

Phase noise and the classification of natural images

Felix A. Wichmann^a, Doris I. Braun^b, Karl R. Gegenfurtner^{b,*}

^a *Max-Planck-Institut für biologische Kybernetik, Spemannstr. 38, 72076 Tübingen, Germany*

^b *Abteilung Allgemeine Psychologie, Justus-Liebig-Universität, Otto-Behagel-Str. 10F, 35394 Gießen, Germany*

Received 14 April 2005; received in revised form 31 October 2005

Abstract

We measured the effect of global phase manipulations on a rapid animal categorization task. The Fourier spectra of our images of natural scenes were manipulated by adding zero-mean random phase noise at all spatial frequencies. The phase noise was the independent variable, uniformly and symmetrically distributed between 0° and $\pm 180^\circ$. Subjects were remarkably resistant to phase noise. Even with $\pm 120^\circ$ phase noise subjects were still performing at 75% correct. The high resistance of the subjects' animal categorization rate to phase noise suggests that the visual system is highly robust to such random image changes. The proportion of correct answers closely followed the correlation between original and the phase noise-distorted images. Animal detection rate was higher when the same task was performed with contrast reduced versions of the same natural images, at contrasts where the contrast reduction mimicked that resulting from our phase randomization. Since the subjects' categorization rate was better in the contrast experiment, reduction of local contrast alone cannot explain the performance in the phase noise experiment. This result obtained with natural images differs from those obtained for simple sinusoidal stimuli where performance changes due to phase changes are attributed to local contrast changes only. Thus the global phase-change accompanying disruption of image structure such as edges and object boundaries at different spatial scales reduces object classification over and above the performance deficit resulting from reducing contrast. Additional color information improves the categorization performance by 2%.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Natural scenes; Phase; Fourier; Contrast; Object categorization; Image classification

1. Introduction

Object classification and categorization is one of the most remarkable achievements of the visual system. Higher primates effortlessly classify a large number of very different and even partly occluded objects in their natural surroundings. Despite the diversity of individual natural images they share, as an ensemble, some statistical structure. It has been found, for example, that natural scenes have a characteristic $1\text{-over-}f$ Fourier-amplitude spectrum implying that most of the power is contained in the low spatial frequency components (Field, 1987; Thomson, 1999a, 1999b; van der Schaaf & van Hateren, 1996).

If projected into a sine-/cosine wave basis (Fourier transform), images of natural scenes thus diverge from each other primarily in terms of their phase and not their amplitude spectra. Every image can, in principle, be synthesized from sine wave gratings. For the generation of a particular image, sine gratings of the correct spatial frequency, amplitude and phase have to be combined. Phase is particularly important for edges, since edges require an alignment of the phase of different spatial frequency components. In a well-known demonstration of the importance of global phase by Piotrowski and Campbell (1982; see also Oppenheim and Lim, 1981) two images were mixed, one contributing its Fourier amplitude and the other its Fourier phase. Invariably, the resulting combination looks much more like the image contributing the phase spectrum and not like the one contributing the amplitude spectrum. Clearly, this results from the aforementioned fact that most natural images

* Corresponding author. Fax: +49 641 99 26 119.

E-mail address: gegenfurtner@uni-giessen.de (K.R. Gegenfurtner).

have similar amplitude spectra with the amplitude decreasing linearly with spatial frequency, typically termed 1-over- f . When the phase spectrum of an image is randomly swapped across frequencies, that is, its Fourier energy is randomly distributed over the image, the resulting image becomes impossible to recognize.

1.1. Human phase-sensitivity using simple stimuli

The importance of global phase in the Fourier representation of image structure in natural images does not imply, however, that the primate visual system necessarily encodes global phase explicitly—this would only be required if the visual system did perform a global Fourier transform of the images mapped onto its retinae, an idea that is certainly not tenable (Derrington & Henning, 1989; Henning, Hertz, & Broadbent, 1975; Westheimer, 2001). This conclusion is in line with Piotrowski and Campbell (1982) who, in addition to their demonstration of the importance of phase for image structure, also found that for recognition human observers are not critically sensitive to the precision of the encoded phase. Similar conclusions were reached by Burr (1980) and Badcock (1984a, 1984b, 1988) who examined the discrimination performance of human observers for phase-shifted gratings composed of several harmonically related gratings. Burr (1980) found phase discrimination thresholds of about 30° for all spatial frequencies tested. Badcock converted the phase discrimination thresholds of his observers to local contrast changes accompanying, willy-nilly, every phase change. Badcock concluded “the visual system does not have direct access to the spatial phase of constituent sinusoidal components in an image but instead codes the local contrast and position of image features (1988, p. 305)”.

1.2. Phase-spectra of natural images and higher-order statistics

In a series of elegant articles Thomson and colleagues (Thomson, 1999a, 1999b; Thomson & Foster, 1997; Thomson, Foster, & Summers, 2000) explored the properties of phase spectra of natural images within a statistical framework. Changing the phase spectrum of an image does not affect its power spectrum (i.e., the autocorrelation function) and thus shows how little of the *content* of an image is contained therein. What phase spectrum manipulations do change, however, are higher-order image statistics (moments and cumulants of degree 3 and above). Thus edges, contours and other visually salient features cannot be captured by first- and second-order statistics but must be contained in the higher-order statistics (Franz & Schölkopf, 2005).

Thomson et al. extend the commonly used first- and second-order statistics analysis by computing higher-order image statistics attempting to find whether regularities in the phase spectra of classes of images—e.g., natural scenes—are reflected in their higher-order image statistics.

Thomson (1999b) showed that (whitened) natural scenes have a strictly positive kurtosis, whereas phase-randomized versions of the very same images have positive and negative kurtosis values very close to zero.

Thomson et al. (2000) conducted a psychophysical experiment relating the structure of natural-image phase spectra to visual perception and notions from efficient coding. Note, however, that their main interest was statistical:

“One obvious threshold psychophysical paradigm would require observers to discriminate a slightly phase-perturbed image from a natural image, but under these circumstances observers’ sensitivities might be determined by one particular ‘feature’ in the natural images, i.e., they may not perform the tasks statistically” (Thomson et al., 2000, p. 1065).

Consistent with their statistical aims human observers in the Thomson et al. (2000) study discriminated completely phase randomized images from images with slightly less phase randomization (or quantization). None of the images looked like “natural images”, i.e., subjects discriminated “cloud-like” images with power spectra derived from natural images. Subsequently Thomson et al. correlated certain higher-order statistics with their observers’ performance and found phase-only kurtosis—kurtosis after removing second-order image structure—to provide a reasonable fit to their empirical data.

The aim of our study is related but different: we are precisely interested in the recognition of individual, natural “real-life” images with their multiple and possibly highly redundant features and how phase perturbation interferes with classification. We thus conducted a study akin to that outlined in the citation of Thomson et al. above: comparing image classification performance for natural images with and without various degrees of phase perturbation.

1.3. Phase-alignment across spatial scales: Edges, lines and contours

Despite the negative results of Piotrowski and Campbell as well as Burr and Badcock on finding phase-sensitive mechanisms using simple stimuli, phase-sensitive mechanisms, or whatever other mechanism such as local energy estimation (Burr, Morrone, & Spinelli, 1989; Morrone & Burr, 1988; Morrone & Owens, 1987; Morrone, Burr, & Spinelli, 1989) may detect the phase-change accompanying image changes, could be highly sensitive to particular phase relationships between harmonically related spatial frequencies. This is particularly true for detecting phase alignment at edges—the notion of alignment across spatial scales (Marr, 1982). If true, we would expect that for natural images with their complex structure of edges and object boundaries a global phase change is more disruptive than a contrast reduction—for natural images phase change could thus not simply be equated with local contrast change.

Rephrasing the above from a higher-order statistics point-of-view, it may be that phase changes in natural

images—affecting higher-order but not the first- and second-order statistics of the images—have a severe impact on the recognition of natural images because phase manipulations destroy particular higher-order regularities (Thomson, 1999b). Burr (1980) and Badcock (1984a, 1984b, 1988) may have found the visual system to be rather insensitive to phase manipulations because simple, artificial stimuli are lacking such regularities.

Thus we investigated the effect of phase manipulations on the processing of natural images rather than the sine gratings. Our goal was to evaluate how robust the human visual object classification and categorization system is with respect to phase manipulations in natural scenes.

1.4. Experimental paradigm: Rapid animal categorization in natural images

Thorpe, Fize, and Marlot (1996) showed that humans are able to rapidly categorize briefly presented natural scenes. Their subjects had to indicate whether the flashed photograph showed an animal by releasing a button. The analysis of the frontal event-related potentials revealed a differential activity between target and non-target trials at about 150 ms, which can be taken as the processing time for high-level object categorization of natural images. In a later study Rousset, Fabre-Thorpe, and Thorpe (2002) compared the human performance when subjects had to do the same task with one or with two simultaneously presented images. Since the behavioural and the electrophysiological data showed no difference between the one- and the two-image condition the authors concluded that high-level object categorization of natural images can be done in parallel. We used the same visual categorisation task and a two-image condition to examine the effect of phase noise on visual processing. In our easy animal versus non-animal categorization task the target and distractor image were presented simultaneously to the left and right of a central fixation point. Since phase noise blurs the contours and reduces local image contrast, shape recognition should become more and more difficult with increasing phase noise.

We also investigated the effect of color on the phase manipulations. Color information has been shown to play an important role for image segmentation and recognition (Gegenfurtner & Rieger, 2000; Wichmann, Sharpe, & Gegenfurtner, 2002). However, the human visual system might be a less sensitive to the phase of colored stimuli compared to luminance stimuli (Troscianko & Harris, 1988). If phase and color information are used in early visual processing for image segmentation then performance should be affected by the addition of small amounts of phase noise even more so when, in addition, color cues are removed from the images.

Sensitivity of the human visual system to phase noise was tested under two conditions: In one experiment the phase spectra of naturally colored images was altered, in the other the same was done with black and white images.

To determine whether the effect of changing global phase is merely that of a reduction in image contrast, subjects had to perform the same task with the same images but now image contrast was varied.

2. General methods

2.1. Experimental set-up

Subjects were seated in a dimly lit room facing a Sony GDM F 520 color monitor. At the viewing distance of 72 cm the active screen area of 1280 by 1024 pixels subtended 32.5° of visual angle on the subjects' retinae. The frame rate of the display was 85 Hz non-interleaved. Natural image stimuli had a size of 256 by 384 pixels and subtended 6.6° by 9.9° of visual angle. Pairs of stimuli were presented for 50 ms 3.5° to the left and right of a small black fixation dot on a grey background. The luminance of the background was adjusted to the overall mean luminance of the images (see Fig. 3). The monitor was carefully calibrated in luminance and color using a Photo Research PR 650 spectroradiometer and a Graseby Electronics model S380 photometer.

2.2. Stimuli

Stimuli were digitized 256 × 512 pixel photographs of natural scenes from the large commercial *Corel Photo Library*. A set of 700 images was selected as targets containing one or more animals (birds, mammals and insects) of a “reasonable” size in their natural environments. Another set of 700 images was selected as distractors. These were photographs of landscapes, plants and rock formations. We excluded photographs when they contained humans or were dominated by buildings or streets because of their regular structure. The same images have been used in previous studies investigating animal categorization (Thorpe, Gegenfurtner, Fabre-Thorpe, & Bülthoff, 2001).

2.3. Phase noise manipulation

The phase spectra of the images were manipulated by adding random phase noise to the images. Random phase offsets from the interval $[-\Phi, +\Phi]$ with zero-mean were added at each spatial frequency. Φ could take values between 0° (unmodified images) up to 180° (0°, 30°, 60°, 90°, 120°, 135°, 150° and 180°). The phase manipulation yields images with fewer edges and features. Fig. 1 illustrates the effect of global phase noise on local features. In this figure, we decomposed a square wave into a series of 100 sinewaves with frequencies 1, 3, 5, ..., whose amplitude decreases in proportion to their frequency. During the recomposition, phase noise was added to each of these component sinewaves. Increasing phase noise results in less correspondence to the original pattern, which manifests itself as an effective reduction in local feature contrast. The contrast reduction is a cosine function of the average magnitude of the phase noise. Phase noise in the interval $[-90, +90]$ has an average phase noise magnitude of 45°, resulting in a contrast reduction of factor $\cos(45) = 0.72$. For our natural image stimuli, we added phase noise to each one of the coefficients of their discrete Fourier transform. Since the DFT of real valued images are symmetric, i.e., at each frequency there is a component at the symmetric frequency with the same real part and an illusory part of the reverse sign, we only added noise to half of the coefficients and retained the symmetry. The effect of the global phase noise on natural images was a reduction in the correlation of the original image and the noisy image, according to the same cosine function mentioned above. The absolute level of RMS contrast across the image remained constant, however, since the Fourier amplitude spectra of the images were unaffected by the phase manipulations. For the color images, the red, green and blue components of the images were Fourier transformed separately, and the same phase noise was applied to each one of the three color components. Note that the color images were not equiluminant. They correspond to the original images. For the black and white condition, the red, green and blue components at

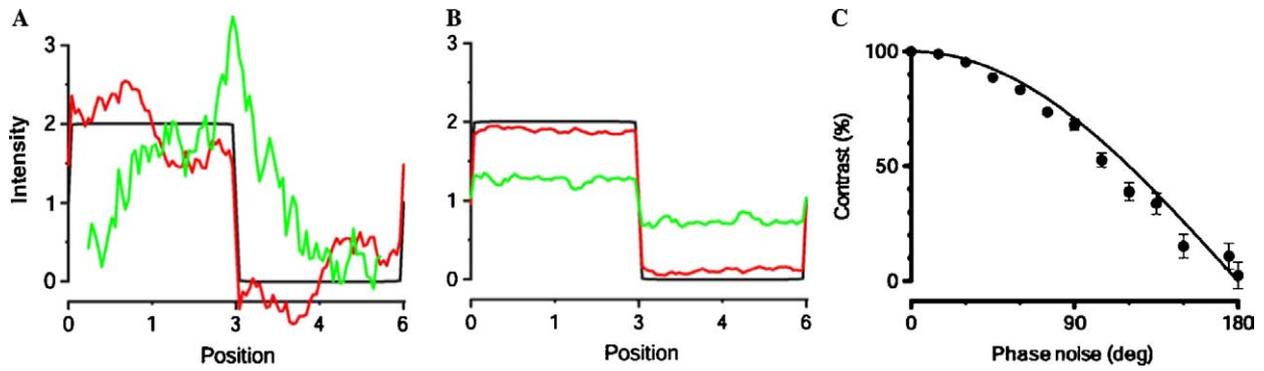


Fig. 1. Phase manipulation. (A) Shows a period of a square wave in black. The square wave can be decomposed into a series of sinewaves. For the red curve, a random phase offset in the interval $[-45, 45]$ was added to each of these component gratings. For the green curve, the phase offset was from the interval $[-135, 135]$. (B) Many of these noisy square waves were averaged. This results in square waves of lower contrast. (C) Shows the correspondence between the amount of phase noise and the local edge contrast.

each pixel were weighted by the luminance of the corresponding phosphors. This way, the colored and black and white images have the same luminance at all corresponding pixels.

2.4. Contrast manipulation

To prevent overflows during the image processing operations, we reduced the contrast of all images to half their original contrast before adding phase noise. Contrast of the images was varied by scaling each pixel relative to the luminance mean of the image. In Experiment 2, where the contrast was varied, five different contrast level were chosen for testing: 5, 10, 25, 50, and 100%, relative to the scaled images, i.e., the contrast levels

were 2.5, 5, 12.5, 25, and 50%, with respect to the original images. Examples for the effect of both image manipulations are given in Fig. 2. For the color images, this manipulation was applied to the red, green and blue components separately.

2.5. Task and subjects

In a two-alternative, spatial forced-choice procedure the subject was asked to indicate which of the two simultaneously presented images contained an animal (the target) by pressing either the left or right mouse button (see Fig. 3). For each subject both the phase and the contrast experiment were tested separately in a randomized sequence. The method

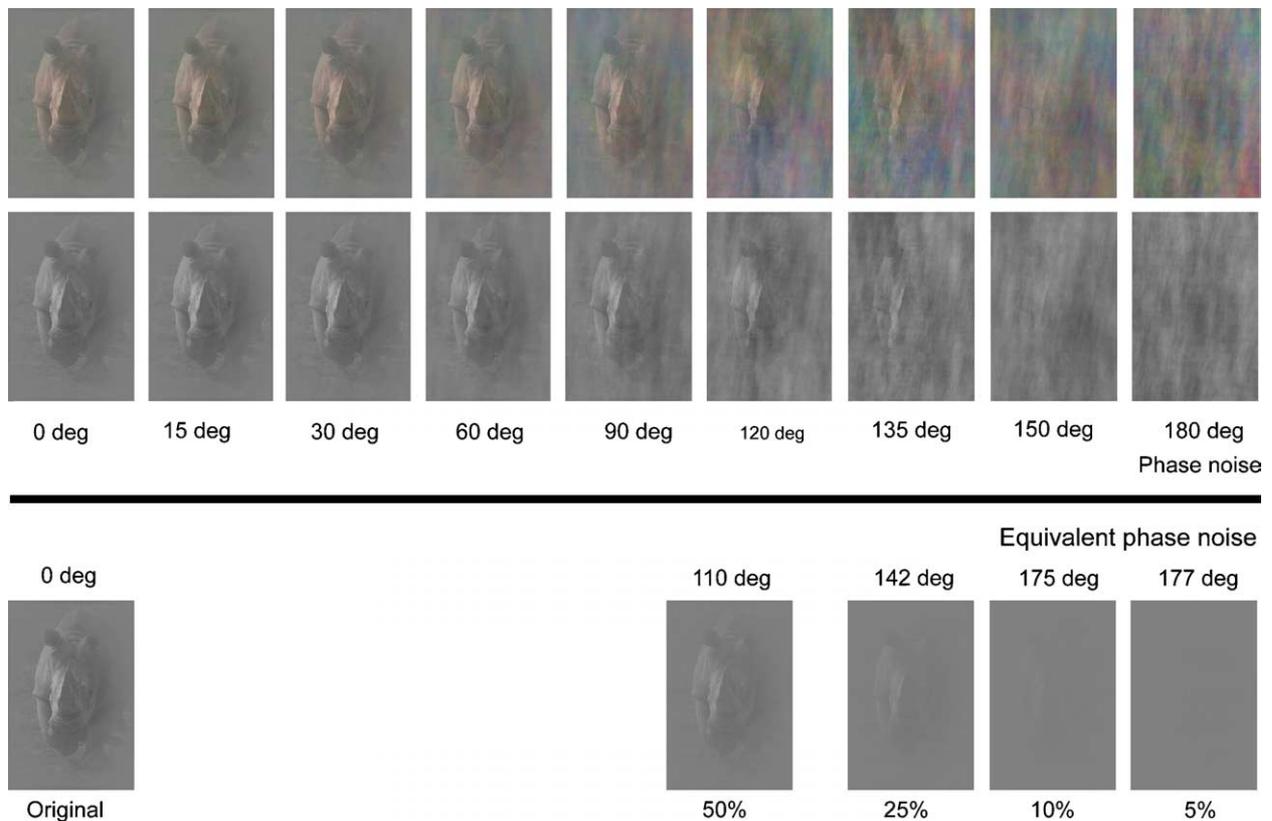


Fig. 2. Example images. The top two rows show an image of a rhinoceros with various amounts of phase noise added. The top row shows the images in color, the second row in black and white. The bottom row shows the same image at various contrasts, with the corresponding amount of phase noise indicated on the scale between the panels.

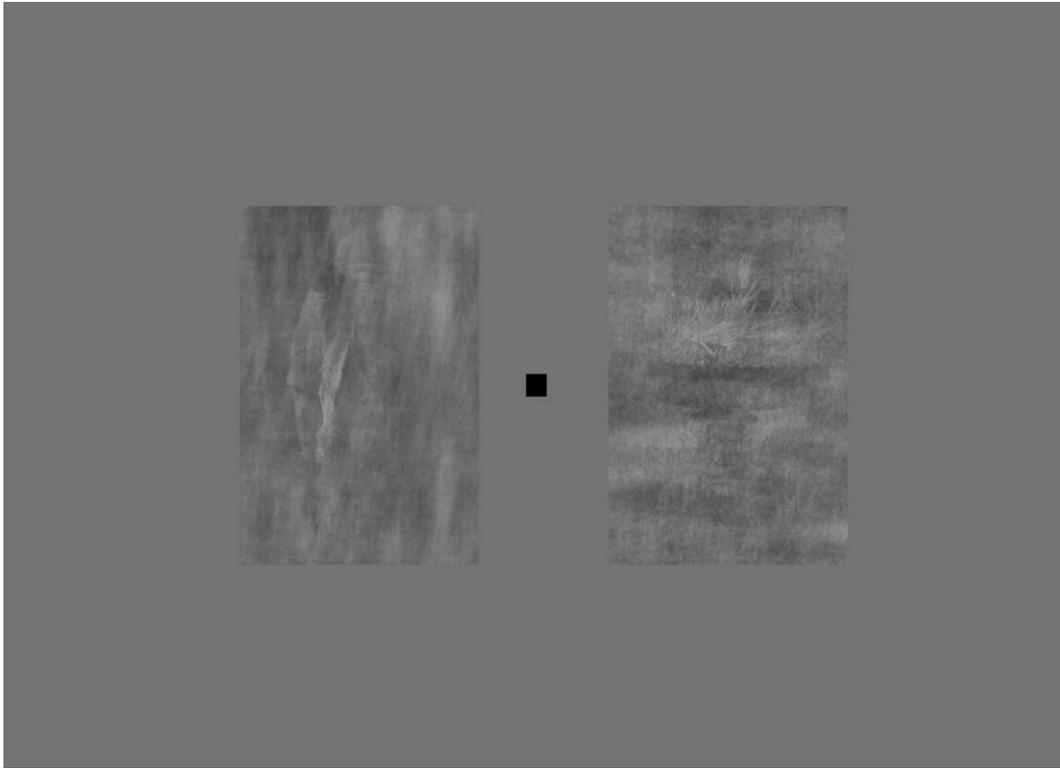


Fig. 3. Experimental stimuli. An animal and a distracter image with an identical amount of phase noise (135°) were displayed side by side on a computer screen. The subject had to indicate the target image containing the animal.

of constant stimuli was used. For each subject we collected the responses of 50 trials for each condition, resulting in 400 trials for the phase experiments with eight phase noise levels and 200 trials for the contrast experiment. The total duration of the experiment was 50–60 min per subject. For both experimental conditions 13 subjects were tested. All subjects were naïve with respect to the aim of the study and had normal or corrected to normal vision. Most of them were students of the University Giessen. The subjects' average age was 29 years for the experiment with black and white images and 24 years for the experiment with colored images. Every subject saw every image only once to exclude any potential confounds stemming from learning or recognition memory for particular images.

2.6. Experimental variations

Two different conditions were tested: in one experimental condition natural images were presented in black and white, in the other they were shown in color (see Fig. 2). For each condition two different experiments were run: In one part of the experiment the phase spectra of the images were manipulated by eight different amounts of phase noise, in the other part contrast was reduced in five steps.

3. Results

3.1. Experiment 1: Phase noise

Fig. 4 shows individual data from four representative subjects for black and white stimuli. All subjects were remarkably resistant to phase noise in our animal categorization task. In Fig. 5 the average of all 13 subjects is shown for both the color and black and white conditions. The proportion of correct answers stays close to 100% correct with $\pm 90^\circ$ noise-distorted images. Even with $\pm 120^\circ$ noise sub-

jects were still well above 75% correct in both experiments. In the color experiment the subjects' categorization rate was slightly better and the differences in performance rate between the color and the black and white Experiment of 2–3% were small in magnitude but statistically significant. The overall pattern of results, i.e., the slight but monotonic decrease in performance with increasing noise was the same, however.

In Figs. 4 and 5, the prediction of the cosine rule is plotted as a solid curve. This prediction shows the level of expected performance if recognition was proportional to local image contrast. It can be seen that performance was in all cases significantly better than the prediction based on contrast reduction alone. This is counterintuitive, since subjects should certainly not be any better than the prediction from the contrast reduction. However, in the above equalization we tacitly assumed that contrast translates linearly into recognition rate. This is, of course, incorrect. To determine the proper amount of equivalent phase noise for each contrast, we determined the nonlinear contrast response function relating contrast to recognition for each individual subject.

3.2. Experiment 2: Contrast reduction

Contrast processing in the visual system is highly nonlinear (Burton, 1981; Henning, Bird, & Wichmann, 2002; Legge, 1981; Legge & Foley, 1980; Wichmann, 1999, 2002). In order to measure the effect of local contrast reduction we

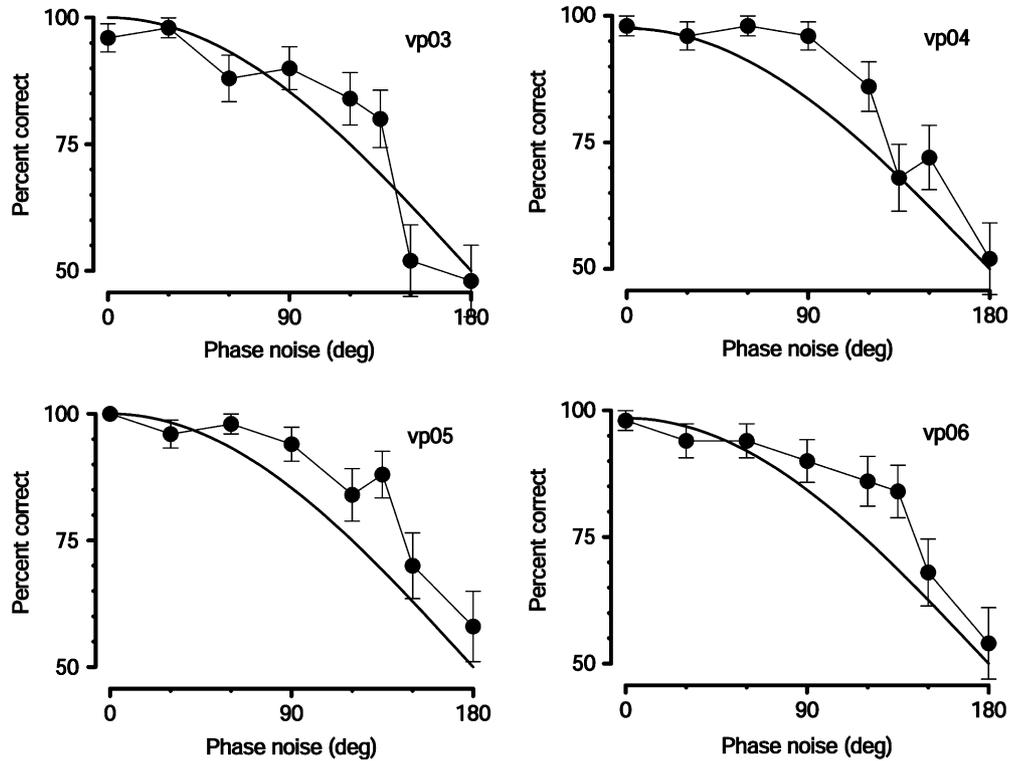


Fig. 4. Data from four representative subjects. Proportion of correctly categorized animals is plotted on the y-axis as a function of the amount of phase noise on the x-axis. See the text for details of the phase manipulation.

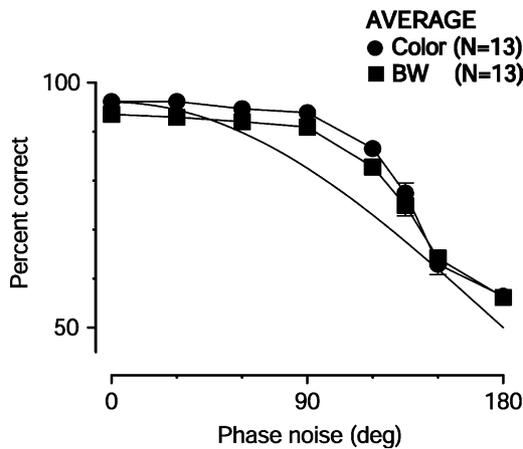


Fig. 5. Average data from 13 subjects for black and white stimuli (squares) and colored stimuli (circles). Proportion of correctly categorized animals is plotted on the y-axis as a function of the amount of phase noise on the x-axis. See the text for details of the phase manipulation.

also tested our subjects' ability to detect animals in natural images whose contrast was reduced. Contrast measures are relative to the original image contrast, as described above. Fig. 6 shows data from four representative subjects.

In Fig. 7 the average correct categorization responses of all 13 subjects is plotted as a function of image contrast—relative to the original contrast of each image. At 5% image contrast subjects were basically guessing, above 25% contrast their categorization rate is better than 75% correct in both experiments. There is essentially no difference for cate-

gorization rates for color and black and white stimuli, except for the highest contrast used. Unlike in the case of the phase manipulation (Experiment 1) color information does not improve the categorization rate when image contrast is reduced.

We can now graph the data from the contrast experiments in the corresponding graphs for the phase noise. Recognition rate from the contrast experiment is plotted at the phase noise angle corresponding to the contrast value according to the cosine rule. Fig. 8 shows the result of this procedure for the average data of all 13 subjects, both for black and white and for colored stimuli. The results from the phase noise manipulation are shown using circles, those from the contrast reduction using squares.

Performance in the contrast experiment is significantly higher than when phase noise is added to the images. In other words, adding phase noise has more of a detrimental effect on classification than just reducing the image contrast.

One interesting feature of all our data is that subjects consistently were above chance even at the highest phase noise of $[-180, 180]$. In this case, all information about phase is lost. However, the amplitude information is still there. That is, the noisy images still have somewhat oriented contours, if these orientations are prevalent in the original images. In our task, where we have feedback to the subjects, subjects seem to be able to use that information to achieve a performance level that is just above chance.

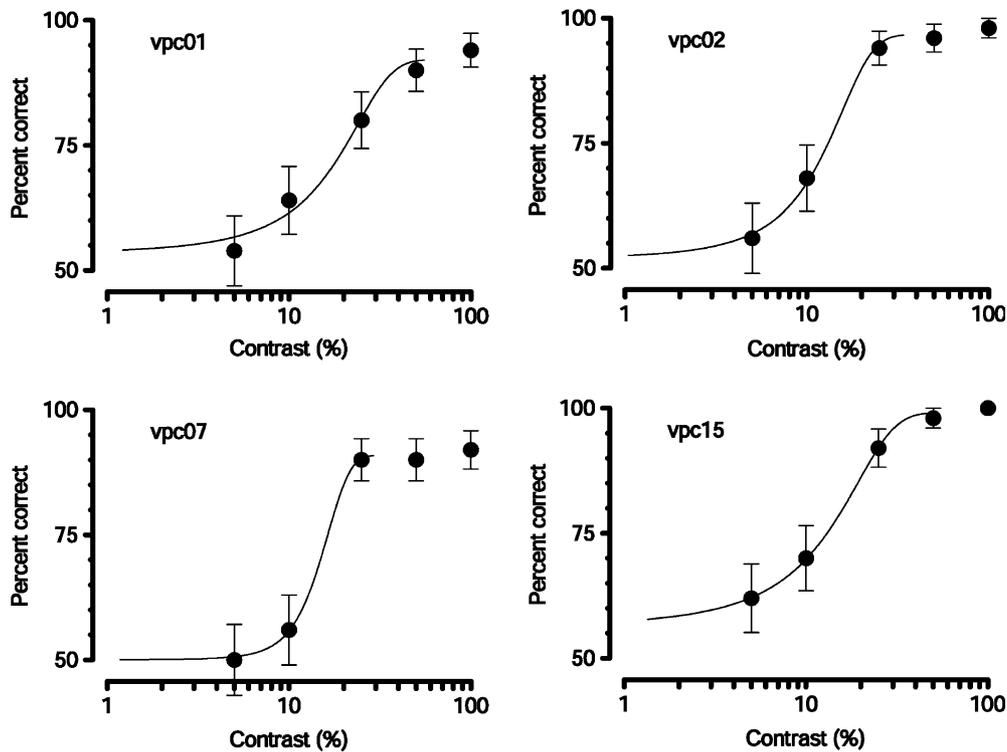


Fig. 6. Data from four subjects. Proportion of correctly categorized animals is plotted on the y -axis as a function of image contrast measured relative to the image contrast of the original image on the x -axis.

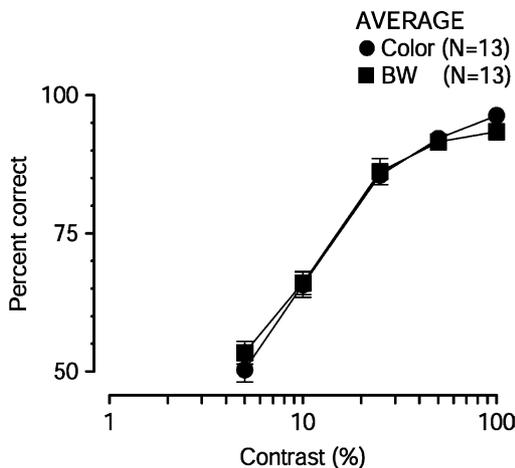


Fig. 7. Average data from 13 subjects for black and white stimuli (squares) and color stimuli (circles). Proportion of correctly categorized animals is plotted on the y -axis as a function of image contrast measured relative to the image contrast of the original image on the x -axis.

4. Discussion

4.1. Summary

We varied the amount of phase noise in natural images to study its effect on classification performance in a rapid animal categorization task. Performance stayed remarkably high for uniformly distributed phase noise up to 90° . When

equating phase-noise and contrast discrimination effects according to the relationship shown in Fig. 1C, the subjects' categorization rate was better in the contrast reduction condition than in the phase noise condition. Thus, reduction of local contrast alone cannot explain the performance in the phase noise experiment. Phase noise is more disruptive than simple contrast reduction because it changes visual features such as local edges, thereby degrading object boundaries.

4.2. Relation to other psychophysical studies

On the one hand, our results agree with Burr (1980) and Badcock (1984a, 1984b) in that we agree on the insensitivity of the human visual system to phase manipulations. Burr (1980) had measured 30° discrimination thresholds for 1-D stimuli, so one may expect even larger resistance to phase noise for the 2-D stimuli we used, assuming the large overlap of filters tuned to different spatial frequencies and orientations resulting in better overall 2-D system estimation of (average) local phase.

On the other hand, our results are in contrast to those obtained using simple, low-level stimuli such as sums of sinusoidal gratings as classification performance with natural images as stimuli cannot be solely attributed to local contrast changes (Badcock, 1984a, 1984b, 1988; Hess & Pointer, 1987; Lawton, 1984). However, even for simple sinewave stimuli a dissociation between contrast and phase has been shown. Tyler and Gorea (1986) varied exposure duration in a task where the observers had to judge the

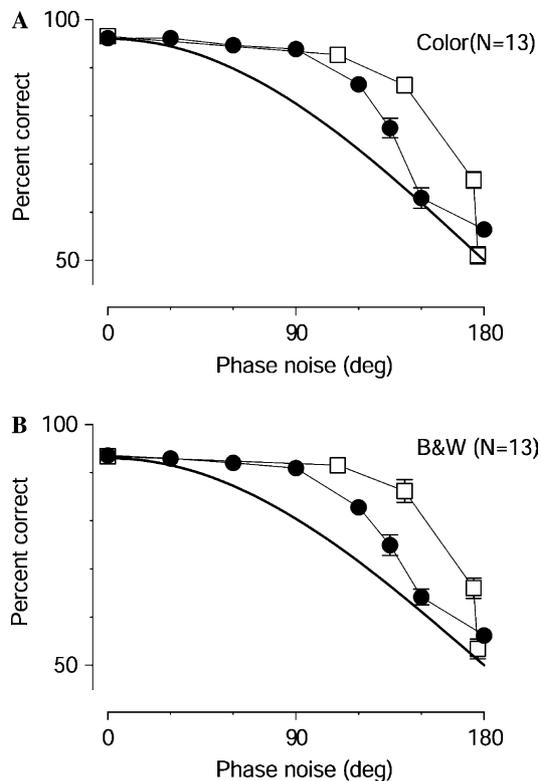


Fig. 8. Contrast corrected average data from 13 subjects for (A) black and white stimuli and (B) colored stimuli. Data from the phase noise condition is shown using circles, from the contrast reduction condition using squares. Proportion of correctly categorized animals is plotted on the y -axis as a function of the amount of phase noise on the x -axis. See the text for details of the conversion from image contrast to equivalent phase noise angle.

position of a sinewave grating relative to a thin line. Phase discriminability increased with exposure duration even when the detectability of stimuli was held constant at all durations. As mentioned in Section 1, our natural images cannot be easily compared to simple $f+3f$ sinewave compounds. It is also difficult to relate our findings to the elegant studies by Victor and Conte (1996) on texture discrimination. In our experiments, subjects were not only required to notice a difference in the animal and distractor images, but they had to actually identify the image that contained the animal.

4.3. Relationship to single unit studies

To a first approximation cells in primary visual cortex provide an oriented spatially band pass filtering of the image. This type of filtering is provided by simple cells, which are sensitive to the phase of the stimuli, and by complex cells, which respond irrespective of spatial phase (DeValois & DeValois, 1988; Hubel & Wiesel, 1968; Movshon, Thompson, & Tolhurst, 1978). While there has been some debate whether these cells really form two separate subpopulations (Martinez et al., 2005; Mechler, Reich, & Victor, 2002), the basic fact that there exist cells sensitive to the spatial phase is without doubt. The question how

well these cells could actually code the absolute phase has been addressed in a number of studies. Based on theoretical grounds, it has been suggested that neurons in V1 would be either odd-symmetric edge detectors or even symmetric line detectors (Morrone & Burr, 1988; Morrone & Owens, 1987). However, recordings from single neurons or pairs of neighboring neurons in cat or monkey visual cortex have shown that simple cells do not fall into these two categories only. Instead, cells have all phases but neighboring cells seem to have a tendency to have phase preferences approximately 90° apart (Field & Tolhurst, 1986; Pollen & Ronner, 1981). These quadrature pairs would be sufficient to calculate the absolute phase of any local pattern but are more likely to encode local contrast energy.

More recently, the neural coding of relative spatial phase in V1 neurons has been determined by Mechler et al. (2002) and Aronov et al. (2003). They found that spatial phase is best coded by single cells and that a population of cells is necessary to achieve phase sensitivity that can approximate psychophysical performance. The picture that is emerging from these studies is that, if expressed in terms of a Fourier basis, phase is coded at a rather coarse level in primary visual cortex. Alternatively, the same fact could be expressed by stating that the Fourier basis may not be a good basis in which to represent images if one is interested in human or animal vision. Which ever way one chooses to express this fact this is in perfect agreement with our psychophysical results.

4.4. Relation to fMRI experiments with phase scrambled images

In a study related to ours, Rainer, Augath, Trinath, and Logothetis (2001) studied the contribution of phase coherence in natural images to the fMRI activation in occipital visual areas in anesthetized monkeys. Unlike perception and recognition, which increased monotonically in humans and monkey with increasing phase coherence, a characteristic V-shaped noise-tuning was found in the BOLD-signal striate and extra-striate visual areas. Like the behavioral data of Rainer et al. we did find categorization performance to be a monotonic function of both phase noise and image contrast, even though the absolute levels of performance cannot be compared because different phase noise manipulations were used. The striking non-monotonic BOLD-signal behavior found by Rainer et al. might, however, have been caused by an artifact in the noise-coherence manipulation of Rainer et al. (2001), as pointed out by Dakin, Hess, Ledgeway, and Achtman (2002). This interpretation is supported by the results of Olman, Ugurbil, Schrater, and Kersten (2004) who did not find a difference between regular and phase scrambled natural images in V1 BOLD response. Thus, it appears as if the BOLD-response in primary visual cortex is mostly driven by RMS contrast of the stimuli, independent of particular phase relationships. Of course, this does not mean that there are no neurons in V1 that are sensitive to phase.

But since the individual neurons in V1 are sensitive to local phase, the sum of all responses might not be affected very much by varying global phase. Also, the fMRI response measures the total activity, which might be dominated by the more numerous complex cells.

4.5. Rapid visual categorization

Our results show that animal detection is robust against adding large degrees of phase noise. In previous studies, it has been observed that performance in this type of task is not at all impaired even if the images are presented upside down (Fabre-Thorpe, Richard, & Thorpe, 1998). Consistent with this observation it has been suggested that observers might make use of rather simple cues based on the amplitude spectrum of the images (Johnson & Olshausen, 2003; Torralba & Oliva, 2003). Indeed, for images taken from the Corel database, combined with images from the world-wide-web and their own digital images, Torralba and Oliva (2003) showed that the global Fourier Amplitude spectrum of a number of natural categories (e.g., forest, beach, animal, man-made object, highway, mountain etc.) are sufficiently different to allow spectral-based categorization based on only the first three principal components of the amplitude spectrum (SPC) of natural versus man-made scenes to reach 80% correct. Somewhat more complex algorithms allow for up to 94% correct classification based on the power spectra of the images. Their results do not exactly tell the level of performance that could be achieved for our task and our images. However, it should be clear from our results that subjects are very limited in their use of the global amplitude spectrum. When the image phase is random, performance is just barely above chance. This indicates that the amplitude spectrum alone is not sufficient to successfully perform this task.

Recently, Wichmann, Rosas, and Gegenfurtner (2005) as well as Wichmann, Rosas, Drewes, and Gegenfurtner (in preparation) explored this issue in greater detail. They tested whether human observers make use of the power spectrum when rapidly classifying natural scenes. In one condition the original images were used for image classification, in the other images whose power spectra were equalized (each power spectrum was set to the mean power spectrum over the ensemble of 1476 images). Thresholds for 75% correct were in the region of 20–30 ms presentation time for all observers, independent of the power spectrum of the images: this result makes it very unlikely that human observers make use of the global power spectrum during rapid image classification, pointing towards individual features such as edges, lines or contours. Thus these studies agree with the results of the current experiments.

4.6. Natural images and higher-order statistics

One very promising line of research is to study higher-order statistics of images and to relate it to visual perception (Franz & Schölkopf, 2005; Thomson, 1999a, 1999b;

Thomson & Foster, 1997; Thomson et al., 2000). Minimally we demonstrated that for natural images phase perturbations, willy-nilly accompanied by higher-order statistics perturbations, are more disruptive than the mere equivalent contrast reduction effect found for simple stimuli. A future goal is to attempt to identify which particular images may have been more or less affected by the phase manipulations and perhaps find correlates of these changes in some higher-order statistics of the images.

4.7. Conclusions

Global phase per se does not seem to be coded in the visual system in the sense of making phase explicit for higher processing areas. Rather local phase—local contrast energy or edge structure—is coded in the visual system, but, because edges are represented across many spatial scales (Marr, 1982), the visual system is very robust against global phase manipulations. These results we obtain with natural images differ from those obtained for simple sinusoidal stimuli where performance changes due to phase changes are explicable by local contrast changes. Thus the global phase-change accompanying disruption of image structure at different spatial scales—edges and object boundaries—reduces object classification over and above the performance deficit resulting from reducing contrast.

Acknowledgments

We are grateful to Matthias Franz, Mark Georgeson, Concetta Morrone and an anonymous reviewer for their helpful comments and suggestions. This research was supported by DFG grants Wi 2103/1 and Ge 879/6 as well as the Max Planck Society.

References

- Aronov, D., Reich, D. S., Mechler, F., & Victor, J. D. (2003). Neural coding of spatial phase in V1 of macaque monkey. *Journal of Neuroscience*, *23*, 3304–3327.
- Badcock, D. R. (1984a). How do we discriminate relative spatial phase? *Vision Research*, *24*, 1847–1857.
- Badcock, D. R. (1984b). Spatial phase or luminance profile discrimination? *Vision Research*, *24*, 613–623.
- Badcock, D. R. (1988). Discrimination of spatial phase changes: Contrast and position codes. *Spatial Vision*, *3*, 305–322.
- Burr, D. C. (1980). Sensitivity to spatial phase. *Vision Research*, *20*, 391–396.
- Burr, D. C., Morrone, M. C., & Spinelli, D. (1989). Evidence for edge and bar detectors in human vision. *Vision Research*, *29*, 419–431.
- Burton, G. J. (1981). Contrast discrimination by the human visual system. *Biological Cybernetics*, *40*, 27–38.
- Dakin, S. C., Hess, R. F., Ledgeway, T., & Achtman, R. L. (2002). What causes non-monotonic tuning of fMRI response to noisy images? *Current Biology*, *12*, R476–R477.
- Derrington, A. M., & Henning, G. B. (1989). Some observations on the masking effects of two-dimensional stimuli. *Vision Research*, *29*, 241–246.
- DeValois, R., & DeValois, K. K. (1988). *Spatial vision*. New York: Oxford University Press.

- Fabre-Thorpe, F., Richard, G., & Thorpe, S. J. (1998). Rapid categorization of natural images by rhesus monkeys. *Neuroreport*, 9, 303–308.
- Field, D. J. (1987). Relations between the statistics of natural images and response properties of cortical cells. *Journal of the Optical Society of America A*, 4, 2379–2394.
- Field, D. J., & Tolhurst, D. J. (1986). The structure and symmetry of simple-cell receptive-field profiles in the cat's visual cortex. *Proceedings of the Royal Society of London B*, 228, 379–400.
- Franz, M. O., & Schölkopf, B. (2005). Implicit Wiener series for higher-order image analysis. In L. K. Saul, Y. Weiss, & L. Bottou (Eds.), *Advances in neural information processing systems* (pp. 465–472, Vol. 17). Cambridge, MA, USA: MIT Press.
- Gegenfurtner, K. R., & Rieger, J. (2000). Sensory and cognitive contributions of color to the recognition of natural scenes. *Current Biology*, 10, 805–808.
- Henning, G. B., Bird, C. M., & Wichmann, F. A. (2002). Contrast discrimination with pulse-trains in pink noise. *Journal of the Optical Society of America A*, 19, 1259–1266.
- Henning, G. B., Hertz, B. G., & Broadbent, D. E. (1975). Some experiments bearing on the hypothesis that the visual system analyzes patterns in independent bands of spatial frequency. *Vision Research*, 15, 887–899.
- Hess, R. F., & Pointer, J. S. (1987). Evidence for spatially local computations underlying discrimination of periodic patterns in fovea and periphery. *Vision Research*, 27, 1343–1360.
- Hubel, D., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology (London)*, 195, 215–243.
- Johnson, J. S., & Olshausen, B. A. (2003). Timecourse of neural signatures of object recognition. *Journal of Vision*, 3, 499–512.
- Lawton, T. B. (1984). The effect of phase structures on spatial phase discrimination. *Vision Research*, 24, 139–148.
- Legge, G. E. (1981). A power law for contrast discrimination. *Vision Research*, 21, 457–467.
- Legge, G. E., & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America*, 70, 1458–1471.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: Freeman.
- Martinez, L. M., Wang, Q., Reid, R. C., Pillai, C., Alonso, J. M., Sommer, F. T., et al. (2005). Receptive field structure varies with layer in the primary visual cortex. *Nature Neuroscience*, 8, 372–379.
- Mechler, F., Reich, D. S., & Victor, J. D. (2002). Detection and discrimination of relative spatial phase by V1 neurons. *Journal of Neuroscience*, 22, 6129–6157.
- Morrone, M. C., & Burr, D. C. (1988). Feature detection in human vision: a phase-dependent energy model. *Proceedings of the Royal Society of London B*, 1280, 221–245.
- Morrone, M. C., & Owens, R. (1987). Feature detection from local energy. *Pattern Recognition Letters*, 6, 303–313.
- Morrone, M. C., Burr, D. C., & Spinelli, D. (1989). Discrimination of spatial phase in central and peripheral vision. *Vision Research*, 29, 433–445.
- Movshon, J. A., Thompson, I. D., & Tolhurst, D. J. (1978). Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *Journal of Physiology*, 283, 53–77.
- Olman, C. A., Ugurbil, K., Schrater, P., & Kersten, D. (2004). BOLD fMRI and psychophysical measurements of contrast response to broadband images. *Vision Research*, 44, 669–683.
- Oppenheim, A. V., & Lim, J. S. (1981). The importance of phase in signals. *Proceedings of the IEEE*, 69, 529–541.
- Piotrowski, L. N., & Campbell, F. W. (1982). A demonstration of the visual importance and flexibility of spatial-frequency amplitude and phase. *Perception*, 11, 337–346.
- Pollen, D. A., & Ronner, S. F. (1981). Phase relationships between adjacent simple cells in the visual cortex. *Science*, 212, 1409–1411.
- Rainer, G., Augath, M., Trinath, T., & Logothetis, N. K. (2001). Nonmonotonic noise tuning of BOLD fMRI signal to natural images in the visual cortex of the anesthetized monkey. *Current Biology*, 11, 846–854.
- Rousselet, G. A., Fabre-Thorpe, M., & Thorpe, S. J. (2002). Parallel processing in high-level categorization of natural images. *Nature Neuroscience*, 5, 629–630.
- Thomson, M. G. (1999a). Visual coding and the phase structure of natural scenes. *Network*, 10, 123–132.
- Thomson, M. G. (1999b). Higher-order structure in natural scenes. *Journal of the Optical Society of America A*, 16, 1549–1553.
- Thomson, M. G., & Foster, D. H. (1997). Role of second- and third-order statistics in the discriminability of natural images. *Journal of the Optical Society of America A*, 14, 2081–2090.
- Thomson, M. G., Foster, D. H., & Summers, R. J. (2000). Human sensitivity to phase perturbations in natural images: a statistical framework. *Perception*, 29, 1057–1069.
- Thorpe, S. J., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381, 520–522.
- Thorpe, S. J., Gegenfurtner, K. R., Fabre-Thorpe, M., & Bülthoff, H. H. (2001). Detection of animals in natural images using far peripheral vision. *European Journal of Neuroscience*, 14, 869–876.
- Torralba, A., & Oliva, A. (2003). Statistics of natural image categories. *Network: Computation in Neural Systems*, 14(2003), 391–412.
- Troschianko, T., & Harris, J. (1988). Phase discrimination in chromatic compound gratings. *Vision Research*, 28, 1041–1049.
- Tyler, C. W., & Gorea, A. (1986). Different encoding mechanisms for phase and contrast. *Vision Research*, 26, 1073–1082.
- van der Schaaf, A., & van Hateren, J. H. (1996). Modelling the power spectra of natural images: statistics and information. *Vision Research*, 36, 2759–2770.
- Victor, J. D., & Conte, M. M. (1996). The role of high-order phase correlations in texture processing. *Vision Research*, 36, 1615–1631.
- Westheimer, G. (2001). The Fourier theory of vision. *Perception*, 30, 531–541.
- Wichmann, F. A. (1999). Some aspects of modelling human spatial vision: Contrast discrimination. PhD thesis. Oxford University.
- Wichmann, F. A. (2002). Modelling contrast transfer in spatial vision (Abstract). *Journal of Vision*, 2, 7a.
- Wichmann, F. A., Sharpe, L. T., & Gegenfurtner, K. R. (2002). The contribution of colour to recognition memory for natural scenes. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 28, 509–520.
- Wichmann, F. A., Rosas, P., & Gegenfurtner, K. R. (2005). Rapid animal detection in natural scenes: Critical features are local (Abstract). *Journal of Vision*, 5, 376a.
- Wichmann, F. A., Rosas, P., Drewes, J., & Gegenfurtner, K. R. (in preparation). Rapid animal detection in natural scenes: Critical features are low-level and local.