Relating Simple and Complex Visual Hallucinations in Charles Bonnet Syndrome

Matthew Caldwell
CoMPLEX, University College London

Supervisors: Richard Abadi, Richard Clement

4787 Words

April 15, 2007

Contents

1 Charles Bonnet Syndrome

1.1 Vision and its Discontents

The human visual system is astonishingly adept at translating light signals received from the external world into internal models of that world. The processes involved are not fully understood, but it is reasonably certain they combine extensive direct analysis of the inputs—termed the *bottom-up* approach—with bringing to bear a vast body of accrued knowledge of how the world both appears and behaves, in order to make plausible inferences as to the *causes* of the signals—known as *top-down* processing.

The consistent accuracy and resolution of vision make it an extremely powerful tool for negotiating one's environment, and for most (sighted) people in most (adequately lit) circumstances it is by far the dominant sensory modality. There is usually a close and dependable correspondence between what is *seen* and what is actually *there*.

Usually—but not *always*.

Cases where there is disagreement between the world as visually perceived and its actual material reality are important for several reasons. Clinically, they may be symptomatic of pathological physical and psychological conditions requiring attention; practically, they may present day-to-day obstacles for those experiencing them that need to be worked around in some way. But beyond that, they are also interesting in the abstract because of what they reveal about the mechanics of vision itself.

It is usual to distinguish between *illusions*, where some entity genuinely present in the received image is incorrectly or ambiguously interpreted by the viewer, and *hallucinations*, where the viewer sees something for which there is no direct visual cue at all. We shall focus on the latter here, although the boundary is far from clear-cut.

Visual hallucinations are experienced in a variety of contexts. They are frequently associated with some form of cognitive impairment, such as dementia or schizophrenia, or with altered mental states induced by psychedelic drugs, and they occur surprisingly often at the boundaries between sleeping and waking. They may also be associated with eye damage and vision loss—a condition known as Charles Bonnet Syndrome.

1.2 Aetiology & Phenomenology

Charles Bonnet Syndrome (CBS) was originally identified in 1936 by de Morsier, named after an 18th Century Swiss naturalist who documented similar symptoms in his grandfather and went on to suffer them himself in later life [14]. It is defined by visual hallucinations in the presence of material loss of visual acuity—usually in the central visual field—and the absence of cognitive impairment.[1] Thus, the hallucinations are taken to be a consequence of the visual damage rather than any mental disturbance. [11, 17, 16, 1, 20, 21, 19, 18]

The syndrome can occur with many different underlying ophthalmic pathologies, including macular degeneration, cataracts, glaucoma, diabetic retinopathy and physical injury to the eyes (usually where damage is binocular and overlapping); the specific cause does not appear to have any bearing on the nature of the hallucinations. Many of these conditions are associated with

[1] There have been a number of variations in definition over the years, to the extent that some have argued the term is of little diagnostic value [7], but it remains in use and represents a useful classification, if only for the distinction drawn between hallucination and psychosis.

ageing, and CBS is indeed most common in the elderly. Although sufferers are (by definition) not mad, they will often be vulnerable, lonely and under-occupied.

Hallucinations tend to occur in familiar surroundings and conditions of low arousal. The subject is usually aware at the time that the things they are seeing are not real, or else becomes so soon afterwards. While the actual experience can be neutral or even pleasurable, people are frequently upset by the *fact* of hallucinating, fearing they are losing their minds.

The actual content of the hallucinations, in CBS and most other hallucinatory conditions, falls into two broad categories:

- **Simple hallucinations** include flashes, lines, abstract shapes, swirls and repeating patterns like grids or tiles. They are usually uninterpreted and move with the eyes rather than with the scene.
- **Complex hallucinations** are coherent objects, people, animals, buildings or entire landscapes, which frequently integrate with the observed environment, maintaining their position within it as the eyes move.

The differences and relationships between these two types is the main subject of this essay, and will be addressed in more detail below.

In most CBS cases, hallucinations do not continue indefinitely but disappear after a year or 18 months. They will also often stop when, or shortly after, the patient's eyesight deteriorates to the point of total blindness, or when it is substantially restored by medical treatment.

Estimates of the prevalence of CBS vary enormously, and it is believed to be significantly under-reported. Many people are reluctant to volunteer the information that they are seeing things that aren't there—and not without justification: clinicians are often under-informed about the syndrome and may indeed diagnose dementia or other mental illnesses when confronted with its symptoms.

2 Simple Visual Hallucinations

2.1 Commonalities of Experience: Klüver's Form Constants

A feature of hallucinations in general—irrespective of CBS—is that they cannot be measured objectively. The only way to explore their nature is to *ask* those experiencing them. Since such people are often in a psychologically or pharmacologically abnormal state, their evidence may be unreliable. Nevertheless, there is a good deal of common ground between subjective accounts of simple visual hallucinations from many different conditions, which must be construed as evidence in their favour.

People experiencing this kind of hallucination as a result of psychotropic drugs, psychological problems, artificial sensory deprivation, organic visual damage or even transient stimuli such as pressure on the eyeballs, frequently describe (and in some cases *draw*) their visions in similar terms: gratings, lattices, fretworks, honeycombs, chequerboards, cobwebs, tunnels, funnels, cones and spirals. While some of this consistency could be due to the limitations of verbal description and shared referents, it is reasonable to suppose that the accounts are informed by a significant commonality of experience.

Drawing on such accounts, Klüver determined a set of recurring patterns that he called *form constants*, examples of which are shown in figure 1 [15]. He proposed that these form con-

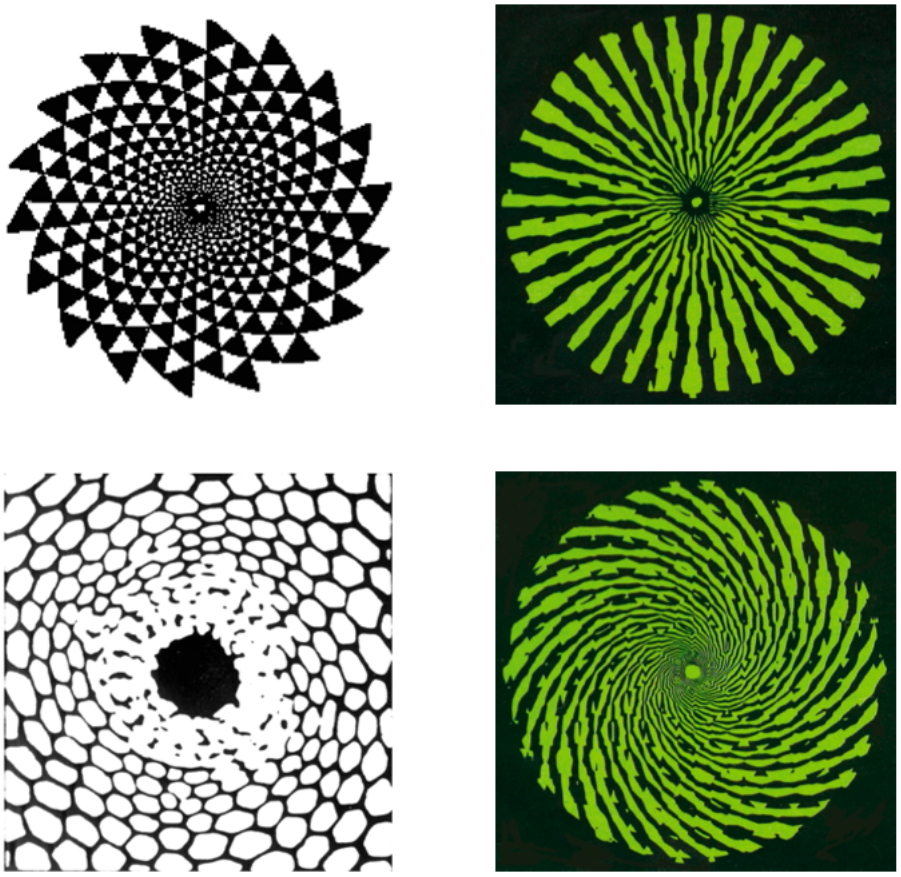

Figure 1: Examples of Klüver's 'form constants' [3]

stants represented consistent patterns of hallucination characteristic of the visual perceptive system.

Klüver's form constants informed a dynasty of hallucination models, begun by Ermentrout and Cowan in the 1970s [10] and more recently refined by Bressloff et al [3, 4]. These models are founded on two ideas: that simple hallucinations derive from internal patterns of neural activity in the visual cortex; and that their subjective appearance is a result of an inverse geometric transform from cortical space to retinal.

2.2 The Retinotopic Transform

The light-sensitive cells in the back of the eye send signals along the optic nerve, via the lateral geniculate nucleus (LGN), to the neural networks of the primary visual cortex (V1). This mapping is spatially coherent, with one significant exception: the signals from each eye are divided along the vertical meridian, with the temporal hemifields of the retina connected to the *ipsilateral* LGN and V1, while the nasal hemifields are connected to the *contralateral* equivalents. Thus, bearing in mind that the retinal image is inverted, the left visual cortex receives images of the right visual field from both eyes, and vice versa.

Although the mapping from retina to cortex preserves locality, it does not preserve angles or areas. The spatial organisation of retinal sensors is essentially polar, and their density is heavily weighted towards the central *fovea*, dropping off with distance in an inverse-square

relationship:[2]

$$\rho = \frac{1}{(w_0 + \varepsilon r)^2}$$

By contrast, cells in the primary cortex are arranged relatively uniformly, with a more or less regular repeating pattern of *hypercolumns* corresponding to particular locations within the visual field. For the purposes of the models discussed here, the arrangement is considered in Cartesian terms, with a mapping between the two geometries $\{r, \theta\} \mapsto \{x, y\}$ of the form

$$x = \frac{\alpha}{\varepsilon} \ln \left[1 + \frac{\varepsilon}{w_0} r \right]$$

$$y = \frac{\beta r \theta}{w_0 + \varepsilon r}$$

for scaling constants α and β. The true mapping is less regular (see figure 2), but the key features are comparable. [3, 2]

This *retinotopy* would be irrelevant were the visual cortex to respond only and always to external stimuli from the eyes: the interpretative processes of the brain 'know' where such signals originate regardless of their subsequent spatial rearrangements, and our perceptions implicitly invert the transform. But where neural activity arises from some other cause—dependent, say, on the spatial arrangement of cells in the cortex—that must also be interpreted according to the same rules, the same retinotopic inversion.

Many of Klüver's form constants exhibit rotational symmetry and decreasing scale towards the centre of vision, corresponding to regular translationally-symmetrical patterns under the retinotopic transform. Thus, if the processes of neural interactions in V1 could give rise to such periodic patterns, we would expect those patterns to be experienced in just the ways that people with simple hallucinations describe.[3]

2.3 Neural Networks & Pattern Formation

V1, like other areas of the cortex, comprises a dense sheet of highly interconnected neurons. Individual neurons may be *excitatory* or *inhibitory*, according to the neurotransmitters they release at their synapses on firing. Their own activity is modulated by excitatory and inhibitory signals received from many other upstream neurons, and they in turn contribute to the modulation of the neurons to which they are connected downstream.

Although discrete models can be used for small networks of neurons, larger scale behaviour is usually represented by continuum models, in which the network is considered as a continuous space, with distance an analogue of connectivity. One or more activity variables—say, neural firing rates—are defined over this space and evolve under the influence of integro-differential equations representing neural interactions.

[2] The weighting constants w_0 and ε must be determined experimentally; their values are not important here.

[3] Bressloff's version of this model takes a more sophisticated approach to the retinotopic mapping that also accounts for the orientation-detecting features of V1, noting that the perceived *edges* in simple hallucinations are typically aligned in a consistent fashion. Conveniently, such consistency mostly corresponds to the stable planforms predicted by the pattern formation part of the model.

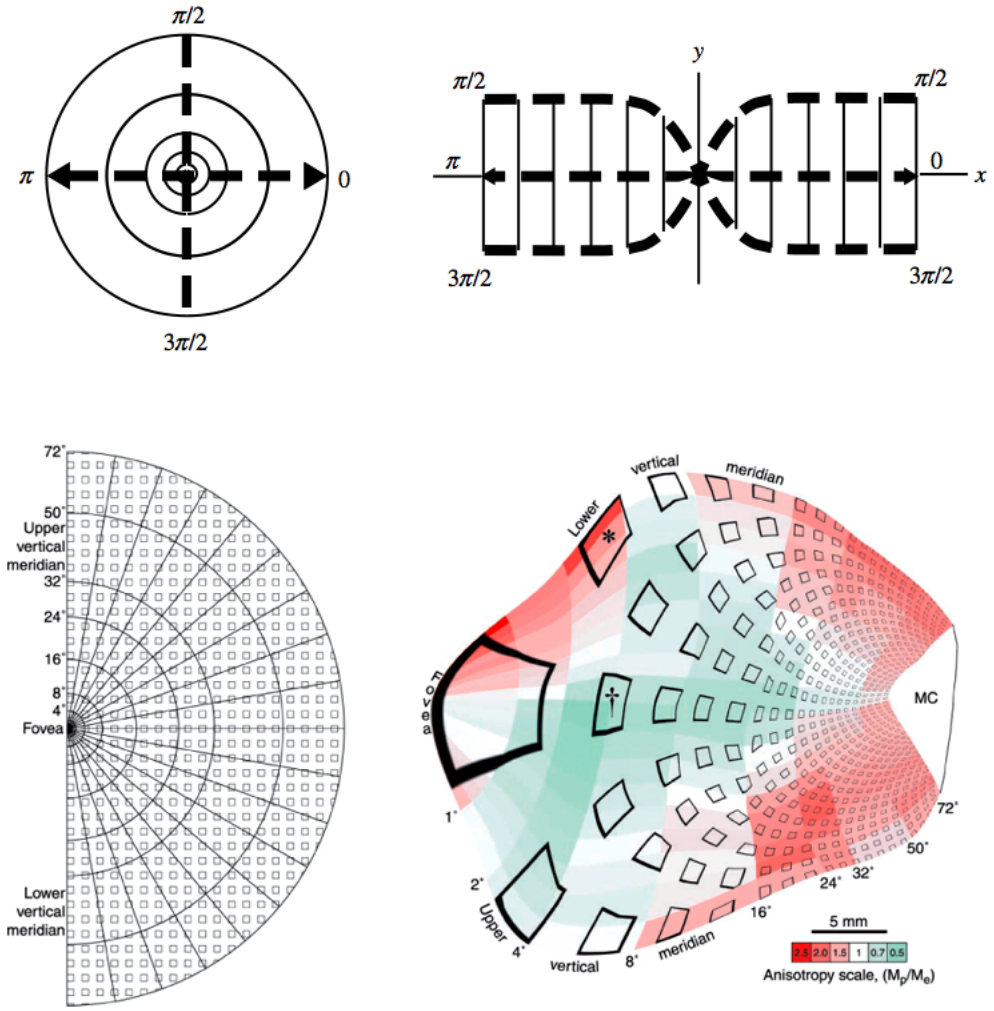

Figure 2: The idealised retinotopic transform, from [3]; and a more physiologically accurate mapping from experiments on the macaque [2]. The latter shows the projection of a grid of squares from one retinal hemifield onto V1.

The governing equations for such models derive from Wilson & Cowan [22, 23], though there are substantial variations between applications of their model (see Appendix A). Although formulated rather differently, they are cognate with reaction-diffusion systems in which the interplay of activation and inhibition on different length scales can lead to a breaking of symmetry and the formation of periodic patterns—which is precisely what is proposed to occur in the visual cortex to generate simple hallucinations.

Linear stability analysis of the (adapted) Wilson-Cowan equations, when restricted to doubly-periodic solutions,[4] gives rise to a finite set of stable pattern types or *planforms* corresponding to the different doubly-periodic symmetry groups for the plane (hexagonal, square and rhombic). When transformed back into the visual field, these do indeed bear a close resemblance to the Klüver form constants (figure 3).

[4]That is, solutions that are periodic in two different directions, constituting a regular tiling of the plane. This restriction is essentially a pragmatic one to make the solution set finite, but there is evidence from other fields that physical systems favour patterns of this kind.

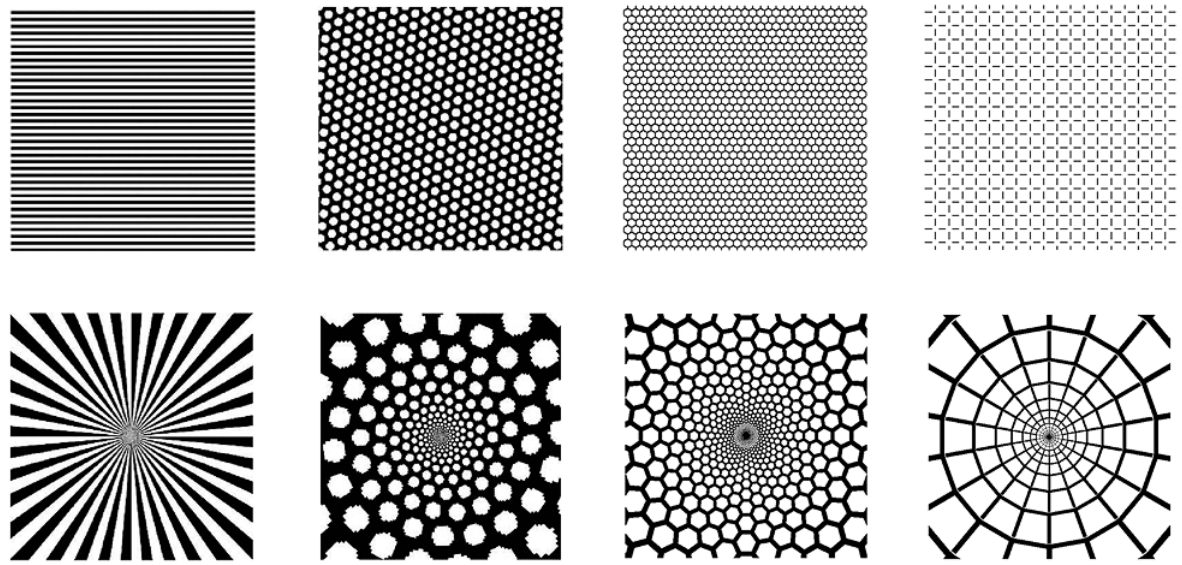

Figure 3: Stable planforms of different symmetries, in the idealised space of V1 (upper row) and after inverse transformation back into retinal space (lower row) [4]

2.4 Deafferentation & Excitability

The Ermentrout-Cowan and Bressloff models demonstrate a possible mathematical mechanism by which simple hallucinations could arise, and propose that the underlying physical causes of this phenomenon are changes in the relative activity of the excitatory and inhibitory neuron populations in V1. This is consistent with what is known of the pharmacology of hallucinogenic drugs—given the abstractions of the models this is necessarily a qualitative rather than quantitative observation—and with neurochemical disturbances in some psychiatric conditions, but it does not immediately explain the similar hallucinations in Charles Bonnet Syndrome.

Burke, drawing on his own experiences of CBS, and ffytche—neither or whom directly reference these pattern formation models[5]—argue that changes in neural excitability are a natural consequence of *deafferentation*—that is, deprivation of normal upstream stimuli. [5, 13, 12]

Neurons exhibit a considerable degree of self-governance, and can tune their receptiveness to particular neurotransmitters as their activities and environment vary. This includes internalising receptors to reduce responsiveness when over-stimulated and increasing the number of receptors when stimulation diminishes, as well as numerous other changes in local biochemistry. In the case of deafferentation, this internal 'gain control' can lead to hyperexcitability. If this occurs over a large enough population it could be sufficient to lead to pattern formation.

Although the specific kinds of stable pattern postulated by the hallucination models seem not to have been observed in the laboratory, isolated patches of cortical tissue do exhibit waves of spontaneous activity in vitro. It is not too far-fetched to suppose that stable patterns might arise in the right circumstances.

[5]Burke's own model, in which his simple hallucinations are taken as direct perceptions of the cortical topography itself, is trivially unconvincing, but that doesn't discount the rest of his argument.

3 Complex Visual Hallucinations

3.1 Higher Visual Processing

Whereas simple hallucinations can plausibly be located in the lower levels of visual processing, complex hallucinations are much harder to pin down. Their contextualised nature implies the involvement of higher faculties such as memory, expectation, attention and even consciousness itself, for which our models are *at best* extremely nebulous. The kind of quantitative grounding in actual physical neural arrangement and activity that was used for V1 no longer applies—indeed, the retinotopic relationship that makes it possible to point to a group of neurons and say with relative confidence what they *do* is quite atypical for the brain, and quickly degenerates as the visual signals propagate further into the perceptual pipeline.

It can be inferred from the experiences of trauma patients and from experiments with imaging techniques such as fMRI that different aspects of the analysis and organisation of partially pre-processed visual inputs are performed by different areas of the brain. Identification of colours occurs in one place, surfaces in a second, motion in a third, etc. Some of these areas appear to be very specialised: a significant area is devoted to the recognition of faces, evidently a matter of great adaptive importance in our highly social species.

All of these analyses are understood to take place in parallel, potentially producing multiple candidate interpretations at every locus of the scene, which must then be integrated into an overall model of the external world by higher-level processes. This integration will obviously involve selection and arbitration, according to criteria drawn from memories and knowledge of the context, but it is further hypothesised to include significant feedback circuits, modulating the behaviour of the lower level analytic subsystems by adding its own inhibitory and excitatory inputs. Such top-down control will thus constrain and direct what the lower levels produce, until the whole system converges on an approximate consensus.

As our everyday experience confirms, this process is both very fast and amazingly accurate across a wide range of operating conditions, but clearly it *can* go wrong. Complex visual hallucinations would seem to be a prime example: the scene model is successfully constructed, but some of the elements selected are incorrect. How and where such mistakes occur is a matter for much speculation and ongoing research.

3.2 Release Phenomena

One popular characterisation of complex hallucinations has been as *release phenomena*. The concept is simple: the integrative processes that build the scene are fed from one direction by a constant stream of candidate building blocks, and from another by memories, knowledge and opinion. In normal operation there is a balance between these sources of information, with the bottom-up inputs exerting an inhibitory pressure on the higher-level cognitive ones. There is thus a feed-forward control system keeping the mind's more fanciful flourishes in check.

According to this model, when the bottom-up inputs falter for some reason—say, a loss of external stimuli due to sensory or neurochemical impairment—there is a corresponding disinhibition of memories and the imagination, which then rush in to fill the gap. [6]

While this explanation is intuitively appealing—and has some correspondence with the simple hallucination model in identifying deafferentation as the proximate cause of hallucinations—it is not clear that it actually explains very much. It is already established that the visual processing pipeline is long and convoluted. If loss of input leads to spurious output, shouldn't those outputs accumulate rather than preserving the loss all the way to the end? Or if there

are some secondary factors that affect whether or not spurious outputs are generated, what are they?

Moreover, by locating the fault so specifically within the process of integration—which is itself essentially mysterious—this model could be argued to abdicate the responsibility of explanation altogether, since there is little demonstrable consequence of it being true.

Nevertheless, if spontaneous patterns can form—and have perceptual consequences—in the lower levels of visual processing, it is likely that they might in other compartments too, and we would expect them then to be interpreted in some way.

3.3 The Perception and Attention Deficit Model

A recent alternative to the release explanation is the Perception and Attention Deficit (PAD) model of Collerton et al, which attempts—perhaps recklessly—to encompass all forms of recurrent complex visual hallucination within a single explanatory framework [8]. Observing that reports of complex hallucinations across a variety of conditions are usually associated with low levels of arousal, the researchers light on *attention* as a crucial ingredient in the hallucinatory—and by implication, the perceptual—process.

Now, attention is a slippery character in psychology and neuroscience, with a number of different but overlapping meanings, none of them especially concrete. Exactly how the everyday notion of 'paying attention' meshes with its technical application will vary, often in apparently unacknowledged ways, among users and circumstances. Invoking it in the context of complex hallucinations might therefore seem—and in some senses actually *be*—just as much an intellectual sleight of hand as the idea of release phenomena.

However, the value of the PAD model is in its identification of the locus of complex hallucination as being the *interaction* between bottom-up and top-down processing. Specifically, the model proposes that a hallucination occurs when an erroneous low-level candidate—Collerton et al use the term *proto-object*—is wrongly incorporated by an inadequately-attended high-level integration process. This misidentification is then reinforced by the usual top-down feedback mechanisms, leading to systematic convergence on an incorrect scene description—that is, seeing something that isn't there.

An advantage of this model is that it is consistent with that for simple hallucinations: the low level patterns already propagate up as false percepts, and all it takes is an attentive failure to lead to them being interpreted as scene constituents and thereby transmuted into complex hallucinations. This could explain why CBS complex hallucination figures are often seen as clad in bizarrely-patterned costumes: it's a cognitively *inexpensive* way to integrate those abstract patterns into the scene model.

Furthermore, the flexibility of the intersection is compelling in accounting for the phenomenological differences between hallucinatory conditions: in schizophrenia, for example, we might expect the errors to be much more heavily weighted toward the top-down side of things than in CBS, where the primary pathology is very clearly bottom-up; and this is borne out by the predominance of complex hallucinations in the former condition and simple ones in the latter.

There remains, of course, the troubling insistence in the PAD model on the ill-defined concept of attention as the sole metric of top-down function; a problem that the authors, to my mind, compound by placing unjustified emphasis on particular cortical pathways and neurotransmitters that do not necessarily accord with what is observed in conditions outside their main area

of research. A number of commentators bring them up on this point—quite rightly—but it is, surely, a distraction.

3.4 Problems of Quantification

Despite its intuitive appeal, the PAD model—certainly in the abstracted form described, or perhaps more honestly *caricatured*, here—subsists more in hand-waving metaphors than in concrete predictions about the functions of the brain. Indeed, this vagueness may contribute to its appeal, since it leaves a good deal of 'wiggle room' for reinterpretation within the readers' own frames of reference.

It would be helpful for both clarity and testability if the model could be recast in more mathematical terms, and indeed several of the open peer commentaries on the article at least gesture in that direction. PAD's conceptual units can readily be analogised to those from statistical inference, signal detection, control theory, etc. However, such suggestions make for an uneasy fit, precisely because they leap to quantify the unmeasurable.

While we can be confident that proto-objects, memories, beliefs, scene descriptions, expectations, attention and so on are embodied in the patterns of connectivity and neurochemical processes of the brain, the *manner* of their encoding remains a mystery. The interactions between these presumptive components are, in consequence, similarly mysterious. In the absence of a convincing physical characterisation of such entities, attempts at quantification are necessarily arbitrary and *ad hoc*. Simplifications and elisions are inevitable. The degree to which they are justified and useful depends on what they contribute to our understanding.

The signal detection argument of Dolgov & McBeath [9] is a case in point. They suggest that hallucinations need to be taken in the context of overall perceptual *value*: as false positives that may be compensated by a decrease in false negatives (figure 4).

Clearly, characterising the visual perceptual process in terms of *signal* and *noise* elides the multidimensional specificities of each; the mathematical space in which the plethora of candidate proto-objects intersect the myriad possible interpretations along a single line would be strange indeed. Nevertheless, the *implications* of such a representation are persuasive and enlightening. Seeing things that aren't there could very well be less disadvantageous than *not seeing* things (for example, predators) that *are*. In the situation of deteriorating vision it may therefore be beneficial to lower the threshold of proto-object acceptance: the apparition of non-existent fairies is a small price to pay for the continued awareness of actual tigers. That any quantification of this threshold is entirely spurious is really by the by.

4 Discussion

4.1 Integrating Models in the Context of CBS

Apart from Burke & ffytche, the hallucination models described above do not address Charles Bonnet Syndrome specifically. Nevertheless, they can be taken together to provide a persuasive explanation for the condition: attenuated vision leads to pattern formation in the lower visual processing areas, and these patterns feed forward through the perceptual pathways to be either perceived directly as uninterpreted artefacts or incorporated into the scene as objects with maximum available congruity, depending on higher level selective thresholds that are tentatively attributed to attention.

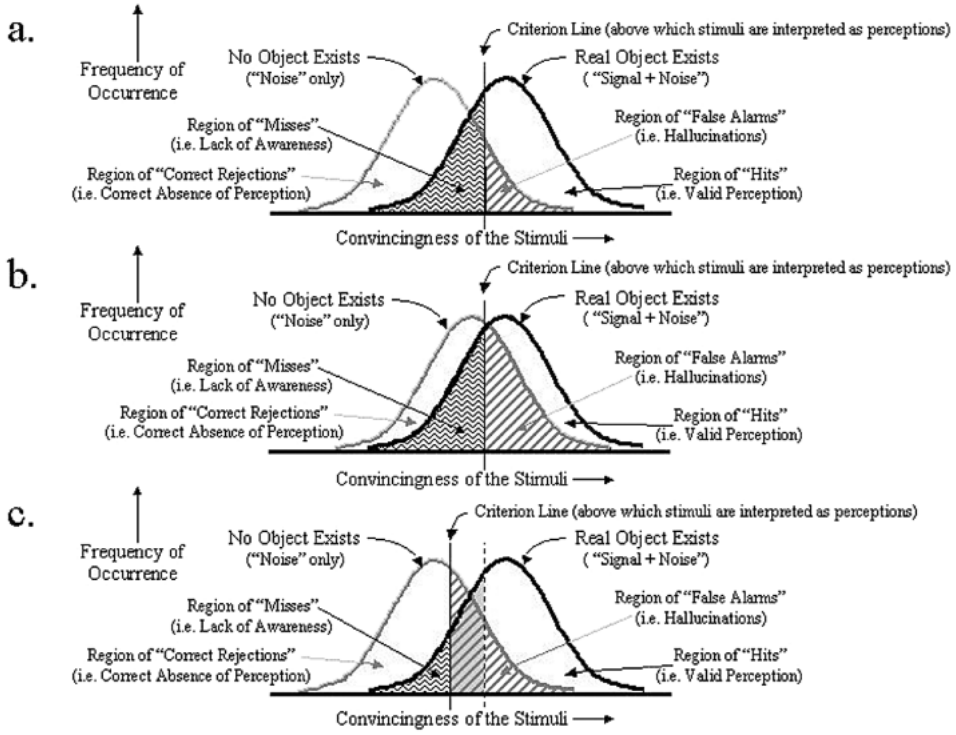

Figure 4: Dolgov & McBeath's signal detection argument: as noise levels increase, false percepts from a lowered acceptance threshold may be compensated by a corresponding increase in veridical ones [9].

The continuity of causation from simple to complex hallucinations in this combined model is particularly attractive given that in CBS both forms manifest regularly with no real aetiological distinction. That Burke's deafferentation can be the driver for stable pattern formation remains speculative, however, and further investigation of this relationship is required. Both the immediate effects of hyperexcitability on the continuum neural network model and the possible contributions from other factors should be considered. Even with substantial visual loss, simple hallucinations are only occasional. Indeed, since their occurrence is associated with low arousal just as are complex ones, it may be that some kind of attentional inhibition normally keeps them in check.

The workings of such top-down action for both simple and complex hallucinations is the main area of the model that remains distinctly sketchy, and this poses a problem for integration. There is an obvious mismatch between the quantitative nature of the pattern formation model and the more qualitative arguments of PAD. For the time being, we have little choice but to incorporate the lower level in similarly nebulous terms. The result certainly has some explanatory benefit, but the loss of rigour is disappointing.

4.2 Testability & Future Work

In seeking to validate or falsify our combined model, we would ideally identify edge cases that distinguish it from any other random fairy tale and can be tested experimentally. Visual hallucinations, by their very nature, are annoyingly unsusceptible to laboratory testing, and the vagaries of subjective description leave plenty of room to misconstrue, so validation is likely to be haphazard at best. Nevertheless, we can suggest some lines of enquiry.

The low-level portions of the model should certainly be amenable to further experimental investigation, attempting to determine the circumstances in which pattern formation may occur in deafferented cortical tissue. In addition to the sort of in vitro experiments cited by Burke, detailed in silico simulation of substantial V1-style neural networks is nowadays well within computational range, and systematic variation of local physical parameters might usefully demonstrate the sort of phenomena we hypothesise occurs in CBS.

The top-down aspects are much harder to pin down, and their is a corresponding dearth of testable predictions. The PAD authors' own proposed tests are frankly unsatisfactory, being in large part just a restatement of the observations from which the model was derived in the first place. Moreover, some specific predictions—for example, that incorrect proto-object selection can only occur by displacement of a correct candidate—do not seem to be required by the model at all; at least, not if we accommodate a low-level pattern formation component, which admittedly Collerton et al themselves do not.

In the absence of a well-defined model of the effects of attention—indeed, a physiological characterisation of top-down processing in general—testable predictions from the PAD model are likely to remain elusive. Thus, the investigation of CBS must in some way be subservient to the broader understanding of cognition.

With this in mind, it may be helpful to consider the question of hallucinations within the context of *constructive* attempts to model visual inference in the field of computer vision. For example, Yuille and Kersten describe a system in which Data-Driven Markov Chain Monte Carlo is used to synthesise a model of a scene from bottom-up candidates according to top-down knowledge embodied in a probabilistic context-free grammar [24]. This description is interesting because it includes top-down rejection of incorrect bottom-up proto-objects (in this case, a face wrongly identified in a pattern of tree bark). It is easy to see how changes in the model parameters might lead to this system suffering the equivalent of complex visual hallucinations.

A Continuum Neural Net Models

Wilson & Cowan [22] framed their continuum model for neural networks in terms of separate subpopulations of excitatory and inhibitory neurons. Their initial presentation ignores spatial variation and instead considers the *aggregate* excitatory (E) and inhibitory (I) activity as proportions of cells firing at any instant within a localised mixed population:

$$E(t+\tau) = \left[1 - \int_{t-r}^{t} E(t')\mathrm{d}t'\right] \cdot \sigma_e \left(\int_{-\infty}^{t} \alpha(t-t')[c_1 E(t') - c_2 I(t') + P(t')]\mathrm{d}t'\right)$$

$$I(t+\tau') = \left[1 - \int_{t-r'}^{t} I(t')\mathrm{d}t'\right] \cdot \sigma_i \left(\int_{-\infty}^{t} \alpha(t-t')[c_3 E(t') - c_4 I(t') + Q(t')]\mathrm{d}t'\right)$$

Here, τ and τ' are response delays, after which cells triggered at time t will be firing. r and r' are recovery times from the refractory state; hence, the expressions in square brackets are the proportion of cells at time t that are sensitive to inputs. The σ functions are sigmoid thresholds indicating the proportion of cells that will fire for a given input level; thus, the expressions in round brackets are the total effective inputs. Connectivity parameters c_1, c_2, c_3 and c_4 define the level of interaction between the neuron types; P and Q are the external inputs to each subpopulation; and α is a decay function for the effect of inputs over time.

In summary: the proportion of cells in each subpopulation that is firing will be the proportion of non-refractory cells for which the cumulative effect of their inputs is above their activation threshold.

Wilson & Cowan generalised this model to include spatial variation in a single dimension, introducing time coarse-graining—a sort of moving-window weighted average—to make the analysis tractable [23]. We shall skip over that version to the subsequent extension to two spatial dimensions in the hallucination work of Ermentrout & Cowan, which proceeds along much the same lines but ignores refractory insensitivity [10]:

$$\begin{aligned} E(x,y,t) = \sigma_e \Bigg(& \int_{-\infty}^{t} \mathrm{d}t'\ h_e(t-t') \int_{-\infty}^{\infty} \mathrm{d}x' \int_{-\infty}^{\infty} \mathrm{d}y' \\ & \cdot \Big[c_{ee} w_{ee}((x-x')^2 + (y-y')^2) E(x',y',t') \\ & - c_{ie} w_{ie}((x-x')^2 + (y-y')^2) I(x',y',t') \Big] \Bigg) \end{aligned}$$

($I(x,y,t)$ is omitted since it is functionally identical.)

In this formulation, h is a temporal response function subsuming both the response delay τ and the decay function α from the first version; c_{ee} and c_{ie} are synaptic connection strengths between cell types; and w_{ee} and w_{ie} are weighting functions characterising the variation in connectivity with distance.

If we specify an exponential decay time function

$$h_e(t) = h_i(t) = e^{-t}$$

define time coarse-grained equivalents of the activity variables

$$\begin{pmatrix} \hat{E} \\ \hat{I} \end{pmatrix} = \int_{-\infty}^{t} e^{-(t-t')} \begin{pmatrix} E(t') \\ I(t') \end{pmatrix} \mathrm{d}t'$$

and take the following notation for the spatial convolution

$$w \otimes z \equiv \int_{-\infty}^{\infty} \mathrm{d}x' \int_{-\infty}^{\infty} \mathrm{d}y' \; w(x - x', y - y') \; z(x, y)$$

we can derive the Ermentrout-Cowan model:

$$\frac{\partial \hat{E}}{\partial t} = -\hat{E} + \sigma_e(c_{ee} w_{ee} \otimes \hat{E} - c_{ie} w_{ie} \otimes \hat{I})$$

$$\frac{\partial \hat{I}}{\partial t} = -\hat{I} + \sigma_i(c_{ei} w_{ei} \otimes \hat{E} - c_{ii} w_{ii} \otimes \hat{I})$$

Bressloff et al [3, 4] dispense with the separation of subpopulations, merging excitation and inhibition into a single real-valued activity variable a, and introduce a new dimension, ϕ, for edge orientation. Further, based on observed patterns of connectivity in animal V1, they include two distinct levels of connectivity: *local*, within the same hypercolumn, and *lateral*, which is longer range and anisotropic, occurring between remote neurons tuned to the same orientation. Substituting vector $\mathbf{r} \in \mathbb{R}^2$ for the spatial variables x and y, they arrive at

$$\begin{aligned} \frac{\partial a(\mathbf{r}, \phi, t)}{\partial t} = &- \alpha a(\mathbf{r}, \phi, t) \; + \; \frac{\mu}{\pi} \int_0^{\pi} w_{loc}(\phi - \phi') \, \sigma\Big(a(\mathbf{r}, \phi', t)\Big) \, \mathrm{d}\phi' \\ &+ \nu \int_{-\infty}^{\infty} w_{lat}(s) \, \sigma\Big(a(\mathbf{r} + s\mathbf{e}_\phi, \phi, t)\Big) \, \mathrm{d}s \end{aligned}$$

In this case, α is a scaling constant for time dependence, μ and ν are coupling constants for the two kinds of connectivity, and w_{loc} and w_{lat} are, respectively, the local and lateral weighting functions, both taken to be 'Mexican hats' of the form

$$\frac{e^{-x^2/2\sigma_1^2}}{\sigma_1} - \frac{e^{-x^2/2\sigma_2^2}}{\sigma_2}$$

It is this system that Bressloff et al use in their analysis of pattern formation and symmetry. Note that if orientation sensitivity is omitted, it reduces to a single-population version of Ermentrout-Cowan.

References

[1] Emily J Abbott, Gillian B Connor, Paul H Artes, and Richard V Abadi. Visual loss and visual hallucinations in patients with age-related macular degeneration (Charles Bonnet Syndrome). *Investigative Ophthalmology and Visual Science*, 48(3):1416–1423.

[2] Daniel L Adams and Jonathan C Horton. A precise retinotopic map of primate striate cortex generated from the representation of angioscotomas. *Journal of Neuroscience*, 23(9):3771–3789, 2003.

[3] Paul C Bressloff, Jack D Cowan, Martin Golubitsky, Peter J Thomas, and Matthew C Wiener. Geometric visual hallucinations, Euclidean symmetry and the functional architecture of striate cortex. *Philosophical Transactions of the Royal Society B*, 356:299–330, 2001.

[4] Paul C Bressloff, Jack D Cowan, Martin Golubitsky, Peter J Thomas, and Matthew C Wiener. What geometric visual hallucinations tell us about the visual cortex. *Neural Computation*, 14:473–491, 2002.

[5] W Burke. The neural basis of Charles Bonnet hallucinations: a hypothesis. *Journal of Neurology, Neuroscience and Psychiatry*, 73:535–541, 2002.

[6] David G Cogan. Visual hallucinations as release phenomena. *Archive of Clinical and Experimental Ophthalmology*, 188:139–150, 1973.

[7] M Cole. Charles Bonnet syndrome: an example of cortical dissociation syndrome affecting vision? *Journal of Neurology, Neurosurgery and Psychiatry*, 71(1):134–, 2001.

[8] Daniel Collerton, Elaine Perry, and Ian McKeith. Why people see things that are not there: a novel Perception and Attention Deficit model for recurrent complex visual hallucinations. *Behavioural and Brain Science*, 28:737–794, 2005.

[9] Igor Dolgov and Michael K McBeath. A signal-detection-theory representation of normal and hallucinatory perception. *Behavioural and Brain Science*, 28:761–762, 2005.

[10] G B Ermentrout and J D Cowan. A mathematical theory of visual hallucination patterns. *Biological Cybernetics*, 34:137–150, 1979.

[11] Antony Fernandez, Gil Lichtshein, and Victor Vieweg. The Charles Bonnet Syndrome: A Review. *Journal of Nervous and Mental Disease*, 185(3):195–200, 1997.

[12] D H ffytche and R J Howard. The perceptual consequences of visual loss: 'positive' pathologies of vision. *Brain*, 122:1247–1260, 1999.

[13] Dominic H ffytche. Two visual hallucinatory syndromes. *Behavioural and Brain Science*, 28:763–764, 2005.

[14] Thomas R Hedges Jr. Charles Bonnet, his life, and his syndrome. *Survey of Ophthalmology*, 52(1):111–114, 2007.

[15] Hans Klüver. *Mescal and mechanisms of hallucinations.* University of Chicago Press, 1966.

[16] G Jayakrishna Menon. Complex visual hallucinations in the visually impaired: a structured history-taking approach. *Archive of Ophthalmology*, 123:349–355, 2005.

[17] G Jayakrishna Menon, Imran Rahman, Sharmila J Menon, and Gordon N Dutton. Complex visual hallucinations in the visually impaired: the Charles Bonnet syndrome. *Survey of Ophthalmology*, 48(1):58–72, 2003.

[18] Peter A Nixon and John O Mason. Visual hallucinations from age-related macular degeneration. *American Journal of Medicine*, 119(3):e1–e2, 2006.

[19] Barry W Rovner. The Charles Bonnet syndrome: a review of recent research. *Current Opinion in Ophthalmology*, 17:275–277, 2006.

[20] Yasuko Shiraishi, Takeshi Terao, Kenji Ibi, Jun Nakamura, and Akihiko Tawara. Charles Bonnet syndrome and visual acuity: the involvement of dynamic or acute sensory deprivation. *European Archive of Psychiatry and Clinical Neuroscience*, 254:362–364.

[21] Yasuko Shiraishi, Takeshi Terao, Kenji Ibi, Jun Nakamura, and Akihiko Tawara. The rarity of Charles Bonnet syndrome. *Journal of Psychiatric Research*, 38:207–213, 2004.

[22] Hugh R Wilson and Jack D Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12:1–24, 1972.

[23] Hugh R Wilson and Jack D Cowan. A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13:55–80, 1973.

[24] Alan Yuille and Daniel Kersten. Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Science*, 10(7):301–308, 2006.