

Vision, an Optical User Interface

Perception

2019, Vol. 48(7) 545–601

© The Author(s) 2019

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/0301006619853758

journals.sagepub.com/home/pec**Jan Koenderink** 

University of Leuven (KU Leuven), Belgium; Justus Liebig University Giessen, Germany; University of California at Berkeley, CA, USA; Utrecht University, the Netherlands

Abstract

I present a personal account of visual awareness, from a biological, that is, ethological and evolutionary perspective. The facts I select are mainly phenomenological, with a sprinkle of psychophysics, and a snuff of brain theory. This is because I consider visual awareness to be a phenomenological *brute fact* that has no equivalent in either physiology or cognitive science. My emphasis is on bare fundamentals and integration. I primarily aim at academic understanding.

Keywords

vision, awareness, user interface, Gestalt, psychogenesis, local sign, scale-space, shape from X, Umwelt

Date Received: 27 November 2018; accepted: 5 May 2019

In this article, there are no new data or dramatic new insights. I give my, at times bold, opinion on a rather large number of topics. Here is a summary of the topics, which might help the reader to find structure in this article:

- Introduction: The introduction sketches a biological account of the psychogenesis of perceptual awareness. It departs from a general discussion of the various scientific approaches relating to the topic of sentience (*Type III Vision, Objectivism and Physicalism*) and ethology (*Ethology: The Sensorimotor Loop, Ethology: The New Loop*). This is followed by a lengthy description of a model of psychogenesis (*Habits, Types and Phantasmatic Self-Stimulation through Distributed awareness*);

Corresponding author:

Jan Koenderink, Laboratory of Experimental Psychology, University of Leuven (KU Leuven), Tiensestraat 102 – Box 3711, Leuven B-3000, Belgium.

Email: koenderinkjan@gmail.com

- How Vision Proceeds: This section discusses vision as an action (*Units of Looking*) and the nature of awareness (*The Specious Present, Concrete Actuality*);
- User Interface Elements and Their Effect on Behaviour: In this important section, I discuss familiar topics from the phenomenology of visual perception (*Fixed Action Patterns of Animal Ethology*), with an emphasis on spatial structures (*The Classical Geometric Illusions through Perspective*). I then consider the coherence of awareness (*Temporal Coherence through Multimodal Coherence*) in various domains. Finally, I consider why perception “makes sense” instead of being “veridical” in the mainstream sense (*Making Sense*).
- Space: I place the topic of space central; thus, I consider “geometry” (*The Notion of a “Point” through Points as Samplers*), then discuss an “ideal model” of geometry (*Scale-Space through The Atlas Model*).¹ As an example, I consider “edges” (*Edges*), which is only one case out of many (*Encapsulated Geometry, Points as Brushes and the Canvas*). The important point here is Lotze’s notion of “local sign” (*Topology of an Array of Brush Touches through Nature of the Array of Brush Points*), an issue that has been ignored in the sciences since the mid-19th century.
- Qualia Dimensions: I discuss a simple formal model that ought to be in everyone’s toolbox, yet is ignored in current vision research (*Pictorial Space as a Fiber Bundle, Image Transformations*);
- Heterotopic Areas: This is a generic issue that requires a major text; here, I mention only the basics (*SFX and Ambiguities*).

In the initial sections, I introduce VISION as an “Optical User Interface” (or OUI for short). Although not such a common notion today (but see Hoffman, 2009; Koenderink, 2011, 2019), it has been a common notion in biology since the early 20th century (Uexküll, 1909).²

Introduction

The question “What is Vision?” is not readily answered by looking up “vision” in a dictionary, but nevertheless, it is fun and perhaps instructive to do so.

The Merriam-Webster dictionary lists numerous meanings, most of which do not seem to answer the query to the satisfaction of a scientist. Some might appeal to the philosopher who takes visions as *judgements*; others might be of interest to the artist who thinks of the power of the *imagination*.

The one definition (listed as the fourth meaning of *The act or power of seeing (sight)*—notice the circularity) that might appeal to the scientist is a longish one, describing a rather speculative theory of how light stimuli lead to representations, probably in the brain. This may be interpreted as roughly the Marrian (Marr, 1982), or INVERSE OPTICS account. Thus, it explains vision as the result of a computation.

The dictionary answers are entertaining but confusing. Might it have to do with the status of “vision science”? Few fields that call themselves a SCIENCE are what they advertise (Christian Science comes to mind; Rocket Science is hardly an academic endeavour).

Indeed, also many, mutually very different, views may be (or are) held in the sciences. To focus the discussion, I propose three short, ontologically distinct characterizations of “vision,” intended as no more than “operational definitions,” but useful in this context because of their succinctness:

- Type I electrochemical activity of the brain caused by irradiation of the retina;
- Type II optically guided behaviour; and
- Type III awareness on opening the eyes in full daylight.

Type I addresses the *physiology*, type II the *ethology*, or behaviourist psychology, and type III the experimental, or observational *phenomenology*. There are also “mixed” cases, such as psychophysics, which is perhaps best characterized as “dry physiology,” say type I’.

A typical, early account of type I is Hartline’s work on the Limulus eye (Hartline, Wagner, & Ratliff, 1956). Contemporary work includes brain scan studies of human subjects (Zeki, 2003). One studies physical reactions caused by stimulating the eye with spatio-temporally structured radiant power spectra. Structural complexity measures may be used to describe information transfer in Shannon’s (1948) sense. Qualities and meanings play no role. Thus, mere “structure” is not distinguished from “data” or even “information” (as in the common understanding). This is a very artificial move, for mere chaos no doubt has structure, but to describe it as data, or information goes against the grain of common sense.

Psychophysics (type I’) is similar, an early example is Maxwell’s (1855) work on the discrimination of compound radiant spectra. This concerns mainly absolute and discrimination thresholds. For instance, Maxwell’s work does not address “colour” as quality. It was irrelevant to his purpose whether his observers were aware of a “yellow” or “blue” paper. Type I’ ignores the awareness of the experimental subject, so even sleeping, or decerebrated animals are routinely used in experiments. This is not an essential issue when awareness is not recognized to exist, or play some causal role. Psychophysics has been very successful as dry physiology and has found numerous important applications in ergonomics.

A typical, early account of type II is von Uexküll’s (1909) *Outer World and Inner World of Animals*. One studies observable behaviour as influenced by ecological optics. There is a broad spectrum from pure behaviourism (“rat racing,” speech as the “movement of air molecules,” see Skinner, 1976) to ethology (Lorenz, 1973; Tinbergen, 1951) and Gibson (1978)’s affordances. Here, the relevant quantities are “information” and “meaning” as aids in improving biological fitness (Wassersug & Wassersug, 1986). Qualities as such play no role. The “meanings” are largely defined by the external descriptions of behaviour, or even the physical stimulus. For Gibson, “throwability” is a property of hand-sized stones, much like their weight or size, although he would grant that humans might pick up this affordance, whereas birds would not.³ If meaning is put in the environment, or ecosystem, there is no phenomenology left, only behaviourism, perhaps in some derivative form.

Most of modern “cognitive science” of vision fits here. The experience of the subject is considered to be of interest, although as a necessarily “hidden” parameter. “Experience” has no objective meaning, so it cannot feature as such. The awareness of the experimenter is considered irrelevant, as it would jeopardize objectivity. Objectivity implies *observations without observers* (in general) and—in experimental psychology—without experimenters too. The ultimate ideal of objectivity is meaningless facts collected in double-blind experiments. I find myself pretty much a dissenter here, although I regularly publish under the banner of cognitive science.

Human subjects are (in principle) not treated differently from rats. The differences are only in the degree of political correctness and acceptable politeness. Human observers are received with a cup of tea, whereas rats are simply set to their task—one has already made sure they’re almost dying from thirst, so they will work till exhaustion for a drop of water. It is considered acceptable to do this to monkeys too, but—in our present society—not to humans.

A typical account of type III is Kanizsa’s (1980) *Grammar of Vision* (Koenderink, 2012f). Here, qualities and meanings play the key role. Both subjects’ and experimenter’s awareness count. Intersubjectivity replaces objectivity. This essentially rules out animal experiments.

Titchener's (1927) *Experimental Psychology: A Manual of Laboratory Practice* starts with the following definition:

A psychological experiment consists of an introspection or a series of introspections made under standard conditions.

(The I-word is not used in Newspeak—Orwell, 1949—but see later.) Both observer and experimenter were trained to understand both themselves and each other, very different from the double-blind experiments on naive subjects, the ideal of type II. Often the experimenter and observer were the same person, strictly a *no-go area* nowadays. Yet Titchener was hardly the deluded victim of his time that he is sometimes made out to be. I can only hope that my brains look as impressive as his (Wilder Brain Collection of Cornell).

Titchener's writings make sense to me. Of course, I slightly correct for the times (just consider what people will think of your writings a century from now!). They make good sense in terms of my own experiences in doing science. It is largely a matter of perspective. I do science because I want to understand, whatever it takes. Understanding the unknown means exploration with an open mind, taking advantage of any tool you can forge or find. Attempting to reject silly hypotheses on the basis of double-blind experiments has no part in that. I'm hunting for surprises; that's where my understanding grows. Instead of feigning hypotheses, I often pursue hunches.

In science, the NEWSPEAK installed after WW-II is still the current standard. The "linguistic turn" as one calls it in philosophy (Ferraris, 2013) had wide-ranging effects in the sciences. Biology excommunicated "vitalism," behaviourism "introspection" and philosophy "phenomenology." Great scientists' careers were sacrificed for the common good. Several lost a reputation gained through a lifetime of dedicated work. (Some of their works are now—after decades—being reprinted.) "Cognitive science" made behaviourism *salonfähig*, and the I-, P- and C-words were removed. The nineties came with minor corrections, for philosophy really needed the C-word, so it was suitably redefined. Words and concepts that were banned decades ago returned with new meanings, a DOUBLESPEAK familiar from the news media (e.g., "We contracted some collateral damage" for "we killed a thousand civilians in serving [=bombing] the city"). It works in science much like it does in politics. Occasionally—and more honestly—one uses "post-" or "neo-" prefixes, as is common in the arts. In fact, one is currently running out of "posts-" or "neo-neo-s."

This article is mainly concerned with type III, although I draw freely and indeed gratefully on the materials gathered by the type I, I' and II communities: In my view, the three approaches ought to be considered mutually *complementary*. However, in actual practice, the relative merits usually remain hidden, and there is not a little cheating in all directions—as when brainscan studies refer to qualities and meanings, or phenomenologists refer to brainscans, whereas, strictly speaking, neither has even the slightest connection to the other. At least no connections that are (or, even worse, can be) scientifically meaningful. Of course, it makes sense to keep the public or the granting agencies happy. I'm all for that. But I feel even esteemed colleagues tend to be deluded themselves at times.

I treat type I, I' and II accounts as "known," because I expect most readers will have roots there (so have I). Throughout the article, I will—perhaps unduly—stress differences between the three approaches, as these may well appear strange, or even blatantly wrong, to researchers with different backgrounds.

Little of what I say is to be considered "novel"; in fact, most is far "outdated." Many of the key ideas were already forgotten—originally *suppressed*—during the first half of the 20th century. As a student (I still am), it took me years to figure that out, because nobody tells

you. I merely put a variety of these concepts into a coherent explanatory framework of “what vision is,” trying to place credit where it belongs.

Lecturing on such lines brought me rather mixed reactions, to put it mildly. So, if it really irritates people that much, the topic cannot be entirely without current interest. I attempt a coherent, critical account of a few of the major conceptual threads. More important, I stick my neck out in suggesting potential roads to progress. Remember that I’ve often been proven wrong.

Type III Vision, Objectivism and Physicalism

From a phenomenological perspective awareness (Koenderink, 2012a, 2012d) is a *brute fact* (Fahrbach, 2005) of life. It is a “brute fact” because it “just happens,” it is not something one does, nor something one could do something about. It does not allow for a scientific explanation.

You may attempt to shut down visual awareness by closing your eyes, but imagery never completely goes away in the waking state, or even during sleep.

Awareness is *presentation*, not representation. What could be “represented”? Only what was originally derived from presentation. Except when you count a footprint on the beach seen by nobody as a “representation,” thus rendering the word meaningless.

Awareness is the ultimate objective, *physical* (from the Greek *φυσική* “nature”) fact, unlike such objects as electrons, which are perhaps partly familiar through remote inference (Dirac, 1928), but are never part of concrete actuality.

Because it is a BRUTE FACT, awareness cannot be understood in terms of anything else. It just irreducibly is. Yet, apart from a few philosophers and scientists (Huxley, 1874), few would deny its existence as a mental entity. Indeed, it appears to me to be logically impossible to do so.

“Presentation” is not reflective thought, for (reflective, logical) thinking one *does*. “Presentation” is not judgement either, for judgement comes from careful logical thinking. Nor is “presentation” introspection, for that is just another activity. No one understands how presentations come about, they simply happen, the phenomenon is known as “visual awareness.” “Scientific explanations” are not called for since awareness is not a scientific fact. There can be no “scientific understanding” of it, like one can’t have a science of square circles. But I take it we’re all familiar with it, or we wouldn’t be here.

How does one study awareness? The only available method is experimental or observational phenomenology (Albertazzi, 2013). This used to have a bad name, because supposed to involve “introspection” (the I-word!). However, awareness does not necessarily involve introspection (or reflective thought) at all.

When seeing a tree on opening the eyes, is that “introspection”? Better call it simply a brute fact. For instance, to say that

“one is aware of oneself looking at and successfully seeing a tree, due to there being a tree in front of one, sufficient diffuse radiative power of a certain kind, and so forth”

is an abomination in that it totally fails to capture the phenomenology. Ask any artist, these are the professionals in these matters (Koenderink, 2018).

How does one reason about awareness, or treat it in discursive thought? This involves looking at oneself, or another person as “from the outside.” This is a problem since the eye cannot see itself in any sense of the term (“...for the eye sees not itself, but by reflection, by some other things,” Shakespeare, *Julius Caesar*). Indeed, such exterospection is far more

problematic than introspection. I think of it as a form of “mindfulness,” trying to bring some aspects of visual awareness into reflective thought without so much as interrupting its natural flow (Koenderink, 2018). This is hard, and it never works if you have to try. It only works if it becomes easy, that means natural, unnoticed.

I glean essentially two gateways in pursuing the matter of making sense of awareness by other means than phenomenology. One is by way of animal ethology, the other human psychiatry, especially the study of agnosias, or “soul blindnesses” (Lissauer, 1890).

In either case, one assumes something akin to “God’s Eye” with respect to the subject (Koenderink, 2014). In the case of the patients, one may sometimes establish some empathic bond too. In the case of animals, this tends to be harder (cats, dogs), or even impossible (octopuses, fleas). Animals may have fully alien kinds of awareness. Think of echolocating bats, or sharks prospecting by electroreception (Nagel, 1974).

The latter may be an advantage in the type I and II modes of vision science because it promotes objectivity. Blindness is no hindrance to being a great vision scientist in these modes. However, it bars the type III approach.

I borrow important concepts from Jakob von Uexküll (1920), the founder of ethology, a biologist, and from Jason Brown (1977, 1991), a psychiatrist–philosopher of mind.

I proceed to introduce a number of key concepts.

Ethology: The Sensorimotor Loop

In his *Theoretical Biology*, von Uexküll (1920) introduced the sensorimotor loop, perhaps the most useful functional account of the simplest systems, like sea urchins or scallops. The loop connects a sensor and an actuator through the animal’s body and through its environment.

The loop is not a physiological entity, but a functional description in the mind of a scientist (here, von Uexküll). The loop captures the essence, its physiological implementation being “accidental.”

A sensor is a transducer that when stimulated by some physical process, inserts a “message” (or token) in the loop. Most sensors are like mousetraps in that they can be triggered by fairly specific physical or chemical processes. Often various types are present. The message is something very different from the stimulus, perhaps a hormonal secretion or an electrochemical process. Its role is to trigger the actuator.

The actuator could be a gland, a muscle, or what have you. Its function is to insert a message (another token) into the causal nexus of the world. Again, that message is very different from the one received from the sensor. It could be a mechanical push, an acoustical disturbance of the air or a photic emission say.

I use “message passing” or “messaging” in the sense commonly understood in object oriented programming (oop, Hewitt, Bishop, & Steiger, 1973). One advantage of oop is that objects are mutually independent. All that is needed, is the (hardwire) syntax of the interface. Messages are one-way and do not involve mutual understanding or a common code or language. All that is needed is that a message be “in the right format,” that is to say, is “a key that fits the lock” of the receiver.

The message in the causal nexus results in a volley of causes and effects, some of which may trigger the sensor and so close the loop. Other effects are lost to the loop, but do change the environment. This may affect other loops, leading to a mutual coupling of loops by way of the environment.

Notice that sensors and actuators do not need to know of each other’s existence. Also notice that the “environment” includes the body of the animal.

Loops are tricky objects because they play havoc with causality. In a closed loop, there is no before or after. The chicken is not before the egg, nor the egg before the chicken (Figure 1); thus, there is no such a thing as a causality dilemma here. The loop just cycles. This is what happens in Gibsonian “resonance” (Gibson, 1951); the perception becomes really “direct” because the temporal order is undefined. “Beyond time” is perhaps even more apt than “direct.”

Ethology: The New Loop

In his *Theoretical Biology*, von Uexküll (1920) introduced a very crucial addition to the sensorimotor loop. He called it the “new loop.” It is slightly more complicated, though still very simple.

In many animals, the path from the sensor to the actuator gains a return path that sends the sensor a “prognosis” of what to expect based on the current action. von Uexküll cites Helmholtz (1867) as the originator (Ernst Mach had the same ideas), mentioning his experiments with the pressed eyeball (see Grüber, 1995 for history).

Much later the principle became known as “von Holst & Mittelstaedt’s *Reafference Principle*” (von Holst & Mittelstaedt, 1950). Curiously, but perhaps afraid of being branded “vitalists” themselves, these authors never mention von Uexküll, although they knew him well.

von Uexküll understood the fundamental importance of the concept and considered it the initial, crucial step in the evolution of brains. He was convinced that minds are not explained or generated by brains: “. . .there is no bridge that leads from the nervous system to the soul. . .” This reflects Leibniz’s notion as exemplified by the example of a mill: stepping into the interior, you see only cogwheels, axles, pulleys and cords. The “soul of the mill” is nowhere in sight. Notwithstanding, von Uexküll *did* speculate on the “inner worlds” of his animals.

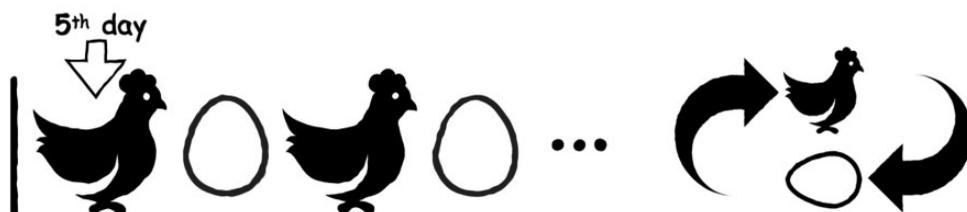


Figure 1. In loops (*right*), there is no temporal order (though, perhaps, a temporal sense) and neither beginning nor end. This solves many problems in biology, as here the infamous “chicken-or-egg dilemma.”

Notice that Genesis (King James Version) simply has: **G1:** [23] . . . the fifth day. [24] And God said, Let the earth bring forth the living creature after his kind, cattle, and creeping thing, and beast of the earth after his kind: and it was so. [25] And God made the beast of the earth after his kind, and cattle after their kind, and every thing that creepeth upon the earth after his kind: . . . , so for the believing mind there is no doubt that the chicken was first (left).

My view is rather that the fast chicken-and-egg cycle is like an eddy (*right*) in the flow of the slow evolution of the chicken tribe. Since the only timescale that is observable is the fast one, the chicken and the egg do not stand in any temporal sequential order. This solves many of the problems raised by “vitalist” notions. Almost all these problems involve “the future acting back on the present” (retrocausation), a kind of teleology.

The basic principle has been reinvented many times and is known under various names. The latest (I think) version of the new loop principle is known as “predictive coding” (Friston & Kiebel, 2009).

In more intricate animals, the new loop involves a “loop controller” (my term) as a two-way connection with both the sensor and the actuator (Figure 2). Such a loop controller might also get additional messages from outside the loop, say current goals, and so forth.

von Uexküll mentions that the new loop “short circuits” the causal nexus of the sensorimotor loop, essentially maintaining a “*counter world*” or “*mirror world*.” Thus, the new loop may generate a “*seek image*” or template for a desirable sensor stimulation. Such seek images are objects of a *user interface* to the causal nexus. Von Uexküll (1909) speculates that they may acquire a “functional tone” when they become habitual, a quality-meaning that is not unlike a Gibsonian “affordance” (Gibson, 1978).

However, Gibson finds the affordance of a stone (say) in the stone, where for von Uexküll it is an element of the animal’s *inner world* (user interface). Their ideas are ontologically distinct. Gibson seeks meaning in the world (or the ecosystem), whereas von Uexküll lets the animal create meaning in its poking of the world (see later).

Habits, Types and Phantasmatic Self-Stimulation

Kant noticed that awareness is not necessarily caused by physical stimulation of the sensitive body surface, but that organisms may also *stimulate themselves* (Lohmar, 1993). In such cases, there are no direct physical causes; the self-stimulations are phantasmata. Yet these phantasmatic self-stimulations are experienced as not less “real” than the regular ones. If I see someone bite a lemon, I cannot help but feel the sour taste on my tongue. The self-stimulation is not distinguishable from a physical stimulation. Kant gives examples of seeing all kinds of objects in randomly shaded bed curtains.

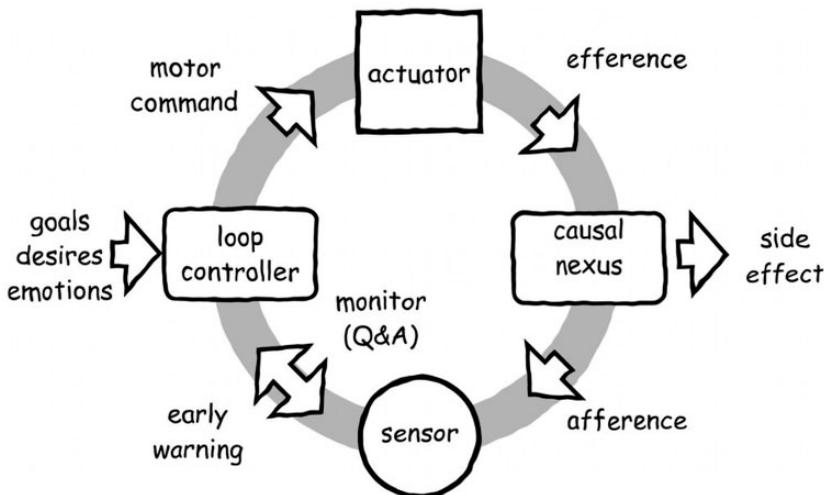


Figure 2. This is von Uexküll’s “new loop.” Notice that the basis is still the simple sensorimotor loop. It runs in the physiology from sensor to actuator (left part) and closes through the causal nexus of the environment from actuator to sensor (right part). The “new loop” adds a loop controller and makes the left part two-way. There are two additional arrows: At the left, the loop controller receives current goals, and so forth; at the right, many causal threads that fork from the causal nexus never come to influence the sensor.

The “lemon” is a familiar type that goes with hand-sized convex things, yellow patches, sour tastes, peculiar smells and many other things. Given one of these, I anticipate the others. This is typical of numerous habits I have; these are usually connected to various types that help me to be prepared for other things. A sound or a faint smell may well prepare me for visual things, say. All this goes on in preawareness. It has nothing to do with concepts, or language, or that kind of thing (Lohmar, 2003).

Habits and types are “sedimented experience” as Husserl (Lohmar, 2014) has it, likely collected through arbitrary association, which may be understood as a kind of analogous (nonlogical) prethought (no “inner talk”). The associations are essentially arbitrary, as was pointed out by Berkeley:

That one idea may suggest another to the mind, it will suffice that they have been observed to go together, without any demonstration of the necessity of their coexistence, or so much as knowing what it is that makes them so to coexist. (Theory of Vision §25)

(Berkeley, 1709), thus “sedimented” appears an apt term. No doubt, this accounts for the bulk of the animal mind, but it is only too easy to miss its fundamental importance for the human mind.

David Hume (1738–1740/1975, 1740/1938) notes three principles of association, namely, *resemblance*, *contiguity in time and place*, and *causation*. Hume considers these to be brute facts, for their causes “are mostly unknown, and must be resolv’d into original qualities of human nature, which I pretend not to explain” (Hume, 1738–1740/1975, pp. 12–13).

He fully understood the “*vast consequences these principles must be in the science of human nature*” and goes so far as to say they “are the only ties of our thoughts, they are really to us the cement of the universe, and all the operations of the mind must, in great measure, depend on them” (Hume, 1740/1938, p. 35).

Thus, Hume is in full agreement with Husserl here.

von Uexküll understood what we now call Gibson’s affordances in terms of habit and type. His “seek images” were a kind of preaware preparation for seeing things, for instance, to spot pubs everywhere when you are dying for a beer.

The importance of this will become clear later; the crux is that in the genesis of awareness the system needs to be able to ignore the bulk of the sensory stimulation. It is not a matter of “attention” but an automatic anticipation, a kind of general preparation for the next moment, putting you—in preawareness—on the right foot, so to speak.

Kant’s phantasmatic self-stimulations are like von Uexküll’s seek images. They tend to disappear as relevant physical stimulation is, or becomes, available (Lohmar, 2003).

Phantasmatic self-stimulation is not just an oddity. You would not be able to perceive *anything* without it, for you would be out of “Optical User Interface” (OUI) elements. Awareness is not an echo of retinal stimulation. It derives from poking on the basis of seek images. The latter may easily give rise to (phantasmatic) self-stimulation, like they may give rise to imagery. When the poking identifies a target, the target soon dominates the process, it is a “real answer,” preferable to the necessarily vague potential answers.

I would venture that this works much like the so-called *nose* of a successful forensic investigator, which leads to *hunches*, or *plots* that apparently come out of the blue, but are decisive in looking for “evidence” while ignoring the bulk of the “data.” Without this, few criminal cases would be solved; one cannot possibly “compute the solution on the basis of the data” as many people believe vision works (so-called inverse optics theories). I’ll return to this important topic later.

Habits and types are formed all the time, often needing only a few shreds of actual experience. They also easily disappear. von Uexküll's experiments with the worm-eating behaviour of the toad yields a simple and impressive example. One sees an "affordance" buildup and subsequently dwindle under experimental control. It shows without a shred of doubt that the "edible" affordance is not in the worm (as Gibson would place it), but in the toad.

The seek images are a kind of *OUT* elements. Yet one should not think of them as rigid templates. A large part of the strength of the types is exactly in their vagueness, or malleability. We can become aware of "something" or of "an animal," without having the least idea whether it is a bird or a cat. There is a vast range of vagueness. We can become aware of the presence of a *Passer domesticus* instead of a mere "small, brownish thing moving in a certain animal pattern," which might be a rat or a dead leaf blown by the wind. I will return to this important topic later in the text, where I have developed some tools that let me handle it.

Habits and types, as sedimented experience through processes of association, allow a kind of mental activity that I would not hesitate to call "thought," that works with feelings instead of symbols, concepts or words.

Notice that—after the linguistic turn—few philosophers would dare to speak of "thought" processes here. It is the kind of mental activity one may expect in the animal mind, but it appears more than likely to me that the bulk of human mental activity is like that too. I would put a far lower emphasis on discursive thought than is conventionally done today.

The kind of mental activity in the creative stage of artistic production seems to be of that kind. Remember Einstein saying "*I never came upon any of my discoveries through the process of rational thinking.*" No doubt, Picasso would have agreed. I do not hold logical reasoning in high esteem myself, although I'm forced to do it at times. I myself also speak and read, I consider language a most effective vehicle of interhuman communication, but I rarely *think* in linguistic terms. I start to do that mainly when people ask me to "explain myself." It doesn't come naturally.

People are different in this respect, at least that's what they tell me. It is perhaps a positive thought that one usually gets by in countries where people speak some (to you) gibberish. For mutual understanding is still possible using a hands-and-feet mode of communication vehicle.

Psychiatry: Controlled Hallucination

Jason Brown (1977, 1991) notably developed a comprehensive theory of psychogenesis based on speculations that were designed for the understanding of the various *agnosias* he empirically studied.

He discards the theories of the "inverse optics" variety as evidently going against the grain of common biological (or evolutionary) wisdom. Evolution articulates novel structures to fine-tune or extend existing structures, whereas the existing structures typically remain to play a nuclear role.

Thus, in neurophysiology, it is a priori rather unlikely that the neocortex would be an independent inverse optics machine, taking over from the existing brain centers. More likely, psychogenesis starts in the evolutionary earlier centers and is fine-tuned by the additions. But this turns the modern notions upside down: Instead of "inverse optics," one deals with "controlled hallucination" (Helmholtz, 1867; Wilkinson, 2014).

Brown (1977, 1991) has shown the viability of these ideas convincingly in a number of books and papers.

In engineering, “controlled hallucination” is known as “analysis by synthesis” (Zhixing & Bir, 2015) and is valued as an extremely robust method. This is very different from psychology, where “controlled hallucination” is generally construed as a childish, unscientific and even ludicrous concept. But in fact, Brown’s theories lend themselves seamlessly to the modern algorithms of “soft computing,” such as “harmony seeking” (Geem, Kim, & Loganathan, 2001).

I would rather call it the adversarial method, that is, the generation of creative imagery to pit against the current optical structure. This is a form of “questioning Nature,” not unlike what one does in experimental physics (my roots).

Formally, psychogenesis in Brown’s sense essentially implements von Uexküll’s *new loop*, especially in its *counter world* and *seek image* aspects. I discuss psychogenesis in more detail later, after the introduction of a few additional elements.

Schrödinger’s Proposition

As explained earlier, one expects no “scientific explanation” for the origin of awareness. Indeed, conventional attempts to do so are nonsensical at even cursory view.

Popular ones are “correlates, or even “centers,” of consciousness” (Crick, 1994) where awareness originates in *deus ex machina* fashion, which is okay (because not science) except that one expects to point out the center or the cause and study the special anatomical/neurophysiological structures, which leads one on an errant ghost hunt—at least, that is what history seems to suggest. That renders it a counterproductive heuristic.

Another one is “emergence,” which is often understood as a kind of causal mechanism. Examples are usually taken from physics. However, I know of no cases where physics uses emergence as a causal mechanism (I may be wrong, but remember I’m a physicist myself, emergence is typically called on by nonphysicists). Rather, one notices instances of awareness after the fact, like science recognized the existence of the oceans as an emergent feature from the atomic theory. I ignore it because it has no explanatory power. There is simply no way to predict the existence of oceans from the atomic theory. It is a relatively harmless heuristic though because it does not suggest any particular direction of exploration. It is an irrelevant heuristic.

A third one—and I’ll stop at that—is to explain one mystery with another, something quantum-theoretical being most popular (Eccles, 1994). This has not led to any useful insights so far, again, history appears to suggest it is hardly good practice. But, who knows? Miracles happen; science acknowledges them as black swans.

Better to understand that there is no “mystery” at all, except for *brute facts*. Brute facts are mysterious because they have no explanation, yet what is more “natural” than the brute fact of awareness? Brute facts stand not in need of explanations, they simply *are*. Nothing one can do about them. It is hardly a heuristic at all, since it essentially terminates the quest.

In my view, any rational understanding must eventually be grounded in brute facts. Otherwise such understanding would be like a closed formal system such as (for instance) projective geometry. Such formal systems allow open sets of mutually distinct interpretations. For instance, in the projective plane, points and lines are distinct objects, but play equivalent roles; thus, points have no such a quality as “punctiformity.” Thus, the closed formal systems are irrelevant here. Of course, the sciences have all *started* from the brute facts of human experience.

Erwin Schrödinger (1944, 1958; Koenderink, 2015a) clearly perceived the issue for what it was, and he understood its far reaching importance: All of *science ultimately builds on brute facts*. He suggests that one needs a heuristic that can never get in conflict with present or future science yet has useful heuristic power in suggesting novel lines of investigation. His idea is that awareness is associated with the learning of the living substance, whereas knowing how is unrelated to awareness.

I refer to it as “Schrödinger’s Proposition.” In my wording, “when an organism’s intentional poking encounters resistance, then the organism ‘meets the world’ in the resistance and experiences a spark of awareness in the process.”

Schrödinger understands this as the ultimate (perhaps unique) *information gaining process*, akin to experimentation in the sciences. This immediately makes contact with two earlier key points, namely Brown’s “controlled hallucination” and von Uexküll’s seek images. Both implement *intentional poking* of the causal nexus. Thus, the sparks of awareness (I call them *psychons*; McCulloch & Pitts, 1943) are associated with the new loops (no relation to Sir John Eccles’s psychons).

Notice that this implies that awareness is “neither in the head, nor in the world.” It is also prepersonal. Awareness is not spatially distributed at all. It lives on another ontological level. I like to think in terms of the TRINITY—World–Body–Mind as a single entity; I’ll touch on that a few times later.

As an aside, awareness should not be confused with “consciousness,” which involves the notion of *self*. The C-word was banned from Newspeak till roughly Orwell’s critical date 1984. About that time (give or take a decade), it was redefined. It isn’t a very useful concept (at least to me), and it is not understood anyway (as is evident from the literature), so I avoid it in this article. The notion of C... is simply too messy and of little relevance to my topic. Discussions usually focus on self-awareness, which has nothing to do with perceptual awareness which is prepersonal.

All sentient beings have awareness, but few are conscious; one needs to be a “social independent” for that (notice that bees are social, but not social independents). Only social independents are in a position to develop selves in a meaningful sense. Again, I’ll stop at awareness here, that being my core topic.

Truth Finding Processes

Notice that you may very well “find” truth, but that it makes no sense to “seek” it, since it isn’t “out there.” *Truths are constructed, not found.*

You keep them when they serve your purpose. In biology, that means that so-called truths are likely to increase your biological fitness and perhaps even serve to postpone your (certain) demise a little longer. No mean feats. Truths should be cherished! But “finding the truth” seems not such a simple matter; it is the topic of this section.

I propose that “controlled hallucination” is a process akin to forensic investigation, reconnaissance or intelligence. It is a “truth-finding” process in an environment whose generic properties are by and large understood.

Such problems cannot be solved by “computation from the data”; in order to know what data are needed, one already needs the solution. For what is usually called “data” is actually just limitless chaos—I call it “structure” here—and of that there is always too much. That is why no computation can even get started; one needs to get rid of the irrelevant structure (the bulk of it!) first. But how to achieve that? Well, as I said, one needs the solution for that.

There is some light in the dark though. One may certainly *check any mere hunch for plausibility*, for that only involves Bayesian reckoning. That is great, for it means that you may be able to recognize a “solution” for what it is.

Perhaps more commonly applicable, you can recognize useless (notice I don’t say “wrong”) ideas for what they are. This is far more important than it is usually made out to be, for how many useless (perhaps even wrong) ideas do you harbor in your mind right now? Lots, at least in my case. If you are not full of useless ideas, then maybe you have no ideas at all. Fortunately, weeding them out you can. But in order to recognize a good solution you first need to have some contenders.

The “truth” I am talking about is not “as seen by God’s Eye,” nor “logical truth,” but it is what promotes efficacious action and biological fitness.

So how to arrive at such a good “hunch” in the first place? Hunches are counterfactual; they can only be “posited,” that is *hallucinated*. In forensic investigation, “controlled hallucination” via “plots” is the only known algorithm. Fortunately, it is very effective in case the environment is sufficiently understood.

The perfect illustration of the efficaciousness of controlled hallucination is the “Game of Twenty Questions” (Walsworth, 1882). Someone takes a word (any word!) in mind, another player has to guess it in less than 20 (two alternative) questions. The game is attractive because good players succeed mostly in about 20 questions, leading to well-balanced loss and gain as players alternate roles. Not surprisingly a popular first question is “mineral or organic?”

Whitman Richards (1982) suggested that perception is much like *Playing twenty questions with nature and win*, an insight that seems so important to me that I do not understand why that paper failed to become an instant classic. Science is not kind and, of course, it shouldn’t be, but some happenings seem counterproductive.

I denote the process “Sherlock Holmes method” after the fictitious detective (Doyle, 1887). Sherlock Holmes is a key example of Alexius Meinong’s (1899) notion that there are objects for which it is true that there are no such objects, in the sense that you can’t do without him (“Sherlock Holmes” I mean). Holmes’ famous “method” has rightly become legendary.

The criminal investigator needs to solve a *whodunnit* starting from scratch. It may not even be clear whether a crime has been committed at all, possibly just a dead body was found, maybe there is only a rumour, maybe an accident, maybe a suicide.

The detective hallucinates a “plot,” using his intuition (e.g., *the butler did it*). There need not be any particular reason for this, call it a hunch. Given the plot he generates questions (*are there prints of flat feet in the neighbourhood of the crime?*) that can be checked in the data file collected by the local (traditionally dull-witted) policeman.

If there is no record, Holmes asks for additional footwork. Possibly the local policeman photographed all discarded cigarette butts in a hundred foot radius, but Holmes simply ignores these. Cases are hardly ever solved by footwork, or forensic science laboratories. Those merely yield more of the (for the most part irrelevant) *structure* (as opposed to “data”).

In this setting, data is called “evidence.” The evidence is only to be found in a very minor part of the structure and is *defined* as such (not *found!*) by the investigator.

Any random structure might be evidence (“by accident” so to speak), but most likely it is not. Evidence has to be created in the context of some plot. The creation implies both finding it and recognizing it for what it is. This implies looking for it—in some very vague sense. For if you don’t look for something, you will not recognize it, and if you can’t recognize it, you will not find it.

“Vague sense” is important here. If you *really* know what to look for, the task is trivial; it can be safely trusted to the lowest rank police officer. But—perhaps surprisingly—the probability of success is less. For if you only vaguely know what to look for, you need to wield *creative sight*, that is what ace investigators are paid for. Your chance of success is larger, because you may mould your target to what you see, as long as it fits your vague idea of what to look for.

Suppose the footwork accidentally turns up impressions of high heels. Noticing this, Holmes switches to another plot. At least for the moment, but eventually he is ready to drop the initial butler idea. After all, hunches are only hunches. So he tries *perhaps the countess did it*. Maybe Holmes noticed the countess smoked a rare brand of cigarette, now suddenly the previously ignored data file becomes important! Some structure is promoted to possible evidence; some previous evidence is demoted to structure and further ignored.

This is purely controlled hallucination, wielding the adversarial method. Plots are hallucinations. They allow Holmes to frame questions. The questions turn mere structure into data for him, all other structure (probably most) being fully ignored. If the plot suggests lacking data, this induces Holmes to ask for additional footwork.

The “Holmes method” of investigation is driven by plots. The plots are creative inventions, merely based on a vague situational awareness and an investigator’s “nose.”

This is entirely analogous to the habits and types as “sedimented experience” noticed by Husserl. The investigator’s “nose” works because of its store of types.

Notice that *structure as such is meaningless*. It only becomes “data” (meaning, or evidence) in the context of some *question* or *plot*.

Although strictly meaningless, a neat ordering of the structure is still useful (all cigarette butt photographs in a single folder, perhaps ordered by brand, and so forth), just in case one needs to find something at the drop of a hat. But no doubt, the large bulk of structures collected at the site will eventually turn out to be meaningless, despite such neat ordering. They will never be addressed. They remain in storage at the police department, gathering dust (who knows, some old case might be “opened” again), but never make it to the court room to be used as “evidence.”

Notice also that *the meaning of an answer is in the question* rather than in the answer. This is Schrödinger’s Proposition. It is the intentional poking (questioning nature) that yields frictions (unexpected answers) that again yield insights. Any question implies possible answers. Answers that don’t fit the “format” implied by the question are willy-nilly ignored.

Holmes is “done” when sufficiently many unlikely and mutually independent facts have been collected so as to put the preferred plot “beyond reasonable doubt.” This is a trivial matter of Bayesian probability calculus (Jaynes, 2003). Sherlock Holmes does that on the guts, not allowing necessary action to be stymied by attempts at fake precision.

Does it mean the “crime has been solved” in terms of “truth found”? Of course, not! But it is the best that can be done, and time is limited when action needs to be taken. Better the occasional mistake than endless indecision, which bars the ability to *act*. In the stories, Holmes is always right, but biology (life!) is different.

Perhaps sadly, mistakes happen (and animals die) but that is not a problem if not too frequent. It is how the species learns (the less successful animals stand smaller chance to procreate). Realistic investigations are *time-limited*.

In perception, *concrete actuality* as given by visual awareness is essential even if not proven “veridical.” At no moment may the ability to *act now* be put in jeopardy. This is of paramount biological importance, often a matter of life and death.

Organisms often have to act on scant evidence; in the worst case, “don’t just stand there, do something” is the best tactic. Just stand there and you’ll soon be absorbed into the food chain.

The only good thing here—at least for your tribe—is that you’re removed from the gene pool.

The Art of Questioning

In forensic investigation, the quest originally derives from society’s cry out for justice. In reconnaissance, it derives from a desire for certain information, and so forth. In biology, an agent’s questions may derive from needs or emotive drives. It doesn’t matter much. In all cases, the drive materializes into plots, focussed searches, hypotheses, you name it. Same things.

When answers are forthcoming, questions will be sharpened, perhaps beyond recognition. The process is “self-focussing.” In the 20 questions game, the initial few questions are necessarily just random poking, whereas the later ones quickly focus so as to “zoom in.” It is an information seeking and gaining process.

How can there be such a “free lunch”? Remember you start from scratch! The reason is that the investigator is thoroughly familiar with the generic environment. The *Umwelt* (Koenderink, 2013a) is thoroughly *modal* on levels of aggregation (Richards, 2018). Holmes would be powerless on Mars or in the Amazon jungle.

In biology, this is the animal’s *Umwelt*. It is a perspective on its “sense world,” all physical processes that may trigger its sensors, and its “action world,” all the ways it may affect its physical environment. The perspective typifies the animal’s lifestyle.

Notice that the “lifeworld” is not the “Umwelt” (Koenderink, 2013a). This is important in understanding that awareness is driven from the inside instead of from the outside (in conventional terms). The tiger and the lamb have roughly comparable *Umwelts*, but very different lifeworlds.

The lifeworld/*Umwelt* is also the “counter world,” which may be understood as its “user interface” to the environment. Any awareness is willy-nilly in terms of the elements of *OUI*. These act as a kind of “templates,” think of the “seek images” or better the “types” I mentioned earlier.

Ethology has come up with spectacular examples. Although the general notion is that humans are essentially different from the mere dumb cattle, it is readily shown that such is not the case. The researches in perceptual Gestalts clearly show that human visual awareness critically depends on template-like structures, or habit and type.

In the generic case, perceptual awareness is likely to come up with hallucinations that are much like the previous moment, this gives the system a head start. Only in relatively rare cases like protopathic “early warnings,” something more focussed is called for.

Think of a sudden beesting. The early warning system makes you jump, but since it is protopathic, modal awareness may kick in only later. The happening focusses the hallucinations though, momentarily switching to a novel mode.

Early questions tend to be general and answers are sought in fairly coarse-grained summary data structures. Later in the process, questions are refined, and answers looked for in more fine-grained data structures.

As a question turns out to be useless, then that line of questioning is terminated. Ask a silly question, get a silly answer. No shame in that if you use it to redirect your questioning, and make a fresh start. As a question turns out to hit interesting structure, then that line of questioning is pursued. In this case, the question is turned into a focussed probing, may fork into a number of refined probings in various directions (Figure 3).

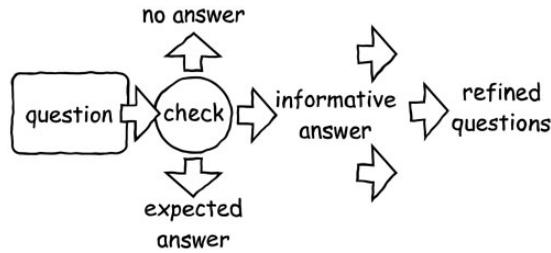


Figure 3. The art of questioning. Notice that the arrows at far right might enter the process again (the arrow at far left). Thus, one generates a possibly endless volley of queries. Only if all answers turn out useless (no answer) or trivial (expected answer) does the process end. In practice, the time limit will nearly always be the effective terminator.

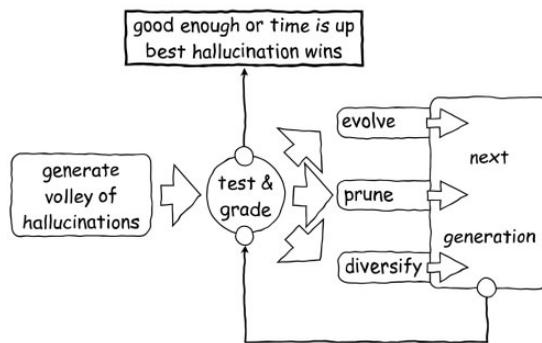


Figure 4. This is the soft computing algorithm known as “harmony seeking.” It has been shown to be able to arrive at useful solutions in cases that would be very hard to tackle in other ways. It is a fitting model for Brown’s model of psychogenesis or Holmes’ method of criminal investigation alike.

Thus, questioning is a process of merciless pruning and successive articulation and forking, constrained (“sculpted” is an intuitive description) by the available structure, where possible turning structure into meaning (Figure 4).

The questioning is a process not unlike the evolution of species in that many plots are initially followed up but few survive. Those that survive may change beyond recognition, an evolution that is focussed (or constrained) by the data identified in the course of questioning so far.

Is meaning “created” or “found”? There is no clean answer to that; it is both created *and* found. What is important is that the meaning that is created or found be useful in guiding successful action. This need drives the creation of initial hallucinations. The subsequent psychogenesis (an apt term for the process of meaning creation) serves to focus and adapt behaviour to the constraints posed by the environment. This is the biological counterpart to the philosophical concept of “veridicality.”

Awareness cannot be “veridical” because that requires an external “God’s Eye” view. Successful biological agents—and all are or they wouldn’t be here—are well adapted to their environments given their lifestyles. This by no means implies that they “perceive the world truthfully”; such a notion cannot be given a biological sense.

The peripheral sensoria are volatile—because continually overwritten—storage areas for sorted and formatted structure. There is neither computation (thus no “inverse optics”

either), nor any mysterious generation of awareness. I would characterize the sensorium as an embodiment of geometry (of various types), which renders it useful to store spatiotemporal structures, not as an area involved in any “computing” (I’ll discuss that later).

One thinks of the questioning process as implemented by a dense nexus of new loops, some perhaps launched in the course of psychogenesis because of a need for a special expertise. This is indeed likely, because when some “inverse optics” is called for, it will almost invariably be needed in limited regions of interest (I’ll return to that later). It is very rare for the environment to be geometrically and materially homogeneous and this is reflected in the part-wise homogeneous structure of generic retinal irradiations.

The Nexus of New Loops

The description of higher animals necessitates one to posit a dense nexus of numerous new loops. This can be seen as a way to enable tackling a complicated whole by breaking it down into weakly interacting parts.

It is the only known way to handle complexity apart from coarse-graining it—which may be of use too. One doesn’t need to make a *choice*, of course.

I start by remarking on some generalities:

- Loops may *share* sensors or actuators;
- A “sensor” may actually involve numerous subsensors (think of the optic nerve), an actuator large muscle groups (as in posture control);
- Loops may live and die. For instance, “expert” loops might be launched when needed. One assumes that they might be “cloned” from some template loop;
- Mutually independent loops may still affect each other by way of side effects they induce in their common environment; and
- The “environment” of a loop needs not be the causal nexus of the environment. It might be the loop controller of another loop.

Especially the latter observation is of crucial importance. It allows the system to *build meaning upon meaning* (Riedl, 1984).

This is extremely important in the case of visual awareness. The peripheral loops will be very local and their counter worlds very limited. Loops that build on a number of these have a more global view; their counter worlds create a context in which the messages from the subloops acquire a “deeper” meaning. And so forth! Many aspects of visual awareness derive from hierarchies of loops.

This is how “meaning” works. The explanation of any complex presupposes that the meaning of the parts is known. The same applies to the meaning of the parts of the parts and so forth. If there is to be an end to this, then some parts have to be brute facts; since these are their own explanation, they simply are what they are.

Templates are pretty much like brute facts in this respect. This is also how dictionaries work. They explain words in terms of words. This works if there are sufficient words not in need of an explanation.

Thus, any meaning is necessarily esoteric. (“Esoteric” denotes the lack of conventional explanation, that is, “known” only to oneself, “core knowledge.”)

In my view, the mind is perhaps best understood—in its functional sense—as a topologically intricate nexus of new loops.

Another important observation that should not be taken lightly is the interaction of loops sharing a common environment by way of some side effects of their actions.

Von Uexküll's famous example is the sea urchin, where each spine has its own independent controller. A measure of synchronization is due to interaction via the common environment, not a central nervous system.

When a sea urchin walks, *the spines move the animal*. When a dog walks, *the animal moves its legs*. Notice the categorical difference. This is how group behaviour is implemented in many social animals, think of schools of fish.

In humans, this is how empathic interactions come to be. Other humans are part of the environment of any single human. Thus, learning the structure of the environment implies learning the structure of other humans. Since all humans are similar, it implies learning the structure of oneself. It is one aspect in the origination of a *self*, which is again one aspect of the origination of a C. . . (as opposed to mere awareness). I will not pursue that here—this is naturally a very extensive topic; this article is restricted to awareness proper.

It is important to remember that there are interpersonal loops, just as intrapersonal loops as well as loops that involve only a tiny subpart of an animal's anatomy/physiology. Scale—in the sense of complexity rather than space-time—is the important parameter.

Distributed Awareness

A single beat of psychogenesis must involve volleys of psychons at various levels in the system. How can there ever be a “single” awareness then?

It depends upon the perspective. From the inside, awareness is necessarily singular, for it exhausts concrete actuality. For an outside observer, it may be useful to reckon with multiple awarenesses, albeit that these are necessarily abstract, whereas the awareness from the inside is concrete.

Some intuitive insight is most easily gleaned from examples of familiar cases of multi-level awareness.

A simple, but useful example is that of the classical regular army. It is structured hierarchically and relies on one-way messaging rather than mutual understanding. The same structure is implemented in smaller structures, like mafia “families” (Figure 5), which are essentially structured after armies.

Messaging differs from communicating in a shared language, because it is one-way. The “meaning” of a message is different for either side. The general, the officers and the soldiers may have very different takes on what they are currently involved in. Yet the army as a whole acts as a single organism and flexibly deals with accidents. The general may never leave his command bunker, whereas the soldiers are in direct contact with the environment. No doubt, their awareness is very different. But the general is aware of higher order data that the men are oblivious of. Yet the general's view is ultimately based on reports by the men, filtered up—and changed in the process—through the chain of command.

When you look at the army from the outside, as a historian, or reporter say, does it make sense to say that the army is an organism? If so, can you guess at its awareness? Awareness and intentionality are often ascribed to the general, like in “Napoleon won the battle of Austerlitz.” Alexander the Great famously said to fear an army of lambs led by a lion more than an army of lions led by a lamb. But that seems not right. The awareness of the army—if there is such a thing—cannot be that of the general, nor that of any soldier or officer. Somehow all of these are involved. The general alone could not win a battle, nor could the men when left to themselves.

Consider that the system of a single human is rather more complicated than that of an army. One meets with the same problem as viewed from the outside perspective. The inside

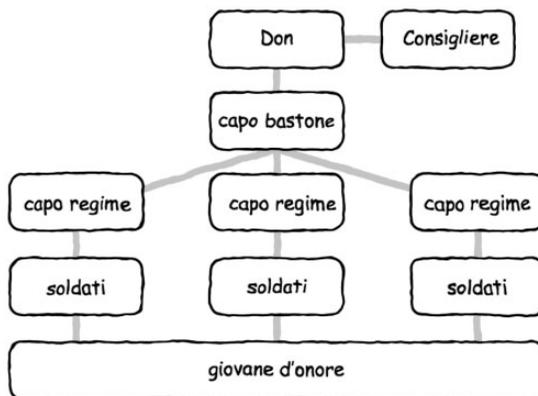


Figure 5. The typical hierarchy in a mafia family. The boss (*Don*, or the *capo famiglia*) is on top, the advisor (*consigliere* also financial advisor [*contabile*], etc.) is not in the active command chain. The lowest level is not truly part of the “family”; these are expendable hirelings, who act as mere muscle and take the brunt of the street action. The underboss (*capo bastone*) screens off the boss, who is really a fully “inner loop.” Most of the work is done by the soldiers (*soldati*), who are grouped under the commanders (*caporegime* or *capo-decina*). The system is a very strong one (much like a smallish army) and can remain effective even when major parts of the structure are eliminated. It seems the perfect model for biological agents. Ethology and physiology suggest that it indeed is. It is instructive to think about the loop structure, which—of course—includes the environment. Many conventional notions from vision science are soon seen to be irrelevant, or counterproductive. Thinking in terms of the mafia structure might be a good way to weed out theoretical notions in brain (or rather mind) theory.

perspective is a singular awareness of concrete actuality, including the body as part of the environment.

Being aware of “blueness” (a blue patch in front of the eye say) no doubt results from the action of myriads of loops acting on loops. The psychons at the various levels have very different significance. The quality of blueness is by no means simply related to the activity of retinal photoreceptors, for all links in the chain somehow contribute to the blueness. It is like the general’s awareness of an imminent attack on the basis of such reports (of course suitably filtered and reformulated by various officers) like “Longbottom’s platoon did not return,” “the men noticed smoke plumes over the western hills,” the general’s present toothache and so forth. The general may order a fake attack or threatening move, just to “probe” the opposition, but the men involved in that need not to know (“... , Theirs not to reason why, Theirs but to do and die. . .”—echoes of Tennyson).

Your awareness does not include photoreceptor events and so forth, they are not part of concrete actuality. They belong to the reality of science.

To conclude this part, the army idea is a useful paradigm in other respects too. It will repay the effort if you ponder it awhile. Consider what the army would be without the enemy, or without the battlefield. It would be meaningless and the men could be sent home. The enemy is a vital part of the army. Likewise, the environment is a vital part of the body–mind.

I think of perception as an aspect of the TRINITY—World–Body–Mind, which is life itself. These aspects codetermine each other, and none of them can be understood in isolation or is a cause of one of the others.

But this only works if you can live with a world without the God’s Eye View and if you forget about the silly notion that perception is the result of a brain-computation of that View.

How Vision Proceeds

Here, I consider the process of looking. I illustrate it with a few topics that have always happened to interest me, but I might have used various alternatives just as exciting.

Units of Looking

Psychogenesis in humans is a legato systolic process that repeats at about a dozen beats per second.

A single beat, or systole, is not worth more than a *glimpse*. In isolation, it may escape awareness, because it necessarily has a protopathic character. Most glimpses are not isolated but member of a sequence of beats with similar outcomes. In that case, they are not noted as glimpses at all, but one is deluded to believe perception is a continuous process.

The smallest intentional unit is a *glance*, which lasts for maybe half a dozen beats and includes one or two fixations. Notice that most fixations are involuntary and part of the psychogenesis. Involuntary fixations show effects of anticipation; they are part of a focussed probing. Glances are part of the questioning and poking that goes on under the hood. They are wielded by the psychogenesis of awareness, of which you only experience the brute fact.

A *good look* would last one or two dozen beats and includes at least one voluntary fixation. A good look is an atomic unit of perception, although it has parts. It is analogous to a “single happening,” like a handshake, which evidently has parts too, yet is a natural unit. The duration of a good look is a “specious present.” You “throw a look,” it is something you do.

Scrutiny is a process that involves awareness, but is embedded in a predominantly cognitive process. From the perspective of perception scrutiny is not a natural unit of awareness, but a heterogeneous, quilted structure, closely intertwined with reflective thought processes.

I do not deal with scrutiny here, because it oversteps the horizon of pure awareness. For the same reason, Adolf Hildebrand (1893) did not consider it worth the attention of the artist.

So isolated glimpses are rare singular events, glances frequently happen to one, whereas good looks are things one does (Koenderink, 2018). They are all atomic entities.

The Specious Present

At the close of a beat of psychogenesis, an investigation from a hardly differentiated dream-like state to an account that fits the constraints of the peripheral sensory areas has been completed, for better or worse. The next systole is already on its way, and the previous one has partly been reenacted. That is why awareness is momentary, yet rooted in the past.

Psychogenesis also “completes” processes it hallucinates, thus has a prognostic character. Awareness has “tentacles” that reach out into the future as it were. There are a present of the past and a present of the future, all rolled into the present moment. These are Husserl’s retention and protention (Husserl, 1893–1917/1964).

This is even more evident at the level of glances or good looks. The spatiotemporal order of experienced events within a specious present may be different from “physical reality” if that “makes more sense” (Koenderink, 2012b).

Psychogenesis does not “find,” but actually *creates* space-time. This involves continually reediting both past and future and their interrelations so as to make current sense. It creates “stories” or “plots” that link immediate past and future in meaningful ways. That is why the specious moment is like a “glue,” one might say it is time itself.

Thus, “lived time” is far from being a movie of which you see parts at any time, but which really (in the “Eye of Eternity” say) exists as a whole. The movie (past and future!) is being *reedited* at any moment, largely in preawareness. It is a moving target, like life itself.

Psychogenesis checks ever refining imagery against the data structures in the sensoria without ever exhausting them. Indeed, psychogenesis is like the forensic investigator in that it manages to ignore the bulk of the available structure in the sensorium. It only uses what serves to “explain” the current plot, or flatly contradicts it.

This is necessary because the available optical structure is almost invariably too rich to be exhausted. This leads to the familiar phenomena of “inattentive blindness” (Mack & Rock, 1998). Only artificial laboratory setups are simple enough to take in at one gulp.

“Attention” (Pashler, 1998) is somewhat of a misnomer here. What “failed to be attended to” simply didn’t exist for the process. In order to exist it would have to be created in psychogenesis first. Psychogenesis is necessarily *focused*, but it is equally necessarily *complete*. It is not that psychogenesis “omits” anything; it is that it *creates* limited complexity.

Each wave of psychogenesis can do only so much. It is an essential bottleneck of vision—as seen from the outside. But animals—including us—cannot boast a God’s Eye view. One cannot be aware of not being aware of something. (Of course, one can—sometimes—in *retrospect*—after all, the story changes while being written!) There is never anything lacking in concrete actuality, that is, logically impossible.

One way to understand this is to consider that animals without eyes are by no means blind. Optical interfaces are simply no part of their *Bauplan* (Riedl, 1975; Uexküll, 1920). I intentionally use “*Bauplan*,” a term that was dropped with the Newspeak of biology that was installed in the late 40s. As with so many of these cases that was rather myopic, it has been suitably redefined by now. (Riedl, 1975 gives a modern account.)

Concrete Actuality

Imagery is a glue that binds the optically mediated “evidence” into a coherent narrative. Notice that “evidence” is *created* rather than *found*. Whether “correctly so” is measured by evolutionary fitness.

Given our median life expectancy, most humans don’t do so badly. But neither do ticks (*Acainae*), which can neither see nor hear. Ticks are not unsuccessful as animals go; they have been around much longer than we have, once they fed on dinosaurs, when most mammals were still nocturnal insectivores, afraid to come out into the light of day. So our basis for anthropocentric pride is thin. It would indeed be wrong to call a tick “blind,” but it has a minimal *OU* when compared with us. (But notice that in many ways “small is beautiful.”) Complexity is not the issue.

This implies that the brute facts will likely be different, even for observers in front of the same scene. Psychogenesis creates actuality. Because this is something modern people find it hard to swallow, I propose to say that psychogenesis creates a spooky object “concrete actuality,” saving “reality” (as understood by analytic philosophers) for such trusty (to them) objects as electrons.

This should be palatable to scientists who consider electrons as far more “real” than something mysterious like even their own awareness. Contemporary intellectuals buy this, whereas people who need to act (politicians, generals, mafiosi and so forth) obviously prefer concrete actuality over such academic reality. Philosophy doesn’t win battles; generals do, relying on concrete actuality.

These latter people also tend to understand that there are manifold actualities, whereas they couldn’t care less about the single reality. That reality is only seen by God’s Eye which

even physicists lack. Practical people understand that the belief in God's Eye doesn't buy you anything. Consider this just good biology, as is evident from the fact that they tend to be survivors.

In the account of science, only *reality* really "exists." In your account, the tree (say) in your awareness is *concrete actuality*. Using such different words, perhaps the two views could coexist—as they should, for science and life play on different levels. Concrete actuality is different for you, your dog, or a dove. Reality is the same for each—though humans will have to testify for the dog and dove here and perhaps God needs to testify for the humans. Such insights lead to treatments like Meinong's (1899) *Theory of Objects*.

But you can't have your cake and eat it. If reality is to pertain to all sentient beings, then it cannot contain such objects as "trees" in the sense of *human actuality*.

Notice that my account is focussed on *psychogenesis*. This does not imply that there are not countless other loops, such as involved in maintaining my posture for instance. Indeed, I rarely fall over spontaneously though I hardly remain in voluntary control all the time. My "robotic infrastructure" almost invariably saves the day, without my even noticing.

Awareness may not even draw on the bulk of the activity in "visual areas" of my brain. The "vision" involved in basic perception-action cycles does not necessarily result in any awareness at all, even though the processes might be complicated. A well-known example is the drive back home after work, which so frequently appears to have been "blind," although driving behaviour was likely to be normal and safe.

The brute facts of awareness are important in yielding the material for potential actions that cannot be handled by mere kinematical extrapolation. Kinematical extrapolation uses current *trends* to predict the immediate future (or—if necessary—the immediate past). All biological systems use kinematical extrapolation big time as it involves no dynamical theory (theoretical physics!) at all.

Next day weather prediction is mainly of this type, the dynamical theories (hydrodynamics, thermodynamics and so forth) not doing one much good. Not that there is anything wrong with these dynamical theories, they are good physics. It is just that they—together with the current observations—fail to yield good predictions, even on a fairly short timescale. Kinematical extrapolation is much simpler and can hardly fail on the short timescales. It is just that such extrapolation assumes a certain smoothness, or inertia. It is unable to deal with sudden events. You cannot predict the next thunderbolt by extrapolation. But in predicting tomorrow's weather as the same as today, you'll be right in the majority of cases—no need for arcane theories.

Thought—not necessarily of a linguistic nature—entertains counterfactual *possible futures*, departing from the stable basis of these brute facts. Such structures may again feed into the structures that launch psychogenesis. Organisms routinely "feel their way into the future" in this manner, automatically selecting desirable alternatives.

Although rarely acknowledged, I think this equally applies to the past. The past is constructed and continuously rewritten no less than the future is; "free will" applies to both, a notion that may surprise some.

Thus, memories are just as much confabulations as anticipations of the future are (Loftus & Davis, 2006). For the long term, this is familiar from the fact that childhood memories become clearer with age. But even on the shortest term, psychogenesis needs to rewrite both past and future continuously in order for its presentation of actuality to make sense. The meaning of the future lies in the past, the meaning of the past lies in the future and accounts of actuality require fitting accounts of what was and what is to be. Your concrete actuality is not an automatically augmented, retrieved record, based on current input, it is a lived presence. Thus, you cannot count on your friends (let alone strangers, or even animals)

to share your concrete actuality. Unfortunately, not all of us reckon with that, reckoning (true, but not entirely relevant) that “the World is the same for all of us.” Well, *the* World (God’s World) *is*, but *our* worlds are *not*. Trivial as the fact may be, it is the root of abundant unfortunate frictions between sentient beings.

User Interface Elements and Their Effect on Behaviour

In this section, I discuss a few examples of the influence of the nature of the OUI on behaviour. Although such cases are not at all rare, they have hardly been systematically collected. Except for a few of the better known instances, there is not much quantitative exploration either. Yet the nature of the OUI often has major consequences for observable behaviour.

These are just a few examples, no doubt the tip of the iceberg. Since many spectacular effects remained largely unnoticed since the start of experimental psychology, there can be little doubt that systematic search will reveal numerous instances of templates in human vision once the search is on.

Fixed Action Patterns of Animal Ethology

Animal ethology (Eibl-Eibesfeldt, 1970; Lorenz, 1973; Tinbergen, 1951, 1952; Uexküll, 1909) has reported spectacular examples of “templates”—dubbed “fixed action patterns”—in vertebrates (our kin) like birds and fishes. Examples include swans feeding fishes, warblers treating cuckoo chicks several times their own sizes as their young, stickleback females preferring gross artifacts over real macho-males and so forth.

Such behaviour triggered by specific interface elements is also present in insects. The case of the Australian jewel beetle (*Julodimorpha bakewelli*) is well documented (Gwynne, 2003; Gwynne & Rentz, 2003). The animal became almost extinct because males prefer humping stray beer bottles found along the roads over (their “proper behaviour,” perhaps, but remember that there is no official jewel beetle manual) humping females.

There can be no doubt whatsoever that “vision” is an idiosyncratic OUI. Only for the stupid cattle, not us? You may very well think so, but I don’t (Spelke & Kinzler, 2007).

The Classical Geometric Illusions

Once one looks for such things, examples abound. Perhaps the most obvious examples are the classical “geometrical illusions,” say the “café wall illusion” (Gregory, 1997). Since these will be familiar to all readers, I skip the discussion. These are strong and numerous instances though (Shapiro & Todorovic, 2017). They tend to be striking because—once noticed—reflective thought *knows* that the OUI is often “off.” But this does not affect visual awareness in any way.

Other examples have to do with “magical thought” leading to rain dancing or things like that. These are reminiscent of the fixed action patterns of ethology. Our modern minds are by no means immune to that. This should make one think. (Knock on wood.)

The Apparent Field of View

The “field of view” is an object of geometrical optics, it is to be distinguished from the “visual field,” which is a mental object. The diameter of the field of view is measured in degrees. The visual field does not have a geometrical diameter. However, there is often a certain “feeling” of its size that may be compared with such feelings as they occur under

various constraints, think of looking through a keyhole. Such a comparison yields a number that I refer to as the “apparent field of view.”

Here is an example. The human eye lets in radiation from the full half-space in front of it (actually a little more), the diameter of its field of view is about 180° . Yet Helmholtz (1867) remarks that it “doesn’t look that way” to him! He reports that, to his amazement, his *apparent field of view* is more like a right-angled cone (90° diameter). The same experience is reported by Kepler (1604).

Such spontaneous observations are extremely rare. Indeed, both Helmholtz and Kepler must have been remarkable observers. Just consider, would you have noticed this spontaneously? Be frank!

Both Helmholtz and Kepler fully understood in reflective thought that their visual awareness did not reflect the facts of geometrical optics. They simply accepted their awareness as a brute fact and even dared to report on it. (The I-word!)

In an investigation involving more than 70 naive persons, I find that it is typical to experience an apparent field of view that is widely different from the optical facts (Koenderink, van Doorn, & Todd, 2009). It is just that people don’t care to report it, but happily live with it. Not all of us are Keplers or Helmholtzes, probably a good thing. I found persons with an apparent visual field of only a few degrees and some who estimate to see far behind their ears (Figure 6). My own field subtends about a right angle.

In view of the limitations of the method used in this study, it is likely that the narrow fields were overestimated, whereas the wide fields were perhaps underestimated. It is very hard to measure such things (most of my colleagues would reckon it impossible). My intuition tells me that the limiting values of the extent of the human visual field are likely to be 0° to 360° .

Many people feel that “everything is in front of them” (narrow apparent field, 0° people), yet that their apparent field “has no bounds,” that there is no gap behind their backs. Others feel great in the conviction that they see everything there is to see (wide apparent field, 360° people). But these people do not feel that there is anything behind their backs either.

Nobody feels there is anything behind their backs (Phillips & Voshell, 2009), no matter the size of their visual field. The 0° and the 360° people behave very similarly. You’d never know the difference. In order to find out you need to experiment with them. Yet it is hard to doubt that the minds of these people are hugely dissimilar.

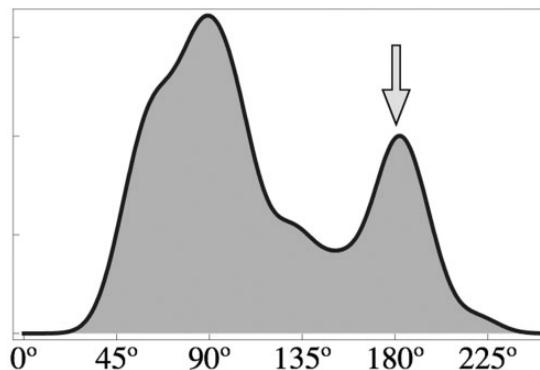


Figure 6. Distribution of the apparent scope of the field of view for more than 70 persons (from Koenderink et al., 2009). The optics would predict 180° (the arrow), but the mode is at about 90° , as reported by Helmholtz. The spread is huge, from about 10° to 240° (but see text).

Both feelings may even exist in the same person and give rise to particular behaviours in different settings.

Parallelism of Objects in the Environment

I have shown that this misperception of the scope of the field of view has immediate consequences for the estimation of parallelism in the environment. I routinely recorded errors in judgement of over a hundred degrees (Figure 7). Such properties of the OUI thus have major consequences for potential behaviour (Koenderink, van Doorn, de Ridder, & Oomes, 2010).

Why are such truly *gigantic* errors not even so much as mentioned in our textbooks, whereas various relatively minor “effects” (a hundred to a thousand times smaller) are discussed in detail? Probably because experimental phenomenology has been a no-go area since (at least) ww-II.

This is very unfortunate, not only to science. When people systematically misperceive important aspects of their physical environment (in this case their own optical sensors, the Euclidean “pyramid of rays”), then that is bound to come with material consequences. When people are often very different in these respects, then it has to have consequences for interhuman communication too.

The Space Behind the Back

Suppose someone points at a person behind you, perhaps meant as a warning. Then, in your awareness, the former person points at *you*, because in your awareness there is *nothing* behind your back (Phillips & Voshell, 2009).

You act exactly like a dog who looks at your finger as you point to something out there. Believe me, you’re as stupid as your dog in that respect. We all are.

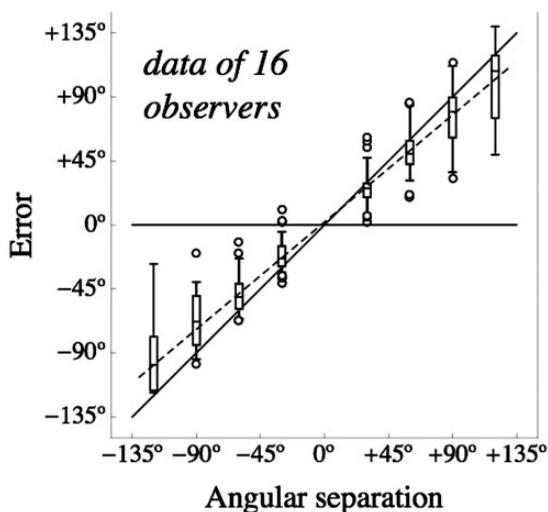


Figure 7. The error in putting two cubes (one presented in random spatial attitude, the other manipulable by remote control) mutually parallel in the environment, as a function of their angular separation in the field of view. Notice that the error roughly equals the angular separation, which suggests that all visual directions are treated as mutually parallel in this setting.

I know about such things, but I'm still as stupid as a dog myself. When someone points at me, I rarely turn my head to look behind me. Well, sometimes I do, but that would not be intuitive, immediate. I might do it because of my empathic reading of the pointer, as when I would feel the person was looking over my shoulder.

This leads to a peculiar inability to deal with full horizon panoramic images, a photographic technique that is rapidly becoming available to the general public. I have shown that this particular mental blindness to the extent of the full optic array applies to virtually all observers (Koenderink, van Doorn, & Wagemans, 2018c).

In such a picture, someone may seem to point at you, the viewer of the picture, whereas the person really points at another person also in the picture. Nothing special if the picture depicts the full horizon (360°) and the two persons in the picture are 180° apart. Then, the camera must be between them, so the pointing is also towards the camera. The person in the picture that is actually being pointed to would be behind your back if you were at the location of the camera. People get really confused in cases like this, even those who have seen many of such pictures. It would seem that humans are not able to get used to them (Koenderink & van Doorn, 2017).

It renders naive users unable to use this modern technology to good advantage. The latter takes training and cognitive "correction" of current awareness, not unlike what one experiences with the classical geometric illusions: One *knows* what one *sees* is not right. There is no doubt that one could learn to use such representations in a "right way," just as there is little doubt that one would still "see them the wrong way" (Koenderink et al., 2018c; Koenderink & van Doorn, 2017).

Perspective

Perspective is heralded as a great achievement of the Renaissance. Does one use it in visual psychogenesis? Visual artists are sceptical, for soon after the Renaissance perspective was generally treated as a gimmick that was largely malleable. Even during the Renaissance, the better artists were aware of that. The ones that use linear perspective "as it should" stand out because of the unnaturally(!) rigid, mechanical nature of their works.

In an extensive study, I found that human observers systematically prefer a template rendering over geometrically correct perspective. These are not minor differences either; they are huge (Figure 8). I'm talking differences of a hundred or more just noticeable differences (Pont, Nefs, van Doorn, Wijntjes, te Pas, de Ridder & Koenderink, 2012).

Observers *never* apply perspective the way "they should," but they experience the three-dimensional environment in terms of simple templates. Thus, they do not bother with "Shape From Perspective" based on inverse optics (Koenderink & van Doorn, 1991, 1997; Szeliski, 2010) at all.

This makes perfect sense in view of the fact that one virtually never will view a perspective rendering with the (single!) eye at the "correct" place. Such would lead to long waiting lines in front of the more popular paintings in art museums. This never happens. Viewers assume any conveniently available position without protest.

This is the reason why wide-angle and narrow-angle (tele) photographs are often considered "distorted" and why the 50-mm focal length on 35-mm ("full frame") cameras is known as "normal." The "normal lenses," used in a "normal way," indeed yield renderings that closely fit the templates of the human OUL. They let you make photographs that are accepted as legal evidence in court, whereas (especially) wide-angle shots are not necessarily admitted without the added interpretation of some professional. This makes good sense as real-tors know.

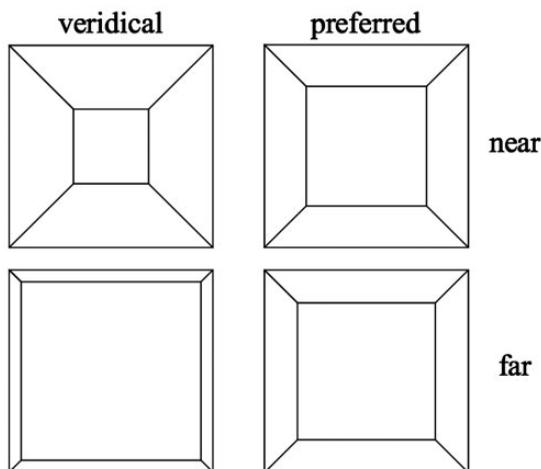


Figure 8. Veridical and preferred aspects of a wireframe cube. The near veridical rendering appears a long corridor, the far veridical rendering as a thin slab to the observers. The “preferred” renderings are means of adjustments by half a dozen observers. Most observers are quite happy with just a single template—that is how a cube looks to them from *any* viewpoint. They are not going to bother with perspective.

Temporal Coherence

The psychogenesis of awareness routinely parses and reorders events in “happenings” in such a way that the happening makes sense, that is to say, is really a “happening.” Happenings are atomic elements in the flow of awareness. An example would be a handshake; there is no such a thing as half a handshake. Happenings are of a short duration, the specious present; otherwise, they could not be elements of awareness proper. Narrative is very different from happening; it is not atomic, but more like a quilt.

Original work is due to Giovanni Vicario, a student of Musatti (himself a student of Benussi) at Trieste (Koenderink, 2012b; Vicario, 1973). Vicario worked on sequences of two or three tones. A sequence like DIH—DAH (high tone, low tone, each of 100 ms) is heard as a simple melody. He also introduced a short noise sound, say SZSZ, of 30 ms, which is atonal. He then noticed that a sequence like DIH—SZSZ—DAH is experienced (in auditory awareness) as DIH—DAH—SZSZ. That is to say, one hears a complete melody (an auditory happening) followed by a meaningless noise. This confirmed Vittorio Benussi’s work at Meinong’s laboratory at Graz.

Psychogenesis thus performs a temporal permutation of the elements of the auditory user interface (AUI) (Koenderink, 2012b). With modern technology, one can do the same with elements of the OUI. Techniques of film cutting rely on that (Dmytryk, 1984).

Spatial Coherence

It is well known in the arts that pictures easily survive rather violent spatial permutation of parts. Artists know this and frequently use it.

It is very hard to produce an incoherent picture (Figure 9). Picasso only succeeded for a few years in his cubist period, for soon enough many people “caught on.”

Psychogenesis produces coherent awareness even in the case of random noise images. In extreme cases, one becomes aware of coherent hallucinations that appear to fit the pictorial



Figure 9. The severe dislocations (check the nose!) in the portrait of Jakob Johann Baron von Uexküll (1864–1944) at right are not particularly bothersome. The grey “grout” is important, if one leaves it out the singularities stand out in focal vision (Koenderink, Richards, & van Doorn, 2012a; Koenderink et al., 2012b). In eccentric vision, the grout is not even necessary.

structure quite well. Leonardo da Vinci compared this with the awareness of the sound of bells in which one may hear all kinds of coherent auditive structures, like voices and so forth.

Spatiotemporal Coherence

When you chop a video sequence into spatiotemporal tiles and shuffle these locally, the resulting video looks awful. However, when you mask the spatial and temporal “seams,” observers report to experience a smooth video sequence “behind” a spatiotemporal grid due to the mask (Koenderink, 2012b; Koenderink, Richards, & van Doorn, 2012b). Severe spatial dislocations and temporal jumps into past and future become unnoticeable when the singularities are masked (Koenderink, van Doorn, & Wagemans, 2017c).

This suggests that psychogenesis imposes remarkable complicated template-like structures to come up with some sense. In this spatiotemporal shuffling experiment—strictly spoken—*there is no sense*, because the physical stimulation is chaotic. However, all psychogenesis is able to come up with sense. That is what it is for! That implies that it turns a blind eye to chaos and even “tampers with the evidence” when that is needed to come up with sense.

This again implies that perceptions are by no means “veridical” reflections of physical structure.

Multimodal Coherence

The ability to impose coherence on disorder is often a necessity in blending multimodal structure into meaningful happenings. For instance, since the speed of sound waves is far less than that of light waves, optical and acoustical structures due to distant events are

necessarily out of sync. Yet psychogenesis routinely lets you experience coherent “happenings.” This also happens in cases where there is no fact of the matter.

It is not at all rare to become aware of “causal events” where there is only chaotic structure. Psychogenesis has little of a choice here; if there is a single *chance* coincidence between different modalities, it is ready to come up with a causal presentation.

Here, you notice the creative nature of psychogenesis. You should cherish that ability! It is one major key to your survival.

Making Sense

Making sense is to one’s biological advantage, being “true to nature” as an ideal is not (Hoffman, 2009). Vision has a strong Procrustean nature. It has been honed by long evolution to render humans survivors rather than philosophers. Well, I’d rather be a survivor than a philosopher, so that suits me fine.

It also means that the spatiotemporal and causal frameworks of sensory awareness are not simply “found.” They are *imposed*, being a kind of OUI templates too.

Space

In this section, I discuss a simple toy model for the creation of the simplest spatial frameworks. A similar discussion, albeit somewhat more intricate, applies to space-time. The topic is the spatial structure of the visual field, which is a twofold extended manifold, filled with a simultaneous order of qualities. I limit the discussion to the simplest qualities such as colours and simple shapes.

The obvious model is the Euclidean plane, although it seems unlikely that anything like the Euclidean plane might make a biologically useful template. Very likely, there should be notions as point-like, line-like and area-like, but constraints with respect to size, rectilinearity and unboundedness seem hardly realistic.

A priori, one guesses that

- absolute size should not matter;
- anything has a finite size; and
- that structure is most likely based on local relations.

These are very strong priors and following such ideas up severely limits the routes one might follow. However, as phrased here, the simplicity of the priors is deceptive. It depends on how one construes the meaning of “point,” “size” or “local.” It should not be assumed that such meanings are obvious, or set once and for all. One has to look for definitions that give these priors a sense.

Any a priori notions are limiting, but one really needs them. It would seem to me that we have a good “nose” for what is of biological importance though. So I’m not going to offer excuses. If anyone is able to come up with useful ideas by breaking these prior notions, I’d be the first to listen, for it would surprise me.

The Notion of a “Point”

A way to start might be to consider the notion of “point” (Koenderink, 2017). According to Euclid (Euclid, 1956[ca. 300BCE]) “*a point is that which has no parts.*” This is an interesting definition, commonly misunderstood in the sense of “points have zero size” or “are very small,” but actually implying that “size” is not a defining property of points.

In my view, the definition is too limiting for biological purposes though. I propose to change it for: “*A point is that of which one ignores the parts.*” This may be reformulated to apply to the phenomenology of vision as: *A point of the visual field is that of which psychogenesis ignores the parts.* Points are *created* by ignoring structure; they don’t exist without such a choice by an observer.

This enables one to say that London is a point on the map of Europe. Thus, highly structured objects become “points” once you consider only their simple being. Points may have arbitrary sizes and may prove to be composed of numerous smaller points when you turn a microscope on them. They may also cease to exist as when London is not marked on a globe. In vision, one needs to consider points of all sizes or sometimes just ignore them at some scales.

Since points lack parts, they lack shape. London, on the map of Europe, is indicated with some conventional sign, which is only symbolic, not iconic (Koenderink, van Doorn, & Pinna, 2015, 2017b; Koenderink, van Doorn, Pinna, & Wagemans, 2016).

An Ideal, Formal Model: Canvas, Blackboard and Image

A point is an object for which one needs some operational definition. I can think of two, which I denote *sampler* and *brush*. Both definitions presuppose a context.

I conceive of a sampler as an operator, a kind of sensor that operates on images. A retinal ganglion cell or a Hubel and Wiesel (1959, 1968; Hubel, 1988) style line detector would be physiological implementations of samplers. The simplest sampler is like the “eyedropper tool” in Photoshop. (I use Photoshop as an example because it is so widely known and used, not because I have relations to Adobe, I have not). It returns the average over a convex area (say a 5×5 pixel area in Photoshop; you can select it in a menu).

The average will typically be weighted; thus, the sampler operator is characterized by a receptive field (RF) profile and location. Notice that the average is not an image but has only a single degree of freedom. That it has no spatial structure renders it a *point*.

For the sake of concreteness, I propose the term *blackboard* (no relation to the blackboard systems of Artificial Intelligence [AI]) for the totality of samples. It is a “representation” of the optical structure at the eye. (Although I use blackboard in the singular, there are actually many blackboards of various types.)

Cortical area V1 may be considered such a representation, but I prefer to stick to an abstract description. The blackboard is like the file of a forensic investigation. It is an orderly representation of structure in some easy to reference format, say *a proxy for the environment in brain readable form*.

The *brushes* are a different story. I conceive of the brushes as a kind of actuators that add local touches to a *canvas*, which is a surface that may receive and accumulate coloured patches. It is a descriptive term for the visual field of awareness.

The brush is like the “brush tool” in Photoshop. When applied, it deposits a touch on the canvas. Its touch is invariably the same, a template. It has a spatial profile that I denote PF (perceptive field). In Photoshop, you can select profile and size from a menu. Since templates have a fixed structure, they represent only a single degree of freedom, which renders the brush a point.

Notice again: A point can be of any size and if you change spectacles, it can even have structure! For instance, a rubber stamp may have a company logo and some text, but since there are no degrees of freedom (except for applying the stamp or not), the stamp will simply paint a “point.”

The canvas is a mental, the blackboard a physical (e.g., neural) entity. A point in the sampler definition operates on an image and yields a local sample. It is a kind of sensor. A point in the brush definition operates on a canvas and deposits a local touch. It is a kind of actuator in the mental realm (a “hallucinator” if you want).

The optical structure is not available to psychogenesis as an image. The blackboard represents the physics, say the *spectral radiance* at the cornea, a twofold extended field of radiant power spectra in such-and-so direction in some standard, brain-readable format.

The blackboard is a volatile buffer that is continually overwritten by the visual front-end. Psychogenesis also deploys involuntary fixations; thus, it has important control over the filling of the blackboard. Poking the world in its struggle to master it (Nietzsche’s “Will to Power” (Nietzsche, 1966/1886)) is the *primum mobile*.

The canvas is refreshed at each beat of psychogenesis. It is like the report of a forensic investigator that uses the blackboard as file. There is no necessary (e.g., causal) map of the blackboard to the canvas though. Neither blackboard nor canvas is to be conceived as some kind of “images.” Think of them in terms of abstract data structures.

Although the blackboard is a representation that depends causally on the optical structure at the eye, the canvas is a presentation that depends as much on the investigator as on the forensic file.

Representations are meaningless structures. Presentations are intentional creations.

Points as Samplers

The samplers are involved in analysis of the optics and the construction of a data structure that can be queried by psychogenesis. I discuss the sampling stage first.

A sampler is a machine that responds to the optical structure with a sample. To keep the discussion simple, I regard only monochrome here (it makes no essential difference), then the sample is a nonnegative number that encodes the number of photons absorbed per square degree of visual angle per second say. (Nothing but a convenient toy model.)

The sampler has some location and a size, but it does in no way know these itself. It is just a dumb machine. A record of a few million samples is fully *dystopic*, lacking any topological structure. The field of view is spatially well structured because “seen by God’s Eye,” I use the external description of physics for the radiance. One cannot assume psychogenesis to have access to that. The “visual field” of awareness and the “field of view” of optics are ontologically distinct entities.

In this section, I introduce a simple formal model. It is a formalism (“scale-space”; Koenderink, 1984a; Lindeberg, 1994a, 1994b; ter Haar Romeny, 2008) that is widely used in image processing, for instance of medical images on a daily basis. This is no small thing, since medical images have become such important diagnostic tools that medical image processing has become literally a matter of life and death.

The model is of a mathematical nature, indeed, and it provides a coherent and useful implementation of the Euclidean plane, in which the elements are geometry, algorithm (operator) and graphics and that has the power of differential geometry on all scales (Koenderink, 1990).

Since the formalism is well known and available in textbooks, I present only the minimal detail here and no formalism at all. This evidently yields a somewhat lopsided perspective. To understand the structure at a little deeper level, one should consult the (easily accessible) formal accounts (ter Haar Romeny, 2003).

Point RF (and also PF) profiles are based on the Gaussian (bell-shaped) function. It has formal properties that make the choice a uniquely necessary one (Koenderink, 1984a).

A sampler at some location, of a certain scale, is thus simply a bell-shaped RF centered at that location and with a width determined by the scale. It needs to be Gaussian, for a variety of reasons, various of which are mentioned throughout the text. One important property is that (I'm only considering isotropic) Gaussians are both centrally symmetric and can be factored in Cartesian coordinates. Gaussians are unique in this respect. Since there are supposed to be samplers all over the place, you can conceive of "the optical structure at a given scale" as represented by many samplers of a given size, each taking a bite out of the optical structure, albeit each at its own location. Technically, this is a convolution.

The result is a "representation" (in the sense of a forensic file, that is an abstract data-structure in the blackboard) that might be thought of as a "map" of the optical structure at a given scale, the size of the RFs. Notice that "map" and "representation" are terms that belong in the description of an external observer. That makes sense, for V1 activity (say) is only ever seen by external observers. The agent itself only has presentations; it cannot see its blackboard.

Scale-Space

One may conceive of an array of samplers, covering the field of view (ontologically distinct from the visual field) with sufficient density. The density will depend on their size, suppose (for the moment) all are of the same size. In a somatotopic setup, I then gain an array of scalar values that may be understood as essentially a (continuous) scalar field. Notice that this applies to an external God's Eye view: The points themselves do not know who their neighbours are.

This setup may be repeated for arrays of points of different size. In fact, to obtain a good representation, you need *all sizes*. If you leave out small sizes you lose resolution, if you leave out large sizes you lose coherency. The former should be obvious, but the latter point perhaps needs some discussion. If the points do not know who their neighbours are, you cannot combine them into larger patches, thus you cannot distinguish features larger than point size. The remedy is to have a supply of points of any size.

The sample array for any given point size is much like the optical structure itself. Thus, it can itself be sampled, although only by points of larger size. This yields an interesting structure, because two points that are each sampled by the same larger point will both be correlated to that larger point, even if they are themselves not mutually correlated. Thus, the correlation structure of the samples contains the neighbourhood relations that the points themselves do not know. Having simultaneously many sizes leads to a structure that implicitly represents the topology of the field of view in the correlations of its activity. This crucial notion was proposed by Helmholtz as a possible solution to Lotze's problem of *local sign* (Koenderink, 1984b).

Such sampling systems have been formalized as *scale-space*. As said, they play a key role in modern diagnostic image processing.

Consistency requires that the samplers have Gaussian RF profiles. In such, a system it does not matter whether you sample the optical structure in one step or by way of an arbitrary number of intermediate steps of samplers sampling samples. This is obviously a very desirable property, technically known as the "semigroup property."

With scale-space, we have constructed an ideal system that sports points of *any* size. In practice, one will surely have a largest and a smallest size, of course. The ideal system is useful because it helps to bring out the essential qualities. Scale-space as a whole is size-invariant; it looks the same at any scale, obviously a useful property for a visual system, which deals with similar objects at different distances. Changing the viewing distance implies

a move in the scale dimension. This is still approximately true in actual implementations; the ideal system serves to open your eyes to that important property *in vitro* so to speak.

The scale-space may be regarded as an implementation of the Euclidean plane, a “geometry engine” (Koenderink, 1990). In image processing, the scale-space is the basic data structure that all kinds of algorithms designed for diagnostic purposes draw upon. It is an almost exact model of the controlled hallucination notion of psychogenesis, except for the fact that these applications do not use the correlation structure. They don’t need to, since they use the external God’s Eye view to skip the local sign problem.

The natural sequence in building scale-space is *fine to coarse*. It can be regarded as a progressive summarizing of spatial structure.

In contradistinction, the natural sequence in using scale-space is from *coarse to fine*. One starts with summaries and only probes into finer detail when necessary.

This saves enormous amounts of unnecessary references. It also enables processing, because going into detail implies already having a context (the summary).

The consistent structure of scale-space guarantees that there will never be surprises (in the sense of novel local structure) as you move to coarser scales; surprises only occur in the other direction. Moving to a finer scale may reveal hitherto unexpected structure.

Of course, this does in no way imply that moving to coarser scales never yields anything new. The point is that it allows you to increase your scope (I’ll discuss the “atlas structure” later). You may notice details that you missed at the finer scales because you failed to see the forest for the trees.

This is how biological systems work. I may tell you where some item (say a key, or a credit card) is by saying “in that country, that city, that street, that house, that room, that desk, that drawer, somewhere near the bottom.” Although the item is hand-size but a thousand miles away, you will surely find it. You can literally jump continents by simply saying “no, not in North America, go to Europe!” Just imagine how much of the structure of the in-between space I ignored in this instruction (Riedl, 1984)!

This is the key to our grasp of spatial structure. Psychogenesis of visual awareness uses it all the time.

Laplacian Scale-Space

In biological systems, absolute physical measurements are problematic. Measures of differences or ratios are much more robust. Hardly surprisingly, one routinely encounters local adaptation, center-surround structures and so forth in neural systems. In addition, in a biological system, it is evidently not efficient to keep redundant data (Barlow, 1972). For instance, it makes little sense to encode the blue sky as a data structure of millions of identical samples. Hence, I consider variations of the basic scale-space setup in order to address such issues.

A simple way to get rid of absolute measures and repetitive inefficiency is to consider only scale-space *slices*. Taking the difference between two scale-space arrays of different scales suffices to get rid of the blue sky problem. You can compute such a slice in one go if you introduce so-called difference of Gaussians (DOG) RFs. Such DOG samplers yield a scalar value, just like a point sampler; thus, they are equally “atomic.” As evident from their samples, they have no visible internal structure. The nature of the RF depends upon the ratio of sizes of the two Gaussians that are point samplers.

The DOG RF is really the “difference of two points.” In the limit of two almost identical points, it becomes the “scale tendency,” or “scale derivative.” This is an important RF by

itself, known as the “Laplacian.” Its sample indicates the change of a point sample for a minor size variation.

As the “difference of two points,” the DOG RF can be understood as a point coming with its own background. It is perhaps not unlike a continuous track vehicle, like a tank, that “brings its own road.” A single-point sampler of any size will react to the blue sky, but the DOG RF will not. It will not, because “the point” has the same colour as “the background” and thus is invisible. That is why DOG RFs are only sensitive to local details that fit their own size.

What happens to the sample as one changes point size of course depends upon the immediate environment of the point. Thus, one expects that the scale tendency of a point, the Laplacian, can be obtained from the local spatial variation of point samples. Such is indeed the case as one easily verifies by explicit calculation. This is very important; one has a partial differential equation that relates local spatial structure to local scale structure. It provides scale-space with a dense, coherent structure (Koenderink, 1984a).

One finds that the Laplacian can also be obtained from the difference of the point sample with the average of the point samples of its neighbours. This is very remarkable and important: The spatial structure at any given resolution level of scale-space determines the change of that layer that occurs when you vary the size of the points.

This renders scale-space the tight spatial nexus that it is. Technically, it is a partial differential equation fondly known by physicists as the “diffusion equation,” or “heat equation.” Scale-space is a very intricate nexus of relations in which spatial variations and variations over scale are tightly coupled.

A final remark on scale-space involves the “ideality” of the formalism. A reviewer remarked that the model only makes sense to the extent that actual RFs would be in the Gaussian family (and he knew for sure they are not). In my view, that is a rather unproductive way to view the matter. Such a model is useful in two distinct ways. First, it is a thought model, a way to rationally understand a type of implementation of the Euclidian plane. Second, it might be posed as a model of some neural structure. Of course, it cannot be expected that such a model will fit perfectly. The issue is how well it does, and if not sufficiently well, how it can be amended. I’m personally convinced it is the most useful model of V1 we have today.

Although the real beauty of the scale-space structure is only evident from the formalism (the well understood diffusion equation), an intuitive notion can be gained by considering the “atlas structure” as explained in the next section. Understanding that is even better than merely understanding the formalism—understanding both is better still, of course.

The Atlas Model

Perhaps the most intuitive way to understand the scale-space structure is to understand it as a (geographical) *atlas*. It is a structure that seamlessly fits the scale-space structure (Griffin, 2019). An atlas is a structure that is foliated into sheets, or pages of different resolution, a scale-space.

There is a page on which London is a mere point, possibly a page where it is not even indicated. However, there are also pages where London appears structured, perhaps distinguishing the House of Parliament and St Paul’s cathedral.

“Paging” the atlas—that is moving through the resolution domain—is not symmetrical. One way (to lower resolution) “summarizes” and thus leaves out detail, whereas the opposite way (towards higher resolution) reveals novel detail.

Because all the pages are of the same finite size, you may not be able to compare details that are far apart on any given map. The map that contains both may not even show them! This will probably happen if you want to compare the House of Parliament to the Moscow Kremlin.

Two places that occur on a single page may be said to exist together *up to the resolution and scope* of the common map. The “size of a page” corresponds to a limiting “scope” in psychogenesis. It may perhaps be estimated from “Bouma’s Law” as known from crowding studies (Bouma, 1970). It will be larger than the various “regions of interest” used in psychogenesis.

You may not compare details that are not on a single map. Some places may not occur together on *any* sheet. They are *heterotopic*; there are no spatial relations between them except those borrowed from places they are contained in. Thus, there are no relations between the House of Parliament and the Kremlin except those borrowed from London (or even the UK) and Moscow (or even Russia). Thus, you may say “the Kremlin is east from the House of Parliament,” but you may not say “the Kremlin is 2498.33 km from the House of Parliament.”

The Laplacian is an interesting operator. I discuss two aspects here, one concerning “edges,” the other concerning the overall structure. The latter aspect is simplest: Since the Laplacian essentially represents a thin scale-space slice, adding up all Laplacian layers must reproduce the original image! That is indeed easily proven formally.

What this means is that the Laplacians used as samplers yield a *complete representation* of the optical import except for the overall average. It reveals the Laplacian samples as (mer-eological) *parts of the image*, the whole bunch of them *adding up to the image*. This equation is a *representation theorem*, expressing the fact that the Laplacians yield a *complete* representation of the structure.

This opens up many possibilities of freedom in a synthesis of the image after analysis by Laplacians. If you can take something apart, then this makes it possible to change it by putting the parts together again, but in different ways, or after changing the parts. Thus, being able to analyse and resynthesize yields a very powerful control. One might selectively ignore, or weight, Laplacian layers so as to obtain a variety of distinct syntheses.

This is of interest to models of psychogenesis because it suggests the kind of “handles” that may be wielded in psychogenesis. The selections may be of various nature and could be a priori (like seek images), or a posteriori (for instance, based on “found” edginess).

I have recently used this to implement algorithms that let you produce various “eidolons” for a given image (Koenderink, Valsecchi, van Doorn, Wagemans, & Gegenfurtner, 2017a). The blackboard representations of such eidolons deviate in controlled ways from the blackboard representation of the original.

The eidolons may find numerous applications in vision research. They are also of interest as tools in the visual arts.

Edges

Laplacians are usually not thought of in connection with “edges”; the latter are thought to be derived from “edge detectors.” There are various mutually distinct misunderstandings in place here. This will take a somewhat longish explanation.

Paintings are distributions of colours over a planar substrate. Painters understand them in terms of coloured patches (“*machie*”; Koenderink, 2013b), made famous by the Macchiaioli (Boime, 1993, the originators of impressionism) and of “edges.” However, a

painter's "edge" is very different from the dividing edge of the colorimetric half field. Edges can have many different qualities, magnitudes and sizes.

It is perhaps best to think of the macchia and edge descriptions as mutually complementary. There are no macchie without edges and no edges without macchie. Edges are like boundary or transition regions of macchie (Koenderink et al., 2016).

In nature, there are few instances of optically uniform areas. (Sure, the blue sky, what next?) The macchie in presentations are extended regions in which psychogenesis ignores internal structure. In that sense, a macchia is like a point; it is lacking (some) parts. However, macchie do have a *shape*; it is one of their important properties.

How does one sample an "edge"? Intuitively, this involves two abutting uniform areas and at least two points, one in each area (Figure 10). The difference of the point samples would be an apt measure, for it would vanish for two points in the same macchia but be appreciable when each point happened to be in a different one. Thus, an "edge detector" should be a *bilocal operator*, two neighbouring points that agree to make only the difference of their samples public.

Such an object would thus be atomic (no spatial structure evident from the outside), but it would have a *direction*. Formally, it would be a directional derivative operator, or, what amounts to the same, a tangent vector.

Such operators are well known and usually denoted "edge detectors." Formally, such directional derivatives are implemented by RFs that are derivatives of Gaussians. One arrives at a very neat description in which derivative operators can both be implemented physically and model the formalism of differential calculus exactly.

What that means is that any formal expression from differential geometry can immediately be implemented as a local neural network. Thus, one obtains a powerful "geometry engine."

I have speculated that cortices are exactly such geometry engines (Koenderink, 1990). One knows that the structure of essentially all cortices is essentially the same, no matter whether visual, auditive or motor cortex. How can that be? One speculation is that they are all implementations of differential calculus. It is like the fact that physics texts on very different topics tend to look alike to lay people, obviously because all those formulas look similar (Figure 11). Indeed they do. Most of it is differential calculus, formally the same although the meaning (thermodynamics, electrostatics, hydrodynamics and so forth) is different.

Notice that the simple bilocal operator in no way "measures the direction of an edge." All an operator has is a single direction, its own. It doesn't have to "know" it, even in relation to its neighbours. Such problems are for the next guy to consider. Another bilocal operator, located at the same spot, may be interested in another direction.

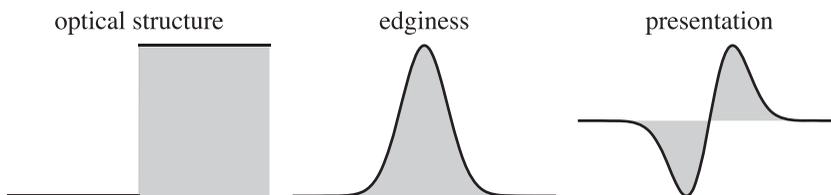


Figure 10. The process of "edge detection." At left, an "ideal edge" in the optical structure. At center, the "edginess" ("how edgy it locally is") sampled by edge detectors. At right, the local presentation of the edge. It is the "touch" of a Laplacian brush. (Formally, it is the same as an edge detector profile: notice the possible confusion!)

Einstein-Maxwell-Gleichungen [Bearbeiten | Que

Wird für $T_{\mu\nu}$ der elektromagnetische Energie-Impuls-Tensor

$$T_{\mu\nu} = -\frac{1}{\mu_0} \left(F_{\mu}^{\alpha} F_{\alpha\nu} + \frac{1}{4} F_{\alpha\beta} F^{\alpha\beta} g_{\mu\nu} \right)$$

in die Feldgleichungen eingesetzt

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R = \frac{8\pi G}{c^4 \mu_0} \left(F_{\mu}^{\alpha} F_{\alpha\nu} + \frac{1}{4} F_{\alpha\beta} F^{\alpha\beta} g_{\mu\nu} \right)$$

so spricht man von den *Einstein-Maxwell-Gleichungen*.

Figure 11. A randomly cut piece from a physics text. To lay people essentially all such material has the same texture. Here, the topic happens to be gravitation and electrodynamics, but hydrodynamics or continuum mechanics look the same. (https://de.wikipedia.org/wiki/Einsteinsche_Feldgleichungen)

The “direction of the edge” cannot be found from the activity of any single bilocal operator. It can only be found from the simultaneous activity of a bunch of them. One needs at least two, if well chosen, in practice a dozen or so. (Notice that this is not to be confused with the so-called aperture problem.) Even so, this assumes that psychogenesis somehow knows the preferred directions of the bilocal operators, or, if not, at least some relations between the preferred directions of a bunch of such things.

There is a world of problems here that modern brain science has in no way understood; worse still, the problems tend to be ignored if they are perceived at all (I don’t think they are).

A “detector” is a machine designed to spot some existing object or state of affairs. Thus, “edge detector” presumes the existence of “edges.” This is not such a great idea, because the only acceptable definition of edge is “that which an edge detector detects.” The bilocal detectors yield some sample at *any* location, but are there edges everywhere? That would render the intuitive notion of edge useless. Edges should be special and occur only at relatively few points of the image.

A simple way out is to define “edges” via some arbitrary threshold value. This is indeed the usual “solution.”

Another way out is to realize that edges only exist in the mind. Thus, you cannot “detect” them, but psychogenesis *creates* them. (On the basis of current blackboard activity, or simply as a phantasm, as often happens in reading paintings.) I would agree, but here, I take a slightly different bend.

I regard the so-called edge detectors as elements of the *differential geometry of the field of view*; they are its tangent vectors. At any point, a basis can be set up composed of two of these in mutually transverse (orthogonal is convenient) directions. Any other direction can be obtained by linear combination. At any point, I find the direction of greatest response and define the largest sample as the local “edginess.” There will be some edginess at any location, but in many cases, one finds longish strips of high edginess that might be identified as “an edge.”

This yields a distribution of “edginess,” that is something like a measure of “*presence*” of some boundary region (Figure 12). It fails to report the *nature of the edge* though.

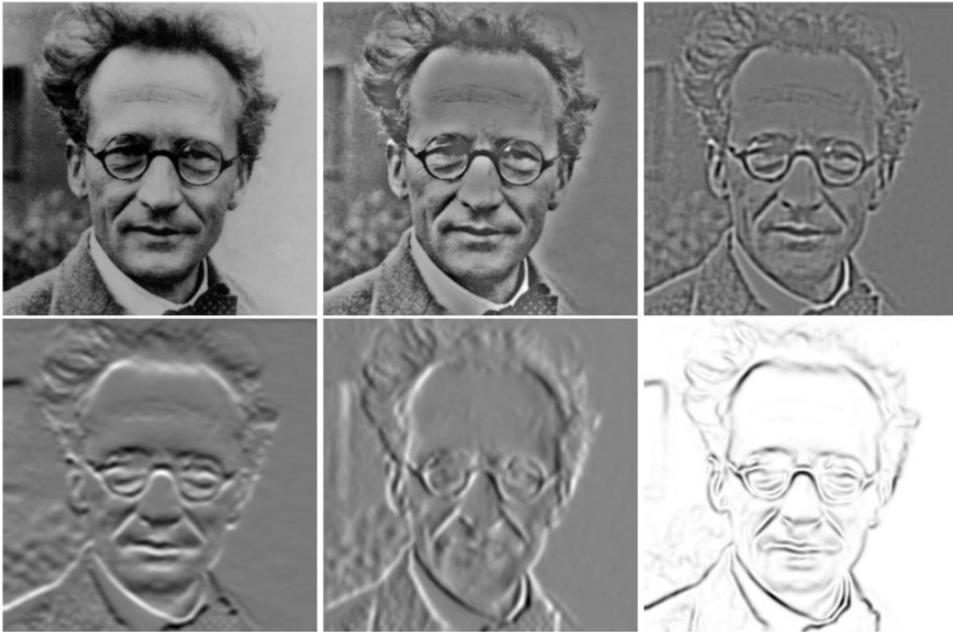


Figure 12. Top left original, top center a rather fat scale slice (this would be a DOG sample), top right the Laplacian at a certain scale (the thinnest scale slice). At bottom from left to right: vertical derivative, horizontal derivative and edginess (of course, all at a certain scale). The Laplacian and the edginess were computed at the same scale; it is of interest to study their differences. (The picture is of Erwin Rudolf Josef Alexander Schrödinger (1887–1961), of quantum mechanical fame.)

In the latter respect, Laplacians do much better than edge detectors (Figure 13). The strange fact is that Laplacians yield zero samples where the edginess is highest! However, that is exactly their strength. They show what the macchie at either side are like and at the location of the edge there is no answer to that. (Do not confuse this with the “Marr-Hildreth edge detector,” which is indeed an “edge detector,” albeit based on the zero-crossings of the Laplacian—it computes “edginess.”)

Notice that boundary regions are local, although perhaps mostly elongated in some direction. They are not necessarily closed so as to outline a region, but you may even have “edginess” at an isolated point! Macchie are usually not fully outlined; they may just peter out or change character somehow. In painting, this is known as the “lost edge,” or the “lost and found” property of edges.

These are the qualities that the painter looks for and intentionally introduces (Jacobs, 1986). They are also the qualities that keep viewers of a painting interested in paying it repeated good looks after the initial impression.

Paintings that fail to achieve that are mere wallpaper at best. A painting should look better than the wall behind it, or it will not be hung.

Encapsulated Geometry

The Laplacian samples something like the difference between the samples of two points (samplers) of slightly different sizes, whereas the edge detector samples something like the



Figure 13. Left the edginess after suitable thresholding (if the threshold is picked inappropriately one could end up with uniform black or white images). The edge detectors report “where the edges are,” but do not reveal their nature. Right the Laplacian at the same scale, the positive and negative values independently thresholded. The Laplacian reveals the nature of the edges, resulting in a clear image. Perception fills in the grey areas; this is Pinna’s “watercolor effect” (Pinna, Brelstaff, & Spillmann, 2001).

difference between the samples of two points (samplers) of the same sizes but slightly different locations.

Both are *atomic* in the sense that one cannot detect the two-point structure from their samples—which are just numbers. Of course, the physiologist routinely measures RF profiles, but that is only possible as external observer in control of the optical structure and combining recordings collected over time.

The Laplacian and the edge detector are examples of “encapsulated geometry,” albeit of the simplest kind—just two points. This is an important concept. Geometrical configurations are multipoint properties. In order to recognize them, one needs to be able to *know where the points are*, at least relative to each other. (In visual awareness, a notion of “absolute location” makes no sense.) This is not a problem for an external, God’s Eye view, but it is a major problem for psychogenesis since the samplers don’t broadcast their location—it is not even evident how to define that notion. For an encapsulated geometry, this problem does not arise since one deals only with scalar samples, not locations.

This is the advantage of *templates* such as edge detectors. Of course, one may consider templates of much higher internal complexity. Indeed, there is physiological evidence for such. However, the problem is that the number of possible templates grows very fast with the number of points. This may not be a major problem if one combines templates of different resolution in a hierarchical fashion though. The “sweet spot” for the size might be arrays of three-by-three samplers (Griffin & Lillholm, 2007). The Laplacian can be implemented that way, so can Hubel and Wiesel’s “line detectors.”

Formally, there are many relations between the various encapsulated entities. For instance, it is easy to show that the Laplacian can be obtained as the mean of line detectors over all orientations. One can use this to prove additional *representation theorems*.

For instance, the optical structure is completely represented through an atlas of line detector maps. This proves that a blackboard that has the structure of cortical area V1 is a complete representation of the optical structure.

In neuro-speech, the average over orientations would be considered a pooling of activity over a cortical pinwheel. The formal representation theorem implies that the totality of pooled multiscale pinwheel activities “represents” the retinal image.

Few people will be surprised by such a statement. I guess it tends to be *assumed*, although I never encountered a formal proof. Neurophysiology is simply not interested in the fundamental roots of their frameworks, very unlike fields such as physics. Coming from physics roots myself, I was surprised to notice this, and I still have difficulties accepting it for a fact.

The “line detector” RFs considered as brushes (thus PFs) paint local edges on the canvas. The representation theorem has that the totality of all these touches simply paints the optical structure; it would be like a detailed photograph. That would really be boring, perfectly fit for robots, but it seems at odds with the facts of life.

The artist, in painting a scene, may also use edge touches (painters often discuss “edge quality”), but they paint only what is necessary. They pick and choose. Some edges are painted, others not. Some edges are accentuated, others played down. A touch may be deposited at a location that is somewhat “off.” They may throw in something that is not in the optics but “should have been there.” That is why paintings can appear more real than optical reality reveals.

Likewise, psychogenesis may well use the line detector PF-profiles representation, yet present something quite different from a straight photograph. More life-like. I mean that quite literally, for one function of visual awareness is no doubt to promote successful living, not necessarily trying to slavishly imitate the God’s Eye view—which doesn’t exist anyway. It is good to know that the blackboard representation in the visual cortex is complete though. Perhaps it makes you feel more secure.

From a formal perspective, the kind of templates I discussed here are elements of the differential geometry. The edge detectors implement the first order, the tangent vectors or gradient. The line detectors implement the second order, the Hessian. And so forth.

An immediate consequence is that a basis of just two edge detectors suffices to predict the samples of edge detectors of *any* direction. This is the trivial consequence of their vectorial character (Freeman & Adelson, 1991).

Similar constructions can be set up for the line detectors, then the minimal basis is three. This makes for a very elegant description of V1 structure from a formal perspective. However, neurophysiology is not interested in such global, formal aspects; it tends to focus on details. In comparison with physics, the field has primarily a myopic perspective.

The advantage of this view is that local geometry can naturally be expressed in terms of these entities. For instance, a “corner detector” can easily be put together from these (Koenderink & Richards, 1988; ter Haar Romeny, 2003). A collection of templates covering the differential geometry up to some (fairly low) order would enable psychogenesis to query the sensorium with “questions” of a geometrical nature.

Points as Brushes and the Canvas

The brushes “paint an image in awareness.” One has to see that as a model of psychogenesis that suggests the resulting presentations. Of course, I do *not* want to push the notion that there is something like a physical “canvas” in the brain! This is an abstract model. Canvas and blackboard exist on different ontological levels. Visual awareness is not like looking at brain images. Where would such image be anyway? The brute fact is the presentation.

The best I can think of to convey the nature of presentations is by way of images that are likely—as judged by your own experience—to evoke the kind of imagery that I intend in the reader.

It is much the same problem as that of the communication of thought through language. Some things cannot easily be communicated in words; thus, formal mathematics and logic are very useful indeed. They extend language in one direction.

But not just anything can be formalized. Some thoughts are better conveyed poetically, perhaps. That extends language in another direction.

In any case, every reader will translate a text differently. It is like that with images. But—like poetry or mathematics—images are a different medium than plain text. Communication by way of images need not rely on an intermediary translation into language, which would defy the purpose. Images are a language of their own.

Again, any observer is likely to pick up something different, for that is the way communication is between biological agents.

The objective of a model of psychogenesis can hardly be a “veridical” image of the spectral radiance at the cornea. Psychogenesis is necessarily focussed, so the presentation will certainly reflect only part of the available structure. There is no reason why all parts of the canvas should receive some paint. The “un-painted” regions are simply not part of concrete actuality.

There is also no reason why everything on the canvas should causally derive from the optics. The “in-painted” regions are part of concrete actuality, but not of reality (if you are a believer in God’s Eye).

There is also no reason why the image should reflect the metric or even the topology of the field of view “veridically.” As in an actual painting, each touch is intentional. There is nothing in a painting that the painter has not painted. This is a categorical difference with photography. Any straight photograph contains detail that the photographer was not aware of at “the moment it clicked” (McNally, 2008). That explains the strength (and weakness) of painting, as well as the strength (and weakness) of photography.

Psychogenesis paints both less and more than is in the sensorium. I may model the sensorium as the Laplacian scale-space derived from the spectral radiance at the cornea. Psychogenesis certainly needs to paint the blue sky, but the blue sky is not represented in the sensorium, except as a certain boundary area, or “edge.”

Thus, the image will contain both more and less than what is in the sensorium. It may even contain structures that were not in the spectral radiance at all, imagery that came out of the depth of the mind, perhaps imagery containing traces of things past, or anticipations of things future.

Without the creative, automatic imagination there can be no vision, for visual awareness is due to a creative process (Sherlock Holmes at work!); it is not the result of some transformation or computation on the retinal irradiance pattern. It is also good to remember Kant’s phantasmatic self-stimulation and von Uexküll’s seek images.

There is literally far more to vision than ever “meets the eye,” that is the physical eye. Presentations present what’s in the mind’s eye (they *are* the mind’s eye).

Topology of an Array of Brush Touches

You cannot reckon on any somatotopic map from the sensorium (likely somatotopically derived from the spectral radiance) to the canvas, because the canvas is not somatotopically implemented. The topology of brush strokes has to be derived from various things, but cannot be considered “given.”

Indeed, it is evident from an agnosia like *tarachopia* (scrambled vision; Hess, 1982) that it is somehow created by psychogenesis. Phenomena like crowding (Bouma, 1970) and the nature of vision in the periphery suggest that *tarachopia* is a fact of life for all of us.

The presentations are certainly *taratopic* to various degrees. On the timescale of good looks, they also contain mutually *heterotopic* areas (Figure 14). These characteristics of visual presentations are also typical for most paintings in some naturalistic tradition.

This brings up Lotze's (1852) problem of *local sign*.

I already mentioned that the geometry of "visual rays" (Koenderink, 1982) fanning out from the eye fails to be a brute fact. For most people, the brute fact is that the presentation is of everything being "in front of them." But that relates to spatial relations exterior to the body.

Might it be different for spatial relations of activities that play inside the skull? Lotze considered the very idea outrageous, but contemporary brain science accepts a close correspondence as a natural consequence of somatotopy (Saladin, 2012).

Brain science apparently conceives of the canvas as a brain area that receives a somatotopic projection and is somehow "watched by the mind." I find it hard to make sense of that, so I tend to side with Lotze that here is an unsolved problem. But many brain scientists

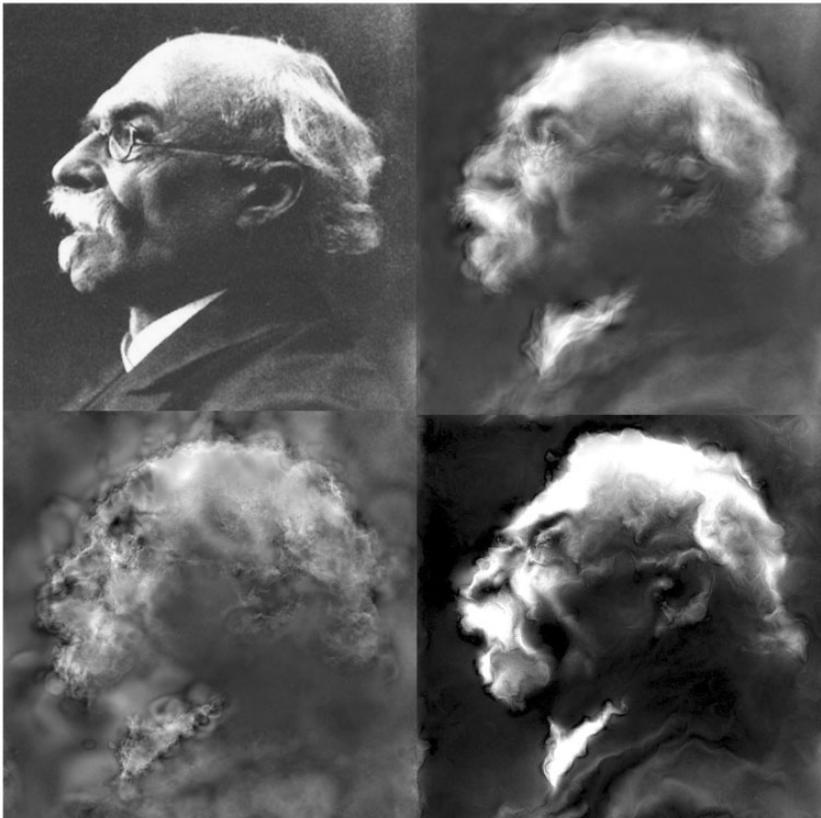


Figure 14. Some examples of taratopic presentation. At top left is the eutopic case. Top right shows the effect of disarray applied to touches of all sizes. (Of course, this is a mild case, increasing the magnitude of the disarray soon renders the image unrecognizable.) At bottom left, the disarray magnitude was scaled with the size of the touches. At bottom right, the disarray was such that large touches draw small ones in their neighbourhood with them. Although a quite large perturbation, this is not a big deal in awareness since local structures survive. (The picture shows Rudolf Hermann Lotze, 1817–1881).

equate mind with brain, so they have no problem. Their problems occur in a different department.

Local Sign

Lotze (1852) speculated that psychogenesis may use a sensorimotor map derived from experience with the effect of eye movements to set up a system of local signs. He may well have been right at that, but it seems not likely that the precision would reach the level of visual acuity.

It might well be sufficient to give psychogenesis a head start though. The correlation mechanism identified by Helmholtz then may take care of the details.

Small random eye movements about the fixation point may “sniff out” the correlations. It might be similar to Ahissar’s (Ahissar & Arieli, 2001) model of the vibrissæ systems of many mammals. (Ahissar’s mechanism uses the eye movements [instead of the whiskers] to map temporal structure to the spatial domain. It is well suited to use in combination with Helmholtz’s mechanism.)

The computational problem of computing topology from correlations soon becomes unmanageable when the number of points involved becomes too large. Thus, I speculate that psychogenesis may have access to precise topology only in the currently fixated area. Of course, it may have such topology on various levels of resolution.

The kind of *eutopic*, coherent region of interest (ROI) may be something like a seven-by-seven points area. Seven is indeed a “magical number” here (Miller, 1956), but much less would be useless and much more unmanageable. Moreover, seven times seven is of the order of the (unilateral) number vibrissæ of many mammals.

Outside this focus of coherence the (much larger) ROI would be subject to uncertain topological relations, so it would be *taratopic*. Such taratopic areas would still be very useful, especially if the fully coherent areas were dynamic, say “real-time” on the scale of a 10th of a second.

Nature of the Array of Brush Points

So what would the presentations be like? Of course, we all should know, since we live by them! However, just because of that they may have become “transparent” to us, which implies invisibility to the mind’s eye. The mind’s eye cannot see itself either.

A lifetime is insufficient to become (cognitively) aware of (precognitive) visual awareness. It makes a great exercise in “mindfulness,” known by artists as “learning to see” (Koenderink, 2018). Most people couldn’t care less, whereas most vision scientists get stuck on the cognitive level.

As I remarked earlier, the canvas is not necessarily completely filled, although the “holes” would never be part of concrete actuality. Concrete actuality is naturally complete as it is.

Taratopic areas would be characterized by a spatial disarray of a stochastic nature, not unlike that seen in the “mongrel” stimuli so elegantly wielded by Rosenholtz and collaborators (Balas, Nakano, & Rosenholtz, 2009). However, any *specific* pattern of touches will hardly be part of concrete actuality. The disarray simply fails to become a brute fact. In that sense, any “frozen” image will be misleading, but so will a dynamic image. The touches simply have locations that are somewhat indeterminate, although it is not very well possible to paint that. Maybe Picasso’s *Girl with mandolin* of 1910 (Figure 15) conveys an impression.

What does appear in concrete awareness are macchie, edges and so forth. Phenomenologically, these are often presented rather more precisely than the disarray



Figure 15. Pablo Picasso, c. 1910, *Girl with a Mandolin (Fanny Tellier)*, oil on canvas, 100.3 × 73.6 cm, Museum of Modern Art, New York (public domain). At least somewhat conveys a local taratopia.
 Note: Please refer to the online version of the article to view the figures in colour.

“allows.” That is to say, the brushes are applied such as to yield as coherent a structure as possible.

One always sees “something,” although partly—if need be even largely—“imaginary.” This is evident from experience in crowding situations where observers give various descriptions to the same stimuli.

This tendency to impose spatial order or *Prägnanz* was empirically studied in the early 20th century (Koenderink, van Doorn, & van Pinna, 2018a). It is currently mainly known from pop-science accounts of “Gestalt Laws.” *Prägnanz* is like a measure of *eutopia*; it applies to the presentations rather than optical structures. (Although it is common enough in the literature to see *Prägnanz* applied to optical structure.)

Although “concrete actuality,” presentations have little to do with “reality.” Visual objects are part of actuality but are beyond reality in the sense of Meinong’s “beyond being and nonbeing” (*G. Außersein*). That is why it is possible to see a PERFECT SQUARE. Even the most expensive computer graphics may at best *approximate* a square, but one can certainly *see* one. Kasimir Malevich famously painted one (in 1915) for everybody to see.

The geometrical configurations I used to draw on the blackboard for my students (while still a professor of physics) were *perfect* entities, despite the dirty blackboard, the cracking chalk and my physiological tremor. At least, they certainly were to the students who had the right eyes for that. That latter remark hits the core of the matter. It reminds one of the story of Archimedes death (c. 212 BCE) at the sword of a Roman soldier who had no eye for his

circles, drawn in the sand (according to some accounts his last words were *Noli turbare circulos meos*).

That concrete actuality is not the same as reality finds spectacular demonstrations in accounts of ethology. The assumption that “humans are different” is painfully anthropocentric (Boddice, 2011).

Attempts to account for visual awareness on the basis of “inverse optics” are based on a mistaken understanding of vision by confusing reality with concrete actuality.

The tree one sees is not in front of the eye, nor is it a physical object. What is in front of the eye is for physics to investigate. It will certainly be different from the presentation, for presentations are idiosyncratic. You, your dog and a bird will have different trees in their concrete actuality, so much is evident from the way they tend to interact with their trees.

If “the” tree is what is common to their awarenesses, it can’t be much! If “the” tree is the totality of the awarenesses of all sentient beings that interact with it, it is an object too complex for science to deal with. The “physical tree” is easier, but it cannot be seen, only appreciated from scientific reports, most of which you’ll be unable to understand unless you are a professional in numerous fields. Moreover, it is not clear in what sense such an “overall AWARENESS” (capitals seem appropriate) might exist.

In God’s Eye perhaps. But as a scientist, I cannot work with that notion, I have no access to it.

As a person, I’m perfectly happy with the tree I see. I know it is only *my* tree, but that’s all I need anyway.

Qualia Dimensions

“Space” as such is not a brute fact, but *spatial relations* in the sense of “simultaneous order” between various qualia are. In the previous section, I essentially used tonal distributions (grey level fields) as an example. Think of chromatic variations and you notice “space” being transformed into something else.

A formal description is readily available from mathematics in the form of “fiber bundles” (Koenderink, 2012h; Seifert, 1933). A fiber bundle is composed of a single *base space* and numerous copies, called *fibers*, of another space. One fiber is attached to each point of the base space. For example, you may use *colour space* for the fibers. A “cross section” of the fiber bundle singles out one point in each fiber. The “bundle projection” assigns a point of the base space to each point of the cross section.

This is just a formalization of the intuitive notions. A painting is a cross section (different cross sections would be different paintings) of the fiber bundle composed of colour space over the plane. The bundle projection is just the plane without the colours; it has only a virtual being, all one knows are various cross sections.

This is what I mean when saying that “space itself is not the brute fact,” whereas the paintings are and “have space.”

This enables one to consider presentations as cross sections of the fiber bundle composed of the canvas with colour space as fibers. You can likewise consider the optical input as a fiber bundle with the space of visual directions as base space and the space of radiant power spectra as fibers. In research, a cross section of the latter would be a stimulus, a cross section of the former a presentation evoked by that stimulus.

Theories of contrast, assimilation and so forth now can be framed in terms of functional relations between these cross sections. It is a general formalism, for which mathematicians have thoroughly investigated the generic properties.

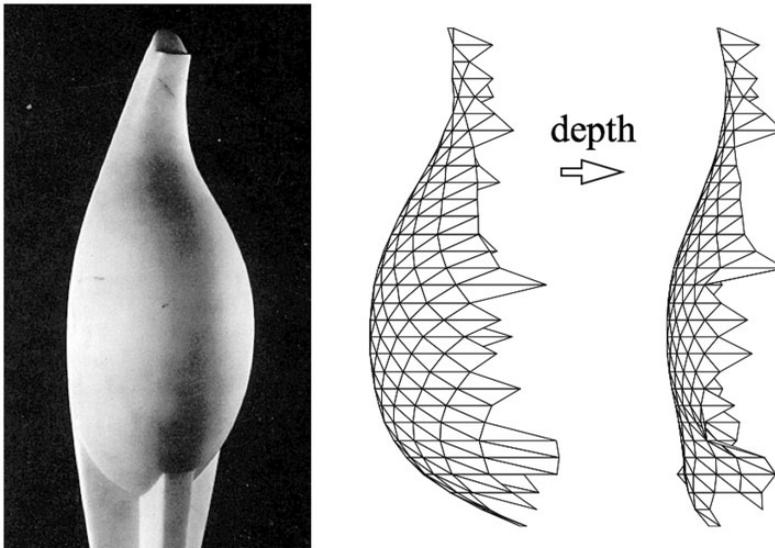
Notice that there is hardly a constraint on the fibers. Thus, the grey tones represent a one-dimensional fiber, the colours a three-dimensional fiber. If you consider qualities like direction, or orientation, the fibers are one-dimensional, but are topological circles. And so forth. The sky is the limit!

Pictorial Space as a Fiber Bundle

An interesting example of such a fiber bundle is “pictorial space” (Koenderink, 2012c), where the fibers represent “depth.” The space of depths has the structure of an affine line. Important cross sections are “pictorial reliefs” (Hildebrand, 1893), which are surfaces in the fiber bundle.

In the past, I have developed a variety of methods to probe the structure of such surfaces (Koenderink, 2012e; Koenderink, van Doorn, Kappers, & Todd, 2001; Koenderink, van Doorn, & Kappers, 1992; Koenderink, van Doorn, & Wagemans, 2018b). One finds that, for a given stimulus, the empirically determined cross sections may vary a lot, certainly between observers, but also for a single observer at different occasions, and so forth (Figures 16 and 17).

However, I find that—of course, within the experimental uncertainty—all these cross sections are related by a group of linear transformations (Koenderink & van Doorn, 2012; Koenderink, van Doorn, & Wagemans, 2011). These transformations involve depth scalings and shears, which appear as a kind of (necessarily non-Euclidian) rotations. Thus, observers experience very similar pictorial reliefs for a given stimulus *modulo* an essentially idiosyncratic transformation of the aforementioned type. They are to be considered *unconstrained hallucinations* that failed to be constrained by psychogenetic check against the contents of the sensorium, simply because such constraints never make it into the optical structure incident on the eye.



5

Figure 16. A stimulus (left) and two responses (center and right) due to two different observers. The response surfaces are seen from the side, so the depth dimensions can be compared. Notice that the responses are qualitatively similar, but differ (mainly) by a depth scaling.

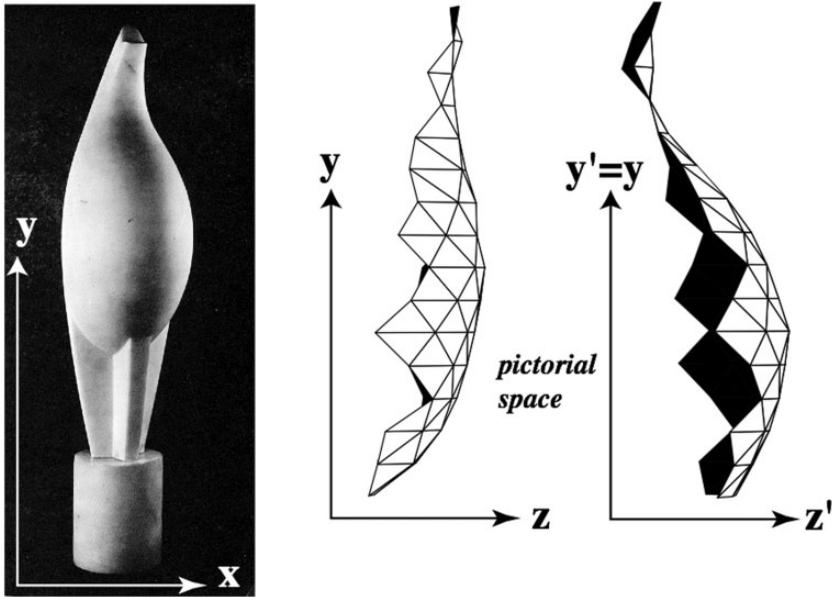


Figure 17. A stimulus (left) and two responses (center and right) due to two different observers. The response surfaces are seen from the side, so the depth dimensions can be compared. Notice that the responses differ (mainly) by a “rotation” about a frontoparallel axis (it is actually a non-Euclidean rotation).

Indeed, from a formal analysis of various “Shape From X” (SFX) algorithms (Koenderink & van Doorn, 1997), one finds that the remaining ambiguities are exactly of that nature. If one considers the group of transformations as the group of “proper movements” or “congruences,” one obtains the formal geometry of pictorial space.

It is a space that is similar to three dimensional Euclidean space, except that it has an “isotropic dimension,” which is the depth dimension. This leads to a full-fledged formalism with very useful quantitative predictive power (Koenderink & van Doorn, 2012; Sachs, 1990; Strubecker, 1942).

This is a rare example of a formal theory that applies to the mental realm—remember that “pictorial depth” need not even have a physical counterpart—yet allows nontrivial, quantitative predictions akin to what one is used to in the setting of classical physics. I consider this encouraging, although I notice that my colleagues in vision research couldn’t care less. This neatly reveals the major cleft between the exact sciences and the humanities. The process of building progress upon progress that drives the former is largely absent in the latter.

Image Transformations

Pictorial space is perhaps the most striking (current) example of the power of the fiber bundle approach in empirical phenomenology. However, there are other examples that have not yet been recognized as such but can easily be framed in the formalism, probably with various advantages.

One example would be the fiber bundle with the space of grey tones as fibers (Figure 18). Its cross sections are monochrome images. In the optical domain, one has radiances, in the presentation “tone” levels. There is a group of transformations that changes the image, yet



Figure 18. Some examples of image transformations. The original is at top left. The various controls are offered under a variety of names in such programs as Adobe Photoshop. (The picture shows Hermann Ludwig Ferdinand von Helmholtz (1821–1894)).

“keeps it the same.” In the old times, photographers deployed such transformations in the darkroom, or people adjusted them using the controls on their monochrome TV sets.

These are transformations that “do not really change” the image, in fact are often hard to detect (Figure 18). I had to use extreme parameter values to demonstrate their effects. In consumer applications such transformations are routinely applied, even without the user’s knowledge, so as to yield acceptable (or even “best possible”) pictures in all cases.

To appreciate this, simply compare “photographic” portraits of women with your visual impressions of the faces of women you are familiar with. “Photoshopping” is similar to a translation into Newspeak. My first impression when seeing such images is that they—in some sense—are “creepy.” They do not look “right,” which—of course—they aren’t. This is apparently a result from “holistic processing” of a kind that I cannot fully describe (that is: understand intuitively). The feeling of creepiness just comes to me. Male portraits are typically less maltreated because all kinds of natural excursions from the mode are considered to reveal “character.” The whole practice is in the worst of taste, yet somehow considered “desirable” in contemporary visual culture. This is even true for the (in my view) victims, possibly because one reckons on the punch of supernormal stimuli (Tinbergen, 1952) and values that.

This is another case in which the fiber bundle formalism might pull many scattered facts together in a common framework.

Heterotopic Areas

It seems to me that biological processes run in local contexts, that is to say, as local as makes sense. Here, “local” does not relate to absolute size. More global interactions tend to be local again, though on a level of summaries.

This makes sense, because there is hardly any “action at a distance” in the natural environment. The few exceptions are the direction of gravity, which has a huge influence on the geometrical structure of the environment, and the illumination direction, say the highly directional sun’s beam or the diffuse beam due to an overcast sky. Of course, in the case of the illumination—different from that of gravity—one already finds numerous local deviations from the global pattern.

Key examples of this need for local processing are the SFX, or “inverse optics” algorithms. It rarely makes sense to run them on the whole field of view. Rare examples include optic flow due to self-motion (Koenderink, 1986). Most optic flows are due to motions of other objects than the self and these tend to be mutually unrelated.

In the case of shape from shading (SFS), one finds that the visual field falls apart in numerous pieces, where in each the light field is very different from that in the next piece (Koenderink, Pont, van Doorn, Kappers, & Todd, 2007).

A good example is Manet’s portrait of Berte Morisot (Figure 19): compare the left and right parts of the face. One side is illuminated by a direct, fairly directional beam, the other by an ambient, diffuse beam. Here, the painter has perfectly succeeded making the fact immediately visible. Just consider the contrasts surrounding the left and right eyes! Well observed and rendered by the painter, it yields the basic tonal distribution for a “true to life” portrait.

Such compartmental structures of environmental light fields are generic. They are due to the fact that the environment is largely made up from opaque objects. This essentially precludes a global application of an SFS algorithm, one would have to segment the image first and treat the segments as distinct environments.

Remarkably, psychogenesis achieves the interpretation on the fly. Whether it actually applies something like SFS remains an open question, of course. I would guess it rather uses the adversarial method and confronts the blackboard structure with its imagery (which is a common way of questioning).

SFX and Ambiguities

In the best of worlds, all SFX algorithms imply groups of ambiguity transformations (Koenderink & van Doorn, 1997). When we just let presentations go by, they may never

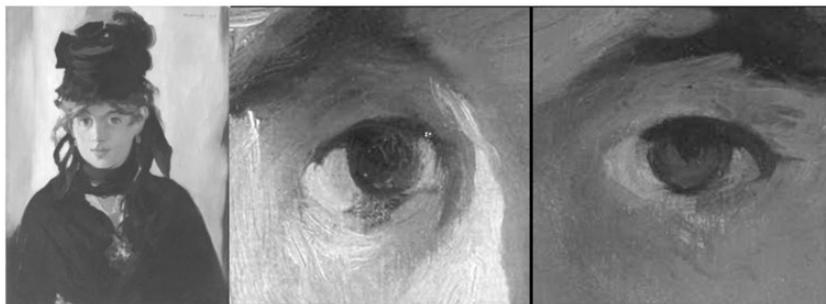


Figure 19. At left, Édouard Manet, *Berthe Morisot au bouquet de violettes* (1872) (part), Musée d’Orsay, oil on canvas (image in public domain). At right, I magnified the regions bordering the eyes. Compare the whites of the eyes and especially the environment of the eye sockets. The eye at the left is surrounded by high contrast light and shade due to a collimated source from the left, whereas the eye at the right is diffusely illuminated. The eyes live in fully distinct luminous atmospheres. No shape-from-shading algorithm will simultaneously apply to both. This is entirely typical for daily life.

Note: Please refer to the online version of the article to view the figures in colour.

need to be resolved, and one happily lives ever after without at any moment noticing ambiguities.

In cases one needs to *act*, there tends to be a need for a unique “solution,” the same holds in laboratory settings where the observer is forced to come up with a unique response. Then, psychogenesis has to select a unique member, for better or worse, from the typically infinite set of solutions left by the ambiguity. In the case of very artificial stimuli, reflective thought may have to step in.

This is a well-known and commonly studied case in experimental psychology. It leads to descriptions in terms of “anchoring” of lightnesses, “regressions” of surface attitudes toward frontoparallelity, “tendency” to sphericity in SFS and so forth. In some cases, large branches of possible solutions never make it into presentations, like negatively curved (hyperbolic) surface forms (Koenderink, 2012g, 2016) in SFS. Some of these may be understood as most convenient templates, others as involving Bayesian priors. No doubt, all these may play a significant role, depending on the context.

Perhaps more important is that there are many cases where it is quite unclear whether a certain SFX algorithm might apply and if so where it would apply. For instance, in the case of SFS, most retinal image gradients are due to physical processes other than those to which the algorithm might apply (Koenderink, 2015b). Moreover, typically such an application would need a prior segmentation, as in a face where one side is illuminated by the sun, the other side by diffuse radiation scattered from the environment.

This practically forces psychogenesis to apply the algorithm (if used at all) only in the phase of testing its imagery against local blackboard structure. As I said before, adversarial methods (“controlled hallucination”) are much more likely. However, the common methodologies of experimental psychology are not aimed at that possibility as I think they should.

Conclusions

In this—rather summarily, abstract and most of all lacunary—account I considered various aspects of the psychogenesis of visual awareness. I failed to specify any algorithms that would enable you to compute the layout of the scene in front of you on the basis of the “optical data.” Indeed, I believe this to be impossible. Moreover, I don’t believe there are such things as “optical data” at all.

Thus, I also do not believe that an eventual outcome might be rated on its “veridicality.” Even worse (depending upon your perspective) I do not think that such a thing as “veridicality” could be given any objective meaning. Biological “truth” is not like common notions of “truth” commonly used in philosophy.

I do believe, however, that vision allows us to deal with our environments in a way that sustains our biological fitness. I would say that this is possible because we live in Umwelts that we *know* how to deal with in terms of guts, blood and bones. Evolution prepared us for it—no doubt accounting for the bulk of our know-how—the rest is our lifetime’s worth of sedimented experience. Notice that I do not imply that we, in the conventional sense, “understand” our environment. From a biological perspective, we don’t need to understand, we need to live.

From the viewpoint of observational phenomenology (the I-word from the dark ages if you want), our Umwelt is the remote horizon of our concrete actuality. No doubt it is very limiting as seen from the perspective of God’s Eye—if there ever were such a thing. But that idea is just as barren as Kant’s notion of the *Ding an Sich*.

Life exploits the regularities it distills from experience, but it cannot count on them. *Anything might happen, anytime*. I mean, things that we could never foresee within the limits of our Umwelt. Call it Black Swan Events (BSEs). We don't happen to live in the trusty casino world, where there always is a fact of the matter, albeit that you don't happen to know it. In real life there need not *be* any fact of the matter. A BSE may strike anytime. Deal with it, or suffer the consequences, that is the biological reality.

Say a coin is tossed, even if you can't see the outcome, you still know the two potential outcomes. But that is a luxury of the casino not offered by nature. Throwing dice or spinning a roulette wheel are very artificial processes. In nature things simply happen, like the alpha decay of a Radium atomic nucleus, or the creation of an electron-positron pair—or something entirely novel intruding your Umwelt.

Shit happening is a *brute fact*, no way to deny that when it happens. Sentient beings had better be ready for that. The touchstone is survival, no matter what happens, anticipated or not. The day the Laws of Physics change, intellectuals are likely to protest, but animals will simply adapt—or die.

The impact that wiped out the dinosaurs was such a (major) BSE. Something entirely novel entered their Umwelt. Maybe God's Eye saw it coming, but the dinos did not and they did not have the time to adapt or the means to cope. Too bad for the dinos, good for us, because finally the early mammals (predominantly small nocturnal insectivores then, it is the reason for the many rods in our peripheral retinas) dared to come out of their burrows to face the light of day.

Thus, BSEs are not necessarily bad things. They are beyond good or evil, being just like events such as the proverbial alpha decay of the radium nucleus. Physics has no problem with things simply happening. One has to accept that what can possibly happen cannot be completely known. All one can do is deal with the present in terms of the past, always trying to learn from mishaps, that is biological reality.

This is essentially Edmund Husserl's view too. Husserl often says that you can only see in terms of what you saw before. He has been attacked for having said that you can never see anything new. Strange? You may very well think so, but I feel he was basically right. We only "see" what we know how to see. History is full of cases where people (for instance in our Western academic world) fairly suddenly gained an eye to some natural phenomenon, whereas prior generations were apparently blind to the obvious (Van den Berg, 1956/1961).

The process of "sedimenting experience" on the basis of a protorational process of association, always goes on. Eventually, we—or our (perhaps remote) offspring—may slowly develop another "functional tone" in von Uexküll's sense. This would perhaps happen many generations after we already "used" that sedimented experience routinely on a daily basis. Most of vision never ends up in visual awareness anyway.

Habits and types (Hume, 1740/1938) are what eventually become functional tones. They are not in any way "defined," but they come to "feel" like something. You can't define "blue" either. In the final instance, all functional tones have to be like that. They are where the buck stops, brute facts, in this case brute feelings. *Qualia* are best understood as brute feelings. They may be like Kant's phantasmatic self-stimulations (Lohmar, 2003), with roots in the physiology. That is probably what von Uexküll hinted at with his—evidently non-scientific, as he very well knew—notion of "functional tone."

Habit and type are the cement of the universe (Hume, 1740/1938). Awareness is the ultimate brute fact. The study of awareness is close to being the study of The World. Our world, anyway.

It would appear a worthwhile endeavour to me.

Acknowledgements

The article was written at the Department of Electrical Engineering and Computer Science, where Jan Koenderink spent a term as a visiting scholar of the Miller Institute for Basic Research in Science, University of California Berkeley.

Reviewers of the article helped to locate it nearer to the mainstream. I thank them all.

The bulk of the work cited here was done in close collaboration with Andrea van Doorn, and I would like to express my gratitude. The only reason she is not a present coauthor here is that this is an “opinion paper” and consequently “all errors are mine.”

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: The work was supported by the program by the Flemish Government (METH/14/02), awarded to Johan Wagemans and by the DFG Collaborative Research Center SFB TRR 135 headed by Karl Gegenfurtner (Justus-Liebig Universität Giessen, Germany). Jan Koenderink was supported by the Alexander von Humboldt Stiftung.

ORCID iD

Jan Koenderink  <https://orcid.org/0000-0002-7623-8336>

Notes

1. As a physicist, I prefer formal theories. “Scale-space” is an example of such an ideal theory.
2. This article reflects my personal opinion, and not all statements are backed up with references, nor will I define common terms. Of course, all shortcomings are mine.
3. Gibson can be read in many ways. On the whole, I tend to be a sympathetic reader. Here, I refer to such statements as (from Gibson, 1979, Chapter 8) “*An affordance is not bestowed upon an object by a need of an observer and his act of perceiving it. The object offers what it does because it is what it is. To be sure, we define what it is in terms of ecological physics instead of physical physics, and it therefore possesses meaning and value to begin with,*” or “*The perception of its affordance should therefore not be confused with the temporary special attraction it may have.*”

References

- Ahissar, E., & Arieli, A. (2001). Figuring space by time. *Neuron*, 32, 185–201.
- Albertazzi, L. (2013). *Handbook of experimental phenomenology: Visual perception of shape, space and appearance*. Chichester, England: Wiley-Blackwell.
- Balas, B., Nakano, L., & Rosenholtz, R. (2009). A summary-statistic representation in peripheral vision explains visual crowding. *Journal of Vision*, 9, 1–18.
- Barlow, H. B. (1972). Single units and sensation: A neuron doctrine for perceptual psychology. *Perception*, 1, 371–394.
- Berkeley, G. (1709). *An essay towards a new theory of vision*. Dublin, Ireland: Printed for Aaron Rhames, at the back of Dick’s Coffee-House for Jeremy Pepyat, Bookseller in Skinner-Row (Numerous modern editions and translations).
- Boddice, R. (Ed.). (2011). *Anthropocentrism: Humans, animals, environments*. Leiden, the Netherlands; Boston, MA: Brill.
- Boime, A. (1993). *The art of the Macchia and the Risorgimento*. Chicago, IL; London, England: The University of Chicago Press.

- Bouma, H. (1970). Interaction effects in parafoveal letter recognition. *Nature*, 226, 177–178.
- Brown, J. W. (1977). *Mind, brain and consciousness*. New York, NY: Academic Press.
- Brown, J. W. (1991). *Self and process*. New York, NY: Springer-Verlag.
- Crick, F. (1994). *The astonishing hypothesis: The scientific search for the soul*. Old Tappan, NJ: Scribner.
- Dirac, P. A. M. (1928). The quantum theory of the electron. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 117, 610–624.
- Dmytryk, E. (1984). *On film editing: An introduction to the art of film construction*. Boston, MA: Focal Press.
- Doyle, A. C. (1887). *A study in scarlet*. London, England: Ward Lock & Co.
- Eccles, J. C. (1994). *How the self controls its brain*. Berlin, Germany: Springer-Verlag.
- Eibl-Eibesfeldt, I. (1970). *Ethology. The biology of behaviour*. New York, NY: Holt, Rinehart and Winston.
- Euclid. (1956[ca. 300BCE]). *The thirteen books of Euclid's elements (translation and commentaries by Heath, T. L.)*. New York, NY: Dover Publications.
- Fahrback, L. (2005). Understanding brute facts. *Synthese*, 145, 449–466.
- Ferraris, M. (2013). Why perception matters. *Phenomenology and Mind*, 4, 36–45.
- Freeman, W. T., & Adelson, E. H. (1991). The design and use of steerable filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13, 891–906.
- Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364, 1211–1221.
- Geem, Z. W., Kim, J. H., & Loganathan, G. V. (2001). A new heuristic optimization algorithm: Harmony search. *Simulation*, 76, 60–68.
- Gibson, J. J. (1951). *The perception of the visual world*. Boston, MA: Houghton Mifflin.
- Gibson, J. J. (1978). The ecological approach to the visual perception of pictures. *Leonardo*, 11, 227–235.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.
- Gregory, R. L. (1997). *Eye and brain*. Princeton, NJ: Princeton University Press.
- Griffin, L. (2019). The atlas structure of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41, 234–245.
- Griffin, L., & Lillholm, M. (2007). Feature category systems for 2nd order local image structure induced by natural image statistics and otherwise. *spie*, 6492, 1–11.
- Grüßer, O. J. (1995). On the history of the ideas of efference copy and reafference. *Clio Medica*, 33, 35–55.
- Gwynne, D. T. (2003). *Encyclopedia of insects*. In V. H. Resh & R. T. Cardé (Eds.). *Ch. Mating mistakes* (pp. 682–686). San Diego, CA: Academic Press.
- Gwynne, D. T., & Rentz, D. C. F. (2003). Beetles on the bottle: Male buprestids make stubbies for females. *Journal of Australian Entomological Society*, 22, 79–80.
- Hartline, H. K., Wagner, H. G., & Ratliff, F. (1956). Inhibition in the eye of limulus. *The Journal of General Psychology*, 39, 651–671.
- Helmholtz, H. L. F. (1867). *Handbuch der physiologischen optik [Handbook of Physiological Optics]*. Leipzig, Germany: Leopold Voss.
- Hess, R. F. (1982). Developmental sensory impairment, amblyopia or tarachopia. *Human Neurobiology*, 1, 17–29.
- Hewitt, C., Bishop, P., & Steiger, R. (1973). A universal modular actor formalism for artificial intelligence. In *IJCAI'73 Proceedings of the 3rd International Joint Conference on Artificial Intelligence* (pp. 235–245). Stanford: IJCAI.
- Hildebrand, A. V. (1893). *Das problem der form in der bildenden Kunst [On the problem of form in the visual arts]*. Strasburg, Germany: Heitz & Münde.
- Hoffman, D. D. (2009). The interface theory of perception. In S. Dickinson, M. Tarr, A. Leonardis, & B. Schiele (Eds.), *Object categorization: Computer and human vision perspectives* (pp. 148–165). New York, NY: Cambridge University Press.

- Hoffman, D. D. (2019). *The case against reality: Why evolution hid the truth from our eyes*. New York, NY: W. W. Norton & Company.
- Hubel, D. H. (1988). *Eye, brain and vision*. New York, NY: W. H. Freeman.
- Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology*, *148*, 574–591.
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Neuroscience*, *195*, 215–243.
- Hume, D. (1738–1740/1975). *A treatise of human nature*. Edited by L. A. Selby-Bigge, 2nd ed. revised by P. H. Nidditch. Clarendon Press, Oxford, 1.1.4.6/12–13.
- Hume, D. (1740/1938). *An abstract of a treatise of human nature* (Reprinted with an Introduction by J. M. Keynes & P. Sraffa). Cambridge University Press, Cambridge, 35.
- Husserl, E. (1893–1917/1964). *The phenomenology of internal time-consciousness* (J. S. Churchill, Trans.). Bloomington, IN: Indiana University Press.
- Huxley, T. H. (1874). On the hypothesis that animals are automata, and its history. *The Fortnightly Review*, *95*, 555–580.
- Jacobs, T. S. (1986). *Light for the artist*. New York, NY: Watson & Guptill.
- Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge, England: Cambridge University Press.
- Kanizsa, G. (1980). *Grammatica del vedere. Saggi su percezione e Gestalt [Grammar of Vision: Perception and Gestalt]*. Il Mulino, NY: Bologna.
- Kepler, J. (1604). *Ad Vitellionem Paralipomena, quibus astronomiae pars optica traditur [Supplement to Witelo, in which is expounded the optical part of astronomy]*. Frankfurt, Germany: Claudium Marnium & Hæredes Ioannis Aubrii.
- Koenderink, J. J. (1982). Different concepts of “ray” in optics: Link between resolving power and radiometry. *American Journal of Physics*, *50*, 1012–1015.
- Koenderink, J. J. (1984a). The structure of images. *Biological Cybernetics*, *50*, 363–370.
- Koenderink, J. J. (1984b). The concept of local sign. In A. J. van Doorn, W. A. van de Grind, & J. J. Koenderink (Eds.), *Limits in perception: Essays in honour of Maarten A. Bouman* (pp. 495–547). Utrecht, the Netherlands: VNU Science Press.
- Koenderink, J. J. (1986). Optic flow. *Vision Research*, *26*, 161–180.
- Koenderink, J. J. (1990). The brain a geometry engine. *Psychological Research*, *52*(2–3), 122–127.
- Koenderink, J. J. (2011). Vision as a user interface. In B. E. Rogowitz & T. N. Pappas (Eds.). *Proceedings of SPIE, human vision and electronic imaging XVI, 7865, 786504*.
- Koenderink, J. J. (2012a). *Visual awareness*. De Clootcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/clootcrans-press>
- Koenderink, J. J. (2012b). *Chrono-geometry*. De Clootcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/clootcrans-press>
- Koenderink, J. J. (2012c). *Pictorial space*. De Clootcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/clootcrans-press>
- Koenderink, J. J. (2012d). *Multiple worlds*. De Clootcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/clootcrans-press>
- Koenderink, J. J. (2012e). *Through the looking glass: On viewing aids*. De Clootcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/clootcrans-press>
- Koenderink, J. J. (2012f). *Experimental phenomenology: Art & science*. De Clootcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/clootcrans-press>
- Koenderink, J. J. (2012g). *Shadows of shape*. De Clootcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/clootcrans-press>
- Koenderink, J. J. (2012h). *Graph spaces*. De Clootcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/clootcrans-press>
- Koenderink, J. J. (2013a). *World, environment, umwelt and innerworld*. De Clootcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/clootcrans-press>
- Koenderink, J. J. (2013b). *Part & whole*. De Clootcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/clootcrans-press>

- Koenderink, J. J. (2014). The all seeing eye? *Perception*, 43(1), 1–6.
- Koenderink, J. J. (2015a). Ontology of the mirrorworld. *Gestalt Theory*, 37, 119–140.
- Koenderink, J. J. (2015b). *Skindeep only*. De Cloutcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/cloutcrans-press>
- Koenderink, J. J. (2016). *About shape*. De Cloutcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/cloutcrans-press>
- Koenderink, J. J. (2017). *Points*. De Cloutcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/cloutcrans-press>
- Koenderink, J. J. (2018). *The way of the eye*. De Cloutcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/cloutcrans-press>
- Koenderink, J. J. (2019). *Sentiencie*. De Cloutcrans Press, Trajectum. Retrieved from <http://www.gestaltrevion.be/en/resources/cloutcrans-press>
- Koenderink, J. J., Pont, S. C., van Doorn, A. J., Kappers, A. M. L., & Todd, J. T. (2007). The visual light field. *Perception*, 36, 1595–1610.
- Koenderink, J. J., & Richards, W. (1988). Two-dimensional curvature operators. *Journal of the Optical Society of America. A*, 5, 1136–1141.
- Koenderink, J. J., Richards, W., & van Doorn, A. J. (2012a). Blow-up: A free lunch? *i-Perception*, 3, 141–145.
- Koenderink, J. J., Richards, W., & van Doorn, A. J. (2012b). Space-time disarray and visual awareness. *i-Perception*, 3, 159–165.
- Koenderink, J. J., Valsecchi, M., van Doorn, A. J., Wagemans, J., & Gegenfurtner, K. (2017a). Eidolons: Novel stimuli for vision research. *Journal of Vision*, 17, 1–36.
- Koenderink, J. J., & van Doorn, A. J. (1991). Affine structure from motion. *Journal of the Optical Society of America A*, 8, 377–385.
- Koenderink, J. J., & van Doorn, A. J. (1997). The generic bilinear calibration-estimation problem. *International Journal of Computer Vision*, 23, 217–234.
- Koenderink, J. J., & van Doorn, A. J. (2012). Gauge fields in pictorial space. *SIAM Journal of Imaging Science*, 5, 1213–1233.
- Koenderink, J., & van Doorn, A. J. (2017). The planispheric optic array. *i-Perception*, 8(3), 1–17.
- Koenderink, J. J., van Doorn, A. J., Kappers, A. M. L., & Todd, J. T. (2001). Ambiguity and the ‘mental eye’ in pictorial relief. *Perception*, 30, 431–448.
- Koenderink, J. J., van Doorn, A. J., & Pinna, B. (2015). Psychogenesis of Gestalt. *Gestalt Theory*, 37(8), 287–304.
- Koenderink, J. J., van Doorn, A. J., & Pinna, B. (2017b). Plerosis and atomic Gestalts. *Gestalt Theory*, 39, 30–53.
- Koenderink, J. J., van Doorn, A. J., Pinna, B., & Wagemans, J. (2016). Boundaries, transitions and passages. *Art & Perception*, 4, 185–204.
- Koenderink, J. J., van Doorn, de Ridder, H., & Oomes, A. H. J. (2010). Visual rays are parallel. *Perception*, 39, 1163–1171.
- Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. L. (1992). Surface perception in pictures. *Perception & Psychophysics*, 52, 487–496.
- Koenderink, J. J., van Doorn, A. J., & Pinna, B. (2018a). Measures of Prägnanz? *Gestalt Theory*, 40, 7–28.
- Koenderink, J. J., van Doorn, A. J., & Wagemans, J. (2018b). Geometry of pictorial relief. *Annual Review of Vision Science*, 4, 451–474.
- Koenderink, J. J., van Doorn, A. J., & Todd, J. T. (2009). Wide distribution of external local sign in the normal population. *Psychological Research*, 73, 14–22.
- Koenderink, J. J., van Doorn, A. J., & Wagemans, J. (2011). Depth. *i-Perception*, 2, 541–564.
- Koenderink, J. J., van Doorn, A. J., & Wagemans, J. (2017c). Trelliswork and craquelure. *i-Perception*, 8(5), 1–5.
- Koenderink, J. J., van Doorn, A. J., & Wagemans, J. (2018c). Magic circle. *i-Perception*, 9(3), 1–26.

- Lindeberg, T. (1994a). Scale-space theory: A basic tool for analysing structures at different scales. *Journal of Applied Statistics (Supplement on Advances in Applied Statistics: Statistics and Images 2)*, 21, 224–270.
- Lindeberg, T. (1994b). *Scale-space theory in computer vision*. Dordrecht, the Netherlands: Kluwer Academic Publishers.
- Lissauer, H. (1890). Ein fall von seelenblindheit nebst einem beitrage zur theorie derselben. *Archiv Für Psychiatrie Und Nervenkrankheiten [A case of soul-blindness and a contribution to its theory]*, 21, 222–270.
- Loftus, E. F., & Davis, D. (2006). Recovered memories. *Annual Review of Clinical Psychology*, 2, 469–498.
- Lohmar, D. (1993). Wahrnehmung als Zusammenspiel von Schematisierung und figürlicher Synthesis: Überlegungen zur Leistung der Einbildungskraft bei Kant [Perception as a mixture of schematizing and figural synthesis: On imagery with Kant]. *Tijdschrift Voor Filosofie*, 55, 100–129.
- Lohmar, D. (2003). Über phantasmatische Selbstaffektion in der typisierenden Apperzeption und im inneren Zeitbewusstsein [On phantasmatic self-stimulation in typifying apperception and inner time consciousness]. In Leitmotiv 3 (Ed.), (pp. 67–80). Retrieved from www.ledonline.it/leitmotiv
- Lohmar, D. (2014). Types and habits. Habits and their cognitive background in Hume and Husserl. *Phenomenology and Mind*, 6, 40–51.
- Lorenz, K. (1973). *Die Rückseite des Spiegels [The backside of the mirror]*. München, Germany: Piper.
- Lotze, R. H. (1852). *Medicinische psychologie, oder physiologie der seele [Medical Psychology, or the Physiology of the Soul]*. Leipzig, Germany: Weidmann'sche Buchhandlung.
- Mack, A., & Rock, I. (1998). *Inattentional blindness*. Cambridge, MA: M.I.T. Press.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York, NY: Freeman.
- Maxwell, J. C. (1855). Experiments on colour, as perceived by the eye, with remarks on colour-blindness. *Transactions of the Royal Society of Edinburgh*, XXI, 275–298.
- McCulloch, W. S., & Pitts, W. H. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 7, 115–133.
- McNally, J. (2008). *The moment it clicks: Photography secrets from one of the world's top shooters*. Berkeley, CA: New Riders.
- Meinong, A. (1899). Ueber gegenstände höherer ordnung und deren verhältniss zur inneren wahrnehmung [On higher order objects and their relation to inner perception]. *Zeitschrift Für Psychologie Und Physiologie Der Sinnesorgane*, 21, 187–272.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81–97.
- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83, 435–450.
- Nietzsche, F. (1966/1886). *Jenseits von Gut und Böse [Beyond Good and Evil] (Walter Kaufmann, Trans.)* New York, NY: Vintage Books.
- Orwell, G. (1949). *Nineteen eighty-four*. London, England: Secker and Warburg.
- Pashler, H. (1998). *The psychology of attention*. Cambridge, MA: M.I.T. Press.
- Phillips, F., & Voshell, M. G. (2009). Distortions of posterior visual space. *Perception*, 38, 1045–1052.
- Pinna, B., Brelstaff, G., & Spillmann, L. (2001). Surface color from boundaries: A new 'watercolor' illusion. *Vision Research*, 41, 2669–2676.
- Pont, S. C., Nefs, H. T., van Doorn, A. J., Wijntjes, M. W. A., te Pas, S. F., de Ridder, H., Koenderink, J. J. (2012). Depth in box spaces. *Seeing and Perceiving*, 25(3–4), 339–349.
- Richards, W. (2018). *Modal worlds: The role of regularities in evolutionary biology*. De Clootcrans Press, Trajectum. Retrieved from <http://www.gestaltrevision.be/pdfs/koenderink/ModalWorlds.pdf>
- Richards, W. A. (1982). *How to play twenty questions with nature and win win*. MIT AI Memo 660. Cambridge, MA.
- Riedl, R. (1975). *Die Ordnung des Lebendigen. Systembedingungen der Evolution [Order in living organisms: A systems analysis of evolution]*. Hamburg, Germany: Paul Parey.

- Riedl, R. (1984). *Biology of knowledge: The evolutionary basis of reason*. Chichester, England: John Wiley & Sons.
- Sachs, H. (1990). *Isotrope geometrie des raumes [Isotropic geometry of space]*. Braunschweig/Wiesbaden, Germany: Vieweg.
- Saladin, K. (2012). *Anatomy and physiology*. New York, NY: McGraw Hill.
- Schrödinger, E. (1944). *What is life?* Cambridge, England: Cambridge University Press.
- Schrödinger, E. (1958). *Mind and matter (Tarner lectures)*. Cambridge, England: Cambridge University Press.
- Seifert, H. (1933). Topologie dreidimensionaler gefaserner Räume [Topology of threedimensional fiber spaces]. *Acta Mathematica*, 60, 147–238.
- Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(379–423), 623–656.
- Shapiro, A. G., & Todorovic, D. (Eds.). (2017). *The Oxford compendium of visual illusions*. New York, NY: Oxford University Press.
- Skinner, B. F. (1976). *About behaviourism*. New York, NY: Random House.
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, 10, 89–96.
- Strubecker, K. (1942). Differentialgeometrie des isotropen Raumes III: Flächentheorie [Differential geometry of isotropic space III: Theory of surfaces]. *Mathematische Zeitschrift*, 48, 369–427.
- Szeliski, R. (2010). *Computer vision: Algorithms and applications*. New York, NY: Springer.
- ter Haar Romeny, B. M. (2003). *Front-end vision and multi-scale image analysis*. New York, NY: Springer.
- ter Haar Romeny, B. M. (2008). *Front-end vision and multi-scale image analysis*. New York, NY: Springer.
- Tinbergen, N. (1951). *The study of instinct*. Oxford, England: Oxford University Press.
- Tinbergen, N. (1952). The curious behaviour of the stickleback. *Scientific American*, 187, 22–26.
- Titchener, E. B. (1927). *Experimental psychology, a manual of laboratory practice, volume I*. London, England: The MacMillan Company.
- Uexküll, J. V. (1909). *Umwelt und innenwelt der tiere [Umwelt and inner world of animals]*. Berlin, Germany: Springer.
- Uexküll, J. V. (1920). *Theoretische biologie [Theoretical biology]*. Berlin, Germany: Springer.
- Van den Berg, J. H. (1956/1961). *Metabletica. Callenbach, Nijkerk, The Netherlands [The changing nature of man]*. New York, NY: Dell.
- Vicario, G. B. (1973). *Tempo psicologico ed eventi [Psychological time and events]*. Firenze, Italy: Giunti Barbera.
- von Holst, E., & Mittelstaedt, H. (1950). Das Reafferenzprinzip: Wechselwirkungen zwischen zentralnervensystem und Peripherie [The Refference Principle: Interactions between the central and peripheral nervous system]. *Naturwissenschaften*, 37, 464–476.
- Walsworth, M. T. (1882). *Twenty questions: A short treatise on the game to which are added a code of rules*. New York, NY: Henry Holt and Company.
- Wassersug, J. D., & Wassersug, R. J. (1986). Fitness fallacies. *Natural History*, 3, 34–37.
- Wilkinson, S. (2014). Accounting for the phenomenology and varieties of auditory verbal hallucination within a predictive processing framework. *Consciousness and Cognition*, 30, 142–155.
- Zeki, S. (2003). The disunity of consciousness. *Trends in Cognitive Neuroscience*, 7, 214–218.
- Zhixing, J., & Bir, B. (2015). Analysis-by-synthesis. *Journal of Computer Vision and Image Understanding*, 134, 48–63.