbal encoding of spatial attitude (eg in terms of a pair of numerically specified angles), though it still decouples the response space from the perceptual space.

We suggest that methods be based upon the generic notion of ‘gauge objects’: introduce a gauge object in the scene—or a picture of a gauge object in a picture of the scene—and let the observer judge its ‘fit’. That such methods might be viable is suggested by cases of spontaneous detection of inconsistencies in photographic montages (eg cases of falsification of historical documents in the former Soviet Union—King 1997). Human observers are quite sensitive to the degree of fit. One easily turns it into a psychophysical method by giving the observer control over the appearance of the test object and requiring the observer to adjust it such as to produce a satisfactory visual fit. The parameters that characterise the appearance represent the result of the measurement. Such methods are not conceptually different from the measurement of spatial extent via yardsticks, the yardstick being used as gauge object. Because the fit is purely visual, one obtains a direct operational method to quantify eye measure, which is exactly what we are after. The methodology is generic and applicable in virtually any perceptual context (Koenderink et al 1992).

We implemented this idea in the simplest possible way. A scene was presented via a pair of stereoscopic images. This allows us to introduce gauge objects at various fiducial positions within the pictorial volume by way of position in the image and disparity. As gauge object we selected a Lambertian (Lambert 1760) sphere, illuminated by a parameterised beam. The parameters considered were the direction of illumination (two angles), the diffuseness of the beam, and the maximum illuminance, in total four degrees of freedom. We decided to refrain from the addition of an ‘ambient term’, because this additional degree of freedom turns out to interfere with the other parameters in a way that renders the interface confusing to the observers. The beam was considered spatially uniform, and, owing to a uniform circular disk source of specified angular subtend, located at infinity. The diffuseness ranges from a fully collimated beam (approximately sunlight) to a Ganzfeld (fully diffuse beam). For intermediary values one has directional but diffuse beams. The hemispherical diffuse beam (eg to a good approximation overcast sky) assumes a central position. In figure 1 we illustrate the nature of the degrees of freedom associated with this probe.

We ignore interactions between the gauge objects and the scene. With one exception (discussed later) objects in the scene throw no cast shadows upon the gauge object and vice versa, nor do they cause reflexes upon each other. In this sense the scene with the gauge object in it is necessarily inconsistent. We do not consider this to yield serious problems (see below), but one should keep it in mind. We do this in order to not disturb the original scene through the introduction of the gauge object. The approximation is better if the gauge object is smaller; on the other hand a larger probe is easier on the observer and may well lead to better performance in this task. Such trade-offs are typical for measurement in the sciences (Kohlrausch 1923; Feynman et al 1965).

2.2 *Implementation and instructions*

The description of the task to the observers was simply “make the test sphere appear like it fits into the scene”. The observers performed the task by adjusting slider controls presented outside the main scene (see below). The observers either looked at the scene while (typically slightly) adjusting a slider, or looked at the scene before and after (in this case often large) adjustments. They kept adjusting the slider controls until they judged the fit of the test object to the scene to be satisfactory.

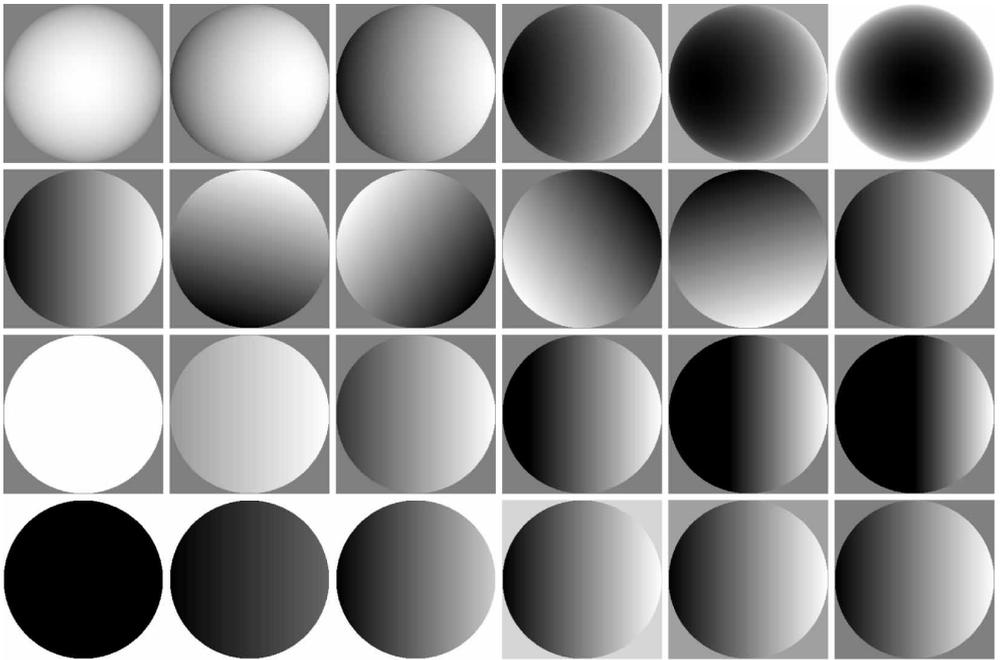


Figure 1. The degrees of freedom of the probe used in the experiment. The observer controls slant, tilt, directedness, and intensity. The rows show variations of slant, tilt, directedness, and intensity (from top row to bottom row). For all but the intensity variations, intensity was set to 1 (maximum); for all but the directedness variations, directedness was set to 0 (hemispherical diffuse); for all but the slant variations, slant was set to 90° ; and for all but the tilt variations, tilt was set to 0° . All parameters are varied over their full range: slant from 0° (frontal illumination) to 180° ('contre jour'), tilt from 0° to 360° , directedness from -1 (Ganzfeld) to $+1$ (collimated beam), and intensity from 0 (black) to 1 (white).

2.3 Ground truth

We set up the scenes in a large studio with walls painted black. This allows maximum control over the light field, but even so the photometric interactions within the scene introduce many complications. The stimuli themselves were straight photographs taken with the usual precautions to guarantee well-calibrated results. These photographs, of course, include the effects of all photometric interactions. It would have been next to impossible to arrive at this degree of realism with computer graphics techniques, although this would certainly have been more convenient.

In order to quantify the ground truth we used two different methods. The first method relies upon standard photometric techniques. We used both (calibrated) irradiance and (1° receptance angle) radiance meters to quantify basic photometric parameters and used extensive position data (obtained with measuring tapes, etc) to describe the scene geometry. We also measured the reflectance data for the various surfaces in the scene. This allows us to estimate the light field at any point with reasonable precision. These methods are standard in illumination engineering as used in interior architecture.

The second method is a direct one. We photographed the scenes both without the gauge objects (artificial gauge objects being introduced into these stimuli in the actual experiment) and the same scenes with actual gauge objects present. These gauge objects were white, Lambertian spheres. Because the gauge objects are small as compared to the dimensions of the scene they hardly perturb the light field. Thus the photographic records of these objects can immediately be used as ground truth.

We checked the results of the two methods and found excellent mutual agreement.

When confronting experimental data with the ground truth it is sometimes preferable to use one or the other set. The first method lets us immediately compare direction, diffuseness, and intensity, whereas the second method is quite independent of these parameters and can be used to check to what extent different parameter settings approximate the pictorial data.

In the experiment, the observers rely upon the pictorial data, although a comparison of objects at different locations in a scene forces them to relate such pictorial data at disparate areas of the picture. In a scientific description, such relations would have to be expressed in terms of the structure of the light field (direction of beams in 3-D space, and so forth). It is not at all obvious what human observers might do in such comparisons; this is what the experiment is about.

Notice that the ground truth is not necessarily implied by the stimulus. For instance, given an image, there exists an infinite set of possible scenes and light fields that might have given rise to that image (Belhumeur et al 1999). For the case of our experiment we do not know the full set of (equally veridical) interpretations. This is not necessarily a problem in the context of the experiment, though. We ask only for a visual fit, but not for scene geometry and so forth. The bas-relief ambiguity (Belhumeur et al 1999) would conserve such a fit; thus the response would be valid for any of the infinite possible interpretations. A comparison with the ground truth remains meaningful. One should not interpret this in an overly restricted manner though.

3 Experiment

3.1 Design

We used three pictorial scenes (see figure 2) and measured the visual light field in about half a dozen locations in each scene (see figure 3).

Eight observers participated in the experiment (three authors—AD, JK, and SP—who are experienced observers, though novel in this type of experiment, and five naive, paid observers—FW, JJ, KF, LIM, and LOM). Observers ranged in age from early twenties to early sixties. Half of the observers were female, half male. The acuity (with proper correction) of all observers was normal. All enjoyed normal binocular stereopsis as checked via a formal procedure.

Each observer repeated all measurements four times (in four sessions). The sequence was randomised per session. No feedback was offered until after the conclusion of the experiment.

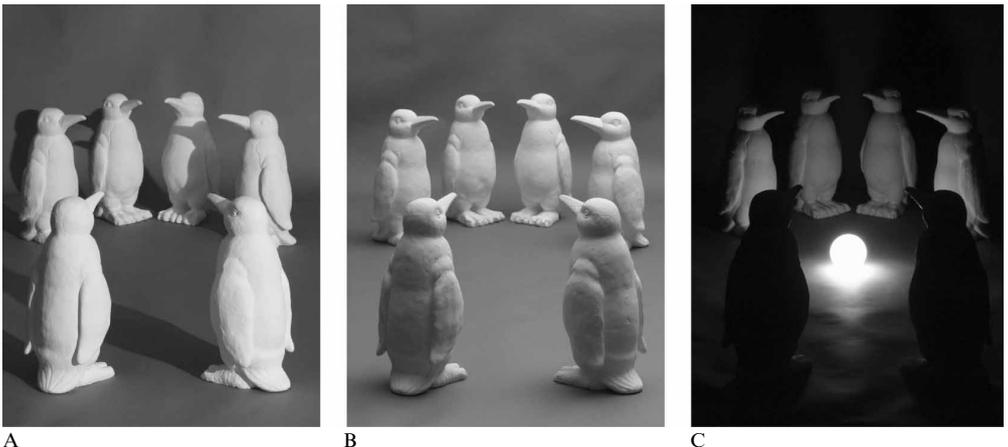


Figure 2. The three scenes. Notice that scene A is like a generic open scene in sunlight; scene B, a generic open scene under an overcast sky; and scene C resembles a generic *Nativity* (or *Adoration of the Shepherds*) painting. Notice differences in direction of illumination, contrast due to shading gradients, and the nature of body and cast shadow edges.

The stimuli (figures 2 and 3) were stereoscopic photographs of scenes set up in a large studio. The objects in the scene were six identical puppets painted matte-white. The background was a uniform gray (about 50% reflectance) paper. The objects were placed

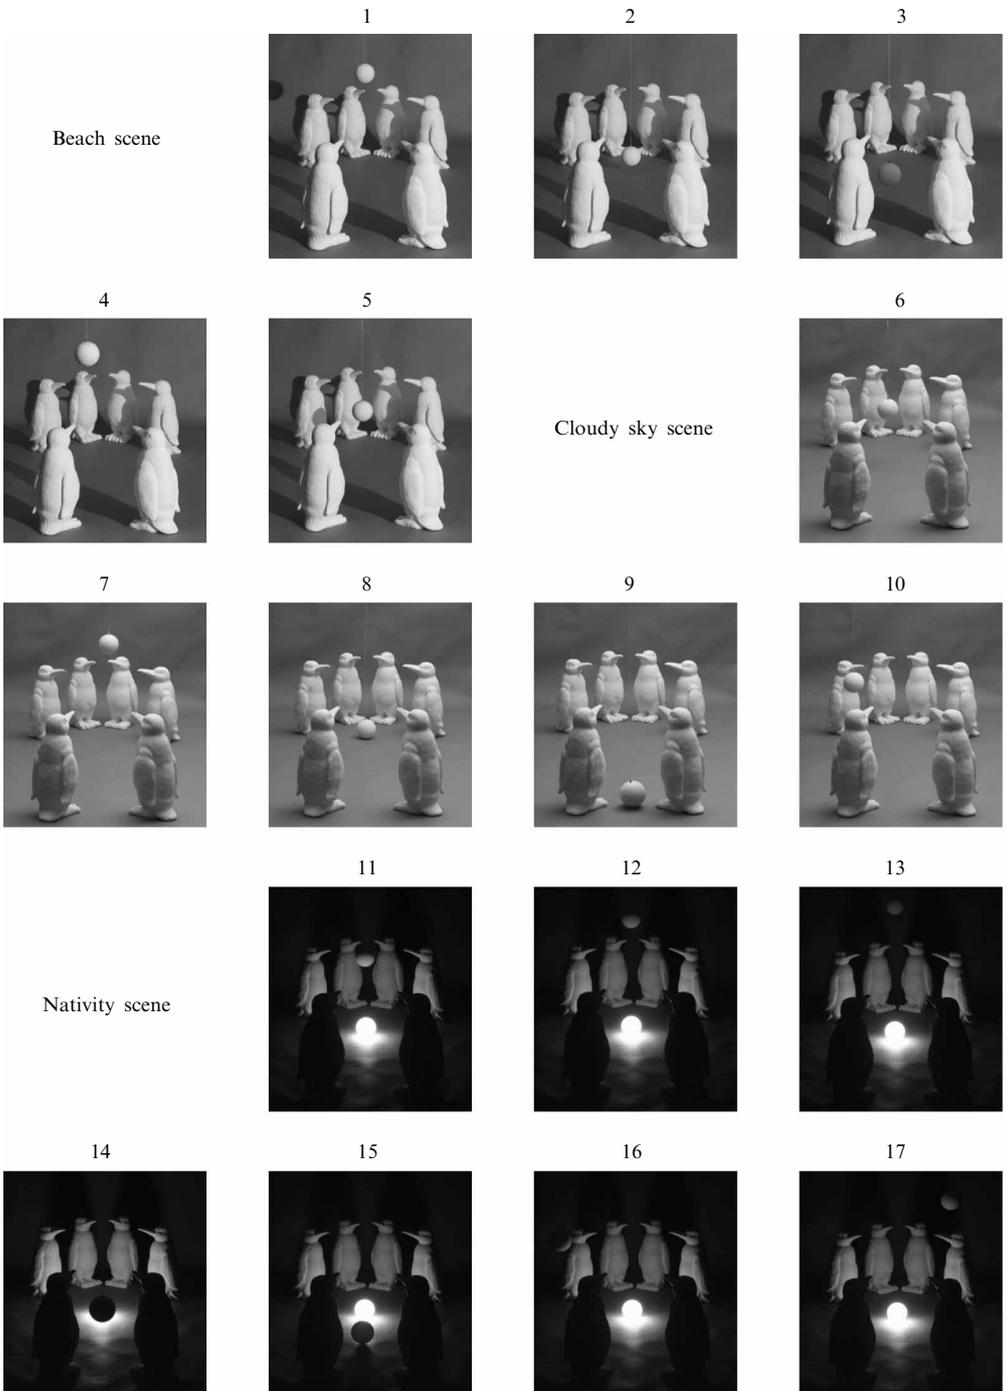


Figure 3. The stimuli are composed of a scene (there are three scenes) and a location in the scene. Here, white spheres are photographed at their respective fiducial locations, yielding the actual ‘ground truth’.

such as to appear positioned in roughly circular fashion. Care was taken to avoid occlusion of too much of the volume of the scene. Three scenes were designed (see figure 2):

(A) The only light source is a small spot at fairly large distance. The inverse-squares law is rendered ineffective because all puppets are at roughly the same distance from the source. All puppets are similarly shaded. Cast shadows are prominent and have sharp edges. This is a typical daylight (bright sunlight) scene.

(B) The scene is illuminated by an extended source vertically above the scene. All puppets are similarly shaded. Cast shadows are very fuzzy and hardly apparent. This is a typical daylight scene under a (heavily) overcast sky.

(C) There is only a small source (a frosted light bulb) at the centre of the group. All other illumination is due to scattering within the scene. Illuminations are roughly determined through the inverse-squares and Lambert's cosine laws. The puppets in the front appear as silhouettes, the puppets in the background appear frontally illuminated. This is the typical 'Nativity' or 'Merry company' scene as known from art history.

In scenes A and B we defined five, in scene C seven, fiducial locations. The direction of the beam varied enormously for the locations in scene C; in some cases the gauge figure appeared 'a contre jour' (thus as silhouette). In the cases of scenes A and B the direction of illumination varies only little. The degree of collimatedness is high in scenes A and C, low in scene B. In scene A, one location was chosen inside the volume cast shadow of a figure.

In the text we refer to the stimuli by index (1, ..., 17) or by scene and subindex [A1–A5 (= 1, ..., 5), B1–B5 (= 6, ..., 10), C1–C7 (= 11, ..., 17)], whichever is most convenient. In any case, figure 3 is convenient for looking up the stimulus of interest.

The photographs were taken with an electronic camera and rendered on a linearised display. A mirror stereoscope with precise geometrical controls was mounted rigidly in front of the CRT. Intraocular distance and ophthalmic corrections were individually adjusted.

The interface consisted of four slider controls (slant, tilt, collimatedness, and intensity) presented monocularly at the bottom of the screen. Observers used either a trackball or a mouse to move the slider controls. At the initiation of each trial the slider controls were set to random positions, though avoiding such cases as the Ganzfeld or total darkness.

3.2 Results

Observers had little trouble with the task and took about a minute per setting. They reported that the gauge object was seen by them

- as part of the scene,
- at a well-defined location in the volume,
- and (after suitable parameter settings) illuminated by the overall light field.

It is of course crucial that the probe be perceptually accepted as belonging to the scene. The very method depends upon it. This even happens when the illumination of the gauge object is judged to be wrong. Apparently the binocular stereo cue is sufficient to make the gauge objects belong.

The probe location in scene A3 that was in a body shadow led to an interesting result: only one of the observers detected the fact. All observers were satisfied with a setting that was close to the light field as it would be in the absence of the body shadow (see figure 4).

Apparently observers have a notion of the structure of the light field, even at locations in empty space, quite remote from any visual object. One may say that the light field in a scene is indeed perceived. That the volume cast shadow region failed to be detected at all (although the corresponding body shadow and cast shadow on the floor were both clearly visible) is a striking exception, though. This finding reminds

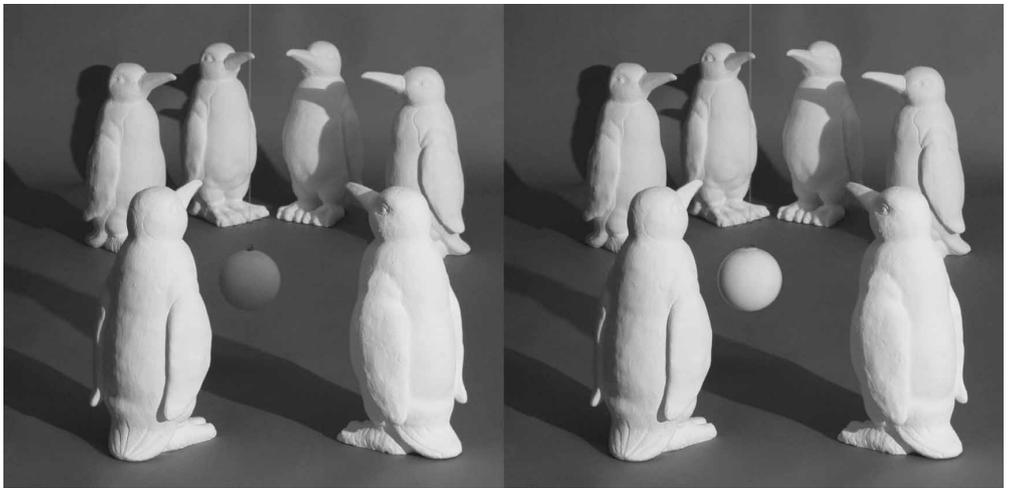


Figure 4. (a) Beach scene A3 with a gauge figure in the volume shadow of the right frontmost puppet. (b) The same scene as it is accepted by the observers. Notice that the physically wrong picture (b) is visually not less immediately acceptable than the actual one (a). (Of course this is less remarkable in this monocular setting than in the actual stereo rendering.) Such examples help to understand the frequent objections against Rembrandt's 1642 treatment (in the famous *Nightwatch*) of the shadow of the captain's hand on the uniform of the lieutenant as a mistaken display of the painter's bravura (Arnheim 1956).

one of the well-known fact that cast shadows are often omitted, or used in an idiosyncratic fashion, in the visual arts (Baxandall 1995; Gombrich 1995; Stoichita 1997). Cast shadows may indeed be important in vision (Mamassian et al 1998; Tarr et al 1998), but their artifactual absence is apparently hardly noticeable.

In figure 5 we plot the R^2 values for all observers for the regression of the slant, tilt, intensity, and directedness settings against the veridical values. (The sequence of observers is in alphabetical order.) The concordance between observers is quite high, with the experienced observers reaching slightly higher values, but not dramatically so:

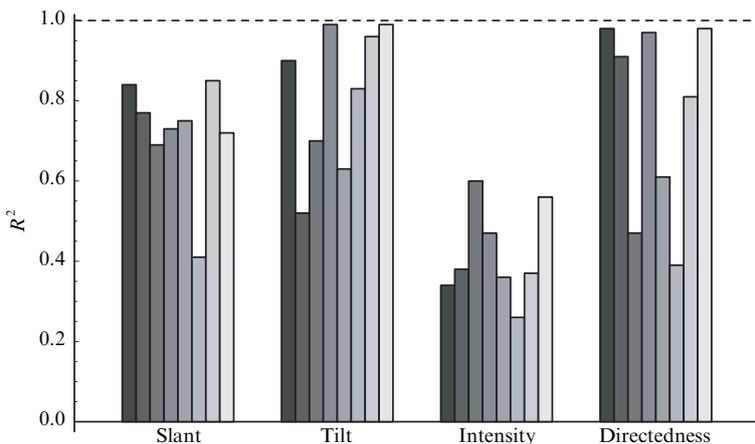


Figure 5. The R^2 values for all observers for the regression of the slant, tilt, intensity, and directedness settings against the veridical values. The sequence of observers is: AD, FW, JJ, JK, KF, LIM, LOM, SP. Notice the high interobserver concordance. The first principal component indeed explains 79% of the variance; thus these data are well summarised via the means: 0.72 for the slant, 0.82 for the tilt, 0.42 for the intensity, and 0.77 for the directedness.

the results of the naive observers are very similar to those of the experienced observers. In all cases the correlations are highly significant. Apparently the observers to a large extent produce settings that vary monotonically with veridical values. Of course, the R^2 values fail to reveal the extent to which these observations reproduce the veridical values quantitatively (but see below).

The parameters used in the probe are in no way ‘orthogonal’ (see figure 1), so it is somewhat difficult to judge the degree to which the observations approach veridicality. In order to address this problem we also did pixelwise regressions of the actual probe images to predicted probe images that were computed for the veridical settings. In this case the idiosyncracies of the parameterisation are immaterial. We refer to these as the ‘image-based correlations’. An example is shown in figure 6 for the (naive) observer JJ. In figure 7 we show the image-based correlations for all stimuli (please refer to figure 3) for all observers (again, in alphabetical sequence). For stimuli 14 through 16 not much correlation can be expected, since the probe is seen a *contre jour*. Indeed, as might be expected, the image-based correlations are quite low for these cases. For the remaining cases the R^2 values for the image-based regressions are quite high. This is the case for all observers, whether experienced in the art of visual observation, or naive. One observer (LIM) is consistently low, the experienced observers perhaps consistently somewhat (but not much) higher than the naive observers. Generic R^2 values for the image-based regressions are in the 0.7–0.9 range.

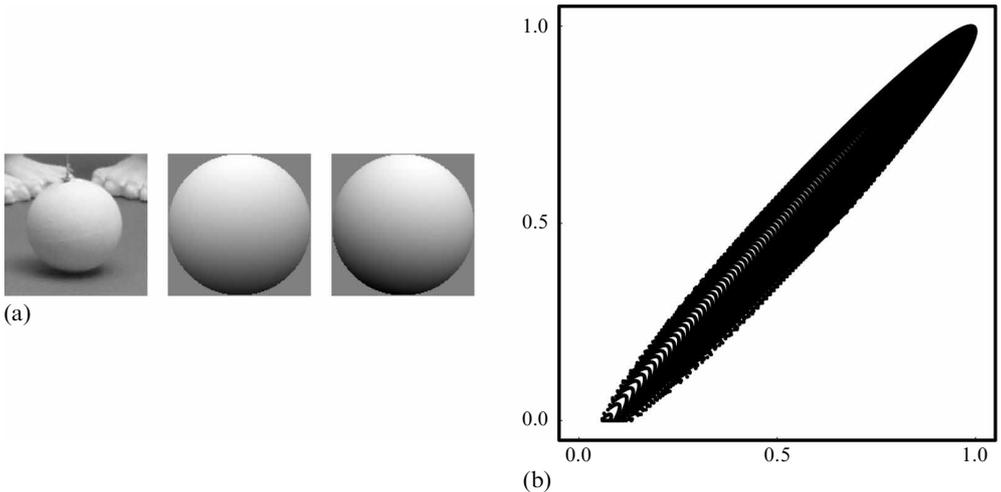


Figure 6. (a) Images of the probe in the case of the third ‘cloudy sky’ scene (stimulus 8). The leftmost image represents the empirical ground truth; it is the photograph of a white sphere introduced into the actual scene. The centre image is the probe rendered for the veridical parameters. Since the intensity range has been normalised, this image cannot be immediately compared to the empirical ground truth. We refer to this image as the theoretical ground truth. The rendering does not include an ambient term and no account is taken of interreflections within the scene, leading to minor differences with the empirical ground truth (the pixelwise correlation is very high though). The rightmost image represents a rendering according to the observer’s (JJ) settings. In this case the settings were rather close to veridical. (b) The graph shows a straight scatterplot of the image intensities of the observation image to the theoretical ground truth. In this case the R^2 value was 0.95.

With regard to the direction of the incident beam, we distinguish between the slant and the tilt, the tilt being the angle in the picture plane, the slant being the angle defined by the fore–aft relations.

The tilt was set very close to the veridical values throughout, by all observers. This need not be any cause for surprise; indeed, given the periodic nature of the tilt, there is no room for either offset or scaling. The only interesting measure is the scatter

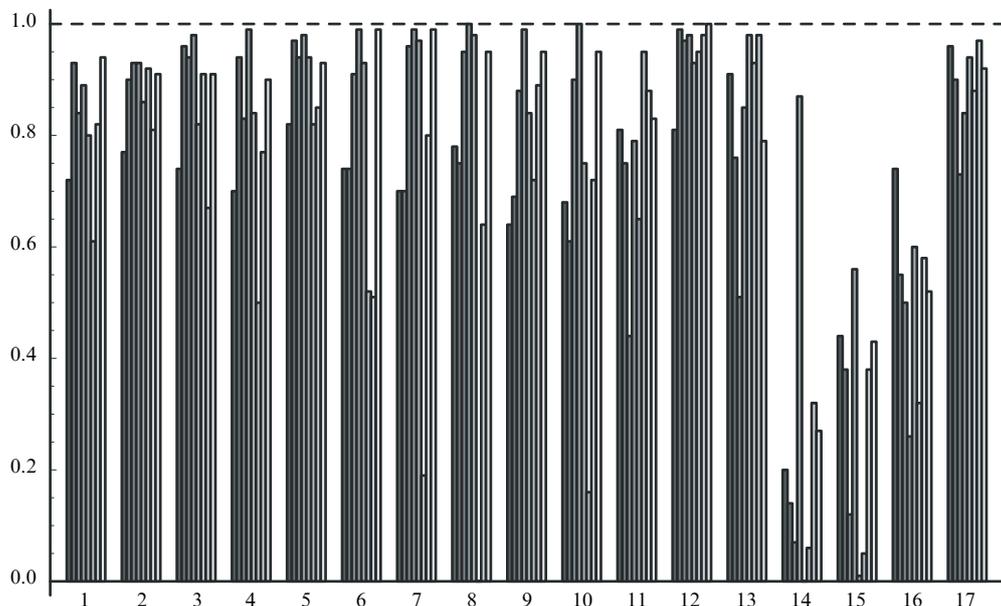


Figure 7. The image-based correlations for all stimuli (refer to figure 3) for all observers. The sequence of observers is again: AD, FW, JJ, JK, KF, LIM, LOM, SP. The first principal component explains 70% of the variance, the second one only 10%; thus almost all of these data may be summarised through the means. For the 17 stimuli the means are: 0.82, 0.88, 0.87, 0.81, 0.91, 0.79, 0.79, 0.76, 0.83, 0.72, 0.76, 0.95, 0.84, 0.24, 0.30, 0.51, 0.89; thus rather high, especially in view of the fact that the three lowest values should really be disregarded.

around the (essentially veridical) mean values. All observers reproduced the tilt within 5° or 10° , thus quite precisely. An example is shown in figure 8—the direction data for the (naive) observer JJ. These are quite typical results.

The slant is not periodic, but ranges between 0° (frontal illumination), over 90° (illumination from the side), to 180° (the contre jour, or backlighting situation). Thus there exists the possibility of offsets and scale compressions. In figure 9 we show the slant ranges for all observers. (These ranges have been calculated from the linear regression on the data.) Although some observers use most of the scale (and thus show close to veridical observations), most observers use less, or even considerably less, of the available range.

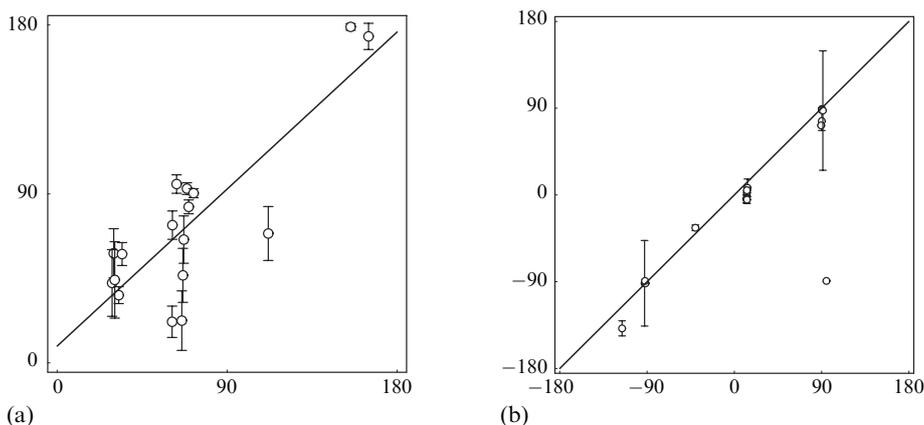


Figure 8. The regression of slant (a) and tilt (b) settings against the veridical values for the (naive) observer JJ. These plots include all data, also for those cases for which either slant or tilt was essentially undefined.

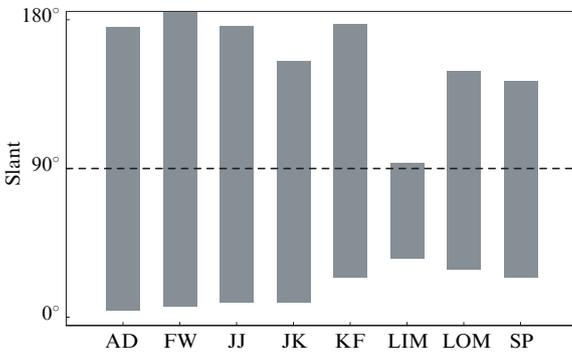


Figure 9. The total slant ranges for all observers. Notice how several of these ranges appear to contract towards the 90° level, representing a regression towards the frontoparallel (sideways illumination).

The settings of these observers tend to be centred about the illuminations from the side. We see a more or less pronounced regression towards the frontoparallel plane in the sense that both illuminated object and illuminating source are pushed closer to a mutual frontoparallel configuration. Such tendencies are also commonly encountered in studies of visual space.

The directedness of the illuminating beam varies from -1 (Ganzfeld) over 0 (hemispherical diffuse, akin to illumination by a heavily overcast sky), to $+1$ (collimated beam, such as direct sunlight). An example is shown in figure 10: the directedness settings of (naive) observer JJ. These results are typical. In figure 11 we show the slopes of the linear regressions of the directedness set by the observers against the veridical values. Apparently all observers, with the possible exception of LIM, are sensitive to the measure of diffuseness of the light field. The slopes are somewhat below 1, though typically over 0.5 .

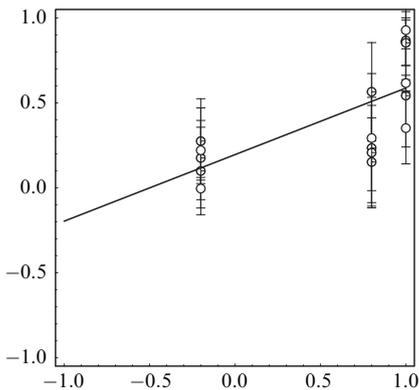


Figure 10. Regression of the directedness settings of (naive) observer JJ against the veridical values. All data have been used, including those (few) for which the directedness is not very well defined (eg the 'contre jour' situations). The highest veridical directedness (1.0) applies to scene C; for scene A the directedness is slightly less (0.8), owing to inter-reflections within the scene; and for scene B it is (for the same reason) even slightly negative (-0.2).

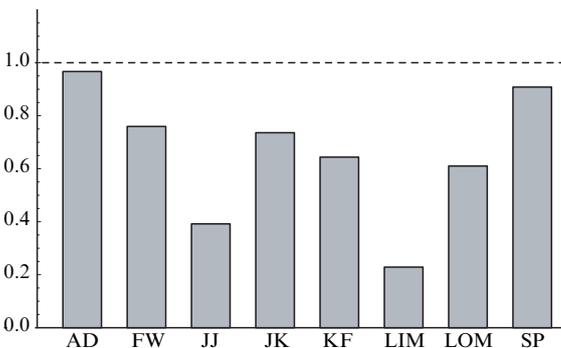


Figure 11. The slopes of the linear regressions of the directedness set by the observers to the veridical values.

The intensity of the light field is most variable in the scene with the source at the centre of the scene, and close to constant for the scene with the collimated beam (actually a small source at large, but finite, distance, leading to small variations of intensity). An example is shown in figure 12 for (naive) observer JJ. These results are typical. In figure 13 we show the slopes of the linear regressions of the observations against the veridical values. Although there is quite a bit of variability between observers, all observers are evidently sensitive to the variations and the observations are monotonically related to the physical variations. Two of the observers are even close to veridical, albeit with quite a bit of variability. The veridical values imply the well-known ‘inverse squares law’ of photometry. Apparently the observers respect this law at least semiquantitatively, albeit not in any precise sense.

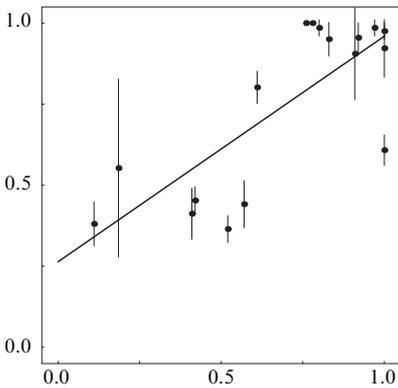


Figure 12. Regression of the intensity settings of (naive), observer JJ against the veridical values. All data have been used, including those (few) for which the intensity is not very well defined (eg the ‘contre jour’ situations).

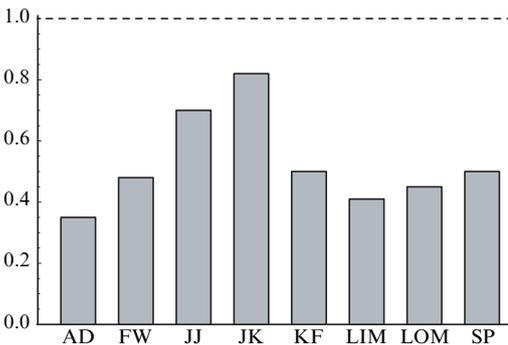


Figure 13. The slopes of the linear regression of the intensity observations against the veridical values.

Whereas it is of evident interest to study the correlation between the veridical parameters and the observers’ settings, it is perhaps not the most important analysis. The reason is that there is no a priori reason to expect that the settings will be anything close to veridical. In fact, in all likelihood we will have a situation like that in the case of ‘visual space’, or ‘pictorial space’, where it is generally accepted that observers routinely commit systematic errors. Thus, it is of interest to study the spread in the settings per observer and the correlation between observers rather than correlations with the veridical values. We studied the directional settings (slant and tilt) and the settings related to ‘quality’ (intensity and directedness) separately.

The slant and tilt settings for a given observer–stimulus pair varied only little from session to session. We used robust statistics because the data contain outliers owing to the fact that the slant or tilt need not necessarily make sense in all cases. Thus, we prefer the median and quartile deviations over the mean and standard deviation. The medians of the quartile deviations over all observers were 5.7° for the slant and 4.9° for the tilt. Thus the directional settings were very reproducible and precise.

The deviations varied between observers, the experienced observers having less spread (minimum 4.0° for the slant and 1.6° for the tilt) and the naive observers more, with observer LIM scoring consistently worse (in all settings, not just the direction) with a quartile deviation of 12.2° for the slant and 13.0° for the tilt. The observers are highly similar, the correlation matrix having a median entry of 0.90, and total range 0.58–0.99 for the slant, a median entry of 0.85 for the tilt, with total range 0.63–1.00. The lowest values are all due to observer LIM.

The median of the quartile deviations for the intensity over all observers is 0.023, the range being 0.00–0.56. The lowest values are due to the experienced observers, the highest value to observer LIM. The correlation matrix has a median entry 0.87 and coefficients ranging from 0.53 to 1.00, the lowest values being due to observer LIM. Observers are indeed very similar and reproducibility over sessions is fair.

The median of the quartile deviations for the directedness over all observers is 0.093, the range being 0.05–0.24. The lowest values are due to the experienced observers, the highest value to observer LIM. The correlation matrix has a median entry 0.76 and coefficients ranging from 0.39 to 0.98, the lowest values again being due to observer LIM. Observers are indeed very similar and reproducibility over sessions is again fair.

4 Conclusions

We have developed a novel method that allows one to probe the ‘visual light field’ and we have demonstrated its viability in a number of cases.

The method avoids the typically strong stimulus and response reduction common (or even cherished) in mainstream research [Palmer (1999) in his excellent overview of the field discusses numerous generic examples]. In our paradigm we address the psychologically relevant issue, that is the ‘eye measure’ operationalisation of the visual light field. The judgment of ‘fit’ avoids introspection as well as verbal (or cross-modal) response. It is as ‘objective’ as can be in view of the fact that the ‘visual light field’ is intrinsically an idiosyncratic, purely mental entity.

That we are able to measure the visual light field at any point in visual space is (of course) no guarantee that the visual light field will be consistent in the sense that the physical light field is, that is to say, the radiometric sense. The physical light field captures the radiance as a function of location and direction. The radiance is necessarily non-negative throughout, and the radiance in any direction is invariant under translations along straight lines in that direction (Gershun 1936; Born and Wolf 1964). Otherwise any function of direction and location specifies a physically possible light field. Whether a visual light field obtained in a location by location fashion will be consistent with the latter constraint is up to empirical verification. A priori we see no reason why it should. We know of examples from visual space where observers were able to achieve local geometrical measurements, but were unable to use such data in a globally consistent manner (Koenderink et al 2002). The same may very well apply to the visual light field. This is one issue that warrants a thorough investigation.

Observers apparently adjusted parameters, such as to cause the gauge object to ‘look right’ in the context of the scene. That was indeed their assigned task. They seem to have optimised the immediate appearance of the gauge object rather than its consistency with distant features of the scene, although overall indicators of the quality of the illumination such as the sharpness of cast shadows throughout the scene, or global illumination trends are, no doubt, used by most observers. This is in good accord with evidence from the art of painting that observers ignore inconsistency in many cases; there is also psychophysical evidence for this (Mingolla and Todd 1986; Ostrovsky et al 2005). It may be due to the ecological fact that most scenes contain compartments with mutually distinct light fields, since this would suggest an evolutionary advantage of avoiding such comparison of distant features. That observers are

satisfied with the ‘eye measure fit’ and do not seem to reason out various relations within the scene perhaps suggests that the visual light field has to be considered to be primarily a precognitive entity, that is to say an element of our visual presentations in Brentano’s (1874) sense. Another way to put this is to understand the visual light field as a particular type of Gestalt.

It is clearly of much interest to develop alternative gauge objects that would force the observer to distinguish between additional degrees of freedom. A case in point would be the introduction of an ‘ambient term’ as is common in computer graphics. We decided against this parameter because pilot experiments revealed interactions between direction, directedness, and ambient fraction, rendering the interface unmanageable. Thus the gauge figure used in this study precluded such a differentiation by design. Yet the ambient fraction is indeed an interesting parameter. Including it would imply a redesign of the gauge figure and the interface, which appears quite feasible. One might do it through the introduction of texture due to surface relief [eg with a golf ball instead of a smooth sphere (Pont and Koenderink 2005)], for example. By using such different designs for gauge objects and the nature of fit one may potentially tap a variety of aspects of the observer’s eye measure. This opens up a wide field of empirical enquiry.

We do not think that the use of stereoscopic presentation has anything to do with the ontology of the visual light field. Although of pragmatic importance, it has little or no conceptual relevance. In this experiment we used binocular disparity solely because we needed a device with which to define the location in 3-D visual space. The same type of experiment can be done in monocular pictorial space, but then one has to find a way to place the gauge object at specific locations in monocular pictorial space. An isolated sphere (as used in the present experiment) is evidently not suitable in such cases. But there is no reason (except from ease of computer rendering, and so forth) to stick to the spherical shape used in the present study. One could equally well use a virtual piece of sculpture for a ‘gauge figure’ and place it on the ‘floor’ in pictorial space (perhaps ‘nailing it to the ground’ with the help of some shadow) in order to specify its location with respect to the other pictorial objects unequivocally. This is clearly an important topic of future research.

In conclusion, we have shown that the ‘visual light field’ is at least as well defined as ‘visual space’, and does not deserve less attention. We have introduced methods that enable one to perform effective measurements of the relevant parameters and thus define the visual light field in an operational sense. This opens up a field of investigation that is likely to be of importance in many contexts of vision science.

Acknowledgments. This research was done under a grant of the European Commission. Sylvia Pont was supported by the Netherlands Organization for Scientific Research (NWO). Hans Kolijn assisted us with virtually all technical problems.

References

- Abrahams M V, Kattenfeld M G, 1997 “The role of turbidity as a constraint on predator–prey interactions in aquatic environments” *Behavioral Ecology and Sociobiology* **40** 169–174
- Adams A, 1950 *Basic Photo* (6 volumes) (New York: Morgan and Lester)
- Adelson E H, Bergen J R, 1991 “The plenoptic function and the elements of early vision”, in *Computational Models of Visual Processing* Eds M S Landy, J A Movshon (Cambridge, MA: MIT Press) pp 3–20
- Adelson E H, Pentland A P, 1996 “The perception of shading and reflectance”, in *Perception as Bayesian Inference* Eds D Knill, W Richards (New York: Cambridge University Press) pp 409–423
- Arnheim R, 1956 *Art and Visual Perception* (London: Faber and Faber)
- Barocci F F, 1597 *Nativity [Painting]* (Madrid: Museo del Prado)
- Baxandall M, 1995 *Shadows and Enlightenment* (New Haven, CT: Yale University Press)
- Belhumeur P N, Kriegman D J, Yuille A L, 1999 “The bas-relief ambiguity” *International Journal of Computer Vision* **35** 33–44

- Bergström S S, 1977 "Common and relative components of reflected light as information about the illumination, colour, and three-dimensional form of objects" *Scandinavian Journal of Psychology* **18** 180–186
- Born M, Wolf E, 1964 *Principles of Optics* (Oxford: Pergamon Press)
- Brainard D H, 1998 "Color constancy in the nearly natural image. 2. Achromatic loci" *Journal of the Optical Society of America A* **15** 307–325
- Brentano F, 1874 *Psychologie vom empirischen Standpunkt* (Leipzig: Duncker und Humboldt)
- Brooks M J, Horn B K P, 1989 *Shape from Shading* (Cambridge, MA: MIT Press)
- Caravaggio, 1600a *The Conversion of St Paul [Painting]* (Rome: Cerasi Chapel, Maria del Popolo)
- Caravaggio, 1600b *The Calling of St Matthew [Painting]* (Rome: Contarelli Chapel, S Luigi dei Francesi)
- Comenius J A, 1658 *Orbis sensualium pictus hoc est omnium fundamentalium in mundo rerum & in vita actionum pictura & nomenclatura* (Nuremberg: M Endter) [Picture from the Introduction]
- Coren S, Komoda M K, 1973 "The effect of cues to illumination on apparent lightness" *American Journal of Psychology* **86** 345–349
- Dror R O, Willsky A S, Adelson E H, 2004 "Statistical characterization of real-world illumination" *Journal of Vision* **4** 821–837
- Feynman R P, Leighton R B, Sands M, 1965 *The Feynman Lectures on Physics* volume 3 (Reading, MA: Addison-Wesley)
- Foley J D, Van Dam A, 1983 *Fundamentals of Interactive Computer Graphics* (Reading, MA: Addison-Wesley)
- Gershun A, 1936 "The light field" translated from Russian by P Moon, G Timkoshenko; 1939 *Journal of Mathematics and Physics* **XVIII** 51–151
- Gibson J J, 1966 *The Senses Considered as Perceptual Systems* (Boston, MA: Houghton Mifflin)
- Gilchrist A, 1999 "Lightness perception", in *MIT Encyclopedia of Cognitive Science* Ed. R W F Keil (Cambridge, MA: MIT Press) pp 471–472
- Gilchrist A, Jacobsen A, 1984 "Perception of lightness and illumination in a world of one reflectance" *Perception* **13** 5–19
- Gombrich E H, 1995 *The Depiction of Cast Shadows in Western Art* (London: National Gallery Publications)
- Hogarth B, 1981 *Dynamic Light and Shade* (New York: Watson–Guptill)
- Hoogstraeten S van, 1678 *Inleyding tot de hooge schoole der schilderkonst: anders de Zichtbaere Werelt. Verdeelt in negen Leerwinkels, yder bestiert door eene der zanggodinnen* Book 7: *Melpomene* (Rotterdam)
- Ikeda M, Shinoda H, Mizokami Y, 1998 "Phenomena of apparent lightness interpreted by the recognized visual space of illumination" *Optical Review* **5** 380–386
- Jacobs T S, 1986 *Drawing with an Open Mind* (New York: Watson–Guptill)
- Kersten D, Yuille A, 2003 "Bayesian models of object perception" *Current Opinion in Neurobiology* **13** 1–9
- King D, 1997 *The Commissar Vanishes: The Falsification of Photographs and Art in Stalin's Russia* (New York: Henry Holt)
- Koenderink J J, Doorn A J van, 2004 "Shape and shading", in *The Visual Neurosciences* Eds L M Chalupa, J S Werner (Cambridge, MA: MIT Press) pp 1090–1105
- Koenderink J J, Doorn A J van, Kappers A M L, 1992 "Surface perception in pictures" *Perception & Psychophysics* **52** 487–496
- Koenderink J J, Doorn A J van, Kappers A M L, Lappin J S, 2002 "Large-scale visual fronto-parallel under full-cue conditions" *Perception* **31** 1467–1475
- Kohlrausch F, 1923 *Lehrbuch der praktischen Physik* (Leipzig: Teubner)
- Kozaki A, Noguchi K, 1976 "The relationship between perceived surface-lightness and perceived illumination" *Psychological Research* **39** 1–16
- Lambert J H, 1760 *Photometria sive de mensura de gratibus luminis, colorum et umbrae* (Augsburg: Eberhard Klett)
- Land E H, 1959 "Color vision and the natural image: Part I" *Proceedings of the National Academy of Sciences of the USA* **45** 115–129
- Langer M S, Bülthoff H H, 2000 "Depth discrimination from shading under diffuse lighting" *Perception* **29** 649–660
- Mamassian P, Knill D C, Kersten D, 1998 "The perception of cast shadows" *Trends in Cognitive Sciences* **8** 288–295
- Mingolla E, Todd J T, 1986 "Perception of solid shape from shading" *Biological Cybernetics* **53** 137–151
- Minnaert M, 1968 *Natuurkunde van het Vrije Veld* (Zutphen: Thieme)

- Minsky M, 1974 *A Framework for Representing Knowledge* MIT-AI Laboratory Memo number 306
- Moon P, Spencer D E, 1981 *The Photoc Field* (Cambridge, MA: MIT Press)
- Nayar S K, Narasimhan S G, 1999 "Vision in bad weather", in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* volume 2, pp 820–827
- Noguchi K, Kozaki A, 1985 "Perceptual scission of surface-lightness and illumination: An examination of the Gelb effect" *Psychological Research* **47** 19–25
- Ostrovsky Y, Cavanagh P, Sinha P, 2005 "Perceiving illumination inconsistencies in scenes" *Perception* **34** 1301–1314
- Palmer S E, 1999 *Vision Science. Photons to Phenomenology* (Cambridge, MA: MIT Press)
- Pentland A P, 1982 "Finding the illuminant direction" *Journal of the Optical Society of America* **72** 448–455
- Pentland A P, 1990 "Linear shape from shading" *International Journal of Computer Vision (Archive)* **4** 153–162
- Pont S C, Koenderink J J, 2005 "Bidirectional texture contrast function" *International Journal of Computer Vision* **62** 17–34
- Rembrandt van Rijn, 1636–1647 *Danaë [Painting]* (St Petersburg: Hermitage)
- Rembrandt van Rijn, 1642 *Night Watch [Painting]* (Amsterdam: Rijksmuseum)
- Rembrandt van Rijn, 1646 *The Adoration of the Shepherds [Painting]* (London: National Gallery)
- Robilotto R, Zaidi Q, 2004 "Limits of lightness identification for real objects under natural viewing conditions" *Journal of Vision* **4** 779–797
- Schöne W, 1979 *Über das Licht in der Malerei* (Berlin: Gebr. Mann)
- Searle J, 1983 *Intentionality: An Essay on the Philosophy of Mind* (New York: Cambridge University Press)
- Spicer J A, Orr L F, 1997 *Masters of Light. Dutch Painters in Utrecht during the Golden Age* (New Haven, CT: Yale University Press)
- Stoichita V, 1997 *A Short History of the Shadow* (London: Reaktion Books)
- Tarr M J, Kersten D, Bülhoff H H, 1998 "Why the visual recognition system might encode the effects of illumination" *Vision Research* **38** 2259–2275
- Todd J T, Norman J F, Mingolla E, 2004 "Lightness constancy in the presence of specular highlights" *Psychological Science* **515** 33–39
- Tricker R A R, 1970 *Introduction to Meteorological Optics* (New York: American Elsevier)
- Zavagno D, 2005 "The phantom illusion" *Perception & Psychophysics* **67** 209–218

PERCEPTION

www.perceptionweb.com

This article is an advance online publication. It will not change in content under normal circumstances but will be given full volume, issue, and page numbers in the final PDF version, which will be made available shortly before production of the printed version.

Conditions of use. This article may be downloaded from the Perception website for personal research by members of subscribing organisations. Authors are entitled to distribute their own article (in printed form or by e-mail) to up to 50 people. This PDF may not be placed on any website (or other online distribution system) without permission of the publisher.