Pergamon

0042-6989(95)00303-7

Vision Res., Vol. 36, No. 16, pp. 2497–2513, 1996

Printed in Great Britain
0042-6989/96 $15.00 + 0.00

Isolating Excitatory and Inhibitory Nonlinear Spatial Interactions Involved in Contrast Detection*

BARBARA ZENGER,† DOV SAGI‡

Received 15 March 1995; in revised form 24 October 1995

Interactions between filters tuned to different orientations and spatial locations were investigated with a masking paradigm. Targets were masked by pairs of Gabor signals presented either at a different orientation ($\pm\Delta\theta$) or at a different spatial location ($\pm\Delta y$). The two mask components were either of equal phase or of opposite phase to each other. Detection thresholds of the target were measured as a function of mask contrast. Typically, the curves obtained showed the following behavior: for increasing mask contrast the threshold first decreased, then reached a minimum and then increased linearly on a log-log scale reflecting a power-law behavior. Mask pairs of equal phase as well as pairs of opposite phase were shown to facilitate detection. Facilitation by mask pairs of equal phase was larger (up to 0.4 log units) and decreased for increasing $\Delta\theta$ and Δy. The facilitation for mask pairs of opposite phase ($\sim$0.1 log units) was observed only for larger $\Delta\theta$ and Δy. Phase independent suppression was observed with higher mask contrasts at smaller $\Delta\theta$ and Δy. The strength of this suppression was shown to decrease with practice. We account for the observed facilitation with an accelerating transducer function applied on a second-stage filter. Suppression is modeled with an additional inhibitory second stage filter that divides the output of this transducer. Selective reduction of the inhibitory gain accounts for the practice effects. Copyright © 1996 Elsevier Science Ltd.

Contrast discrimination Gain control Learning Masking Second stage filters

INTRODUCTION

A widely used model for early visual processing suggests linear filtering of the image as a first processing stage. Many spatially local filters, each selectively tuned to a specific orientation and spatial frequency, are assumed to act in parallel over the whole visual field. Psychophysical evidence for this model is provided by a variety of experimental paradigms, such as selective adaptation (Blakemore & Campbell, 1969; Blakemore & Nachmias, 1971), simultaneous masking (Campbell & Kulikowski, 1966; Legge & Foley, 1980) and sub-threshold summation (Kulikowski *et al.*, 1973). Filters, though followed by a nonlinear transducer function, have been treated as linear and independent. However, independence holds only to a first approximation and interactions between filters with different tuning properties have been described. Olzak and Thomas (1991, 1992) demonstrated that information from tuned pathways is not always used directly in making spatial judgments, but in some case is combined across wide regions of the Fourier domain prior to the discrimination decision. Lateral inhibition between orientation detectors was suggested as a mechanism that can account for the apparent tilt of a line in the presence of a line of somewhat different orientation (Blakemore *et al.*, 1970; Carpenter & Blakemore, 1973) or after adaptation to a line of somewhat different orientation ("tilt aftereffect") (Magnussen & Kurtenbach, 1980; Kurtenbach & Magnussen, 1981). Inhibitory and facilitatory interactions were found between neighboring filters on the spatial (Sagi & Hochstein, 1985; Polat & Sagi, 1993; Polat & Norcia, 1995) and spatial frequency (Tolhurst & Barfield, 1978) dimensions, possibly accounting for human performance on texture segmentation (Rubenstein & Sagi, 1990) and perceptual grouping (Ben-Av & Sagi, 1995) tasks.

Interactions between filters can be studied with contrast masking experiments (Tolhurst & Barfield, 1978; Polat & Sagi, 1993; Foley, 1994a). In these experiments, contrast thresholds for a target are measured in the presence of a pattern (mask). Nonlinear masking effects can be quantitatively characterized by the curves

*Part of this paper was presented at the 17th ECVP conference, Eindhoven, The Netherlands (September 1994).

†Section of Visual Science, Department of Neuroophthalmology, Waldhörnlestr. 22, D72076 Tübingen, Germany.

‡To whom all correspondence should be addressed at the Department of Neurobiology, Brain Research, The Weizmann Institute of Science, Rehovot 76100, Israel. [*Email* dubi@nisan.weizmann.ac.il; *Tel* 972-8-343747; *Fax* 972-8-344140].

describing target thresholds as a function of mask contrast (pedestal). For increasing pedestal contrast, thresholds typically first decrease, then reach a minimum and then increase linearly on a log-log scale [reflecting a power-law behavior (Legge, 1981; Swift & Smith, 1983)]. "Dipper"-shaped curves of this type have been described in many studies (Nachmias & Sansbury, 1974; Legge & Foley, 1980; Wilson, 1980; Bradley & Ohzawa, 1986; Ross & Speed, 1991; Foley, 1994a) and, with target and mask having the same orientation and spatial frequency, they were used to derive nonlinear contrast response functions (Nachmias & Sansbury, 1974; Wilson, 1980). These functions typically have a positive second derivative (i.e. acceleration) at low stimulus contrasts and a negative second derivative (i.e. suppression, compression) at high contrasts.

A simple model for contrast detection assumes that target detection is mediated by a single filter, the most sensitive for the target. Masks that are presented within the bandwidth of this filter provide some input to it and thus shift the operating point on its transducer function (Legge & Foley, 1980; but see Nachmias, 1993). The predicted curves will be dipper-shaped but, since the masks contribute only a certain ratio of their contrast to the target filter, the curves will be scaled (or shifted on a logarithmic scale). Results reported in the literature do not follow this prediction. Detection thresholds of gratings were measured in the presence of mask gratings of various contrasts and orientations (Campbell & Kulikowski, 1966; Ross & Speed, 1991). These studies show that the facilitatory effect is tuned narrowly as practically no facilitation is observed when masks differ from the target by more than 10–15 deg. Foley (1994a) measured detection thresholds of Gabor patches that were masked by gratings of different orientations. His results also show a reduction of facilitation for increasing orientation difference between target and mask. A reduction of facilitation is observed also for masks that differ from the target in spatial frequency (Tolhurst & Barfield, 1978; Legge & Foley, 1980; Ross & Speed, 1991). Legge and Foley (1980) accounted for the reduction of facilitation by assuming that response pooling across spatial filters is effective only at low mask contrasts and not at high mask contrasts. Ross and Speed (1991) developed a quantitative model in which they assume that masks have two effects differing in bandwidth: first, they directly stimulate the detecting mechanism (narrow tuning) and secondly, they shift the contrast response function towards higher mask contrasts (broad tuning). Their model represents a parametric description of their data, but mechanisms are not suggested. Foley (1994a) accounts for facilitation with an accelerating transducer function and for suppression with broad-band divisive inhibition. His model is similar to a model for cat striate cell responses proposed by Heeger (1992).

Results from previous masking studies do not allow us to separate local spatial interactions from orientation dependent interactions, as these studies used wide field masking gratings. In fact, Foley (1994a) notes that Gabor masks can lead to a larger facilitation than grating masks, making the additional spatial masking effect of gratings evident. In the experiments described here, this problem was avoided by using localized target and mask stimuli. A Gabor target was masked with two Gabor signals differing from the target either in orientation ($\pm\Delta\theta$) or in spatial location ($\pm\Delta y$). We further tried to isolate two different processing stages at which masks can affect detection. Masks can provide direct input to the target filter (as it was assumed in the single-filter model) or they may affect detection indirectly by stimulating another filter which then interacts with the target filter. In order to separate these two types of processes, the experiments described here were performed for two different mask-phase relationships; the two masks were presented either with equal phase or with opposite phase (see Methods section). For masks of equal phase, direct as well as indirect masking effects should be observed. Masks of opposite phase cancel each others input to the target filter and therefore do not affect detection "directly". In this condition, only indirect masking effects are expected.

Phase dependency of masking effects on grating detection was investigated recently by Lawton and Tyler (1994). Their results show that suppression of detection does not depend on whether the mask grating is presented in phase or in quadrature (90 deg) phase shift with the target, a finding that may indicate a major "indirect" masking source. As a possible explanation for their experimental observation they suggest that the "self-masking effect is pooled over a local region of cells of various positions and types" (including, in particular, cells sensitive to different phases). Foley (1994b) finds phase independence of masking effects at high mask contrast (suppression), but not at low mask contrast (facilitation), indicating nonlinear inhibition. Morgan and Dresp (1995), using a luminance detection task in the presence of a lateral mask also failed to find (in two out of three observers) detection facilitation for mask and target of opposite contrast polarity.

An intriguing aspect of the masking literature is the reports on interobservers' variability of the experimental results (Morgan & Dresp, 1995; Olzak & Thomas, 1992) and of practice effects (Swift & Smith, 1983). Such behavior can be accounted for by plasticity of the mechanisms involved in the masking process, in agreement with recent experimental results (Karni & Sagi, 1991; Polat & Sagi, 1994b) indicating long-term modifications in early stages of visual processing. In particular, it is possible that filters involved in the detection task are modified due to the presence of the mask or, alternatively, interactions may change with time and may depend on the observer state of experience. The results of the experiments described here allow for a rough characterization of the filters involved and their nonlinear interactions. The data make it further possible to separate inhibitory interactions that account for suppression from excitatory interactions that account for facilitation. We show also that results change with

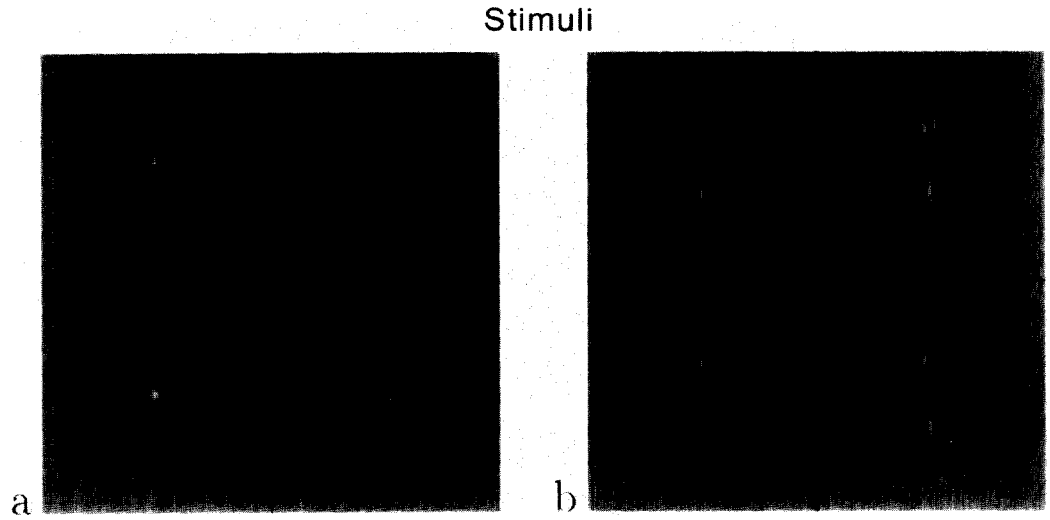

FIGURE 1. (a). Example of stimuli used in the orientation masking experiments where masks differed from the target in orientation ($\Delta\theta = 45$ deg). Masks of equal phases are presented in the upper quadrant on the left-hand side. Masks of opposite phases, with the phase of one of the two mask signals reversed, can be seen in the upper quadrant on the right-hand side. In the lower quadrants, a vertical target is added to the mask stimulus. The observers task is to detect this target. (b) Same as (a), but with spatially displaced masks ($\Delta y = 3\lambda$).

practice and point to plasticity of specific interactions in accounting for the learning effects.

METHODS

Apparatus

Stimuli were displayed as a gray-level modulation on an Hitachi HM-3619A color monitor, using an Adage 3000 raster display system. The video format was 56 Hz noninterlaced, with 512×512 pixels occupying a 9.6×9.6 deg area. The mean luminance was 50 cd/m^2. Stimulus generation was controlled by a Sun-3/140 workstation and the stimulus display by the Adage local processor. The stimuli were viewed from a distance of 1.5 m.

Stimuli

Stimuli consisted of one target signal and two mask signals. The spatial luminance distribution of target and mask signals is described by a Gabor function, which can be interpreted as a cosine grating with its amplitude modulated by a Gaussian envelope:

$$G_{y_0,\theta_0}(x,y) = cos\left(\frac{2\pi}{\lambda_0}((x-x_0)cos\theta_0 + (y-y_0)sin\theta_0)\right) \times exp\left(-\left(\frac{(x-x_0)^2 + (y-y_0)^2}{\sigma^2}\right)\right),$$

with x and y being the horizontal and vertical coordinates. The spatial location of the Gabor signal is determined by x_0 and y_0, its orientation by θ_0 and its wavelength by λ_0. The standard deviation of the Gaussian envelope is given by σ. For all stimuli used in these experiments, $\lambda_0 = 0.15$ deg, x_0 (at the center of the screen) and $\sigma = \lambda_0$ were kept constant. The target signal was presented at the fixation point (x_0, y_0), with either vertical ($\theta_0 = 0$ deg) or horizontal ($\theta_0 = 90$ deg) orientation.

Two different sets of experiments were performed. In the first set, mask signals and target signal were presented at the same location, but mask orientation differed from the target orientation by $\pm\Delta\theta$. The luminance distribution was thus:

$$L(x,y) = C_t G_{y_0,\theta_0} + C_m(G_{y_0,\theta_0+\Delta\theta} + \varphi G_{y_0,\theta_0-\Delta\theta})/2,$$

with C_t as target amplitude, C_m as mask amplitudes and φ as the relative polarity of the second mask (being 1 or -1 for same and opposite phase patterns). In the second set of experiments mask and target orientation were the same, but the mask was displaced vertically by $\pm\Delta y$. Here the luminance distribution was given by:

$$L(x,y) = C_t G_{y_0,\theta_0} + C_m(G_{y_0-\Delta y,\theta_0} + \varphi G_{y_0-\Delta y,\theta_0})/2.$$

Both sets of experiments were performed in two conditions: the two mask components were either of equal contrast polarity ($\varphi = 1$) or of opposite polarity ($\varphi = -1$), with C_t, $C_m \geq 0$. Examples of stimuli presented in the experiments are shown in Fig. 1.

Experimental procedures

A two alternative forced choice procedure was used. Observers activated a trial sequence by pressing a key, after fixating a small cross in the center of the screen. Each trial consisted of a blank period of 500 msec, followed by two sequential stimulus presentations (90 msec each) that were separated by 1000 msec. Only one of the two stimulus presentations contained the target (but both contained the mask). The stimulus intervals were marked by two peripheral high contrast crosses. Observers had to determine which of the two presentations contained the target. The decision was indicated by pressing a key and auditory feedback was given for incorrect response.

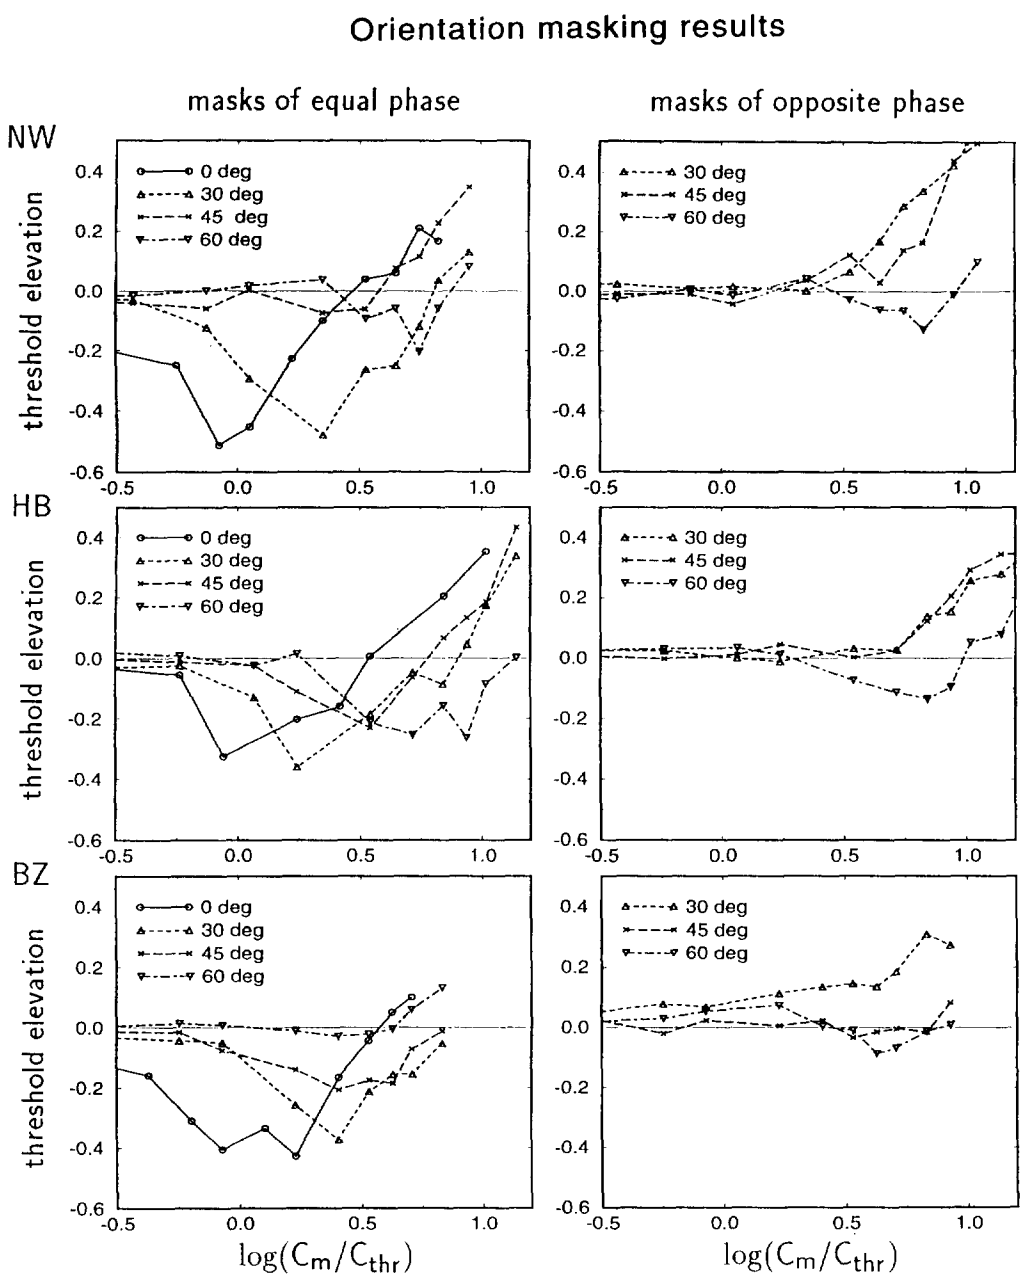

FIGURE 2. Detection thresholds were measured as a function of mask contrast for different mask orientations and mask-phase relationships (both detection thresholds and mask contrasts are normalized to the observer threshold average C_{thr}). Each datum point is the average of several threshold estimates (at least three, on average five–six). Results are presented for three different observers. The magnitude of facilitation by masks of equal phase decreases for increasing orientation difference between target and mask. Masks of opposite phase can facilitate and suppress detection.

Detection thresholds for the target were estimated using the following staircase procedure: C_t is increased by 0.1 log units after every incorrect response and decreased by 0.1 log units after three consecutive correct responses. A block was terminated after 10 reversals of C_t and the geometric mean of the last eight reversal points was used as a threshold estimate. This staircase procedure was shown to converge to a level of 79% correct (Levitt, 1971). Apart from C_t, all stimulus parameters were kept constant within one block. During one session (which lasted approximately 50 min) and between different blocks mask amplitudes were varied while $\Delta\theta$, Δy and all Gabor phases were kept constant.

Observers

Five observers (including the first author) with normal or corrected to normal vision took part in the experiments. Four observers performed the orientation masking experiments, that included seven different conditions (masks of equal phase for $\Delta\theta = 0$, 30, 45 and 60 deg, masks of opposite phase for $\Delta\theta = 30$, 45 and 60 deg). For two observers (HB, BZ) the target was horizontal, for one observer (AD) it was vertical and another observer (NW) performed both sets of experiments. Three observers (AD, AL and BZ) participated in the spatial masking experiments that also included seven conditions (masks of equal phase for $\Delta y = 0\,\lambda$, $2\,\lambda$, $3\,\lambda$ and $4\,\lambda$, masks of

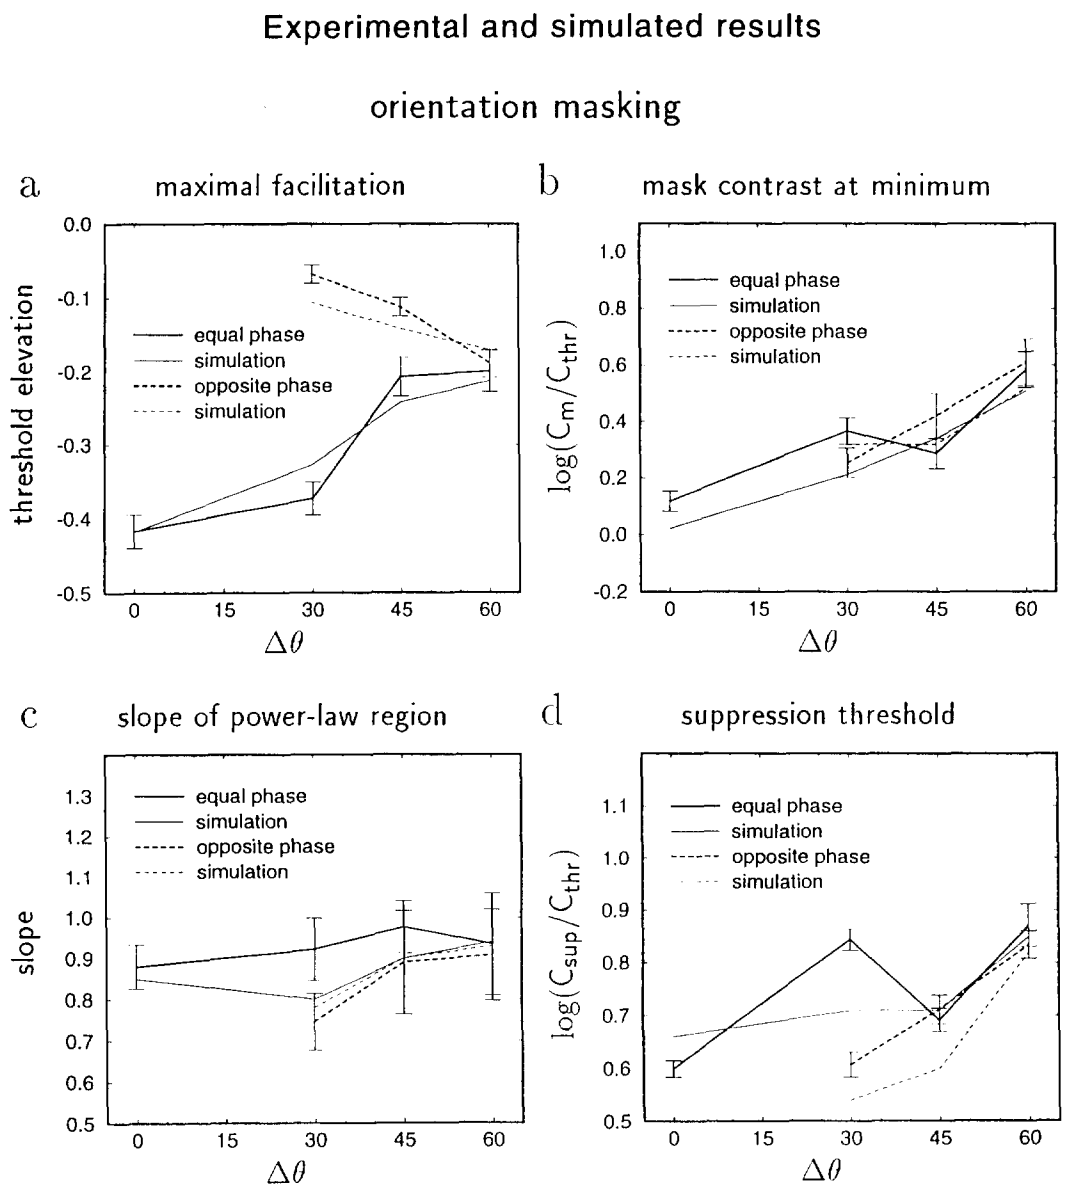

FIGURE 3. Data analysis for the orientation masking experiments. For each session, several parameters were estimated. The values of these parameters are presented here as a function of $\Delta\theta$ for masks of equal phase and for masks of opposite phase. The average over all sessions and observers is presented. Error bars indicate standard errors. The psychophysical results can be compared directly with the results of a model simulation.

opposite phase for $\Delta y = 2\lambda$, 3λ and 4λ). One of the observers (AD) did not perform the two conditions at 4λ. In the spatial masking experiments the target was always vertical.

RESULTS

Detection thresholds of a Gabor target were measured as a function of mask contrast C_m. The masks differed from the target either in orientation or in spatial location and the two mask components were either of equal phase or of opposite phase.

Orientation masking

Data for the orientation masking experiments are presented in Fig. 2 for three different observers. As is evident, the curves show the following general behavior: for increasing mask contrast the thresholds first decrease, then reach a minimum and then increase linearly on a log-log scale (which corresponds to a power-law behavior). The magnitude of maximal facilitation and the mask contrast at which the minimum occurs depend on the mask orientation and on the mask-phase relationship. Though individual differences between observers can be seen, in all cases the result pattern clearly deviates from the prediction of the single filter model. For masks of equal phase, the curves are not simply shifted relative to each other, but the magnitude of facilitation decreases for increasing orientation difference between target and mask. Masks of opposite phase cancel each others input to the target filter and the single-filter model, assuming detection by the most sensitive filter, would predict that this mask pattern does not affect detection. However, the

psychophysical results show that masks of opposite phase can suppress and facilitate detection.

In order to make the data more easily accessible for analysis, the following data features were extracted from each session:

(i) The threshold of the isolated target; the average detection threshold C_{thr} of each observer was used for normalization;

(ii) The minimum C_{min} (= lowest threshold across the C_{mask} range obtained in a session); the threshold elevation log (C_{min}/C_{thr}) served as an estimation of *maximal facilitation*;

(iii) The mask contrast C_m at which the minimum occurred; log (C_m/C_{thr}) gives an estimate of the *mask contrast at minimum* (cases where $C_m = 0$ were not considered);

(iv) The *slope of the power-law region* was estimated by fitting a line (on log-log scale) through all the data points of the power-law region; the beginning of this region was defined as the lowest mask contrast from which on all threshold estimates were at least 0.1 log units above the minimum;

(v) All data points of the power-law region were fitted by a line of slope 0.89 (which was obtained as the average value); the mask contrast C_{sup} at which the fitted line equals the observer's threshold average C_{thr} reflects the mask contrast at which masks start to suppress detection; log (C_{sup}/C_{thr}) was thus used as an estimate of the *suppression threshold.*

For each condition (masks of equal phase: 0, 30, 45 and 60 deg; masks of opposite phase: 30, 45 and 60 deg) the parameters described above were averaged across all observers and all sessions. The results are shown in Fig. 3.

Both masks of equal phase and masks of opposite phase can facilitate detection [Fig. 3(a)]. For masks of equal phase the facilitation decreases with increasing $\Delta\theta$. A particularly strong decrease is observed between $\Delta\theta = 30$ deg and $\Delta\theta = 45$ deg. Interestingly, masks of opposite phase can also enhance target sensitivity. The magnitude of this facilitation is smaller and increases with increasing orientation difference. For $\Delta\theta = 60$ deg the facilitation effect is independent of mask phase. (Note that, due to noise in the data, maximal facilitation is somewhat overestimated.)

As described, maximal facilitation was estimated separately for each session as we also wanted to analyze practice effects. The method has the disadvantage that noise in the data alone can produce minima below threshold. In order to show that the observed facilitation for masks of opposite phase is real we selected for each subject the region that included those two tested mask contrasts where the average facilitation (across all sessions) was maximal. For masks of opposite phase at $\Delta\theta = 60$ deg four out of five observers showed significant facilitation in this region. Interestingly, one of these observers had comparatively strong suppression in the first four sessions and shows significant facilitation only in the last three sessions. The development of facilitation with practice is well consistent with the practice-dependent decrease in the suppression thresholds that we observed (see "Practice effects" section, below) and it might further explain why one observer (who performed only in three sessions in this condition) had practically no facilitation.

Though the magnitude of facilitation depends (in most cases) on mask phase, the mask contrast at which the minimum occurs appears to be mask phase independent [Fig. 3(b)]. With increasing $\Delta\theta$, the minima (and the start of the power-law region) shift towards higher mask contrasts. Such a shift is also predicted by the single filter model.

The slope of the power-law region is practically the same in all conditions with an average value of 0.89 [Fig. 3(c)].

The suppression threshold is the mask contrast at which mask presentation starts to suppress target detection. In general, the suppression threshold seems to increase for increasing $\Delta\theta$. However, there is one interesting exception: for masks of equal phase, the suppression threshold for $\Delta\theta = 45$ deg is significantly lower than for $\Delta\theta = 30$ deg. This corresponds well to the fact that facilitation for $\Delta\theta = 45$ deg is much weaker than for $\Delta\theta = 30$ deg while the minima occur at very similar mask contrasts. Since the functions rise with the same slope, the suppression threshold for $\Delta\theta = 45$ deg is expected to be smaller. In a separate analysis, suppression thresholds were found to increase significantly with practice (see "Practice effects" section). Therefore, the average values that are presented in Fig. 3(d) have to be treated with caution.

For $\Delta\theta = 60$ deg none of the parameters showed significant phase dependency.

Spatial masking

The results of the detection threshold measurements in the presence of spatially displaced masks are presented in Fig. 4 for three different observers. Curves for masks of equal phase at a distance of 2λ appear to be shifted relative to the curve for masks presented at target location—as is expected if the single-filter model is valid. For masks of equal phase at larger distances (3λ and 4λ), the behavior is less clear. There is quite strong facilitation but no evident power-law region. The absence of the power-law region can also be seen in the results for masks of opposite phase (however, some suppression is observed for masks at 2λ). Masks of opposite phase at larger spatial distances facilitate detection. Facilitation at 3λ and 4λ (again averaged over a region including two mask contrasts) is significant for all observers, showing that the single-filter model also fails to account for the spatial masking experiments.

An analysis similar to the one performed for the orientation masking data was also carried out for the spatial masking experiments. However, as many curves

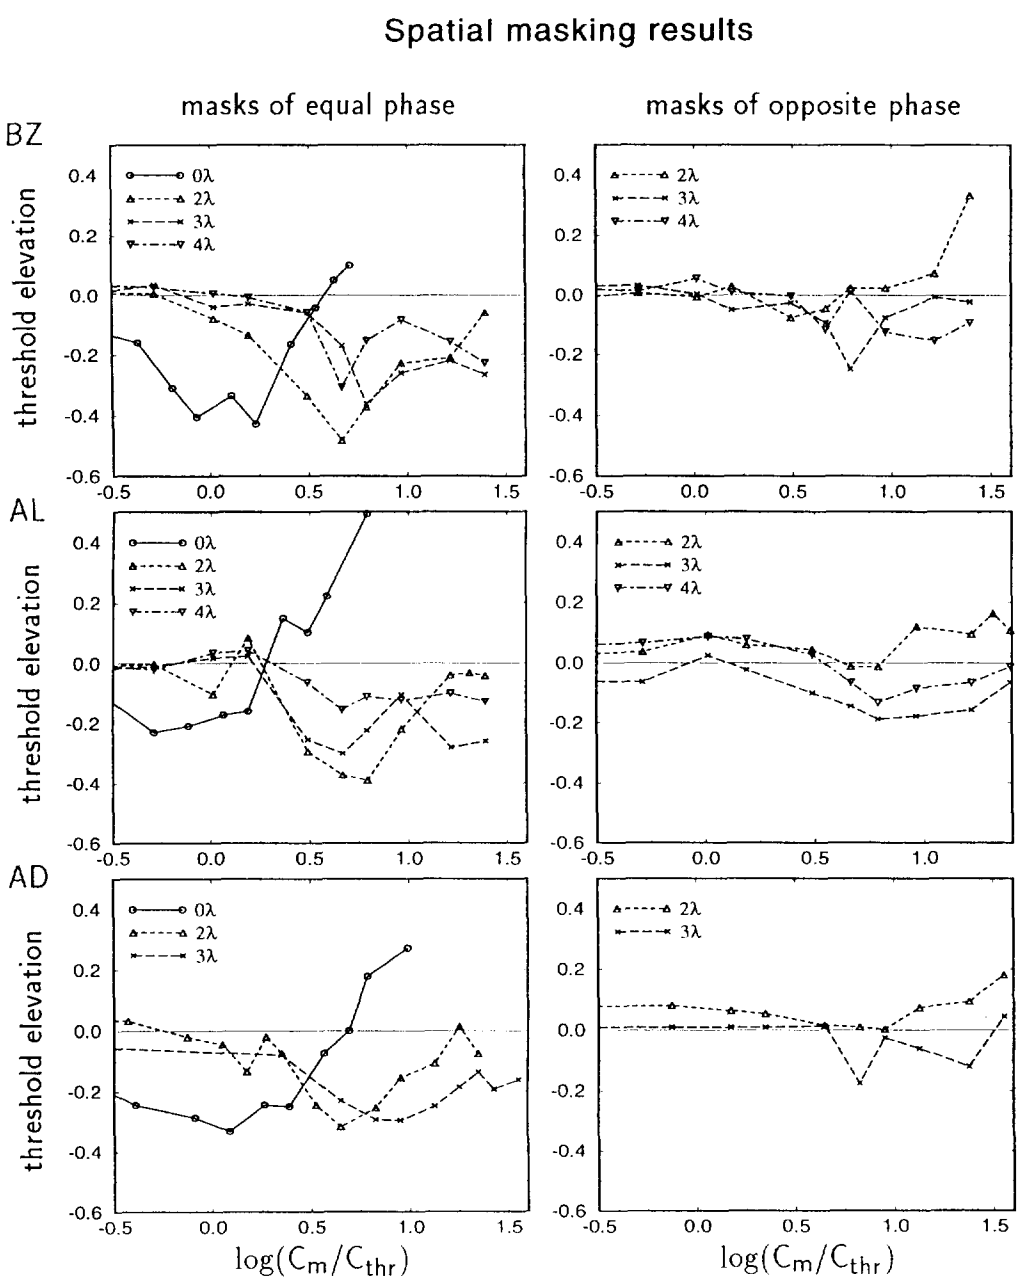

FIGURE 4. Detection thresholds were measured as a function of mask contrast for different mask distances (given in units of the Gabor wavelength) and mask–phase relationships. Each datum point is the average of several threshold estimates (at least two, on average four). Results are presented for three different observers. The curves for masks of equal phase at 2λ are shifted relative to the curves for masks at 0λ. At larger distances (3λ and 4λ) the behavior is less clear. There is quite a strong facilitation but often no evident power-law region. Masks of opposite phase at larger distances ($>3\lambda$) can facilitate detection.

did not show a clear power-law region, no estimates were obtained for the slope of the power-law region and for suppression thresholds. Figure 5 shows the results of this analysis. As already pointed out, facilitation can be observed for masks of both phase relationships. Interestingly, the minima occur at very high mask contrasts compared to those of the orientation masking results. This might partly explain the absence of the power-law region for masks at large spatial distances. Perhaps a power-law behavior could be observed if higher mask contrasts could be tested. Furthermore, it is possible that the actual minima occur at these higher mask contrasts and that the actual facilitation is larger than the one estimated here.

Practice effects

As mentioned before, some of the measured parameters appeared to change during the course of the experiments. To test this phenomenon, the parameter stability was analyzed. The analysis was done only for the orientation masking experiments since here the average session number per experimental condition was large enough to find significant effects.

As the experiments were not designed originally to study temporal changes of the result patterns, the performance order of the various conditions was not systematic, thus placing limitations on the information that can be extracted from the data. In the analysis carried

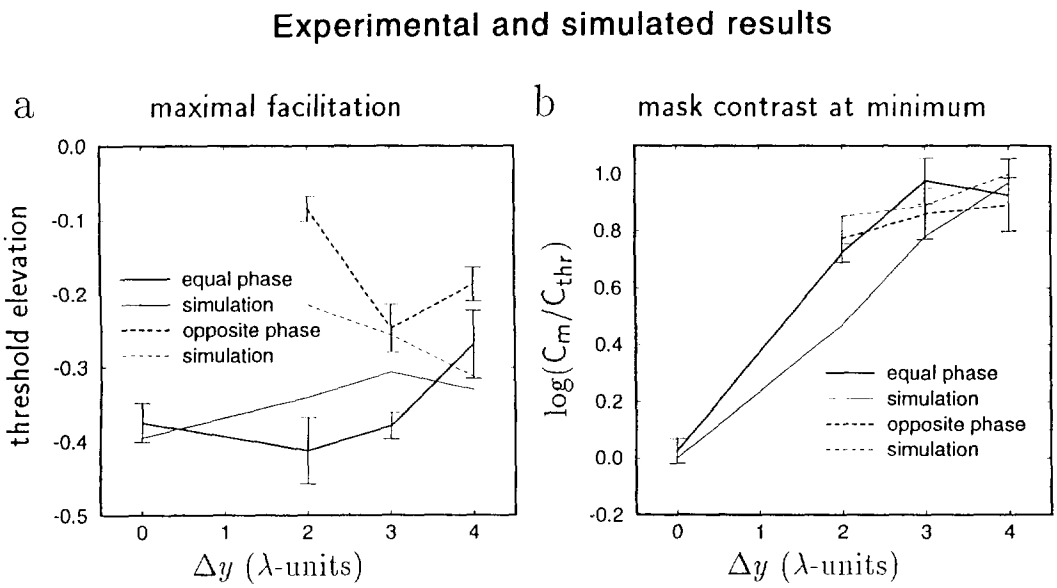

FIGURE 5. Same as Fig. 3(a) and (b), but for spatial masking experiments.

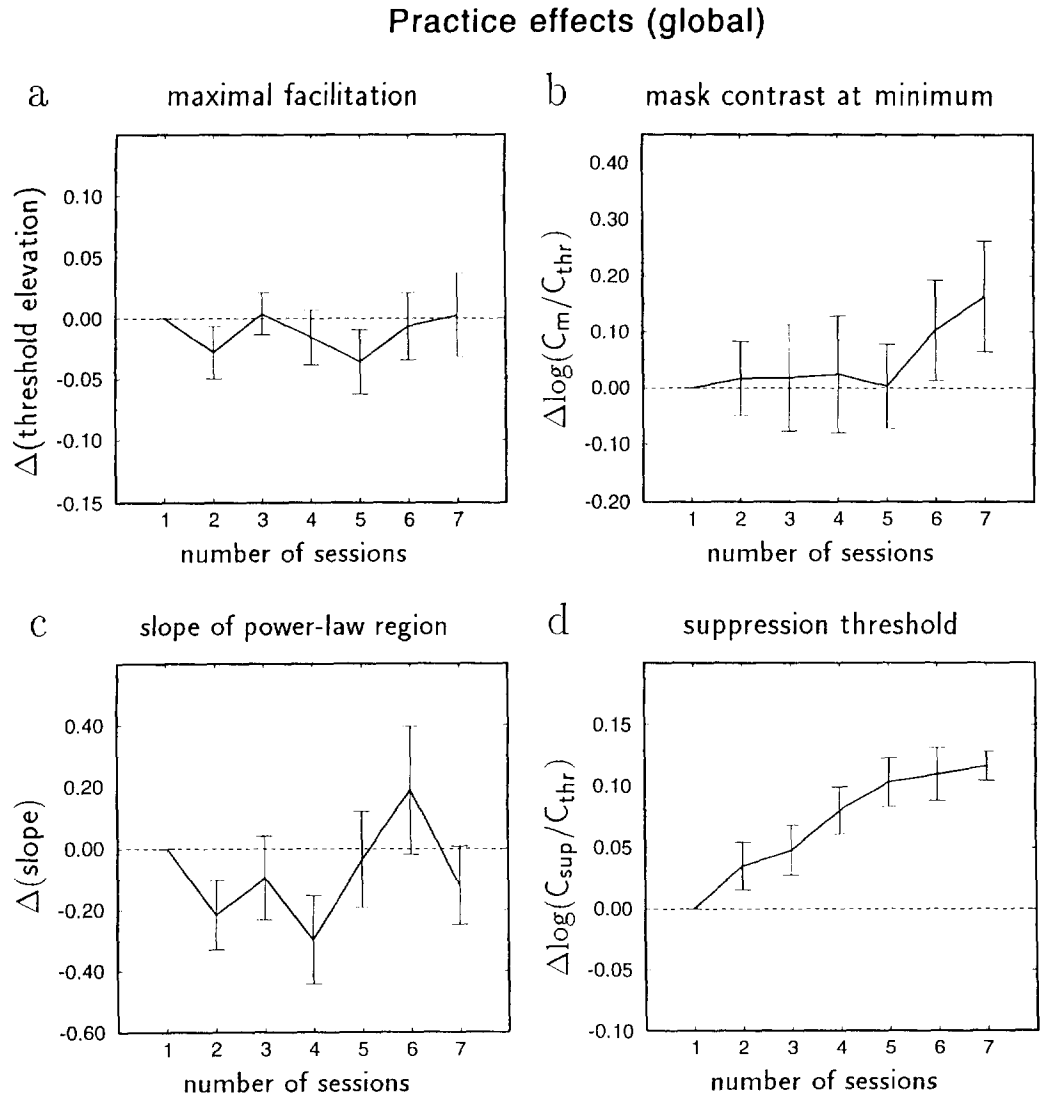

FIGURE 6. Effects of practice on performance. For each condition, the parameter value is plotted as a function of session number in that condition (normalized to the value obtained in the first session). Global averages across all conditions and observers are presented. Error bars indicate standard errors. Suppression thresholds are shown to increase with practice, reflecting a performance improvement for high contrast masks.

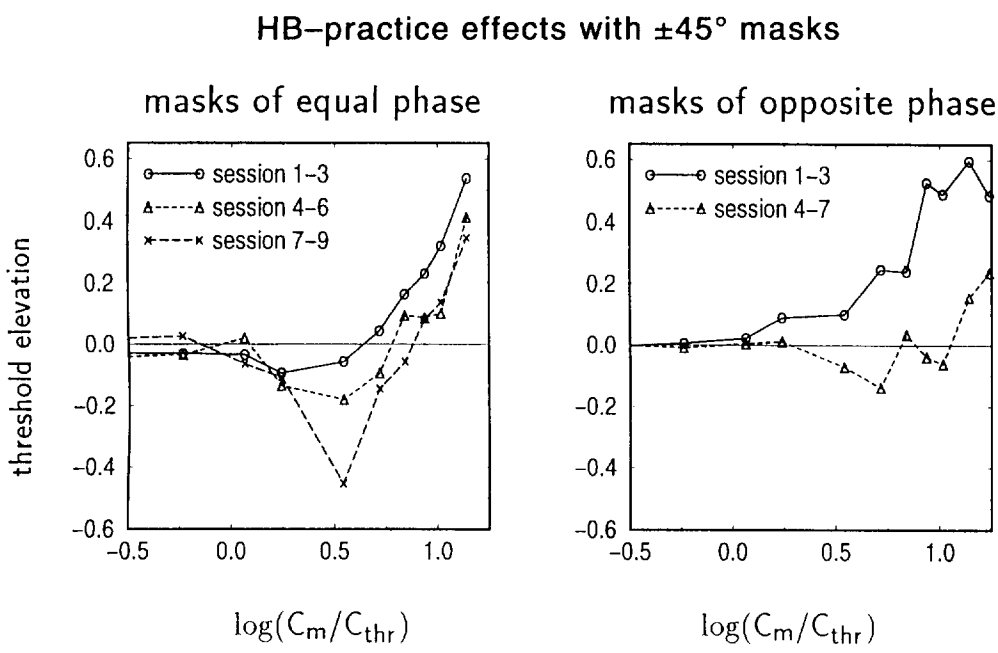

FIGURE 7. Practice effects of observer HB for masks of equal and of opposite phase at $\Delta\theta = 45$ deg. Detection thresholds were measured as a function of mask contrast. Each curve represents the average of three or four sessions. With practice, the suppression threshold increases and facilitation is developed.

out here, each condition was considered separately while the absolute ordering of conditions was ignored. Parameters were normalized to the value obtained during the first session in that condition. The normalized values were plotted vs time (where number of sessions always refers to the number of sessions in the respective condition). The results of this analysis (averaged across all observers and all conditions) are presented in Fig. 6. A strong practice effect was found for suppression thresholds which increased with time (linear correlation: $P > 0.001$). The suppression threshold increase reflects a decrease in contrast detection thresholds (improvement) for targets masked with high contrast masks. The threshold improvement was slow and continued for at least several sessions. The effect is consistent across observers (with the exception of observer BZ, who was highly trained on contrast detection tasks) and it also appeared to be rather consistent across experimental conditions.

The practice effect was exceptionally strong for one observer (HB) at $\Delta\theta = 45$ deg. The respective curves are presented in Fig. 7, where each line is the average of three or four sessions. It should be noted that the sessions were not performed sequentially, but that other conditions were tested in between. The increase in the suppression threshold was highly significant ($P > 0.001$) for both mask patterns and is combined here with a significant increase in facilitation. Note that for mask components of opposite phase an initial suppression turns after three sessions into enhancement with target thresholds decreasing to less than half of their initial values. A "development of facilitation" was also seen in a few other cases, e.g., for masks of opposite phase at $\Delta\theta = 60$ deg. At a mask contrast of approx. 0.7 log units above threshold, the four observers that had between six and seven sessions in that condition have insignificant suppression for the first four sessions [threshold elevation = 0.013 ± 0.033 (SE) log units], but have a very clear facilitatory effect of −0.15 ± 0.022 (SE) log units for the remaining two to three sessions.

Summary

The results show that masks of equal phase can facilitate detection. The magnitude of this facilitation was found to decrease with increasing $\Delta\theta$ but it decreases less with increasing Δy. Facilitation was also observed for masks of opposite phase when they were presented at larger orientation differences and spatial distances. For high contrast masks detection thresholds can be described with a power-law (with the exception of masks presented at large spatial distances). The pattern of results was found to change with practice. Suppression thresholds increase with practice, reflecting a performance improvement for high contrast masks. In some cases, enhancement was shown to increase dramatically with practice, reflecting a performance improvement for low to medium contrast masks.

A TWO-STAGE FILTERING MODEL FOR DETECTION

Architecture

The data presented above provide further evidence for the inadequacy of models assuming linear filters tuned to different orientations and spatial locations in accounting for human detection data.

Such models predict that the maximal facilitation is independent of $\Delta\theta$ and Δy (for masks of equal phase). Our results do not follow this prediction and indicate that

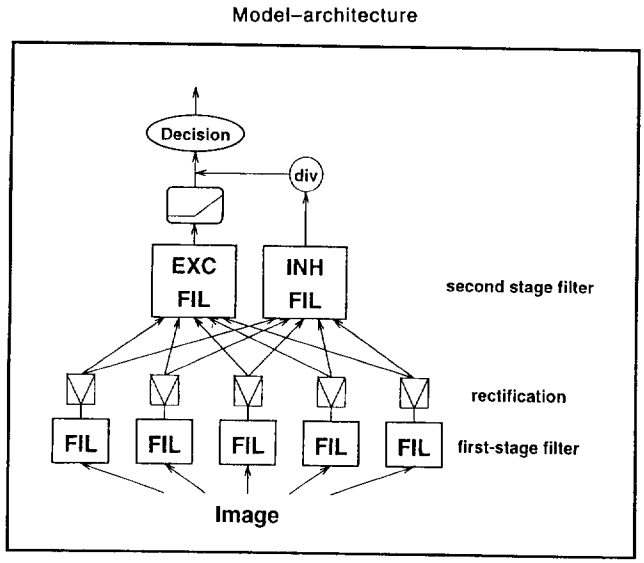

FIGURE 8. The model consists of two filtering stages: linear filtering of the image is followed by a full-wave rectification. The first-stage output provides input to two second-stage filters. The excitatory filter is followed by an accelerating transducer function that accounts for the facilitation observed for masks of both mask–phase relationships. The inhibitory second-stage filter accounts for the observed suppression by dividing the excitatory transducer output.

the facilitatory and suppressive effects have different tuning behavior. In the model presented here, therefore, we attribute facilitation and suppression to two different filters, allowing us to define the tuning for both effects separately. One of the filters, which is called the "excitatory filter", is followed by an accelerating transducer function, leading to facilitation (Nachmias & Sansbury, 1974; Legge & Foley, 1980; Wilson, 1980; Ross & Speed, 1991; Foley, 1994a). The output of this filter is divided by the output of the "inhibitory filter", leading to a compression of the resulting transducer function that can account for the power-law region (Foley, 1994a).

Our second important finding is that facilitation and suppression can both be observed also for masks of opposite phase. Since first-stage filters are insensitive to masks of opposite phase these results require a second processing stage to be incorporated into the model. The excitatory as well as the inhibitory filter are consequently described as second-stage filters. The two processing stages are separated by a nonlinearity. We find full wave rectification at the first-stage filter output sufficient in accounting for phase independence.

In short, the following processes are suggested (for a schematic diagram of the model see Fig. 8):

(i) Linear filtering of the image with Gabor (or alike) filters;

(ii) Full-wave rectification of the filter output;

(iii) Two second-stage filters that integrate over the first-stage output:
—the "excitatory" filter is followed by an accelerating transducer function;
—the "inhibitory" filter divides the transducer output;

(iv) A decision that is based on the divided output signal.

Tuning

The model should describe the observed tuning of facilitation and suppression. Basically, all model parameters influence its tuning behavior; however, the following parameters appear to be of particular importance:

(i) Bandwidths of the first-stage filters;

(ii) Shape and bandwidth of the "excitatory" second-stage filter;

(iii) Shape and bandwidth of the "inhibitory" second-stage filter.

The bandwidth of the first-stage filter affects mask-phase dependent differences. In the orientation masking experiments, results for $\Delta\theta = 60$ deg appeared to be mask-phase independent reflecting a small first-stage filter (full) bandwidth of less than 30 deg, in agreement with earlier studies (Campbell & Kulikowski, 1966; Phillips & Wilson, 1984). The spatial masking results show that masks presented at a target-to-mask distance of 2λ shift the minimum towards high mask contrasts, suggesting that the excitatory input from 2λ is rather small. Therefore, a small spatial bandwidth for both the first and second stage (excitatory) filters is implicated (see Table 1 for model parameters).

The shape of a specific masking curve depends on the sensitivity of the two second-stage filters to the particular mask configuration. It turns out that the magnitude of the maximal facilitation is determined by the ratio of the mask input to the excitatory and the inhibitory filter, whereas the mask contrast at which the minimum occurs is, to a good approximation, determined by the mask input to the excitatory filter.* Our data for maximal facilitation and mask contrast at minimum for different

TABLE 1. Parameter values used for the simulation

Model parameters		
σ_{or}	10 deg	Orientation: bandwidth of first-stage filter
α_1	0.4	Orientation: excitatory second-stage filter
$\sigma_{or;exc}$	40 deg	
α_2	0.3	Orientation: inhibitory second-stage filter
α_3	0.2	
θ_{inh}	45 deg	
$\sigma_{or;inh}$	5 deg	
σ_{sp}	$0.9\,\lambda$	Space: bandwidth of first-stage filter
β_1	0.06	Space: excitatory second-stage filter
$\sigma_{sp;exc}$	$4\,\lambda$	
β_2	0.07	Space: inhibitory second-stage filter
$\sigma_{sp;inh}$	$2\,\lambda$	
c	2.5	Parameters of accelerating transducer function
μ	1.0	
n	4.0	

mask configurations consequently allow for an estimation of the second-stage filter shapes.

The observed facilitation at large orientation differences such as 60 deg suggests that the excitatory filter integrates over a broad range of orientations. The reduction of maximal facilitation with increasing $\Delta\theta$ shows that the inhibitory second-stage filter is more broadly tuned than the excitatory second-stage filter. However, monotonically decreasing broad-band inhibition cannot account for the observed results in a quantitative way, as a strong reduction of facilitation was observed for masks of equal phase at $\Delta\theta = 45$ deg. Within the present theoretical framework, two alternative accounts can be offered: a decrease in excitatory input, or an increase in inhibition. A decrease in the excitatory input (for $\Delta\theta < 30$ deg) would account for the reduction of maximal facilitation, but it would also lead to an enormous shift of the minimum towards higher mask contrasts. This shift is not seen in the experimental data. Therefore, we account for the reduction of facilitation at $\Delta\theta = 45$ deg with increasing the inhibitory input, suggesting side inhibition from around $\Delta\theta = 45$ deg.

The results of the spatial masking experiments also suggest integration over a large range of spatial distances, as facilitation for masks of both phase-relationships is observed at large distances. The reduction of maximal facilitation is small as compared with the results for orientation masking, implying a rather small inhibitory input to account for the observed power-law behavior at small spatial distances.

Model simulation

The model behavior was tested with a computer simulation. In order to keep the number of free parameters small, the filters were described only as one dimensional filters, separately defined for orientation and space (only one spatial dimension). We use linear filters with Gaussian sensitivity profiles in space and in orientation. They were modeled as:

Orientation:

$$F_i(\theta) = G(\theta - \theta_i^0 \mid \sigma_{or}) \quad for \quad \theta_i^0 = \ldots, -15°, 0°, 15°, 30°,$$

Space:

$$F_j(y) = G(y - y_j^0 \mid \sigma_{sp}) \quad for \quad y_j^0 = \ldots, -1\lambda, 0\lambda, 1\lambda, 2\lambda, \ldots$$

$$with \quad G(x \mid \sigma) = exp\left(\frac{-x^2}{2\sigma^2}\right).$$

Second-stage filters are assumed also to be linear. Although it would seem natural to describe the excitatory second-stage filter with a Gaussian function, test simulations showed that the model could be improved significantly with an additional excitatory input from the first-stage target filter added. This additional term in the excitatory filter description may indicate two mechanisms involved in the excitatory process, one being a 'self-excitation' and the other providing lateral integration (a hint toward a single-layer feed-back network). The mathematical description of these filters is:

Orientation:

$$F_{\epsilon_i}(\theta) = \begin{cases} 1 & for \quad \theta = \theta_{\epsilon_i}^0 \\ \alpha_1 G(\theta - \theta_{\epsilon_i}^0 \mid \sigma_{or;exc}) & for \quad \theta \neq \theta_{\epsilon_i}^0 \end{cases}$$

Space:

$$F_{\epsilon_j}(y) = \begin{cases} 1 & for \quad y = y_{\epsilon_j}^0 \\ \beta_1 G(y - y_{\epsilon_j}^0 \mid \sigma_{csp;exc}) & for \quad y \neq y_{\epsilon_j}^0 \end{cases}$$

where θ and y are the orientation and location, respectively, of the first-stage filter, providing input to the second-stage filter, $\theta_{\epsilon_i}^0$ and $y_{\epsilon_j}^0$ are the second-stage filter orientation and location, respectively.

The suggested inhibitory second-stage filter consists of two components: broad-band inhibition independent of input orientation and side inhibition that is modeled with two additional Gaussians. The spatial masking results do not provide evidence for side inhibition and the inhibitory spatial second stage filter was thus modeled with a single Gaussian:

Orientation:

$$F_{\iota_i}(\theta) = \alpha_2[\alpha_3 + G(\theta - \theta_{\iota_i} - \theta_{inh} \mid \sigma_{or;inh}) + G(\theta - \theta_{\iota_i} + \theta_{inh} \mid \sigma_{or;inh})]$$

Space:

$$F_{\iota_j}(\Delta y) = \beta_2 G(y - y_{\iota_j}^0 \mid \sigma_{sp;inh}).$$

The model response R is a function of the input-stimulus $S(\theta)$ or $S(y)$. The Gabor stimuli are described

*This is expected if the divisive inhibition is applied after a threshold-type transducer function (with a continuous derivative everywhere but at threshold), as the threshold, the point where maximal facilitation occurs, is not affected by division (assuming a smooth inhibitory transducer function), unlike the gain.

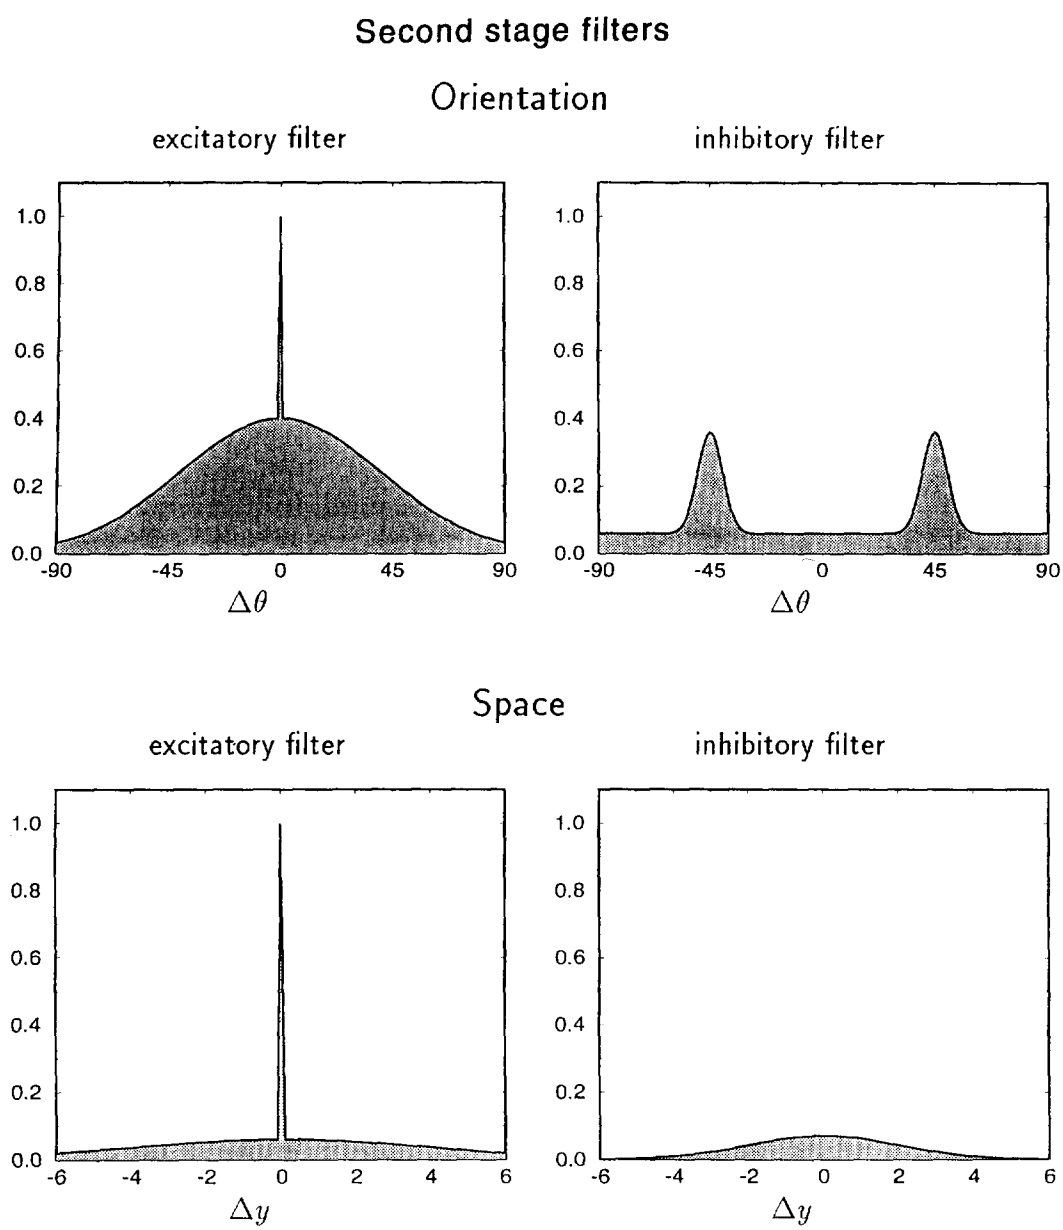

FIGURE 9. Second-stage filters used in the computer simulation (for the mathematical description see text).

here as Gaussians with an amplitude proportional to their contrast (the amplitude is negative in the case of phase reversal). Assuming a vertical target ($\theta = 0$) at the origin ($y = 0$) and masks with orientations of $\pm\Delta\theta$ located at $y = 0$, or with vertical orientation located at $\pm\Delta y$:

Orientation:

$$S(\theta) = C_{m1}G(\theta - \Delta\theta \mid 13^\circ) + C_t G(\theta \mid 13^\circ) + C_{m2}G(\theta + \Delta\theta \mid 13^\circ)$$

Space:

$$S(y) = C_{m1}G(y - \Delta y \mid \lambda) + C_t G(y \mid \lambda) + C_{m2}G(y + \Delta y \mid \lambda).$$

The first-stage responses r_i and r_j (after rectification) are then given by:

Orientation: $$r_i = \left| \int_0^{2\pi} F_i(\theta) \cdot S(\theta) d\theta \right|$$

Space: $$r_j = \left| \int_{-\infty}^{\infty} F_j(y) \cdot S(y) dy \right| .$$

Next, we assume that decision is based on the output of the vertical second-stage filters corresponding to target location. These second-stage responses r_{ϵ_0} and r_{ι_0} are given by:

Orientation:

$$r_{\epsilon_0} = \Sigma_i F_{\epsilon_0}(\theta_i^0) \cdot r_i \quad \textit{and} \quad r_{\iota_0} = \Sigma_i F_{\iota_i}(\theta_i^0) \cdot r_i$$

Space:

$$r_{\epsilon_0} = \Sigma_j F_{\epsilon_0}(y_j^0) \cdot r_j \quad \textit{and} \quad r_{\iota_0} = \Sigma_j F_{\iota_j}(y_j^0) \cdot r_j$$

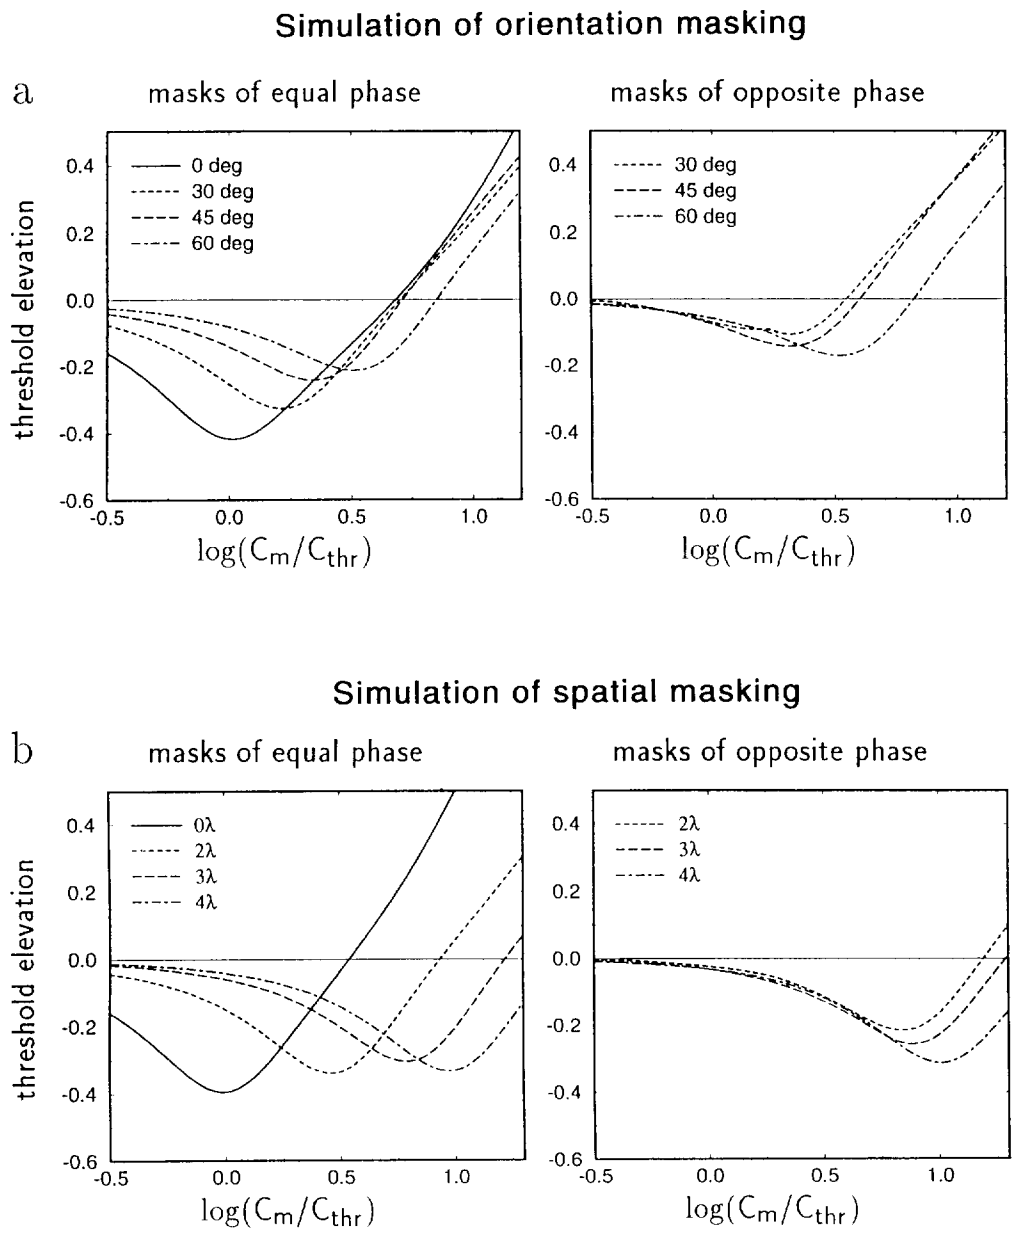

FIGURE 10. Model simulation results for masks presented at different orientations (a) or different spatial locations (b). The curves can be compared to the psychophysical results presented in Figs 2 and 4.

Finally, the model response R is given by:

$$R = \frac{trd(r_{\epsilon_0})}{1 + r_{\epsilon_0}} \quad \text{with} \quad \text{trd}(r) = \frac{cr^n}{\mu^{n-1} + r^{n-1}}.$$

The transducer function trd(r) is similar to the widely used Naka–Rushton function (Naka & Rushton, 1966), with the exponent in the denominator being reduced by 1. The transducer is thus still accelerating for $r \ll \mu$ (and can account for facilitation), but it does not saturate for $r \gg \mu$ and converges to a linear function. It is important to note that this transducer function predicts constant detection thresholds for large inputs and that the power-law behavior in the simulations is entirely due to divisive inhibition. Detection thresholds can be evaluated by assuming that two stimuli are discriminable if and only if $\Delta R \geqslant 1$.

The values of the parameters used in the simulation are given in Table 1, and the shapes of the second-stage filters are presented in Fig. 9. Figure 10(a) shows the simulation results for orientation masking. By comparing simulated and experimental results (see Fig. 2) one can appreciate that the main data features are captured well by the model; namely, the decrease of facilitation for masks of equal phase and the increase of facilitation for masks of opposite phase (for increasing $\Delta\theta$). To show this further, the analysis that was performed on the experimental data was also carried out for the simulated results. It can be seen in Fig. 3 that the obtained fit is quite good. Suppression thresholds were shown to change with practice and a good fit is not necessarily expected.

The simulation results for the spatial masking experiments are presented in Fig. 10(b). Psychophysical and

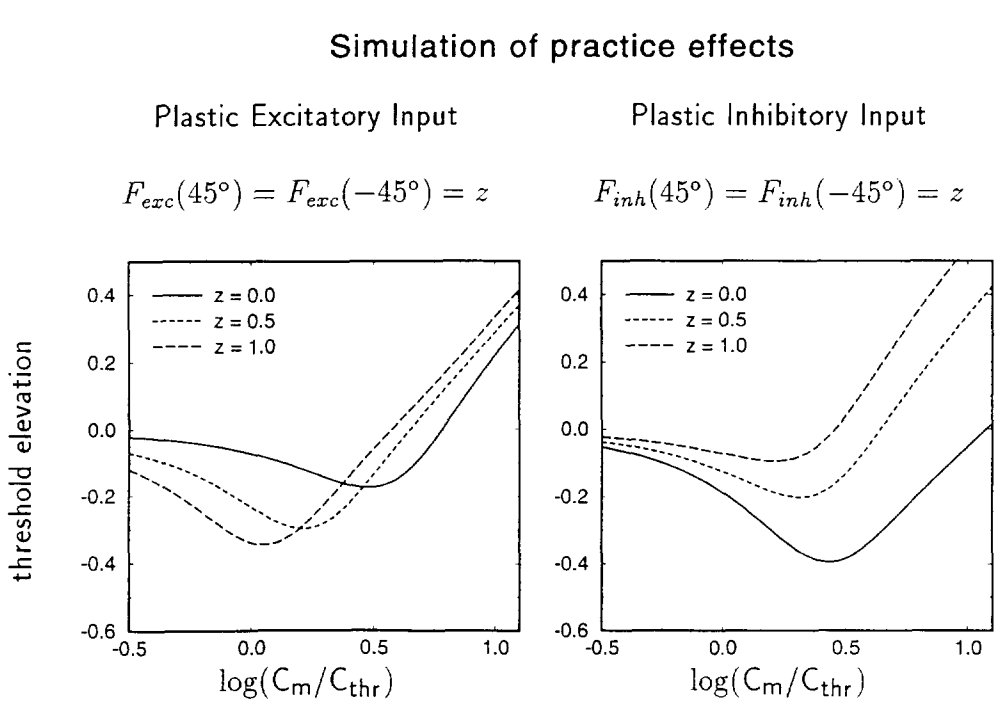

FIGURE 11. Simulation of model plasticity for masks of equal phase at $\Delta\theta = 45$ deg. A decrease of excitatory input from 45 deg and a decrease of inhibitory input from 45 deg both can account for an increase of suppression thresholds. Decrease of inhibitory input can further account for the practice effect of observer HB (see Fig. 7).

simulated results are plotted together in Fig. 5. Clearly, the fit is less accurate than for the orientation masking results. One problem is that the data do not follow a clear general behavior, especially at larger mask distances. More experimental data are necessary (also for masks at closer spatial distances such as 1λ) in order to obtain a more accurate estimation of the second-stage filters described here. In any case, the model can account for facilitation by masks of both phase relationships and it can simulate approximately the mask contrast at minimum.

The model described here was deliberately kept simple: the first stage is not very different from a linear "stage", as neither a threshold nor saturation are assumed for the first-stage units. Moreover, both facilitation and suppression are accounted for by only one mechanism each: facilitation by an accelerating transducer function applied on a second-stage filter and suppression by divisive inhibition. Because of its simplicity, the model provides a useful basis for further investigations as it allows for various modifications. For example, nonlinear transducer functions might be applied on the first-stage filter output or on the inhibitory second-stage filter. In addition, the temporal dynamics of the system might be described, possibly allowing for discrimination between feed-forward and feed-back structures.

Plasticity

The experimental results were found to change with practice, a finding that implies that the model described above has to be modified. Namely, it has to account for the global increase of suppression thresholds [see Fig. 6(d)]. Simple modifications to the second-stage filters were examined and tested by computer simulations. We suggest that the input weights to the second-stage filters can be modified by experience and that these modifications apply locally to the particular first-stage filter used. We consider here the case where masks of equal phase are presented at $\Delta\theta = 45$ deg and we modulate independently the 45 deg-input weight to the excitatory and to the inhibitory second-stage filter, while the filters remain otherwise unchanged.

The result of the simulation is presented in Fig. 11. A decrease in the excitatory input and a decrease in the inhibitory input can both account for the observed increase in suppression thresholds. The development of facilitation that we observed in some cases (see Fig. 7), however, is not consistent with a decrease in excitatory input and suggests a decrease in inhibitory input as a possible learning mechanism.

DISCUSSION

A contrast masking paradigm was used to study nonlinear interactions between filters tuned to different orientations and spatial locations. The experiments were carried out for two different mask-phase relationships allowing for an isolation of two separate processing stages.

We find that, for increasing mask contrast, thresholds usually first decrease, then reach a minimum and then increase linearly on a log-log scale (which corresponds to a power-law behavior). The magnitude of maximal facilitation and the mask contrast at which the minimum occurs depends on the mask orientation, the spatial displacement of the mask and the mask-phase relation-

ship. For masks of equal phase, facilitation was shown to decrease with increasing orientation difference. A particularly strong decrease was found between $\Delta\theta = 30$ deg and $\Delta\theta = 45$ deg. A facilitatory effect also was observed for masks of opposite phase when masks were presented at larger orientation differences (such as 60 deg) or larger spatial distances ($< 2\lambda$). The power-law behavior with an exponent of 0.89 was observed independently of mask phase in all conditions (except for masks at large spatial distances).

The results are accounted for by two filtering stages. Linear filtering of the image is followed by a full-wave rectification. The first-stage output provides input to two second-stage filters, an excitatory filter that is followed by an accelerating transducer function and an inhibitory second-stage filter that provides divisive inhibition to the output of the excitatory transducer function. Facilitation is accounted for by the accelerating transducer function and the divisive inhibition accounts for the observed suppression. The model is similar to a model recently published by Foley (1994a). However, an important difference is that, in the model presented here, excitatory and inhibitory filters are described as second-stage filters rather than first-stage filters. This was motivated by the results obtained for masks of opposite phase showing that both facilitatory and suppressive mask effects can be observed independently of mask phase. The data obtained allow for an estimation of the second-stage filter parameters.

Second-stage filters

Two alternatives have been discussed concerning the tuning of inhibitory interactions: broadly tuned inhibition (more or less insensitive to orientation) and orientation selective inhibition. Broad-band inhibition was suggested as a mechanism that effectively normalizes cell responses and helps to avoid response saturation (Heeger, 1992). This type of inhibition is consistent with physiological data: for example, the contrast independence of orientation tuning in cat striate cells (Sclar & Freeman, 1982). However, there is also evidence for orientation-selective inhibition (Hata *et al.*, 1988; Bonds, 1989; Volgushev *et al.*, 1993), which would serve as a mechanism for sharpening the orientation tuning curves of cortical cells. Possibly, both mechanisms act together (Bonds, 1989).

The data presented here indicate that the inhibitory input coming from 45 deg is much larger than the inhibitory input coming from 30 deg and that inhibition is, therefore, not independent of orientation. Accordingly, side inhibition was assumed in the model, supporting the hypothesis that inhibitory interactions do play a role in sharpening orientation tuning functions. Foley (1994a) accounts for his results from masking experiments by broad-band inhibition. However, as he was using gratings as mask stimuli, spatial inhibition (surround inhibition) from orientations similar to the target orientation may have also affected the data, by contributing strong inhibition around the target orientation.

In addition to the inhibitory second-stage filter, the model also describes an excitatory second-stage filter. This filter integrates over neighboring orientations and neighboring spatial distances. Anatomical models of the visual cortex suggest that cells tuned to the same spatial location but to different orientations are located close to each other within a "hyper-column", whereas neighboring spatial locations are encoded in neighboring hyper-columns. This could explain the fact that the excitatory input coming from spatially displaced masks appears to be smaller than the input coming from neighboring orientations. Neurons with long axons that could mediate long-range interactions were described in the visual cortex of the cat (Gilbert & Wiesel, 1979, 1983) and, furthermore, there is evidence for facilitation among cells whose receptive fields are co-aligned and co-oriented (Nelson & Frost, 1985; Ts'o, Gilbert & Wiesel, 1986). The spatial integration of the excitatory filter thus might be non-isotropic. Psychophysical evidence for this anisotropy was provided by Polat and Sagi (1994a) using a lateral masking paradigm (similar to the one used in the present study). They found that facilitation of detection by masks is maximal when target and masks were presented co-linearly.

Feed-forward/feed-back

The data presented here do not allow for a decision between feed-forward or feed-back structure and the feed-forward structure was chosen entirely for the sake of simplicity. Foley (1994a) suggests a feed-forward structure, based on the observation that masks presented for only 33 msec give rise to large inhibition. However, the processing time in the cortex might not be restricted to the actual stimulus presentation. Heeger (1992) argues for a feed-back structure of inhibition, as only then response saturation could be avoided. Different architectures can be suggested, as excitation and inhibition do not necessarily follow the same interaction pattern (Stemmler *et al.*, 1995). A feed-back architecture for excitation was suggested to account for the increased range of excitatory interactions with practice (Polat & Sagi, 1994b). The excitatory second-stage filters derived here (see Fig. 9) can be viewed as the sensitivity pattern of weak lateral excitatory inputs to a first-stage filter, with the first-stage filter response dominating. Further psychophysical experiments are necessary, for example, for testing feed-back specific effects like dis-inhibition (Carpenter & Blakemore, 1973; Kurtenbach & Magnussen, 1981) or investigation of time course of different interactions (Wilson & Humanski, 1993).

Plasticity

Some of our results showed significant practice effects. These findings agree with observations made before, as for the existence of learning effects in masking experiments. Swift and Smith (1983), using eight-component noise gratings, described a reduction of the discrimination function slope at the suppression region from 1 to 0.65 (the slope they obtained without practice for single component gratings), which took place each time they

changed the mask combination. They placed the learning effect at the decision stage, with practice affecting decision criteria. Here learning was shown to take place with two-component masks, the most consistent practice effect being an increase in suppression thresholds. For both single-component and double-component masks used, we found fairly stable slopes (0.89). The model described here can account for the observed practice effects if the second-stage filters are modified, suggesting plasticity at an early stage of visual processing. The performance improvement that was seen for observer HB strongly indicates that inhibition is reduced due to practice.

Evidence for low-level plasticity has been reported (Karni & Sagi, 1991; Poggio *et al.*, 1992; Polat & Sagi, 1994b). Practice effects have been described that are specific for eye, stimulus location and stimulus orientation. The high specificity of practice effects indicates that plasticity is present at early processing stages. Texture learning was found to be task-specific (Ahissar & Hochstein, 1993), implying that stimulus presentation alone does not lead to plasticity but that high-level processes are also necessary for learning. However, it is possible that learning is mainly a low-level process and that a high-level signal simply enables or gates synapses (Karni & Sagi, 1991) in a certain brain region to change their efficacy. In the experiments described here, a high-level signal could be sent to the second stage target filters, thus allowing for their modification. The actual modifications might then be completely stimulus-dependent. Local learning rules could be described, similar to the rules suggested for excitatory synapses by Hebb (1946) and for inhibitory synapses by Barlow (1990). In both cases, the learning rules assume an increase of synaptic efficacy with correlated activity on the two synaptic sides and a decrease in efficacy for uncorrelated activities. Within the context of the model presented here, a slow decrease in the efficacy of divisive inhibition seems to take place with repetitive stimulation and task performance. Assuming local learning rules, the decrease of synaptic strength can be a result of uncorrelated activities in the corresponding excitatory and the inhibitory second-stage filters, as these two filters have different tuning profiles (the inhibitory filter receives a strong input from first-stage filters at 45 deg, while the input to the excitatory filter is dominated by the target orientation). Alternative accounts are possible if a feed-back design is adopted, enabling indirect effects due to increased mutual inhibition between mask responding second-stage filters (thus producing a reduced effective inhibition on the target filter). However, it is possible that learning is supervised and synapses can be modulated independently of input correlations so as to optimize and reduce discrimination thresholds, with network architecture being the limiting factor. Further experiments, using paradigms similar to the one described here, might provide an answer to these open questions and may help in understanding the principles governing learning.

REFERENCES

Ahissar, M. & Hochstein, S. (1993). Attentional control of early perceptual learning. *Proceedings of the National Academy of Science, U.S.A., 90*, 5718–5722.

Barlow, H. B. (1990). A theory about the functional role and synaptic mechanism of visual after-effects. In Blakemore, C. (Ed.), *Vision: coding and efficiency*, Ch. 32 (pp. 363–375). Cambridge, UK: Cambridge University Press.

Ben-Av, M. B. & Sagi, D. (1995). Perceptual grouping by similarity and proximity: Experimental results can be predicted by intensity autocorrelations. *Vision Research, 35*, 853–866.

Blakemore, C. & Campbell, F. W. (1969). On the existence of neurons in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology, London, 203*, 237–260.

Blakemore, C., Carpenter, R. S. H. & Georgeson, M. (1970). Lateral inhibition between orientation detectors in the human visual system. *Nature, 228*, 37–39.

Blakemore, C. & Nachmias, J. (1971). The orientation specificity of two visual after-effects. *Journal of Physiology, London, 213*, 157–174.

Bonds, A. B. (1989). Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Visual Neuroscience, 2*, 41–55.

Bradley, A. & Ohzawa, I. (1986). Comparison of contrast detection and discrimination. *Vision Research, 26*, 991–997.

Campbell, F. & Kulikowski, J. (1966). Orientational selectivity of the human visual system. *Journal of Physiology, London, 187*, 437–445.

Carpenter, R. S. H. & Blakemore, C. (1973). Interactions between orientations in human vision. *Experimental Brain Research, 18*, 287–303.

Foley, J. M. (1994a). Human luminance pattern-vision mechanisms: masking experiments require a new model. *Journal of the Optical Society of America, A11*, 1710–1719.

Foley, J. M. (1994b). Spatial phase sensitivity of luminance pattern mechanisms determined by masking. *Investigative Ophthalmology & Visual Science (Suppl.), 35*, 1900.

Gilbert, C. D. & Wiesel, T. N. (1979). Morphology and intracortical projections of functionally identified neurons in cat visual cortex. *Nature, 280*, 120–125.

Gilbert, C. D. & Wiesel, T. N. (1983). Clustered intrinsic connections in cat visual cortex. *Journal of Neuroscience, 3*, 1116–1133.

Hata, Y., Tsumoto, T., Sato, H., Hagihara, K. & Tamura, H. (1988). Inhibition contributes to orientation selectivity in visual cortex of cat. *Nature, 335*, 815–817.

Hebb, D. O. (1949). *Organization of behavior*. NY: Wiley.

Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience, 9*, 181–197.

Karni, A. & Sagi, D. (1991). Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. *Proceedings of the National Academy of Science, U.S.A., 88*, 4966–4970.

Kulikowski, J., Abadi, R. & King-Smith, P. E. (1973). Orientation selectivity of grating and line detectors in human vision. *Vision Research, 13*, 1479–1486.

Kurtenbach, W. & Magnussen, S. (1981). Inhibition, disinhibition, and summation among orientation detectors in human vision. *Experimental Brain Research, 43*, 193–198.

Lawton, T. B. & Tyler, C. W. (1994). On the role of x and simple cells in human contrast processing. *Vision Research, 34*, 659–667.

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.

Levitt, H. (1971). Transformed up-down methods in psychoacoustics. *Journal of the Acoustical Society of America, 49*, 467–477.

Magnussen, S. & Kurtenbach, W. (1980). Adapting to two orientations: Disinhibition in a visual aftereffect. *Science, 207*, 908–909.

Morgan, M. J. & Dresp, B. (1995). Contrast detection facilitation by spatially separated targets and inducers. *Vision Research, 35*, 1019–1024.

Nachmias, J. (1993). Masked detection of gratings: The standard model revisited. *Vision Research*, *33*, 1359–1365.

Nachmias, J. & Sansbury, R. V. (1974). Grating contrast: Discrimination may be better than detection. *Vision Research*, *14*, 1039–1042.

Naka, K. I. & Rushton, W. A. H. (1966). S-potentials from luminosity units in the retina of fish (*Cyprinidae*). *Journal of Physiology, London*, *185*, 587–599.

Nelson, J. I. & Frost, B. J. (1985). Intracortical facilitation among co-oriented, co-axially aligned simple cells in cat striate cortex. *Experimental Brain Research*, *61*, 54–61.

Olzak, L. A. & Thomas, J. P. (1991). When orthogonal orientations are not processed independently. *Vision Research*, *31*, 51–57.

Olzak, L. A. & Thomas, J. P. (1992). Configural effects constrain Fourier models of pattern discrimination. *Vision Research*, *32*, 1885–1898.

Phillips, G. C. & Wilson, H. R. (1984). Orientation bandwidths of spatial mechanisms measured by masking. *Journal of the Optical Society of America*, *A1*, 226–232.

Poggio, T., Fahle, M. & Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. *Science*, *256*, 1018–1021.

Polat, U. & Norcia, A. M. (1996). Neurophysiological evidence for contrast dependent long-range facilitation and suppression in the human visual cortex. *Vision Research*, *36*, 2099–2109.

Polat, U. & Sagi, D. (1993). Lateral interactions between spatial channels: Suppression and facilitation revealed by lateral masking experiments. *Vision Research*, *33*, 993–999.

Polat, U. & Sagi, D., (1994a). The architecture of perceptual spatial interactions. *Vision Research*, *34*, 73–78.

Polat, U. & Sagi, D. (1994b). Spatial interactions in human vision: From near to far via experience dependent cascades of connections. *Proceedings of the National Academy of Science, U.S.A.*, *91*, 1206–1209.

Ross, J. & Speed, H. D. (1991). Contrast adaptation and contrast masking in human vision. *Proceedings of the Royal Society of London*, *B246*, 61–69.

Rubenstein, B. S. & Sagi, D. (1990). Spatial variability as a limiting factor in texture discrimination tasks: Implications for performance asymmetries. *Journal of the Optical Society of America*, *A7*, 1632–1643.

Sagi, D. & Hochstein, S. (1985). Lateral inhibition between spatially adjacent spatial frequency channels? *Perception & Psychophysics*, *37*, 315–322.

Sclar, G. & Freeman, R. D. (1982). Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Experimental Brain Research*, *46*, 457–461.

Stemmler, M., Usher, M., & Niebur, E. (1995). Lateral interactions in primary visual cortex: A model bridging physiology and psychophysics. *Science*, *269*, 1877–1879.

Swift, D. J. & Smith, R. A. (1983). Spatial frequency masking and Weber law. *Vision Research*, *23*, 495–505.

Tolhurst, D. J. & Barfield, L. (1978). Interactions between spatial frequency channels. *Vision Research*, *18*, 951–958.

Ts'o, D. Y., Gilbert, C. D. & Wiesel, T. N. (1986). Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *Journal of Neuroscience*, *6*, 1160–1170.

Volgushev, M., Peo, X., Vidyasagar, T. R. & Creutzfeldt, O. D. (1993). Excitation and inhibition in orientation selectivity of cat visual cortex neurons revealed by whole-cell recordings *in vivo*. *Visual Neuroscience*, *10*, 1151–1155.

Wilson, H. R. (1980). A transducer function for threshold and suprathreshold human vision. *Biological Cybernetics*, *38*, 171–178.

Wilson, H. & Humanski, R. (1993). Spatial frequency adaptation and contrast gain control. *Vision Research*, *33*, 1133–1149.

Acknowledgements—We thank Alexander Cooperman, Alain Dorais, Alumit Ishai, Yasuto Tanaka and Marius Usher for helpful comments on an earlier version of the manuscript. This work was supported by the Basic Research Foundation administered by the Israel Academy of Sciences and Humanities and by Foundation Mordoh Mijan de Salonique.