Forthcoming in *Philosophy and Phenomenological Research.* Please cite published version.

The Nature, Structure, and Perception of Illumination

Will Davies

In the Blue Marble, a famous photograph taken by the crew of Apollo 17, Earth appears as a magnificent, glowing orb in the void. You can make out the continent of Africa, with oceans either side. Vast swirls of cloud lie over Antarctica. You can also see the atmosphere, enveloping Earth like a blue blanket. The atmosphere seems bright and luminous, as if it were resonating with radiant energy. This is in contrast to the space around, which seems utterly dark. Yet this space is replete with light. The amount of light energy reaching Earth from the sun is 1361W/m^2. This is only moderately higher than the energy at sea level, which averages around 1000W/m^2. Why, then, does the atmosphere appear luminous, whereas space seems dark? The obvious point is that the atmosphere contains gases that scatter light. Scattering alone, however, does not explain Earth's glow, as 'empty' space contains plenty of dust, ice, and gases that also scatter light. The atmosphere does not merely scatter light: it traps it in a self-sustaining system, where it is continuously dispersed between sky and ground. We no longer have just light rays travelling in straight lines, as in space. We have a dense, structured field of illumination.

Illumination is a defining characteristic of natural environments, yet its nature and spatial structure remain poorly understood. As Section 1 explains, illumination is not simply light, but emerges from the constant and cumulative interplay between radiation and matter. These accretions of light form fields of illumination, which possess novel features and self-organising structures that persist through the continual flow of light itself. These include shadows and shadow volumes, which have already attracted some philosophical attention.[1]

[1] A recent sample includes Sorenson (2008), Phillips (2018), and Casati and Cavanagh (2019). Casati & Cavanagh (2019: 215-234) discuss related issues concerning illumination.

This paper advances the debate by characterising the nature and structure of illumination in generality. The account developed here encompasses not only shadows, but also light enclosures and sunbeams; amorphous formations, like dappled illumination; and more uniform expanses, like regions of sunlight, skylight, and twilight. These structures are immaterial and ephemeral, radically unlike the solid material objects that dominate our naïve ontologies. Yet they are no less real for that. They are bona fide entities, falling squarely within the remit of the special sciences.

Section 2 argues that these illumination entities are objects of perception, things to which features are perceptually attributed.[2] This contrasts with the majority view, which explains illumination-dependent variations in appearance solely by appeal to the perceived properties of material objects. On this view, illumination is not perceptually represented, but only inferred from its effects on object appearance. Hilbert (2005: 151), for example, claims that 'what we see as changing with the illumination is an aspect of the object itself, not the light source or the space surrounding the object.' Casati and Cavanagh (2019: 161-62) maintain that 'shadows are a product of illumination and thus are properties of the surface they fall on.' Likewise, Chalmers (2006: 87) holds that 'it is best to talk of shadow properties instantiated at locations on objects, rather than talking of shadows: while we sometimes have the phenomenology of seeing shadows as objects, it is arguable that more often we do not.'[3]

[2] Matthen (2010, 2018), Brown (2014), and Gert (2017) also endorse illumination perception. Allen (2016: 35) is sympathetic, and Burge (2022: 218-219, 245-247) is open but noncommittal. The view developed in Davies (2018) is that colour vision constitutively involves capacities to represent spectral features of the illumination, though did not consider spatial-structural features, hence underplayed the extent of illumination perception.

[3] Many hold that material objects appear to have certain illumination-dependent colours, in addition to surface colours, variously called *viewpoint-relative colours* (Tye, 1996), *perspectival-colours* (Noë, 2004), *situation-dependent colours* (Schellenberg, 2008), or *apparent colours* (Allen, 2016). Hilbert (2005), Jagnow (2010), and Davies (2016) propose instead that surface colour appearance includes dimensions beyond hue, saturation, and lightness, which capture aspects of the incident illumination. This idea dates to Katz (1911/1935: 76, 176), though Katz took a more expansive view of illumination perception, as now do I.

Chalmers is presumably right that we rarely see shadows 'as objects,' in the sense that applies to perceptions of medium-sized dry goods.[4] As Section 1 shows, however, illumination entities have distinctive structures—more like continuous masses of fluid than discrete solids. Expanses of sunlight, sunbeams, dappled lighting, and light enclosures are therefore better compared to stretches of seawater, currents, ripples, and waves. These entities are individuated not by rigid boundaries, but by stable patterns of intensity and flow. With this structural account in place, questions concerning the perceptual representation of illumination can be approached with greater precision. I argue that this fluid-like structure is mirrored by a distinctive type of perceptual organisation, whereby scenes are parsed into zones and flows of illumination at multiple scales. This organisation involves attributing spatial-structural features that apply distinctively to illumination, which warrants treating it as an object of perception.[5]

1. The Nature and Spatial Structure of Illumination

1.1. Light, Matter, and Illumination

This section clarifies the nature and spatial structure of illumination. What kind of thing is illumination, and how is it organised? Clearly, illumination is intimately related to light. It originates from light emitting objects, like the sun. Brightly illuminated regions receive more light than dimly illuminated regions. Shadows form where objects block the light, causing reductions in its intensity. Illumination and light, though, are different kinds of stuff. Light, conceived as electromagnetic radiation, exists at every point in space and time, whereas

[4] De Wit et al. (2012) found that perceived shadows are sometimes associated with object-specific preview benefits, a mark of 'object-based' attention. In contrast, Rensink & Cavanagh (2004) found that visual search for shadow shape is relatively slow, suggesting these shapes are 'discounted' after an initial detection and classification phase. See Casati & Cavanagh (2021: 98) for discussion.

[5] This is consistent with the pluralistic view that we perceive illumination itself *and* certain illumination-dependent properties of material objects, as suggested by Allen (2016: 35ff). As noted, the majority view appeals solely to the latter, and thus disregards illumination perception.

illumination exists only where light interacts continuously with macroscopic masses of matter. As Gibson (1986: 50) says, 'radiation becomes illumination by *reverberating* between the earth and the sky and between surfaces that face one another.' Once light enters the atmosphere, it starts a cascade of light-matter interactions. These interactions reshape the flow of light, producing a structured system with a stable organisation. Much as the prevailing winds and tides shape sand into dunes, so the constant interplay with matter shapes light into illumination.

The result, Gibson (1986: 50) says, is that at each point, 'light… converges from all directions and, moreover, has different intensities in different directions.' This convergence is not a mere agglomeration of rays, but produces emergent flow structure. Compare a room with a heater at one end and an open window letting in cold air at the other. The warm air tends to rise, whereas the cool air sinks. At each point in the room, air of different temperatures arrives from multiple directions. This produces complex, emergent flow patterns at macroscopic scales, called 'convection.' Similarly, take a room with a lamp at one end, and a window at the other. The lamp produces warm, directed light. The window lets in cool, diffuse skylight. At each point, light with different properties arrives from every direction. The light ricochets within the space, rather than simply passing through it. This yields emergent patterns of flow at macroscopic scales, which characterise the illumination within the room.

I said that illumination exists only where light interacts with matter. More precisely, illumination is generically existentially dependent on matter.[6] Defined modally, the condition is that, necessarily, illumination exists in *C*, only if some matter exists in *C*. The dependence is generic, as it does not entail the existence of some specific bit of matter, just some matter.

[6] See Tahko & Lowe (2020: §4.3).

How much matter, in what sort of arrangement? Crucially, the matter must form an environment of some kind. An environment defines the conditions and surroundings in which an animal or plant lives. As Gibson (1986: 8) notes, organism and environment thus form 'an inseparable pair.' This makes the existence of illumination an organism-dependent matter, falling under ecology rather than physics.[7] For humans, an environment typically encompasses a ground surface, surrounded by air, populated by vegetation and medium- to large-sized structures, both natural and man-made. This could be a forest, city, building, or room, for example, depending on the explanatory context.

Illumination is generically existentially dependent on matter, but not partly constituted by it. By comparison, rivers existentially depend on solid matter, in that a river exists in *C*, only if some solid matter forms a channel in *C*. The channel is merely a conduit for the river; it does not partly constitute it. The channel causally determines some of the river's features, such as its shape and flow; but these are properties of the river, not the channel. Similarly, illumination exists in *C*, only if some matter forms an environment in *C*. The illumination is constituted by the light, however, not the matter. The matter causally determines some of the illumination's features, such as its intensity and flow; but these are properties of the illumination, not the matter.

I described illumination as 'emergent.' More precisely, it is *weakly emergent*, in the sense that illumination is metaphysically necessitated by fundamental physical phenomena. Illumination falls under the special sciences, but like rivers and weather, poses no threat to physicalism. More positively, the nature of illumination is partly characterised by *novelty*, *self-organisation*, and *replacement autonomy*.[8] Briefly, concerning novelty, illumination has properties not possessed by quanta of light, just as rivers have properties not possessed by

[7] Thompson (1995: 102-103), Matthen (2010: 238-241), Allen (2016: 34), and Burge (2022: 218) discuss the ecological significance of illumination.
[8] This follows Humphreys (2016: 35, 153-156).

individual water molecules. These include scalar and vector magnitudes related to intensity and flow, along with certain mereological and topological properties. As for self-organisation, illumination systems contain patterns or structures that develop endogenously, without any interventions from outside, except for light flowing into them. These patterns are non-accidental, conforming to special scientific laws that describe relations between the light input, objects within the scene, and the resulting structure. They tend to be relatively stable across time, at scales of minutes and hours, though usually not whole days. They are also scalable, with similar structures emerging at small, medium, and large scales. Finally, per replacement autonomy, these structures persist through continual changes in the light quanta, much as river flows persist through changes in water molecules. Understanding the nature of illumination therefore requires close attention to these emergent properties and structures. Now, time for details.

1.2. Illumination Fields

The emergent properties of illumination can be characterised by Gershun's (1939) theory of the *light field*.[9] Despite the name, light fields are not strictly fields, in the sense of functions defined on all space-time points. Confusingly, light fields are not even fields of light, in the sense of electromagnetic radiation. Electromagnetic radiation passes right through opaque objects, whereas illumination stops at the surface. Gershun thus describes light fields as part of '*geometrical optics* rather than physical optics,' features of the 'macrocosmos' or lived environment, (1939: 56). To avoid confusion, I call them *illumination fields*.

Illumination fields are represented by five-dimensional spherical functions that describes the radiance arriving at each point (x, y, z) from every direction (θ, ϕ).[10] Adding a

---

[9] Cf. Adelson & Bergen (1991) on the 'plenoptic function.'

[10] See Kartashova et al. (2016: 2) and Xia et al. (2014: 613). Pont (2019: 507) defines *radiance* as the 'amount of radiation that passes through, is emitted from, or is reflected from a particular area and falls within a given solid angle.' Note that one can technically define such a function in outer space, even though, ex hypothesi, illumination exists only in natural environments.

wavelength parameter gives the spectral characteristics of this radiation. The value of the function at a point can be regarded as the image formed by unfolding the surface of a perfectly reflective sphere placed at that point.

The spherical function can be analysed as the sum of various components, each representing distinctive properties of illumination.[11] The 0th-order component yields (nonnegative) scalar magnitudes, which represent the radiance averaged over all directions at each point. This is sometimes called ‘ambient’ or ‘Ganzfeld illumination.’ It is rare for ambient magnitudes to dominate in daylight conditions, in the sense of being the primary determinants of the character of the illumination. This occurs only in situations like a polar white-out or heavy fog. At night, though, the illumination is characterised primarily by low ambient magnitudes.

Nocturnal illumination, like most low-intensity illumination, lacks significant flow features. These are captured by the 1st-order component, which yields vector magnitudes representing the dominant direction and amount of illumination flow at each point. These magnitudes dominate in sunlit scenes, where illumination flows predominantly in one direction, and indoor scenes with multiple light sources, where illumination may flow in contrary directions at different locations. Where these flows meet, the vectors may cancel out, producing a clamp or local minimum, like two waves colliding. These structures are characterised by the 2nd-order component, which represents pairs of directions and magnitudes of illumination flow, or 4D tensors. This is relevant to the discussion of saddles below, so I park it for now.

Following Xia and colleagues (2017), diffuseness is defined by the ratio of magnitudes of the 1st- and 0th-order components, subtracted from 1. The more diffuse the

[11] These components are derived from spherical harmonics, as detailed by Mury et al. (2007: 7309-7312).

illumination, the smaller the ratio of directed-to-undirected components. Values range from completely collimated illumination, containing parallel rays with no divergence, to completely diffuse, coming from all directions with equal intensity. The two commonest illumination conditions, sunlight and skylight, are both directed strongly downwards, but skylight is more diffuse than sunlight. Snow storms and fog produce nearly maximally diffuse illumination. Bright moonlight can be quite directional, but nocturnal illumination is generally diffuse, with light from the moon, stars, and terrestrial sources being scattered throughout the sky.

1.3. Surface Illumination Fields

Illumination fields are defined where radiance converges from all directions. These fields are discontinuous in regions containing opaque material objects.[12] Opaque objects *screen* the illumination, in the sense that they induce the presence of surface boundaries in the illumination field. Call these *surface illumination fields*.[13] Surface illumination fields are *illumination surfaces*: two-dimensional boundaries of volumes of illumination. These boundaries are major structural features of illumination, and have many interesting properties, which I now explain.

In some ways, illumination surfaces are like certain material surfaces, such as surfaces of water. First, as liquid stuff, water takes on the shape of any container in which it is placed. Its surfaces are shaped from the outside, by the vessel it fills. Illumination surfaces are similarly shaped by the objects that screen the illumination, precisely following their contours and textures. Illumination is not liquid, though. It is an immaterial stuff, constituted by light. Illumination is thus not deformed, in the sense of changing shape and size due to mechanical

[12] See Pont & Koenderink (2004: 2) and Mury (2009: 77).
[13] Cf. Koenderink et al. (2007: 896), Van Doorn et al. (2011: 13), and Sorensen (2008: 92).

force. Material things rather exert a kind of 'reflective tension' on the illumination, redirecting and moulding its flow.

Second, as boundary entities, material surfaces are existentially dependent on the bodies that own them.[14] There are no free-floating water surfaces, only surface parts of water volumes. Similarly, illumination surfaces are existentially dependent on the fields that own them. By definition, illumination surfaces form only where an illumination field is screened by matter. Hence, no illumination field, no illumination surface. More precisely: necessarily, if some illumination surface *S* exists, then there is some illumination field of which *S* is a boundary part.

Third, water surfaces have features not possessed by the interiors of water volumes. For example, because these surfaces mark interfaces with different media, they will reflect some light, and can be glossy. They also exhibit surface tension, allowing some small objects to float, despite being denser than water. Similarly, illumination surfaces have features not possessed by the interiors of illumination fields. These include *irradiance* magnitudes, often called *shading*. These are scalar magnitudes—surface analogues of ambient illumination—giving the amount of radiation received by a surface, per unit of time and area. There are also *irradiance flow* magnitudes, which reflect the net directional bias in irradiance. These are vector magnitudes, giving the dominant direction of illumination flow tangent to the surface at each point. As van Doorn and colleagues (2012: 477) say, 'irradiance flow exists only on the surfaces of objects in the scene.' Like surface tension properties, then, these magnitudes characterise the boundaries of illumination fields, not the interiors.

Fourth, fluid flow vectors are defined at the surfaces of a water volume, as throughout its interior. Unlike interior vectors, however, surface vectors are defined by specific boundary

[14] For a related point, see Stroll (1988: 36-37).

conditions on the system. For example, the 'no slip condition' requires that fluid in contact with a rigid surface has the same velocity as that surface, where they make contact. If the rigid surface is stationary, the flow velocity is zero. Analogously, flow magnitudes exist at illumination surfaces, as throughout the illumination field. Like surface fluid flows, irradiance vectors are defined by boundary conditions on illumination fields. To see how this works, first consider how scalar irradiance varies with features of the scene. Imagine holding a completely flat object with the surface normal (the line perpendicular to the surface tangent) pointing directly at a source of collimated light. The irradiance will be at its maximum value. Now rotate the object, so the angle between the illumination direction and surface normal increases to ninety degrees. The irradiance will decrease as the illumination hits at more grazing angles, eventually reaching zero. Irradiance thus depends on the properties of the surrounding illumination field, and the geometry of the material surface. More compactly, shading depends on both illumination and form.

Similar points apply to irradiance flow. Take the stone sphere in Figure 1, which is illuminated from above and to the left by highly directional, nearly collimated beams from the sun. Whereas the irradiance peaks where the surface faces the source, the irradiance flow is zero here. The flow increases as the surface curves away, reaching a maximum where tangent to the sun's beams. This is similar to fluid flow: if you hold a flat object with the surface normal pointing at running water, the water splashes off, generating no flow. As you tilt the object, the water clings more to its surface, reaching maximum flow at ninety degrees. The flow abruptly stops around this midline, leaving the lower hemisphere in near total shadow. In contrast, diffuse illumination would flow right down over the surface, converging on another flow minimum at its base.[15] In sum, irradiance flow is defined by boundary

[15] These shading patterns are due to *vignetting*, an interaction effect between surface orientation and variations in the local illumination field. See Koenderink & Pont (2002), Pont & Koenderink (2004), and van Doorn et al. (2011).

conditions that relate the surrounding illumination field and object geometry to the direction and magnitude of flow across the illumination surface.

Figure 1: Surface illumination field on sphere in sunlight. Source: Creative Commons CC0.

Fifth, and finally, although surface fluid flow depends on features of the rigid surface, the flow is a feature only of the fluid surface. The dependence is causal and computational, not constitutive. As for causal dependence, molecular interactions between the surfaces generate adhesive forces, which cause the fluid to have zero flow at its boundary. Computationally, surface fluid flow is modelled as a function of rigid surface velocity. Yet the resulting vectors characterise the boundary of the fluid flow system, not the rigid material. Similarly, although irradiance flow depends on the geometry of the material surface, the flow characterises only the illumination surface. As for causal dependence, intervening on surface curvature toggles the net energy received from the illumination, which defines irradiance. This also toggles the dominant direction of illumination flow tangent to the surface, hence irradiance flow. Computationally, irradiance and irradiance flow are modelled as functions of the illumination field and surface geometry. This is significant, as taken in inverse, if you know the properties of the illumination field and the irradiance at a point, you can compute the surface normal at that point. This is how you get 'shape-from-shading.'[16] The main point, though, is that irradiance magnitudes characterise illumination surfaces, not the surfaces of opaque objects. While shading undoubtedly occurs at the surface of the sphere, it is strictly not a feature of the sphere.

In some ways, then, illumination surfaces are like certain material surfaces. In other ways, they are sui generis. First, illumination is an immaterial stuff, thus not subject to the forces acting on matter. Illumination surfaces do not arise from mechanical forces, like the shear stresses exerted on water by the sides of a glass. They arise from the reflective tension exerted by matter on illumination. This reflective tension is not a mechanical force, rather a structural constraint on illumination that emerges from the continuous interactions between light and matter. This tension redirects and redistributes illumination across the material surface, forming patterns of irradiance flow.

Second, illumination surfaces lack the marks of materiality borne by ordinary surfaces. As Stroll (1988: 32) notes, an illumination surface 'cannot be said to be dark or light, thick or thin, smooth or rough, wet or damp; nor can it be removed, sanded, polished, and so forth.' It cannot be glossy, possessing shine or sheen. Though immaterial, illumination surfaces are not mere abstracta. They are not 'Leonardo surfaces,' in Stroll's (1988: 44) sense of purely 'conceptual entities, … logical limits that mark a theoretical distinction.' They are as real as surfaces of water, despite being untouchable.

Third, at least naively, when you put water in a glass, the liquid surface that is formed is not co-located with the glass.[17] Illumination surfaces, in contrast, are precisely co-located with material surfaces. Intriguingly, however, these surfaces share no parts. Take the

[16] Ramachandran (1988) provides compelling demonstrations of the influence of perceived illumination direction on perceptions of surface geometry.
[17] Smith (2007) offers a sophisticated account of contact between material bodies, drawing on continuum mechanics, which challenges this naïve perspective.

irradiance field formed over the surface of the stone sphere. No part of this field is stone: it is composed of irradiance magnitudes. Likewise, no part of the sphere is an irradiance field: it is all stone. The field and the sphere are co-located, but do not overlap, hence they interpenetrate. This plausibly reflects the fact that these are fundamentally different kinds of thing—one immaterial, the other material. Since nothing immaterial is material, and vice versa, no part of one can be part of the other.

1.4. The Mereotopology of Illumination

Illumination fields have complex mereological and topological (*mereotopological*) properties. The previous subsection highlighted properties relating to boundedness and closure. Illumination fields are bounded by illumination surfaces that are induced by material objects; they are closed in these regions, containing their surfaces as parts.[18] The internal structure of illumination is characterised by scalar field and vector field topologies. The 0^{th}-order component of illumination comprises a three-dimensional scalar field. The 1^{st}-order component comprises a three-dimensional vector field. In the surface illumination field, irradiance and irradiance flow magnitudes comprise two-dimensional scalar and vector fields, respectively. These magnitudes form complex patterns at multiple scales. The patterns form topological structures, which survive under continuous geometrical transformations, such as stretching or bending – so-called 'rubber sheet' transformations. These structures are ephemeral yet genuinely existent *illumination entities*, composing the parts of illumination fields, as I now explain.

1.4.1. Scalar Field Topology

Scalar field topologies characterise the structures of undirected magnitudes distributed across space. The key structures are *critical points* (lines, surfaces) where the gradient of the scalar

[18] An object is *closed* (at a region) just in case it contains its boundary (at that region), and *open* (at a region) just in case it does not.

function vanishes. In the illumination field, these include *local illumination maxima* and *minima*, the 'hills' and 'valleys' where the ambient illumination is decreasing or increasing everywhere in the neighbourhood of the point. Critical points are thus nonlocal features, depending on the illumination nearby. There are also *ambient illumination saddles*, where the illumination increases toward a point along one axis, and decreases along the orthogonal axis.

The scalar component of the surface illumination field similarly contains *local irradiance maxima* and *minima*, where the irradiance is decreasing or increasing everywhere in the neighbourhood of the point. It also contains *irradiance saddles*, where the irradiance increases towards a point along one axis, and decreases along the orthogonal axis.

Also important are *separation contours* (*surfaces*), which connect saddles to local minima and maxima. *Isocontours* connect points with equal scalar values. Together with critical points, these structures suffice to characterise the *scalar topological skeleton* of an illumination field. Figure 2 represents the skeleton of a simple, two-dimensional scalar field. The left and right circles represent local maxima, while the central circle represents a saddle. The solid lines are separation contours connecting the saddle to the two maxima left and right, and local minima above and below, lying out of view. The arrows on these lines point in the direction of increasing gradient; the dotted lines are isocontours.

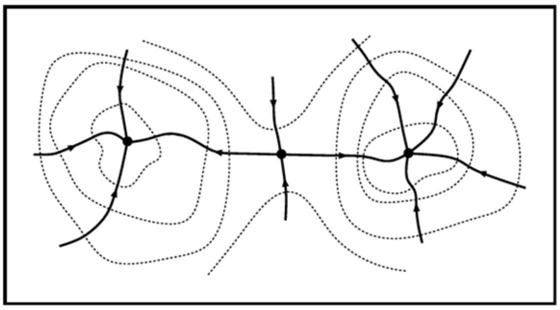

Figure 2. Representative scalar topological skeleton.

These structures are commonplace in surface illumination fields, like the dappled illumination in Figure 3a. Figure 3b overlays a scalar skeleton on this field. The two black dots are irradiance maxima, where the illumination is brightest. The central blue dot is an irradiance saddle, where the two bright blobs meet at the fringes. Although these blobs merge seamlessly, the saddle captures the fact that two distinct structures make ‘contact’ here. The solid lines are separation contours connecting these critical points, with arrows pointing in the direction of increasing gradient, from shadowy to bright illumination. The dotted red lines are isocontours, which convey the bulbous shapes of the blobs. These structures proliferate across the scene, creating a dense network of light and shadow. At larger scales, the entire field composes a lattice of blobs and tendrils, like an ephemeral tissue.

Figure 3. (a) dappled illumination, (b) overlaid scalar topological skeleton. Stock Adobe Image, Standard License

Such skeletons capture structures in illumination that are invariant under rubber sheet transformations. Illumination surfaces do not strictly stretch or bend. Nevertheless, the light spots might swell, contract, or elongate as the trees sway overhead. This may change the locations of irradiance maxima and saddles, along with the shapes of the separation contours and isocontours. The surface illumination may nonetheless retain its skeletal structure, with two local maxima meeting at a saddle. This structure could be broken, for instance, if a patch of shade formed at the centre of one of the light spots, replacing the irradiance maximum with a local minimum; a trough, where once was a peak. Similarly, if a new light spot formed between the original two, this would introduce a further irradiance maximum, like a new swell in the sea, thus changing the overall structure of the field.

In Figure 3a, the irradiance mostly varies gradually. This is typical of natural scenes, which tend to have relatively smooth transitions in illumination. There are some sharper gradations, such as at the thick, roughly diagonal shadow near the top. These are often called 'illumination edges.'[19] This might suggest that surface illumination involves another first-order structure, *edge*, in addition to those discussed above. On the contrary, edges are just limiting cases of a more fundamental gradient structure. They occur where the irradiance gradient is very steep or discontinuous. These structures can be derived from the scalar field, by taking the second derivative of irradiance, or some suitably smoothed function thereof. Edges are thus second-order features, not first-order. The first-order structure is fully described by maxima, minima, saddles, and gradient contours. These features are more fundamental to understanding the scalar structure of illumination.

Casati and Cavanagh (2019: 1) claim that shadows are 'holes in the light, created by an opaque or semiopaque object blocking the light falling on a surface.' Casati and Varzi (1994: 6) similarly claim that shadows are constituted by space, just like holes in material things. In one respect, the analogy is apt. Like holes, shadows are existentially dependent on the presence of some surrounding stuff. The hole in a doughnut exists only at the pleasure of the dough. Likewise, shadows exist only when surrounded by brighter illumination.[20] As Casati and Cavanagh (2019: 215) remark, 'shadow and light celebrate an unbreakable marriage.'

In other, more fundamental respects, however, the analogy is mistaken. Firstly, shadows are not holes in light: they are structures in surface illumination fields, formed where local irradiance minima gradate to local irradiance maxima. The gradient may be steep, as with shadows cast in strong sunlight, or shallow, as in diffuse skylight. In the most

[19] See Gilchrist et al. (1983), Adelson & Pentland (1996), Burge (2010: 352-354), and Akins & Hahn (2014: 140ff).
[20] Sorenson (2008: 102-103) disagrees, though Phillips (2018: 181-182) offers a convincing reply.

diffuse conditions, these gradients vanish altogether, leaving no shadows. Secondly, shadows are not constituted by space, but by irradiance magnitudes. In strong and directional illumination, these magnitudes may be close to zero. On the moon, for example, shadows are truly dark, due to the near-complete lack of atmospheric scattering. On Earth, shadows are never completely dark, as the illumination is always at least somewhat diffuse. In highly diffuse illumination, shadow irradiance can be comparable to white surfaces under strong artificial lighting. Shadows, then, are not just absences or voids, like holes. They are scalar structures defined by local minima and connecting gradients, like cold wells in the temperature field.

Illumination fields also have significant scalar topological structure. Figure 4a shows a simple case involving two, large structures: a volume of dim illumination in the hallway, and bright illumination in the conservatory.[21] The local minima are deep in the corners of the hallway. Little illumination ventures in that far, and any that does is weakened by reverberations between the walls. The ambient illumination maxima are just inside the windows. These minima and maxima are connected by a fairly smooth gradient, which is steepest around the door. This gradient captures the balloon-like curvature of the bright illumination, which seems to 'bulge' into the hallway. Similarly, in Figure 4b a storm creates a vast expanse of murk, abutting the daylight. The ambient illumination maxima cluster near the centre, where sunlight radiates through the clouds; the local minima are to the left. There is a clear gradient from light to dark, which is steepest where the storm clouds end.

[21] The example evokes Gilchrist's (1977) famous experiments, where subjects matched the lightness of a patch that could be made to appear either in a nearer, darker room, or a farther, lighter room. The perceived location had a substantial effect on subjects' matches, suggesting that the visual system takes the spatial structure of illumination into account when computing lightness. See also Gilchrist (2006: 297ff), Casati & Cavanagh (2019: 66-69).

Figure 4. (a) Hallway into conservatory, (b) storm over sea. Stock Adobe Image, Standard License

Illumination fields contain scalar saddle structures, where the ambient illumination increases towards a point along one axis, and decreases along an orthogonal axis. In Figure 5, each streetlight generates a local, orb-shaped maximum. The sky forms another maximum, extending indefinitely to the horizon. The minima lie off to the sides, in the buildings and trees. These structures induce multiple conflicting gradients. The overlaid scalar skeleton shows a blue saddle line formed between two street lights. Ambient illumination increases away from this line in the horizontal axis, towards each of the street lights. The ambient increases towards this line in the vertical axis, as we move away from the shadowy areas nearer the ground and in the trees. The saddle captures the fact that the light-orbs collide, like two plumes of hot air merging.

Figure 5. Streetlights at dusk with overlaid scalar topological skeleton. Stock Adobe Image, Standard License

1.4.2. Vector Field Topology

Scalar field structures reflect distributions of undirected magnitudes. Vector field topologies, in contrast, characterise the flows of fluids or forces through space. The key structures are *singularities*, which are points (lines, surfaces) where the vector is zero. In illumination and surface illumination fields, singularities occur where there is no dominant direction of illumination or irradiance flow, respectively.

Like critical points, singularities are nonlocal features, typed by the behaviour of the surrounding flow. *Illumination sources* are points around which illumination flows only outwards, producing patterns of *divergent illumination flow*. These should not be confused with light sources in the ordinary sense, like light bulbs: these are glass things containing a filament, not singularities. Nonetheless, when a bulb is switched on, it introduces a source in the illumination field. *Irradiance sources* are points in the surface illumination field, from which irradiance flows solely outwards, producing *divergent irradiance flow*. These commonly occur where a convex object directly faces an illumination source, redirecting the illumination flow down over its surface.

Illumination sinks are points in the illumination field towards which illumination flows only inwards, yielding *convergent illumination flow*. Likewise, *irradiance sinks* are points in surface illumination fields towards which irradiance flows solely inwards, producing *convergent irradiance flow*. These commonly occur in concave depressions in a surface, which redirect irradiance flow towards the base, much as water collects at the bottom of a cup.

Illumination flow saddles are points where illumination flows inwards along one axis, and outwards along an orthogonal axis. Paradigm examples include points where the illumination flows from two opposing sources meet in empty space, and the vectors cancel out. In the streetlight scene above, for instance, there are flow saddles between each neighbouring pair of streetlights, where their spheres of illumination collide. *Irradiance flow saddles* are points where irradiance flows inwards along one axis, and outwards along another axis. An example would be a sphere illuminated from above and below by diffuse sources. The resulting irradiance flows would travel north and south from the poles, meeting at the equator, where they cancel out.

Other important flow-related features include *flux lines* and *tubes*. As defined by Gershun (1939: 107), flux lines are everywhere tangent to the local flow vector within the illumination field. The lines emerge from luminous bodies, and either spread out to infinity, meet another singularity, or terminate at material bodies. Unlike light rays, these lines are usually curved. Following Gershun (1939: 109ff), flux lines can be bundled into tubes, which provide a compact representation of the illumination flow through the scene. We shall see examples below.

Real-world illumination is exceedingly complex. In the simplest case, all light flows into the system from a single source. In reality, there are typically at least two sources of light, including the sun and sky. The quality and respective contribution of these sources varies throughout the day. Objects also scatter light, causing interreflections and highlights, effectively adding further sources. This is captured in higher-order components (n>2) of the illumination field, which vary significantly across scenes, producing 'light textures.'[22] At lower orders, illumination is usually more stable. Empirical analyses have shown that steady flow patterns typically emerge at medium to large scales. These form structures composed of singularities – primarily sources, but also saddles – and connecting flux tubes. These can be represented by *flow topological skeletons*, similar to the scalar skeletons above. These

[22] See Mury et al. (2007), Mury (2009), and Pont (2019: 509). These textures are important in perceiving local variations, like highlights, which provide cues to features like glossiness.

structures are non-accidental, being predicted by characteristics of the illumination source(s) and scene geometry, together with the absorbing and reflecting characteristics of the objects.

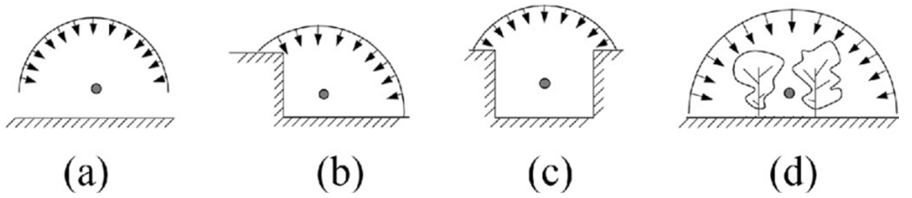

Figure 6. Schematic illumination flow in (a) open field, (b) wall, (c) street, and (d) forest. Mury et al. (2007)

Mury and colleagues (2007) modelled illumination flow fields across four scene-types, schematised in Figure 6: an open landscape, wall, city street, and forest. The primary source was always uniformly bright and hemispherical skylight. To focus on two examples, in the city street, the tall buildings on either side create an opening that channels skylight downwards. This primary flow is oriented away from the centre of the opening, curving downwards and outwards towards the sides of the buildings. Illumination is also reflected and channelled back up between the buildings. This creates a strong, uniform upward flow, which produces saddles where it meets the downward flow. In the forest, the primary flow is also oriented uniformly downward, though with more random variation due to the stochastic nature of openings in the foliage. Unlike the street scene, the forest contains many randomly oriented surfaces, which produce more complex scattering and interreflections, thus less consistent upward flow.

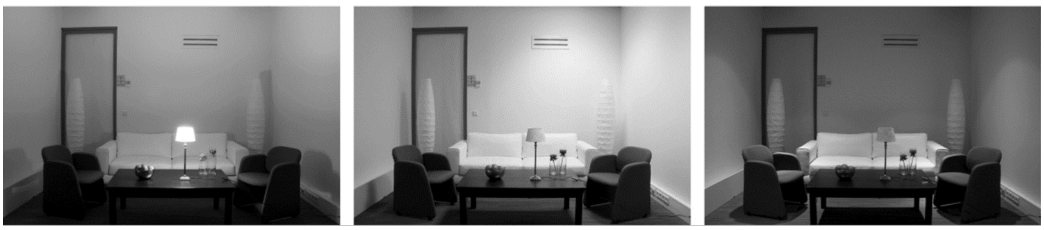

Figure 7. Room illuminated by (a) lamp, (b) diffuse light, (c) collimated spotlights. Source: Kartashova et al. (2016), under Creative Commons.

Kartashova and colleagues (2016) analysed the flow fields within the room in Figure 7, illuminated by (a) a centrally located lamp, (b) diffuse lighting from above and to the right, or (c) collimated spotlights in the left and right of the ceiling. Figure 8a shows the flux tubes generated from physical measurements of the illumination in each condition. In the first condition, the tubes arc out from the position of the lamp, curving down at the top, due to the reflection from the ceiling and between the corners of the walls. The tubes get thicker the further one gets from the lamp, indicating weakening illumination flow.[23] In the diffuse lighting condition, the tubes are straighter and more consistent in width, indicating a smaller effect of reflection from the furnishings and floor, thus more constant illumination flow. In the spotlights condition, there are two clusters of thin, almost vertical tubes on the left and right, representing the collimated beams flowing down from the ceiling. Elsewhere, the tubes are thicker and irregularly shaped and oriented, due to scattering.

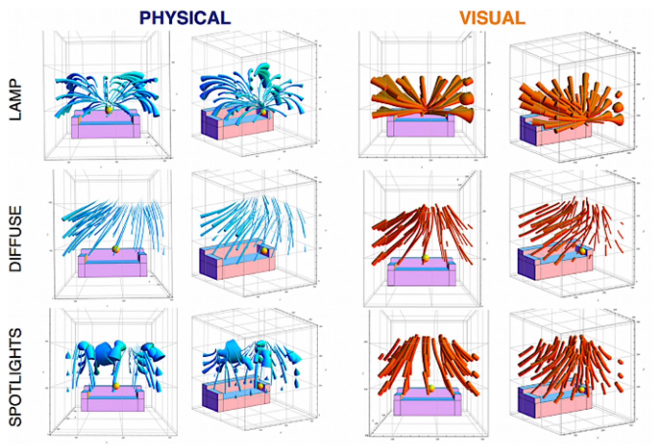

Figure 8. Flux tube representations of (a) physically measured illumination, (b) interpolated visual data. Source: Kartashova et al. (2016), under Creative Commons.

These studies evince the stable, self-organising character of illumination flow fields. Once you fix the source(s) of illumination and scene geometry, predictable patterns emerge. The same generic structures form at large scales, as in the open field, city, and forest; and

[23] Flux tubes are defined such that the total radiant flux received at any cross section is the same. Accordingly, the thicker the tube, the weaker the average flow at each point in the cross section.

medium scales, as in the room. These structures are stable under small changes in illumination and local perturbations in geometry, such as caused by trees swaying in the wind or small moving objects. These are major environmental constancies, which covary with frequently encountered types of illumination and layout.

In the inverse, illumination flow structures carry significant information about the position of light sources, the type of illumination they provide, and the broad layout of objects in the scene. At the largest scales, generic schemas for illumination flow help identify the type of environment one is in. Urban grids channel light between buildings in the same way in New York as New Delhi. Forests form patchy swathes of illumination, partially recreated in city parks. At medium scales, tunnels and hallways produce characteristic tubes of illumination flow along a dominant horizontal axis. Bedrooms, living rooms, and city streets canonically contain orb-like structures, where illumination flows out in all directions from lamps and street lights. Illumination structures are thus diagnostic of the places where they form.

Distributions of illumination flow also help assess the limits of an enclosed space. Flow may be interrupted at walls and corners, producing discontinuities that signal a material boundary. Flow may be reduced by absorption or scattering as it passes through a room, or reflected strongly, giving the illumination a correspondingly soft or hard edge. This is familiar to artists, interior designers, and architects, who exploit these effects to create illumination 'atmospheres.'[24] These pursuits plausibly exploit capacities to perceive distributions and flows of illumination, though these capacities remain under-theorised. The next section addresses this dearth.

[24] On 'atmosphere' perception, see Katz (1911/35: 224-232), Descottes & Ramos (2013), and Albertazzi et al. (2018).

2. The Visual Perception of Illumination

2.1. The Structure-Object Assumption

I now argue that illumination is an object of perception. More precisely, visual perception represents illumination entities: emergent structures in scalar and vector illumination fields. These include paradigm ephemera, like sunbeams, shadow volumes, surface light enclosures, and shadows; amorphous structures, like dappled illumination and the flows of street lights; and spatially homogenous flow and diffusion structures, like large expanses of sunlight, skylight, and twilight. I provide phenomenological and empirical evidence for the perceptual attribution of mereotopological features, which are distinctively associated with such structures. A key assumption is that evidence for such mereotopological attributions is evidence that illumination entities are themselves perceived. Call this the *Structure-Object Assumption*. I start, therefore, by motivating this assumption.

Attributions of mereotopological features are hallmarks of ordinary object perception. Mereological properties—like part-whole relations—are key to explaining the way objects appear as coherent or unified entities, rather than mere collections of qualities. Objects usually appear as having stable organisations of parts, even as they move or change shape. Empirical work suggests that such perceptual organisation follows several psychological principles. For example, the principle of *uniform connectedness* holds that representations as of closed figures or volumes of homogenous colour, lightness, texture, or other material features, tend to form primitive units for an entry level organisation of the perceived scene.[25] Whole units are organised into complex structures, according to Gestalt grouping principles such as proximity, similarity, and common fate.[26] Units are 'parsed' into parts, via principles such as the *minima rule*, whereby interior boundaries between parts tend are attributed where

[25] Palmer & Rock (1994).
[26] See Wagemans et al. (2012a,b) and Wagemans (2015) for synoptic coverage of perceptual organisation.

the surface of the object is locally most concave.[27] These boundaries tend to cut the shortest possible path between surface points, per the *short-cut rule*.

Green (2019) argues that perceived parthood structures are organised hierarchically, representing relations between objects and their major parts; these major parts and their parts; and so on.[28] For example, a chair may be organised into a back, a seat, and four legs. The back may be further parsed into five vertical slats, each connected to a horizonal slat at the top and bottom. The resulting decompositions are closely approximated by medial axis structures, representing points that have two or more closest points on the boundary of the figure. Representations of these kinds account well for perceptual constancies relating to geometrical structure, where objects appear to maintain invariant relations among their parts through transformations caused by motion, stretching, and bending.

Similar points apply to topological features, like connection, closure or boundedness, and holes. Although abstract, attributions of connection relations are implicated in the parthood structures discussed above. For instance, the vertical slats of the chair are plausibly represented as connected to the horizontal slats. Attributions of closure and boundaries are implicated by uniform connectedness, which posits representations as of closed figures. In a similar vein, Goldman (1977: 280) takes perceiving certain 'natural units' to require a 'representation of distinctness,' which 'might be a representation of an edge or boundary...' Explanations of figure-ground also implicate attributions of boundary ownership. As Koffka (1936: 181) explains,

[27] For discussion, see Green (2019: 567ff).

[28] See also Feldman (2003: 252-254), Hummel (2013), Leek et al. (2009), Marr & Nishihara (1978), and Palmer (1977).

> [A] closed contour line, although separated from the rest of the field on either of its sides by the same leap of stimulation, *belonged* to the enclosed figure and segregated it from the surrounding field.

Figural status is thus achieved when contours formed by contrasts between the object and its background drive the formation of representations as of an environmental edge belonging to the object.[29]

In sum, attributions of mereotopological properties are central to empirical and philosophical explanations of object perception. Where such attributions occur, there is good, if defeasible, reason to think that an object is perceived. While property instances or tropes may be assigned to perceived locations, they are not typically attributed part-whole relations, connection relations, or boundaries. The Structure-Object Assumption is thus well-motivated, given the role played by parallel assumptions in explaining ordinary object perception.

I exploit the Structure-Object Assumption in arguing in support of illumination perception. I present evidence that subjects perceptually attribute certain structural properties, as part of a distinctive kind of perceptual organisation relating to illumination. The evidence is phenomenological, concerning how things seem to subjects of visual experience when viewing scenes containing variable illumination; and empirical, concerning certain visually-based discriminations and judgements that subjects can make concerning illumination. The arguments are informed and guided by the preceding ontology of illumination. This ontology allows questions concerning the perception of illumination to be approached with greater precision, in the following respects.

[29] Other studies suggest that representations of topological features facilitate shape discrimination, motion perception, object tracking, and coding in visual working memory. See Chen (2005), Zhou et al. (2010), Wei et al. (2019), and Rips (2020), and discussion in Casati (2009), Davies (2021, 2022), Green (2017: 366ff; 2019), and Burge (2010: 444ff; 2022: 92ff, 297-312).

Firstly, the ontology indicates the kinds of structure inherent in environmental illumination, which could, in principle, be represented in perception. For each type of structure—including boundary, scalar, and flow structure—we can assess whether such structure is apparent to subjects of visual experience in certain conditions, and whether they can discriminate its features. This approach has precedents in the Gestalt psychology, where Koffka (1936: 67) claims that one must 'study the organisation of the environmental field,' to 'find out the forces which organise it into separate objects and events,' and 'how they produce behaviour in all its forms.' Within computational vision science, the approach is exemplified by Marr (1982: 43), who holds that 'the structure of the real world… plays an important role in determining both the nature of the representations that are used and the nature of the processes that derive and maintain them.'[30] I assume familiarity with this approach, thus with the importance of attending to environmental structure in developing perceptual theories.

Secondly, it is widely held that visual systems are adapted to the structural regularities of the environment, and that these regularities constrain and guide perceptual representation. Canonical examples include the tendency for objects to be convex and for illumination to come from above. Such structure is *exploited* by the visual system in the form of priors or assumptions about likely scene configurations, which in turn shape the transformations that produce perceptual representations. These regularities are also *mirrored* in the structure of perceptual representations themselves: they help explain why certain features—such as surface curvatures or patterns of illumination flow—are perceptually attributed in the ways they are, often under conditions of ambiguity. The structure of the environment thus shapes

[30] A major philosophical source is Burge's (1986, 2010, 2022) anti-individualism.

the development and operation of visual mechanisms, and provides a rationale for assigning specific representational capacities to organisms that live in the environment.

I therefore assume that, if illumination entities are perceptually represented, these representations will exploit and mirror the characteristic structure of illumination. Since this structure is defined by scalar and flow field topologies, these topologies should inform the development of perceptual theory. To be clear, the claim is not that visual systems implement the full representational resources of such topologies: that is obviously implausible. Likewise, nobody thinks perceptual representations of parthood structures mimic formal mereology. Rather, the claim is that vision represents simplified, schematic illumination structures, composed via as yet unknown principles. As in other domains of perceptual organisation, these principles are likely to emerge only through an integrated approach: combining knowledge of the actual structure of the environment, phenomenological reflection, and evidence from psychophysics and cognitive psychology. With these assumptions in place, we may now proceed.

2.2. The Perceived Spatial Structure of Illumination

The paradigm objects of perception are solid material things, bounded by a surface. These objects appear apart from their backgrounds, and separate from any surface on which they may rest. Though they may make mechanical contact with such surfaces, they do not appear connected to them, as if carved from a single piece. When they move, they also do not appear to displace the air or change the overall structure of the environment. Instead, they behave as independent actors, interacting through mechanical force, bumping, colliding, or resting against each other, like billiard balls.

In evidencing illumination perception, we must attend to altogether different kinds of structure. As Section 1 showed, in many ways, these structures are more similar to those of large fluid masses than medium solid objects. To fix ideas, I therefore start with a liquid analogy. Imagine being lost at sea, water as far as you can see. The sea appears as a vast, undifferentiated expanse, bounded only by the horizon, at some indiscernible distance. It appears effectively shapeless; it is not perceptually indexed and tracked, or apt to engage 'object-based' attention.[31] It appears stuffy, rather than thing-like. More precisely, every visible part, large or small, appears as of the same liquid kind. This is unlike the perception of a chair, for instance, where only the whole appears as a chair, while its proper parts do not.[32]

Despite its stuffy appearance, the sea appears highly structured, across multiple scales. At the largest scales, some regions appear darker than others, due to clouds of plankton beneath the surface. The water seems to flow in a dominant direction, dictated by the prevailing current. At smaller scales, more intricate patterns emerge, such as waves, swells, and ripples. Latching onto one, you can follow its trajectory, though the structures are transient and ephemeral, liable to dissipate at any time. These structures do not seem to move like billiard balls through space: they propagate through the medium, like vibrations along a string. They do not appear as discrete, disconnected things; rather as integral, connected parts of a liquid mass.

Sometimes, more pronounced structures appear. Imagine a whirlpool forming nearby. You clearly discriminate this structure, attributing rough shape, size, and motion features. The whirlpool nonetheless appears unlike any ordinary object. It appears as composed of a revolving flow pattern, converging on a sink hole. It appears unbounded or open at its extremities, manifesting solely through its interior flow. You could define an outer boundary,

---

[31] Huntley-Fenner et al. (2002) found that infants are worse at tracking poured portions of sand, than rigid, cohesive objects descending along the same path. VanMarle & Scholl (2003) found the subjects are comparatively poor at tracking stuffs that expand and contract as they pour, suggesting that tracking depends on rigid motion trajectories. Gibson (1986: 19-22) discusses distinctions between media, substances, and surfaces in visual perception.

[32] This point mirrors a condition on the semantics of mass nouns, known as 'divisiveness.'

by connecting the farthest points of the spiral, but no such boundary is perceived. The whirlpool appears continuous with the surrounding chop, a flow structure embedded within the sea.

We plainly see such liquid masses, and the structures within.[33] These entities do not appear structured in the same ways as ordinary solid objects. Nonetheless, as any sailor or canoeist will tell you, safely navigating seas and rivers requires 'parsing' their currents, swells, eddies, and whirls, at large and small scales. This is a helpful model on which to think about illumination perception. Imagine standing atop Mont Ventoux on a sunny day, looking south towards the Luberon. The landscape appears bathed in a broad, evenly distributed expanse of bright illumination, washing over every visible surface. The surface illumination appears uniform, and all of one piece, like an unbroken luminous sheet. This sheet appears bounded only by the horizon, at some indefinite distance. Yet it does not appear formless: it seems shaped from beneath, inheriting its contours from the countryside. Indeed, the sheet seems to lie precisely at the surface of the land, not at some remove from it. In short, it appears as a continuous field of surface illumination.

The illumination does not, however, appear merely surface-bound. The countryside appears quite unlike the moon when lit by the sun. Intuitively, the moon seems to intercept only light, not a field of illumination, as the space around appears void. In contrast, the entire space over the countryside seems suffused with illumination, forming a sensible, radiant plenum.[34] It appears as though the land is interrupting this illumination, channelling it over the hills and valleys. The atmospheric illumination thus appears connected to the surface illumination, as all of one piece, showing large-scale structural uniformities. For instance, the brightness and yellowishness of the surface illumination coheres with the bright and

[33] On liquid perception, see Hespos et al. (2009) and Kawabe et al. (2015).
[34] For related perspectives, see Koenderink et al. (2007) and Xia et al. (2014) on the 'visual light field,' Schirillo (2013) on 'inferring light in space,' and Ikeda et al. (1998) on the 'recognized visual space of illumination.'

yellowish atmospheric illumination. The sharply defined shadows of trees and buildings mesh with the highly directional quality of the sunlight. As daylight fails, the environment undergoes a systematic shift in brightness, tending more towards red. The illumination palpably becomes more diffuse. This shift manifests not just in the sky or land alone, but the whole sky-land illumination structure.

These points echo Katz (1911/1935: 42), who claims that,

> there can be no doubt that illumination is perceived with objects and by means of objects; it is equally certain that illumination is perceived in the *empty space* which lies before objects…

Katz considers these 'co-variant phenomena' (1911/1935: 44), in that 'the [perceived] lighting of empty space corresponds to the [perceived] illumination of the objects enclosing and bounding it,' (1911/1935: 272). That is often true, though fine details in surface illumination often are not mirrored in the space around. I expand on this later. In claiming that we perceive illumination 'by means of objects,' Katz means that illumination is perceived, only in situations where we also perceive material objects as such. Looking up only at the blue sky, you see a luminous expanse, but no material things. In these contexts, Katz (1911/1935: 39) says, 'the factor of illumination is not given.' This tracks, as the atmospheric illumination manifests in experience through the ways in which it shapes, and is reciprocally shaped by, the illumination over the landscape. One sees the total volume-surface illumination structure, only if one sees some features of its surface boundaries, much as one typically sees a mass of liquid, only if presented with part of its surface.

This phenomenology aligns with the major topological structure of environmental illumination. First, recall that illumination fields are bounded by two-dimensional illumination surfaces, induced by the surfaces of material objects. These volume-surface structures form coherent, unified systems, where atmospheric light irradiates material surfaces. These illumination and material surfaces are precisely co-located, interpenetrating one another. Second, the features of illumination surfaces—chiefly, irradiance and irradiance flow—are causally determined by the surrounding illumination and geometry of the material environment. In the other direction, the material environment shapes and redistributes the illumination flow within it. This introduces systematic covariations in the intensity and flow of atmospheric illumination, and the irradiance and irradiance flow over the scene. Illumination fields and surface illumination distributions are thus interdependent, with reciprocal constraints.

These gross structural features are borne out in the Ventoux experience. First, the luminous sheet of surface illumination seems to be laid down by the atmosphere, creating an unmistakable unity in the volume-surface structure. The surface illumination appears as precisely following the contours of the landscape, thus as co-located with it. It nonetheless displays its own distinctive intensive, chromatic, and structural features, marking it as an immaterial illumination entity, rather than as part of the landscape. Second, the distribution of these surface illumination features does not appear random, as if from nowhere. It appears as an imprint of the surrounding field, thus constrained by its uniform intensity, colour, and flow. To reinforce this point, suppose a divine hand carved out a two-dimensional slice of the landscape—maintaining its exact surface illumination field—and transported it to outer space. Stripped of its atmosphere, the illuminated terrain would appear strikingly incongruous, with the irradiance distributions unmoored from the field that sustained them. Although the landscape would look similar in some ways, the entire scene would appear structurally different. This difference lies in the illumination surface having been severed from the field of which it was once part.

I have been arguing that the phenomenology of illumination comports with the structural unity of illumination fields. Illumination also appears internally structured, in ways that approximate the scalar topological skeletons of these fields. Let us start with some simple gradient structures. Imagine standing in the gloomy hallway pictured above, looking out to the conservatory. You have the sense of being enveloped by darkness, looking out on a brightly illuminated space. The dark manifests at the walls and floor, in a blanket of shade. Yet this does not exhaust your experience of the gloom, which seems to hang oppressively around you. Light bursts in through the doorway, like a flood of water through a breached dam, spilling out in a widening pool of brilliance. Beyond the doorway, the space appears as a luminous dome, not a mere arrangement of bright surfaces. As you move towards it, you seem to step into the light, not just onto it.

Besides the architectural boundary imposed by the door, the illumination itself seems to partition the scene into two distinct zones, differing markedly in appearance. Indeed, architects sometimes talk in terms of 'light-zones' and 'shadow-zones.' As Madsen (2007: 54) explains, 'varying daylight-levels mean that the space's light-zones are defined individually and perceived as special kinds of "rooms within the room."' On the present account, these zones evince the perceptual organisation of illumination into coarse scalar structures. The illumination itself appears as divided into large, connected expanses, organised around perceived maxima and minima in intensity, joined by an intensity gradient. This perceived structure is rough and ready, and by no means precise. It nonetheless approximates the large-scale structure of the field.

To reinforce this point, imagine the scene unfolding over time, as the sun sets. The illumination in the conservatory gradually dims, its radiance receding. The dark of the hallway expands outward, eating into the space beyond the doorway. The bright zone contracts systematically, both in the retreating pool of light on the floor, and the correlative deflation of the bright dome that once bulged into the hallway. Throughout this change, however, the gross organisation of the zones remains intact. You still perceive two connected expanses of light and dark. Though previously comparable in size, now it seems as if a balloon is being squeezed, constricting the light portion, and making the darkness bulge. Despite these changes, the illumination appears invariant in its global, maximum-minimum structure. To contrast, imagine a dark well suddenly opening up in the middle of the conservatory. Though this has the same net effect of decreasing the illumination there, the change is markedly more salient. It introduces a new local minimum, changing the overall illumination structure. It is like a sinkhole appearing on a smooth slope leading from a valley to a peak. This is in contrast to the sunset, which is more like a valley of darkness gradually expanding, reducing the size of the luminous peak—a noticeable change, but not a gross structural change.

Continuing the analysis of perceived scalar structure, imagine standing on a street at dusk, lit by streetlights. The scene is uniformly and dimly illuminated by the evening sky. The most evident structure, though, appears around the lamps, which form focal points of strong, radiating illumination. Each lamp generates a spherical zone, which gradates into twilight. These zones organise the scene into a series of light-orbs, like glowing pearls on a necklace. As Rasmussen (1986: 208–9) observes, 'light alone can create the effect of an enclosed space. A campfire on a dark night forms a cave of light circumscribed by a wall of darkness.' Unlike pearls or caves, these orbs do not appear precisely bounded. They fade out, eventually meeting at the fringes, like intermingling plumes of mist. The orbs appear as distinct entities, though spatially continuous, their descending gradients meeting in a basin of darkness. This perceived structure loosely approximates the actual topological skeleton, defined by a series of local maxima, connected at saddles. Of course, you do not see the contours that define this structure, any more than you see the medial axis structures that define the skeletons of ordinary solid objects. You rather see the illumination as having an invariant structure, in approximately the conditions under which its skeleton is, in fact, unperturbed. The skeleton thus captures a stable environmental structure, which is mirrored in visual perception.

The 'zone' metaphor builds on Katz (1911/1935: 190), who claims that 'objects, with their colour qualities, are apprehended as belonging in a *chromatically illuminated* field.' To expand, Katz held that variations in illumination prompt a distinctive type of perceptual organisation. As well as grouping and parsing material objects into their major parts, subjects perceptually organise scenes into 'fields' of illumination. Gilchrist (2006) and Zdravković and colleagues (2012) claim that the visual system represents illumination 'frameworks,' which group areas to be treated as under equivalent illumination. These groupings are intended to operationalise Katz's phenomenological fields.[35] On Gilchrist's (2006: 297) account, though, frameworks are subpersonal and pre-perceptual, grouping patches in the retinal image. These contrast with standard Gestalt groupings, which relate to perceptions as of ordinary objects. Frameworks thus do not function to represent illumination entities.[36] This differs from the present account, which posits perceptual representations as of illumination structures, attributing features that approximate scalar and flow field structures. While consistent with the framework approach, it moves well beyond it.

The perceptual organisation of illumination zones is usually coarse, operating at low spatial frequencies. More intricate structures are often perceived in surface illumination, where interactions with materials induce patterns at higher frequencies.[37] Imagine standing in

[35] Cf. van Doorn et al. (2011: 10).
[36] Gilchrist & Solzano (2019) recognise this deficit and address illumination perception more directly, though still incompletely. Compare Adelson & Pentland (1996) on 'adaptive windows' for illumination.
[37] It nonetheless seems that surface illumination phenomena are processed at lower frequencies than material objects and textures. Lovell et al. (2009) found that discrepancies in shadow orientation were detected at lower frequencies when stimuli were perceived as shadows, than when they were perceived as material objects. They infer that shadows are processed by a 'functionally separate, spatially coarse, mechanism.'

a field sheathed in dappled light. You perceive an array of amorphous blobs of brightness, bleeding into a fibrous network of shadows cast by branches above. Many blobs seem to merge, creating bulbous agglomerations of radiance. Although these forms are indistinct and continuous, they do not appear as noise, rather as highly structured. Plausibly, this sense of structure emerges from the appearance of waymarkers of maximum and minimum irradiance, and the gradients connecting them. Like undulating terrain, these peaks and troughs form slopes, leading to pathways or channels, joining structures together. Just as a skilled navigator reads the contours of the landscape, so a skilled perceiver discerns the ebb and flow of surface illumination.

Although transient, these waymarkers and contours appear as relatively stable features of the surface illumination. Suppose the trees are swaying slightly, gaps widening and narrowing in the foliage, so the bright blobs gently expand and contract. This changes the heights of peaks and troughs, but not their overall structure. The most pronounced maxima and minima remain, as do the major channels connecting them, though not all minor pathways. As the wind picks up, these contours morph more violently and randomly, like the surface of a pond in heavy rain. The illumination appears noisy and unstructured, like a random luminance mask. Structure is definitely lost, but despite the noise, you perceive an invariance in certain gross scalar features, such as the spatial average illumination. The overall balance of light and dark remains constant, and perceptibly so.

In dappled illumination, surface illumination appears intricately structured. You do not, however, see correlatively fine structure in the atmospheric illumination. The illumination filling the space beneath the trees lacks the high frequency patterns of the forest floor.[38] Again, though, this illumination does not appear structureless. You see a rhythmic

[38] One possible explanation is that large volumes of illumination are processed in different ways than surface illumination. For example, the former might be integrated more closely with scene or place perception, than surface and object perception. See Wright (2013: 446-447) for related speculation. Such a division might also

modulation of brightness, thrumming with the swaying trees. These modulations vary stochastically, but the pervading sense is that illumination flows consistently and strongly downwards, counterposed by weaker, more variable flow emanating upwards from the grass and leaves. This mirrors the actual structure of forest illumination, as described by Mury and colleagues (2007).

Flow-related patterns in illumination are visible at multiple scales. Whole scenes often appear organised by loosely structured patterns of illumination flow and diffusion, as in the forest. In the Luberon, there appears a vast mass of strongly directional illumination, emanating from a focused source. Unlike diffuse illumination, which can seem thick and claggy, the atmosphere appears crisp and taut, like a river in full spate, surging forcefully in one direction. Unlike sunbeams, which appear as marked ephemera, the atmospheric structure is vast and amorphous. Although it lacks differentiating form, the expanse appears structured by the uniform flow of illumination from the sun.

As another example, take the hallway-conservatory scene. The major flow of illumination is written all over the floor and walls: it manifestly pours in from outside and pools around the doorway. The illumination seems to come at an angle, rather than straight down, as it is redirected across the floor into the hallway. This coheres with the major pattern of surface irradiance flow, which points away from the windows, right at you. The actual illumination flow also correlates with the light appearing to flood through the open door, not just along the floor and walls. This flood emanates from the windows, which appear as large, planar sources of illumination. The windows form unmistakable singularities in the flow: surfaces from which all light flows only inwards.

explain subjects' insensitivity to inconsistencies in surface illumination at different locations within the scene, as discussed by Ostrovsky et al. (2005) and Lovell et al. (2009).

The streetlight scene appears densely structured by intersecting flows of illumination from the sky and lamps. The downward flow from each lamp manifests in an arcing pattern of irradiance spreading over the ground. The area directly beneath the lamp appears to receive the most intense flow. From there, light redistributes outwards, fading into the shadows beneath the trees. A steady, diffuse illumination descends from the sky, filling the space with a low-intensity glow. This ambient flow is punctuated from below by the more forceful, radiating streams of light from the lamps. Though subtler than ocean currents, these patterns are perceptible, organising the scene into a stable structure of intersecting illumination paths. As throughout, this structure mirrors the topological skeleton of the environmental flow field—not in every detail, but in the broad pattern of sources, divergent flows, and their intersections.

In summary, in many conditions, the structure of the illumination within a scene is apparent to subjects of visual experience. This includes the structural unity of illumination fields and their surface boundaries, and scalar and flow-related structures at multiple scales. Many of these structures are associated with perceptual constancies, whereby the illumination appears approximately invariant in respect of its organisation into intensity zones or aggregate flow patterns, despite perturbations across time and space.

Nonetheless, it may be hard to accept that such structures are truly seen, particularly in what is usually considered 'empty' space. It is, then, significant that psychophysical evidence suggests subjects may perceptually represent certain features of the illumination field, even at points containing no objects. Koenderink and colleagues (2007: 7) examined this by introducing a spherical gauge object at random points in the empty space within a scene. Subjects could control the illumination of the sphere, adjusting parameters of horizontal and vertical direction, diffuseness, and intensity. These varied the irradiance and irradiance flow patterns over the sphere. Subjects were told to set the parameters so that the sphere appeared to 'fit' into the scene.[39] Interestingly, their settings tended to approximate the actual illumination field around that point. When interpolated across all probe locations, the settings closely aligned with the global structure of the illumination field.

Kartashova and colleagues (2016) used the same methods to assess subjects' representations of the room in Figure 7, illuminated by a lamp, diffuse light, or collimated spotlights. Subjects' settings again approximated the illumination field immediately surrounding the probe. The interpolated settings were representing using flux tubes, shown in Figure 8b. When compared with the actual measured illumination, they found that the interpolated tubes 'grasp the basic structure of the physical [illumination] field and converge at light sources, but ignore subtle changes due to (inter)reflections,' (2016: 13). In the lamp condition, for example, the tubes fit a coarse 'template' for a diverging illumination flow field, centred on the lamp. The template had fairly straight flux tubes radiating out from the source, in contrast to the curved tubes found in actuality. It also did not incorporate the complex effects of interreflections between the ceiling and walls, nor the increased diffuseness around these surfaces.[40]

One explanation of these results is that, across trials, subjects perceptually represent the illumination field throughout the scene, including at empty locations. On each trial, subjects set the illumination of the probe to cohere with their representation of the field's properties at that location. In this vein, Koenderink and colleagues suggest that (2007: 7),

[39] The methodology was introduced by Gilchrist & Radonjić (2010) to test for functional frameworks of illumination.
[40] Murray & Adams (2019: 50) report that 'people tend to overestimate diffuseness, especially in scenes with relatively directional light, which is consistent with… a prior for diffuse light.' On diffuseness perception, see Pont & Koenderink (2007), Xia et al. (2014, 2017), and Morgenstern et al. (2014).

observers have a notion of the structure of the [illumination] field, even at locations in empty space, quite remote from any visual object. One may say that the [illumination] field in a scene is indeed perceived.

The claim is not that subjects represent the fine-grained structure of the illumination field, but rather coarse, simplified intensity zones and flow patterns. These representations plausibly involve generic templates for illumination from collimated or diffuse sources, adapted and fitted to the particularities of the scene. These templates might function like structural schemas for familiar types of objects, like bodies; or, perhaps more likely, types of scenes, like rural, city, and beach scenes. Just as these scenes tend to contain predictable configurations of objects, thus distributions of colour and lightness, so illumination templates may conform to average tendencies in the structure found in different types of environments, at different times of day, weather conditions, and so forth.

The results are, however, consistent with more deflationary explanations.[41] On the most deflationary explanation, subjects do not perceptually represent the illumination field at all, only illuminated objects. Subjects perform the task simply by inferring how the sphere should look, given the appearance of other objects in the scene. Another, moderately deflationary explanation is that subjects do represent the illumination field, but only in the vicinity of objects perceived—and perhaps attended—in the scene. On each trial, then, subjects form a transient, expendable representation of the illumination field around the sphere, wherever it may appear. Subjects set the illumination of the sphere to cohere with this representation.

The moderately deflationary explanation, if true, still would be significant in establishing illumination perception. While subjects may not perceptually represent full-field

[41] Compare Koenderink and colleagues (2007: 14), Kartashova et al. (2016: 14), and Pont (2019: 520).

illumination structures on this account, they would at least represent the illumination in the space around material objects, when perceiving—and perhaps attending—those objects. Indeed, this is one way to finesse Katz's (1911/1935: 42) claim that 'illumination is perceived with objects and by means of objects… [but also] perceived in the empty space which lies before objects.' As noted above, Katz thus holds that illumination is 'given,' only in situations where subjects perceive material objects as such. The moderately deflationary interpretation similarly holds that illumination is perceptually represented, only in the spatiotemporal neighbourhood of perceived objects.

The most deflationary explanation undercuts the inference from the psychophysical data to perceptual representations of illumination. While it is impossible to discount this explanation completely, it is markedly worse than those that appeal to illumination perception. First, subjects' settings tended to approximate the direction, diffuseness, and intensity of the illumination, even at locations relatively distant from other objects. If subjects were inferring these settings from the appearance of other objects, one might expect performance to decline as the probe is placed at greater distances from them. Yet this is not what was found. Second, the explanation implies that subjects compute a new local illumination setting from scratch on each trial. Each computation will rely on different cues, or weight cues differently, based on the probe's location. Such computations are extremely complex, thus highly susceptible to error. Yet subjects performed the task easily, with little apparent cognitive effort, and high accuracy. Third, as Section 1 attests, illumination fields contain stable, self-organising structures, which emerge from the complex interplay of radiation and matter. Explicit knowledge of this structure is acquired via the special sciences, not apriori. Yet subjects' settings approximate this structure, despite their ignorance of the science. This suggests they have internalised some of the principles and regularities governing these fields. It is certainly possible that this internalised structure forms part of a cognitive 'folk physics' module, for instance, which subjects deploy in inference. Given subjects' reliance on perceptual cues, however, it is more plausible that this internalised structure underlies visual capacities to represent the illumination field. It remains to be seen whether these representations are truly global, or more local, as on the moderate deflationary explanation. For what it is worth, my hunch is that both types of representation ultimately will play a role.

I have been discussing empirical evidence that subjects represent coarse patterns of illumination flow in empty space. Further studies suggest that subjects represent more intricate, high-frequency patterns of irradiance flow. Van Doorn and colleagues (2011, 2012), for example, found that subjects can discriminate between divergent, convergent, and uniform irradiance flow patterns over textured surfaces. Divergent and convergent patterns are defined by flow singularities—irradiance sources and sinks, respectively. These structures covary systematically with surface curvature, as discussed above. The discrimination task was to judge whether an array of disks appeared uniformly convex, concave, or mixed. The irradiance pattern over the disks was varied, yielding either a globally divergent, convergent, or uniform flow; rotational flow; or random flow. Subjects accurately discriminated the first three patterns. The rotational flow patterns, though consistent with uniform convexity or concavity, were treated the same as the random patterns—namely, as mixed.

Van Doorn and colleagues (2012: 468) conclude that subjects perceptually 'synchronise' or 'group' the individual irradiance textures across the disks, into coherent flow structures. These structures create a compelling impression of surface curvature, with disks appearing variously convex or concave. These are clearly perceived features of the disks, if illusory features. The organisation induced by the irradiance groupings, in contrast, does not relate to the disks. It characterises the structure of a surface illumination field, distributed across the entire array of disks. While the disks appear as separate entities, they appear to lie within a unified, structured field of illumination. This perceptual organisation mirrors the actual topological structure of the illumination, defined by flow singularities and the unified patterns of irradiance surrounding them. This organisation is invoked to explain the coherence of subjects' curvature attributions across the array. Yet it is not reducible to the perception of illuminated objects: the explanation requires that subjects represent an overarching illumination structure, which shapes the perceived curvature of the disks in a coordinated way. The perception of shape here depends not just on what lies on the surface, but on what flows across it.

3. Seen the Light?

It is natural to assume that illumination is simply light, thus within the province of quantum physics. Yet illumination is an emergent, ecologically significant kind, with novel features and self-organising structures at macroscopic scales. Illumination fields have distinctive kinds of surface boundaries, which interpenetrate material surfaces. Illumination fields and surfaces also have complex internal structures, characterised by scalar and vector topologies. These structures comprise bona fide entities, on any reasonable interpretation of the term.

This ontology provides a framework for addressing questions about illumination perception in a rigorous, empirically-informed manner. First, it clarifies the kind of structure contained within the illumination—structure that, in principle, might be perceptually represented. Second, given the assumption that perceptual representations exploit and mirror environmental structure, the ontology guides the formulation and testing of hypotheses concerning the representation of specific structural features. I have presented evidence that subjects perceptually attribute structural features distinctive of illumination, and that these attributions yield novel forms of perceptual organisation. The account spans wide range of phenomena, from largescale, homogenous structures—such as sunlight, skylight, and twilight—to small-scale, intricate formations—such as dappled illumination and surface shading patterns. Though markedly unlike ordinary objects, these structures are genuine objects of illumination perception.

Much work remains to be done to develop our understanding of illumination perception. There are potentially fruitful though underexplored connections with natural language expressions for illumination. Strikingly, illumination is often described with dynamic, fluid terms, like ‘flood,’ ‘stream,’ ‘glimmer,’ ‘flash,’ ‘burst,’ and ‘pool.’ Darkness is usually described in more static terms, as ‘still’ and ‘heavy,’ or as a ‘veil,’ ‘cloak’, or ‘blanket.’ Language also makes extensive use of partitive constructions that pick out structured portions of illumination with distinctive sizes, shapes, durations, intensities, and flows. For example, ‘expanse of daylight’ picks out a large, unshaped, persisting portion; ‘shaft of light’ denotes an angular, intensely flowing, transient portion; ‘spot of shadow’ picks out a roughly circular, low intensity, often mobile portion; and so forth. These expressions reveal a shared understanding of illumination as a spatially and temporally structured medium—marked especially by patterns of intensity and flow—which complements the present account of perceptual individuation and organisation.

The semantics of such expressions further supports this picture. Whereas ‘shadows’ is count, ‘illumination’ and ‘shadow’ are mass nouns, like ‘water’ and ‘fog.’ This aligns with the view that illumination perception involves representing the structure of continuous fields, rather than standalone objects. It suggests that illumination perception is best understood as a kind of stuff perception—more akin to our perception of liquids or gases than of solid, bounded things. These semantic features offer evidence of a rich conceptual structure relating to illumination, which mirrors the account developed here, and opens new avenues for inquiry.[42]

References

Adelson, E. H. & Bergen, J. (1991). The plenoptic function and the elements of early vision. In *Computational Models of Visual Processing,* M Landy, JA Movshon (eds.), pp. 3–20. MIT Press.

Adelson, E. H., & Pentland, A. P. (1996). The perception of shading and reflectance. *Perception as Bayesian inference*, *409*, 423.

Akins, K. A., & Hahn, M. (2014). More than mere colouring: The role of spectral information in human vision. *The British Journal for the Philosophy of Science*, 65(1), 125–171.

Albertazzi, L., Canal, L., Chiste, P., Micciolo, R., & Zavagno, D. (2018). Sensual light? Subjective dimensions of ambient illumination. *Perception* 47(9):909-926.

Allen, Keith. (2016). *A Naïve Realist Theory of Colour*, Oxford University Press.

Brown, Derek. (2014). Colour layering and colour constancy. *Philosophers' Imprint*, 14(15): 1-31.

Burge, Tyler (1986). Individualism and psychology. *The Philosophical Review* 95.1: 3-45.

——— (2010). *Origins of Objectivity*. Oxford University Press.

——— (2022). *Perception: First Form of Mind*. Oxford University Press.

[42] Thanks to audiences at the University of Toronto and the European Society for Philosophy and Psychology in Prague, and attendees of a graduate seminar at Oxford for discussion of earlier versions of this work; to two reviewers for this journal for their stimulating suggestions and questions; and to Patrick Cavanagh, Mohan Matthen, Eliot Michaelson, Richard Murray, and Ian Phillips for written comments. Special thanks to Mike Martin for continued guidance and inspiration in bringing the project to life.

Casati, R. (2009). Does Topological Perception Rest Upon a Misconception About Topology. *Philosophical Psychology*, 22:1, 77-81.

Casati, R. & Cavanagh, P. (2019). *The Visual World of Shadows*. MIT Press.

Casati, R. & Varzi, A. (1994). *Holes and Other Superficialities*. MIT Press.

Chen, L. (2005). The topological approach to perceptual organization. *Visual Cognition*, 12:4, 553-637.

Davies, W. (2016). Color constancy, illumination, and matching. *Philosophy of Science*, 83(4), 540-62.

——— (2018). Colour vision and seeing colours. *The British Journal for the Philosophy of Science*, 69(3), 657–690.

——— (2021). Colour relations in form. *Philosophy and Phenomenological Research*, *102*(3), 574-594.

——— (2022). The paradox of colour constancy: Plotting the lower borders of perception. *Noûs*, *56*(4), 787-813.

Descottes, H, and Ramos, C. (2013). *Architectural lighting: designing with light and space*. Princeton Architectural Press.

de-Wit, Lee, David Milner, and Robert Kentridge. (2012). Shadows remain segmented as selectable regions in object-based attention paradigms. *i-Perception* 3.2:150-158.

Feldman, J. (2003). What is a visual object? *Trends in Cognitive Sciences,* 7(6), 252–256.

Gershun, A. (1939). The light field. *Journal of Mathematical Physics.* 18:51–151. Transl. P Moon, G Timoshenko.

Gert, J. (2017). *Primitive colours*. Oxford University Press.

Gibson, James. J. (1986) *The ecological approach to visual perception*. Psychology Press.

Gilchrist, A. (1977). Perceived lightness depends on perceived spatial arrangement. *Science* 195.4274: 185-187.

——— (2006). *Seeing black & white.* New York, NY: Oxford University Press.

Gilchrist, A., Delman, S., & Jacobsen, A. (1983). The classification and integration of edges as critical to the perception of reflectance and illumination. *Perception & Psychophysics*, *33*, 425-436.

Gilchrist, A., & Radonjić, A. (2010). Functional frameworks of illumination revealed by probe disk technique. *Journal of Vision* 10.5:1-12.

Gilchrist, A., & Soranzo, A. (2019). What is the relationship between lightness and perceived illumination. *Journal of Experimental Psychology: Human Perception and Performance* 45.11: 1470-1483.

Goldman, Alvin. (1977). Perceptual Objects. *Synthese*, 35:257-84.

Green, EJ. (2017). A Layered View of Shape Perception. *British Journal for the Philosophy of Science*, 68:355-387.

——— (2019). On the perception of structure. *Noûs* 53.3: 564-592.

Hespos, S. J., Ferry, A. L., & Rips, L. J. (2009). Five-month-old infants have different expectations for solids and liquids. *Psychological Science*, *20*(5), 603-611.

Hilbert, David. (2005). Color constancy and the complexity of color. *Philosophical Topics* 33(1):141-158.

Humphreys, P. (2016). *Emergence: A Philosophical Account*. Oxford University Press.

Hummel, J.E. (2013). Object recognition. In D. Reisburg (ed.), *Oxford Handbook of Cognitive Psychology,* pp.32–46. Oxford: Oxford University Press.

Huntley-Fenner, G., Carey, S., & Solimando, A. (2002). Objects are individuals but stuff doesn't count: Perceived rigidity and cohesiveness influence infants' representations of small groups of discrete entities. *Cognition* 85(3):203-221.

Ikeda, M: Shinoda, H; & Mlzokami, Y. (1998). Three dimensionality of the recognized visual space of illumination proved by hidden illumination. *Optical Review*, 5: 200-205.

Jagnow, R. (2010). Shadow-experiences and the phenomenal structure of colors. *Dialectica, 64(*2), 187–212.

Kartashova, T; Sekulovksi, D; de Ridder, H; Pas, S; Pont, S. (2016). The global structure of the visual light field and its relation to the physical light field. *Journal of Vision* 16(10):9, 1-16.

Katz, David. (1911/1935). *The world of colour*. R.B.Macleod & C.W.Fox (trans.) Routledge.

Kawabe, T., Maruya, K., Fleming, R. W., & Nishida, S. Y. (2015). Seeing liquids from visual motion. *Vision research*, *109*, 125-138.

Koenderink, J. J., & Pont, S. C. (2002). Texture at the terminator. In *Proceedings. First International Symposium on 3D Data Processing Visualization and Transmission*, pp. 406-415. IEEE.

Koenderink, J; Pont, S; van Doorn, A; Kappers, A; Todd, J. (2007). The visual light field. *Perception* 36:1595–1610.

Koffka, K. (1936). *Gestalt psychology.* NY: Harcourt, Brace & World.

Leek, E.C., Reppa, I., Rodriguez, E., & Arguin, M. (2009). Surface but not volumetric part structure mediates three-dimensional shape representation: Evidence from part–whole priming. *The Quarterly Journal of Experimental Psychology,* 62(4), 814–830.

Lovell, P. G., Gilchrist, I. D., Tolhurst, D. J., & Troscianko, T. (2009). Search for gross illumination discrepancies in images of natural objects. *Journal of Vision*, *9*(1), 37-37.

Madsen, M. (2007). Light-zone(s): as Concept and Tool. *Enquiry The ARCC Journal for Architectural Research*, *4*(1):50-59.

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information.* San Francisco: W.H. Freeman.

Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London, B: Biological Sciences,* 200, 269–294.

Matthen, M. (2010). How things look (and what things look that way). In *Perceiving the World*, ed. Nanay, B., 226-253. New York: Oxford University Press.

——— (2018). Ephemeral Vision. In *Perceptual Ephemera*, Thomas Crowther & Clare Mac Cumhaill (eds.), Oxford University Press. pp.312-339.

Morgenstern, Y; Geisler, W.; & Murray, R. (2014). Human vision is attuned to the diffuseness of natural light. *Journal of Vision* 14.9: 15-15.

Murray, R. F., & Adams, W. J. (2019). Visual perception and natural illumination. *Current Opinion in Behavioral Sciences*, *30*, 48-54.

Mury, A. (2009). *The light field in natural scenes*. PhD Thesis, Delft University of Technology.

Mury, A; Pont, S; & Koenderink, J. (2007). Light field constancy within natural scenes. *Applied Optics* 46.29: 7308-7316.

Noë, A. (2004). *Action in Perception*. Cambridge.: MIT Press.

Ostrovsky, Y; Cavanagh, P; & Sinha, P. (2005). Perceiving illumination inconsistencies in scenes. *Perception*, 34:1301-1314.

Palmer, S.E. (1977). Hierarchical Structure in Perceptual Representation. *Cognitive Psychology,* 9, 441–474.

Palmer, S. & Rock, I. (1994). Rethinking perceptual organization: The role of uniform connectedness. *Psychonomic bulletin & review* 1.1: 29-55.

Phillips, I. (2018). No more than meets the eye: shadows as pure visibilia. In *Perceptual Ephemera*, Thomas Crowther & Clare Mac Cumhaill (eds.), Oxford University Press. pp.172-193.

Pont, S. (2019). Light: Toward a Transdisciplinary Science of Appearance and Atmosphere. *Annual Review of Vision Science*, 5, 503-527.

Pont, S. & Koenderink, J. (2004) Surface illuminance flow. In *Proceedings. 2nd International Symposium on 3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004,* pp. 2-9. IEEE.

——— (2007). Matching illumination of solid objects. *Perception & Psychophysics*, 69(3), 459-468

Ramachandran, Vilayanur S. (1988). Perception of shape from shading. *Nature* 331.6152: 163-166.

Rasmussen, S. E. (1986). *Experiencing Architecture.* Cambridge, MA: MIT Press.

Rensink, R. A., & Cavanagh, P. (2004). The influence of cast shadows on visual search. *Perception*, *33*(11), 1339-1358.

Rips, L. (2020). Possible objects: topological approaches to individuation. *Cognitive Science*, 44(11), 1-35.

Schellenberg, S. (2008). The Situation-Dependency of Perception. *The Journal of Philosophy*, 105(2), 55–84.

Schirillo, J. (2013). We infer light in space. *Psychonomic Bulletin & Review*, 20:905–915.

Smith, S. R. (2007). Continuous bodies, impenetrability, and contact interactions: The view from the applied mathematics of continuum mechanics. *The British Journal for the Philosophy of Science*, 58 (3):503-538.

Sorensen, R. (2008). *Seeing Dark Things*. Oxford University Press.

Stroll, A. (1988). *Surfaces*. University of Minnesota Press.

Tahko, T. & Lowe, E. J. (2020). Ontological Dependence. *The Stanford Encyclopedia of Philosophy* (Fall 2020 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2020/entries/dependence-ontological/>.

Thompson, E. (1995) *Colour Vision: A Study in Cognitive Science and the Philosophy of Perception*. London and New York: Routledge.

Tye, M. (1996). Perceptual Experience is a Many-Layered Thing. *Philosophical Issues*, 7:117–26.

van Doorn, A; Koenderink, J; & Wagemans, J. (2011). Light fields and shape from shading. *Journal of Vision*, 11(3):21, 1-21.

van Doorn, A; Koenderink, J; Todd, J; & Wagemans, J. (2012). Awareness of the light field: the case of deformation. *i-Perception*, 3:467–480.

VanMarle, Kristy, and Brian J. Scholl. (2003). Attentive tracking of objects versus substances. *Psychological Science* 14(5):498-504.

Wagemans, J. (2015). *The Oxford Handbook of Perceptual Organization*, (ed.) Oxford University Press.

Wagemans, J., Elder, J., Kubovy, M., Palmer, S., Peterson, M., Singh, M., & Von der Heydt, R. (2012a). A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure–ground organization. *Psychological bulletin*, 138(6):1172-1217.

Wagemans, J., Feldman, J., Gepshtein, S., Kimchi, R., Pomerantz, J., Van der Helm, P., & Van Leeuwen, C. (2012b). A century of Gestalt psychology in visual perception: II. Conceptual and theoretical foundations. *Psychological bulletin* 138(6) 1218-1252.

Wei, N; Zhou, T; Zhang, Z; Zhuo, Y; & Chen, L. (2019) Visual working memory representation as a topological defined perceptual object. *Journal of Vision*, 19(7):1-12.

Wright, W. (2013). Color constancy reconsidered. *Acta Analytica*, *28*(4), 435-455.

Xia, L; Pont, S; & Heynderickx, I. (2014). The visual light field in real scenes. *i-Perception* 5:613–629.

——— (2017) Light diffuseness metric part 1: theory. *Lighting Research and Technology* 49:411–427.

Zdravković, S., Economou, E., & Gilchrist, A. (2012). Grouping illumination frameworks. *Journal of Experimental Psychology: Human Perception and Performance* 38.3: 776-784.

Zhou, K., Luo, H., Zhou, T., Zhuo, Y., and Chen. L. (2010). Topological change disturbs object continuity in attentive tracking. *Proceedings of the National Academy of Sciences* 107(50):21920-21924.