VISUAL COGNITION, 2006, 13 (4), 401–480

Psychology Press
Taylor & Francis Group

A neural network model of inhibitory processes in subliminal priming

Howard Bowman

Computing Laboratory, University of Kent at Canterbury, UK

Friederike Schlaghecken

Department of Psychology, University of Warwick, Coventry, UK

Martin Eimer

Department of Experimental Psychology, Birkbeck College, University of London, UK

Masked priming experiments have revealed a precise set of facilitatory and inhibitory visuomotor control processes. Most notably, inhibitory effects have been identified in which prime–target compatibility induces performance costs and prime–target incompatibility induces performance benefits. We argue that this profile of data is commensurate with an ''emergency braking mechanism'', whereby responses can be retracted as a result of changing sensory evidence. The main contribution of this paper is to provide a neural network-based explanation of this phenomenon. This is obtained through the use of feedforward inhibition to implement backward masking, lateral inhibition to implement response competition, and opponent processing mechanisms to implement response retraction. Although the model remains simple, it does a very good job of reproducing the available masked priming data. For example, it reproduces a large spectrum of reaction time data across a number of different experimental conditions. Perhaps most notably, however, it also reproduces lateralized readiness potentials that have been recorded while subjects perform different conditions. In addition, it provides a concrete set of testable predictions.

Please address all correspondence to: Howard Bowman, Computing Laboratory, University of Kent at Canterbury, Canterbury, Kent CT2 7NF, UK. Email: H.Bowman@kent.ac.uk

We would like to thank Dietmar Heinke, Trevor Robbins, Phil Barnard, Chris Barry, and two reviewers of the paper; all of whom gave valuable suggestions on the work presented here. Adam Aron was also involved in preliminary research on this model.

http://www.psypress.com/viscog DOI:10.1080/13506280444000823

INTRODUCTION

Direct perceptuomotor links and response inhibition

In a constantly changing environment, the successful control of behaviour requires an organism to respond quickly and flexibly to novel stimuli. This includes the ability to rapidly initiate a motor response even before the respective stimulus has been consciously identified, and to quickly interrupt and change ongoing behaviour as soon as sensory evidence indicates that it is no longer appropriate. The former is assumed to be mediated by so-called ''direct perceptuomotor links''—straight pathways from sensory to motor systems which allow sensory information to trigger motor responses directly, i.e., without the need for prior conscious stimulus recognition (e.g., Neumann, 1990; Neumann & Klotz, 1994). The latter is assumed to be brought about by inhibitory processes controlled by central executive mechanisms in the prefrontal cortex (for a recent review, see Band & van Boxtel, 1999).

Initial evidence for the existence of direct perceptuomotor links stems from the observation that perceptually salient stimulus features (e.g., their spatial location) will trigger the initiation of a corresponding motor response even if these features are—per instruction—totally response-irrelevant (Simon, 1969; see also de Jong, Liang, & Lauber, 1994; Eimer, 1995). A more spectacular demonstration of direct perceptuomotor links can be found in certain neurological impairments such as ''blindsight''. This disorder is caused by a lesion to the primary visual cortex, and results in a total loss of visual experience in the affected field. Surprisingly, though, patients are often able to respond accurately to stimuli appearing in their blind field—despite denying that they can see anything at all in this area (Weiskrantz, Warrington, Sanders, & Marshall, 1974). Similar phenomena can also be observed in visual-illusion experiments with neurologically unimpaired subjects. Results indicate that although conscious experience is ''fooled'' by a perceptual illusion, pointing and grasping movements frequently are not (Agliotti, DeSouza, & Goodale, 1995; Bridgeman, Kirch, & Sperling, 1981).

The most exhaustive investigation of direct perceptuomotor links, however, has been conducted in the context of the ''masked prime'' paradigm. In this paradigm, participants make a speeded two-alternative choice response to a clearly visible ''target'' stimulus (e.g., a left- or right-hand key press to an arrow pointing to the left or right). On each trial, the target is preceded by a briefly presented and subsequently masked ''prime'' stimulus. The prime can either be one of the two possible targets (e.g., a left- or right-pointing arrow), or a neutral stimulus that is not associated with any motor response (e.g., a plus sign). Due to presentation duration and masking, primes are ''subliminal'', i.e., they remain below the threshold of perceptual awareness, as evidenced by various detection and identification procedures (e.g., Dehaene et al., 1999;

Eimer & Schlaghecken, 1998; Klotz & Wolff, 1995; Schlaghecken & Eimer, 1997).

However, although participants are not consciously aware of their presence, primes systematically affect responses to the subsequent target. Compared to neutral-prime trials, responses are usually faster and more accurate on ''*compatible*'' trials, where prime and target are mapped to the same response, and slower and less accurate on ''*incompatible*'' trials, where they are mapped to opposite responses (Dehaene et al., 1999; Jaskoski, van der Lubbe, Schlotterbeck, & Verleger, 2002; Klotz & Wolff, 1995; Leuthold & Kopp, 1998; Naccache & Dehaene, 2001; Neumann & Klotz, 1994; Schlaghecken & Eimer, 1997, 2000). These behavioural results, together with converging evidence from electrophysiological and haemodynamic measures (e.g., Dehaene et al., 1999; Eimer, 1999; Eimer & Schlaghecken, 1998; Jaskowski et al., 2002; Klotz & Wolff, 1995; Leuthold & Kopp, 1998), indicate that the primes, despite being inaccessible to conscious awareness, trigger activation of their corresponding motor response.

Being able to activate a motor response to a stimulus that has not (or not yet) been fully analysed has the obvious advantage of allowing extremely quick reactions to novel and potentially important stimuli. On the other hand, it has the equally obvious disadvantage that these reactions might turn out to be inappropriate. Therefore, it has been suggested that all behaviour is continually monitored by a ''central executive'', which—if necessary—will activate inhibitory control processes to interrupt the ongoing motor response (e.g., Dehaene, Posner, & Tucker, 1994; Kopp, Rist, & Mattler, 1996; Sasaki, Gemba, Nambu, & Matsuzaki, 1993; and for a review, see Band & van Boxtel, 1999).

Evidence for this comes from event-related brain potential (ERP) studies of response inhibition. In a typical experiment, participants have to make a speeded motor response on most trials, but are occasionally required to withhold any overt response. On these ''stop'' trials, activity in the prefrontal cortex increases, as evidenced by an enlarged negative-going shift (N2) at anterior electrode sites on ''stop'' as compared to ''go'' trials (e.g., Bokura, Yamaguchi, & Kobayashi, 2001; Bruin, Wijers, & van Staveren, 2001; Eimer, 1993; Falkenstein, Hoormann, & Hohnsbein, 1999; Heil, Osman, Wiegelman, Rolke, & Henninghausen, 2000; Jodo & Kayama, 1992; Kok, 1986; Kopp, Mattler, Goertz, & Rist, 1996). The same N2 effect has also been observed when, instead of withholding all responses, participants have to change from an inappropriately preactivated response to the response actually required by the target (Gehring, Gratton, Coles, & Donchin, 1992; Heil et al., 2000; Kopp, Mattler, et al., 1996; Kopp, Rist, & Mattler, 1996). Converging evidence from functional magnetic resonance imaging (fMRI) studies (e.g., Carter et al., 1998) and from patients with prefrontal lesions (e.g., Godefroy, Cabaret, Petit-Chenal, Pruvo, & Rousseaux, 1999), further supports the hypothesis that these response inhibition processes depend on prefrontal inhibitory control

mechanisms (see also Baddeley, 2000; Baddeley & Hitch, 1974; Fuster, 1989; Shallice, 1988).

There is, however, a problem with this notion: The prefrontal cortex is generally regarded to mediate *voluntary* control processes (e.g., Luria, 1992; Posner & DiGirolamo, 1998; Shallice, 1988; Spence & Frith, 1999), but may not be involved in *automatic* behaviour (e.g., Jueptner et al., 1997). Processes of subliminal perception and automatic motor control rather appear to be a function of parietal cortex (Frith & Dolan, 1996; Pisella et al., 2000)—an area that did not show any increased activation in the above-mentioned motor inhibition tasks. Thus one might speculate that if inhibition is a function of the prefrontal cortex, and if the prefrontal cortex is not involved in the perception of subliminal stimuli, then inhibitory control might not be available for response tendencies triggered by subliminal stimuli.

Evidence to support this hypothesis has indeed been found in a number of different experimental paradigms (e.g., Allport, Tipper, & Chimel, 1985; Marcel, 1980; McCormick, 1997; Neill, Valdes, & Terry, 1995), suggesting that subliminal stimuli trigger only ''passive'' activation processes, but not active inhibition. Results from the masked prime experiments reported above are also generally in line with the activation-only hypothesis. Behavioural positive compatibility effects (PCEs—better performance on compatible than on incompatible trials) have been found to be accompanied by electrophysiological and haemodynamic evidence of motor activation of the primed response (e.g., Dehaene et al., 1999; Jaskoski et al., 2002; Klotz & Wolff, 1995; Leuthold & Kopp, 1998), without evidence for subsequent inhibition.

Despite this compelling evidence, we now have reason to believe that the ''activation-only'' hypothesis of subliminal motor priming is incorrect. In a recent series of masked prime experiments (Eimer, 1999; Eimer & Schlaghecken, 1998, 2001; Eimer, Schubö, & Schlaghecken, 2002; Klapp & Hinkley, 2002; Schlaghecken & Eimer, 1997, 2000, 2001, 2002) behavioural and electrophysiological evidence for the existence of inhibitory control processes in response to subliminal stimuli has been provided. At the behavioural level, the basic finding of these experiments is that an initial PCE, obtained when the target follows the masked prime immediately, i.e., mask–target stimulus onset asynchrony (SOA) of 0 ms, turns into a *negative compatibility effect* (NCE—benefits on incompatible trials, and costs on compatible trials, relative to neutral trials) when the mask–target SOA is sufficiently long (about 80 ms or more; see Figure 1).

At the electrophysiological level, this sequence of PCEs and NCEs is reflected in a specific biphasic pattern of the lateralized readiness potential (LRP—an online measure of unimanual motor preparation, obtained by computing the difference in activation levels between contralateral and ipsilateral motor cortex relative to target direction). Figure 2 (taken from Eimer, 1999) depicts LRP waveforms from an experiment where mask–target SOA was 100

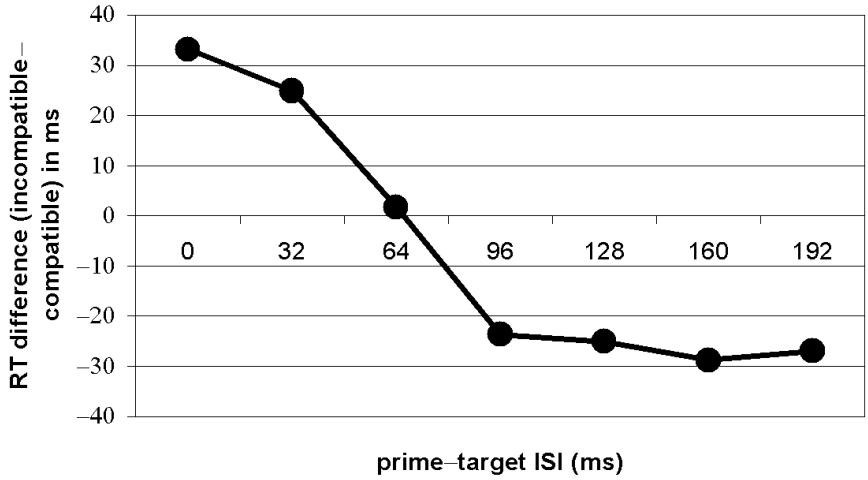

Figure 1. Masked priming task response times for different prime–target ISIs. The figure is a re-representation of data in Figures 1a and 1b of Eimer and colleagues (Eimer, 1999; Eimer & Schlaghecken, 1998; Eimer & Schlaghecken, 2001; Eimer et al., 2002; Klapp & Hinkley, 2002; Schlaghecken & Eimer, 1997, 2000, 2001, 2002). Note that in these experiments prime and mask are presented in the same spatial location, but the target is spatially offset to enable it to be presented simultaneously with the mask.

ms, and behavioural NCEs were obtained. In this graph, the origin corresponds to prime onset, and potentials are measured up to 600 ms after prime onset. To reiterate, the direction of deflection (above or below the x-axis) indicates the activation difference between contra- and ipsilateral motor cortices relative to target direction, i.e., the difference between motor cortex activations corresponding to the correct and incorrect hand movement. As a reflection of the fact that the LRPs indicate how the difference in activation between left and right hemispheres changes over time, we will throughout the paper talk in terms of *separation* between response alternatives or, in other words, about the extent to which one response dominates the other.

In Figure 2 the black arrow indicates the time (around 250 ms after prime onset) where the LRP shows significantly increased initial separation of the contralateral motor cortex (corresponding to a correct response preparation) on compatible trials and of the ipsilateral motor cortex (corresponding to an incorrect response preparation) on incompatible trials. Since compatible and incompatible trials—by definition—differ with respect to whether the prime is mapped to the correct or the incorrect response, these initial separations can safely be interpreted as motor separations triggered by the masked prime. Interestingly, the initial separation subsequently returns to baseline and is replaced by an increased separation in the opposite direction. This is most

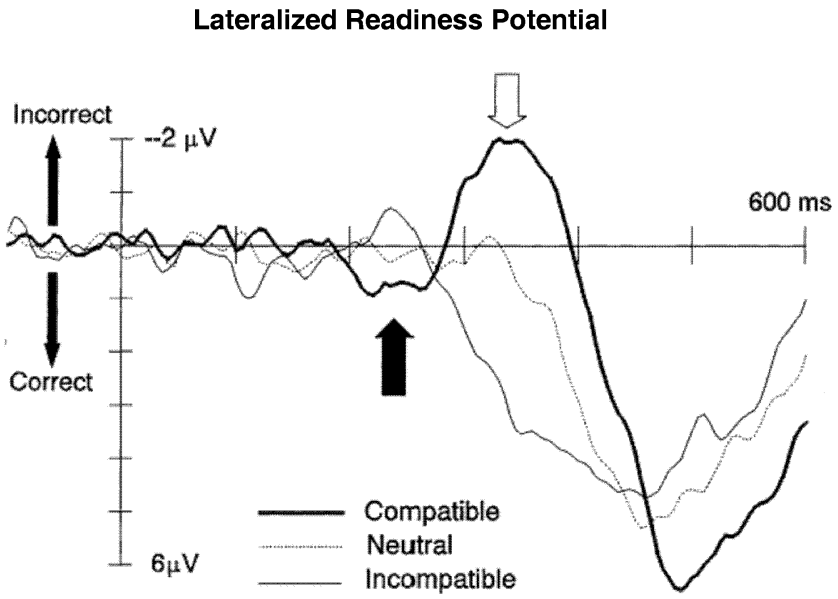

Figure 2. LRP for masked priming task (from Eimer, 1999).

clearly seen on compatible trials, where this ''reversal phase'' results in a substantial separation towards the ipsilateral (i.e., incorrect) motor cortex. On incompatible trials, the reversal results in a separation towards the contralateral (i.e., correct) motor cortex. The white arrow indicates the maximum of the reversal phase, i.e., the point at which target-related activation begins to cancel out the reversal in the compatible condition and to accentuate it in the incompatible condition. On neutral trials (prime not mapped to any response), neither initial separation nor subsequent reversal occurs.

The LRP data thus clearly illustrates why PCEs occur with short SOAs, and NCEs with longer SOAs. If motor preparation in response to the target starts early, it will benefit from the correct direction separation triggered by the prime on compatible trials, but will have to overcome the incorrect direction separation on incompatible trials. If, however, motor preparation in response to the target starts late, the opposite is true: Now it will benefit from a reversed—and hence correct—separation on incompatible trials, and will have to overcome a reversed—and hence incorrect—separation on compatible trials. If the mask–target SOA is chosen so that motor preparation in response to the target starts after the initial separation phase, but before the subsequent reversal phase, no behavioural effects of prime–target compatibility should be obtained. Exactly this has been demonstrated in two experiments where SOA was increased stepwise from 0 ms to 100 ms (Schlaghecken & Eimer, 1997) or to 192 ms

(Schlaghecken & Eimer, 2001): In both experiments, initial PCEs disappeared with a mask–target SOA of about 60 ms, and turned into NCEs with longer SOAs.

PCEs turning into NCEs and the reversal of prime-induced motor separation strongly suggest the operation of inhibitory processes. Moreover, direct evidence for response inhibition has been obtained in an experiment using a go/nogo variant of the masked prime task (Eimer & Schlaghecken, 1998, Exp. 3). In addition to finding NCEs on reaction times (faster responses on go trials when the prime was a nogo stimulus than when it was a go stimulus), there was also a substantial reduction of false alarms on incompatible nogo trials (i.e., when the prime was a go stimulus). This indicates that the response triggered by the prime had been actively inhibited.

NCE and self-inhibition

The interpretation of the NCE as reflecting motor inhibition is in obvious conflict with the, previously discussed, assumption that the prefrontal cortex—a structure presumably not involved in processing subliminal information—is always implied in inhibitory control. Four different explanations seem possible:

(a) The prefrontal cortex is, despite the evidence to the contrary, involved in the processing of subliminal stimuli.
(b) The NCE does not, after all, result from inhibition.
(c) The prefrontal cortex is not involved in the processing of subliminal stimuli, and the NCE does result from inhibition, but primes were not sufficiently masked, i.e., not subliminal.
(d) Primes were presented subliminally, the prefrontal cortex is not involved in the processing of subliminal stimuli, and the NCE does result from inhibition, but the inhibition process responsible for the NCE in the masked prime task is different from the top-down inhibition process responsible for the interruption of ongoing motor activity triggered by supraliminal stimuli.

We disregard explanation (a) on the basis of the available evidence both from published work (e.g., Frith & Dolan, 1996; Jueptner et al., 1997; Luria, 1992; Pisella et al., 2000; Posner & DiGirolamo, 1998; Shallice, 1988; Spence & Frith, 1999), and from as yet unpublished data from our own laboratories, indicating that the prefrontal cortex is not specifically involved in inhibitory control of subliminally primed motor responses.

We disregard explanation (b) for two reasons: First, on the basis of the go/nogo data reported above, and second, because there is no acceptable alternative explanation available. If the NCE is not due to inhibition, it can only be due to a selective activation of the opposite response *above* the (uninhibited) activation

level of the initially primed response (this interpretation does not conflict with the LRP data because the LRP is a *difference* waveform, i.e., it reflects differences in relative activation levels of contra- and ipsilateral motor cortices, not absolute activation levels). How would such a strong selective activation be brought about? The only conceivable explanation seems to be that a stimulus presented after the prime, i.e., the mask, actively triggers preparation of the opposite response. In the earlier experiments (e.g., Eimer & Schlaghecken, 1998, Exp. 1; Schlaghecken & Eimer, 1997), left- and right-pointing arrows served as prime and target stimuli, and the mask was composed of superimposed left- and right-pointing arrows and was presented immediately after the prime. One might argue that this effectively ''added'' the opposite stimulus to the prime. This could have resulted in the required selective activation of the opposite response. However, an increasing number of experiments have employed masking stimuli that do not share any features with primes or targets (e.g., Eimer & Schlaghecken, 1998, Exp. 2; Schlaghecken & Eimer, 2002). Importantly, similar (although somewhat smaller) behavioural and electrophysiological results were obtained, thus effectively disproving the hypothesis that a mask-induced activation of the opposite response underlies the NCE. It might thus be argued that masking with the opposite stimulus artificially increases the NCE, but it cannot be argued that it causes the NCE.

Finally, we disregard explanation (c) on the basis of the large number of various stimulus detection and identification tests, all demonstrating that participant's performance was at chance level (e.g., Dehaene et al., 1999; Eimer & Schlaghecken, 1998; Klotz & Wolff, 1995; Schlaghecken & Eimer, 1997). Furthermore, Klapp and Hinkley (2002) have demonstrated that when primes are unmasked (i.e., supraliminal), PCEs rather than NCEs occur, thus making an explanation in terms of ''inhibition triggered by residual prime visibility'' highly implausible.

This leaves explanation (d), that the NCE is caused by an inhibition process different from the one responsible for controlling motor activations in response to supraliminal stimuli. We have argued elsewhere (Schlaghecken & Eimer, 2001, 2002) that the NCE reflects the operation of a local self-inhibition process, which acts as an ''emergency brake'' system so that a strongly preactivated response becomes actively inhibited if the sensory evidence for this response is suddenly removed. Self-inhibition mechanisms, where the activation of a component directly causes its subsequent inhibition, have been described repeatedly in the literature (see, e.g., Arbuthnott, 1995; Houghton, Tipper, Weaver, & Shore, 1996). They are characterized by an activation-followed-by-inhibition sequence (similar to that obtained in masked priming), and generally show a close relationship between the amount of initial activation and the strengths of subsequent inhibition.

The same relationship was observed with masked primes. Only responses triggered by perceptually ''strong'' masked primes—presented at the centre of

fixation for about 16–33 ms—seem to be subject to inhibition once the respective sensory evidence is removed. Perceptually weak masked primes, in contrast, apparently fail to elicit inhibition. If masked primes are not presented foveally, but in the periphery of the visual field, PCEs instead of NCEs are obtained (Schlaghecken & Eimer, 1997, 2001). The same holds for primes presented foveally, but overlaid with visual noise (Schlaghecken & Eimer, 2002). Conversely, if primes are presented peripherally, but are made perceptually more salient by increasing prime–mask interstimulus interval (ISI), NCEs will reappear (Schlaghecken & Eimer, 2002).

In Schlaghecken and Eimer (2002), a simple functional model of subliminal motor control was presented. It comprised direct perceptuomotor links (both excitatory and inhibitory), local self-inhibition circuits in the form of interconnected ON/OFF-nodes, and an inhibition threshold from OFF- to ON-nodes. This model informally explains results from a wide variety of priming tasks, including PCEs turning into NCEs in subliminal priming and PCEs in supraliminal priming. It can also account for PCEs in the flanker paradigm (a central target ''flanked'' by distractors; Eriksen & Eriksen, 1974) and the decrease of the flanker PCE with increasing retinal eccentricity of the flankers (Goolkasian, 1997; Miller, 1991). However, this model has been formulated at a purely descriptive level. Thus, its full computational consequences have not been explored.

However, related computational models can be found in the literature. In particular, an important influence on our work has been the connectionist modelling of selective attention by Houghton and Tipper (1994). They have provided connectionist models of two selective attention phenomena: Negative priming (Tipper, 1985) and inhibition of return (Klein, 2000). Although different in nature to masked priming (which is not subject to high-level attentional control), inhibitory effects not unlike those arising in masked priming, can be observed in both inhibition of return and negative priming. In addition, Houghton and Tipper's model was based upon self-inhibitory ON–OFF circuits not unlike those used in (Schlaghecken & Eimer, 2002). Houghton and Tipper call these *opponent processes*.

Due to this common role for inhibition we were drawn to investigate whether Houghton and Tipper's (1994) connectionist modelling principles could be applied to the masked priming task. However, as previously suggested, the existing Houghton and Tipper model is not suitable for modelling masked priming. This is because, in both their negative priming and inhibition of return implementations, the release of inhibition is driven by (higher level) attentional mechanisms, which would classically be viewed to have their locus in prefrontal areas. Consequently, in this paper we seek a computational explanation that is ''dumb'' in the sense that recourse is not taken to high-level processes. Rather a model that reflects the characteristics of a direct nonconscious link from perception to motor action is developed. In undertaking this endeavour, we will

though use a number of mechanisms underlying Houghton and Tipper's modelling work.

Thus, the aim of the present study is to explore how principles from connectionist modelling, particularly those employed by Houghton and Tipper (1994), can be used to explain the available masked priming data. This will enable us to provide a computationally prescribed explanation of the mechanisms that underlie the masked priming effects. In order to do this we will develop a simple, but surprisingly powerful, computational model of human early motor control. This model has also been influenced by the explanation of masked priming effects given in (Schlaghecken & Eimer, 2002) and can thus also be viewed as a computational realization of that earlier descriptive model.

The benefits of constructing computational models are numerous, but most significantly such models force researchers to think hard about their theories, and prevent them from being able to hide behind imprecise natural language explanations. A running model can be objectively tested against available and newly arising data. This approach is particularly worthwhile when a large amount of empirical data is available to constrain the model; this is exactly the case with masked priming. Also, a major benefit of the existence of a computational realization is that it can be used to systematically generate predictions that can then be fed into further empirical work. The modelling work undertaken here will yield just such a set of concrete predictions.

We begin by summarizing the masked priming data in the next section. Subsequent sections describe the theoretical principles that underlie our model, present the model itself, and document the results of our modelling. The final two sections contain a discussion and concluding remarks respectively.

MASKED PRIMING

The basic paradigm

From amongst the various experimental paradigms, we particularly focus here on the masked prime task of Eimer and Schlaghecken (1998)—hereafter referred to as the ''basic paradigm''—which provides behavioural and electrophysiological data on priming, as well as data on prime visibility.

In the priming task, trial structure is as follows:

1. *Prime phase*. A prime is presented for 16.667 ms at the centre of fixation. The prime is either a response-mapped left- or right-pointing double arrow (''<<'', ''>>''), or a neutral (not response-mapped) inward- or outward-pointing double arrow (''><'', ''<>'').
2. *Mask phase*. A masking stimulus—consisting of superimposed left- and right-pointing double arrows—is presented centrally, immediately after prime-offset, for 100 ms.
3. *Target phase*. A left- or a right-pointing double arrow is presented

centrally, immediately after offset of the mask, for 100 ms. A left- or right-hand key press has to be executed as quickly and accurately as possible in response to target-arrow direction.

This experimental set-up yields three compatibility conditions:

- *Compatible*. Where prime and target are mapped to the same response, i.e., a directional arrow is presented in the prime phase, and the direction of the arrows in the prime and target phases is the same.
- *Incompatible*. Where prime and target are mapped to opposite responses, i.e., a directional arrow is presented in the prime phase, and the direction of the arrows is reversed between prime and target phase.
- *Neutral*. Where the prime is not mapped to any response, i.e., a neutral (nondirectional) stimulus is presented in the prime phase.

In the basic paradigm, compatible, incompatible, and neutral trials are presented with equal probability. These three conditions have been extensively investigated and have yielded NCEs on reaction times (RTs) and error rates as shown in Figure 3 (taken from Eimer, 1999) and LRP waveforms as shown in Figure 2.

In accordance with the ''landmarks'' provided by the LRP waveforms, we will use the following terminology in the remainder of the paper:

- *Prime-induced activation onset*. The point, around 200 ms after prime onset, at which compatible and incompatible waveforms show a significant enough deflection from the origin that it cannot be attributed to background activation fluctuations.
- *Reversal onset*. The point at which suppression starts to take affect, somewhere around the black arrow.
- *Inhibition-induced crossover*. The crossover point that is just before 300 ms, at which suppression has taken sufficient hold that direction of response separation changes.
- *Target-induced activation onset*. The point at which target activation starts to take affect, somewhere around the white arrow.
- *Target-induced crossover*. The compatible case crossover just before 400 ms at which target activation has taken sufficient hold that direction of response separation changes again.

The spectrum of experiments

Basic experiment

The starting point for our modelling is to reproduce the data arising from the basic experiment discussed in the previous subsection. One particularly notable point about this LRP profile—most clearly seen in the waveforms from

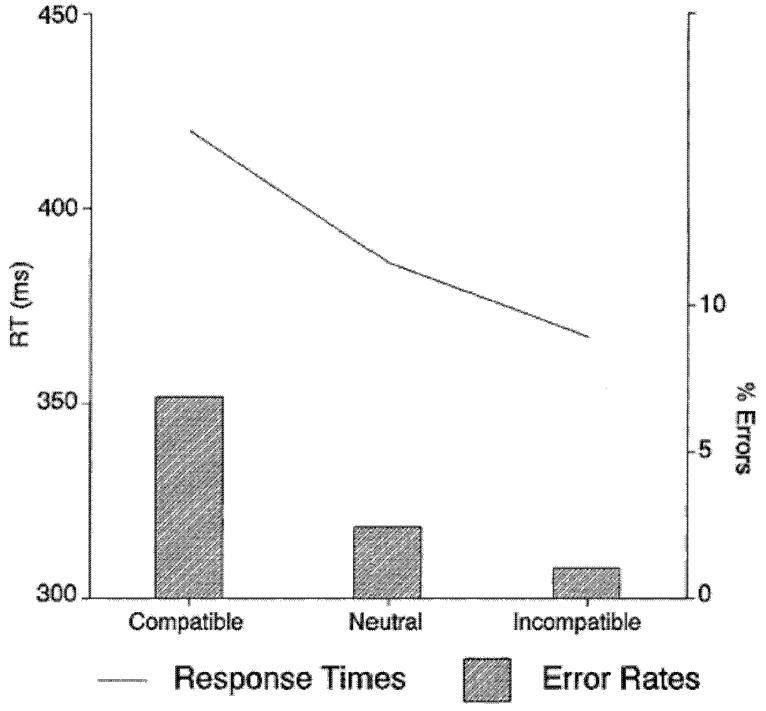

Figure 3. Reaction times and error rates for masked priming task (from Eimer, 1999).

compatible trials—is that the reversal deflection is substantially larger than the prime-triggered preactivation. The need to reproduce such a response profile will impose strong constraints on our model.

However, in order to further constrain the model we wish to reproduce the full spectrum of data discussed in the introduction. In the following, we will discuss the implications of these findings for computational models of masked priming.

Influence of SOA

PCEs obtained with short prime–target SOAs turn into NCEs with longer SOAs, reflecting the activation-followed-by-inhibition process typical for self-inhibitory circuits. The model thus has to be able to produce a biphasic temporal pattern, where initial separation in a particular direction is followed by a separation reversal.

Influence of prime strength

Only perceptually ''strong'' primes lead to an NCE (indicating the presence of self-inhibition), while perceptually ''weak'' primes result in PCEs only (indicating the absence of self-inhibition). This suggests the operation of some sort of threshold mechanism, which controls the release of inhibition back onto response nodes—only activation triggered by high strength primes will cross this threshold and trigger the inhibitory reversal.

Prime visibility results

As noted above, a number of different prime visibility tests have yielded strong evidence that masked primes are not available to conscious awareness. In the forced choice variant of the basic paradigm, for example, masked directional arrow primes are presented with different durations. Results clearly indicated that for 16.666 ms primes, identification performance was at chance level (Eimer & Schlaghecken, 1998). This is in itself a revealing outcome since, as is evident from the LRPs shown in Figure 2, the prime does cause an activation of its corresponding motor response. That is, prime-induced activation propagates right through to response systems. Indeed there is a considerable profile of motor cortex activation before the target is presented and there is no a priori reason to believe that this activation will fail to arise in the forced-choice experiment. The implication then of the forced-choice data is that in the absence of a target stimulus and/or in the absence of speeded RT instructions, these deflections do not have any effect on response outcomes. In other words, residual activation induced by the prime only has an effect on overt responses if it is built upon by target-induced activation.

This suggests that there is some form of *selection criterion* working at the response end, which could be defined in terms of motor cortex activation levels. At first glance a simple threshold might seem to be a sufficient selection criterion, i.e., once the separation of the two response alternatives crosses a particular threshold, an overt response is released, while activation below this threshold has no effect on outcomes. However, for reasons that will become clear in the following subsection, a simple threshold is not sufficient.

Masked and unmasked priming

Perhaps the strongest constraints imposed on our model come from four as yet unpublished experiments, which varied the masking of the primes. The experimental set-up was similar to the basic experiment, with the exceptions outlined below.

Full mask. This experiment is essentially a replication of the basic experiment described above. However, instead of a superimposed-arrows mask, a scrambled pattern mask (a rectangular array of overlapping lines of

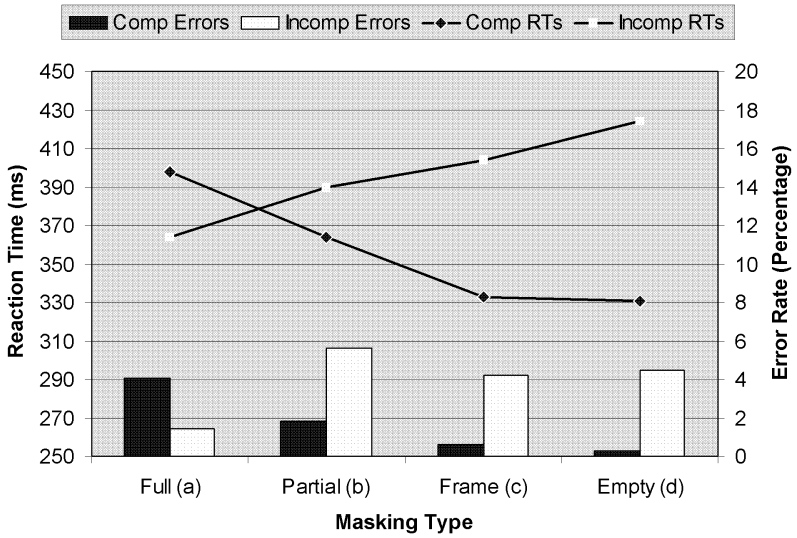

Figure 4. Reaction times and error rates on compatible and incompatible trials as a function of masking conditions: (a) Full mask, (b) partial mask, (c) frame, and (d) empty (no mask) 150 ms ISI experiments.

different length and random orientation) was employed, and a 50 ms empty interval, inserted between mask offset and target onset, increased prime–target ISI (mask–target SOA) to 150 ms. Forced choice and staircase tests have confirmed that primes are successfully masked with this procedure, and behavioural and LRP results (cf. Figure 4a and Figure 5a) replicate the earlier findings (cf. Figure 2 and Figure 3).

There is one small but pertinent feature of these data that it is worth emphasizing. With the longer ISI, there appears to be sufficient time between prime and target presentation for the inhibitory reversal to complete and for activation to reverse again (''*double reversal*''). This is most evident in the incompatible condition, where the LRP shows a tendency to return to baseline prior to the final separation in the target direction. Having to reproduce this double reversal will impose strong constrains on our model.

It is also noteworthy that the behavioural NCE is smaller in this experiment (approximately 30 ms) than in the basic experiment (approximately 50 ms). This might be due either to an artificially increased NCE with arrow masks in the basic experiment (see the ''NCE and self-inhibition'' subsection on p. 391) or to the onset of the double reversal. Of course, a combination of both factors might be possible. We will return to this issue later (cf. the ''Comparing ISI 100 and ISI 150 experiments'' subsection on p. 440).

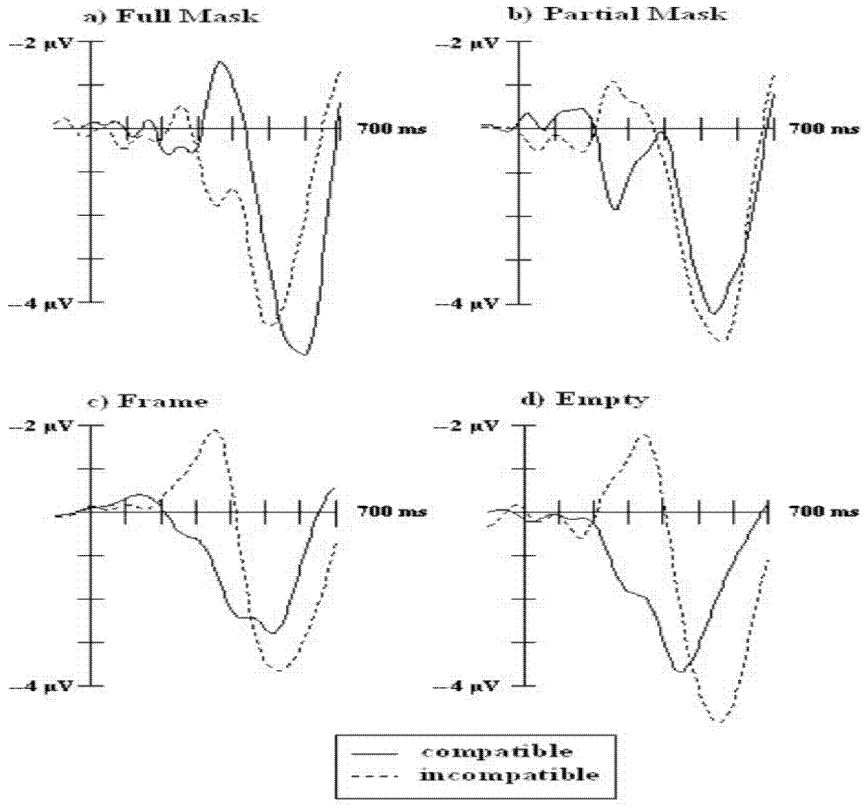

Figure 5. LRPs for 150 ms ISI experiments: (a) Full mask, (b) partial mask, (c) frame, and (d) empty (no mask).

Partial mask. This experiment is identical to the one described above, only that the number of lines constituting the mask has been substantially reduced (3 instead of 30). This results in incomplete masking and—correspondingly—supraliminally presented primes. The LRP (Figure 5b) shows that, unsurprisingly, the initial prime-triggered activation is much stronger than with fully masked primes. Also, it can be seen that the partial mask causes some inhibitory effect (cf. the return of the initial separation to baseline on compatible trials). However, this inhibition is insufficient to cause a full reversal of relative activation levels. Behavioural PCEs are obtained in this experiment (cf. Figure 4b), as can be expected with supraliminal primes.

However, from a closer inspection of the LRP one might have expected NCEs to occur, as the incompatible condition separates towards the final minimum earlier than the compatible condition does. Its minimum point is later

(in accordance with the observed PCE), but it certainly separates faster. It is this effect, which, we argue, suggests that the selection criterion for response execution cannot be a simple threshold (cf. the ''Prime visibility results'' subsection on p. 397). If the criterion for judging when response separation generates an overt response was a simple threshold, we should obtain NCEs, since the incompatible condition would reach the threshold first!

We therefore suggest that the selection criterion is an accumulation of separation over time, in which ''wrong direction'' separation counts negatively. Thus, PCEs are obtained in this experiment because the incompatible condition initially accumulates wrong direction separation, which has to be overcome when separation moves in the correct direction. In contrast, the compatible condition only accumulates correct direction separation before target activation takes effect, thus reaching the point of selection criterion satisfaction earlier. If we locate the point of selection criterion satisfaction at the final minimum of a separation curve, this accumulation-based interpretation of selection is consistent with the observation that the incompatible waveform has a lower minimum than the compatible waveform. In fact, we will argue that such an accumulation-based separation criterion is consistent with neural network response selection mechanisms that have previously been employed (e.g., Cohen, Dunbar, & McClelland, 1990).

Frame. Again, the general experimental set-up was as described above, but an empty rectangular ''frame'' was presented instead of a mask. Thus there was effectively no masking, although the same outer contour as the mask was presented for the usual 100 ms interval. Since the prime was clearly visible on each trial, participants were explicitly instructed not to respond to it, i.e., to withhold their response until the target presentation. As expected, behavioural PCEs were obtained (cf. Figure 4c), and the LRP waveforms show no evidence of an inhibitory reversal (cf. Figure 5c).

This data gives strong justification for the role of a (successful) mask in generating the inhibitory reversal, as without a successful mask, the reversal does not arise. Unsurprisingly, the initial separation build-up is even larger than in the partial mask experiment. Broadly speaking, the difference between contra- and ipsilateral motor cortex activation—starting around 200 ms after prime onset—continues to increase throughout the next 150 ms (although the rate of increase slows down somewhat over time). In addition, as can be seen clearly for the incompatible condition, around 350 ms after prime onset, the initial separation reverses sharply. Both the timing and the slope of this reversal suggest strongly that it reflects the onset of target-related motor processes, rather than a mere decay of prime-triggered motor processes.

Obtaining the sustained prime-induced separation that we see for this condition will impose strong constraints on our model. We argue that it is suggestive of two mechanisms. First, the perceptual trace of an unmasked prime must decay relatively slowly in order to maintain prime-direction response

separation throughout the 150 ms ISI (discussed in more detail in the ''Sustained perceptual trace'' subsection on p. 402). Second, there has to be a mechanism by which separation at the response end is accentuated over time, because a decaying perceptual trace seems insufficient to actively increase the motor activation separation until target separation takes effect. Consequently, we propose that response nodes are competing and that selected directions will be accentuated over time through this process of competition. The details of this mechanism will be discussed later (cf. the ''Response competition'' subsection on p. 404).

Empty interval. In this experiment, a 150 ms blank screen is presented between the prime and the target. This produces an almost identical LRP to the frame experiment (Figure 5d), giving further evidence that it is the sudden removal of sensory evidence with successful masking that causes the inhibitory reversal. In addition, the behavioural data is largely also in line with that for the frame experiment (cf. Figure 4d), i.e., a (somewhat larger) PCE occurs, with similar mean RTs and error rates.

The role of the task set

In several experiments, we have investigated whether compatibility effects in masked priming are automatically elicited by the directional information inherent in the arrow-primes, or whether they require the presence of a particular task set, i.e., knowledge about particular stimulus–response (S–R) mappings, and the intention to use them. Results strongly supported the latter assumption. Arrow-primes did not elicit compatibility effects when targets were letters instead of arrows (Eimer & Schlaghecken, 1998), and arrow-primes that were mapped to one response modality (e.g., hand responses to arrow-targets) failed to elicit compatibility effects in a different response modality (e.g., foot responses to nonarrow-targets). Although, compatibility effects with nonarrow-targets were reliably found when these targets were mapped to the same response modality as arrow targets (e.g., Eimer, 1999; Eimer & Schlaghecken, 2001). These findings clearly demonstrate that compatibility effects in masked priming represent a mixture of voluntary and automatic processes. They are automatic to the extent that they are triggered by stimuli the participant is not consciously aware of, but at the same time, they depend on S–R mappings as defined in the instructions.

THE THEORY

The neural network model that we describe in the section after this is a concrete realization of our theory of masked priming. However, before we discuss this formal realization, we first motivate and informally introduce the central mechanisms of this model.

Sustained perceptual trace

As noted above, we assume that the initial prime-related phase of motor response separation is not the result of a one-off (16.666 ms) input to the motor system, but is supported by input from a sustained, though decaying, ''perceptual trace''. The psychological justification for this is the literature on (preattentive visual) iconic memory (Baddeley, 1997), for which time bounds as much as 500 ms have been suggested (Baddeley, 1997). In addition, it is known from single cell recordings that briefly presented, unmasked stimuli result in an increased firing rate of subcortical (Schiller, 1968) and cortical (Rolls, Tovee, & Panzeri, 1999) neurons, which outlasts stimulus offset for several hundred milliseconds. Although, it is obviously the case that presentation of a successful backward mask will cut off the sustained perceptual trace, as it will do in our model.

The LRP data we are working from supports the assumption of such a decaying perceptual trace. With unmasked primes (Figures 5c and 5d), the separation between motor cortex activation levels—which starts 200 ms after prime onset—increases for approximately 150 ms. This suggests that there is either a very powerful response competition mechanism at work, or that a substantial ''afterimage'' of the prime is preserved in the unmasked conditions.

Although, we will include response competition in our model (cf. the ''Response competition'' subsection on p. 404), we have found that if response competition is made sufficiently strong to simulate the initial separation profile in unmasked conditions, it will severely disrupt the models' overall performance. This is especially the case when testing successful masking conditions (since the opponent process would have to fight against it when separation sign reversal occurs in strongly masked experiments). Thus from a modelling perspective, as well as from the available physiological data, assuming a sustained perceptual trace appears to be the best option. In the present model, we obtain the required decaying trace using a double application of time averaging, which we describe in the ''Double time averaging'' subsection on p. 415.

Masking and competition between percepts

The next question to answer is how to implement masking. A number of competing theories for this are available (see, for example, the conflicting proposals discussed in the articles by Enns & di Lollo, 2000, 2002; Keysers & Perrett, 2002a, 2002b). However, it is beyond the scope of this paper to enter fully into this debate. Thus, we have selected a mechanism that is both consistent with one of the main theories of masking and with the empirical data we are striving to model.

Specifically, we will implement masking as a *neural competition mechanism* (Keysers & Perrett, 2002a). In an intuitive sense the masked and masking stimuli compete for limited resources. Over time, this competition drives the system to emphasize a single percept, and the process only stabilizes once this single

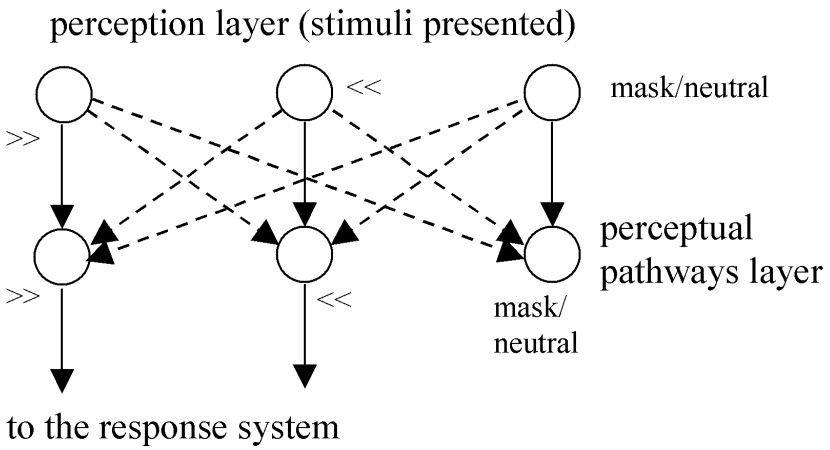

Figure 6. Perceptual system.

percept has come to dominate. This competition mechanism will be implemented using feedforward inhibition, which will be located between a ''perception'' layer at which input stimuli are presented, and a ''perceptual pathway'' layer, at which the decaying trace of the percept is sustained. Using simple localist representations, as we do, the resulting mechanism is as shown in Figure 6 (where each node in each layer codes for the percept indicated and inhibitory links are shown as broken lines).

We also investigated placing lateral inhibition between perceptual neurons, which has often been proposed as a masking mechanism (e.g., Rolls & Tovee, 1994). However, we were unable to generate the required ''sharp'' masking effect.

Note that—in accordance with task demands—the mask percept is not mapped to a response. Indeed the mask/neutral node in the lower (perceptual pathways) layer actually plays no role in the remainder of the model since it has no outgoing links, and it could be removed without affecting our simulation results (we merely include it for completeness of presentation).

The key aspect of the mechanism is that activation of a stimulus at the perception layer suppresses the traces of previously perceived (competing) stimuli. In particular, excitation of the mask perception layer node will both excite the corresponding mask percept and suppress the perceptual trace of any other (previously) excited stimuli (e.g., the prime).

Precedents for such a feedforward inhibition-based mechanism can be found in the modelling literature. For example, Grossberg, Mingolla, and co-worker's influential boundary contour system (Francis, Grossberg, & Mingolla, 1994; Grossberg & Mingolla, 1985a, 1985b) makes liberal use of feedforward inhibition. For example, in their first competitive stage, complex cells (which have receptive fields corresponding to complex visual cortex cells) feed inhi-

bition forward on to hypercomplex cells at competing spatial locations. Through such mechanisms masking effects (in particular, metacontrast masking) have been modelled (Francis, 2000; Francis et al., 1994). Thus, although they use feedforward inhibition at a much more microscopic (and biologically detailed) level than we do, the principle of feedforward-induced competition between percepts is the same. In addition to the literature supporting such neural competition-based accounts of masking, we will also show that the pattern of unmasked and masked activation we obtain in perceptual areas of our model are broadly consistent with results from single cell recordings (see the ''Emergent sustained perceptual trace'' subsection on p. 415).

Response competition

We also postulate that there is a competitive mechanism at work at the response end of the system. From a theoretical perspective this is an entirely reasonable assumption: Multiple responses can be excited at the same instant (as some level of evidence is available for each response); however, only one response will be executed during performance of the experimental task. Competition ensures this by emphasizing the dominating response at the expense of competitors.

Furthermore, the LRP data makes clear that throughout the time course of the experiment, different responses become dominant. For example, in the compatible condition depicted in Figures 2 and Figure 5a, the dominant response changes a number of times. Initially, the ''correct'' response dominates, then the ''incorrect'' response, and finally it switches back to the ''correct'' response. Such switching between responses is consistent with some form of competition between response outcomes—as evidence for one response increases, its capacity to dominate other responses increases. However, the available evidence may then change, and an alternative response becomes dominant.

Perhaps the strongest justification for response competition comes from the unmasked priming experiments (''frame'' and ''empty'', Figures 5c and 5d). As already discussed in the ''Masked and unmasked priming'' subsection on p. 397, and the ''Sustained perceptual trace'' subsection on p. 18, it is notable that on the basis of only a 16.666 ms prime, response separation continues and even increases throughout the 150 ms prime–target ISI. In addition to the assumption of sustained perceptual traces, we also assume that the small excitatory ''push'' given by the prime initiates a temporally sustained competition between responses, with the primed response slowly but in a sustained fashion suppressing competing responses and in turn being disinhibited (i.e., released from strong inhibition) by their suppression. The mechanism we use to implement this competitive process is lateral inhibition between response nodes. Since there are only two response nodes, the competitive mechanism is rather simple; it can be depicted as shown in Figure 7, where each node represents a particular response, e.g., left or right. The relative strength of competing responses determines which will eventually dominate, with the strongest suppressing the other response over

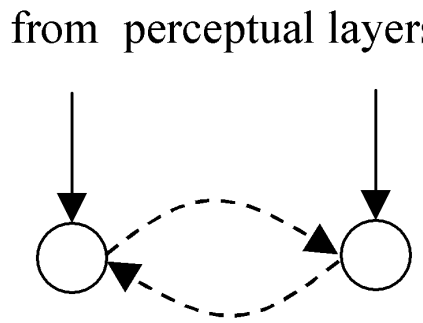

Figure 7. Response system 1.

time, i.e., a ''winner take all'' dynamics. Of course a human has the capacity to make more than two responses—however, we will typically only depict the two task relevant responses. This is because a mechanism we introduce later (cf. the ''Response-set maintenance, response foregrounding, and selection'' subsection on p. 409) implements a foregrounding mechanism that is functionally consistent with this two-response perspective.

Opponent processing

As previously discussed, our modelling work has been strongly influenced by Houghton and Tipper's (1994) model of negative priming and inhibition of return. The central element of their model that we have inherited is their opponent network. A number of different incarnations of opponent processing can be found in their papers (see Houghton & Tipper, 1994; Houghton et al., 1996; Jackson & Houghton, 1994). However, the central ideas are the same. First, nodes are designated to reflect response activation build-up. For the masked priming task, a node would be allocated to each of the two possible responses—left- and right-hand selections. Activation build-up at one of these nodes reflects increasing preparation of that response. In fact, in our masked priming model the interpretation of such activation build-up is slightly more complicated. This is because the difference in activation between the two possible responses will actually be the indicator of increasing evidence for a particular response. We will make this mechanism precise shortly. However, for the moment a broad understanding of the mechanism is sufficient.

The key element though of an opponent process is that an opponent (OFF) node is associated with each response, and these two nodes are linked via an excitatory link from the response to the OFF node and an inhibitory link from the OFF node back to the response, see Figure 8.

In terms of function, the opponent node regulates activation in the associated response node through the release of inhibition. Thus, as activation builds up at a

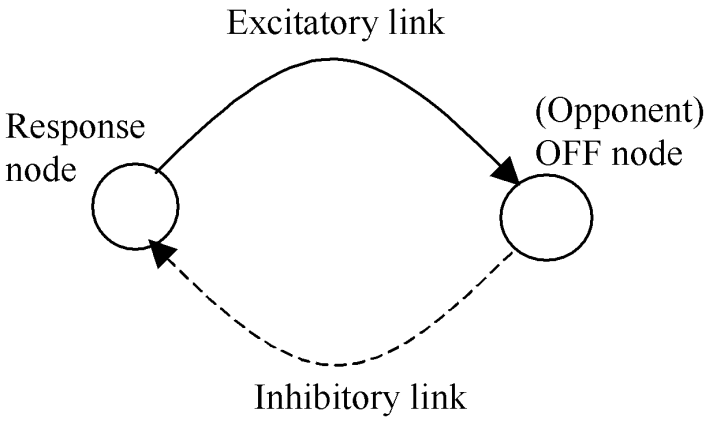

Figure 8. Response system 2: Opponent process.

response node there is a delayed build-up of activation at the OFF node. Eventually the OFF node feeds inhibition back onto the response node. In addition, by adapting the activation function applied in the OFF node, e.g., how quickly it saturates or how responsive it is to low levels of input (which amounts to a graded input threshold), the time course with which inhibition is released onto the response node can be regulated.

The available empirical findings on the masked priming task make it clear that it is the presentation of the mask that causes the inhibitory reversal (see the ''Masked priming'' section on p. 394). The inhibitory pressure is initiated by the removal of sensory evidence: The mask suppresses the primed percept, thus removing sensory evidence for the corresponding response and initiating its suppression. The hypothesis that we investigate here is that this behaviour arises from a low-level ''emergency braking'' mechanism, which can be realized by an opponent processing circuit. We can depict the mechanism as shown in Figure 9, where we also include the lateral inhibition previously introduced.

This mechanism, representing the basis for inhibition-controlled direct perceptuomotor links, naturally satisfies two notable requirements:

1. Inhibitory forces should not be so strong that they interfere with ''appropriate'' response activation build-up. Thus, if the build-up and persistence of sensory evidence is strong and stable, the corresponding response activation build-up should dominate any inhibitory forces (cf. the ''frame'' and ''empty'' experiments).
2. When sensory evidence changes (e.g., due to mask presentation), inappropriate responses should be suppressed quickly and efficiently (cf. ''basic'' and ''full mask'' experiments).

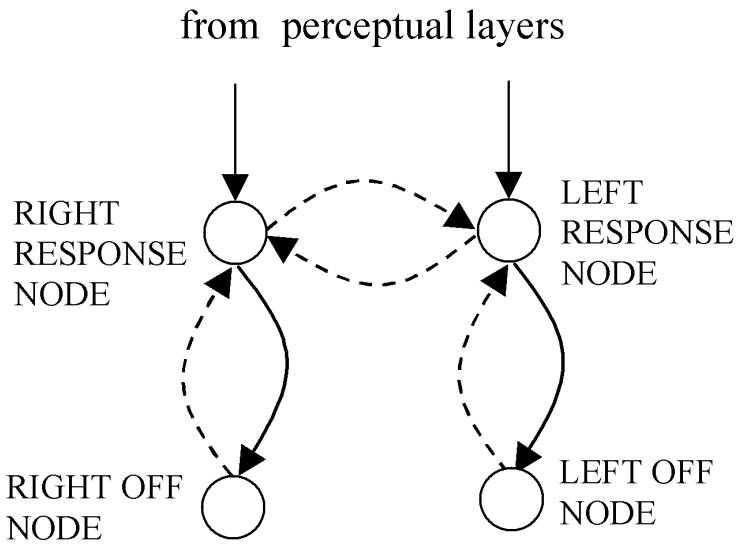

Figure 9. The full response system.

Opponent processing provides an elegant realization of both these requirements. The former is obtained because strong and stable ON (response) node build-up will counteract any inhibition being fed back from the corresponding OFF node. The latter arises because removal of sensory evidence for a particular response causes the (already) built-up OFF node activation to be released (since it is no longer counteracted by ''bottom-up'' excitation), so that the response is rapidly suppressed.

We would also argue that the need for the second of these requirements to be satisfied rules out a solution based purely upon lateral inhibition between responses. Although we postulate lateral inhibition between response nodes (see the ''Response competition'' subsection on p. 404), we believe that on its own, this mechanism is insufficient to obtain the effects described above. Critically, lateral inhibition can only work to emphasize an existing activation difference between competing nodes—however, removal of sensory evidence does not in itself have any excitatory effect on the competing (nonprimed) response. Thus, in the absence of opponent mechanisms, lateral inhibition between responses would not generate the separation sign reversal that is characteristic of the NCE. We provide simulations to justify this argument in the ''Insufficiency of lateral inhibition alone'' subsection on p. 424.

The selection criterion

We argued earlier (cf. the ''Masked and unmasked priming'' subsection on p. 397) that instead of a simple threshold, accumulation of separation over time

(with ''wrong direction'' separation counting negatively) should act as the final response selection criterion. This argument is based upon the observation that behavioural benefits for one type of trial might be obtained even when LRP waveforms indicate that it is on the *other* type of trial that response separation starts earlier (see ''partial mask'' experiment, Figures 4b and 5b). Furthermore, it is noteworthy that there is a systematic relationship between behavioural effects, separation direction immediately before target-related motor separation takes effect, and size and latency of the final LRP minimum. With fully masked primes (behavioural costs on compatible trials), compatible-trial LRPs show wrong direction separation prior to target-related motor activity and their final minimum is larger (and occurs later) than that of incompatible trials (Figures 2 and 5a). Conversely, with unmasked or partially masked primes (behavioural costs on incompatible trials), it is the incompatible-trial LRP that shows wrong direction separation and the larger and later final minimum (Figures 5b, 5c, and 5d). We therefore assume that wrong direction separation occurring prior to target-related separation has to be compensated for in the size of the final correct direction deflection. Since the separation slope seems to be relatively unaffected by experimental conditions, a larger final minimum will also occur later than a smaller one.

In order to reproduce these data, we have employed an accumulator-based selection criterion. Thus, in addition to activation driven response nodes, an accumulator node is included which is updated according to the difference in activation level between response nodes. Since we only have two responses, we only need a single accumulator node, for which positive accumulated evidence indicates one direction of response and negative accumulated evidence the opposite response. Generalization of this technique to more than two nodes would require a dedicated accumulator for each response (see, e.g., Ratcliff, 1978).

With the single accumulator neuron approach employed here, if a particular response, say the left response, is more highly active than the competing (right) response at a particular time point, a positive difference is added to the accumulator. A more strongly active right response, in contrast, will cause a reduction in the accumulator value. We also employ a discounting mechanism, whereby the significance of older separation information is progressively discounted. In effect, accumulator node activation decays over time and more recent separation information is more influential upon the decision process than older separation information. Thus, this node yields a measure of the *discounted accumulated separation between response alternatives*. In particular, negatively valued accumulation has to be compensated for by a corresponding sustained level of positive separation over time before the accumulator node becomes positive, and vice versa. The actual selection criterion is then a threshold on the accumulator node—once accumulated evidence crosses this threshold, either in a positive or a negative direction, the corresponding response is adjudged to have been released for execution. This particular approach has its roots in random walk and diffusion process models (see for example, Ratcliff, 1978), and

a number of neural network models (e.g., Cohen et al.'s, 1990, Stroop model) employ related mechanisms.

Response-set maintenance, response foregrounding, and selection

The results discussed earlier (cf. ''The role of the task set'' subsection on p. 401), demonstrating that compatibility effects in masked priming depend on the currently active task set, clearly suggest an important role for higher level cognitive processes even at this automatic level of motor control. However, results also suggest that these processes operate at the general task level rather than at the individual trial level. That is as long as, for example, left- and right-pointing arrows are mapped to left- and right-hand responses, manipulations of the probability of individual trial types do not affect compatibility effects (see Schlaghecken & Eimer, 2001).

We aim to account for these findings by introducing a ''response-set maintenance node'' (see Figure 10). This node provides excitatory input to the task-relevant ON and OFF nodes, thereby ''foregrounding'' the relevant response channels from the remaining channels, which stay at baseline activation levels. Note that this foregrounding is symmetric between the response alternatives, so

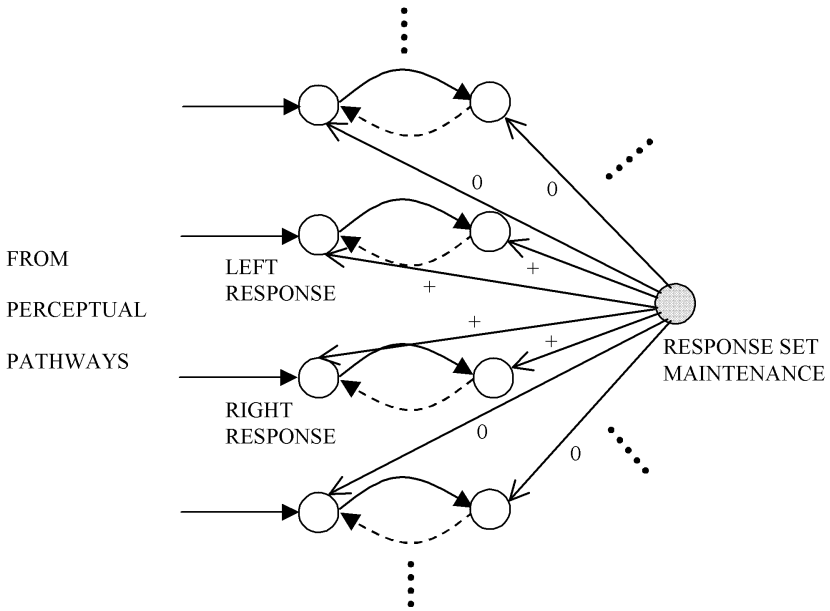

Figure 10. Response-set maintenance (for simplicity of presentation, lateral inhibition between response nodes is not depicted). Links labelled ''0'' are inactive in this particular experimental condition, thus maintaining the ''backgrounding'' of irrelevant pathways; links labelled ''+'' are relaying sustained foregrounding excitation.

that even though response nodes will become more active, the difference in response node activation (i.e., separation) will remain at zero until external stimulation triggers the build-up of asymmetric activation.

Similar preactivation mechanisms have been used in other neural network models (e.g., Cohen et al., 1990). In that model, nodes are by default placed in a very unresponsive part of their activation function. This is achieved by giving nodes a strong negative bias, thus placing them below the graded threshold built into the low net input end of a sigmoidal activation function. However, additional ''task demand units'' relay compensatory positive activation into task-relevant nodes, thereby placing them in the most responsive region of their activation function. This effectively foregrounds task-relevant from task-irrelevant nodes, in a similar manner to the mechanism employed here.

We have also endeavoured to tie response selection in with the response-set maintenance mechanism. As previously stated, the point at which a response is selected is determined by the selection criterion, i.e., it corresponds to the point at which discounted accumulated evidence crosses a particular threshold. However, it is reasonable to assume an additional ''switching off'' mechanism, which would ensure that separation quickly returns to zero once a response has been released for execution (cf. LRP profiles in Figures 2 and 5). In the present model, this mechanism is implemented by an inhibitory link from the accumulator node to the response-set maintenance node. Once the response criterion has been met (i.e., a motor response has been released), the accumulator node inhibits the response-set maintenance node, thus enabling the relevant response channels to return to background activation levels (see Figure 11).

Notice that there are alternative ways in which we could obtain this effect. For example, we could build inhibitory links from the accumulator system to the relevant ON and OFF nodes which would suppress these response channels when a response execution is reached, as signalled by the accumulator outputting a 1 (see the ''Response-set maintenance and response selection'' subsection in the Appendix on p. 463 for more details of the accumulator mechanism). This would allow the response-set maintenance node to continually feed excitation into the relevant pathway and would be consistent with S–R mappings being sustained throughout a sequence of trials. Notice though that within the context of the work being presented here, which is only seeking to reproduce single trial data, this accumulator inhibition approach is functionally equivalent to the approach implemented.

THE MODEL

We now move to a discussion of how the model is actually put together. A number of our models have been constructed using the Stuttgart Neural Network Simulator (version 4.2; SNNS Team, 2001). However, the data presented in this paper has been taken from an Excel implementation of our model.[1]

[1] All these models are available on request.

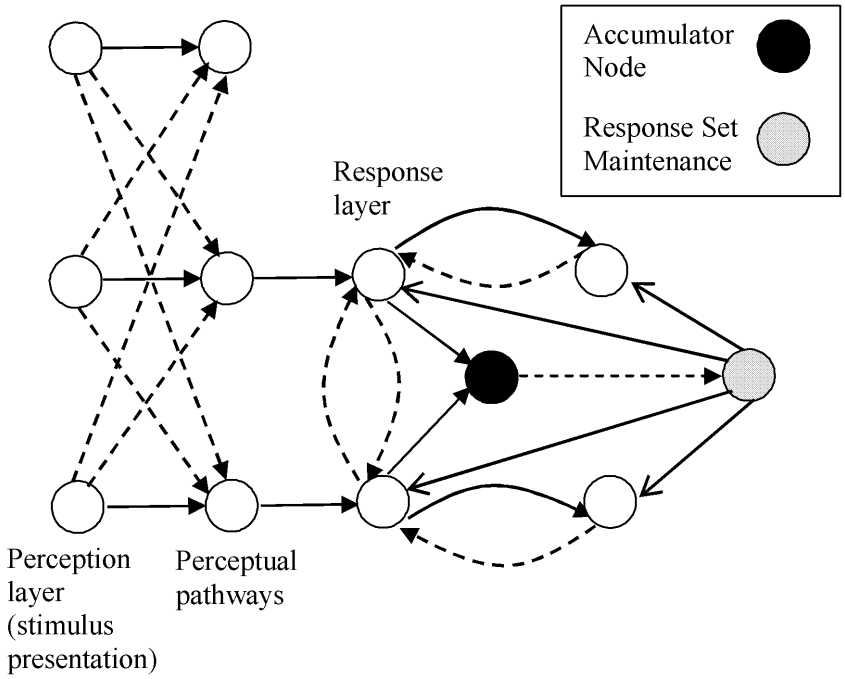

Figure 11. Accumulator and response-set maintenance system (for simplicity of presentation, nonrelevant response channels are not depicted).

There are two types of node in the network—input nodes and hidden nodes. An execution of the net corresponds to a sequence of cycles. On each cycle a new activation pattern is presented to the input nodes of the net and activation of all nodes is updated accordingly. In particular, hidden nodes are updated as a function of current activation and the new activation flowing into the node; an activation function, which will be introduced shortly, is defined for this purpose.

Configuration

The basic configuration of the model is as shown in Figure 12. Although, for simplicity of presentation, accumulator node, response-set maintenance node, and task irrelevant response pathways are not shown.

In terms of high-level structural configuration, the model contains three layers—a perception layer, a perceptual pathways layer, and a response selection layer. These layers reflect the distinction between input and hidden nodes: All perception layer nodes are input nodes and all other nodes are hidden.

It is also important to note that, in terms of our perception and perceptual pathway layers, the process of object recognition is not directly modelled.

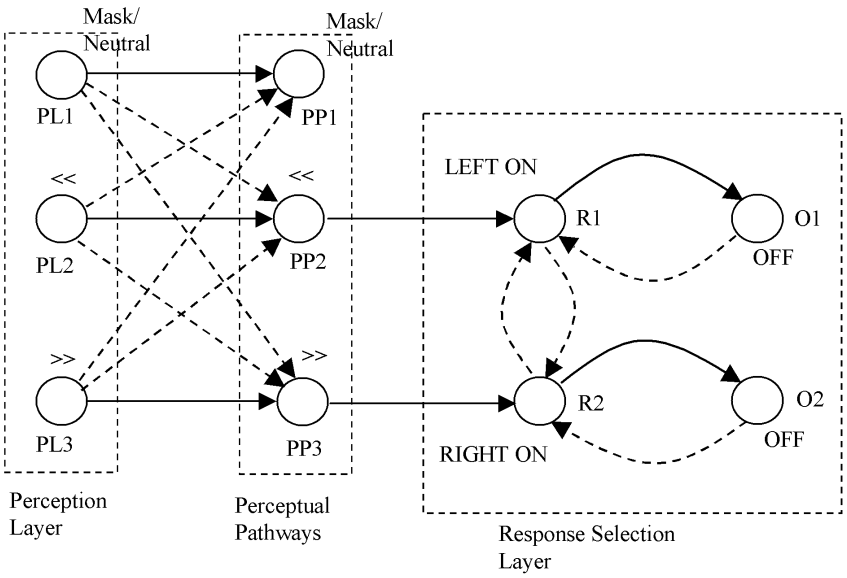

Figure 12. The model.

Indeed our model largely abstracts from the question of the level of cognitive representation of stimuli that is built during performance of the task. In effect, the perceptual pathway nodes denote (through their activation level) *the extent of the perceptual evidence* for a particular stimulus, and we make no assumption about how activation levels correspond to construction of object representations.

Activation of node PL1 is used to model the mask or the neutral prime being presented, while activation of node PL2 corresponds to a left double arrow stimulus and activation of node PL3 to a right double arrow stimulus. Notice one pair of perceptual nodes (PL1 and PP1) represents both neutral and mask stimuli. In terms of the particular stimulus sequence we use, this is functionally equivalent to including two pairs of nodes and adding an extra set of inhibitory links. Also notice that there is no path from mask or neutral stimuli to a response. This is consistent with the previously discussed foregrounding mechanism and the task demands.

Nodes PP1, PP2, and PP3 implement what we will call the perceptual pathways from perceptual stimuli to (corresponding) response nodes. Each perceptual pathway node implements a temporally sustained trace of the corresponding stimulus—the mask/neutral stimulus for PP1, << for PP2, and >> for PP3.

Nodes R1 and R2 represent a response layer, i.e., each node denotes a particular motor response channel. Node R1 corresponds to a left-hand response and node R2 to a right-hand response. Perception of the left and right stimuli

causes excitation of their corresponding response node via the perceptual pathways, through excitatory links from PL2 to R1 via PP2 and from PL3 to R2 via PP3. O1 and O2 are the respective OFF (opponent) nodes for R1 and R2. Thus, when, for example, activation builds up in R1, it also causes activation to build up in O1.

Basic activation functions

As previously explained, the activation level of input nodes is determined externally through the presentation of input patterns. In contrast, the activation level of hidden nodes is evaluated internally. Time averaging over a sigmoidal activation function is used for this purpose. This is a standard approach in neural networks that model activation change over time and is closely related to McClelland's (1979) cascade approach.

First, we denote the input to an arbitrary node i on cycle t+1 as $net_i(t+1)$. This is defined, in standard fashion, as a weighted sum of the activation currently being input to the node,

$$net_i(t+1) = \sum_{1 \le j \le K} (w_{ij} \times a_j(t)) \tag{1}$$

where j indexes node i's K predecessor nodes in the network, w_{ij} is the weight on the link from node j to node i, and $a_j(t)$ is the activation of the jth predecessor node on the tth cycle. In particular, activation arriving along negatively weighted links will reduce net_i and will hence have an inhibitory effect. In addition, as indicated by them not having a cycle parameter, weights are fixed throughout our simulations. We discuss this issue later (cf. the ''Interpretation of the model'' subsection on p. 420).

Using these concepts, we can now define the activation function that we use for hidden nodes.

$$a_i(t+1) = \tau\, a_i(t) + (1-\tau)\sigma(net_i(t+1)) \tag{2}$$

where τ is a constant between zero and one, which regulates the temporal dynamics of the time averaging.

The function expresses the new activation level of a hidden node, i.e., $a_i(t+1)$, in terms of its current activation, i.e., $a_i(t)$, and any new activation being input into the node, i.e., $net_i(t+1)$. A level of temporal activation stability is preserved by the term $\tau a_i(t)$. However, the function ''leaks'' at a rate determined by the constant τ. Thus, broadly speaking, temporal stability of the activation function increases as τ increases. In other words, when τ is large (i.e., approaches 1) the activation level on previous iterations becomes more significant in determining the new activation level. Thus, large fluctuations in net input will have less effect on the value of a_i than they would with small values of τ.

The second summand of our activation function—(1–τ) σ($net_i(t+1)$)—determines the influence of new activation entering the node. σ is a standard sigmoidal function (see Figure 13), which has three defining parameters denoted, *bs*, *sp*, and *rg* (which for simplicity of presentation we often omit). It is defined as follows,

$$\sigma_{bs,sp,rg}(X) = \frac{rg}{1 + e^{-\frac{(bs+X)}{sp}}} \tag{3}$$

where, *bs* is the bias term, which regulates the basic excitability of the node (see Figure 13), *sp* is the steepness of the activation function, i.e., how responsive it is to changes in its input (see Figure 13), and *rg* determines the output range of the function, i.e., the corresponding node can be activated between zero and *rg* (see Figure 13). We will discuss the settings we have chosen for these parameters and their implications shortly.

Perceptual pathways and masking

As previously motivated, perceptual pathway nodes maintain a short-term (decaying) trace of sensory stimuli. This sustained trace is implemented using a double time averaging mechanism.

It would though be wrong of us to associate our perceptual pathway nodes with a single visual region. Rather we would argue that our definitions produce

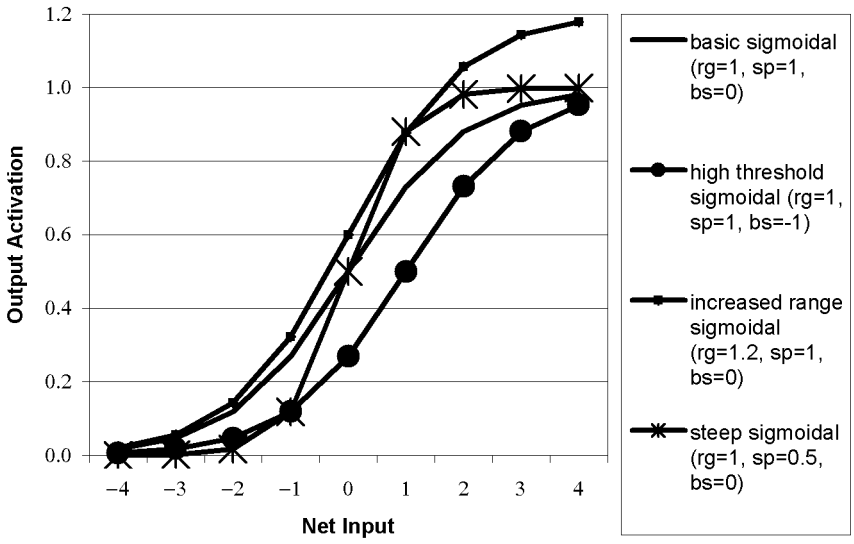

Figure 13. Sigmoidal activation function.

an overall (emergent) effect that is commensurate with the composite global behaviour of the relevant visual regions. In fact, Rolls and Deco (2002, p. 159) have argued that perceptual traces are sustained through mutual excitation exchange between the neurons that constitute a particular percept, this activation exchange being supported via the recurrent collaterals in an autoassociative (attractor) network. In a fully localist representation like ours, such a solution is not possible since there is no distribution of processing between neurons and thus, no possibility of recurrent collaterals sustaining the trace. The time averaging dynamics we implement could be viewed as an abstraction of such a more detailed mechanism of distributed encoding.

Double time averaging

In order to obtain the required temporal stability of representation we assume that time averaging is applied to both the input conductances entering the perceptual pathway nodes and the activation build-up itself. Thus, net input is sustained using the following time averaging mechanism:

$$NET_i(t+1) = \tau_1 NET_i(t) + (1-\tau_1) net_i(t+1) \tag{4}$$

where τ_1 regulates the temporal stability of the function in the same manner as τ does in equation 2. McClelland (1979) introduced this mechanism in what he called Cascade networks. Also, related time averaging techniques are applied to input conductance build-up along dendrites by O'Reilly and Munakata (2000) in an effort to obtain biologically plausible activation dynamics.

In line with the activation functions used throughout the network we also apply a second time averaging function at perceptual pathway nodes. This is ostensibly equation 2, although for perceptual pathway nodes we feed the result of equation 4 into this definition, i.e.,

$$a_i(t+1) = \tau\, a_i(t) + (1-\tau)\sigma(NET_i(t+1)) \tag{5}$$

This double time averaging mechanism allows us to sustain the response of perceptual pathway nodes, with the time averaging of equation 5 building an extended trace upon the trace yielded by equation 4.

Emergent sustained perceptual trace

As evidence that the temporal dynamics of perceptual pathways that emerges from our definitions is appropriate we offer the activation profiles presented in this subsection. These profiles show how the level of activation of a perceptual pathway node fluctuates over time in response to stimuli being presented at perception layer nodes.

Perceptual pathway nodes have a baseline activation level of 0.5, which they stabilize at during preactivation of the network. This ensures that they are at a responsive part of their activation function and will consequently respond rapidly to net input change. This is consistent with findings by Rolls et al. (Rolls & Tovee, 1994; Rolls et al., 1999) that in order to explain the very rapid processing in visual areas, background activation levels place neurons close to their firing threshold (and hence they respond rapidly to input change).

First consider the unmasked prime profile shown in Figure 14a. This is a recording of the activation at one of the perceptual pathway nodes in response to one cycle of stimulation of the corresponding perception node and no other stimuli. In our model one simulation cycle corresponds to 16.666 ms of experiment time. Thus, this stimulus sequence is simulating presentation of a 16.666 ms prime on its own, in the absence of mask or target stimuli, which allows us to directly observe the temporal dynamics of the perceptual trace under ''normal'' conditions, i.e., in the absence of masking. What we observe is a rapid increase in activation over an approximately 30 ms period (two simulation cycles, i.e., 33.333 ms) followed by a slow decay back to baseline. We would argue that the trace has the temporal dynamics we seek—it responds rapidly to stimulus presentation, but then sustains that excitation subject to a slow decay. Furthermore, the profile has effectively returned to baseline between 18 and 20 cycles after the stimulus started having an effect, which corresponds to 300–333 ms. This is somewhat less than what would be suggested by the iconic memory literature but is broadly in line with the single cell recordings made by Rolls et al. (Rolls & Deco, 2002; Rolls & Tovee, 1994; Rolls et al., 1999). The shape of the activation profile that we obtain is also a reasonable approximation to the profiles presented in Rolls et al. (Rolls & Deco, 2002; Rolls & Tovee, 1994; Rolls et al., 1999).

Now consider the perceptual pathway profile that is obtained if we add a target stimulus. Thus, this corresponds to the ''empty'' compatible condition: A 16.666 ms prime followed by a 150 ms blank screen and then a 100 ms (compatible) target. The resulting trace is the unmasked profile shown in Figure 14b and is as one would expect: The initial segment is the same as the Figure 14a unmasked profile, but then presentation of the target builds upon the primed activation to generate a substantial activation increase. The reason that this increase is much larger than that for the prime being that the target stimulus lasts much longer than the prime. Once the target stimulus is removed representation of the percept slowly decays.

Now we consider the perceptual pathway profile that is obtained through masking. Figure 14a shows the profile for a masked prime (in the absence of a target stimulus). The first cycle of prime-induced activation is the same as was observed in the unmasked Figure 14a profile. However, feedforward inhibition induced by mask onset rapidly suppresses the trace and the perceptual pathway

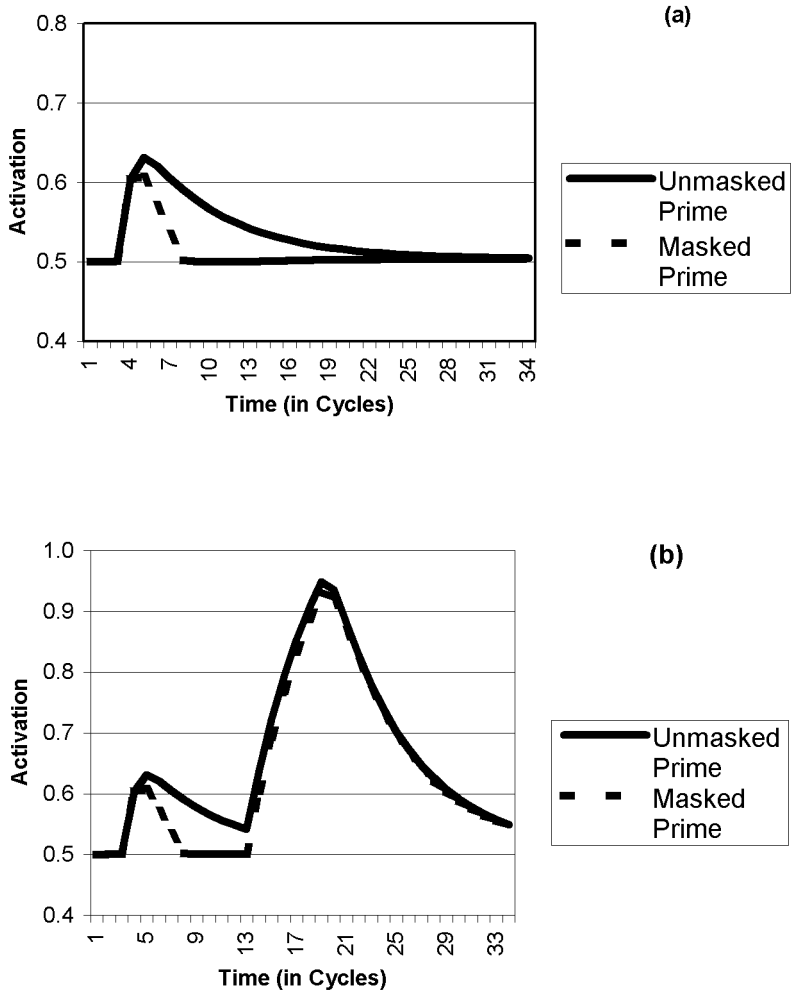

Figure 14. Activation trace from a perceptual pathway node as result of (a) an unmasked and masked prime, without a target; and (b) an unmasked and masked prime with a compatible target (note change of y-axis scale between figures).

node returns to baseline activation. Finally, Figure 14b shows the perceptual pathway profile for a masked prime followed by a target; this is the ISI 150 ms compatible condition. It is exactly as one would expect—a brief priming activation followed by mask-induced baseline activation and finally significant target-induced activation.

Response selection system

The response system comprises two classes of nodes—(ON) response nodes and (OFF) opponent nodes. We discuss these in turn.

Response nodes

The model provides a simple implementation of response competition between nodes with time averaged activation dynamics. As previously discussed, response competition is implemented by including lateral inhibition between R1 and R2.

Opponent nodes

Due to the rather specific role taken by OFF nodes their activation dynamics are somewhat different from other nodes; they have the following characteristics:

Thresholds. As previously emphasized, the purpose of the opponent networks is to suppress the activation leftover in response nodes when evidence for the response is removed. Consequently, it would be inappropriate for the opponent mechanism to inhibit response node activation *before* any activation from the perception layer has arrived. Note, we are not arguing that suppression only starts once a response has been fired, since this would contradict the basic findings of the masked priming task, i.e., that a subliminal prime can induce a separation reversal, while not inducing an overt response. It is though one thing to accept that inhibition can be initiated in the absence of an overt response, but it is quite another to think that inhibition is initiated even before excitation induced by perceptual stimuli has filtered through to the response system.

Thus, in order to prevent the opponent system cutting in too early, we have added nonzero thresholds on the entrance to OFF nodes. These are set to be above the level of activation of response nodes when their foregrounding has reached equilibrium. Consequently, OFF nodes remain at baseline activation levels (i.e., almost zero) throughout response circuit preactivation. The OFF node threshold can only be exceeded (and OFF nodes excited) when perception layer activation reaches the opponent system.

In fact this threshold is set somewhat above the level of response node preactivation. This is because the threshold also ensures that we obtain positive compatibility with perceptually weak primes, which was discussed earlier (cf. the ‘‘Prime visibility results’’ subsection on p. 397). Our interpretation of this effect is that a weak prime only excites its corresponding response node a small amount above its preactivation equilibrium, which does not yield sufficient OFF node excitation to cross its threshold. As a result, an inhibitory reversal is not initiated with weak primes and positive compatibility results ensue.

This threshold is built into the OFF node activation function by giving such nodes a large negative bias (the *bs* parameter in equation 3), which shifts their sigmoidal to the right and ensures that only high levels of net input are mapped to responsive ranges of the sigmoidal activation function.

Steep activation dynamics. Related to this choice of a high OFF node threshold is our use of a steep OFF node activation function (the *sp* parameter in equation 3). This ensures that once OFF nodes cross their threshold, they build up activation very rapidly (and also saturate rapidly). The reversal shape typically observed (see Figures 2 and 5a) justifies this choice. The reversal is sharp and deep and thus OFF nodes must build up activation very rapidly. The actual parameter settings for OFF nodes are presented in the Appendix.

Response-set maintenance and accumulators

Response-set maintenance

This mechanism has two components. First, it preactivates the response selection system according to task demands, by exciting R1, O1, R2, and O2 in order to ''foreground'' them from the set of possible response circuits. Note that even when fully foregrounded OFF nodes nonetheless stay at a low activation level during preactivation. This is because the threshold implementing negative bias on these nodes is so strong. Secondly, the layer maintains this delineation of response set, by continuing to feed activation into these nodes. In operational terms, the response system is preactivated by running the network for a number of cycles before the experimental sequence starts (i.e., before the prime is presented) with a pattern which, from amongst the input nodes, only excites the response-set maintenance node. The number of cycles is chosen in order to ensure that the preactivation has stabilized and reached an equilibrium level before the experimental sequence starts.

The second aspect to the response-set maintenance mechanism is that when a response is released, which occurs when the selection criterion is satisfied, the response-set maintenance node stops foregrounding the relevant pathway. Operationally, there is an inhibitory link from the accumulator mechanism to the response-set maintenance node, which is inactive preresponse selection. However, when the accumulator node crosses its (selection) threshold, inhibition is released and the response-set maintenance node becomes deactivated. The relevant equations are detailed in the Appendix.

Accumulators

The accumulator node operates according to the following equation:

$$a_{acc}(t+1) = \tau_{acc}\ a_{acc}(t) + (a_{R1}(t+1) - a_{R2}(t+1)) \qquad (6)$$

where $a_{acc}(t)$ denotes the activation of the accumulator node at time t, $a_{R1}(t+1)$ denotes the activation of the left response at time $t+1$, $a_{R2}(t+1)$ denotes the activation of the right response at time $t+1$, and τ_{acc} is the discounting term (which is set between zero and one). Thus at any time point the activation of the accumulator node is a sum of the current response separation and discounted activation of the accumulator on the previous cycle. a_{acc} is initially set to zero and remains at zero throughout response system preactivation, until priming induces response separation.

Interpretation of the model

A number of issues impinge upon how the model should be interpreted. First, it is important to reiterate that even though, in broad terms, activation build up on a particular response node (i.e., R1 or R2) suggests increased excitation of that direction, responses are not selected via constraints local to response nodes. The key criterion is the difference in activation between response nodes (i.e., response separation), which when plotted over time serves as a direct comparator for the LRP data. From this measure of response separation, an RT measure is determined via a threshold on accumulator nodes, which, as previously discussed, determines how separation accumulates over time or, in other words, the degree to which one response is (stably) dominating the other response.

We should also make clear the bounds of our model. Most significantly, unlike a large swath of connectionist research, there is no learning. This is because there is no empirical evidence concerning how learning affects the masked priming task discussed here or indeed whether it would at all.

In fact, Klapp and Hinkley (2002) report a strong effect of learning on the NCE: in their study, participants usually produced reliable NCEs only on the second day of testing. However, these experiments differed from all the others reported above in that the target duration was only 16 ms. Since no learning effect was ever observed with 100 ms targets, it seems likely that in the Klapp and Hinkley experiments learning was required for successful target discrimination, not for establishing direct perceptuomotor links or their inhibitory control.

Thus, in the model presented here link weights are hard-wired and fixed during our simulations. This is not to say that we have not explored the consequences of setting weights to different values. In fact, one element of our work has been to use the model to explore parameter settings that can reproduce the available empirical data and central parameters in this respect are the link weights. It is indeed important to note that the model is sensitive to parameter change. Thus, the parameter settings used are integral and central elements of our proposal and are documented in the Appendix.

Finally, the simulations presented in this paper correspond to mean empirical data values. For example, the separation profiles the model generates aim to reflect grand mean LRP values, i.e., averaged across the mean values of 10 or more participants per experiment. We have not attempted to generate a distribution of response times, although by applying Gaussian noise to the model, it would not be hard to do so. For example, we could add Gaussian noise to the response selection process in a manner similar to Cohen et al. (1990). However, we would only then average across the resulting distributions in order to compare with the central tendencies reflected in grand mean values. Thus, currently, we have avoided this added complexity.

Patterns and reaction times

As previously stated, activation is input into the network via activation patterns. These define the level of excitation of input (i.e., perception layer) nodes on each cycle. As an illustration, the pattern for a typical cycle of the presentation of a left pointing target would enforce the following activations:

PL1 – 0.0
PL2 – 1.0
PL3 – 0.0

i.e., the mask is unactivated, left-pointing double arrows are activated, and right-pointing double arrows are unactivated.

For the basic (ISI 100) experiment, the three pattern sequences are as follows:

1. *Compatible*. An activation of 1.0 is presented for one cycle at PL2 (this corresponds to the prime). Then six cycles of an activation level of 1.0 are presented at PL1 (this corresponds to the mask). This is followed by six cycles of activation at PL2 again, and finally a number of cycles in which all input activation is zero.
2. *Incompatible*. An activation of 1.0 is presented for one cycle at PL3. Then the sequence proceeds identically to the compatible case.
3. *Neutral*. An activation of 1.0 is presented for one cycle at PL1—the mask/neutral location. Following this neutral prime, the pattern sequence is the same as the compatible and incompatible cases.

Notice that these timings are in the same proportions as those used in the basic experiment, where the prime is presented for 16.666 ms and the mask and the target are presented for 100 ms each. One of our cycles corresponds to 16.666 ms of ''human'' experiment time. Because of this *one cycle = 16.666 ms* relationship, we can move between cycles and millisecond timings when relating model and human data.

In order to implement the response-set maintenance mechanisms we have discussed, 30 cycles in which just the response-set maintenance node is activated are prepended prior to these sequences, and the response-set maintenance node continues to be activated throughout. This prepended 30 cycles is sufficient to ensure equilibrium has been reached before the prime is presented.

We would like to obtain results from our simulations that correspond to the LRPs reproduced earlier. However, our graphs will start at some point to the right of the origin in Figures 2 and 5. This is because there is a lag between stimulus onset and the point at which this stimulus starts to affect activation in motor cortex. This lag, which we call the *onset delay*, arises because it takes some time for the sequence of processing performed in visual pathways to have completed before the motor system starts to be affected. This onset lag is not modelled in our implementation. However, from Figures 2 and 5, it can be estimated as somewhere around 180–200 ms.

We also wish to reproduce RT data. Obviously, our model does not yield this measure directly. Human RTs are a measure from stimulus onset to the registration of a button press, while our model only generates the separation profile between the point at which motor cortex activation starts to be affected by a stimulus and the point at which a response selection is indicated in motor cortex. We will call the time gap between these two points the separation delay. In order to compare it with RTs, we also need to factor in what we will call the residual delay. This includes both the onset delay and the execution delay (the time from when the response decision is made at motor cortex—assumed to be reflected in the final LRP minimum—to when a button press is registered). This residual delay should, broadly speaking, be fixed across conditions and experiments. Furthermore, for human performance, we can obtain a measure of residual delay in a particular experimental condition by subtracting the separation delay (which we can determine from the relevant LRP) from the overall RT. We can then obtain a central tendency for the residual delay by averaging across all conditions of all experiments. We have done this and included the condition and experiment specific separation and residual delays in the Appendix.

The result of the process is an overall central tendency for the residual delay of 263 ms. Using this, we extract RTs from our model by first measuring the target separation delay (which is directly extractable from our model). To this, we add the fixed residual delay to obtain an RT measure. Although this is rather a crude approach, it produces good results, as will become clear in the next subsection.

MODEL RESULTS

The main contribution of this section is to document how our model reproduces the masked priming effects outlined at the beginning of this paper (cf. the ''Introduction'' and ''Masked priming'' sections on p. 386 and p. 394, respectively). However, as a ''sanity check'' of our model, we begin by con-

sidering how the model would behave under ''normal'' processing circumstances and also we justify our argument that lateral inhibition alone could not produce the required effect.

Continuous target

It is important to ascertain that our model is well behaved in ''normal'' processing situations. Experimental work is often focused on rather specific phenomena and frequently, on situations in which the system behaves in a surprising or even counterintuitive manner. This is true of the masked priming experiments considered here, which encapsulate a sequence of stimuli that would rarely (if ever) occur in natural environments. It is possible to develop models that are too strongly constrained by such rare occurrences, which successfully reproduce the specific phenomena at hand, but which are not well behaved in more typical situations.

This section verifies that our model is indeed well behaved during such ''normal'' processing. That is, we demonstrate that when a single target stimulus is presented, i.e., without any priming or masking, the system builds up evidence at the corresponding response node, fires the response (as modelled by satisfaction of the selection criterion), and then resets itself. Such behaviour is shown in Figure 15, which presents the separation trace, i.e., difference between response node activations, that arises from continuously stimulating (with a value of 1) the right double arrows input neuron, i.e., neuron PL3. It is clear

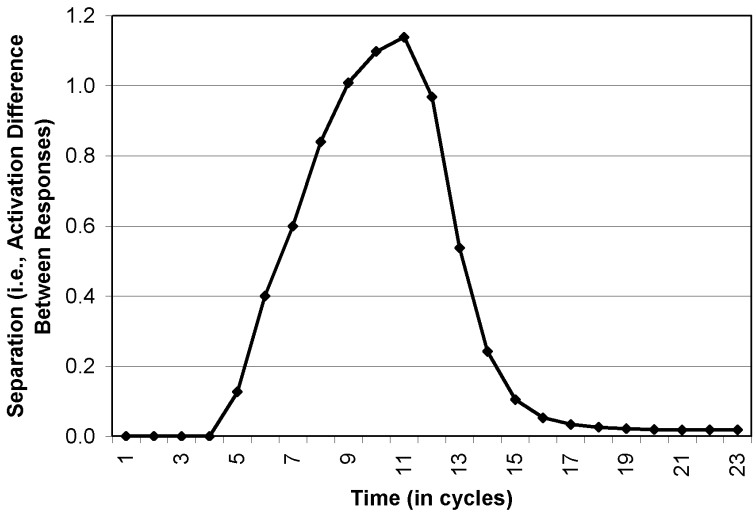

Figure 15. Separation of left and right response nodes, i.e., the activation difference between R1 and R2, plotted against time, in response to continuous stimulation of the network with a right arrows target.

from Figure 15 that the system is indeed well behaved in this situation: Continuous stimulation ensures that the required response has a sustained advantage in the winner-take-all race between response nodes. This advantage begins at cycle 5 and continuous until cycle 11, at which point sufficient evidence has accumulated for the right response to satisfy the selection criterion. Such satisfaction denotes release of the response. As discussed earlier (cf. ''The selection criterion'' subsection on p. 407, and the ''Response-set maintenance, response foregrounding, and selection'' subsection on p. 409), once the selection criterion is satisfied, response-set maintenance is withdrawn, removing the foregrounding of the two responses; as a result, response and off neurons tend over time towards an activation value of zero, which, in turn, causes separation between right and left responses to rapidly decline between cycles 12 and 18. By cycle 19 the system has affectively reset itself.

While this separation profile is completely to be expected, it serves as an important sanity check with regard to the inhibitory reversal. Once the right response becomes markedly activated above its baseline level, it will excite its corresponding OFF node, which will in turn feed inhibition back onto the response. If parameter settings were selected which made the inhibitory reversal stronger than the excitatory pressure being brought to bear, the response node could be prevented from becoming strongly active, which would in turn prevent sufficient evidence to accumulate to satisfy the selection criterion. Figure 15 demonstrates that this is not the case; it ensures that our off nodes are not feeding *too much* inhibition back onto response nodes. In addition, our later simulations will demonstrate that our off nodes do not feed *too little* inhibition back onto response nodes, since they will show that an inhibitory reversal does occur when evidence for a particular input stimulus is removed. Such removal of evidence for a stimulus is the key difference between the sequence of stimuli considered here and the sequence that occurs in masked priming.

Insufficiency of lateral inhibition alone

As a further sanity check, in this section we consider whether the inhibitory reversal central to the NCE could be obtained from a system that only employs lateral inhibition between response alternatives and that does not employ opponent processing. This will give further justification for our position that an opponent process is required in order to generate the NCE.

The results of our investigation are shown in Figure 16 (which considers a compatible target) and Figure 17 (which considers an incompatible target). These simulation results were obtained from networks in which both opponent processes had been cut out of the model and the strength of lateral inhibition was varied. The opponent process was removed by setting the weights of excitatory links from response nodes to OFF nodes, i.e., from R1 to O1 and R2 to O2, to zero. In addition, as a result of these changes, in order that response nodes still

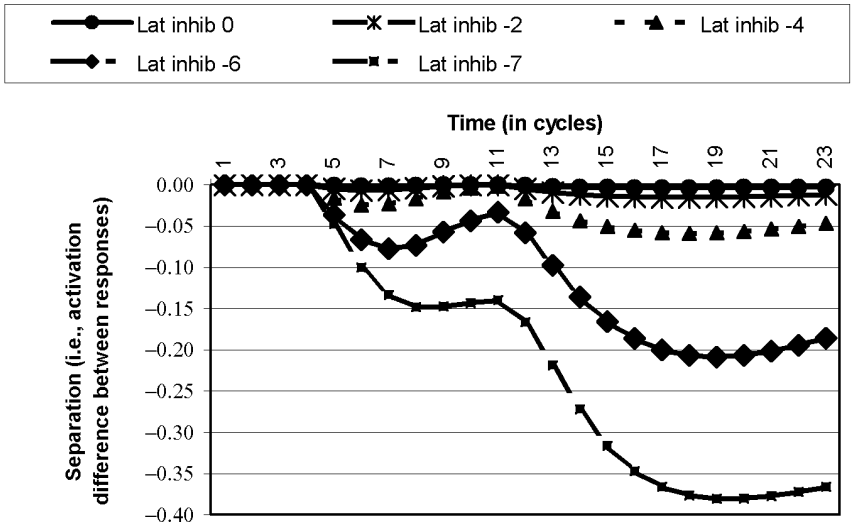

Figure 16. 16.666 ms prime, 100 ms mask, and 100 ms (compatible) target without opponent process.

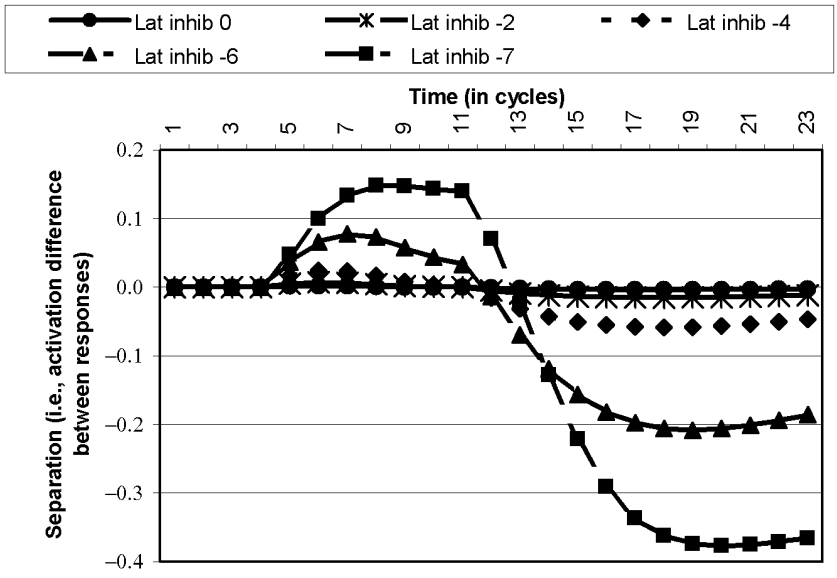

Figure 17. 16.666 ms prime, 100 ms mask, and 100 ms (incompatible) target without opponent process.

stabilize at 0.5 during network preactivation, response node biases had to be varied; the values used are documented in the Appendix.

In interpreting these simulations, the reader should note that cycle 5 (see, the x-axis) is the point at which the prime starts to have an effect on the response nodes, i.e., the start of prime-induced separation. In addition, the prime only feeds activation into the system for one cycle, as discussed earlier (cf. the ''Patterns and reaction times'' subsection on p. 421). Thus, the profile of separation between cycles 6 and 11 results from the prime being masked. Note also that cycle 12 is the point at which target-induced activation reaches the response nodes and causes a corresponding separation in the direction of the target, which, in the compatible case (Figure 16), generates separation in the same direction as the prime and, in the incompatible case (Figure 17), generates separation in the opposite direction to that resulting from the prime.

These simulations make clear that none of the lateral inhibition values explored generated the required inhibitory reversal. Specifically, with very low strength lateral inhibition (0 and –2) the prime did not even generate a significant prime-induced separation, which should commence on cycle 5. In contrast, with intermediate lateral inhibition values (–4 and –6) the prime did induce separation (see the separation commencing at cycle 5), but in this case, removal of sensory evidence initiated a decay of response node activation back to baseline, resulting in a corresponding decline of separation towards zero, which starts at (approximately) cycle 8 and continues through to cycle 11. In particular, there was no axis reversal pretarget onset, which should occur at some point between cycles 6 and 11 in order to be consistent with the human data; see the axis crossover between the black and white arrows in Figure 2. In addition, with high strength lateral inhibition (–7) separation continued in the same direction even after sensory evidence had been removed (between cycles 6 and 11). That is, the system is still striving to emphasize ''winner'' nodes (indeed this aspect of lateral inhibition between responses is at the heart of our modelling of the unmasked ''frame'' and ''empty'' experiments), but again removal of sensory evidence does not generate an axis reversal.

These simulations confirm our intuition that on its own lateral inhibition is insufficient to obtain the NCE. Critically, lateral inhibition can only work to emphasize an existing activation difference between competing nodes—however, removal of sensory evidence does not in itself have any excitatory affect on the competing (nonprimed) response. Thus, in the absence of opponent mechanisms, lateral inhibition between responses would not generate the separation sign reversal that is characteristic of NCE.

Basic experiment

With the model having successfully passed our two sanity checks, we now move on to consider how it responds to the basic masked priming sequence. With the pattern sequence discussed earlier (cf. the ''Patterns and reaction times'' sub-

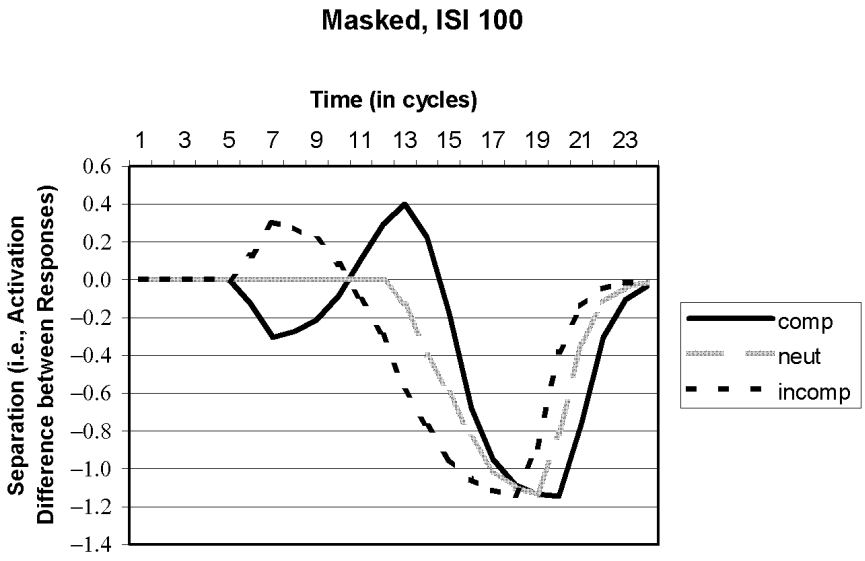

Figure 18. Model separation curves for the basic experiment.

section on p. 421), we obtain the separation curves shown in Figure 18. In each condition, the figure shows the time course of the difference between R1 and R2 activation. Note that we are not plotting activation of the accumulator node; however, the point at which the accumulator node crosses its threshold determines when a response is adjudged to have occurred, and hence (through deactivation of the response-set maintenance node) the final minimum in these profiles.

As with the LRP waveforms shown above, a downward-going deflection represents separation (i.e., the activation difference between the two responses) in the correct direction, and an upward-going deflection separation in the incorrect direction. The critical behaviour of the model in the basic experiment is as follows:

1. Response-set maintenance provides a stable preactivation of response circuits without separation—the two relevant response nodes are equally excited (only a few of these time steps are shown in the separation curves).
2. The prime excites the corresponding perceptual pathway (let's say PP2) and consequently pushes response node separation a small amount in a particular direction, i.e., R1 is active.
3. This in turn causes excitation of opponent nodes. However, only the primed response (R1) is generating enough activation on its ON to OFF node link to cross the OFF node threshold. Consequently, O1 becomes highly activated while O2 remains at baseline.

4. When the mask is presented it rapidly suppresses the prime (i.e., PP2).
5. The build-up of activation at O1 is relayed back to R1, but now in an inhibitory form. When PP2 becomes inactive the inhibitory pressure is no longer counteracted and R1 is strongly suppressed. Furthermore, once this suppression is strong enough to push R1 activation below R2 activation, lateral inhibition begins to work in the opposite direction; rapidly emphasizing the reversed direction of separation.

In the compatible case, this reversal yields (pretarget onset) separation in the incorrect direction, while in the incompatible case it is in the correct direction.

From Figure 18 it is clear that we have reproduced the basic masked priming effect, and that the separation profile has a good fit to the corresponding human data (see Figure 2). Compatible and incompatible, but not neutral, curves show initial separation in the direction of the prime. In the compatible case the following reversal (as inhibition from the OFF node cuts in) results in strong separation in the incorrect direction. In the incompatible case, it results in equally strong separation in the correct direction. If we consider the profile from the point of target-induced activation onset (cycle 13), the incompatible has already separated in the correct direction, followed by the neutral and then the compatible condition.

RTs extracted from the model (cf. the ''Patterns and reaction times'' subsection on p. 421) are also consistent with the human data (see Figure 19). In particular, in both human and model RTs, the incompatible condition is faster than the neutral condition, while the compatible condition is slower. This said, there is an important discrepancy between model and human RTs for this experiment. Model RTs do not quite show as large an effect as the human data and are on average a little faster (see Figure 19). This is an interesting issue since, as discussed previously (cf. the ''Masked and unmasked priming'' subsection on p. 397), a similar difference has been found between the basic experiment and the ISI 150 ''full mask'' experiment, with RT effect size in the former considerably larger than in the latter. We provide an explanation for this discrepancy later (cf. the ''Comparing ISI 100 and ISI 150 experiments'' subsection on p. 440).

Finally, it is also worth observing that for almost all the human–model RT comparisons we make, the model RTs will show a similar effect size to the human RTs. However, model RTs will also generally be a little faster than human RTs. It is of course possible that an alternative parameter setting exists, which resolves this problem. However, the main issue is that the model reproduces the same pattern of priming effect.

Low strength prime

As discussed earlier (cf. the ''NCE and self-inhibition'' subsection on p. 391), direction of compatibility effects in the masked prime paradigm depend on the perceptual strength of the primes, with very low strength primes yielding

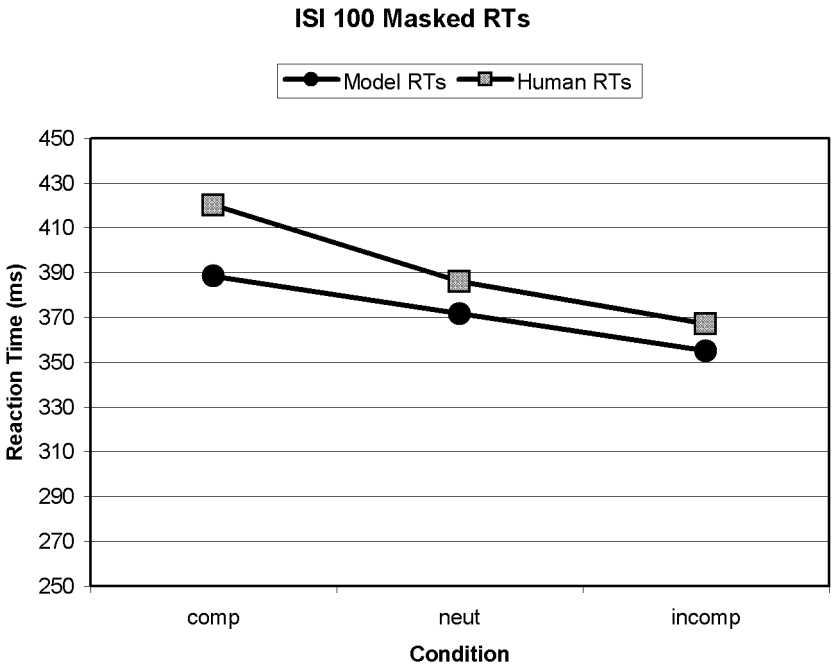

Figure 19. Model and human RT data for basic experiment (human RTs from Eimer & Schlaghecken, 1998, Exp. 1a).

positive compatibility effects (PCEs) rather than NCEs (cf. Schlaghecken & Eimer, 2002). This has been reproduced in the model. As shown in the separation profiles of Figure 20, if the activation strength of the prime is reduced sufficiently (a value of 0.095 is used in the simulations presented here), no inhibitory reversal is initiated with the basic experiment pattern sequence, and a PCE is obtained (Figure 21).

It is interesting to note that whenever we have reduced the strength of stimuli in our simulations (cf. also the reduced masking strength conditions), surprisingly small values have been required in order to yield the desired effect. This suggests that there is a nonlinear relationship between the strength of stimuli in the model and of stimuli in the physical world. Although not as yet explained, this aspect does not affect the interpretation of our model's data since the stimulus strength reduction is systematic throughout experiments and across stimulus types. It is interesting that in Francis' (1997, p. 578) metacontrast masking model, in order to generate the desired target luminance data he also had to use very small luminance values of simulated targets.

Notice also in Figure 21 that the fact that the neutral RT is the same as the incompatible RT in the model is artefactual: The difference between compatible and incompatible RTs is only one model cycle, so that without a more fine-

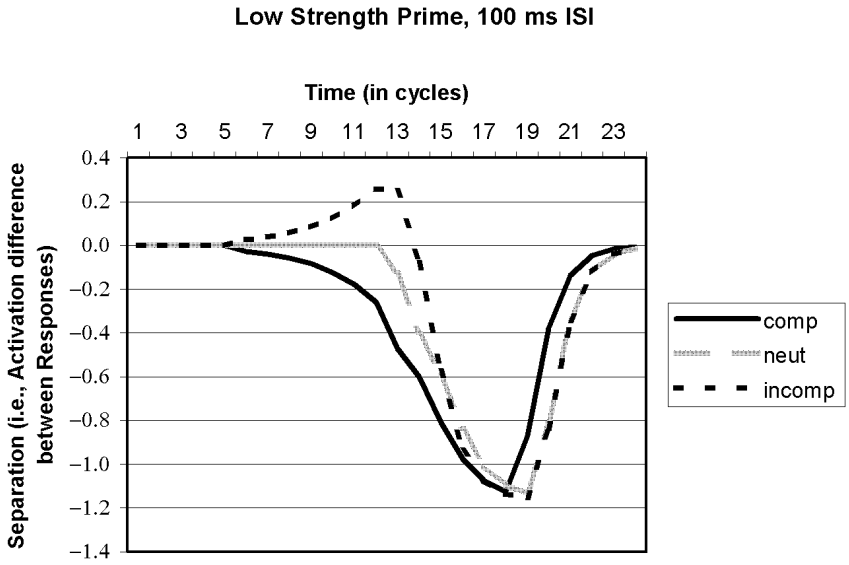

Figure 20. Model separation curves for low strength prime experiment.

grained discretization of time, the model is unable to distinguish further between these conditions. For the same reason, model RT effect size (16.67 ms) cannot come any closer to human RT effect size (7.5 ms). This effect reflects the strong negative bias in the sigmoidal activation function for OFF nodes. Thus, OFF nodes do not respond to low strength ON node activation.

It should also be noted that the separation profile in Figure 20 highlights the response competition in our model. The prime is so weak that it initially causes very little separation. However, between cycles 7 and 12 lateral inhibition between responses works on the small separation and accentuates it, such that when target presentation starts having an effect, it is at such a level that the compatible condition reaches its final minimum first. No LRPs for these experiments have been recorded. Thus, these separation profiles are a concrete (testable) prediction of the model.

ISI 150 experiments

Full mask

This experiment is the same as the basic experiment except that the prime–target ISI is somewhat longer. In reproducing this effect we present a sequence of patterns corresponding to the following series of stimuli:

1 cycle of the prime (16.666 ms)
6 cycles of the mask (approximately 100 ms)

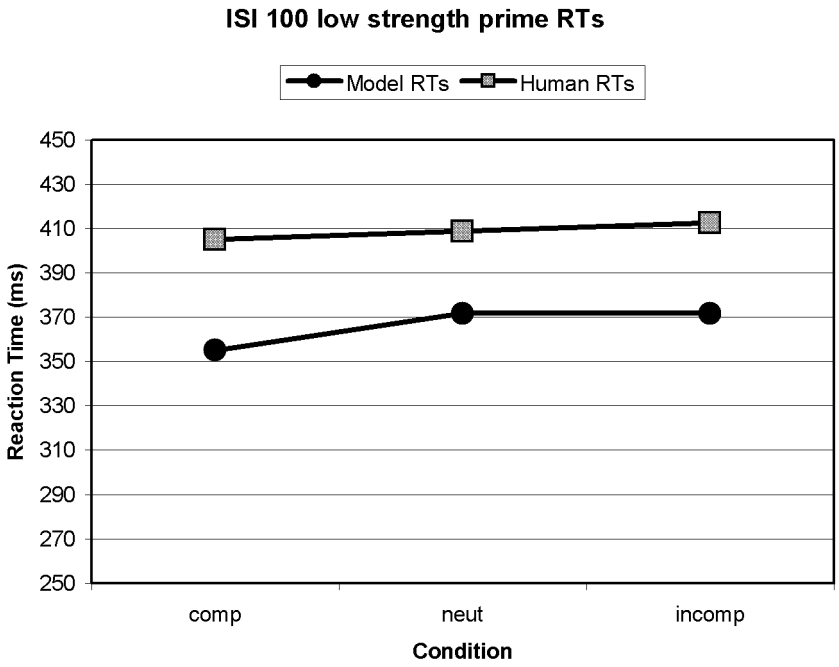

Figure 21. Model and human RT data for low strength prime experiment (human RTs from Schlaghecken & Eimer, 2002, Fig. 5; a neutral condition is not available for human RTs).

3 cycles of no stimulation (approximately 50 ms)
6 cycles of the target (approximately 100 ms)

The resulting separation profiles are shown in Figure 22. A comparison of these profiles with the full mask LRP in Figure 5a suggests that (broadly speaking) we have successfully reproduced the LRP. Model RT data (Figure 23) also reproduced the required NCE.

One interesting aspect is that, although it is not as pronounced as it is in the human data, our model shows the same double reversal around time points 14 and 15 (cf. the ''Masked and unmasked priming'' subsection on p. 397). That is, the separation curves turn back towards zero just before target-induced activation starts taking effect, as can be most clearly seen in the incompatible condition. Operationally this occurs when the nonprimed response has come to dominate the primed response (through lateral inhibition) to such an extent that the opponent node of the nonprimed response crosses threshold and starts to relay inhibition back on to it. Consequently, the nonprimed response begins to be suppressed and hence (through lateral inhibition) disinhibits the primed response. If this process was allowed to continue (by delaying target onset) the primed response should ultimately start to dominate and to win the response

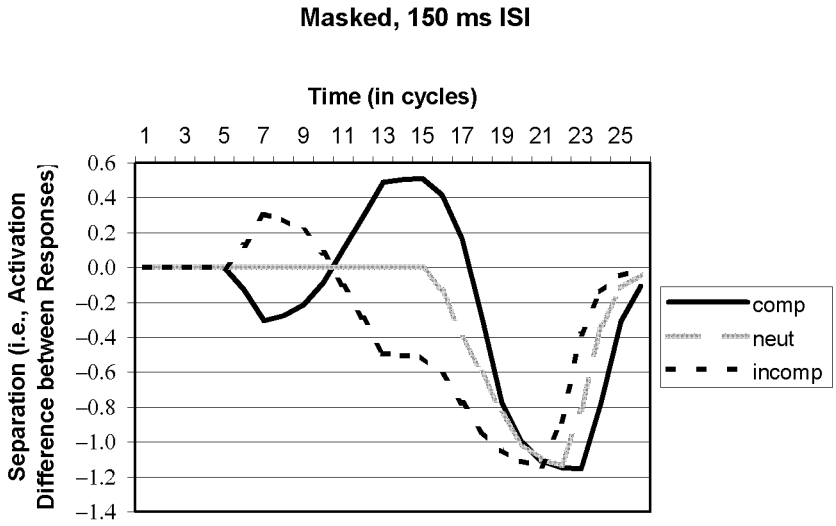

Figure 22. Model separation curves for full mask 150 ms ISI experiment.

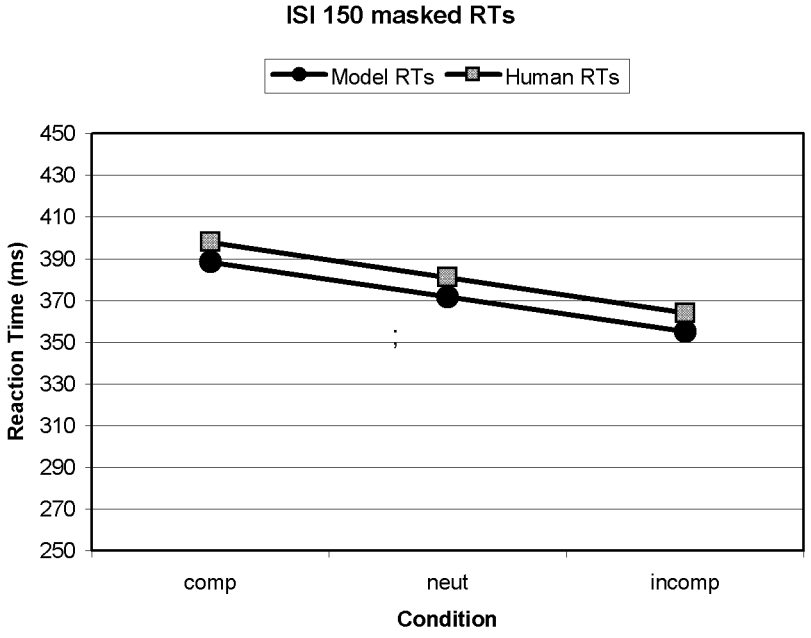

Figure 23. Human and model RTs for full mask 150 ms ISI experiment (a neutral condition is not available for human RTs).

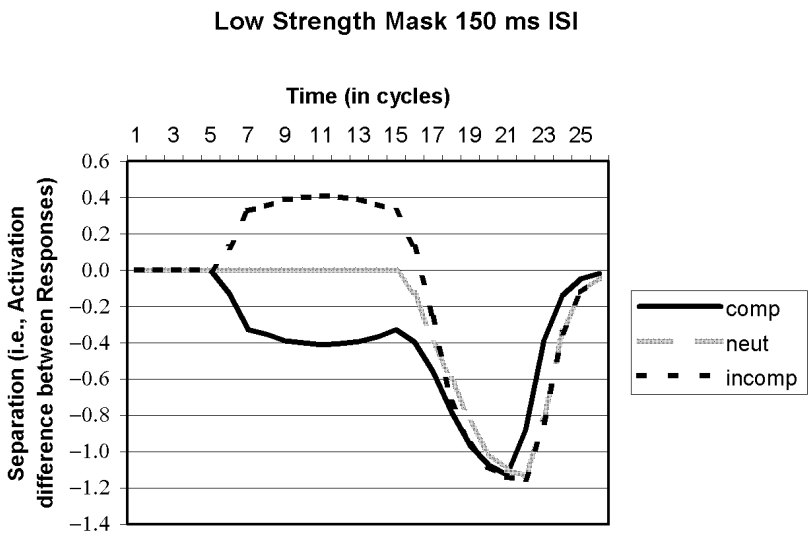

Figure 24. Model separation curves for partial mask, 150 ms ISI experiments.

competition race. Thus, a further prediction of the model is that if prime–target ISI is sufficiently long, NCEs will be nullified, since separation will have returned to zero by onset of target-induced activation. We discuss this prediction in further detail later (cf. the ''Simulation predictions'' subsection on p. 449).

Partial mask, ISI 150 experiment

In order to simulate this experiment, the same stimulus sequence as just described above was used, except that the mask stimulus had a lower activation level (0.152). The resulting separation curves and RTs are shown in Figures 24 and 25, respectively.

The separation profile is a compromise between unmasked priming (in which separation continues to increase until target activation takes effect, see Figure 28), and the ISI 150 ''full mask'' experiment (see above). That is, opponent inhibitory pressures are at work, but they are slower acting and less powerful than in the fully masked experiments. In particular, they are not strong enough to push separation back to zero and across the x-axis.

Obviously, however, the separation profile obtained in this simulation is not an exact match to the LRP profile of the partial mask experiment in Figure 5b. In particular, separation only shows a weak tendency to return to zero at the time of target-induced activation onset (cycle 15). A possible explanation for this is that in the experiment, the partial mask was randomly constructed on each trial. Consequently, it might have acted as a weak mask only on some trials

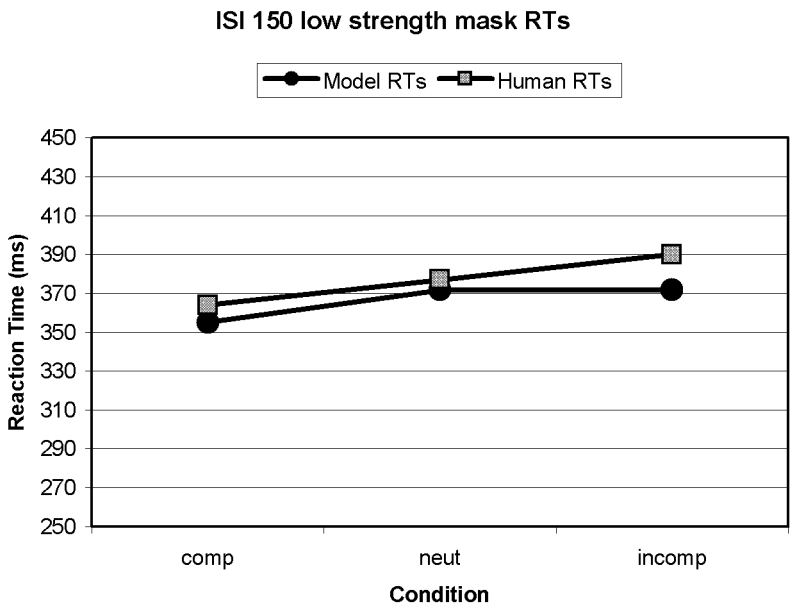

Figure 25. Human and model RT data for ISI 150 ms with partial mask (a neutral condition is not available for human RTs).

(corresponding to the situation simulated here). On other trials, however, it might have fully masked the prime (triggering complete separation reversal and behavioural NCE), and on yet another set of trials, it might not have acted as a mask at all (i.e., failing to trigger even a slight reversal, and resulting in behavioural PCE). There is in fact evidence for this from a reanalysis of the compatible trials in the partial mask experiment. Separate averaging of LRPs from slow response trials and fast response trials (relative to median RT) reveal that slow responses are accompanied by a complete reversal (identical to the one observed for fully masked compatible primes). No such reversal occurred for fast compatible trials. If the LRP depicted in Figure 5b results from such a mixture, then reproducing this profile with a single mask input activation level would be problematic.

However, a key aspect that has been reproduced here is that despite the fact that the incompatible separation crosses the compatible separation (around cycle 19), we still obtain PCEs. This is due to the discounted accumulation of separation evidence—even though at this time point the incompatible condition has separated further than the compatible condition, its final minimum is still farther away since it has to compensate for the legacy of previous incorrect direction separation.

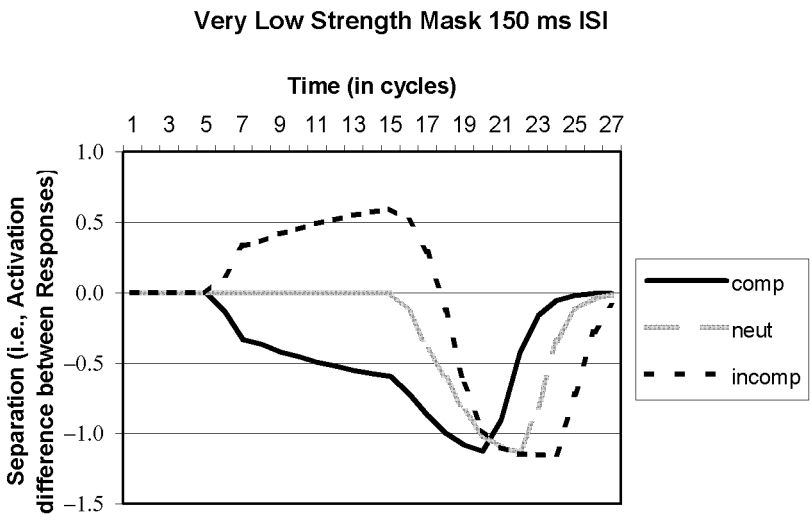

Figure 26. Model separation curves for 150 ms ISI and very low strength masking.

Unmasked ''frame'' experiment, ISI 150 ms

The data arising from the ''frame'' experiment (Figure 5c) was reproduced by employing the same stimulus sequence as above, but using a mask activation strength of only 0.01. The resulting profile is shown in Figure 26. With such low strength ''masking'', almost no inhibitory reversal is induced. More specifically, the perceptual trace of the prime is preserved throughout the masking period because the mask is too weak to suppress it, and this preserved activation counteracts the build-up of opponent inhibitory pressure. The resulting strong PCE (cf. Figure 27) is as expected.

Unmasked ISI 150 ms experiment

Finally, data arising from the ''empty'' experiment (Figure 5d), has been simulated by including a complete blank (i.e., zero activation) between prime and target. The effect (cf. Figures 28 and 29) is similar to that which is obtained in the ''frame'' condition, and does not require additional justification.

Short prime–target ISI conditions

As discussed at the beginning of this paper (cf. the ''Direct perceptuomotor links and response inhibition'' subsection on p. 386) and encapsulated in Figure 1, short prime–target ISIs (between 0 and 32 ms) yield PCEs, which turn into NCEs as prime–target ISIs increase beyond 64 ms. In order to reproduce this effect we need to make some additions to the model, this is because with short

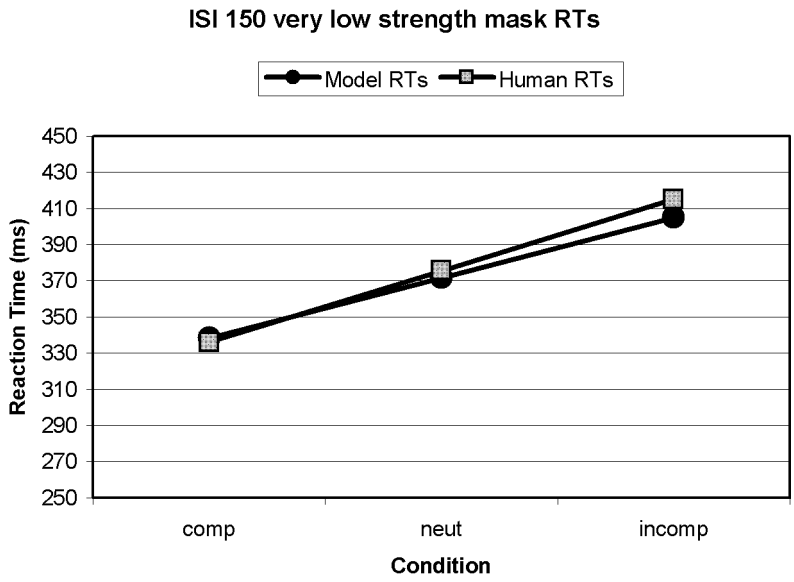

Figure 27. Human and model RTs for 150 ms ISI and very low strength mask (a neutral condition is not available for the human data).

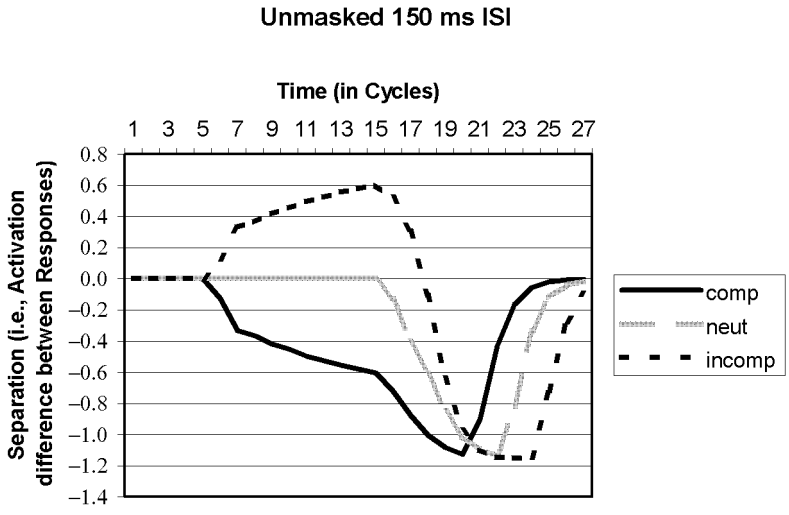

Figure 28. Model separation curves for 150 ms ISI and no masking.

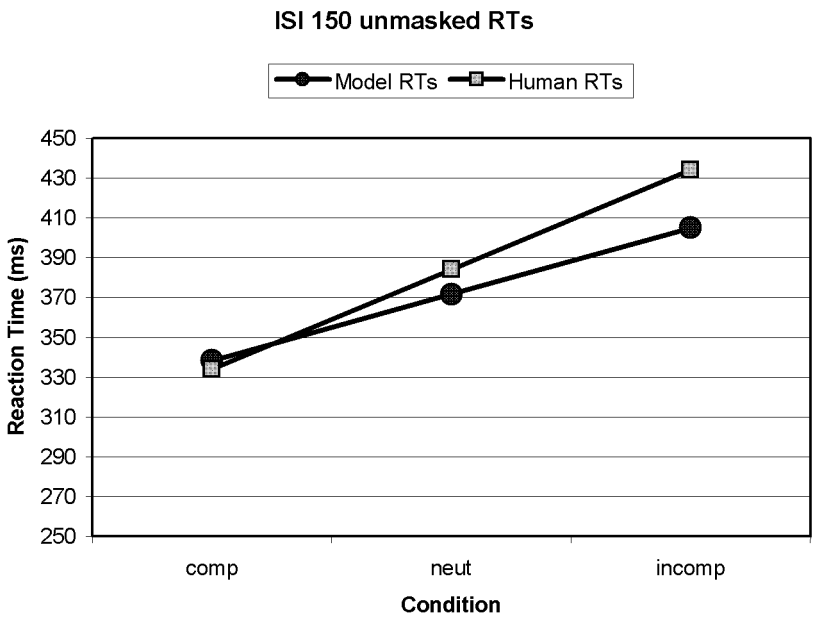

Figure 29. Human and model RTs for ISI 150 ms and no masking (a neutral condition is not available for the human data).

prime–target ISIs, the mask and target presentations overlap in time, which in the experiments that yielded Figure 1 was obtained by presenting the target in an offset spatial location. As a result, inhibitory forces change, that is, in the existing model, the feedforward inhibition between the perception and perceptual pathway layers ensures that each new input stimulus backward masks previous stimuli. In particular, not only does the mask backward mask the prime but in fact, the target has a backward masking effect on the mask. However, this will not occur if the target is spatially offset from the target. To reflect this change, we add a second input pathway, which comprises a separate set of perception and perceptual pathway nodes, as depicted in Figure 30. This new pathway is used to model stimuli that are spatially offset, i.e., the target in short prime–target ISI conditions.

The key point to note is that the PP′ nodes are not affected by feedforward inhibition from PL nodes. Rather, the only input that PP′ nodes receive is excitatory and from corresponding perception nodes. In addition, the two pathways join at response nodes, which they both have excitatory projections to. Parameter settings for the additional input pathway are documented in the appendix.

Figure 31 shows that at short prime–target ISIs of between 0 and 33.333 ms, the model indeed generates PCEs, which turns into NCEs as prime–target ISIs

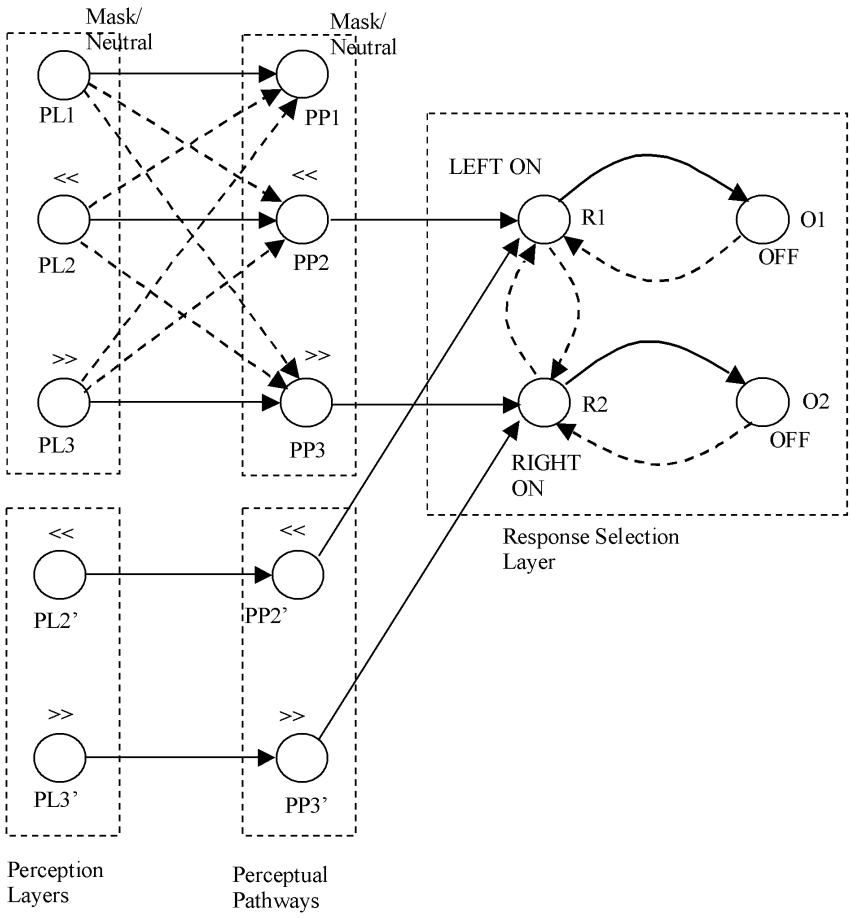

Figure 30. Short prime–target ISI model, with an additional input pathway, which is not subject to feedforward inhibition.

are increased to 66.666 ms and then 83.333 ms. Although, PCEs change into NCEs slightly earlier, in respect of prime–target ISI, in the model than they do in human data (see Figure 1), the key point is that this PCE to NCE transition is found. As a testable prediction of our model we present Figure 32, which shows the separation profiles that arise with a 16.666 ms prime–target ISI.

Forced choice conditions

As noted previously (cf. the ''Prime visibility results'' subsection on p. 397), forced choice identification tasks require participants to respond to the direction of the masked prime, which is normally presented without subsequent target.

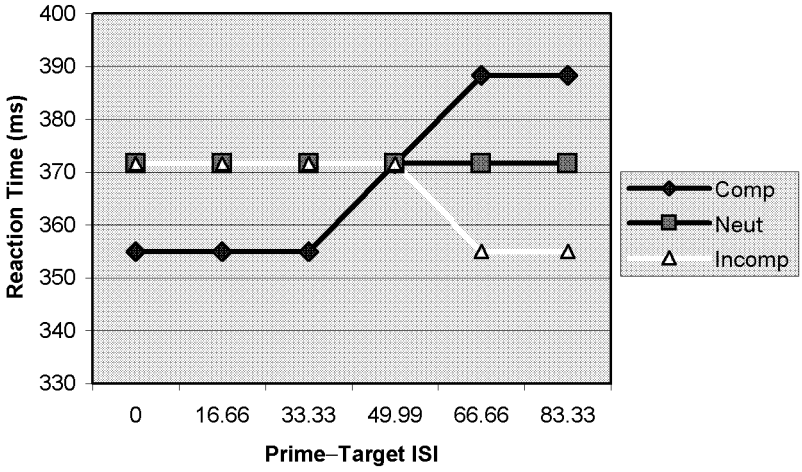

Figure 31. Reaction times across conditions as a function of (short) prime–target ISI.

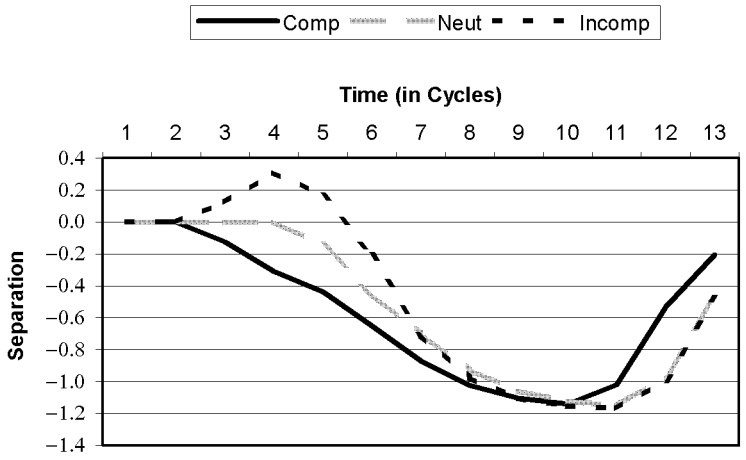

Figure 32. Separation profiles (i.e., response activation differences) across conditions, in which a 16.666 ms prime is followed by a 100 ms mask and a 100 ms target with a short prime–target ISI of 16.666 ms.

Earlier in the paper (cf. the ''Prime visibility results'' subsection on p. 397) we offered an explanation for why participants' performance is at chance level in these tasks, even though prime-induced activation does affect motor cortex activation levels. We argued that residual activation induced by the prime only has an effect on response outcomes if it is built upon by target activations.

The same explanation can be extrapolated to our model. Specifically, in order to reproduce these results we can run the model without target phase activation. As required, separation fails to reach a point at which the selection criterion is satisfied. We would interpret absence of selection as reflecting that the prime is not available to conscious experience and hence, that participants default to a random selection between response alternatives. The separation profiles that the model produces in this situation are shown in Figure 33, which serve as a further testable prediction from the model. Thus, there is significant response layer activation in this situation, including the expected inhibitory reversal, however, none of this activation is strong enough to yield response node separation sufficient to satisfy the selection criterion.

Comparing ISI 100 and ISI 150 experiments

As discussed earlier (cf. the ''Masked and unmasked priming'' subsection on p. 397), one issue that arises from the human data is that, in the basic experiment, RT effects were larger than in the ISI 150 experiments. One would have thought that, if it had any effect at all, a longer prime–target ISI would produce stronger

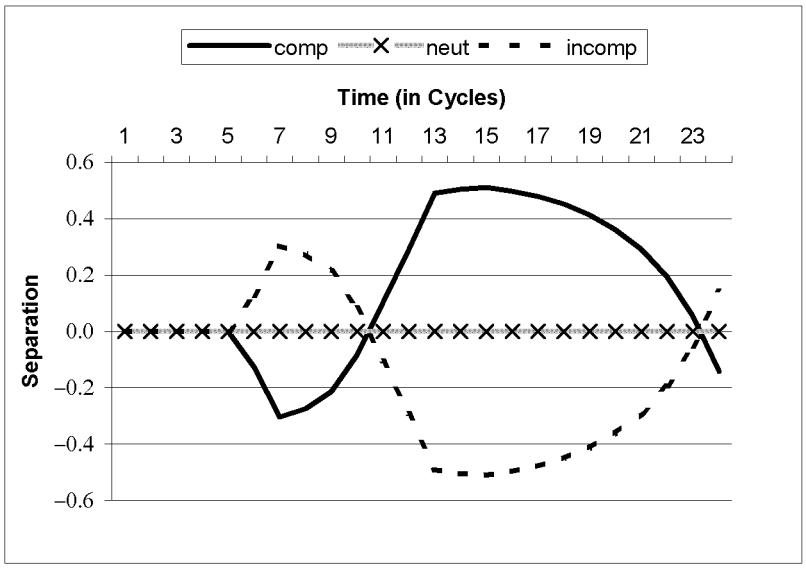

Figure 33. Separation profiles across conditions in the forced choice condition, in which a 16.666 ms prime is followed by a 100 ms mask and no target is presented.

opponent suppression and thus increase RT differences. We offer two explanations, both of which have been illuminated by the model construction: The double reversal and the arrows mask.

Double reversal

As discussed above, a double reversal (i.e., separation turning back to zero shortly before target activation takes effect) was obtained in both the simulation and the human LRP data of the ISI 150 ''full mask'' experiment. This suggests that the mask-induced reversal has a temporal limit. With longer prime–target ISIs, target-induced activations start when inhibitory processes have already declined, thus yielding less pronounced RT effects. Indeed, although RT effects size in the model do not increase between ISI 100 and ISI 150 ''full mask'' experiments, they do stay the same and with other parameter settings we can obtain a greater RT effects size for ISI 100 compared to ISI 150 ''full mask'' experiments. Furthermore, it is exactly this aspect that underlies our prediction that negative compatibility effects will disappear at long prime–target ISIs (see the ''Simulation predictions'' subsection on p. 449).

However, despite the successful reproduction of the double reversal, since (with the selected parameter settings) the model does not actually reproduce the reduced RT effects, we suggest that—although double reversal may play a role—other factors are also at work. As noted above (cf. the ''NCE and self-inhibition'' subsection on p. 391, and the ''Masked and unmasked priming'' subsection on p. 397), effects in the basic experiment might have been particularly large as a result of the specific masking stimulus employed. We will discuss this possibility in the following subsection.

Arrows mask

It has been argued that the use of overlapping right and left double arrows in the basic experiment mask effectively resulted in ''adding'' the opposite stimulus to the prime. This, in turn, might result in the perception of ''induced motion'', i.e., an illusionary movement ''away'' from the initial prime direction (Eimer & Schlaghecken, 1998). Conceivably, this may trigger corresponding response activation. In other words, the release of inhibition triggered by successful masking, and the response activation triggered by induced motion, would both facilitate motor response separation in opposite direction to the prime. Obviously, this would increase the size of the NCE relative to a situation where response separation is driven only by successful masking.

Although we cannot provide a full exploration of these possibilities within the context of our model, we can explore the consequences of a superimposed arrows mask. We investigated this by stimulating both left and right arrow perception nodes (PL2 and PL3) during masking. However, these nodes were stimulated asymmetrically with the nonprimed direction being more strongly excited than the primed direction, as a reflection of the proposed apparent

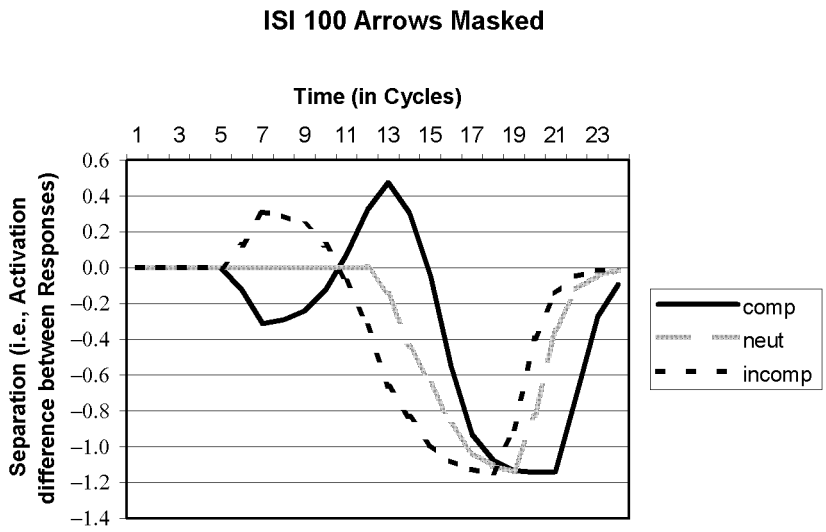

Figure 34. Separation curves with arrows mask.

motion towards the unprimed response. Thus, the masking part of the pattern sequence corresponds to stimulation of PL1 to a level of 1 and PL2 and PL3 to levels of 0.15 and 0.05 respectively, with the former value used for the non-primed stimulus (these values generate the largest increase in effect size; however, other values in this range also generate the desired effect). We still stimulate PL1 to reflect the fact that even the superimposed arrows mask does function as a mask and suppresses the perceptual trace of the prime.

Comparing the resulting separation profiles (shown in Figure 34) with those from the basic nonarrow-mask condition (Figure 18) reveals an increase in reversal size (cf. time point 13 in the compatible condition). This is in line with the explanation given above, and has indeed had the effect of increasing the resulting RT effect size (Figure 35). The finding that model RTs now closely match human RTs from the basic experiment provides further support for our model (see also the ‘‘General Evaluation’’ subsection on this page).

DISCUSSION

General evaluation

In support of our claim that the model accurately reproduces the available data, we offer Table 1 and Figure 36. The former of these documents how model and human RTs differ across experiment, both in terms of RT per condition and overall effect size. Then Figure 36 demonstrates that human and model RTs are very strongly correlated across experiment and condition.

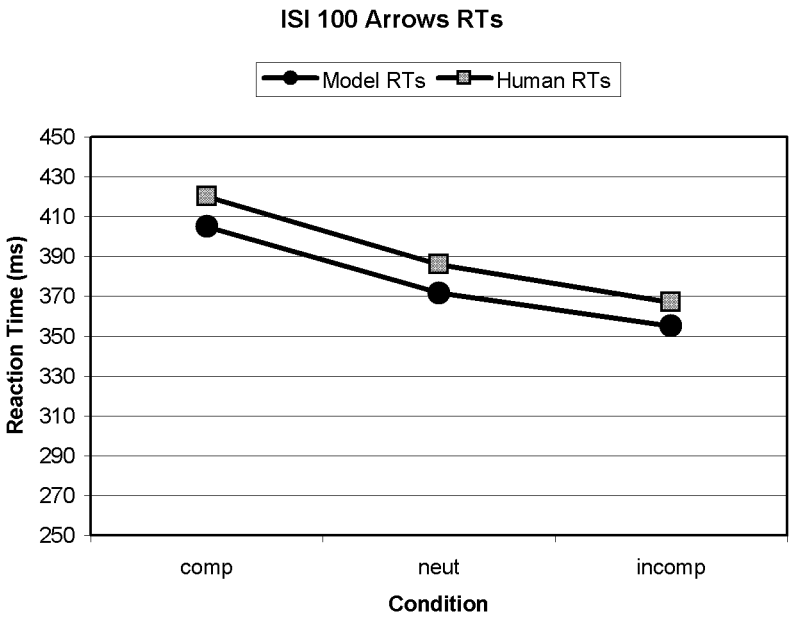

Figure 35. Human RTs from basic experiment compared with model RTs with arrows mask.

TABLE 1
Differences between model and human RTs

Condition	*Compatible*	*Incompatible*	*Neutral*	*Total*	*Effect*
Basic experiment (ISI 100)	–32	–12	–14		–20
Arrows mask (ISI 100)	–15	–12	–14		–3
Low strength prime (ISI 100)	–50	–40	n.a.		+9
Full mask (ISI 150)	–10	–9	n.a.		–1
Partial mask (ISI 150)	–9	–18	n.a.		–9
''Frame'' (ISI 150)	+2	–10	n.a.		–12
''Empty'' (ISI 150)	+4	–29	n.a.		–33
Average (ms)	–15.71	–18.57	–14.00	–16.09	–9.86
Average (cycles)	–0.94	–1.11	–0.840	–0.97	–0.59

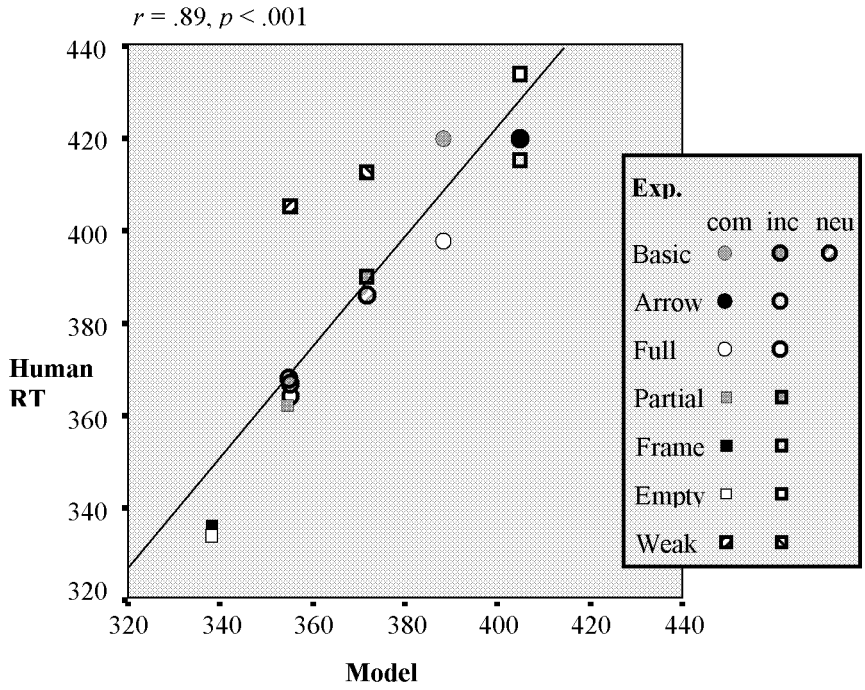

Figure 36. Correlation between human and model RTs.

Relationship to other models

Relationship to Houghton and Tipper's model

Although the work of Houghton and Tipper (1994) heavily influenced our prototype models (cf. Bowman, Aron, Eimer, & Schlaghecken, 2001b), the current model is significantly different. This reflects the different purposes of the two models. As previously stated, Houghton and Tipper's work was concerned with high-level attentional control processes. In contrast, our model is concerned with low-level motor control processes. In fact, response-set maintenance and switch-off is the closest we come to high-level control, providing the interface between voluntary, instruction-induced set-up of particular direct perceptuomotor links and involuntary automatic activation and inhibition processes controlling these links can be modulated by task demand (cf. Neumann & Klotz, 1994). The key difference here is that in Houghton and Tipper's work attention controls the direction of response selection, while our response-set maintenance mechanism excites both potential responses equally.

A further illustration of the difference in purpose of the two models is that here we use a simpler opponent network configuration. While, for example, in

Houghton and Tipper (1994) a gain system is used which associates both an ON and an OFF node with each property node, here we associate a single OFF node with each response node. Thus, in comparison with the Houghton and Tipper model, we have (in effect) conflated ON cells and property cells into a single response node (which is actually in line with simplifications made by Houghton in other papers—e.g., Houghton, 1994). The central reason for this simplification is that the Houghton and Tipper architecture reflected the need to have a repository for activation redirected from high-level ''target fields''. Obviously, such selective redistribution of activation is not a concern of our model.

A further difference between our model and that of Houghton and Tipper is that we have used a standard sigmoidal activation function. In contrast, Houghton and Tipper's function could accumulate negative activation. Elsewhere (Bowman, Aron, Eimer, & Schlaghecken, 2001a), we have clearly motivated why we have made this change to the function we inherited from Houghton and Tipper, and the change has enabled us to reproduce the LRP data more faithfully. Furthermore, it could certainly be argued that we have moved to a more standard (and biologically accepted) activation function—probably the most common activation function to be found in the literature is a sigmoidal function with a minimum activation of zero.

Relationship to Grossberg and Mingolla's boundary contour system

Perhaps the most complete computational model of masking available is that which has arisen from Grossberg and Mingolla's (1985a, 1985b) boundary contour system (BCS). Francis (1997) has demonstrated how a large spectrum of the available data on metacontrast masking can be reproduced using BCS. The BCS model provides a more detailed explanation of masking than our model does. In particular, a detailed treatment of spatial aspects of stimulus presentation is provided. However, although our model is much less detailed, there are a number of points of similarity in respect of coarse grain mechanisms employed. For example, as previously mentioned, the use of feedforward inhibition is at the heart of masking in both models. In fact, we would argue that our model could be viewed as a high-level abstraction of the more detailed BCS model. As an illustration of this, note how our use of double time averaging in the context of our coarse-grain localist representations produces an emergent after image, and how this is obtained in a more detailed manner through resonating feedback excitation in BCS.

An important avenue for further research would be to replace the perception layers of the present model with a more detailed BCS style model of the perceptual end of the system. This would have the added benefit of investigating how the style of backward pattern masking considered in this paper can be

reproduced in the BCS model, which has to date only been used to explain metacontrast masking effects.

Earlier versions of the present model

A preliminary version of our model was described in Bowman et al. (2001a, 2001b). This earlier model reproduced the separation reversal, which underlies the basic experiment. However, it failed to reproduce the spectrum of data (both separation profiles and RTs) that we have done here, especially the ISI 150 experiments. There are a number of aspects of the model presented here that differ from the preliminary model; we list these largely without discussion, since the inclusion of each has been justified elsewhere in the paper.

1. Bowman et al. (2001a, 2001b) used a gating mechanism to implement masking. This could be viewed as an abstraction of an inhibitory mechanism and the feedforward inhibition we have employed here is essentially a refinement of the earlier approach.
2. The response-set maintenance mechanism we employ here has evolved from the response-set maintenance used in Bowman et al. (2001a, 2001b). The present mechanism preactivates and maintains excitation of both ON and OFF nodes in the relevant response channels, while in previous models it was solely focused on response (ON) nodes. In addition, the discounted accumulator node mechanism, which is tied to response-set maintenance, and the manner in which response channels are ''backgrounded'' when a response is released, are both new to this model.
3. As just discussed, we employ a more standard, positively ranged, sigmoidal activation function than that which we inherited from Houghton and Tipper and used in (Bowman et al., 2001a, 2001b).
4. The treatment of perceptual pathway activation dynamics and sustained traces is new to the model presented here.
5. The rather crude handling of opponent circuit thresholds of (Bowman et al., 2001a, 2001b) has largely been refined here by integrating them into the OFF node activation function.

Also, although unimplemented, the model discussed in Schlaghecken and Eimer (2002) predicted a number of aspects of the model described in this paper. In particular, although not couched in the same terminology, a form of opponent process was at the heart of the model described in Schlaghecken and Eimer (2002).

Relationship to physiology

First, we must emphasize that, in its current form, our model is best viewed as providing a ''cognitive-level'' explanation of masked priming. In particular, the relationship between our model and neurophysiology is not completely clear.

Nonetheless it is valuable to speculate on how the mechanisms we have employed could relate to brain structures, thus providing a springboard to the development of a more biologically plausible model of the masked priming data. Such a consideration is made here.

We divide our discussion into two main issues:

- biological justifications for the micro elements of our model, in particular, neuron activation functions and activation transfer across links
- how the macro elements of our model, i.e., layers and neural circuits, relate to anatomical structures that could be involved in the masked priming effect.

Micro comparison

Perhaps the key element of the microlevel of our model is the activation functions employed. Central aspects of these functions are (a) time averaging and (b) the sigmoidal function, which were introduced earlier (cf. the ''Basic activation functions'' subsection on p. 413).

Time averaging. Time averaging is built into the definition of $a_i(t+1)$ with τ playing a key role in regulating the responsiveness of the function to fluctuations in net input. Time averaging can biologically plausibly be explained as simulating the gradual build-up of postsynaptic generator potentials through temporal summation, and is frequently used in the computational modelling of brain functions. For example, in O'Reilly and Munakata (2000) the membrane potential at a time $t+1$ is expressed as a function of the potential at time t and the newly arriving electrical charge in that time interval. In addition, in order to reflect the ''sluggish'' propagation and aggregation of synaptic inputs, in O'Reilly and Munakata time averaging is used to express the excitatory input conductance to a neuron. This is similar to the double time averaging we use in perceptual pathway nodes.

The sigmoidal function. The sigmoidal function (see Figure 13) that we use is, as has been argued many times (e.g., O'Reilly & Munakata, 2000), in broad terms biologically justifiable. Node activation levels represent the firing rate of neuron populations (not individual neurons). This is a standard approach that does not in itself preclude biological plausibility. In particular, although there is some debate surrounding this issue (see Abeles, Bergman, Margalis, & Vaadia, 1993), it has been argued that because of the high noise level of individual neurons' firing rates, the computationally most meaningful level of neural activity is the firing rate of neuron populations (O'Reilly & Munakata, 2000).

If we consider the depiction of the function in Figure 13, in standard manner our function saturates at high positive net inputs. This is justified by neuronal refractory periods, which ensure that there is an upper bound on the firing rate of neurons. In addition, the function contains a graded threshold at low levels of net

input. This is in line with the known threshold mechanisms in biological neurons, which will produce action potentials only if their integrated input is sufficient to produce a generator potential that exceeds threshold values.

Connectivity. At the microlevel of comparison we can also identify biological mechanisms that support activation transfer in the manner we have been considering here. First, links in our model are broadly related to coarse biological interconnections such as large-scale neural pathways or projections between cortical areas. Consequently, link weights denote accumulated synaptic efficacy across a set of individual connections.

Furthermore, the distinction between excitatory and inhibitory connections mimics an identical distinction in the nervous system, which can be found at all levels from individual synapses to types of neurons to neurochemical pathways. With regard to the opponent processing mechanism it is important to note that our OFF node is no more than a set of inhibitory interneurons. Neural populations that when receiving excitatory input have an inhibitory effect back onto the same neurons are commonplace in the brain. Furthermore, such structures have increasingly been employed in the move towards biologically plausible neural network modelling. For example, in O'Reilly and Munakata (2000), layers of inhibitory interneurons are interconnected with hidden layers. This is done in order to regulate the level of excitation in hidden layers, since the presence of (excitatory) recurrent collaterals and autoassociator (attractor) dynamics can yield unconstrained excitatory self-stimulation and unstable settling dynamics. The inclusion of inhibitory interneurons stabilizes this process. Consistent with this perspective we believe that each localist ON or OFF node representation that we employ in our model would be implemented in the brain as a neural population, which, in the former case, is composed of excitatory neurons, and, in the latter, inhibitory interneurons. With, in the same manner as in our model, excitation of the ON population counteracting the inhibitory pressures from the OFF interneuron population, until excitation of the ON population is removed and the inhibitory reversal is initiated.

There is though one potential glitch with such a population encoding, namely, the degree to which the response encoding is distributed. In a fully distributed encoding, each response would be encoded across *all* the available neurons. However, response neurons need to be suppressed by directing inhibitory pressure at them through link interconnections. In a fully distributed situation, directing inhibitory pressure onto the appropriate ON neurons is not straightforward, as the inhibitory pressure would act upon *all* the available neurons and thus would also suppress competing responses. Consequently, in order that inhibitory interneurons can be appropriately interconnected and their action localized, it is important that encodings are sparse (even localist in the sense of Page, 2000) and that the representation of competing responses is (functionally) nonoverlapping.

Macro comparison

At this level of comparison, biological plausibility is less clear. The question we wish to address—beyond simply saying that single nodes represent populations of neurons—is how components of our model relate to neuroanatomical structures. Although we can make ''broad brush'' statements such as that our perception layer is playing the role of early visual areas, the more interesting and difficult question to answer is where the opponent circuits reside in the brain.

While it is clear that response execution is ultimately triggered in the primary motor cortex, it is also clear that response selection occurs at an earlier stage. Although the precise mechanisms are as yet unknown, it is generally assumed that the basal ganglia play a crucial role in this process (Band & van Boxtel, 1999). These nuclei possess extensive projections to motor cortex and are known to be heavily involved in controlling motor responses (cf. Parkinson's and Huntington's diseases, which are pathologies of the basal ganglia). Furthermore, the basal ganglia are known to have an inhibitory action. For example, GABAergic neurons in the Globus Pallidus and Substantia Nigra Pars Reticulata are tonically active and hold the thalamus (amongst other structures) in a state of tonic inhibition (Redgrave, Prescott, & Gurney, 1999). Furthermore, imaging and neuropsychological work on the masked priming task has implicated corticostriatothalamic circuits in the inhibitory reversal (Aron et al., 2003) and previous work by Jackson and Houghton (1994) has related opponent network processing to basal ganglia architecture. In addition, there is evidence that lateral inhibition arises in the striatum and subsequent basal ganglia areas (such as the globus pallidus) (Rolls & Treves, 1998). Furthermore, it has been postulated that response competition is implemented by such neural circuitry (Rolls & Treves, 1998), which would be consistent with the view that some analogue of our response layer resides in the basal ganglia.

Simulation predictions

There are many minor predictions implicit in our model, e.g., the neutral condition RTs for a number of experiments in which human RT measures have not been made. However, we would particularly emphasize four predictions. The first three are separation profiles for conditions where LRPs have not yet been recorded: (1) The short prime–target ISI shown in Figure 32; (2) the separation profiles that we derived for the low strength prime (cf. Figure 20), a notable feature of which is that separation builds up slowly but in a sustained fashion throughout the masking period; and (3) the forced-choice profiles shown in Figure 33. All of these are concrete testable predictions.

Our fourth main prediction is that as prime–target ISIs increase (from 100 ms, through 150 ms, and beyond) the NCE will decrease and eventually disappear. We have run a number of simulations to provide a prediction of the mask length at which this happens. In the simulations, the following sequence of events is

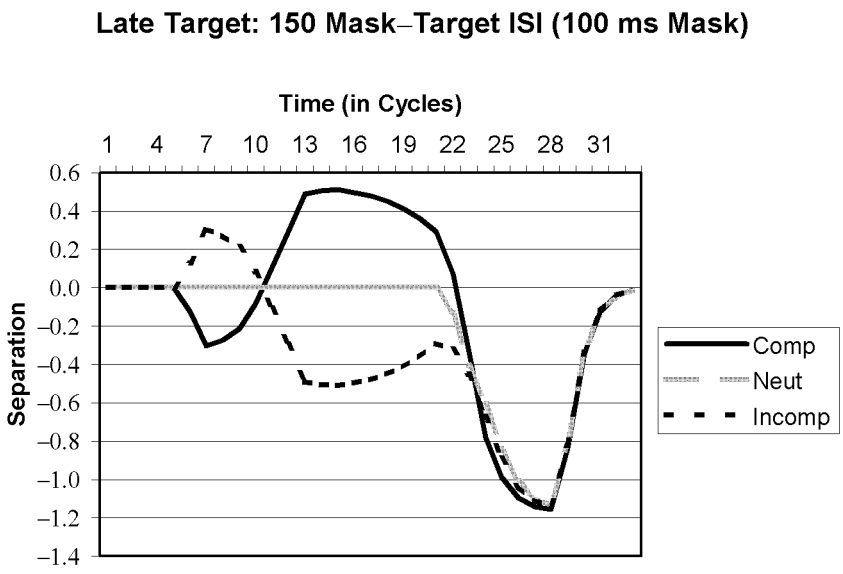

Figure 37. Separation profiles for the sequence: 16.666 ms prime, 100 ms mask, 150 ms blank, 100 ms target.

presented to the network—1 cycle (16.666 ms) prime; 6 cycles (100 ms) mask; 3 cycles (50 ms) blank; *N* cycles (16.666 ms * N) blank; 6 cycles (100 ms) target. We are interested in how the separation profiles and resulting RTs change as *N* varies. We have run simulations for values of *N* between 0 and 10. Figures 37 and 38 are notable separation profiles from amongst these simulations. When *N* = 6 the NCE has disappeared (Figure 37), as a reflection of which the final minima in the three separation profiles have closed up and are almost overlaid upon one another. In contrast, when *N* = 8, we obtain the profiles shown in Figure 38, in which the double reversal is so strong that the compatible condition separates towards its final minimum first and we obtain positive compatibility.

To summarize these changing RT effects, Figure 39 plots RTs for each of the three conditions against mask–target ISI, i.e., 50 + (N*16.666) ms. The switch from negative compatibility (*N* = 0,1,2,3,4,5); to null compatibility (*N* = 6); to positive compatibility (*N* = 7,8,9,10) is made clear. The separation profiles of Figure 37 and Figure 38 and the resulting change in RT effect depicted in Figure 39 are a concrete testable prediction of our model.

With regard to this prediction it is also important to note that the kind of phasic interleaving of facilitation and inhibition that causes first PCEs to turn into NCEs at short prime–target ISIs and then to turn back into PCEs at long prime/mask–target ISIs is characteristic of an opponent process, such as the one we have employed here. In our model, the time course of these phasic changes is

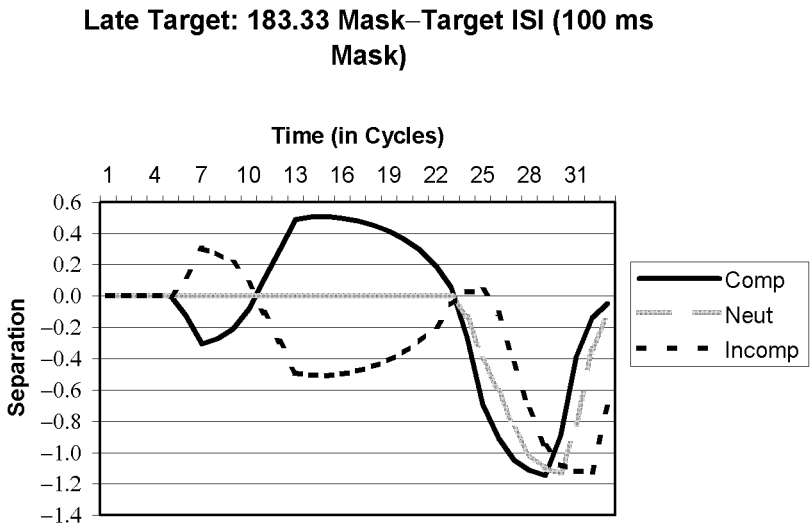

Figure 38. Separation profiles for 16.666 ms prime, 100 ms mask, 183.333 ms blank, 100 ms target.

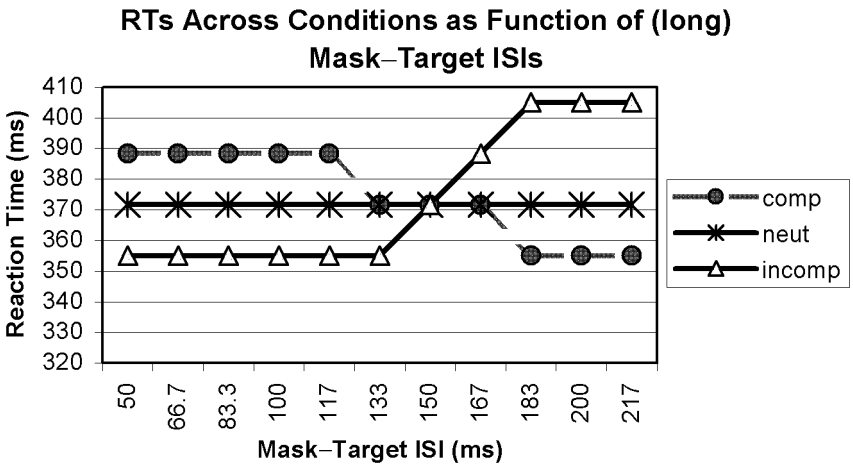

Figure 39. RT effect size as mask–target blank period is increased (50 corresponds to $N = 0$, 66.666 corresponds to $N = 1$, 83.333 to $N = 2$, etc.).

complicated by the additional inclusion of lateral inhibition between response nodes. A typical sequence of states that would arise with our response system is a follows. First, if say the LEFT response receives bottom-up stimulation it will initially win the response race (due to the lateral inhibition). However, the LEFT off node will also become strongly excited and will suppress the response if its counteracting bottom-up excitation is removed. Assuming no further bottom-up stimulation, the system then enters a phasic pattern of activation dynamics, whereby the previously winning response (LEFT) becomes strongly suppressed, releasing the competitor response (RIGHT), which then starts to use lateral inhibition to its advantage; however, as it becomes more active, its off node will eventually cut in and suppress it below the LEFT response, which again begins to win the race and the cycle repeats itself from this point. Thus, confirmation of this prediction that PCEs change into NCEs and then back to PCEs as prime–target ISIs increase will give strong support for some form of opponent system, since the cardinal marker of such a mechanism would have been identified.

Theoretical justification

Although at first sight the self-inhibition mechanism that we have modelled may seem a little odd, it is in fact consistent with, and indeed a natural consequence of, the well accepted ''direct perceptuomotor link'' hypothesis (Neumann & Klotz, 1994). Specifically, it can certainly be argued that responding to a stimulus, which has not (yet) been fully consciously perceived, could bring evolutionary benefits; for example, in respect of responding rapidly to impending threat. However, it is clear that dangers could also arise from such a necessarily ''quick and dirty'' and thus, error prone, low-level system. Therefore it is hypothesized that an emergency brake mechanism is needed, which ensures that we are not slaves to these nonconscious activations. In particular, if the sensory evidence supporting such an action is suddenly removed, and the activation has already come dangerously close to overt response execution, then the system switches into self-inhibition mode. Furthermore, as previously argued, the opponent process mechanism that we have presented here has exactly the necessary characteristics to generate such self-inhibition and is likely to be widely used in the nervous system for the control of ongoing activity (see for example, Houghton, 1994; Houghton & Tipper, 1994; Houghton et al., 1996).

The slightly surprising aspect of this emergency braking mechanism is that the sequence of events initiated by retracting a response (through self-inhibition) causes the opposite response to become activated above baseline levels. However, this is a natural consequence of the two responses being connected via lateral inhibition: If one of the responses becomes strongly suppressed below baseline, the other will be disinhibited and will thus be pushed above baseline. It is also interesting to note that the degree to which nonprimed responses benefit from suppression of a primed response is dependent upon the number of

competing responses. Specifically, as the number of mutually inhibitory responses is increased, the degree to which nonprimed responses benefit from suppression of a primed response decreases. This is because the inhibitory and disinhibitory forces acting on a particular response are progressively watered down as the number of response alternatives increases. This raises the possibility that significant disinhibition of a competing response could be an artefact of there only being two response alternatives at play in the masked priming experiment. In contrast, in ''natural'' real world situations, there would always be a huge number of competing responses.

To explore this issue, we have run a set of (as yet) unpublished experiments (Schlaghecken, Bowman, & Eimer, 2005) and the findings are indeed consistent with the view that increasing the number of response alternatives reduces the extent to which a (retracted) primed response disinhibits competing responses.

Finally with respect to the theoretical justification for our model, there has recently been debate in the masked priming literature on whether NCEs reflect self-inhibition or sensory/perceptual interactions between primes and masks (Lleras & Enns, 2004). Specifically, Lleras and Enns have used the object substitution masking framework to suggest that visual features contained in the mask could be exciting the nonprimed stimulus and hence the nonprimed response. However, Klapp (in press) has completed a series of experiments that demonstrate that the masked priming NCE is not based on perceptual interactions of the type suggested by Lleras and Enns. That said, we do accept that the nature of the mask is critical to this issue and indeed that with a mask that contains arrows as elements, there might in fact be two mechanisms at play: One of motor activation and inhibition and one of perceptual interaction. We provided simulations to this effect and associated discussion earlier (cf. the ''Arrows mask'' subsection on p. 441).

CONCLUSIONS

We have given further evidence for an inhibition-based account of the NCE in masked priming. We have done this by presenting a neural network model of inhibitory processes, the behaviour of which is consistent with the available data. The underlying mechanisms we use are related to those employed in Houghton and Tipper's (1994) influential models of negative priming and inhibition of return. Although, as stated previously, the notable difference between Houghton and Tipper's model and ours is the role that high-level attention plays. Our network has the right flavour to reflect a low-level visuomotor link. In this respect the mechanisms used could be described as ''dumb''. The only role that goal driven processes play is in response-set maintenance. However, this is not a selective process and its effect is fixed. Our model proposes a low-level inhibitory mechanism, the initiation of which is controlled by fluctuating excitatory pressure on response nodes.

It would be wrong of us to claim that what we have presented ensures that our model is the only one that could reproduce the masked priming data. Indeed, such assurances are never available in modelling work. However, we would argue that the mechanisms employed are strong candidates for those underlying masked priming. In particular, we can point to the following arguments in favour of our approach:

1. We have built upon the precedent set by Houghton and Tipper's (1994) work. Furthermore, similar opponent systems have played an important role in models of serial order recall such as (Burgess & Hitch, 1999). Thus, when viewed as a body of research, one might suggest that opponent effects are ubiquitous—arising at many levels of cognition, which makes it more likely that they underlie the masked priming effect. Indeed this is not surprising considering the ubiquity of inhibitory interneuron populations in the brain (cf. O'Reilly & Munakata, 2000).
2. Our model is simple and canonical. This is always an advantage in modelling work.
3. The techniques employed are psychologically plausible and reflect a level of cognitive mechanism commensurate with the dynamics of the masked priming task.
4. The model accurately reproduces the available masked priming data. In particular, not only have we accurately reproduced RT data, we have also reproduced the spectrum of available LRP data. It is common for computational models to reproduce RT data, but our success in modelling electrophysiological data is much less common. Thus, in contrast to other models we also obtain an indication that the ongoing time course of our model's behaviour is consistent with the human system. This is a significant strength of our work, which at the least has allowed us to assess our model against a very large number of data points, e.g., the separation profiles for each experiment have around 60 data points.
5. In assessing the value of a computational model, it is generally accepted that a qualitative fit to the experimental data (i.e., reproducing the broad pattern of results) is sufficient to demonstrate the worth of a model. The data presented in this paper demonstrates that our model definitely succeeds in this qualitative sense. However, we have also been able to obtain a relatively accurate quantitative fit to the human data, as demonstrated by the correlation shown in Figure 36 and the close match between our model's separation profiles and the LRPs recorded from humans. Such a more precise fit does bring benefits. For example, the main prediction arising from our model, which was encapsulated in Figure 39, has both a qualitative and a quantitative component. Specifically, not only do we assert that NCEs will turn into PCEs as mask–target ISIs are increased—a qualitative prediction—our model also asserts a particular mask–target ISI

at which NCEs turn into PCEs (a mask–target ISI of 150 ms)—a quantitative prediction. Although it is extremely unlikely that NCEs will turn into PCEs at exactly the point predicted by the model, it does suggest a region where experiments should direct their focus and, most significantly, without a close quantitative fit between our model and the human data, such a quantitative claim could not be made. Finally, it is important to emphasize again that although the model contains a large number of parameters, there is only one set (of parameters) being proposed and this is used to generate the, aforementioned, quantitatively precise fit across the spectrum of conditions.

A key objective of our work has been to provide a testable theory of activation followed by inhibition effects arising from subliminal priming. We have done this via a concrete implementation of our theory. This implementation is based upon a number of key mechanisms—opponent processes (as implemented by inhibitory interneurons); response competition (as implemented by lateral inhibition); competitive masking mechanisms (as implemented by feedforward inhibition); delineation of response channels through response-set maintenance; and response selection via discounted evidence accumulation. The value of the computational model is that it allows testable predictions to be made, which we have documented earlier (cf. the ''Simulation predictions'' subsection on p. 449).

As previously suggested, a potential weakness of our model is that a number of parameters need to be set in order to obtain the behaviour sought (all of which are documented in the Appendix). However, we would argue that these parameter settings are integral to the model being proposed; testing our predictions also tests these parameter settings.

Our theory is consistent with a cognitive-level explanation of activation followed by inhibition. This ''emergency brake hypothesis'' argues that the suppression implements a low-level response retraction mechanism. In a direct (below conscious) parameter specification setting there could be considerable fluctuation of sensory evidence and responses may be inappropriately activated on the basis of transient sensory evidence. The opponent inhibitory circuit is a mechanism to retract such inappropriately excited responses.

REFERENCES

Abeles, M., Bergman, H., Margalis, E., & Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology*, *70*, 1629–1638.

Agliotti, S., DeSouza, J. F. X., & Goodale, M. A. (1995). Size-contrast illusions deceive the eye but not the hand. *Current Biology*, *5*, 679–685.

Allport, A., Tipper, S. P., & Chimel, N. R. J. (1985). Perceptual integration and postcategorical filtering. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance XI: Attention and neuropsychology* (pp. 107–132). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Arbuthnott, K. D. (1995). Inhibitory mechanisms in cognition: Phenomena and models. *Current Psychology of Cognition*, *14*, 3–45.

Aron, A., Schlaghecken, F., Fletcher, P. C., Bullmore, E. T., Eimer, M., Barker, R. A., et al. (2003). Inhibition of subliminally primed responses is mediated by the caudate and thalamus: Evidence from fMRI and Huntington's disease. *Brain*, *126*, 713–723.

Baddeley, A. D. (1997). *Human memory, theory and practice*. Hove, UK: Psychology Press.

Baddeley, A. D. (2000). The episodic buffer: A new component of working memory? *Trends in Cognitive Sciences*, *4*(11), 417–423.

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. A. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 8, pp. 47–89). New York: Academic Press.

Band, G. P., & van Boxtel, G. J. (1999). Inhibitory motor control in stop paradigms: Review and reinterpretation of neural mechanisms. *Acta Psychologica*, *101*, 179–211.

Bokura, H., Yamaguchi, S., & Kobayashi, S. (2001). Electrophysiological correlates for response inhibition in a go/nogo task. *Clinical Neurophysiology*, *112*, 2224–2232.

Bowman, H., Aron, A., Eimer, M., & Schlaghecken, F. (2001a, November 2001). *Connectionist modelling of inhibitory processes in motor control* (Tech. Rep. No. 14-01). Canterbury, UK: Computing Laboratory, University of Kent at Canterbury.

Bowman, H., Aron, A., Eimer, M., & Schlaghecken, F. (2001b). *Neural network modelling of inhibition in visuo-motor control*. Paper presented at the seventh Neural Computation and Psychology Workshop on Connectionist Models of Cognition and Perception, Brighton, UK.

Bridgeman, B., Kirch, M., & Sperling, A. (1981). Segregation of cognitive and motor aspects of visual function using induced motion. *Perception and Psychophysics*, *29*(4), 336–342.

Bruin, K. J., Wijers, A. A., & van Staveren, A. S. J. (2001). Response priming in a go/nogo task: Do we have to explain the go/nogo N2 effect in terms of response activation instead of inhibition? *Clinical Neurophysiology*, *112*, 1660–1671.

Burgess, N., & Hitch, G. J. (1999). Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review*, *106*(3), 551–581.

Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, *280*(5364), 747–749.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, *97*(3), 332–361.

Dehaene, S., Naccache, L., Le Clec'H, G., Koechlin, E., Mueller, M., Dehaene-Lambertz, G., et al. (1999). Imaging unconscious semantic priming. *Nature*, *395*, 597–600.

Dehaene, S., Posner, M. I., & Tucker, D. M. (1994). Localization of a neural system for error detection and compensation. *Psychological Science*, *5*, 303–305.

De Jong, R., Liang, C., & Lauber, E. (1994). Conditional and unconditional automaticity: A dual-process model of effects of spatial stimulus–response correspondence. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 731–750.

Eimer, M. (1993). Effects of attention and stimulus probability on ERPs in a go/nogo task. *Biological Psychology*, *35*, 123–138.

Eimer, M. (1995). Stimulus–response compatibility and automatic response activation: Evidence from psychophysiological studies. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 837–854.

Eimer, M. (1999). Facilitatory and inhibitory effects of masked prime stimuli on motor activation and behavioural performance. *Acta Psychologica*, *101*, 293–313.

Eimer, M., & Schlaghecken, F. (1998). Effects of masked stimuli on motor activation: Behavioural electrophysiological evidence. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1737–1747.

Eimer, M., & Schlaghecken, F. (2001). Response facilitation and inhibition in manual, vocal, and oculomotor performance: Evidence for a modality-unspecific mechanism. *Journal of Motor Behaviour*, *33*, 16–26.

Eimer, M., Schubö, A., & Schlaghecken, F. (2002). The locus of inhibition in the masked priming of response alternatives. *Journal of Motor Behaviour*, *34*(1), 3–10.

Enns, J. T., & di Lollo, V. (2000). What's new in visual masking? *Trends in Cognitive Sciences*, *4*(9), 345–352.

Enns, J. T., & di Lollo, V. (2002). What competition? *Trends in Cognitive Sciences*, *6*(3), 118.

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in visual search. *Perception and Psychophysics*, *16*, 143–149.

Falkenstein, M., Hoormann, J., & Hohnsbein, J. (1999). ERP components in go/nogo tasks and their relation to inhibition. *Acta Psychologica*, *101*, 267–291.

Francis, G. (1997). Cortical dynamics of lateral inhibition: Metacontrast masking. *Psychological Review*, *104*(3), 572–594.

Francis, G. (2000). Quantitative theories of metacontrast masking. *Psychological Review*, *107*(4), 768–785.

Francis, G., Grossberg, S., & Mingolla, E. (1994). Cortical dynamics of feature binding and reset: Control of visual persistence. *Vision Research*, *34*(8), 1089–1104.

Frith, C., & Dolan, R. (1996). The role of the prefrontal cortex in higher cognitive functions. *Cognitive Brain Research*, *5*, 175–181.

Fuster, J. M. (1989). *The prefrontal cortex: Anatomy, physiology, and neuropsychology of the frontal lobe*. New York: Raven Press.

Gehring, W. J., Gratton, G., Coles, M. G. H., & Donchin, E. (1992). Probability effects on stimulus evaluation and response processes. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 198–216.

Godefroy, O., Cabaret, M., Petit-Chenal, V., Pruvo, J.-P., & Rousseaux, M. (1999). Control functions of the frontal lobes: Modularity of the central-supervisory system? *Cortex*, *35*(1), 1–20.

Goolkasian, P. (1997). Size scaling and spatial factors in visual attention. *American Journal of Psychology*, *110*, 397–415.

Grossberg, S., & Mingolla, E. (1985a). Neural dynamics of form perception: Boundary completion, illusory figures, and neon colour spreading. *Psychological Review*, *92*, 173–211.

Grossberg, S., & Mingolla, E. (1985b). Neural dynamics of perceptual grouping: Textures, boundaries and emergent segmentations. *Perception and Psychophysics*, *38*, 141–171.

Heil, M., Osman, A., Wiegelman, J., Rolke, B., & Henninghausen, E. (2000). N200 in the Eriksen-task: Inhibitory executive process? *Journal of Psychophysiology*, *14*, 218–225.

Houghton, G. (1994). Inhibitory control of neurodynamics: Opponent mechanisms in sequencing and selective attention. In M. Oaksford & G. D. A. Brown (Eds.), *Neurodynamics and psychology* (pp. 107–155). London: Academic Press.

Houghton, G., & Tipper, S. P. (1994). A model of inhibitory mechanisms in selective attention. In D. Dagenback & T. H. Carr (Eds.), *Inhibitory processes in attention, memory and language* (pp. 53–112). San Diego, CA: Academic Press.

Houghton, G., Tipper, S. P., Weaver, B., & Shore, D. I. (1996). Inhibition and interference in selective attention: Some tests of a neural network model. *Visual Cognition*, *3*(2), 119–164.

Jackson, S., & Houghton, G. (1994). Sensorimotor selection and the basal ganglia: A neural network model. In J. Houk & J. Davis (Eds.), *Models of information processing in the basal ganglia* (pp. 337–369). Cambridge, MA: MIT Press.

Jaskowski, P., van der Lubbe, R. H. J., Schlotterbeck, E., & Verleger, R. (2002). Traces left on visual selective attention by stimuli that are not consciously identified. *Psychological Science*, *13*(1), 48–54.

Jodo, E., & Kayama, Y. (1992). Relation of a negative ERP component to response inhibition in a go/no-go task. *Electroencephalography and Clinical Neurophysiology*, *82*, 477–482.

Jueptner, M., Stephen, K. M., Frith, C. D., Brooks, D. J., Frackowiack, R. S. J., & Passingham, R. E. (1997). Anatomy of motor learning: I. Frontal cortex and attention to action. *Journal of Neurophysiology*, *77*, 1313–1324.

Keysers, C., & Perrett, D. I. (2002a). Visual masking and RSVP reveal neural competition. *Trends in Cognitive Sciences*, *6*(3), 120–125.

Keysers, C., & Perrett, D. I. (2002b). What competition? (Response from Keysers and Perrett). *Trends in Cognitive Sciences*, *6*(3), 119.

Klapp, S. T. (in press). Two versions of the negative compatibility effect: A reply to Lleras and Enns (2004). *Journal of Experimental Psychology: General*.

Klapp, S. T., & Hinkley, L. B. (2002). The NCE: Unconscious inhibition influences reaction time and response selection. *Journal of Experimental Psychology: General*, *131*, 255–269.

Klein, R. M. (2000). Inhibition of return. *Trends in Cognitive Sciences*, *4*(4), 138–147.

Klotz, W., & Wolff, P. (1995). The effect of masked stimulus on the response to the masking stimulus. *Psychological Research*, *58*, 92–101.

Kok, A. (1986). Effects of degradation of visual stimuli on components of the event-related potential (ERP) in go/nogo reaction tasks. *Biological Psychology*, *32*, 21–38.

Kopp, B., Mattler, U., Goertz, R., & Rist, F. (1996). N2, P3 and the lateralized readiness potential in a nogo task involving selective response priming. *Electroencephalography and Clinical Neurophysiology*, *99*, 19–27.

Kopp, B., Rist, F., & Mattler, U. (1996). N200 in the flanker task as a neurobehavioral tool for investigating executive control. *Psychophysiology*, *33*, 282–294.

Leuthold, H., & Kopp, B. (1998). Mechanisms of priming by masked stimuli: Inferences from event-related brain potentials. *Psychological Science*, *9*, 263–269.

Lleras, A., & Enns, J. T. (2004). Negative compatibility or object updating? A cautionary tale of mask-dependent priming. *Journal of Experimental Psychology: General*, *133*, 476–493.

Luria, A. R. (1992). *Das Gehirn in Aktion—Einführung in die Neuropsychologie* [The brain in action—introduction to neuropsychology]. Reinbeck, IA: Rowohlt.

Marcel, T. (1980). Conscious and preconcscious recognition of polysemous words: Locating the selective effects of prior verbal context. In R. S. Nickersen (Ed.), *Attention and performance VIII* (pp. 435–457). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

McClelland, J. L. (1979). On the time-relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, *86*, 287–330.

McCormick, P. A. (1997). Orienting attention without awareness. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 168–180.

Miller, J. (1991). The flanker compatibility effect as a function of visual angle, attentional focus, visual transients, and perceptual load: A search for boundary conditions. *Perception and Psychophysics*, *49*, 270–288.

Naccache, L., & Dehaene, S. (2001). Unconscious semantic priming extends to novel unseen stimuli. *Cognition*, *80*, 215–229.

Neill, W. T., Valdes, L. A., & Terry, K. M. (1995). Selective attention and the inhibitory control of cognition. In F. N. D. C. J. Brainerd (Ed.), *Interference and inhibition in cognition* (pp. 207–261). San Diego, CA: Academic Press.

Neumann, O. (1990). Direct parameter specification and the concept of perception. *Psychological Research*, *52*, 207–215.

Neumann, O., & Klotz, W. (1994). Motor response to nonreportable, masked stimuli: Where is the limit of direct parameter specification. In C. Umiltà & M. Moskovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 123–150). Cambridge, MA: MIT Press.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

Page, M. P. A. (2000). Connectionist modelling in psychology: A localist manifesto. *Behavioral and Brain Sciences*, *23*, 443–512.

Pisella, L., Grea, H., Tilikete, C., Vighetto, A., Desmurget, M., Rode, G., et al. (2000). An ''automatic pilot'' for the hand in human posterior parietal cortex: Toward reinterpreting optic ataxia. *Nature Neuroscience*, *3*, 729–736.

Posner, M. I., & DiGirolamo, G. J. (1998). Executive attention: Conflict, target detection, and cognitive control. In R. Parasuraman (Ed.), *The attentive brain* (pp. 401–424). Cambridge, MA: MIT Press.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108.

Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the selection problem? *Neuroscience*, *89*, 1009–1023.

Rolls, E. T., & Deco, G. (2002). *Computational neuroscience of vision*. Oxford, UK: Oxford University Press.

Rolls, E. T., & Tovee, M. J. (1994). Processing speed in the cerebral cortex and the neurophysiology of visual masking. *Proceedings of the Royal Society of London, Series B*, *257*, 9–15.

Rolls, E. T., Tovee, M. J., & Panzeri, S. (1999). The neurophysiology of backward visual masking: Information analysis. *Journal of Cognitive Neuroscience*, *11*(3), 300–311.

Rolls, E. T., & Treves, A. (1998). *Neural networks and brain function*. Oxford, UK: Oxford University Press.

Sasaki, K., Gemba, H., Nambu, A., & Matsuzaki, R. (1993). No-go activity in the frontal association cortex of human subjects. *Neuroscience Research*, *18*, 249–252.

Schiller, P. H. (1968). Single unit analysis of backward visual masking and metacontrast in the cat lateral geniculate nucleus. *Vision Research*, *8*, 855–866.

Schlaghecken, F., Bowman, H., & Eimer, M. (2005). *Dissociating local and global levels of perceptuo-motor control in masked priming*. Manuscript submitted for publication.

Schlaghecken, F., & Eimer, M. (1997). The influence of subliminally presented primes on response preparation. *Sprache und Kognition*, *16*, 166–175.

Schlaghecken, F., & Eimer, M. (2000). A central/peripheral asymmetry in subliminal priming. *Perception and Psychophysics*, *62*(7), 1367–1382.

Schlaghecken, F., & Eimer, M. (2001). Partial response activation to masked primes is not dependent on response readiness. *Perceptual and Motor Skills*, *92*, 208–222.

Schlaghecken, F., & Eimer, M. (2002). Motor activation with and without inhibition: Evidence for a threshold mechanism in motor control. *Perception and Psychophysics*, *64*(1), 148–162.

Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge, UK: Cambridge University Press.

Simon, J. R. (1969). Reactions towards the source of stimulation. *Journal of Experimental Psychology*, *81*, 174–176.

SNNS Team. (2001). *Stuttgard Neutral Network Simulator*. Retrieved from http://www-ra.informatik.uni-tuebingen.de/SNNS/

Spence, S. A., & Frith, C. D. (1999). Towards a functional anatomy of volition. *Journal of Consciousness Studies*, *6*, 11–29.

Tipper, S. P. (1985). The negative priming effect: Inhibitory effects of ignored primes. *Quarterly Journal of Experimental Psychology*, *37A*, 571–590.

Weiskrantz, L., Warrington, E. K., Sanders, M. D., & Marshall, J. (1974). Visual capacity in the hemianopic field following a restricted occipital ablation. *Brain*, *97*, 709–728.

Manuscript received April 2003
Revised manuscript accepted November 2004

APPENDIX

Accumulated reaction times

In Tables 2, and 3 we document both human and model RTs for the set of experiments and conditions considered. Basic experiment human RTs are taken from Eimer and Schlaghecken (1998, p. 1740, Exp. 1a); low strength prime experiment human RTs are taken from Schlaghecken and Eimer (2002, Fig. 5); and ISI 150 experiment human RTs are new to this paper.

Residual delay calculation

Table 4 shows how the residual delay is calculated for each of the experiment conditions (see the main text for a description of this process, i.e., the ''Patterns and reaction times'' subsection on p. 421). The separation delays are read off from the corresponding LRP profiles (see Figures 2 and 5) and then the residual delay is calculated by subtracting the separation delay from the corresponding RT. We then take the average residual delay over the values in the last column, which yields the value 263 ms, and feed this into our calculation of model RTs.

TABLE 2
Human RT data

Condition	*Compatible (ms)*	*Incompatible (ms)*	*Neutral (ms)*
Basic experiment (ISI 100)	420.0	367.0	386.0
Low strength prime (ISI 100)	405.0	412.5	
Unmasked ''empty'' (ISI 150)	334.0	434.0	
Unmasked ''frame'' (ISI 150)	336.0	415.0	
Partial mask (ISI 150)	364.0	390.0	
Full mask (ISI 150)	398.0	364.0	

TABLE 3
Model RT data

Condition	*Compatible (ms)*	*Incompatible (ms)*	*Neutral (ms)*
Basic experiment (ISI 100)	388.3	355.0	371.7
Low strength prime (ISI 100)	355.0	371.7	371.7
Unmasked ''empty'' (ISI 150)	338.3	405.0	371.7
Unmasked ''frame'' (ISI 150)	338.3	405.0	371.7
Partial mask (ISI 150)	355.0	371.7	371.7
Full mask (ISI 150)	388.3	355.0	371.7
Arrow mask (ISI 100)	405.0	355.0	371.7

TABLE 4
Residual delay calculations for all conditions for which LRPs are available

Experiment	*Condition*	*Human RTs (ms)*	*Separation delay (ms)*	*Residual delay (ms) (RTs—separation delay)*
Basic experiment	Compatible	420.00	162.50	257.50
(ISI 100)	Incompatible	367.00	101.50	265.50
	Neutral	386.00	118.75	267.25
Unmasked empty	Compatible	334.00	73.50	260.50
(ISI 150)	Incompatible	434.00	153.00	281.00
Unmasked ''frame''	Compatible	336.00	92.00	244.00
(ISI 150)	Incomatible	415.00	147.00	268.00
Partial mask (ISI 150)	Compatible	364.00	134.75	229.00
	Incompatible	390.00	153.00	237.00
Full mask (ISI 150)	Compatible	398.00	171.50	226.50
	Incompatible	364.00	95.00	269.00

Details of model

Basic sigmoidal settings

Consider the parameters used in sigmoidal functions (cf. equation 3). First, bias values (the *bs* parameter) differ between layers (this is because we wish all nodes to asymptote at the same value—0.5—and we ensure this by varying biases). However, for all sigmoidals in the network (apart from those in OFF nodes, see the ''OFF node dynamics'' subsection on p. 463), we set the steepness (*sp*) parameter to 1 and the range (*rg*) to 1.2. The choice of a range of 1.2 (rather than 1) arose in order to ensure that the preactivation asymptote value (0.5) was just below the most responsive part of the sigmoidal (which is 0.6). However, there is no technical reason why the whole network could not be rescaled into a 0–1 range.

Weight settings

Link weight settings are documented in Table 5. The link being referred to is most easily identified with reference to Figure 12.

Parameters per layer

Biases (*bs* in equation 3) and tau parameters in activation time averaging (τ in equation 2) are set per layer. These settings are documented in Table 6. Layers have different biases because the connectivity of the network ensures that different layers have different basic excitatory pressures. Unless these differences in excitatory pressures are compensated for, which is the role of the bias term, different layers will asymptote during preactivation of the network at different levels. The bias settings employed here ensure that perceptual pathways and response nodes all asymptote at a value of 0.5 (due to their input threshold, OFF nodes asymptote at a much lower level, see the ''OFF node dynamics'' subsection on p. 463). Different values of tau are employed because different layers have different temporal dynamics. As previously discussed, perceptual pathways need to sustain their

TABLE 5
Weight settings

Links	*Purpose*	*Value*
PLi to PPj (i not equal to j)	Feedforward inhibition	–0.705
PLi to PPi	Feedforward excitation	6.140
PPi to R(*i*–1)	PP to corresponding response	6.000
Ri to Rj (i not equal to j)	Lateral inhibition between responses	–5.800
Ri to Oi	Excitation to OFF node	2.500
Oi to Ri	Inhibition from OFF node	–4.250
Response-set maintenance node to Ri and Oi	Response-set maintenance preactivation and deactivation	11.300
Accum to response-set maintenance node	Response selection control of response channel foregrounding	–10.000

activation. Hence a high value of tau is employed. In contrast, response nodes are more responsive to fluctuations in their net input. We discuss response-set maintenance node settings in the ''Response-set maintenance and response selection'' subsection on p. 463.

Perceptual pathways

As discussed in the main text (cf. the ''Double time averaging'' subsection on p. 31), perceptual pathway nodes employ a double time averaging, as regulated by the constant tau. That is, both net input and activation are time averaged. The value of tau for the latter of these was presented in Table 6, while the value for τ_1 in equation 4 is set to 0.2. Thus, net input time averaging is very responsive to changes in net input, while the activation time averaging built over this responsive net input is much more sustained and stable (i.e., a high value of tau).

We also build into the perceptual pathways a mechanism by which once they have crossed the preactivation asymptote of 0.5, they cannot drop back below this level (as made apparent in the perceptual pathway traces shown in the ''Emergent sustained perceptual trace'' subsection on p. 415). This is to reduce the degrees of freedom in our model. When, in earlier versions of our model, we allowed both perceptual pathways and response nodes to freely fluctuate below preactivation levels, a very complex dynamics emerged from the network and it was difficult to set parameters correctly. In the current version it is only response nodes that are free to become suppressed below 0.5, in this way the strong inhibitory pressures are localized to response nodes, which is where our theoretical argumentation has positioned them.

TABLE 6
Parameter settings per layer

Layer	*Bias*	*tau*
Perceptual pathways	–0.325	0.850
Response layer	–11.201	0.315
OFF nodes	–12.300	0.650
Response-set maintenance node	3.000	0.700

OFF node dynamics

As previously discussed the activation dynamics of OFF nodes are somewhat different from those of other nodes. This is consistent with our theoretical position that such nodes have a high input threshold, are very responsive (i.e., build up activation very rapidly) and also saturate quickly, see our discussion in the ''Opponent processing'' subsection on p. 405. Thus, the steepness parameter (*sp*) is set to 0.02, the bias is set to –12.3, and OFF nodes saturate at 0.273 (i.e., this is an absolute upper bound on the activation level of such nodes). It is important to note that with such a strongly negative bias, OFF node preactivation asymptotes just above zero. Thus, medium levels of net input will only have small affects on the activation level of the node. However, large levels of net input will have a dramatic (because of the setting of *sp*), but bounded (by saturation level) effect.

Response-set maintenance and response selection

As discussed in the main text (cf. the ''Response-set maintenance, response foregrounding, and selection'' subsection on p. 409, and the ''Response-set maintenance and accumulators'' subsection on p. 419), the response-set maintenance node feeds a constant level of activation into the relevant response pathway. However, this delineating activation is dependent upon the strength of activation of the response-set maintenance node. For example, when the node has an activation of zero, the response channel will become deactivated and will return to the background.

The response-set maintenance node has a strong positive bias (see Table 6) and only one incoming link, which is from the accumulator node. This link is inhibitory (see Table 5). Thus, as long as the accumulator node outputs zero activation, due to its bias, the response-set maintenance node will be strongly activated. The accumulator node outputs zero activation, unless the selection criterion is satisfied, at which point it outputs a 1. This has a strongly inhibitory pressure on the response-set maintenance node and causes the response pathway to be unbound, i.e., returned to the background.

In order to be consistent with the activation dynamics throughout the model, response-set maintenance node activation builds up gradually, via a time averaged sigmoidal of the net input (Table 6 has parameter settings). Thus, during preactivation, response pathway foregrounding builds up gradually and after response selection, returning of pathways to the background progresses gradually. Activation of the accumulator node is regulated by equation 6, where τ_{acc} is set to 0.665. As long as activation (a_{acc}) is below the selection threshold, this node outputs zero. However when the threshold is crossed, the node outputs 1. The selection threshold is set at 2.72.

Sufficiency of lateral inhibition investigations

As discussed in the main text (cf. the ''Insufficiency of lateral inhibition alone'' subsection on p. 424), we explored cutting the OFF node circuits and varying the strength of lateral inhibition. Thus, in these investigations links R1 to O1 and R2 to O2 are both set to zero. In addition, as we vary the lateral inhibition weight between responses we also vary the response node bias. This is because changing lateral inhibition between responses also changes the inhibitory forces on response nodes, thus in order to ensure that response nodes asymptote at 0.5 we have to correspondingly adjust response node biases in order to compensate. These settings are documented in Table 7. Finally, more extreme values of lateral inhibition (i.e., beyond –7) are not considered since we could not obtain stable network preactivations with these values, i.e., the network went into an oscillatory pattern of activation.

Second sensory pathway

As discussed in the main text (cf. the ''Short prime–target ISI conditions'' subsection on p. 435), in order to investigate short prime–target ISIs we also explore a variant of our basic network in which

TABLE 7
Bias and lateral inhibition settings

Biases of response nodes (R1 and R2)	*Lateral inhibition (between R1 and R2)*
–8.10	0
–7.10	–2
–6.10	–4
–5.10	–6
–4.65	–7

a second sensory pathway is added, as shown in Figure 30. The activation equations and parameter settings for the duplicate perception layer (PL2′ and PL3′) and perceptual pathways (PP2′ and PP3′) are identical to those in the corresponding layers in the basic network. The final settings in the duplicate sensory pathway are as follows: The PLi′ to PPi′ weight is set to 6.14, the PPi′ to R(i–1) weights are set to 7, and the response node bias is set to –14.208 (to compensate for the increased excitation it obtains during preactivation from the second set of perceptual pathway nodes). Note, that this second sensory pathway is only used when running short prime–target ISI conditions and all other conditions are run on the basic network.