

A Neural Basis for Perceptual Dynamics

Howard S. Hock and Gregor Schöner

Abstract. Perceptual stability is ubiquitous in our everyday lives. Objects in the world may look somewhat different as the perceiver's viewpoint changes, but it is rare that their essential stability is lost and qualitatively different objects are perceived. In this chapter we examine the source of this stability based on the principle that perceptual experience is embodied in the neural activation of ensembles of detectors that respond selectively to the attributes of visual objects. Perceptual stability thereby depends on processes that stabilize neural activation. These include biophysical processes that stabilize the activation of individual neurons, and processes entailing excitatory and inhibitory interactions among ensembles of stimulated detectors that create the "detection instabilities" that ensure perceptual stability for near threshold stimulus attributes. It is shown for stimuli with two possible perceptual states that these stabilization processes are sufficient to account for spontaneous switching between percepts that differ in relative stability, and for the hysteresis observed when attribute values are continually increased or decreased.

The responsiveness of the visual system to changes in stimulation has been the focus of psychophysical, neurophysiological, and theoretical analyses of perception. Much less attention has been given to the role of persistence, the effect of the visual system's response to previous visual events (its prior state) on its response to the current visual input. Perceiving an object can facilitate its continued perception when a passing shadow briefly degrades its visibility, when attention is momentarily distracted by another object, when the eyes blink, or when a random fluctuation within the visual system potentially favors an alternative percept. Having perceived an object's shape from one viewpoint can facilitate its continued perception despite changes in viewpoint that distort its retinal projection, potentially creating a non-veridical percept. These examples highlight the importance of the visual system's prior state, not just for perceptual stability, but also for perceptual selection; i.e., for the determination of which among two or more alternatives is realized in perceptual experience.

In this essay we discuss three neural properties that form a sufficient basis for a theory of perceptual dynamics that addresses the relationship between persistence, responsiveness to changes in stimulation, and selection. These neural properties are: 1) Individual neurons have the intrinsic ability to stabilize their activation state. 2) Neurons responsive to sensory information (i.e., detectors) are organized into ensembles whose members respond preferentially to different values of the same attribute (e.g., motion direction). Members of such ensembles have overlapping tuning functions; i.e., a detector responding optimally to one stimulus

value will also respond, though less strongly, to similar attribute values. 3) The activation levels of a detector affects and is affected by nonlinear excitatory and inhibitory interactions with other detectors.

On this basis, we examine the persistence of steady-state detector activation despite the presence of random perturbations, the effect of neural stabilization on a detector's response to stimulation, the crucial role of "detection instabilities" in minimizing perceptual instability and uncertainty for near-threshold stimuli, and the importance of differences in the rate-of-change in activation for perceptual selection. Finally, we demonstrate that the signature features of perceptual dynamics, spontaneous switching between percepts differing in relative stability, and hysteresis, follow from the same three neural properties.

1 Perceptual Stability: Natural or Otherwise

Natural, everyday percepts are almost invariably monostable. The same percept occurs each time a stimulus is presented. It rarely happens that two qualitatively different percepts are formed for the same stimulus (this would constitute bistability), and the experience of spontaneous switching between alternative percepts is likewise rare. Because everyday experiences of monostability are so pervasive, stability is not always recognized as an important perceptual property. Not so for James Gibson (1966), who attributed the stability of real-world percepts to the tuning of our visual system to unambiguous, invariant properties of stimulation.

Although Yuille and Kersten (2005) take a different position, maintaining that natural images are inherently ambiguous, they join Gibson (1966) and others in disdaining the usefulness of artificial stimuli for an understanding of perception in the natural environment. It is arguable, however, that many natural objects are potentially bistable (e.g., bumps and holes), but there is sufficient disambiguating contextual information in the natural environment to over-ride the potential of such objects to exhibit the dynamical behavior that is readily observed in the laboratory. Indeed, it is the exceptional situations accessible in the laboratory that most clearly bring the fundamental indeterminance of perceptual bistability into the domain of phenomenal perception.

Our dynamical research has taken place in well-controlled laboratory settings, where we have studied single-element apparent motion (Hock, Kogan & Espinoza, 1997; Hock, Gilroy & Harnett, 2002), displaced targets embedded in noise (Eastman & Hock, 1999), displaced rows of evenly spaced dots (Hock & Balz, 1994; Hock, Balz & Smollon, 1998; Hock, Park & Schöner, 2002), and the motion quartet (which is described below).

Single element displacements result in unique motion percepts, and many stimuli with multiple element displacements result in the perception of unique motion patterns. For instance, parallel horizontal motions are perceived for two vertically aligned elements alternating with two vertically aligned elements that are horizontally displaced, as in Figure 1a. This percept uniquely solves the motion correspondence problem; i.e., how visual elements presented during successive points in time are "paired" with respect to the start and end of

perceived motion paths (Ullman, 1979). This is the case even though diagonal motions are in principle also possible for this stimulus. That is, despite single element motion being easily perceived for each independently presented diagonal displacement (Figures 1b and 1c), intersecting diagonal motions are never seen when the two diagonal displacements are combined in the same stimulus, as in Figure 1a.

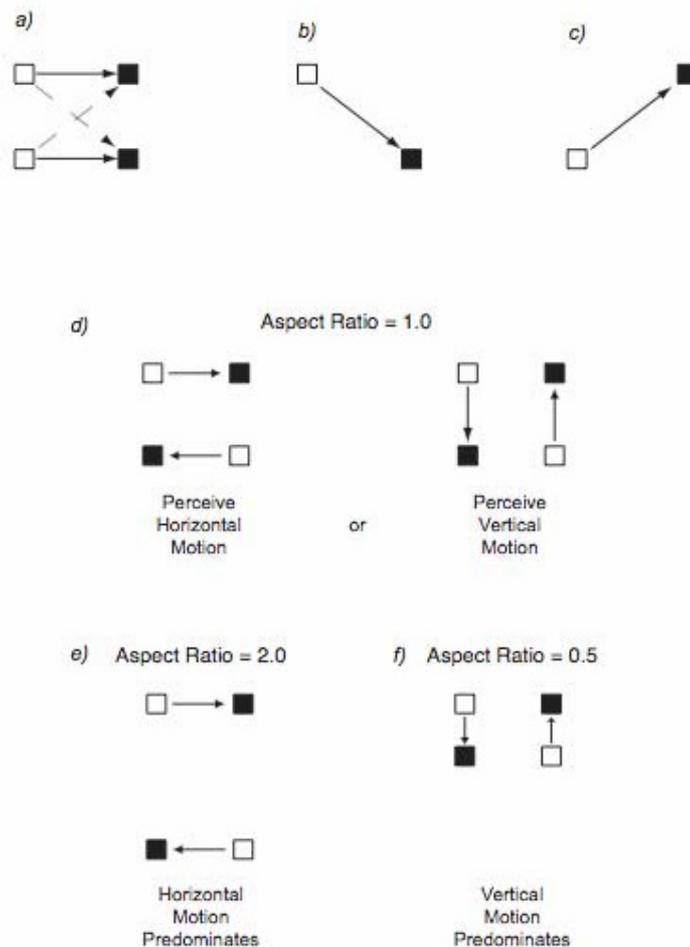


Fig. 1 (a) An illustrative apparent motion stimulus for which there is a unique solution to the motion correspondence problem (horizontal motion always is perceived), even though diagonal motions are possible (b and c). (d) The motion quartet, an apparent motion stimulus for which there are two qualitatively different solutions to the motion correspondence problem. For intermediate aspect ratios (the vertical divided by the horizontal distance between the elements), either horizontal or vertical motion is perceived. Horizontal motion predominates for relatively large aspect ratios (e) and vertical motion predominates for relatively small aspect ratios (f).

In contrast to the stimulus in Figure 1a, the motion quartet is an apparent motion stimulus for which two qualitatively different solutions can be realized in experience; either horizontal or vertical motion is perceived, the proportion of each depending on the aspect ratio of the quartet (Figures 1d-1f). Our research with this bistable stimulus has included psychophysical experiments and the dynamical modeling of spontaneous switching, hysteresis, selective adaptation, and activation-dependent detector interactions (Hock, Kelso & Schöner, 1993; Hock, Schöner & Hochstein, 1996; Hock, Schöner & Voss, 1997; Hock, Schöner & Giese, 2003; Hock, Bukowski, Nichols, Huisman, & Rivera, 2005; Hock & Ploeger, 2006; Nichols, Hock & Schöner, 2006). We currently are studying the effect of neural feedback on the stabilization of global motion patterns for stimuli composed of multiple motion quartets (Hock, Brownlow & Taler, in preparation).

It is for stimuli like the motion quartet that it is possible to directly observe the nonlinear mechanisms that bind stimulus specification with the ongoing neural activity resulting from preceding visual events, revealing fundamental properties of the processing mechanisms that are the basis for perception, not just in the laboratory, but in the natural environment as well. The most fundamental of these properties is neural self-stabilization.

2 Neural Stabilization

Whether an individual neural detector is activated by a stimulus or not, random events (perturbations) will cause its activation to fluctuate randomly with respect to some steady-state value. However, the variability of these fluctuations does not increase indefinitely over time. Although at first glance this may not be surprising, the "boundedness" of variability reflects a crucial, though often unrecognized feature of neural behavior, namely, that a neuron's activation is actively stabilized.

This idea can be made concrete by starting with any activation level for a neuron at any moment in time, and assuming that there is no interaction with other neurons. A random perturbation, if unconstrained, with equal probability will increase or decrease the neuron's activation. Assume it increases activation. The next and all following random perturbations will again with equal probability increase or decrease activation. Thus, there is nothing that systematically returns the activation from its increased level. Similarly, if an initial perturbation decreases activation, there is nothing that returns the activation from its decreased level. The same logic applies to any activation state generated by perturbations. Over time, states further and further removed from the initial activation state can be reached (e.g., by the chance event of a number of consecutive random increases in activation) and nothing drives the system systematically back from such states. It is intuitively clear, therefore, that the variance of activation would increase indefinitely over time. A formal argument of this kind led to an account for Brownian motion and the increase in time of the uncertainty about the location of a Brownian particle (Einstein, 1905).

The essential feature that keeps the variance of random fluctuations bounded is that successive random perturbations do not increase or decrease the neuron's activation level with equal probability. That is, the effects of random perturbations on activation are not unconstrained. When a random perturbation causes a fluctuation in activation, the change is opposed by the neuron's intrinsic ability to stabilize its activation, which reduces the size of the fluctuation. It is because of this resistance to the effects of random perturbations that there is an upper bound to the variance of random fluctuations in activation. The steady-state activation value of a neuron (or population of neurons) that is thus stabilized against the effects of random perturbations is referred to as an attractor.

2.1 The Biophysical Basis of Neural Stabilization

The biophysics of individual neurons provides a mechanism for achieving this stabilization of neural activation (Trappenberg, 2002). Specifically, the electrical potential across the membrane that separates the interior of a nerve cell from its inter-cellular environment is kept stable through the mechanisms of osmotic pressure. Ion pumps keep the concentration of different kinds of ions unequal on both sides of the membrane, the resulting flow of ions being in equilibrium when the electrical potential across the membrane just counterbalances the difference in ion concentration. If the equilibrium is perturbed (e.g., by an electrical current injected into the cell), the flow of ions quickly re-establishes the steady-state membrane potential.

With a neuron's membrane potential thus stabilized, synaptic input to the neuron increases the potential, increasing the probability that the neuron's activation will be transmitted to other neurons through action potentials traveling down its axon. In our account of neural dynamics (and most other such accounts) the stabilized membrane potential, averaged over local neural populations composed of hundreds or thousands of individual neurons, is sufficient to account for the mapping of psychophysical events onto patterns of neural activation. To be sure, the mathematical relationship between ion flows that stabilize the membrane potential of individual neurons and the stability properties of neural populations is not well understood. Eggert and van Hemmen (2001) have provided one such account, but it is limited by the simplifying assumptions that the constituents of a neural population are both identical in their responsiveness to stimulation and non-interactive. This notwithstanding, it is reasonable to proceed based on the principle that stability properties of neural populations are inherited from the dynamics through which individual neurons stabilize their membrane potential (Jancke, Erlhagen, Dinse, Akhavan, Giese, Steinhage & Schöner, 1999).

2.2 The Time Scale

The extent to which a neuron or population of neurons resists fluctuations in activation caused by random perturbations depends on how quickly fluctuation-opposing changes emerge within the neurons. This determines the time scale of

stabilization. If there were only one instantaneous perturbation, the pre-perturbation activation level (i.e., the average membrane potential) would be restored over an interval determined by the time scale. This is called the “relaxation time.” However, random perturbations occur continually, so depending on the time scale, there is sufficient time only for the partial restoration of the fluctuation in activation caused by one perturbation before the next one occurs. The faster the time scale, T , the greater the restoration of activation, and therefore, the greater the resistance to the effects of the random perturbation.

2.3 The Core Dynamical Concept

Neural stabilization provides the basis for the core concept of a theory of perceptual dynamics. That is, whatever causes a change in the current neural activation, u , will be opposed in the immediate future by a change in activation, du/dt , in the opposite direction. Activation increases in the immediate future when the change in activation, du/dt , is positive (because u has decreased) and it decreases in the immediate future when the change in activation, du/dt , is negative (because u has increased). This relationship among current levels of activation, u , and changes in activation that will occur in the immediate future (du/dt) can be expressed as:

$$du/dt = -u/T$$

where T is the time scale of perturbation-opposing reactions within the neuron. T determines the size of the change in activation in opposition to random fluctuations, with larger compensating changes (larger values of du/dt) occurring when T is smaller/faster.

2.4 Stable Activation States in the Absence of Stimulation

In classical, non-dynamical approaches to the study of perception, unstimulated detectors are simply inactive, and although there are numerous dynamical accounts of perception, they generally do not address the status of detectors when they are unstimulated. In our dynamical conceptualization, however, a neuron's ability to stabilize its activation means that populations of detectors have stable activation states even when they are unstimulated, and irrespective of their connectivity to other detectors. This means that unstimulated detectors can maintain activation near an attractor value that is below the threshold level required for perception, thereby minimizing the likelihood that random fluctuations would cause the activation of unstimulated detectors to rise above this threshold. The stabilization of activation in the absence of stimulation (i.e., in the vicinity of the no-stimulus, resting level) can be characterized by adding h to the dynamical equation:

$$du/dt = (-u + h)/T$$

where h is the detector's resting level.

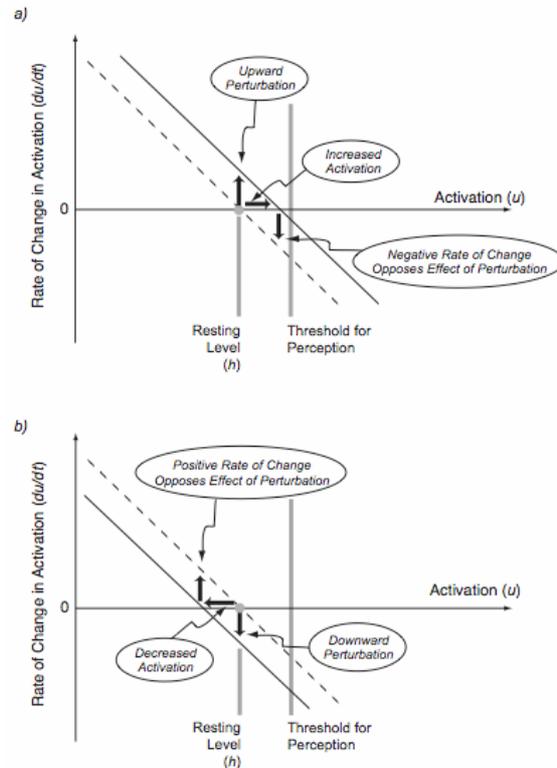


Fig. 2 The straight lines with negative slope represent the stabilization of activation (u) in the absence of stimulation (i.e., in the vicinity of the no-stimulus, resting level). This is determined by the dynamical equation: $du/dt = (-u + h) / T$, where the rate-of-change, du/dt , determines whether and by how much activation will increase or decrease in the immediate future. h is the detector's resting level and T is the time scale of activation change. Because of the negative slope, random fluctuations in activation are resisted by changes in the opposite direction, with activation stabilizing at the attractor for the no-stimulus/resting state ($u^* = h$). Panels (a) and (b) differ with respect to the time scale of the dynamics

As can be seen in the graphical representation of the equation (Figure 2), the attractor for the no-stimulus (resting) state is the activation value $u^* = h$. By simple calculation from the above equation, it is the steady-state activation level when the rate-of-change, du/dt , is 0. When a perturbation occurs, it imposes a randomly determined upward or downward rate of change in activation; i.e., it creates a tendency for activation to change in a particular direction. Such perturbations cannot be instantaneous. They must last long enough for activation to reach a value different from h . It can be seen in Figure 2a that the value of du/dt is negative when the perturbation imposes a positive rate of change, so neural stabilization opposes the effect of the perturbation by proportionally decreasing activation following the perturbation. It similarly can be seen in Figure 2b that the

value of du/dt is positive when a random perturbation imposes a negative rate of change in activation, so neural stabilization opposes the effect of the perturbation by proportionally increasing activation after the perturbation. Irrespective of direction, the larger the change in activation caused by the perturbation, the greater the opposing change, du/dt . The latter, together with randomness with respect to whether perturbations have positive or negative effects, stabilizes activation in the vicinity of the resting level.

2.5 Response to Stimulation

What happens when a stimulus is presented for which detectors are responsive? In classical non-dynamical approaches, the steady-state activation level of a detector is determined by the strength of the stimulus, and there is little concern with how activation evolves over time toward these steady-state values. In our dynamical conceptualization, however, the initial response to the presentation of a detector-activating stimulus occurs in the context of activation states (at the resting level) that are stabilized with respect to the effects of random perturbations. This stabilization mechanism therefore determines how activation evolves from the resting level toward the steady-state activation level determined by the stimulus.

This can be made intuitive by imagining that a stimulus presented for a finite period of time is equivalent to a dense sequence of activation-increasing perturbations (rather than a random sequence of positive and negative perturbations, as in the preceding section). Each of these "perturbation-induced" upward fluctuations in activation is partially opposed by neural stabilization, so at a rate determined by the time scale, T , of the neural stabilization mechanism, successive excitatory perturbations incrementally move a detector's activation away from its (no-stimulus) resting level, toward the stimulus-determined activation level. By imagining the stimulus as a dense sequence of partially restored excitatory fluctuations, it can be understood that the time course of the activation as it rises from the resting level depends on the same neural stabilization mechanism that keeps the activation of unstimulated detectors from randomly fluctuating above the threshold level for perception. The activational effect of a stimulus on a detector ensemble therefore can be characterized by adding S to the previously introduced dynamical equation:

$$du/dt = (-u + h + S) / T$$

When the stimulus is presented, it imposes a positive rate-of-change on activation; i.e., activation increases immediately after the stimulus is presented, which moves it from the no-stimulus attractor, $u^* = h$, toward the stimulus-determined attractor, $u^* = h + S$. (It can be seen from the graphs in Figure 3 that du/dt is positive in relation to the stimulus-determined attractor when activation has been at the resting level.) A comparison of Figures 3a and 3b shows that the time scale determines how quickly activation changes as it moves away from the resting level, toward the stimulus-determined attractor; smaller/faster time scales result in larger, more rapid shifts in activation toward the attractor.

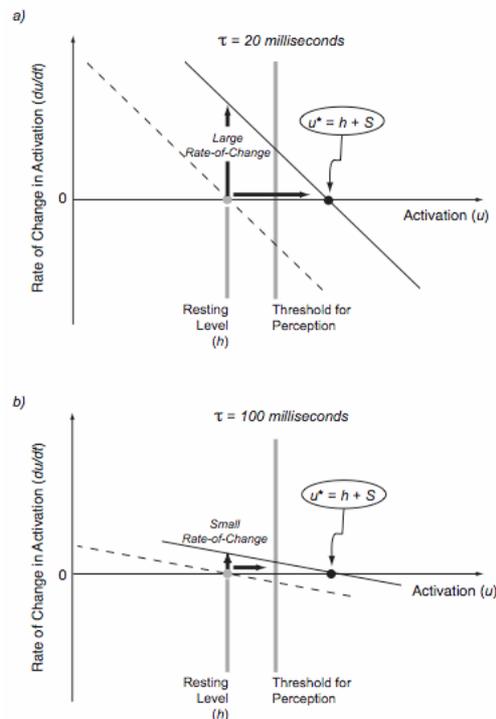


Fig. 3 The broken straight lines with negative slope represent the stabilization of activation (u) in the absence of stimulation, as in Figure 2, and the parallel solid lines represent the stabilization of activation (u) in the presence of stimulation. The latter is determined by the dynamical equation: $du/dt = (-u + h + S) / T$, where the rate-of-change, du/dt , determines whether and by how much activation will increase or decrease in the immediate future. h is the detector's resting level, T is the time scale of activation change, and S is the stimulus-initiated activation. Activation increases to the steady-state attractor value, $u^* = h + S$, at a rate determined by the time scale of the dynamics, which differs in panels (a) and (b).

This evolution of activation for the stimulated detector is illustrated in Figure 4 for two different time scales, which shows that there is a greater rate-of-change in activation for the faster time scale as activation rises from the resting level. In addition, random fluctuations are less variable (with the same level of random noise perturbations) for the faster time scale. So long as the detector's activation is not influenced by interaction with other detectors (or adaptation), activation would settle near the attractor, $u^* = h + S$, for both time scales. For the simulations in Figure 4, $h = -8$ and $S = 16$, so it is readily calculated from the above equation that when du/dt is zero, the attractor is at $u^* = -8 + 16 = 8$.

It is important to note that the comparisons in Figures 3 and 4 are made in order to provide an intuitive understanding of the effects of time scale on changes in activation. At this juncture, it is not possible to directly determine time scales for perceptual dynamics (slower time scales cannot be discriminated from higher levels of noise) or to separate contributions to time scale from the stabilization of

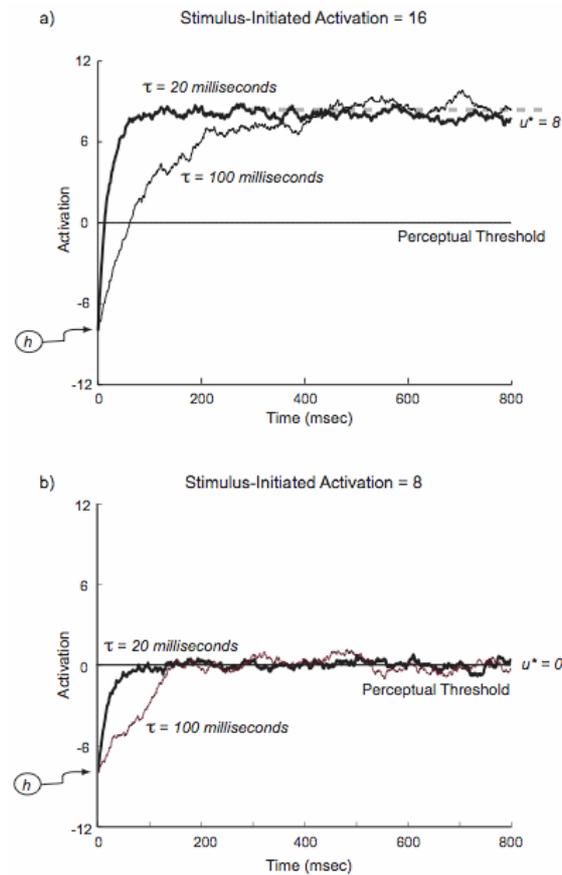


Fig. 4 This evolution of activation as the result of stimulation for two time scales, demonstrating the faster change in activation for the briefer time scale. Activation settles near the attractor, $u^* = h + S$, for both time scales, with $S = 16$ in panel (a) and $S = 8$ in panel (b). Activation for the latter would straddle the threshold for perception, $u^* = -8 + 8 = 0$, so even small random fluctuations would render perception highly unstable.

membrane potentials for individual neurons, and from to-be-discussed interactions between detectors that compose neural ensembles.

2.6 Perceptual Thresholds, Perceptual Stability and Uncertainty

In accounts of perception based on neural dynamics, it is assumed that a stimulus attribute is perceived when the average activation of detector populations responsive to that attribute is stabilized at a level that exceeds a particular threshold value (perhaps determined by the level of membrane potential that results in the transmission of action potentials). However, neural stability does not

guarantee perceptual stability. If the stimulus-initiated activation in the preceding example were weaker ($S = 8$), the stabilized activation (i.e., the attractor) would straddle the threshold for perception, $u^* = -8+8 = 0$. As a result, even small random perturbations would rapidly shift activation back and forth across the threshold, rendering perception highly unstable, and therefore, highly uncertain (Figure 4b).

This kind of near-threshold uncertainty, which is classically the domain of signal detection theory (Green & Swets, 1966), parallels the dynamical account presented thus far. That is, both entail the detection of a signal (the steady-state activation level) embedded in noise (random fluctuations in activation), and a criterion (the perceptual threshold) that determine whether or not the attribute is present (perceived). While signal detection theory is sufficient to account for the near-threshold uncertainty that occurs in many contexts, near-threshold uncertainty does not generally obtain in motion perception. For example, Hock, Kogan and Espinoza (1997) found values of luminance contrast change that result in the perception of single-element apparent motion for half the trials and the perception of nonmotion for the other half. For almost all the trials, either motion or nonmotion was clearly perceived. There were only occasional trials for which subjects were uncertain regarding what they perceived. Consistent with such experimental results, perceptual uncertainty does not occur in our dynamical conceptualization because as activation rises from the resting level for individual detectors, interaction with other detectors emerges, and what would be near-threshold activation levels for the individual detectors are boosted to above-threshold levels and further stabilized at those levels. How this occurs is discussed next.

3 The Stabilization of Activation within Detector Ensembles

Up until this point in the discussion we have focused on the stabilization properties of individual detectors (or populations of detectors) in the absence of interaction among them. In this section we step back to establish the neuro-anatomical basis for detector interactions and the particular interactions that are the basis for detection instabilities.

3.1 Neural Connectivity

In a neural network conceptualization of perception, stimulus specification is roughly characterized by the feedforward path through the network. Activation induced in this path is largely stimulus determined, the preferential responding of different detectors occurring by virtue of their receptive fields being structured to realize various feature extraction filters (for motion, orientation, line length, texture, color, etc.). However, most neuronal activity entails more than the feedforward stream. Braitenberg (1978) has estimated that 95% of the input to each cortical neuron comes from its connectivity with other cortical neurons, and Felleman and Van Essen (1991) have determined that there is more feedback than feedforward connections between higher- and lower-level areas in the brain.

Given this neuro-anatomical evidence, one cannot expect the visual system to simply “compute” the perceptual output from the stimulus input. When a stimulus is presented, it is necessary to also take into account what already is occurring within highly interconnected neural networks. This is the domain of neural dynamics, which was first introduced into the study of perception by Stephen Grossberg and his colleagues (Grossberg, 1973; Grossberg & Mingolla, 1985; Grossberg & Rudd, 1992; Francis, & Grossberg, 1996; Chey, Grossberg, & Mingolla, 1997; Baloch & Grossberg, 1997). In our analysis, perceptual dynamics describe how activation within a network of detectors evolves in time under the influence of both current input (stimuli consistent with the preferential responding of detectors in the network) and ongoing activity (persistent activation due to earlier stimulus input).

3.2 Interaction

The most important consequence of neural connectivity is that a detector, when activated, can influence the activation levels of other detectors, either by increasing their activation through excitatory interaction or decreasing their activation through inhibitory interaction. An essential feature of such interactions is that they are activation dependent. That is, the more strongly a detector is activated, the greater its interactive influence on the detectors with which it is connected. This is illustrated by the sigmoidal function in Figure 8b.

Thus, each detector in an interconnected network is subject not only to the time-varying activational effect of the stimulus and random perturbations that produce fluctuations in its activation, but also to activation-dependent interactive influences from other detectors. (Another kind of influence, which we will ignore in this chapter, is activation-dependent adaptation; but see Hock et al. 2003; Nichols et al. 2006.) Because of this interaction, all contributions to a detector's activation change its influence on other detectors, which ultimately comes around again by affecting their interactive influence on the detector whose activation is the source of the interaction. This ongoing re-cycling of change, or recurrence, reflects the state-dependence of the network; i.e., the evolving activation state of a detector depends on its own previous activation state as well as the activation state of the detectors with which it interacts. In this way, the presentation of a stimulus initiates a recurrent cycle of activation change in networks of detectors that moves activation toward steady-state values, values for which rates-of-change in activation are near zero.

3.3 Interaction within Detector Ensembles

As indicated earlier, a detector's activation can be stabilized at a steady-state value, but this will not result in perceptual stability if the steady-state value is near the threshold for perception. Even small random perturbations would then be sufficient to rapidly shift activation back and forth across the threshold, rendering

perception unstable and uncertain. This "problem" can be resolved by assuming that perception requires the stimulus-initiated activation of individual detectors, but that it is ultimately determined by the pattern of activation within ensembles of interacting detectors that respond preferentially to different values of the same attribute. Because these detectors have overlapping tuning functions, some will respond optimally to one attribute value, others will respond, though less strongly, to that attribute value, and still others will not respond at all to that attribute value. It is within the population of stimulated detectors that excitatory and inhibitory interactions create detection instabilities.

3.4 Detection Instability

When an appropriate stimulus is presented, the activation of responding ensemble members rises from their resting levels (at a rate determined by each detector's neural stabilization mechanism), and an activation level is approached for which the population of detectors will boost each others' activation through mutual excitatory interaction. Up until this point, activation is still subthreshold for perception, so whether or not further growth in stimulus-initiated activation engages excitatory interactions will have dramatic effects on activation levels for the population of stimulated ensemble members. If stimulus-initiated activation is insufficient to initiate excitatory interactions among ensemble members, activation for the stimulated detectors will be stabilized below the threshold level required for perception. However, if stimulus-initiated activation is sufficient to initiate excitatory interactions, detectors that respond optimally will increase the activation of detectors that prefer similar, but different attribute values. And the latter, in turn, will increase the activation of the detector already responding optimally to the stimulus.

This reciprocal excitation will result in detector activation rapidly passing through a "detection instability" (Bicho, Mallet, & Schöner, 2000; Schöner, 2008; Schneegans & Schöner, 2008). Activation will accelerate well past levels that are near-threshold for perception because detector interactions are activation dependent; i.e., the more strongly a detector is activated the greater its excitatory effect on other ensemble members. The indefinite increase in activation for these self-excited detectors is prevented by the accompanying rise in activation-dependent inhibitory interactions, resulting in activation settling at a stable value (attractor) for the population of detectors that respond preferentially to attribute values at or near to the attribute value of the stimulus. The inhibitory interactions that limit the growth in self-excited activation, though weaker than the excitatory interactions, come from ensemble members with preferences for a wider range of attribute values (Amari, 1977; Wilson & Cowan, 1973). As a result, pairs of detectors with similar preferences will more strongly excite than inhibit each other, whereas the reverse will be the case for pairs of detectors with dissimilar preferences.

Inhibitory interactions are thus comparatively long-range. They can be effective over longer distances (in attribute space) than short-range excitatory interactions,

so they prevent the spread of activation to inter-connected detectors with preferences for attribute values that are much different than the attribute value of the stimulus. It will be seen in Section 4 that this long-range inhibition also is important for perceptual selection, which comes into the picture when stimulus-initiated activation for much different attribute values makes it possible for the activation of more than one population of ensemble members to pass through a detection instability.

3.5 Stabilization of Activation within Detector Ensemble

When a self-excited population of detectors passes through a detection instability, its activation is boosted well past the threshold level required for perception, and in addition, its stability is enhanced with respect to the fluctuations in activation produced by random perturbations. The reason is as follows. If the net effect of random perturbations over the population of detectors is for activation to fluctuate upwards, the increased activation will spread through detector pairs with similar preferences via short-range excitatory interactions. However, detector activation due to the net-upward fluctuations is compensated for by increased inhibitory interactions. In the case of inhibition there is both short-range inhibition between pairs of detectors with similar preferences, and long-range inhibition between pairs of detectors with dissimilar preferences. The net effect will be to oppose the net-increase in activation by increasing activation-reducing inhibitory interactions.

The opposite occurs when random perturbations result in a net downward fluctuation in activation for the population of self-excited detectors. The reduced activation will spread through the reduction in short-range excitation between detector pairs with similar preferences. However, the reduction in activation also will decrease both short-range inhibition between pairs of detectors with similar preferences, and long-range inhibition between pairs of detectors with dissimilar preferences. The net effect will be for the population of self-excited detectors to oppose the net-decrease in activation by decreasing activation-reducing inhibition, which effectively increases activation.

Stability therefore is asymmetrical for the two states associated with the detection instability. If stimulus-initiated activation is too weak to engage the excitatory interactions that boost the ensemble's activation, activation will remain below the threshold level required for perception. Stability in the presence of random perturbations then would be determined by the intrinsic neural stabilization properties of the stimulated detectors (i.e., by the balance of membrane potentials and ion concentrations). If stimulus-initiated activation is sufficient for activation to produce a self-excitatory boost in activation, stability in the presence of random perturbations would then be determined by both the neural stabilization properties of the stimulated detectors and the balance of excitatory and inhibitory interactions that results from their passing through a detection instability, as described above.

3.6 *Removing the Stimulus*

The new stable state induced by the detection instability resists decay when the activation-initiating stimulus is removed. This is because excitatory interactions keep boosting detector activation even as the removal of the stimulus induces a reduction in activation. Although this resistance to the removal of the stimulus ultimately fails, it slows the return of activation to the detectors' no-stimulus resting levels, so that activation persists for time intervals that are much longer than those needed to initially stimulate the detectors and boost activation through the detection instability. This lingering of activation (below the threshold for perception, but above the resting level) potentially accounts for percepts exhibiting stability over surprisingly long temporal intervals (Leopold, Wilke, Maier & Logothetis, 2002).

4 Perceptual Selection

Thus far we have discussed how activation within a detector ensemble is stabilized around detectors that optimally respond to a particular attribute value. The principle is that above-threshold activation for these self-excited detectors would signify the perception of this attribute value. However, in natural perception (and in the laboratory), activation can be simultaneously initiated around more than one attribute value. For example, both horizontal and vertical motions are simultaneously stimulated by the motion quartet described in Figures 1d-1f. When the detectors responsive to one attribute value are more strongly stimulated than the detectors responsive to other attribute values, the long-range inhibitory interactions that stabilize activation in one part of the attribute space (Sections 3.3 and 3.4) can reduce stimulus-initiated activation in other parts of the attribute space, preventing activation at that location from passing through an activation-boosting detection instability. This can result in one overwhelmingly dominant percept (the horizontal motion in Figure 1a), even though the alternative attribute values (the diagonal motions in Figures 1b and 1c) would be perceived in the absence of competition from the detectors that respond optimally to horizontal motion. The characteristics of this inhibitory competition, which would contribute to the monostability typical of natural perception, is best studied in the case of bistability; i.e., when the competing stimulus attributes are relatively similar with respect to the level of activation that they stimulate.

4.1 *Perceptual Bistability*

Perception is said to be bistable when two different percepts are possible for the same stimulus. When the competing percepts are based on different values of the same attribute (we will call the values A and B), the value selected for perception depends on activation-dependent inhibitory interactions between the populations of detectors responsible for the perception of each attribute value. When the stimulus-initiated activations for A and B are equal, random fluctuations occurring as activation rises

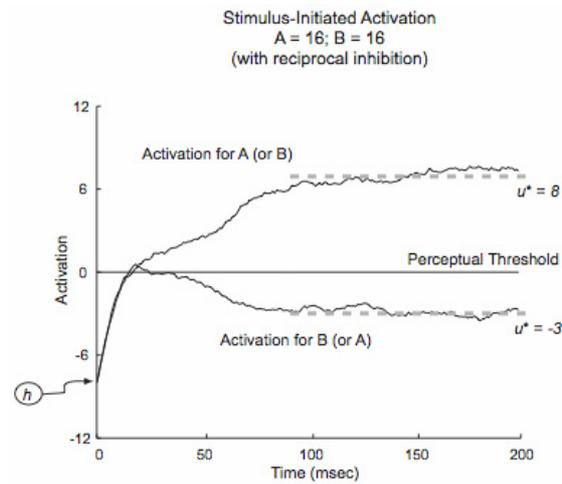


Fig. 5 When the stimulus-initiated activations for attributes A and B are equal, random fluctuations occurring as activation rises from the resting level result in detector populations that preferentially respond to one attribute having a momentary activational advantage over detector populations that preferentially respond to the other attribute. Reciprocal, activation-dependent inhibitory interactions between the two populations drive their activations apart such that the randomly advantaged attribute has its activation stabilized at an above-threshold level ($u^* = 8$) while its competitor is suppressed to a sub-threshold activation level ($u^* = -3$).

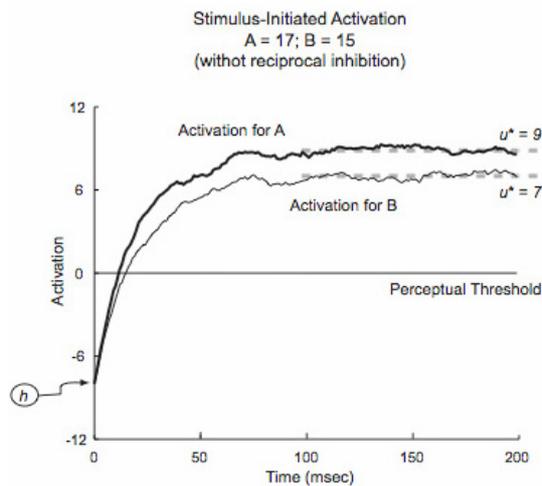


Fig. 6 When the stimulus-initiated activations for attributes A and B are unequal, and there are no activation-dependent inhibitory interactions between the detector populations responsive to these attributes, both would reach levels of activation that are above the threshold value required for perception, so both would be perceived simultaneously.

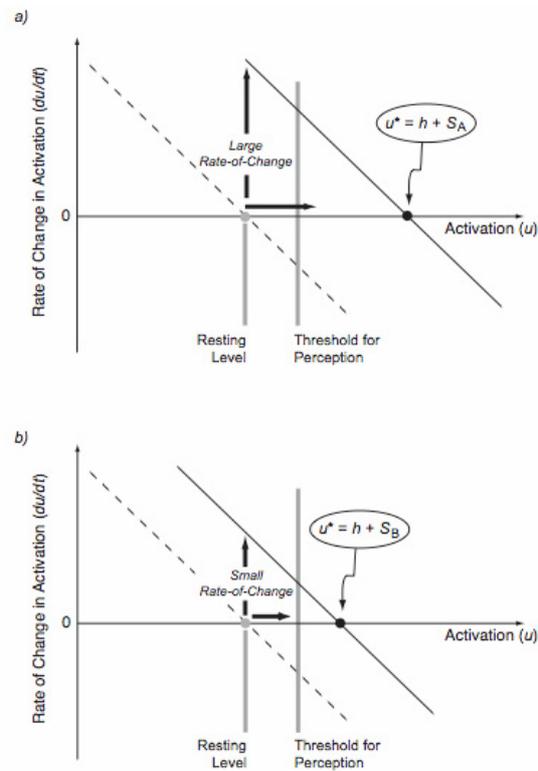


Fig. 7 Graphical representations for the stabilization of activation for non-interactive detector populations responsive to attributes A (panel a) and B (panel b), which differ with respect to their stimulus-initiated activation, as in Figure 6. Activation departs more quickly from the resting level (h) for the stronger (S_A) compared with the weaker attribute (S_B). This follows from the rate-of-change, du/dt , being larger for attribute A because its new attractor ($u^* = h + S_A$) is further from the initial, no-stimulus attractor ($u^* = h$) compared with the new attractor for B ($u^* = h + S_B$).

from the resting level result in detectors that preferentially respond to one attribute value having a momentary activational advantage over detectors that preferentially respond to the other attribute value, enabling the former to inhibit the latter more than vice versa (Figure 5). As activation increases, reciprocal, activation-dependent inhibition drives their activation levels further and further apart, with activation stabilizing at an above threshold level for one alternative and at a subthreshold level for the other. In this way, random fluctuations in activation lead to the perceptual selection of one of the equally stimulated alternatives. (In all the simulations, the inhibitory interaction has a maximum value of 11. With $S_A = S_B = 16$ for the simulation in Figure 5, the attractor for the above-threshold attractor is near $u^* = -8 + 16 = 8$, and the attractor for the subthreshold attractor is near $u^* = -8 + 16 - 11 = -3$.)

4.2 Rates of Change in Activation

We next examine selection between competing attribute values, A and B, when the detectors that respond preferentially to A are more strongly stimulated than the detectors that respond preferentially to B. If there were no activation-dependent inhibitory interactions between these detector populations, both would reach levels of activation that are above the threshold value required for perception. Both would be simultaneously perceived, which is contrary to true bistability. This is illustrated by trajectories (without interaction) for the evolution of activation toward the attractors for A, $u_A^* = h + S_A$, and for B, $u_B^* = h + S_B$, are indicated in Figure 6. (With $S_A = 17$ and $S_B = 15$, $u_A^* = -8+17 = 9$ and $u_B^* = -8+15 = 7$.) It can be seen in Figure 6 that upon stimulus presentation, the rate of change of activation as it rises from the no-stimulus resting state is greater for the detectors responding preferentially to A than for the detectors responding preferentially to B. This is shown in a different way by the graphical representations in Figure 7.

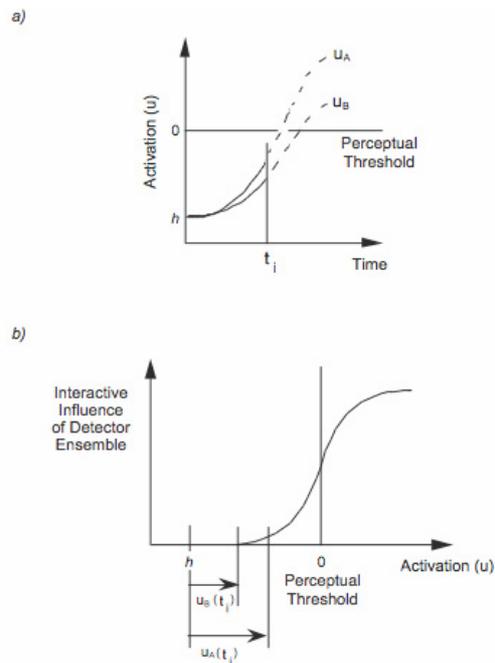


Fig. 8 When there are reciprocal, activation-dependent inhibitory interactions between detector populations, differences in the rate-of-change of activation (panel a) are of critical importance. As illustrated in panel b, this is because the detector population (for attribute A), with the faster rate-of-change, will reach an activation level where it will begin to inhibit its competing detector population (for attribute B), with a slower rate-of-change, before there can be an inhibitory influence in the opposite direction.

That is, activation for the stronger stimulus departs more quickly from the resting level, h , because the value of du/dt is greater when the new attractor is further from the resting level.

4.3 Perceptual Selection of the Favored Stimulus Alternative

As illustrated in Figure 8, the key to the perception of attribute value A when it is favored by the stimulus is that its detectors reach an activation level where they can begin inhibiting the activation of the detectors for attribute value B before the reverse occurs. This is why differences in the rate-of-change of activation are of critical importance. As a consequence of its faster rate-of-change, the activation for A increases to its above-threshold attractor value near $u^* = h + S_A = -8 + 17 = 9$, while at the same time activation for B is suppressed by inhibition from A, stabilizing below the perception threshold at attractor value near $u^* = h + S_B - I_A = -8 + 15 - 11 = -4$ (Figure 9). A comparison of Figures 6 and 9 shows that the reciprocal inhibitory interaction substantially separates the alternative activation states, facilitating the perceptual selection of the alternative that is most strongly specified by the stimulus. (Note: The simulations in this section and the remainder of the essay do not include the additional divergence in activation due to the self-excitation that occurs when the more strongly activated alternative passes through a detection instability.)

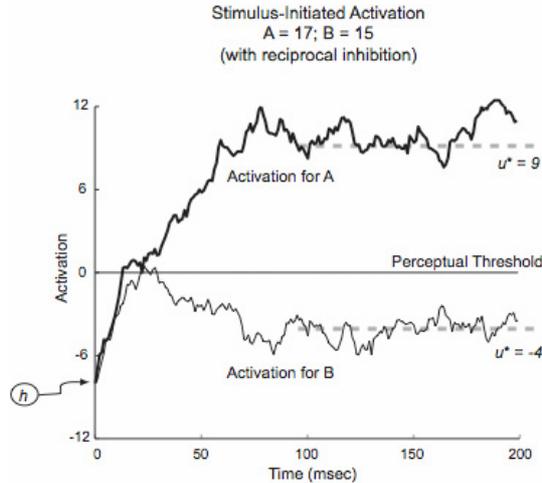


Fig. 9 The evolution of activation in the presence of reciprocal, activation-dependent inhibitory interactions when the more strongly activated detector population (for attribute A) suppresses the activation of its competitor, the detector population for attribute B.

4.4 Perceptual Selection of the Unfavored Stimulus Alternative

The detectors for A are more strongly stimulated than the detectors for B, so its perception is signified most of the time when the stimulus is presented. However, bistability means that B also can be perceived, though not necessarily as often as A. As illustrated by the activation trajectories in Figure 10, a sufficiently large fluctuation can result in B's activation becoming larger than A's as activation for both increase from the no-stimulus, resting level, h . B then can begin inhibiting its stimulus-favored competitor, A, before the reverse occurs. B's activation would then rise to an attractor value near $u_B^* = h + S_B = -8 + 15 = 7$, which lies above the threshold for perception, while A's activation decreases toward an attractor value near $u_A^* = h + S_A - I_B = -8 + 17 - 11 = -2$, which is subthreshold for perception. The perceptual selection of B is signified, but because this requires a random fluctuation that reverses the relative activation of the detectors for A and B as they rise from the resting level, it is perceived less often than A.

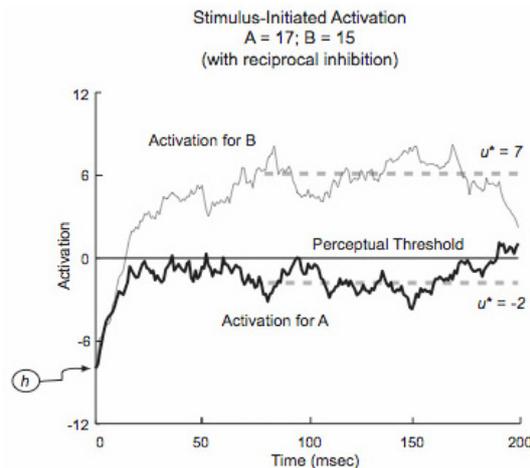


Fig. 10 The evolution of activation in the presence of reciprocal, activation-dependent inhibitory interactions when the more weakly activated detector population (for attribute B) gains a momentary activational advantage as a result of random fluctuations in activation, and suppresses the activation of its competitor, the detector population for attribute A.

5 Objects and Feature Integration

We have focused in this essay on the dynamical basis for the perception of individual attributes, arguing that it is determined by the pattern of activation over an ensemble of detectors responsive to different values of the same attribute. However, natural objects and even artificial objects created in the laboratory are multi-dimensional. The integration of the attributes (or features) belonging to the same object, called the

binding problem, has been the subject of numerous empirical and computational studies (see Treisman, 1998). The dynamical framework described in this essay suggests that short-range excitatory interactions among attribute/feature ensembles at similar retinal locations might be sufficient to account for their integration. Alternatively, higher-level object units or units in working memory might receive input from attribute ensembles at similar retinal locations, and in turn provide feedback to those ensembles that maintains precise spatial information and affects activation within each attribute ensemble in a manner consistent with the activated higher-level units (Johnson, Spencer & Schöner, 2008).

Evidence consistent with the latter alternative comes from a recent study with stimuli composed of four motion quartets, as illustrated in Figure 11 (Hock, Brownlow & Taler, in preparation). For this arrangement, the motion directions for the individual quartets are integrated into a global rotational pattern (alternating clockwise and counterclockwise rocking motion; the rotation of the outer elements predominates), likely due to the activation of global motion detectors in Area MSTd (Tanaka & Saito, 1989). Feedback from the global to motion detector ensembles with different directional preferences was indicated by the motion for the quartets being in directions consistent with global rotation. For example, horizontal motion for the quartets on the top and bottom of the configuration even though their aspect ratio would otherwise favor motion in vertical directions, as in Figure 1f. This kind of result is consistent with Treisman's (1998) prediction that "The strongest evidence will come when changes in neural activity are found to coincide with perceived changes in binding, perhaps in ambiguous figures or attentional capture (page 35)."

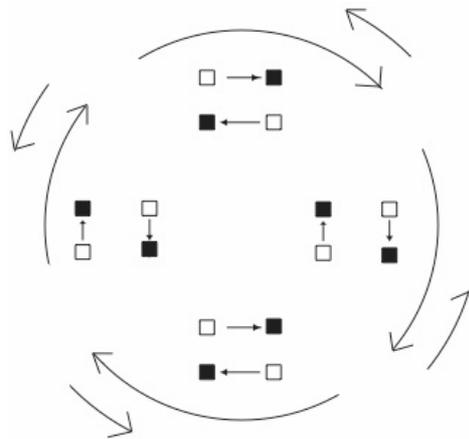


Fig. 11 Stimuli composed of four motion quartets arranged in a configuration for which the motion directions for the individual quartets can be integrated into a global rotational rocking pattern (alternating clockwise and counterclockwise rotation, with the perception of rotation predominating for the outer elements).

6 The Signature Features of Perceptual Dynamics

The stated purpose of this essay was to show that three neural properties are sufficient to provide the basis for a theory of perceptual dynamics that accounts for the relationship between persistence, responsiveness to stimulation, and selection. We have shown how perceptual persistence results from the intrinsic neural stabilization of individual detectors (through the balance of the membrane potential and ion concentrations; Section 2.1) and from the balance of excitatory/inhibitory interactions within a population of self-excited detectors with preferences for similar attribute values (Section 3.4). We have shown how change-resistant neural mechanisms affect the rate at which activation changes in response to changes in stimulation (Section 2.5), and how near-threshold perceptual uncertainty is minimized when stimulus-initiated activation is amplified by self-excitation (Section 3.3). Finally, we have shown how bistability arises as a result of reciprocal, activation-dependent inhibitory interactions between populations of detectors that preferentially respond to different values of the same attribute, and why differences in the rate-of-change activation are critical for the perceptual selection of the alternative percept that is most strongly specified by the stimulus (Sections 4.2 and 4.3). We show in this section that a perceptual dynamics based on these neural properties can result in the signature features of a dynamical system: spontaneous switching between percepts differing in relative stability, and hysteresis.

6.1 *Spontaneous Switching*

When two percepts are possible for the same stimulus, sufficiently large random perturbations can cause spontaneous switches between the percepts. The probability of a switch is inversely related to the likelihood of perceiving one of the alternatives when the bistable stimulus is presented. That is, the more likely it is for a percept to occur when a stimulus is presented, the less likely it is that there will be a spontaneous switch to the alternative percept.

This is illustrated by our earlier example in which stimulus-initiated activation is greater for the detectors responsible for the perception of A than for the detectors responsible for the perception of B. A is more likely to be perceived because it is more strongly specified by the stimulus, but sometimes B is perceived instead. Comparing Figures 9 and 10, it can be seen that the activation difference between the perceived and unperceived (subthreshold) alternatives is greater when the more likely of the alternatives is perceived (Figure 9) than when the less likely of the alternatives is perceived (Figure 10). A larger, less probable perturbation would be required to overcome the larger activation advantage of the perceived alternative when it is favored by the stimulus (A) compared with when it is not favored by the stimulus (B).

The relative stability of the two percepts is further illustrated in Figure 12 by continuing the simulation for a longer period of time, which provides more opportunity for the occurrence of random perturbations large enough to produce

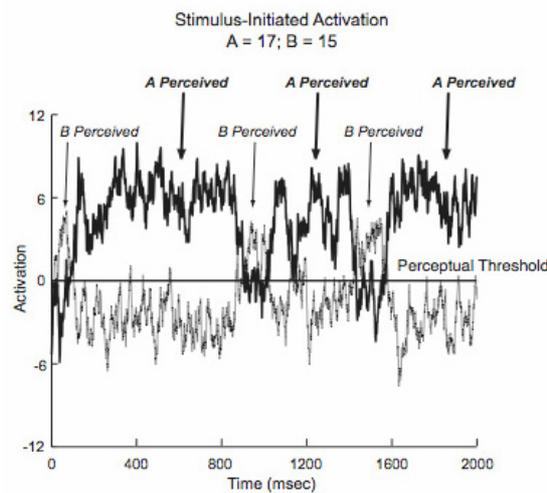


Fig. 12 A simulation demonstrating the relative stability of the percepts for which attributes A and B are perceived. Stimulus-initiated activation is greater for attribute A, so its perception is more stable than the perception of attribute B. This is indicated by the longer temporal intervals over which attribute A is perceived compared with the perception of attribute B.

perceptual switches (noise strength was increased in order to increase the frequency of switching). As is evident in Figure 12, both A and B are perceived, but the temporal intervals over which A is perceived are much longer than the temporal intervals over which B is perceived. The perception of A is more stable.

6.2 Hysteresis

Persistence (or its lack) in the presence of non-systematic, random events (passing shadows, distractors of attention, eye blinks, neural fluctuations) is observed dynamically in Section 6.1 as the dependence of spontaneous perceptual switching on the relative stability of the alternative percepts that are possible at a given moment. Such switching is rare in the natural environment because of the strong dominance of one percept, and because of the presence of disambiguating contextual information.

Persistence also facilitates the continuation of a previously established percept despite systematic changes in stimulus input, as might occur when the retinal projection of an object is distorted by gradual changes in viewpoint (due to the motion of an object or the egomotion of the perceiver). Objects in the world are invariant despite changes in viewpoint, so it is of obvious benefit to maintain the percept of an object that is established with less distorting projection angles.

To frame this in a manner consistent with our earlier examples, assume that perceiving the veridical shape of an object depends on the perception of attribute value A, but changes in viewpoint distort the retinal projection, favoring the

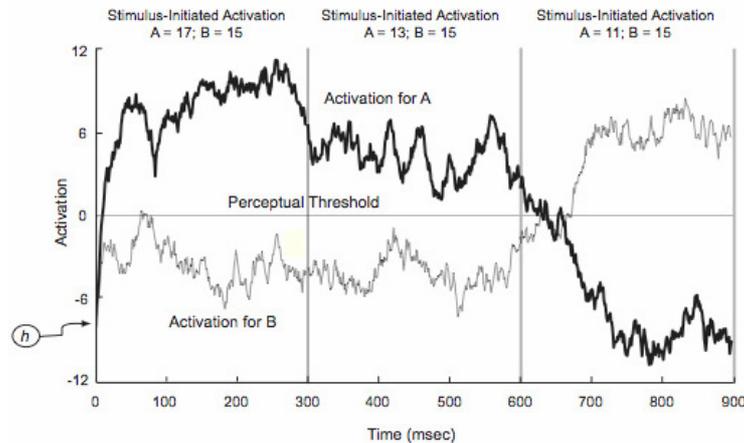


Fig. 13 A simulation demonstrating hysteresis. The perception of attribute A is established during the first 300 millisecond interval, when it is the stronger of the two attributes. Hysteresis is indicated by the perception of attribute A persisting during the second 300 millisecond interval despite the stimulus changing such that B becomes the stronger attribute. It is only during the final 300 millisecond interval, when the perception of attribute B is even more strongly favored by the stimulus, that there is a switch to its perception.

perception of attribute value B, and thus, a different, nonveridical shape. Hysteresis, the persistence of percept A despite systematic changes in the relative strength of attributes A and B, is illustrated in Figure 13. It can be seen in the figure that the perception of A, established when it is the stronger of the two attributes (during the first 300 millisecond interval of the simulation), is maintained despite the stimulus changing, and B becoming the stronger attribute (during the second 300 millisecond interval of the simulation). It is only when the advantage of stimulus-initiated activation more strongly favors attribute B that the initial perception of A gives way to the perception of B (during the final 300 millisecond interval of the simulation).

7 Conclusion

Over the years, there have been many physics- and mathematics-based entry points into the study of perceptual dynamics (e.g., Ditzinger & Haken, 1990; Poston & Stewart, 1978). As indicated earlier, Grossberg and his colleagues were the first to specifically study brain functions in terms of dynamics, building complex neural architectures in order to account for a wide range of perceptual and cognitive phenomena. (See also Wilson & Kim, 1994.) We have taken a different approach in this essay and in an earlier article (Hock, Schöner & Giese, 2002). That is, we have shown that a few basic neural properties are sufficient to provide the foundation for an understanding of the dynamical characteristics of

perception, characteristics often over-looked or "taken for granted" by most investigators of perceptual phenomena and perceptual behavior.

References

- Amari, S.: Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics* 27, 77–87 (1977)
- Baloch, A.A., Grossberg, S.: A neural model of high-level motion processing: Line motion and formation dynamics. *Vision Research* 37, 3037–3059 (1997)
- Bicho, E., Mallet, P., Schöner, G.: Target representation on an autonomous vehicle with low-level sensors. *The International Journal of Robotics Research* 19, 424–447 (2000)
- Braitenberg, V.: Cortical architectonics: General and areal. In: Brazier, M.A.B., Petsche, H. (eds.) *IBRO Monograph Series, Architectonics of the Cerebral Cortex*, vol. 3, pp. 443–466. Raven Press, New York (1978)
- Chey, J., Grossberg, S., Mingolla, E.: Neural dynamics of motion grouping: From aperture ambiguity to object speed and direction. *Journal of the Optical Society of America A* 14, 2570–2594 (1997)
- Ditzinger, T., Haken, H.: The impact of fluctuations on the recognition of ambiguous patterns. *Biological Cybernetics* 63, 453–456 (1990)
- Eastman, K., Hock, H.S.: Bistability in the perception of apparent motion: Effects of temporal asymmetry. *Perception & Psychophysics* 61, 1055–1065 (1999)
- Eggert, J., van Hemmen, J.L.: Modeling neuronal assemblies: Theory and implementation. *Neural Computation* 13, 1923–1974 (2001)
- Einstein, A.: On the movement of small particles suspended in a stationary liquid demanded by the molecular-kinetic theory of heat. *Annalen der Physik* 17, 540 (1905)
- Felleman, D.J., Van Essen, D.C.: Distributed hierarchical processing in primate visual cortex. *Cerebral Cortex* 1, 1–47 (1991)
- Francis, G., Grossberg, S.: Cortical dynamics of form and motion integration: Persistence, apparent motion, and illusory contours. *Vision Research* 36, 149–173 (1996)
- Gibson, J.J.: *The senses considered as perceptual systems*. Houghton Mifflin, New York (1966)
- Green, D.M., Swets, J.A.: *Signal detection theory and psychophysics*. Wiley, New York (1966)
- Grossberg, S.: Contour enhancement, short term memory, and constancies in reverberating neural networks. *Studies in Applied Mathematics* 52, 213–257 (1973)
- Grossberg, S., Mingolla, E.: Neural dynamics of motion perception: Direction fields, apertures, and resonant grouping. *Psychological Review* 92, 173–211 (1985)
- Grossberg, S., Rudd, M.E.: Cortical dynamics of visual motion perception: Short-range and long-range apparent motion. *Psychological Review* 99, 78–121 (1992)
- Hock, H.S., Balz, G.W.: Spatial scale dependent in-phase and anti-phase directional biases in the perception of self-organized motion patterns. *Vision Research* 34, 1843–1861 (1994)
- Hock, H.S., Balz, G.W., Smollon, W.: The effect of attentional spread on self-organized motion perception. *Vision Research* 38, 3743–3758 (1998)
- Hock, H.S., Brownlow, S., Taler, D.: The role of feedback in the formation of global motion patterns (in preparation)

- Hock, H.S., Bukowski, L., Nichols, D.F., Huisman, A., Rivera, M.: Dynamical vs. judgmental comparison: Hysteresis effects in motion perception. *Spatial Vision* 18, 317–335 (2005)
- Hock, H.S., Gilroy, L., Harnett, G.: Counter-changing luminance: A non-Fourier, non-attentional basis for the perception of single-element apparent motion. *Journal of Experimental Psychology: Human Perception and Performance* 28, 93–112 (2002)
- Hock, H.S., Kelso, J.A.S., Schöner, G.: Bistability and hysteresis in the organization of apparent motion patterns. *Journal of Experimental Psychology: Human Perception & Performance* 19, 63–80 (1993)
- Hock, H.S., Kogan, K., Espinoza, J.K.: Dynamic, state-dependent thresholds for the perception of single-element apparent motion: Bistability from local cooperativity. *Perception & Psychophysics* 59, 1077–1088 (1997)
- Hock, H.S., Nichols, D.F., Espinoza, J.: When motion is not perceived: Evidence from adaptation and dynamical stability. *Spatial Vision* 17, 235–248 (2004)
- Hock, H.S., Park, C., Schöner, G.: Self-organized pattern formation: Experimental dissection of motion detection and motion integration by variation of attentional spread. *Vision Research* 42, 991–1003 (2002)
- Hock, H.S., Ploeger, A.: Linking dynamical decisions at different levels of description in motion pattern formation: Psychophysics. *Perception & Psychophysics* 68, 505–514 (2006)
- Hock, H.S., Schöner, G., Giese, M.: The dynamical foundations of motion pattern formation: Stability, selective adaptation, and perceptual continuity. *Perception & Psychophysics* 65, 429–457 (2003)
- Hock, H.S., Schöner, G., Hochstein, S.: Perceptual stability and the selective adaptation of perceived and unperceived motion. *Vision Research* 36, 3311–3323 (1996)
- Hock, H.S., Schöner, G., Voss, A.: The influence of adaptation and stochastic fluctuations on spontaneous perceptual changes for bistable stimuli. *Perception & Psychophysics* 59, 509–522 (1997)
- Jancke, D., Erlhagen, W., Dinse, H.R., Akhavan, A.C., Giese, M., Steinhage, A., Schöner, G.: Parametric population representation of retinal location: Neuronal interaction dynamics in cat primary visual cortex. *Journal of Neuroscience* 19, 9016–9028 (1999)
- Johnson, J.S., Spencer, J.P., Schöner, G.: Moving to higher ground: The dynamic field theory and the dynamics of visual cognition. *New Ideas in Psychology* 26, 227–251 (2008)
- Leopold, D.A., Wilke, M., Maier, A., Logothetis, N.K.: Stable perception of visually ambiguous patterns. *Nature Neuroscience* 5, 605–609 (2002)
- Nichols, D.F., Hock, H.S., Schöner, G.: Linking dynamical decisions at different levels of description in motion pattern formation: Computational simulations. *Perception & Psychophysics* 68, 515–533 (2006)
- Poston, T., Stewart, I.: Nonlinear modeling of multistable perception. *Behavioral Science* 23, 318–334 (1978)
- Schneegans, S., Schöner, G.: Dynamic field theory as a framework for understanding embodied cognition. In: Calvo, P., Gomila, T. (eds.) *Handbook of Cognitive Science: An Embodied Approach*, pp. 241–271 (2008)
- Schöner, G.: Dynamical systems approaches to cognition. In: Sun, R. (ed.) *Cambridge handbook of computational cognitive modeling*, pp. 101–126. Cambridge University Press, Cambridge (2008)

- Tanaka, K., Saito, H.-A.: Analysis of motion of the visual field by direction, expansion/contraction, and rotation cells clustered in the dorsal part of the medial superior temporal area of the macaque monkey. *Journal of Neurophysiology* 62, 643–656 (1989)
- Trappenberg, T.P.: *Fundamentals of computational neuroscience*. Oxford University Press, Oxford (2002)
- Treisman, A.: The binding problem. In: Squire, L.R., Kosslyn, S.M. (eds.) *Findings and current opinion in cognitive neuroscience*. The MIT Press, Cambridge (1998)
- Ullman, S.: *The interpretation of visual motion*. MIT Press, Cambridge (1979)
- Wilson, H.R., Cowan, J.D.: A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik* 13, 55–80 (1973)
- Wilson, H.R., Kim, J.: A model for motion coherence and transparency. *Visual Neuroscience* 11, 1205–1220 (1994)
- Yuille, A.L., Kersten, D.: Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Science* 10, 301–308 (2006)