The Broadbent Lecture at the XVth Congress of the European Society for Cognitive Psychology, Marseilles, France, 2007

Claus Bundesen[a]

[a] University of Copenhagen, Copenhagen, Denmark

First published on: 01 March 2008

PLEASE SCROLL DOWN FOR ARTICLE

Ψ **Psychology Press**
Taylor & Francis Group

The Broadbent Lecture at the XVth Congress of the European Society for Cognitive Psychology, Marseilles, France, 2007

The nature of visual attention: Formulas bridging cognition and neurophysiology

Claus Bundesen

University of Copenhagen, Copenhagen, Denmark

The neural theory of visual attention (NTVA) introduced by Bundesen, Habekost, and Kyllingsbæk (2005) is presented. NTVA is a neural interpretation of Bundesen's (1990) formal theory of visual attention (TVA). In NTVA, *filtering* (selection of *objects*) changes the number of cortical neurons in which an object is represented so that this number increases with the behavioural importance of the object (reallocation of processing resources by dynamic remapping of receptive fields). Another fundamental mechanism, *pigeonholing* (selection of *features*), scales the level of activation in individual neurons coding for a particular feature. The theory accounts both for a wide range of attentional effects in human performance (reaction times and error rates) and for a wide range of effects observed in firing rates of single cells in the primate visual system. NTVA provides a mathematical framework to unify the two fields of research—formulas bridging cognition and neurophysiology.

In this paper I present the Neural Theory of Visual Attention (NTVA) introduced by Bundesen, Habekost, and Kyllingsbæk (2005; see also Bundesen & Habekost, in press). By use of the same basic equations, NTVA accounts for a wide range of attentional effects in human

performance (reaction times and error rates) and a wide range of effects observed in firing rates of single cells in the visual system of primates. NTVA provides a mathematical framework to unify the two fields of research—formulas bridging cognition and neurophysiology.

NTVA is a further development of the Theory of Visual Attention (TVA) proposed by Bundesen (1990). TVA is a formal, computational theory that accounts for a wide range of attentional effects in mind and behaviour reported in the psychological literature. At the heart of TVA are two equations, and NTVA is a neural interpretation of these equations.

The central equations of TVA jointly describe two mechanisms of attentional selection: one for selection of objects and one for selection of features or categories. As a tribute to Broadbent (1971), I use the term *filtering* for selection of objects and the term *pigeonholing* for selection of features or categories. Selection of a feature is the same as selection of a category, namely, a category formed by all those objects that possess the given feature.

In NTVA, filtering affects the number of cortical neurons in which an object is represented, whereas pigeonholing is a multiplicative scaling of the level of activation in neurons coding for particular features (see Figure 1). The total activation representing a visual categorisation of the form "object x has feature i" is directly proportional to both the number of neurons representing the categorisation, which is controlled by filtering, and the level of activation of the individual neurons representing the categorisation, which is controlled by pigeonholing. The *rate equation* of TVA essentially expresses this fact.

Filtering is done in such a way that the number of cells in which an object is represented increases with the behavioural importance of the object. Thus, processing is parallel with differential allocation of resources so that important objects are represented in many cells. More specifically, the probability that a cortical neuron represents a particular object within its classical receptive field equals the attentional weight of the object divided by the sum of the attentional weights across all objects in the receptive field.

The *weight equation* of TVA describes how attentional weights are computed, and logically this computation must occur before processing resources (cells) can be distributed in accordance with the weights. Accordingly, in NTVA, a normal perceptual cycle consists of two waves: a wave of unselective processing, in which attentional weights are computed, followed by a wave of selective processing, when processing resources have been allocated according to the attentional weights. During the first wave, cortical processing resources are distributed at random (unselectively) across the visual field. At the end of the first wave, an attentional weight has been
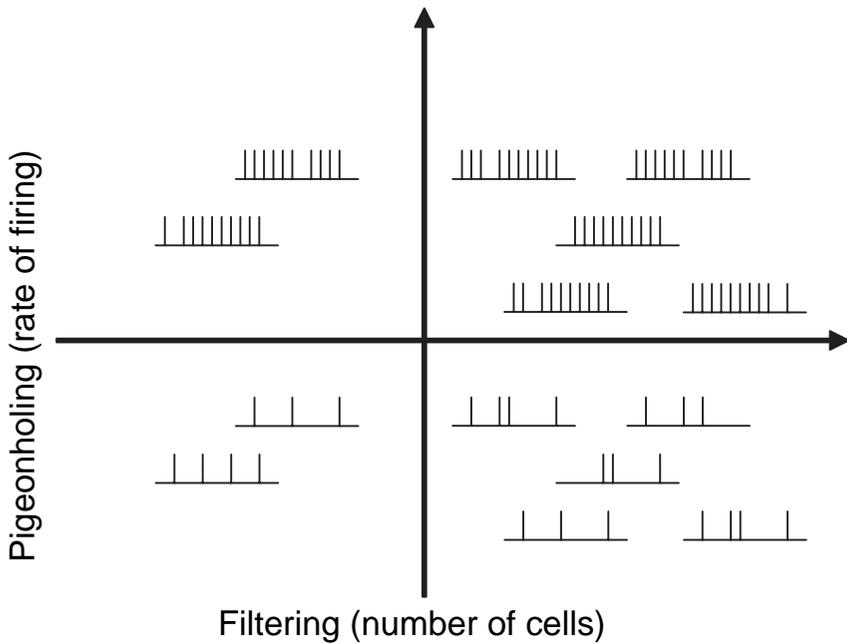
Figure 1. Attentional selection in NTVA: Combined effects of filtering (selection of objects) and pigeonholing (selection of features) on the set of cortical spike trains representing a particular visual categorisation of the form "object *x* has feature *i*". The four conditions (quadrants) correspond to the factorial combinations of 2 levels of filtering (weak vs. strong support to object *x*) ×2 levels of pigeonholing (weak vs. strong support to feature *i*). Filtering changes the number of cortical neurons in which an object is represented. Pigeonholing changes the rate of firing of cortical neurons coding for a particular feature. From "A Neural Theory of Visual Attention: Bridging Cognition and Neurophysiology", by C. Bundesen, T. Habekost, and S. Kyllingsbæk, 2005, *Psychological Review, 112*, p. 292. Copyright 2005 by the American Psychological Association. Adapted with permission.

computed for each object in the visual field and the weight has been stored in a saliency map. The weights are used for reallocation of attention (visual processing capacity) by dynamic remapping of receptive fields of cortical neurons so that the number of neurons allocated to an object increases with the attentional weight of the object. Thus, during the second wave, cortical processing is selective in the sense that the amount of processing resources allocated to an object (the number of neurons in which the object is represented) varies with the attentional weight of the object. As more processing resources are devoted to behaviourally important objects than less important ones, the important objects are more likely to become encoded into visual short-term memory (VSTM). The VSTM system is

conceived as a feedback mechanism that sustains activity in the neurons that have won the attentional competition.

The rest of the paper falls in three parts. First, TVA is presented as a formal theory of visual attention. Then the neural interpretation of TVA, NTVA, is presented. Finally, a few examples are given of how NTVA has been applied to single-cell studies.

A FORMAL THEORY OF VISUAL ATTENTION (TVA)

Basic assumptions

In TVA, both visual identification and selection of objects in the visual field consist in making visual categorisations. A visual categorisation has the form "object x has feature i", or equivalently, "object x belongs to category i". A visual categorisation is *made* when the categorisation is encoded into visual short-term memory (VSTM). When one makes the visual categorisation that x belongs to i, object x is said to be selected and object x is also said to be identified as a member of category i.

When a visual categorisation of an object completes processing, the categorisation enters VSTM, provided that memory space for the categorisation is available in VSTM. The capacity of VSTM is limited to K different objects, where K is about 4 (cf. Luck & Vogel, 1997).

Clearing VSTM effectively starts a race among objects in the visual field to become encoded into VSTM. An object is encoded in VSTM if and when any categorisation of the object is encoded in VSTM, so each object x is represented in the encoding race by all possible categorisations of the object.

Rate equation.   By the rate equation of TVA, the rate, $v(x, i)$, at which a particular visual categorisation, "x belongs to i", is encoded into VSTM is given by a product of three terms:

$$v(x,i) = \eta(x,i)\beta_i \frac{w_x}{\sum_{z \in S} w_z} \tag{1}$$

The first term, $\eta(x, i)$, is the strength of the sensory evidence that x belongs to category i. The second term, β_i, is a perceptual decision bias associated with category $i (0 \le \beta_i \le 1)$. The third term is the relative attentional weight of object x—that is, the weight of object x, w_x, divided by the sum of weights across all objects in the visual field, S.

Technical note.   In the original presentation of TVA (Bundesen, 1990), I interpreted $v(x, i)$ as the rate of point events generated by a Poisson process

and I also assumed that a single point event suffices for encoding a categorisation into VSTM. (Here, a point event may be interpreted as a single neural spike or as a single volley of r or more synchronised neural spikes.) In this case, $v(x, i)$ equals the probability density that the categorisation "x belongs to i" completes processing at a certain point in time given that the categorisation has not completed processing before that time. In later articles in the *Psychological Review*, Logan (1996, 2002) also interpreted $v(x, i)$ as the event rate of a Poisson process but explored the implications of assuming that, depending on threshold settings, many point events from the Poisson process representing a particular categorisation may be needed for encoding the categorisation into VSTM.

Weight equation.   The attentional weights in the rate equation of TVA are derived from *pertinence* values. Every visual category is supposed to have a certain pertinence. The pertinence of a category is a measure of the current importance of attending to objects that belong to the category. The weight of an object x in the visual field is given by the weight equation of TVA,

$$w_x = \sum_{j \in R} \eta(x, j)\pi_j \tag{2}$$

where R is the set of all visual categories, $\eta(x, j)$ is the strength of the sensory evidence that object x belongs to category j, and π_j is the pertinence of category j. By Equation 2, the attentional weight of an object is a weighted sum of pertinence values. The pertinence of a given category enters the sum with a weight equal to the strength of the sensory evidence that the object belongs to the category.

Mechanisms of selection

The rate and weight equations jointly describe two mechanisms of selection: a mechanism for selection of objects (filtering) and a mechanism for selection of categories (pigeonholing). The *filtering* mechanism is represented by pertinence values and attentional weights. As an example, if selection of red objects is wanted, the pertinence of *red* should be high. The weight equation implies that, when *red* has a high pertinence, red objects get high attentional weights. Accordingly, by the rate equation, processing of red objects is fast, so red objects are likely to win the processing race and be encoded into VSTM.

The *pigeonholing* mechanism is represented by perceptual decision bias parameters. Pertinence values determine which objects are selected (*filtering*), but perceptual biases determine how the objects are categorised (*pigeonholing*). If particular types of categorisations must be reported or otherwise responded to, the bias values of the corresponding categories

should be high. By Equation 1, then, the desired types of categorisations are likely to be made.

When the selection system is coupled to a sensory system that supplies appropriate η values, and when pertinence and bias parameters have been set, both filtering and pigeonholing are accomplished by a race between visual categorisation processes whose rate parameters are determined through the simple algebraic operations of the rate and weight equations. Thus, the theory yields a truly computational account of selective attention in vision.

Applications

TVA has been applied to findings from a broad range of paradigms concerned with single-stimulus identification and selection from multi-element displays. In addition, TVA has been applied in studies of attention deficits after brain damage, and the scope of the theory has been extended to other cognitive domains such as memory and executive control.

Single-stimulus identification.   For single-stimulus identification, TVA provides a mathematical derivation of the classical biased-choice model of Luce (1963; see Bundesen, 1990; also see Bundesen, 1993). The biased-choice model has been successful in explaining many experimental findings on effects of visual discriminability and bias (see, e.g., Townsend & Ashby, 1982; Townsend & Landon, 1982).

Selection from multielement displays.   For selection from multielement displays, TVA provides a derivation of the Fixed-capacity Independent Race Model (FIRM) of Shibuya and Bundesen (1988). Figure 2 shows a fit to the data of Shibuya and Bundesen: probability distributions of partial and whole report scores with varying numbers of targets and distractors and varying exposure durations. Other well-established findings closely fitted by TVA include effects of object integrality in selective report (e.g., Duncan, 1984), effects of the spatial position of targets in studies of divided attention (e.g., Posner, Nissen, & Ogden, 1978; Sperling, 1967), effects of selection criterion and effects of the number of distractors in studies of focused attention (e.g., Bundesen & Pedersen, 1983; Estes & Taylor, 1964; Treisman & Gelade, 1980; Treisman & Gormican, 1988), and effects of consistent practice in search (Schneider & Fisk, 1982; also see Kyllingsbæk, Schneider, & Bundesen, 2001).

Attention deficits after brain damage.   Recently, the principles of TVA have been extensively applied in the study of attention deficits after brain
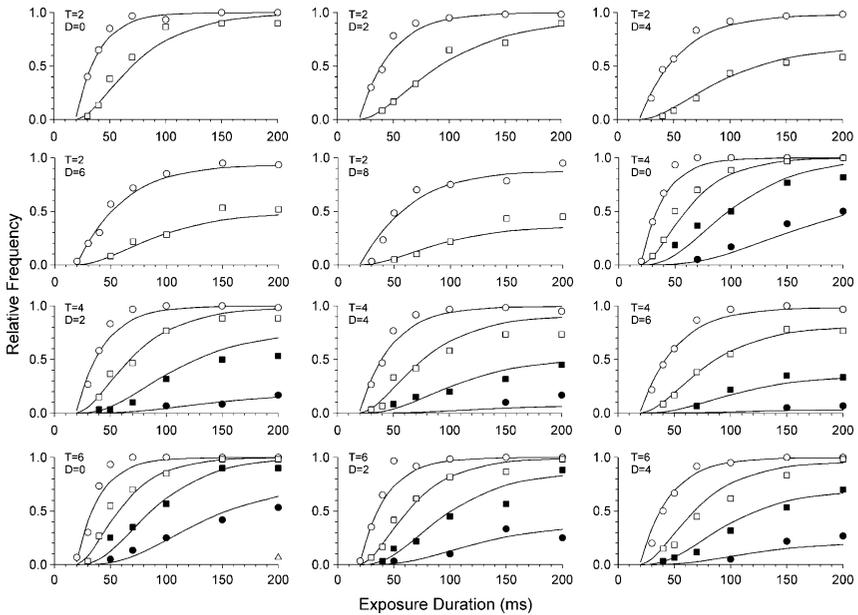
Figure 2. Relative frequency of scores of *j* or more (correctly reported targets) as a function of exposure duration with *j*, number of targets *T*, and number of distractors *D* as parameters in partial report experiment of Shibuya and Bundesen (1988). Data are shown for Subject MP. Parameter *j* varies within panels; *j* is 1 (open circles), 2 (open squares), 3 (solid squares), 4 (solid circles), or 5 (triangle). *T* and *D* vary among panels. Smooth curves represent a theoretical fit to the data by the FIRM model. For clarity, observed frequencies less than .02 were omitted from the figure. Adapted from "Visual Selection from Multielement Displays: Measuring and Modeling Effects of Exposure Duration", by H. Shibuya and C. Bundesen, 1988, *Journal of Experimental Psychology: Human Perception and Performance*, *14*, p. 595. Copyright 1988 by the American Psychological Association.

damage. Duncan et al. (1999) showed how analysis in terms of parameters defined by TVA enables a very specific measurement of attentional deficits in visual neglect patients. TVA-based assessment has also been used in studies of simultanagnosia (Duncan et al., 2003), Huntington's disease (Finke, Bublak, Dose, Müller, & Schneider, 2006), Alzheimer's disease (Bublak, Redel, & Finke, 2006), alexia (Habekost & Starrfelt, 2006), and subclinical attention deficits (Habekost & Bundesen, 2003). Currently, research groups in England, Ireland, France, Germany, and Denmark are extending these investigations to other patient groups (e.g., ADHD, dyslexia, depression, schizophrenia, and age-related changes in white matter).

Other cognitive domains. Logan (1996) proposed an extension of TVA, the CODE theory of Visual Attention (CTVA), which combines TVA with

the COntour DEtector theory of perceptual grouping by proximity (van Oeffelen & Vos, 1982, 1983). CTVA explains a wide range of spatial effects in visual attention (see Logan, 1996; Logan & Bundesen, 1996).

Logan and Gordon (2001) extended CTVA into a theory of executive control in dual-task situations that accounts for crosstalk, set-switching cost, and concurrence costs. The theory, ECTVA, assumes that executive processes control subordinate processes by manipulating their parameters. TVA is used as the theory of subordinate processes, so a task set is defined as a set of TVA parameters that is sufficient to configure TVA to perform a task. Set switching is viewed as a change in one or more of these parameters, and the time taken to change a task set is assumed to depend on the number of parameters to be changed.

Logan (2002) proposed an Instance Theory of Attention and Memory (ITAM) that combines ECTVA with the exemplar-based random walk model of categorisation proposed by Nosofsky and Palmeri (1997). The exemplar-based random walk model itself is a combination of Nosofsky's (1986) generalised context model of categorisation and Logan's (1988) instance theory of automaticity. By integrating theories of attention, categorisation, and memory, the development of ITAM seems to be an important step toward a unified account of visual cognition.

A NEURAL INTERPRETATION OF TVA (NTVA)

NTVA gives the equations of TVA an interpretation at the level of individual neurons. A typical neuron in the visual system is assumed, first, to be specialised to represent a single feature and, second, to respond to the properties of only one object at any given time. The feature for which the neuron is specialised can be a more or less simple "physical feature" or a "microfeature" in a distributed representation. The object selection of the neuron occurs by dynamic remapping of the cell's receptive field so that the functional receptive field contracts around the selected object. The remapping is done so that the probability that the neuron comes to represent a particular object equals the attentional weight of the object divided by the sum of the attentional weights of all objects in the classical receptive field.

The *activation* of a neuron by the appearance of an object in its receptive field is defined as the increase in firing rate above a baseline rate representing the undriven activity of the neuron. If the baseline rate is zero, the activation is just the firing rate. Independently of the distribution of neurons among objects (i.e., independently of filtering), the activation in neurons representing particular features (say, the activation in the set of neurons representing feature i) is scaled up or down (pigeonholing). The rate equation of TVA,

$$v(x, i) = \eta(x, i)\beta_i \frac{w_x}{\sum\limits_{z \in S} w_z},$$

describes the combined effects of filtering and pigeonholing on the total activation of the population of neurons representing the categorisation "object x has feature i". The v value on the left-hand side of the rate equation is the total activation of the neurons that represent the categorisation at the level of processing where objects compete for entrance into VSTM. At this level of processing, the classical receptive fields of neurons are assumed to be so large that each one covers the entire visual field. On the right-hand side of the equation, both the bias for assigning objects to category i, β_i, and the relative attentional weight of object x,

$$\frac{w_x}{\sum\limits_{z \in S} w_z},$$

range between 0 and 1, so the strength of the sensory evidence that x is an i, $\eta(x, i)$, equals the highest possible value of $v(x, i)$. Thus, $\eta(x, i)$ equals the total activation of the set of all neurons coding feature i when all of these represent object x (say, when x is the only object in the visual field) and when the featural bias in favour of i is maximal (i.e., $\beta_i = 1$). When the proportion of feature-i coding neurons representing object x,

$$\frac{w_x}{\sum\limits_{z \in S} w_z},$$

is less than 1, then the total activation representing the categorisation "object x has feature i" is scaled down by multiplication with this factor on the right-hand side of the equation. The total activation representing the categorisation depends not only on the number of neurons representing the categorisation but also on the level of activation of the individual neurons representing the categorisation. The bias parameter β_i is a scale factor that multiplies activations of all individual feature-i coding neurons, so the total activation representing the categorisation "object x has feature i" is also directly proportional to β_i.

Thus, in the neural interpretation we propose for the rate equation of TVA, the total activation representing the categorisation "object x has feature i" is directly proportional to both the number of neurons representing the categorisation and the level of activation of the individual neurons representing the categorisation. The number of neurons is controlled by the relative attentional weight of x (filtering), but the activation of the individual neurons is controlled by β_i (pigeonholing).
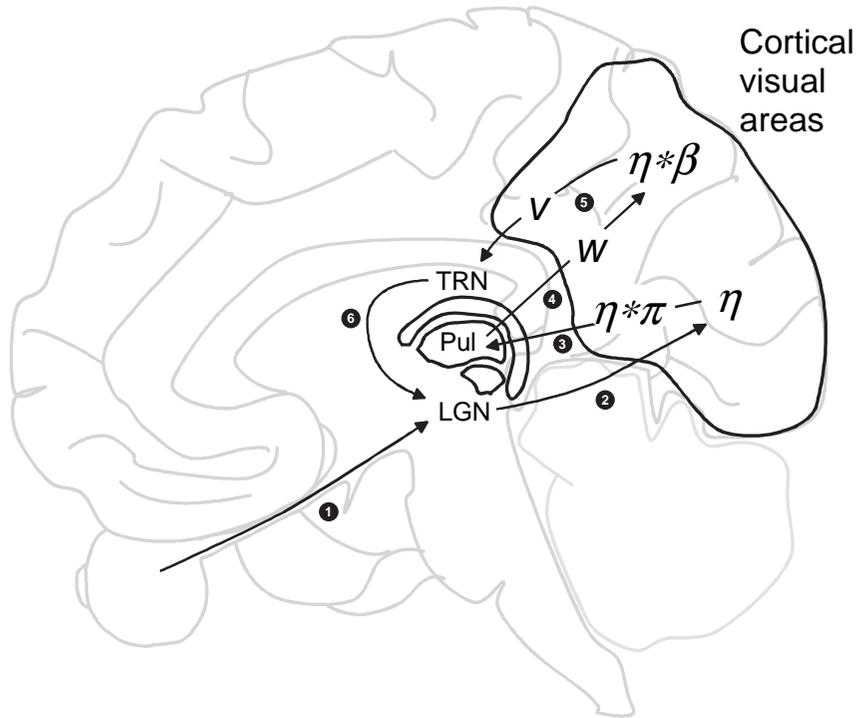
Figure 3. Possible distribution of visual processing across the human brain. Visual information from the eye enters the lateral geniculate nucleus (LGN) of the thalamus (1) and is transmitted to striate and extrastriate cortical areas where η values (strengths of evidence that objects at particular scales and positions have particular features) are computed (2). The η values are multiplied by π (pertinence) values, and the products are transmitted from the cortex to a saliency map in the pulvinar (Pul) nucleus of the thalamus, where the products are summed up as attentional weights of the stimulus objects (3). After the first (unselective) wave of processing, cortical processing capacity is redistributed by means of attentional weight signals (w) from the pulvinar to the cortex, so that during the second (selective) wave of processing, objects with high attentional weights are processed by many neurons (4). The resulting η values are multiplied by β (bias) values, and the products are transmitted from the cortex to a multiscale VSTM map of locations, which is tentatively localised in the thalamic reticular nucleus (TRN) (5). When the VSTM map is initialised, objects in the visual field effectively start a race to become encoded into VSTM. In this race, each object is represented by all possible categorisations of the object, and each possible categorisation participates with an activation (v value) proportional to the corresponding η value multiplied by the corresponding β value. For the winners of the race, the TRN gates activation representing a categorisation back to some of those cells in LGN whose activation supported the categorisation (6). Thus, activity in neurons representing winners of the race is sustained by positive feedback. From "A Neural Theory of Visual Attention: Bridging Cognition and Neurophysiology", by C. Bundesen, T. Habekost, and S. Kyllingsbæk, 2005, *Psychological Review, 112*, p. 295. Copyright 2005 by the American Psychological Association. Adapted with permission.

NTVA does not depend critically on a particular anatomical localisation of the proposed computations. However, we have suggested some plausible ways in which visual processing may be distributed across the human brain. One possibility is illustrated in Figure 3.

APPLICATIONS TO SINGLE-CELL STUDIES

Neural filtering

NTVA's notion of attentional filtering was originally inspired by a study of Moran and Desimone (1985). Moran and Desimone presented macaque monkeys with stimuli at two locations. The monkeys were trained to attend only to stimuli presented at one of the locations in the display and ignore any stimuli shown at the other location. The monkeys performed a match-to-sample task, in which they encoded a sample stimulus that was shown at the attended location, retained it for a delay period, and then matched it to a test stimulus shown at the same location. During the presentation of both sample and test displays, Moran and Desimone recorded the responses of single neurons to the stimuli. They found that, when the target and distractor stimuli were both present at the same time within the RF of the recorded cell in visual areas V4 or IT, the cell's rate of firing depended on the properties of the target rather than the properties of the distractor. For example, they recorded the response of a cell to a pair of stimuli consisting of (a) a stimulus that elicited a high rate of firing in the cell when the stimulus was presented alone (an effective sensory stimulus) and (b) a stimulus that had little or no effect on the rate of firing in the cell when the stimulus was presented alone (an ineffective sensory stimulus). On trials in which the effective sensory stimulus was the target and the ineffective sensory stimulus was the distractor, the cell showed a high rate of firing. However, on trials in which the ineffective sensory stimulus was the target and the effective sensory stimulus was a distractor, the cell showed a low rate of firing. Moran and Desimone remarked that the typical cell responded "as if the RF had contracted around the attended stimulus".

The effect depended on both the target and the distractor being located within the recorded neuron's RF. However, in IT the RFs were very large, covering at least the central 12 degrees of both the contralateral and ipsilateral field. For these neurons, the distractors were always located inside the RF. On the contrary, in V1 only one stimulus could be fitted into each neuron's RF. In this case, no significant effects of attention were found. Also in V4, no effect of attention was found when only a single stimulus was placed in the RF. These negative findings fit with the predictions of NTVA. Overall, the results of Moran and Desimone clearly support NTVA's notion of filtering, and in fact suggested the interpretation in the first place.

Reynolds, Chelazzi, and Desimone (1999) corroborated and extended the findings of Moran and Desimone (1985). Recording from neurons in monkey areas V2 and V4, they found that the mean firing rate of a cell to a pair of stimuli in its RF approximated a weighted average of the firing rates of the cell to each of the stimuli in the pair when they were presented alone. When the monkey's attention was directed to one of the stimuli in the pair, this increased the weight on the target stimulus so that the mean response of the neuron was driven (up or down) towards the response elicited when the target stimulus was presented alone. The results fit in with the dual conjectures that (a) at any time, a cell was driven by only one of the stimuli in its RF, and (b) the probability that the cell was driven by a given stimulus was proportional to the attentional weight of the stimulus.

Neural pigeonholing

Whereas the filtering mechanism of selection affects the number of neurons in which an object x is represented, the *pigeonholing* mechanism of selection affects the way in which the element is represented in each of those neurons that are allocated to the object. The activation in a neuron representing the categorisation that object x has feature i should be proportional to the multiplicative bias, β_i, which is applied to neurons preferring stimuli with feature i. Recent recordings from single cells in the visual system of monkeys have provided evidence of such a mechanism of selection: a *feature-based* mechanism of attention that selects groups of neurons with similar stimulus preferences for a multiplicative enhancement in response strength.

Treue and Martinez-Trujillo (1999; also see Martinez-Trujillo & Treue, 2004) recorded from area MT, which contains cells that are selective to direction of motion. In the basic task, Treue and Martinez-Trujillo presented monkeys with two coherently moving random dot patterns, one placed inside the RF of the neuron being recorded and the other one in the opposite visual hemifield (see Figure 4). At the start of each trial, the monkey was shown a cue at one of the locations. Following this, the random dot patterns appeared and the monkey was required to detect small changes in the speed or direction of movement of the pattern at the cued location. These changes occurred after a random delay ranging between a few hundred milliseconds and several seconds.

In one of their experiments, Treue and Martinez-Trujillo demonstrated an effect of pigeonholing with respect to a given direction of movement ("feature-based attention"); probably the first clear demonstration of pigeonholing in the single-cell literature. In this experiment, the recorded neuron's RF was stimulated by a pattern moving in the direction preferred by the neuron. Outside the RF, a pattern was presented that moved either in
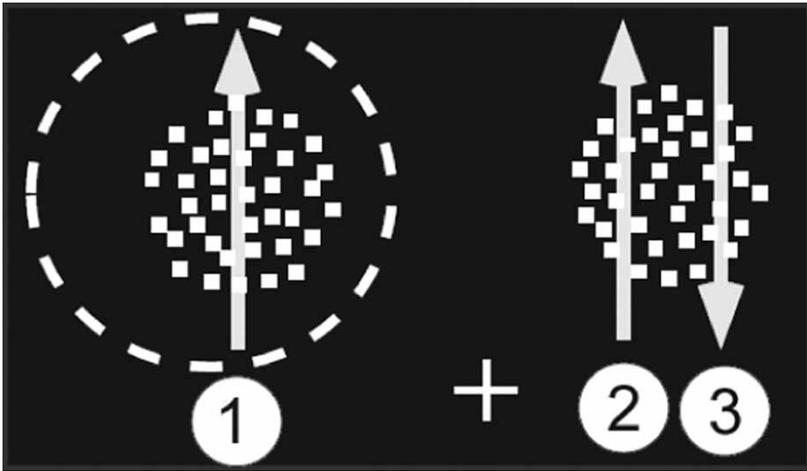
Figure 4. Stimuli used by Treue and Martinez-Trujillo (1999). The random dot pattern inside the receptive field (dashed circle) always moved in the cell's preferred direction (upwards pointing arrow, (1)); the stimulus outside moved in either the same (2) or the opposite direction (3). From "Feature-Based Attention Influences Motion Processing Gain in Macaque Visual Cortex", by S. Treue and J. C. M. Martinez-Trujillo, 1999, *Nature*, *399*, p. 577. Copyright 1999 by Nature Publishing Group. Adapted with permission.

the same direction or in the opposite direction. When the monkey attended to the pattern outside the RF, the response of the recorded neuron varied with the direction of movement being attended. The firing rate went up when the attended pattern (outside the RF of the recorded neuron) moved in the preferred direction of the recorded neuron. The firing rate went down when the attended pattern (outside the RF of the recorded neuron) moved in the direction opposite to the preferred direction of the recorded neuron. The spatial location of the attended pattern was exactly the same in the two conditions. Thus, a nonspatial, feature-based mechanism of attention seemed to be at work: pigeonholing.

The result is readily explained by assuming that when the monkey attended to movement in a particular direction, the β value for movement in that direction was high while the β value for movement in the opposite direction was low. As the monkey was monitoring the pattern for hundreds to thousands of milliseconds, there should be plenty of time to adjust the β values in accordance with the display. Consider the situation in which the preferred movement of the recorded neuron was upwards but the monkey attended to movement in the opposite direction (downwards), trying to detect a small change in the speed or direction of movement of the pattern to be attended. Since the monkey had learned that only small changes in the direction of movement of the target would occur, and the target was moving
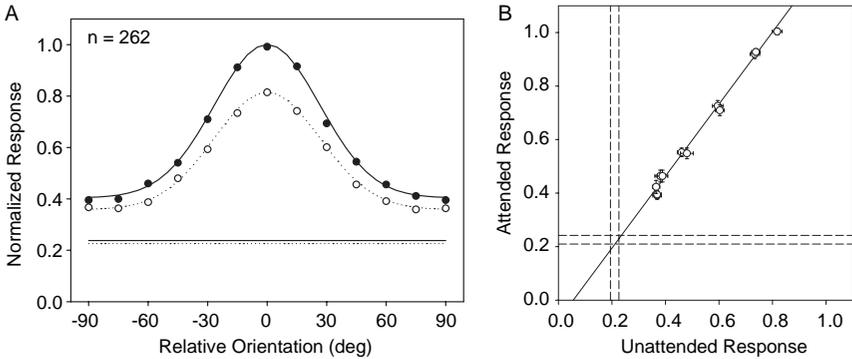
Figure 5. Effects of attention on mean rates of firing in the experiment of McAdams and Maunsell (1999). A: Normalised population tuning curves for all V4 neurons. Solid circles fitted by a solid Gaussian curve show the normalised response as a function of the angular deviation between the stimulus and the preferred orientation when the monkey was attending to the stimulus inside the receptive field. The solid horizontal line represents the undriven activity, measured as the mean firing rate during the fixation period before stimulus presentation in the same attention condition. Corresponding data for the condition in which the monkey was attending to the stimulus outside the receptive field are shown by open circles, the dotted Gaussian curve, and the dotted horizontal line. B: The attended response versus the unattended response for each of the tested orientations. The results are fitted by a least-squares line with a slope of 1.32. The pairs of dashed lines show undriven activity ± 1 SE. The close fit and the finding that the line very nearly passes through the point where attended response = unattended response = undriven activity show that attention very nearly effected a multiplicative scaling of the activation (the total firing rate minus the level of undriven activity). From "Effects of Attention on Orientation-Tuning Functions of Single Neurons in Macaque Cortical Area V4", by C. J. McAdams and J. H. R. Maunsell, 1999, *Journal of Neuroscience*, *19*, p. 437. Copyright 1998 by the Society for Neuroscience. Adapted with permission.

downwards, $\beta_{upwards}$ should be low (perceptual bias should generally reflect expectations). Because the activation of the recorded cell should be proportional to $\beta_{upwards}$, the recorded activation also should be low. By contrast, when the monkey attended to movement in the preferred direction of the recorded cell, β values should be high for upwards and nearby directions. In this case, $\beta_{upwards}$ being high, the activation of the recorded neuron also should be high.

McAdams and Maunsell (1999) presented related evidence of multiplicative scaling of activations in individual neurons with particular orientation preferences. Their paradigm was more complex, involving both filtering and pigeonholing (see Bundesen et al., 2005, pp. 317–319, for a detailed analysis). They used coloured Gaussians and Gabor patterns as stimuli, and they always presented the Gabor pattern in the RF of the recorded V4 neuron and the coloured Gaussian outside the RF. Figure 5A shows the responses of neurons when attention was directed to the Gabor pattern inside the RF (see

the solid circles and the solid bell-shaped curve) and when attention was directed to the Gaussian outside the RF (see the open circles and the dotted bell-shaped curve). Based on the same data, Figure 5B is a plot of the attended response against the unattended response at each of the tested orientations. The dashed lines show undriven activity ± 1 SE. A strikingly good fit to the data is provided by a straight line with a slope of 1.32 through a point representing the undriven activity in both conditions. The goodness of fit suggests that the effect of attention on the mean rate of firing of a recorded V4 neuron was a multiplicative enhancement by about one third of the activation of the neuron—a proportional scaling of the activation of the neuron in response to stimuli in all orientations.

CONCLUSION

Being a neural interpretation of TVA, NTVA provides quantitative accounts of human performance (reaction times and error rates) in a broad range of experimental paradigms of single-stimulus recognition and attentional selection from multielement displays. By use of the same basic equations, NTVA also accounts for a broad range of attentional effects observed in firing rates of single cells in the primate visual system. NTVA provides a mathematical framework to unify the two fields of research.

REFERENCES

Broadbent, D. E. (1971). *Decision and stress*. London: Academic Press.

Bublak, P., Redel, P., & Finke, K. (2006). Spatial and non-spatial attention deficits in neurodegenerative diseases: Assessment based on Bundesen's theory of visual attention (TVA). *Restorative Neurology and Neuroscience, 24*, 287–301.

Bundesen, C. (1990). A theory of visual attention. *Psychological Review, 97*, 523–547.

Bundesen, C. (1993). The relationship between independent race models and Luce's choice axiom. *Journal of Mathematical Psychology, 37*, 446–471.

Bundesen, C., & Habekost, T. (in press). *Principles of visual attention: Linking mind and brain*. Oxford University Press.

Bundesen, C., Habekost, T., & Kyllingsbæk, S. (2005). A neural theory of visual attention: Bridging cognition and neurophysiology. *Psychological Review, 112*, 291–328.

Bundesen, C., & Pedersen, L. F. (1983). Color segregation and visual search. *Perception and Psychophysics, 33*, 487–493.

Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General, 113*, 501–517.

Duncan, J., Bundesen, C., Olson, A., Humphreys, G., Chavda, S., & Shibuya, H. (1999). Systematic analysis of deficits in visual attention. *Journal of Experimental Psychology: General*, *128*, 450–478.

Duncan, J., Bundesen, C., Olson, A., Humphreys, G., Ward, R., Kyllingsbæk, S., et al. (2003). Dorsal and ventral simultanagnosia. *Cognitive Neuropsychology*, *20*, 675–701.

Estes, W. K., & Taylor, H. A. (1964). A detection method and probabilistic models for assessing information processing from brief visual displays. *Proceedings of the National Academy of Sciences of the USA*, *52*(2), 446–454.

Finke, K., Bublak, P., Dose, M., Müller, H. J., & Schneider, W. X. (2006). Parameter-based assessment of spatial and non-spatial attentional deficits in Huntington's disease. *Brain*, *129*, 1137–1151.

Habekost, T., & Bundesen, C. (2003). Patient assessment based on a theory of visual attention (TVA): Subtle deficits after a right frontal-subcortical lesion. *Neuropsychologia*, *41*, 1171–1188.

Habekost, T., & Starrfelt, R. (2006). Alexia and quadrant-amblyopia: Reading disability after a minor visual field deficit. *Neuropsychologia*, *44*, 2465–2476.

Kyllingsbæk, S., Schneider, W. X., & Bundesen, C. (2001). Automatic attraction of attention to former targets in visual displays of letters. *Perception and Psychophysics*, *63*, 85–98.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492–527.

Logan, G. D. (1996). The CODE theory of visual attention: An integration of space-based and object-based attention. *Psychological Review*, *103*, 603–649.

Logan, G. D. (2002). An instance theory of attention and memory. *Psychological Review*, *109*, 376–400.

Logan, G. D., & Bundesen, C. (1996). Spatial effects in the partial report paradigm: A challenge for theories of visual spatial attention. In D. L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 35, pp. 243–282). San Diego, CA: Academic Press.

Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review*, *108*, 393–434.

Luce, R. D. (1963). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 1, pp. 103–189). New York: Wiley.

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*, 279–281.

Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology*, *14*, 744–751.

McAdams, C. J., & Maunsell, J. H. R. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience*, *19*, 431–441.

Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, *229*, 782–784.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, *104*, 266–300.

Posner, M. I., Nissen, M. J., & Ogden, W. C. (1978). Attended and unattended processing modes: The role of set for spatial location. In H. L. Pick & I. J. Saltzman (Eds.), *Modes of perceiving and processing information* (pp. 137–157). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Reynolds, J. H., Chelazzi, L., & Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience*, *19*, 1736–1753.

Schneider, W., & Fisk, A. D. (1982). Degree of consistent training: Improvements in search performance and automatic process development. *Perception and Psychophysics*, *31*, 160–168.

Shibuya, H., & Bundesen, C. (1988). Visual selection from multielement displays: Measuring and modeling effects of exposure duration. *Journal of Experimental Psychology: Human Perception and Performance, 14*, 591–600.

Sperling, G. (1967). Successive approximations to a model for short-term memory. *Acta Psychologica, 27*, 285–292.

Townsend, J. T., & Ashby, F. G. (1982). Experimental test of contemporary mathematical models of visual letter recognition. *Journal of Experimental Psychology: Human Perception and Performance, 8*, 834–864.

Townsend, J. T., & Landon, D. E. (1982). An experimental and theoretical investigation of the constant-ratio rule and other models of visual letter confusion. *Journal of Mathematical Psychology, 25*, 119–162.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12*, 97–136.

Treisman, A. M., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review, 95*, 15–48.

Treue, S., & Martinez-Trujillo, J. C. M. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature, 399*, 575–579.

Van Oeffelen, M. P., & Vos, P. G. (1982). Configurational effects on the enumeration of dots: Counting by groups. *Memory and Cognition, 10*, 396–404.

Van Oeffelen, M. P., & Vos, P. G. (1983). An algorithm for pattern description on the level of relative proximity. *Pattern Recognition, 16*, 341–348.