CHAPTER 9

Behavioral, Neuroimaging, and Neuropsychological Approaches to Implicit Perception

Daniel J. Simons, Deborah E. Hannula, David E. Warren, and Steven W. Day

Abstract

For well over a century, the idea that rich, complex perceptual processes can occur outside the realm of awareness has either intrigued or exasperated researchers. Although popular notions of implicit processing largely focus on the practical consequences of implicit perception, the empirical literature has addressed more focused, basic questions: (a) Does perception occur in the absence of awareness? (b) what types of information are perceived in the absence of awareness? and (c) what forms of processing occur outside of awareness? This chapter discusses recent advances in the study of implicit perception, considering the ways in which they do and do not improve on earlier approaches. We contrast the conclusions a skeptic and a believer might draw from this literature. Our review considers three distinct but related classes of evidence: behavioral studies, neuroimaging, and brain-damaged patient case studies. We conclude by arguing that qualitative differences between perceptual mechanisms are interesting regardless of whether or not they demonstrate the existence of perception without awareness.

Introduction

> *. . . [T]here is now fairly widespread agreement that perception can occur even when we are unaware that we are perceiving. (Merikle & Joordens, 1997a, p. 219)*
>
> *Unconscious cognition is now solidly established in empirical research (Greenwald, 1992, p. 766).*
>
> *My contention is that most, if not all, claims for SA/CI [semantic activation without conscious identification] in dichotic listening, parafoveal vision, and visual masking are in reality based on the failure of these experimental methods to reveal whether or not the meaning of the critical stimulus was available to consciousness at the time of presentation (Holender, 1986, p. 3; brackets added)*

For well over a century, the idea that rich, complex perceptual processes can occur outside the realm of awareness has either intrigued or exasperated researchers. The notion that many of the cognitive processes that occur with awareness might also occur without awareness is both exciting and frightening; it would not only reveal untapped or unnoticed powers of mind but would also raise the specter of undesirable mechanisms of mind. If implicit cognitive processes are rich and powerful, then given the right tools, we might be able to exploit these resources – we might be capable of using far more information than reaches awareness. Alternatively, implicit processes might counteract our explicitly held attitudes, thereby changing our behavior without our knowledge (Greenwald & Banaji, 1995).

This fear has its roots in psychodynamic views of unconscious processing that attribute many psychological problems to unconscious conflicts and motivations (Freud, 1966). It manifests itself in the fear that subliminal advertising can affect our beliefs against our will (Pratkanas, 1992). These desires and fears drive a large market in subliminal self-help tapes as well as public outcry about apparent attempts at implicit influence. Yet, evidence for subliminal persuasion of this sort is scant at best (Greenwald, Spangenberg, Pratkanis, & Eskenazi, 1991; Pratkanis, Eskenazi, & Greenwald, 1994).

Although popular notions of implicit processing focus largely on the practical consequences of implicit perception, the empirical literature has addressed more focused, basic questions: (a) Does perception occur in the absence of awareness? (b) what types of information are perceived in the absence of awareness? and (c) what forms of processing occur outside of awareness? Few researchers question the idea that some perceptual processing occurs outside of awareness. For example, we are not usually aware of the luminance changes that lead to the perception of motion. Rather, we just perceive the motion itself. Some processing of the luminance boundaries occurs outside of awareness even if we are aware of the stimulus itself.

The more subtle, more interesting question is whether the meaning of a stimulus is processed without awareness. This problem is of fundamental theoretical importance because any evidence of semantic processing in the absence of awareness strongly supports late-selection models of attention and awareness (Deutsch & Deutsch, 1963). Presumably, implicit processes occur independent of explicit attentional selection, so if the meaning of a stimulus can be perceived implicitly, selective attention is not necessary for semantic processing. Each of these questions, at its core, asks how implicit perception is like explicit perception.

For more than a century, strong claims for the existence of complex perceptual processes in the absence of awareness have been dismissed on methodological grounds. In one early study, for example, observers viewed a card with a letter or digit on it, but their viewing distance was such that the character was hard to see – it was reported to be blurry, dim, or not visible at all. Although subjects could not consciously report the nature of the stimulus, they accurately guessed whether it was a letter or digit, and they could even guess its identity better than chance (Sidis, 1898). This lack of a clear conscious percept combined with better performance on an indirect, guessing task might provide evidence for implicit perception. However, alternative interpretations that require no implicit perception are equally plausible. For example, observers might simply be more conservative when asked to produce the name of a digit or letter than they would be when making a forced-choice decision (see Azzopardi & Cowey, 1998, for a similar argument about blindsight). This bias alone could account for better performance on a forced-choice task even if there were no difference in conscious perception. Moreover, the forced-choice task might just be a more sensitive measure of conscious awareness, raising the possibility that the dissociation between the two tasks is a dissociation within conscious perception rather than between conscious and non-conscious perception. Finally, the measure of awareness – the ability to recognize the character from a distance – might be

P1: JzG
0521857430c09 CUFX049/Zelazo 0 521 85743 0 printer: cupusbw November 6, 2006 15:48

inadequate as an assessment of awareness, leaving open the possibility that some conscious perception had occurred.

This example illustrates some of the weaknesses inherent in many studies of implicit perception. Although the behavioral and methodological tools for studying implicit perception became far more sophisticated toward the end of the 20th century, and despite some claims to the contrary (e.g., Greenwald, 1992; Merikle, Smilek, & Eastwood, 2001), the controversy over the mere existence of implicit perception persists (Dulany, 2004). Often, the same data are taken by some as convincing support for the existence of implicit perception and by others as unpersuasive (see the critique and responses in Holender, 1986).

In fact, theoretical reviews of the existing literature often arrive at strikingly different conclusions. Whereas Holender (1986) concludes that most demonstrations of implicit semantic processing are unconvincing, others consider the converging support for implicit effects to be overwhelming (e.g., Greenwald, 1992). In part, these divergent conclusions simply reflect different default assumptions. "Skeptics" assume the absence of implicit perception unless definitive evidence supports its presence. "Believers" assume the presence of implicit perception given converging evidence, even if none of the evidence is strictly definitive. At its core, the debate often devolves into little more than arguments over parsimony or over the criteria used to infer implicit processing.

The goal of this chapter is not to resolve this controversy. Nor is it to provide a thorough review of this century-old debate. Rather, we discuss recent advances in the study of implicit perception, considering the ways in which they do and do not improve on earlier approaches. We also contrast the conclusions a skeptic and a believer might draw from this literature. Since the mid-1980s, claims about implicit perception have become more nuanced, focusing less on the mere existence of the phenomenon and more on the nature of the information that might be implicitly perceived and on the mechanisms underlying implicit perception. Our review considers three distinct but related classes of evidence: behavioral studies, neuroimaging, and brain-damaged patient case studies.

Limits on the Scope of our Chapter

Given the availability of many excellent and comprehensive reviews/critiques of the early literature on implicit perception (e.g., Greenwald, 1992; Holender, 1986; Merikle, 1992), our chapter focuses primarily on the theoretical and methodological innovations introduced in recent years. Many disciplines include claims about implicit processing, and incorporating all of them in a single overview would be impractical. Instead, we highlight claims for implicit perceptual or semantic processing of discrete stimuli, largely overlooking implicit skill learning, artificial grammar learning, or other forms of procedural knowledge that might well be acquired without awareness. Our neglect of these areas does not imply any denigration of the evidence for implicit perception they have produced. Although we limit our review to the possibility of semantic processing without awareness and closely related questions, we also consider recent arguments about how best to study implicit perception. Finally, we discuss how qualitative differences in the nature of perceptual processing may be of theoretical significance even without a clear demonstration that processing occurs entirely outside of awareness.

Early Evidence for and against Implicit Perception

Claims for and against implicit perception received extensive empirical attention starting in the late 1950s, with sentiment in the field vacillating between acceptance and skepticism. Many early studies used a dichotic listening method in which observers attend to a stream of auditory information in one ear and verbally shadow that content while simultaneously ignoring another stream in their other ear (Cherry, 1953; Moray, 1959; Treisman, 1960, 1964). If the ignored channel is actually unattended and information from the ignored channel intrudes into awareness, then the ignored information must have been processed implicitly. With this technique, observers occasionally hear their own name in an ignored channel (Moray, 1959), and they sometimes momentarily shift their shadowing to the ignored channel when the auditory information presented to each ear is swapped (Treisman, 1960). If ignored information is truly unattended, then these findings support a strong form of late selection in which unattended information is processed to a semantic level and sometimes intrudes on awareness. Other studies using this dichotic listening technique found evidence for skin conductance changes to words in the ignored stream that were semantically related to shock-associated words (Corteen & Dunn, 1974; Corteen & Wood, 1972).

Of course, the central assumption underlying these conclusions is that an ignored auditory stream is entirely unattended. If participants periodically shift attention to the "ignored" channel, then the influence of semantic information in the ignored channel might occur only with attention. To conclude that perception of the semantic content of the ignored stream was caused by implicit processing, the experimenter must show that it did not result from explicit shifts of attention at the time of presentation. The difficulty of verifying that attention was never directed to the ignored channel gave meat to skeptics (Holender, 1986). In fact, this critique can be applied far more generally. The vast majority of studies of implicit perception, including those in the past 20 years, rely on what is commonly known as the dissociation paradigm (Merikle, 1992). To demonstrate the existence of implicit perception, experimenters must eliminate explicit perception and show that something remains. Applied to dichotic listening, the task for experimenters is to rule out attention to the ignored stream and then show that something remains. The failure of the premise, that ignored means unattended in the case of dichotic listening, weakens evidence for implicit perception. Given the fairly convincing critiques of evidence based on dichotic listening (Holender, 1986), few current studies use dichotic listening to study implicit perception. The dissociation paradigm, however, remains the dominant approach to studying implicit perception.

The modern use of the dissociation paradigm in the study of implicit perception was triggered by a series of experiments in the 1980s in which masked primes were shown to influence subsequent processing of a target stimulus even though observers did not notice the primes themselves (Marcel, 1983a, b). This approach is a classic application of the dissociation paradigm: Rule out explicit awareness of the prime stimulus and show that it still influences performance in some other way. Importantly, these studies provided evidence not just that something was perceived but also that its meaning was processed as well; the semantic content of a masked word served as a prime for a subsequent response to a semantically related target word (Marcel, 1983b). Many of the recent behavioral studies of implicit perception use variants of this masked prime approach.

The Merits and Assumptions of the Dissociation Paradigm

The dissociation paradigm is particularly appealing because it requires no assumptions about the nature of or mechanisms underlying implicit perception. In its purest form, the dissociation paradigm has a single constraint: Implicit perception can only be demonstrated in the absence of explicit perception. Superficially, this constraint seems straightforward. Yet, it amounts to confirming the null hypothesis – demonstrating no effect of explicit perception – leading some to decry its usefulness for the study of implicit perception (Merikle, 1994). Given that most claims for implicit perception are based on the dissociation paradigm, most critiques of these claims focus on violations of this assumption, often producing evidence that some contribution from explicit perception can explain the residual effects previously attributed to implicit perception (see Mitroff, Simons, & Franconeri, 2002 for a similar approach to critiquing evidence for implicit change detection). For example, critiques of dichotic listening studies typically focus on the possibility that subjects devoted some attention to the ignored stream (Holender, 1986). Given that the dichotic listening paradigm does not allow a direct measure of the absence of attention to the ignored channel, it cannot rule out the possibility that explicit factors contributed to perception of ignored material. More subtle critiques raise the possibility that observers were momentarily aware of ignored or unattended material, but rapidly forgot that they had been aware. If so, then explicit awareness could have contributed to any effects of the "unattended" information. This "amnesia" critique has been applied more recently to such phenomena as inattentional blindness (Wolfe, 1999).

To meet the assumptions of the dissociation paradigm, the measure of explicit perception must be optimally sensitive – it must exhaustively test for explicit influences on performance (Merikle, 1992). If a maximally sensitive measure reveals no evidence of explicit perception, we can be fairly confident that explicit factors did not contribute to performance, and any residual effects can be attributed to implicit perception. This criterion was adopted by some of the more ardent critics of the early literature on implicit perception (Holender, 1986). The explicit measure most typically adopted as a sensitive measure of explicit awareness is the simple detection of the presence of a stimulus. If subjects cannot detect the presence of a stimulus, but the stimulus still has an effect on performance, then that effect presumably resulted from implicit perception. In essence, this approach served as the basis for early work on priming by masked stimuli (Marcel, 1983b). If a masked prime cannot be detected but still influences performance, it must have been implicitly perceived. Note, however, that even a simple detection task may not exhaustively measure all explicit influences on performance and residual effects of a stimulus that cannot be detected might still reflect some explicit processing (Merikle, 1992). Later in this chapter, we review new behavioral studies that attempt to meet these assumptions, but we also note that few of them systematically demonstrate null explicit sensitivity to the presence of a stimulus.

Objective vs. Subjective Thresholds – What Is the Appropriate Measure of Awareness?

One recurring controversy in the study of implicit perception concerns whether the threshold for explicit perception should be based on an *objective* or *subjective* criterion. Although the notion of thresholds has fallen into disfavor with the advent and increased use of signal detection theory in perception (e.g., Green & Swets, 1966; Macmillan, 1986), it still has intuitive appeal in the study of implicit perception. Later in this chapter, we discuss the importance of using signal detection to measure awareness in the dissociation paradigm. In the interim, the distinction between *objective* and *subjective* thresholds may still provide a useful rubric for explaining some of the continuing controversy in the literature.

Most studies of implicit perception rely on a subjective threshold to determine whether or not a stimulus was explicitly noticed; this approach assumes that observers will report a stimulus if they are aware of it and will not if they are unaware. For example, blindsight patients typically will report no awareness of a static stimulus presented to their blind field – the stimulus falls below their subjective threshold. Use of the subjective threshold to rule out explicit perception essentially treats the observers' reports of their experiences as the best indicator of whether or not they were aware. More often than not, studies using subjective thresholds are interested in performance on each individual trial, and claims about implicit perception are derived from the consequences of a specific stimulus that was not reported. This approach is appealing because it treats observers' reports of their own mental states as more legitimate than the experimenter's ability to infer the observers' state of awareness.

P1: JzG
0521857430c09 CUFX049/Zelazo 0 521 85743 0 printer: cupusbw November 6, 2006 15:48

Objective thresholds are based on the idea that observers might fail to report a stimulus even if they did have some explicit awareness of its presence. They might adopt a conservative response bias, responding only when certain. Or, they might lack the means to express verbally what they saw. Typically, objective thresholds are measured across a large set of trials. The threshold is that level at which a stimulus is *not perceivable* rather than simply not perceived. In using this approach, experimenters often adopt the standard of null explicit sensitivity required by the dissociation paradigm, assuming that if a series of trials show that a stimulus is not explicitly perceivable, then it could not have been perceived on any individual trial. Consequently, any influence of that stimulus must be implicit. Unlike the subjective threshold approach, objective thresholds are based on the idea that observers might fail to report a stimulus not because they failed to see it, but because they adopted too conservative a criterion. This approach does not trust an observer's subjective experience on a given trial to be a true indicator of his or her actual awareness of the stimulus.

In a sense, the terms "objective" and "subjective" are misnomers. Both approaches rely on explicitly reported experiences, so both are subjective. Subjective thresholds are based on experiences on each trial, whereas objective thresholds are based on cumulative experiences across a larger number of trials. Thus, when measuring an objective threshold, responses on individual trials do not necessarily indicate the observer's awareness. Observers might respond that they saw a stimulus, but that response might simply be a guess. Similarly, they might report having no conscious experience, even if they had some vague inkling that failed to surpass their criterion for responding. Finding an objective threshold requires manipulating the stimulus presentation such that judgments of stimulus presence are no better than chance over a reasonably large number of trials. If responding to this sort of explicit task is at chance over a set of trials, then presumably any individual trial is based on a guess. The challenge is in demonstrating that explicit performance was truly random and not somewhat better than would be expected by chance alone.

The use of an objective threshold can lead to a seeming paradox wherein subjects report no conscious awareness of a stimulus (i.e., they report guessing) but still show better than chance performance; their performance exceeds the objective threshold even though their subjective impression is of guessing. Those adopting a subjective threshold approach would conclude that such a finding reflects implicit processing. The appeal of relying on the subjective threshold is that it accepts what the observer reports at face value. If observers report no awareness, then they had no awareness. However, it also relies on the observer's ability to judge probabilities over a series of trials. Does the subjective report of no awareness really mean that they were guessing, or does it mean that they *thought* that they were guessing? If observers lack precise access to their probability of a successful response, they might report guessing when in actuality, they were slightly, but significantly performing better than chance.

The primary difference between the objective threshold approach and the subjective threshold approach is that objective thresholds take the responsibility of estimating the extent of correct responding out of the observer's hands. Rather than relying on the observer to estimate when they felt they were guessing, the objective threshold technique objectively measures when their actual performance across a series of trials reflected guessing. In both cases, though, the subjects' subjective experience on a given trial contributes to the assessment of whether or not they were aware of the critical stimulus.

Differential reliance on objective and subjective thresholds underlies much of the controversy in the field. Most critiques of implicit perception simply show that performance actually exceeded an objective threshold for awareness. For example, evidence for implicit priming from masked stimuli was premised on the idea that subjects were no better than chance at determining whether or not the prime was present – explicit performance did not exceed the objective threshold (Marcel, 1983b). Yet critiques of those studies suggest that the thresholds were not adequately measured and that explicit performance might well have exceeded threshold (Holender, 1986). Even studies that do attempt to demonstrate that explicit detection was no better than chance rarely meet the statistical requirements necessary to infer null explicit sensitivity (Macmillan, 1986). Many studies, especially those of patients, make no attempt to measure an objective threshold, but instead rely entirely on the observer's self-assessment of awareness, much as early behavioral studies did (e.g., Sidis, 1898). Such studies are open to the criticism that explicit perception might well affect performance even when subjects do not consciously report the presence of a stimulus.

As we discuss later in this chapter, this issue is only of importance when questioning whether or not an example of perception is entirely implicit. Finding a dissociation in the types of processing that occur above and below a subjective threshold would still be of theoretical (and practical) import even if explicit perception contributed to both types of processing. For example, in studies of inattentional blindness, observers view a single critical trial and quite often fail to notice the presence of salient but unexpected objects and events (Mack & Rock, 1998; Most et al., 2001; Simons, 2000; Simons & Chabris, 1999). When counting the total number of times one team of basketball players passes a ball and simultaneously ignoring another team of players passing a ball, approximately 50% of observers fail to notice a person in a gorilla suit who walks through the display (Simons & Chabris, 1999). The interesting aspect of these studies is that observers can fail to notice or consciously detect surprisingly salient unexpected events. Most people expect that they would notice such events, and the fact that they do not report objects as unusual as a gorilla is startling (see Levin, Momen, Drivdahl, & Simons, 2000 for similar examples from the change blindness literature).

Unfortunately, the studies are not ideal for demonstrating implicit perception. Imagine, for instance, that observers in this study reported not noticing the gorilla, but then showed priming for the word "monkey." Would that provide evidence for implicit perception of the gorilla? The study uses the dissociation paradigm, and subjects subjectively report no awareness of a gorilla. This finding suggests that any priming effects might be implicit. However, observers might have had some awareness of the gorilla, or they might have had momentary awareness of some furry object, even if they failed to report noticing anything unusual. Given that the method only allows one critical trial and the "gorilla" is demonstrably perceivable (i.e., it is above the objective threshold), the possibility of some residual explicit awareness cannot be eliminated.

Arguments for implicit perception on the basis of such one-trial studies rest on the plausibility of the alternative explanations for the priming effects. As the measure of explicit awareness becomes less "objective" and more reliant on the observer's self-assessment, it is more likely to miss some aspect of explicit processing. The sufficiency of the measure of explicit awareness, regardless of whether it is considered objective or subjective, rests on the plausibility of the possibility that some explicit awareness was not tapped by the measure. Of course, even if the gorilla exceeded an objective threshold for awareness, this hypothetical finding would still be interesting because it would reveal a discrepancy between what people see and what they can explicitly report. Moreover, their surprise at having missed the gorilla suggests that their awareness of it likely was limited. Consequently, evidence for inattentional blindness may have important practical consequences even if some residual awareness of the unexpected event exists.

Rather than viewing the objective-subjective difference as a dichotomy, we prefer to characterize it as a continuum that varies along the dimension of the *experimenter's* confidence in the accuracy of the subjective judgments. With a subjective judgment on a single trial, the experimenter should lack confidence in the veracity of the observer's claim of no explicit awareness. One-trial approaches do not systematically eliminate the possibility that the stimulus was perceived and then forgotten, that some less-easily-reportable aspect of the stimulus was consciously perceived, or that the stimulus was explicitly detected and partially but not completely identified.

Critiques of the Dissociation Paradigm

Although the dissociation paradigm has intuitive appeal, some critics argue that the exhaustiveness requirement is a fatal shortcoming – that no task can fully satisfy the exhaustiveness assumption (Merikle, 1992). Even if a task were optimally sensitive to explicit perception and even if it showed null sensitivity, some other unmeasured aspect of explicit perception could still influence performance. Logically, this view is unassailable. Even if a task showed null sensitivity for all known explicit influences, it might neglect some as yet unknown and unmeasured explicit influence. Practically, however, if a task eliminates all known, plausible explicit influences, then claims of implicit perception might be more parsimonious than defaulting to some unknown explicit factor.

A second critique of the dissociation paradigm rests on the idea that no task measures just explicit or just implicit perception (Reingold & Merikle, 1988). Performance on any task involves a mixture of implicit and explicit influences. Consequently, finding null sensitivity on an "explicit" task might also eliminate implicit perception because the task likely measures aspects of both. By analogy, a sledgehammer to the head would eliminate all explicit awareness, but it also would eliminate most implicit effects on performance. Any manipulation that leads to null explicit sensitivity might simply be so draconian that no measure would be sufficiently sensitive to detect any implicit processes.

This *exclusivity* critique is based on the premise that tasks do not provide a pure measure of either implicit or explicit perception. Whether or not this premise is valid, the exclusiveness critique carries less force than the exhaustiveness critique. The failure to use an exclusive measure of explicit awareness is one reason why studies using the dissociation paradigm might *fail* to find evidence for implicit perception. The lack of exclusivity can only *decrease* the probability of finding implicit perception, and it should not spuriously produce evidence for implicit perception. Thus, positive evidence for implicit perception derived from the dissociation paradigm cannot be attributed to the lack of pure measures of implicit and explicit processing. If evidence for implicit perception using the dissociation paradigm is not forthcoming, failed exclusivity would provide a plausible explanation for how implicit perception might occur but be undetectable via the dissociation paradigm.

Recent Behavioral Approaches to Studying Implicit Perception

Despite concerns about the need for exhaustive measures of awareness, most recent studies of implicit perception have relied heavily on the dissociation logic. The approaches to studying implicit perception have become somewhat more refined in their treatment of the problem. In this section, we review several relatively new behavioral approaches to studying implicit perception. In some cases, these approaches follow the dissociation logic, but with improved attempts to exhaustively measure explicit influences. Others dismiss the dissociation paradigm as flawed and propose new approaches to measuring implicit perception. For each topic, we consider possible criticisms of the evidence for implicit perception, and at the end of the section, we provide contrasting conclusions that might be drawn by a believer and by a skeptic.

Modern Applications of the Dissociation Paradigm

Since the mid-1980s, the tools and techniques used to measure implicit perception have developed substantially, largely at the goading of skeptics (Holender, 1986). However, straightforward applications of the dissociation logic still dominate studies of implicit perception, and many (if not most) of them neglect to address the standard critiques of the dissociation paradigm. This shortcoming is particularly true of neuroimaging work and of studies using patient populations, where failures to provide an adequate exhaustive measure of awareness are commonplace (see Hannula, Simons, & Cohen, 2005 for a detailed discussion of neuroimaging evidence for implicit perception). In part, the methods in these studies are constrained by the need to include imaging measures or by the nature of the patient's deficit. However, behavioral studies of implicit perception are not limited in these ways, and a number of new techniques have emerged to provide sensitive and relatively rigorous tests of the existence of implicit perception.

Some of the simplest approaches are based closely on early studies of masked priming, focusing on the ability to perceive a target as a function of an unseen prime (e.g., Bar & Biederman, 1998; Watanabe, Nanez, & Sasaki, 2001). For example, one study examined naming accuracy for briefly presented line drawings (Bar & Biederman, 1998). For the first time a stimulus was presented, subjects were only able to name it correctly approximately 15% of the time. However, when the same stimulus was presented a second time, subjects were far more successful, suggesting that having seen the stimulus before, even without being able to name it, facilitated subsequent processing. This priming benefit only occurred when the same object was presented (a different exemplar of the same category received no priming) and was maximal when the object was presented in the same location. These results suggest that implicit processing of the prime stimulus led to facilitated naming of the target stimulus even when subjects typically were unsuccessful at naming the prime. Although this study is consistent with implicit perception, critics might well raise the objection that the explicit measure (naming) was not an exhaustive test of explicit awareness. Given that the logic of this task follows from that of the dissociation paradigm, unless explicit awareness of the prime is eliminated, naming improvements could result from residual explicit awareness.

Other studies adopted the repetition approach with a more rigorous measure of awareness of the initial stimulus (Watanabe et al., 2001), although these studies focused on perceptual learning rather than priming per se. While subjects performed a primary task involving the perception of letters in the center of a display, a set of dots behind were organized into somewhat coherent motion; most of the dots moved randomly, but a subset moved in a coherent direction. Critically, a small enough subset of the dots (5%) was coherent that subjects could not reliably discriminate the coherent motion displays from displays in which all dots moved randomly. The dots were entirely irrelevant to the primary task during the first phase of the experiment. Then, in a later phase, subjects attempted to judge the direction of coherent motion of another set of dot arrays, this time with somewhat more coherence (10%). Subjects were reliably better at determining the direction of these dot displays if they moved in the same direction as the previously viewed displays. Thus, even though subjects were unable to determine that the dots were moving coherently at all in the first phase of the experiment, the frequent repetition of a particular motion direction led to better performance with a somewhat easier judgment task. This indirect test provides evidence for implicit perception of the coherent motion of dots in the first phase, even though subjects had no conscious awareness of their motion. This approach is an elegant instance of the dissociation paradigm; subjects could not reliably detect the presence of coherent motion in the prime stimulus, but the motion coherence still affected subsequent judgments. Perceptual learning approaches like this one have distinct advantages over typical priming experiments in that awareness of the prime stimulus can be psychophysically eliminated. Other more recent priming studies have attempted to adopt more rigorous measures of awareness as well.

Many of these recent studies exploit response compatibility as an indirect, but sensitive measure of perceptual processing. For example, an experiment might measure response latency to a supraliminal target preceded by a supposedly subliminal prime. If the target would require a different response than the prime, subjects might be slowed by the presence of the prime. If subjects do not consciously detect the prime, then response compatibility effects likely resulted from implicit processing of the prime. One large advantage of this approach over traditional semantic priming studies in the dissociation paradigm is that response compatibility effects can be positive, negative, or absent, allowing additional ways to measure the effects of an unseen stimulus.

Given that this approach adopts the dissociation logic, experiments must provide direct evidence for the invisibility of the prime. As for studies of masked semantic priming, most decrease detectability by limiting presentation times and by adding masking stimuli before and/or after the prime (e.g., Eimer & Schlaghecken, 2002; Naccache & Dehaene, 2001b). Others have used small differences in contrast to camouflage primes against a background of a similar color (e.g., Jaskowski, van der Lubbe, Schlotterbeck, & Verleger, 2002). Even within the masked presentation approach, however, studies vary in terms of how systematically they manipulate the visibility of the prime. Some studies use a single stimulus duration, contrast level, or type of masking for all subjects (e.g., Naccache & Dehaene, 2001b), whereas others adjust the stimulus presentation to account for individual differences in perceptibility (Greenwald, Draine, & Abrams, 1996). Both approaches can work provided that neither shows any evidence of explicit detection of the prime stimulus. Unfortunately, many of the studies using a constant prime and mask across subjects do not entirely eliminate explicit perceptibility for all subjects, raising some concerns about the exhaustiveness assumption.

Although early studies of priming by masked stimuli focused on semantic priming by words, more recent studies using response compatibility have adopted a host of different stimuli and judgment tasks, including left/right discrimination of arrows (Eimer, 1999; Eimer & Schlaghecken, 2002; Klapp & Hinkley, 2002); concrete/abstract word discrimination (Damian, 2001); lexical decision (Brown & Besner, 2002); words and pictures in animacy judgments (Dell'Acqua & Grainger, 1999; Klinger, Burton, & Pitts, 2000); words and nonword stimuli in Stroop interference tasks (Cheesman & Merikle, 1984; Daza, Ortells, & Fox, 2002); words in positive/negative valence judgments (Abrams & Greenwald, 2000; Abrams, Klinger, & Greenwald, 2002); numerals and number words in relative magnitude judgments (Greenwald, Abrams, Naccache, & Dehaene, 2003; Naccache & Dehaene, 2001b; Naccache, Blandin, & Dehaene, 2002;); names in male/female judgment (Greenwald et al., 1996); and diamonds and rectangles in shape categorization (Jaskowski et al., 2002). Despite the varied stimuli and judgment tasks, the results of these studies are remarkably consistent.

Moreover, all of these approaches to compatibility effects fall into roughly four types: (1) centrally presented masked primes followed by a target, (2) centrally presented masked primes followed by a target with a limited interval for an allowed response (i.e., a "response window"), (3) masked flanker tasks, and (4) Stroop tasks. Findings from the first two approaches are reviewed below. A few of these studies were accompanied by neuroimaging results, some of which are discussed in this section and some of which are considered in the section on neuroimaging evidence for implicit perception.

Masked Priming without a Response Window

The influence of masked primes on response time and accuracy to subsequently presented

P1: JzG
0521857430c09 CUFX049/Zelazo 0 521 85743 0 printer: cupusbw November 6, 2006 15:48

target items varies as a function of the compatibility of the responses mapped to those items (Dehaene et al., 1998; Eimer, 1999; Eimer & Schlaghecken, 1998; Koechlin, Naccache, Block, & Dehaene, 1999; Naccache & Dehaene, 2001b; Neumann & Klotz, 1994). In many cases, target items elicit faster, more accurate responses when the target and prime require the same, compatible response than when they require different or incompatible responses (Dehaene et al., 1998; Koechlin et al., 1999; Naccache & Dehaene, 2001b; Neumann & Klotz, 1994). In one task, subjects judged whether an Arabic numeral or number word target was greater than or less than 5 (Dehaene et al., 1998; Koechlin et al., 1999; Naccache & Dehaene, 2001a; Naccache & Dehaene, 2001b). Target numbers were preceded by a compatible or incompatible number prime (e.g., if 6 were the target, a prime of 7 would be compatible and a prime of 4 would be incompatible). In this case, compatible primes benefited performance regardless of whether or not the prime was masked (Koechlin et al., 1999). Moreover, the compatibility effects persisted even when the notation of the target and prime were different (i.e., Arabic numerals primed both Arabic numerals and number words), suggesting that the priming effect must be more abstract than feature-based visual matching.

Not all studies show a positive effect of compatibility, however. In fact, some studies show a negative compatibility effect (NCE) in which responses are slower and more error prone for compatible primes (Eimer & Schlaghecken, 1998)! For example, when a post-masked priming arrow pointed in the same direction as a subsequent target arrow, subjects were slower and less accurate than when the prime arrow pointed in the opposite direction (Eimer & Schlaghecken, 1998). One explanation for these contradictory results appeals to the effects of delays between the prime and the response on compatibility effects. In one experiment that systematically manipulated the delay, positive compatibility effects were found for short delays between the prime and the response, but negative effects of compatibility resulted from delays longer than 350–400 ms (Eimer, 1999; Eimer & Schlaghecken, 1998).

The transition from positive to negative effects has been characterized more completely using recordings of ERPs. The lateralized readiness potential (LRP), detected via ERP recording, measures the activation from motor cortex of the hemisphere opposite the response hand (Coles, Gratton, & Donchin, 1988) and provides a direct way to determine whether a stimulus leads to activation of motor cortex. On incompatible trials, the prime should elicit transient activation of motor cortex ipsilateral to the responding hand followed by contralateral motor cortex activation in response to the target. With a compatible prime and target, this ipsilateral activation should be absent. In fact, the behavioral compatibility studies often incorporated ERP recording and consistently found LRPs in response to masked primes (Dehaene et al., 1998; Eimer, 1999; Eimer & Schlaghecken, 1998). Ipsilateral activation was evident shortly after the prime, both for arrow primes and numerical stimuli. Assuming that the masked primes were not consciously perceived, these LRPs provide evidence of processing in the absence of awareness.

The time course of neural activation corresponding to a masked prime might also help explain the paradoxical negative compatibility effect sometimes observed with longer lags between the prime and response (Eimer, 1999; Eimer & Schlaghecken, 1998). The burgeoning neural activity associated with a subliminal prime diminishes rapidly when observers do not make an overt response. If inhibitory mechanisms, not yet fully characterized, are responsible for preventing an overt motor response to the masked prime (Eimer, 1999), then they might also induce a refractory period during which activation consistent with the prime is suppressed. Thus, activation in response to the consistent target would overlap temporally with this refractory period, leading to the paradoxical result of slowed responses with compatible primes. Regardless of whether the prime produces a positive or negative compatibility effect, these studies confirm that masked primes activate corresponding motor cortices.

Together, the behavioral and ERP evidence for compatibility effects suggests that unseen primes influence both performance and neural activity. However, these studies still follow the logic of the dissociation paradigm, and any claims for implicit perception must satisfy the exhaustiveness assumption. Otherwise, differences between visible and "subliminal" primes might just reflect different levels of explicit activation rather than a dissociation between explicit and implicit perception. In most response compatibility studies, the perceptibility of the prime is measured not during the primary task, but in a separate set of trials or separate control experiments (e.g., Naccache & Dehaene, 2001b). Although few subjects report having seen the primes after the primary task, performance in these separate prime perceptibility trials implies some awareness of the "subliminal" primes. For example, sensitivity using signal detection measure (d') ranged from 0 for some subjects to as high as 1.3 for other subjects (Naccache & Dehaene, 2001b). Given that d' levels of as low as .3 can reflect some reliable sensitivity to the presence of a prime and a d' level of 1 represents fairly good sensitivity, these studies do not adequately eliminate explicit awareness of the prime stimuli. Consequently, claims of compatibility effects that are devoid of any explicit awareness are not entirely supported; the masked primes might well have been explicitly detected by some of the subjects on some trials.

Masked Priming with a Response Window

One recent refinement of the masked priming approach involves the use of a speeded response to maximize the effects of implicit processing (Draine & Greenwald, 1998; Greenwald et al., 1996; Greenwald, Schuh, & Klinger, 1995). In this approach, subjects must make their judgment within a fixed temporal window after the presentation of the target (e.g., between 383 and 517 ms instead of a more typical response latency of about 600 ms). The goal in forcing speeded responses is to maximize any implicit compatibility effects based on the premise that such implicit compatibility effects might be short-lived. As is typical of implicit response compatibility studies, prime visibility was measured by asking subjects to detect the masked prime stimulus either in a simple detection task or in a discrimination task (e.g., distinguish between a word prime and a random string of digits). Not surprisingly given the fixed prime presentation durations, a number of subjects had d' levels above 0. However, these studies did not simply look at performance on the compatibility task and then presume that explicit awareness was nil. Rather, a new analytical approach was adopted: Regression was used to predict the level of the compatibility effect when explicit awareness was absent ($d' = 0$). If the intercept of the regression of the compatibility effect on explicit sensitivity is greater than 0, then the study provides evidence for implicit perception. That is, implicit processing is revealed when the indirect measure reveals some consequence of the perception of the prime even when explicit sensitivity is extrapolated to $d' = 0$. This approach revealed significant response compatibility effects for prime durations ranging from 17–50 ms when explicit sensitivity was extrapolated to $d' = 0$.

This approach was premised on the assumption that the response window was necessary to detect implicit compatibility effects. Another experiment tested the validity of this assumption by varying the stimulus onset asynchrony (SOA) between the prime and the target (Greenwald et al., 1996). If more than 67 ms elapsed between the prime and target onsets, masking the prime eliminated the compatibility effect. In contrast, unmasked primes produced compatibility effects at a wide range of SOAs. This finding represents an important qualitative difference between visible and subliminal primes. Moreover, the regression technique and the response-window methodology are valuable contributions to the study of implicit perception.

More importantly, the findings raise some important limitations on implicit processing. Findings from this response-window technique suggest that implicit effects are extremely short-lived and are disrupted by even slight increases to the delay between the prime and the target. If this form of perception proves to be the only reliable way to find evidence for implicit perception, it would undermine more radical claims about the pervasiveness of implicit processes, especially implicit persuasion. Most hypothesized processes of implicit persuasion would require a much longer delay between the priming stimulus and the changed belief or action. These findings also provide an explanation for why studies of implicit perception often fail to replicate – the effects are ephemeral.

Although the response compatibility effects seem to provide evidence for implicit semantic processing, many of the findings could be attributed to motor interference rather than to semantic priming. Subjects learn responses to a stimulus, and it is the responses that conflict, not the abstract or semantic representations of those stimuli. In the response-window approach, semantic priming effects are difficult to produce, and most results can be attributed to response compatibility rather than any more abstract priming (Klinger et al., 2000). In fact, the effects, at least in some cases of word primes, seem not due to the semantic content of the word, but rather to response associations formed earlier in the experiment. In one striking example, subjects were asked to make positive/negative valence judgments about words. In the critical trials, words that previously had been used as targets were recombined in a way that changed the valence and then were used as a prime word. For example, the targets "smut" and "bile" would become the prime word "smile." Although the semantic representation of "smile" should lead to a compatibility benefit for a positive target, it instead facilitated processing of negative words (Abrams & Greenwald, 2000; Greenwald et al., 2003)! Moreover, unless a word or part of a word had been consciously perceived as a target during an earlier phase of the experiment, it produced no priming at all. In fact, other evidence not using a response-window technique suggests that not only must a word be consciously perceived to later serve as an effective prime but it must also have been used as a target such that the word would be associated with a motor response (Damian, 2001). In the example above, the word "smile" without prior exposure to "smut" and "bile" did not prime positive or negative words (Abrams & Greenwald, 2000; Greenwald et al., 2003). This claim directly contradicts other evidence of implicit semantic processing (Dehaene et al., 1998; Koechlin et al., 1999; Naccache & Dehaene, 2001b) by suggesting that only fragments of words are processed implicitly and that the associations they prime are developed through conscious experience as part of the experiment. Yet, evidence for priming of Arabic numerals by number words implies priming of more abstract representations, and such studies also showed priming from stimuli that had not previously been the target of a judgment (Naccache & Dehaene, 2001b). Moreover, switching the required response did not eliminate priming, so the effect cannot be entirely due to some automated form of response priming (Abrams et al., 2002).

In a recent intriguing paper, the primary adversaries in the argument over the nature of priming in the response-window paradigm combined their efforts to determine whether the effects were due to more than response compatibility (Greenwald et al., 2003). These experiments adopted numerical stimuli (Naccache & Dehaene, 2001b) in a response-window task. The stimuli were all two-digit Arabic numerals, and the judgment task required subjects to determine whether the target was greater or less than 55. Unfortunately, the use of two-digit numbers precluded the assessment of cross-notation priming, which was one of the strongest arguments for semantic processing in earlier experiments (Naccache & Dehaene, 2001b). The experiment replicated the finding of response compatibility effects with stimuli that had not previously been used as targets, refuting the argument that subjects must have formed a response association to a stimulus for it to produce priming (Abrams & Greenwald, 2000; Damian, 2001). However, the study

P1: JzG
0521857430c09 CUFX049/Zelazo 0 521 85743 0 printer: cupusbw November 6, 2006 15:48

also replicated the counter-intuitive finding that prior judgments affect the directionality of priming (Abrams & Greenwald, 2000). If 73 had served as a target, then subsequently using 37 as a prime enhanced response times to numbers greater than 55! Taken together, these findings imply that previously unclassified primes can produce compatibility effects and that they do so based on long-term semantic representations. However, such representations are overridden once a prime has been consciously classified, and then its features lead to priming based on the response association formed during the experiment.

Interim Conclusions

The response-window and regression approach lend new credibility to the traditional dissociation technique, and they show exceptional promise as a way to produce consistent evidence for priming by masked stimuli. Although some of the findings using this method are counter-intuitive and others are contradictory, the basic approach represents one of the best existing attempts to meet the challenges of critics of implicit perception (e.g., Holender, 1986). The approach is firmly couched in the dissociation logic, and most experiments make a laudable attempt to eliminate explicit sensitivity. The regression approach, in particular, is a clever way to examine performance in the absence of awareness. However, the existing literature does leave plenty of wiggle room for skeptics unwilling to accept the existence of implicit perception. First, because the approach adopts the dissociation paradigm, the measures of explicit sensitivity might fail to measure explicit sensitivity exhaustively (Merikle & Reingold, 1998).

Perhaps of greater concern to those who are otherwise willing to adopt the dissociation logic is the nature of the regression approach itself. The approach has been criticized for making assumptions about the nature of the relationship between the direct and indirect tasks and measures (Dosher, 1998). For example, the conclusions from the regression approach often depend on extrapolation, with relatively few subjects (e.g., 25%) performing at chance on the explicit detection task and a sizable minority of subjects (25%) showing substantial explicit sensitivity with d′ levels greater than 1 (Greenwald et al., 1996). If most subjects show greater than chance explicit sensitivity, the extrapolation to zero sensitivity might not be appropriate. A skeptic could easily imagine a non-linearity in the relationship between implicit and explicit measures when explicit performance is just barely above $d' = 0$. Perhaps there is a qualitative difference between minimal sensitivity and fairly good sensitivity. If so, then extrapolating to no sensitivity from fairly good sensitivity would not allow a clear conclusion in favor of implicit effects. Of course, this concern could be remedied with a more systematic manipulation of prime visibility within rather than across subjects, thereby obviating the need for any extrapolation. Given the trend toward progressively more sophisticated analyses and methodologies in this literature, this new approach shows great promise as an effective use of the dissociation paradigm.

Alternatives to Dissociation

The concerns about exhaustiveness and the possible role of failed exclusivity in minimizing evidence for implicit perception have spurred a new approach to studying implicit perception: Concentrate on qualitative or quantitative differences between tasks that purportedly measure implicit perception to different degrees. Examining differences in performance on these tasks as a function of an experimental manipulation can reveal the operation of distinct implicit and explicit processes. Two types of "relative differences" methodologies have used this logic: (1) the *relative sensitivity procedure*, which looks for greater sensitivity to stimulus presence with indirect measures than with direct measures of awareness, and (2) the *process dissociation procedure*, which looks for qualitatively different performance for implicit and explicit perception. Neither methodology requires

P1: JzG
0521857430c09 CUFX049/Zelazo 0 521 85743 0 printer: cupusbw November 6, 2006 15:48

process-pure measures of implicit or explicit perception. Rather, both assume that all tasks have implicit and explicit components. Carefully designed experiments can pull apart the underlying processes, revealing differences between the implicit and explicit processing in a given task. Both approaches also assume that implicit and explicit processes underlie functionally different types of behavior; explicit processes underlie intentional actions, whereas implicit processes govern automatic (non-intentional) behaviors. A behavior may result from a conscious, deliberate decision or from an automatic predisposition, or a combination of the two. The relative differences methodologies attempt to show that such automatic and deliberate processes can lead to qualitatively different performance.

Relative Sensitivity: An Alternative to Dissociation

The goal of the relative sensitivity procedure is to reveal implicit processes by showing instances in which indirect measures are more sensitive than comparable direct measures in making a given discrimination. This approach assumes that performance of any task involves both implicit and explicit contributions, neither of which can be measured exclusively by any task. Direct tasks measure performance when subjects are instructed to use their percept of a critical stimulus to make a judgment or discrimination. Indirect tasks involve an ostensibly unrelated behavior or judgment that nevertheless can be influenced by perception of a critical stimulus. Although the direct task might not exclusively measure explicit contributions, on its face it is demonstrably more explicit than the indirect task. Any decision-making process that relies on conscious awareness of the critical stimulus should lead to better performance on a direct measure than on an indirect measure because subjects should optimally rely on their conscious percept. In contrast, indirect measures do not require conscious perception of the critical stimulus, so subjects are unlikely to rely on conscious processing of that stimulus in making their judgment. Therefore, if indirect measures reveal better performance than direct measures, implicit processes must have influenced performance.

One critical component of this paradigm is that the two tasks must be equated in most respects. Unless the visual displays are equivalent and the task requirements comparable, any performance differences could be caused by the differences between the displays or the task demands and requirements. Proponents of this approach rightly take pains to make sure that the only difference between the direct and indirect tasks are in the instructions (a similar approach has been adopted in the study of implicit memory; see Schacter, 1987).

Note that this criterion – equivalency across direct and indirect measures – is not often met in studies of implicit perception. Many experiments use entirely distinct indirect and direct measures, making comparability more difficult. When observers report no awareness of a stimulus on a direct measure, indirect measures such as eye movements, patterns of neural activation, skin-conductance changes, or ERPs might reveal sensitivity to the presence of a stimulus. Although these sorts of indirect measurements certainly provide important insight into the nature of the processing of the stimulus, they do not provide conclusive evidence for processing in the absence of awareness. They might only reveal greater sensitivity of the measure itself; using such measures to provide corroborating evidence for qualitative differences in implicit and explicit processing may prove more fruitful (see the neuroimaging evidence section below). For the inference of implicit processing to follow from the relative sensitivity of direct and indirect measures, however, the measures must be comparable.

In one of the first experiments to adopt the relative sensitivity approach for the study of implicit perception (Kunst-Wilson & Zajonc, 1980), subjects viewed a series of briefly presented pictures of geometric shapes. Then, they either performed an old/new recognition task (the direct task) or they picked which of two shapes they preferred (the indirect task). When performing the direct task, subjects performed no better than chance at discriminating previously viewed from novel shapes. In contrast, when performing the indirect preference judgment task, they preferred the previously studied shape over a novel shape at rates significantly above chance levels. In other words, a direct measure of conscious recognition showed less sensitivity to the presence of a representation than an indirect measure of preference. This experiment meets the standards necessary for inferring implicit processing in the relative sensitivity approach: (a) the experimental environment was constant across tasks, with only the task instructions changing across conditions, and (b) performance on the indirect task exceeded that on the direct task. By the logic of the relative sensitivity approach, the direct task represents a putatively better measure of conscious awareness, so the relatively increased sensitivity of the indirect task must have resulted from implicit processes.

One possible concern about this conclusion derives from the use of separate study and test phases rather than testing performance at the time of presentation; subjects may have perceived and forgotten the consciously experienced shape even if a vague, explicitly generated preference persisted longer and affected performance on the indirect task. Furthermore, subjects might have been less motivated to make the more difficult intentional recognition judgments than the simpler preference judgment, so they were more likely to select responses randomly. If so, then responding in both cases might reflect access to an explicit representation, with the relative increase in sensitivity for the indirect task resulting not from implicit processing but from a differential effect of motivation (cf., Visser & Merikle, 1999).

One representative experiment pitted a recognition task (direct) against a perceptual contrast judgment (indirect) in which subjects judged the contrast of a word against the background (Merikle & Reingold, 1991). In a study phase, subjects viewed pairs of words and were asked to read the cued one. Then, in the test phase, they viewed individual words against a noise background and either judged whether it was old or new (a direct recognition task) or judged whether it was presented in high or low contrast (an indirect measure). Performance on the contrast judgment task revealed greater sensitivity to the presence of a word in the study phase than did the direct recognition task (at least for the first block of trials). Presumably, the prior presentation reduced the processing demands, leading to a subjective impression that the words were easier to see against a noisy background even if the words were not recognized. Once again, the study used comparable stimuli in the direct and indirect tasks and found greater performance for the indirect task, suggesting implicit processing.

Problems with Relative Sensitivity as an Approach

Although this approach is touted as an alternative to the classic dissociation paradigm, any positive evidence for implicit perception is subject to many of the same assumptions. Positive evidence for implicit perception requires some task to have a greater implicit contribution than explicit contribution. Otherwise, performance on the indirect task could not exceed that on any task with a greater explicit component. If some task has a greater implicit than explicit component, then it should also be possible to make the task sufficiently difficult that the explicit component would be eliminated, leaving only the residual implicit component. That is, the "indirect > direct" approach is a superset of the standard dissociation paradigm that does not require the elimination of an explicit component. Yet, any case in which the indirect > direct approach reveals implicit perception would also support the possibility that the dissociation paradigm could reveal implicit perception. In essence, this approach amounts to a more liberal variant of the dissociation paradigm in which explicit processing need not be eliminated. However, as critics of early work on implicit perception have noted, whenever a stimulus is consciously perceptible, explicit factors may contaminate estimates of implicit processing.

A more general concern about this paradigm is that it assumes a unitary explicit contribution and a unitary implicit contribution. In arguing that a direct measure involves a greater explicit contribution than an indirect measure, the assumption is that the explicit contributions to each task are of the same sort. If more than one sort of explicit contribution exists, then a "direct" task might exceed an "indirect" task on some forms of explicit contribution but not others. Unless the direct task exceeds the indirect task on all explicit contributions, the logic underlying the paradigm fails. Just as the dissociation paradigm suffers from the problem of exhaustively eliminating all possible explicit contributions, the indirect > direct approach requires that the two tasks measure the same explicit component and only that explicit component. Consequently, for the logic of the paradigm to hold, the experimenter must exhaustively eliminate any extraneous explicit contributions to the indirect task that might explain superior performance on the indirect task. Given that this exhaustiveness assumption applies to the relative sensitivity approach, its advantage over the standard dissociation paradigm is somewhat unclear.

Qualitative Differences

One criterion often used to infer the existence of implicit perception relies on differences in the patterns of performance derived from implicit and explicit processes. When the pattern of performance diverges from what would be expected with explicit perception, then the processes leading to this qualitative difference might well be implicit. Qualitative differences in performance for implicit and explicit tasks or measures often provide an intuitive way to infer the existence of implicit perception. The negative compatibility effects described earlier provide one illustration of the importance of such differences for inferring implicit processing (Eimer, 1999). However, the interpretation of qualitative differences is often muddied by the challenge of determining whether differences in performance are qualitative rather than quantitative. An effect that initially appears to reflect a qualitative difference might simply be a difference along a non-linear dimension.

More importantly, though, qualitative differences in performance can occur even when subjects are aware of the stimulus (Holender, 1986). That is, qualitative differences are possible within explicit perception, so the existence of a qualitative difference in performance alone does not unequivocally demonstrate implicit perception. Rather, the qualitative difference must be accompanied by an exhaustive measure of explicit awareness. Consequently, qualitative differences can provide converging evidence for the existence of implicit perception, but they are not definitive in and of themselves (Holender, 1986). Perhaps the best example of the use of qualitative differences in studies of implicit perception comes from the use of the Process Dissociation paradigm (otherwise known as the "exclusion" paradigm).

Process Dissociation

In the process dissociation technique, implicit and explicit performance are put in opposition (Jacoby, 1991). As in the relative sensitivity approach, direct and indirect measures are thought to rely differentially on explicit and implicit processing. In this approach, intentional actions are assumed to be under explicit control, whereas automatic responses are thought to reflect implicit processing. Presumably, people will use consciously available information to guide their intentional actions. In contrast, information available only to implicit processes will be less subject to intentional control. Consequently, when subjects produce responses that differ from those associated with intentional actions, they may have been influenced by non-conscious processes. The critical difference between the process dissociation procedure and the relative sensitivity procedure is that the task instructions and goals are constant. Subjects always perform the same task. Rather than manipulating the task across conditions, the perceptibility of the critical stimulus itself is varied so that some responses are consistent with awareness of the stimulus and others are not.

Subjects are instructed to respond one way if they are aware of a stimulus, but implicit or indirect influences lead them to respond the opposite way by default. For example, in the original instantiation of this approach in the memory literature, subjects studied a list of words and then were asked to complete word fragments with words that had *not* been on the studied list (Jacoby, 1991). Presumably, if they remembered the studied word, they would successfully avoid it in the fragment task. If they did not explicitly remember the word, they might automatically or implicitly be more likely to complete a fragment with a studied word than a non-studied word. Implicit influences should increase the likelihood of completing fragments with studied words, whereas explicit influences should decrease the likelihood of completing fragments with studied words. The same logic can be applied to implicit and explicit perception: If subjects explicitly detect the presence of a word, they should avoid using it to complete a word fragment. However, if they do not detect it and it still influences performance implicitly, they should be more likely to complete a fragment with a studied word than a non-studied word.

Studies using this procedure have been taken to support the existence of implicit perception. For example, one study varied the presentation time for words. Immediately after viewing each word, subjects were given a word stem completion task in which they were asked to complete the stem with a word other than the one that had been presented (Debner & Jacoby, 1994). With long presentation durations, subjects were aware of the words and successfully avoided completing stems with the "studied" words relative to the baseline performance of subjects who had never been shown the word. In contrast, with shorter presentations, subjects completed the stems with the "studied" word more often than the baseline condition. Even when they were unable to use their memory for the word to guide their intentional actions (i.e., choose another word), the briefly presented word still received enough processing to increase its availability in the stem completion task.

A similar pattern emerges when attentional focus rather than presentation duration is manipulated (see Merikle et al., 2001 for an overview). In one study, subjects viewed a briefly presented cross and judged which of its two lines was longer. During a subset of trials, a word was presented briefly along with the cross (see Mack & Rock, 1998 for the origins of this method). Depending on the condition, subjects were asked to focus attention either on the cross judgment or on the words. In both conditions, subjects subsequently attempted to complete a word stem with a word that had not been presented (this study was described by Merikle et al., 2001). Those subjects who focused on the words performed well, rarely using the presented words to complete the stem. In contrast, those who focused attention on the cross judgment completed the stem with presented words more often than would be expected based on a previously determined baseline (see also Mack & Rock, 1998). When the words were the focus of attention, they presumably were available to awareness, and subjects could use that information to exclude them in the stem completion task. However, when subjects focused attention on the cross judgment, they were less aware of the words, but automatic processing of the words biased them to use the presented words in the stem completion task.

A variety of exclusion tasks have been used to study implicit perception. For example, subjects show differential effects of interference in a variant of the Stroop task when aware and unaware of a stimulus. Typically, when color patches are incongruent with a preceding word (e.g., a green patch preceded by the word "red"), subjects are slower to identify the color of the patch than if the two matched. However, if mismatches occur on a large proportion of trials (80%), subjects use this information to perform faster when the word and patch *mismatch* than when they match (see Merikle & Joordens, 1997b for a discussion of these studies). When words were presented long enough to be consciously detected, subjects used their explicit knowledge to override Stroop interference. In contrast, briefly presented words were not consciously detected, and subjects were significantly slowed when a word-color patch mismatch occurred (Merikle & Joordens, 1997a,b). Stroop interference can be counteracted if subjects are aware of the word and of the predictiveness of the word, but if the word is not consciously perceived, subjects cannot override these automatic interference effects.

Other exclusion studies have used similar manipulations of target predictability in response compatibility paradigms (e.g., McCormick, 1997). Subjects were asked to decide whether an X or O was in the display, and this target item was presented either on the right or left side of the fixation cross. Before the presentation, a cue appeared on the left or right side of the display. On approximately 80% of trials, the cue was on the side opposite where the target would appear. Thus, the cue predicted that the target would be on the opposite side of the display. When the cue was presented for long enough to be consciously detected, subjects responded more rapidly to targets on the side opposite the cue. In contrast, when the cue was presented too briefly to be consciously detected, subjects were faster to respond when the cue and target were on the same side of the display (McCormick, 1997). Presumably, the cue automatically attracts attention, and only with awareness can subjects override this automatic shift of attention. Without awareness, the cue automatically draws attention, leading to better performance when the target appears at the cued location. Although this finding does not involve semantic processing without awareness, it does suggest that attention shifts can be induced without awareness of the inducing stimulus.

Problems with Process Dissociation as a Measure of Implicit Perception

This approach has promise as a means of studying implicit perception. One concern about this approach, however, is that it might be subject to biases and motivational factors that affect the criterion that subjects adopt. If so, estimates of implicit processing might be inflated (Visser & Merikle, 1999). Any case in which the subject's criterion is differentially affected by exclusion and inclusion instructions can produce a change in the criterion that could then influence estimates of unconscious processing. For example, increasing incentives to exclude studied items led to improved performance, thereby decreasing estimates of unconscious processing (Visser & Merikle, 1999). More broadly, variations in the degree of confidence or certainty in a representation or a percept can lead to different degrees of success on the exclusion task. Given that the exclusion task provides the basis for inferring implicit representations, such variations are problematic. A word that is explicitly detected, but with low confidence, might lead to a failure to exclude that item on a stem completion task even though there was an explicit contribution to perception. In terms of signal detection, if subjects were conservatively biased when reporting explicit detection, estimates of implicit perception would be inflated. Thus, as in the dissociation paradigm, the explicit task must demonstrably eliminate all explicit detection and must not be subject to conservative response biases for this paradigm to provide a clear estimate of implicit perception.

A Believer's Interpretation

The past 15 years have seen tremendous improvements in the behavioral methods used to study implicit perception. More importantly, many of the early critiques of the implicit perception literature have been addressed. Most studies using the dissociation paradigm now use signal detection theory to determine the explicit perceptibility of the prime stimulus, thereby providing a more convincing demonstration that priming results from implicit processing rather than from explicit contamination. The recently introduced technique of regressing performance on an indirect measure (e.g., a response compatibility effect) on performance on an explicit detection task provides a more nuanced approach to the dissociation technique. Even when performance on the explicit task is extrapolated to null sensitivity, performance on some indirect measures is still better than chance. The use of response compatibility allows an indirect measure that can, under the right circumstances, reveal implicit semantic processing. For example, priming persists even when the format of a number (text vs. Arabic numeral) changes from prime to test. The combination of the regression technique and response compatibility paradigms provides a powerful new tool to study implicit perception, one that has produced consistent and replicable evidence for implicit perception. Finally, work using process dissociation and relative sensitivity approaches reveals evidence for qualitative differences between implicit and explicit processing. These qualitative differences suggest that different mechanisms underlie implicit and explicit perception, thereby providing further evidence for the existence of implicit perception. In sum, evidence from a wide variety of tasks and measures provides support for implicit perception and even for semantic processing in the absence of awareness. Given the wide variety of tools used in the study of implicit perception, the converging evidence for the existence of implicit perception is overwhelming.

A Skeptic's Interpretation

The tools and techniques used to study implicit perception have improved immensely over the past 20 years. Many studies have adopted signal detection theory as a way to verify the absence of explicit perception, thereby making evidence from the dissociation paradigm less subject to the standard criticisms. Moreover, seeking evidence of qualitative differences using the process dissociation paradigm or other relative sensitivity approaches is a promising avenue for the exploration of implicit perception. However, none of these approaches or studies provides airtight evidence for implicit perception, and all are subject to fairly plausible alternative explanations that rely solely on explicit mechanisms. For example, in studies using the regression approach, the direct measure often reveals sensitivity to the presence of the prime stimulus at levels far above $d' = 0$ (Naccache & Dehaene, 2001b); the prime is readily visible to some subjects. Consequently, the inference for implicit perception relies on extrapolation of performance to $d' = 0$ from a number of subjects who show positive sensitivity to the stimulus. This extrapolation is potentially hazardous, particularly if the distribution of subjects is not centered on d' of 0. If the relationship between explicit perception and the indirect measure is non-linear, the extrapolation may be invalid (Dosher, 1998). Moreover, the presence of a positive indirect effect might require only a minimal amount of explicit sensitivity. No published studies have examined the effect of varying explicit sensitivity systematically (within subjects) on the magnitude of the indirect response compatibility effect. Any application of the dissociation paradigm, including the regression approach, depends critically on demonstrating null sensitivity to the presence of the critical stimulus. None of the studies to date have done so adequately.

Evidence from the process dissociation paradigm suggests a qualitative difference between implicit and explicit perception, something that would be more difficult to explain via explicit contamination. Most studies of implicit perception simply reveal "implicit" effects that are weaker versions of what would be expected with explicit processing. The process dissociation procedure, in contrast, suggests that implicit and explicit mechanisms differ. However, as accurately noted in critiques of the implicit perception literature, qualitative differences alone are insufficient to claim evidence for implicit perception. The qualitative difference could simply be a dissociation between two forms of explicit perception rather than between implicit and explicit perception. Moreover, inferences of implicit perception depend on the extent to which the explicit, intentional response fully measures all of the explicit processing. The only studies to address this question suggest that performance on the explicit task can be enhanced via motivation manipulations, thereby decreasing the evidence for implicit perception (Visser & Merikle, 1999).

In sum, the new tools introduced to study implicit perception may be promising, but the evidence for implicit perception is not yet convincing. Moreover, the implicit effects that have been reported are small and tend to vary with the extent of demonstrated explicit awareness, hinting that the "implicit" effects might well be driven by residual explicit processing. For a study using the dissociation paradigm to make a strong claim for implicit perception, no subject should show explicit sensitivity to the visibility of the critical stimulus; no study to date has met this strict criterion. Converging solid evidence from a variety of techniques can provide powerful support for a claim of implicit perception, but the convergence of weak and controvertible evidence for implicit perception does not merit strong support for the claim. If all of the evidence can be explained by plausible explicit confounds, then there is no need to infer the existence of a separate mechanism or set of mechanisms.

Evidence for Implicit Perception – Neuroimaging Data

Neuroimaging approaches provide several distinct advantages over behavioral approaches in the study of implicit perception. First, the effects of a subliminal stimulus can be assessed without an overt response; neuroimaging techniques provide an additional dependent measure of the consequences of perception, one that may allow dissociations that would be impossible with strictly behavioral measures. Moreover, differences in the pattern of activation for explicit and implicit perception might reveal additional qualitative differences between these forms of processing even if behavioral responses show no difference; neuroimaging might simply provide a more sensitive measure. Finally, the known functions of various brain regions can be mapped onto the pattern of activation produced in response to seen and unseen stimuli, allowing yet another way to determine the richness of implicit percepts.

Although such approaches have great promise as a new tool for the study of implicit perception, in many respects the existing research on the neural bases of implicit perception falls prey to the same critiques leveled at the behavioral research. Perhaps more importantly, as our review suggests, the neural activity elicited by implicit perception often is similar to that corresponding to overt perception, just diminished in amplitude. In the absence of qualitative differences in the pattern of activation, such diminished effects might well result from low-level overt perception. In such cases, the same standards and criteria applied to the use of the dissociation paradigm in behavioral research must be applied to the neuroimaging methods (see Hannula et al., 2005 for a detailed treatment of these issues).

A wide array of neuroimaging tools, most notably functional magnetic resonance imaging (fMRI) and event-related potentials (ERPs), have been adapted to the study of implicit perception. Most often, these investigations draw on existing knowledge of the functional brain regions likely to be involved in overt perception of a class of stimuli (e.g., emotional faces, words, etc.), and then try to determine whether those same regions are active even when observers report no awareness of the stimuli. Neuroimaging studies of implicit perception typically rely on several different types of processing and stimulus classes, and for the sake of organizing this rapidly expanding field, we consider three types of evidence for implicit perception: implicit perception of faces, implicit perception of words and numbers, and ERPs

P1: JzG
0521857430c09 CUFX049/Zelazo 0 521 85743 0 printer: cupusbw November 6, 2006 15:48

in priming studies. Within the neuroimaging literature, most inferences about implicit perception depend critically on the pattern of neural localization or the magnitude of activation resulting from explicitly detected and implicitly perceived stimuli.

Implicit Perception of Faces

Face processing represents one of the more promising avenues for the study of implicit perception because the neural regions activated in response to faces are fairly well described in the neuroimaging literature. A slew of recent neuroimaging studies of face perception reveal an area in the fusiform gyrus that responds relatively more to faces than to other stimuli (the Fusiform Face Area, or FFA; Kanwisher, McDermott, & Chun, 1997; McCarthy, Puce, Gore, & Allison, 1997). Is this area active even when observers are unaware of the presence of a face stimulus? Also, fearful faces are associated with activation of the amygdala (Breiter et al., 1996; Morris et al., 1996). Do fearful faces lead to amygdala activation even when they are not consciously perceived? Finally, recent neuroimaging studies using the phenomenon of binocular rivalry have explored the areas that are activated by stimuli when they are consciously perceived and when rivalry removes them from awareness.

Recent neuroimaging studies of visual extinction patients have explored whether an extinguished face leads to activation in the FFA (Rees et al., 2000; Vuilleumier et al., 2001). Unilateral brain lesions, particularly those located in the right posterior inferior parietal lobe, are associated with spatial neglect of the contralesional visual field. Many neglect patients exhibit visual extinction, accurately detecting isolated stimuli presented in either visual field, but failing to identify a contralesional stimulus when items are presented simultaneously in both visual fields. Behavioral research (discussed in later sections of this chapter) provides evidence for residual processing of extinguished stimuli, perhaps due to intact striate and extrastriate cortex along with ventral inferotemporal areas that process object identity.

One study required a patient to respond differently to a stimulus presented solely on the left, solely on the right, or simultaneously on the left and right (Rees et al., 2000). Given that the patient had right inferior parietal lobe damage, extinction would be revealed by incorrect "right-side" responses when a stimulus was presented on both the left and the right simultaneously. By comparing residual neural activity corresponding to correct right-side responses and incorrect right-side responses on extinction trials, residual neural activity associated with extinction could be revealed. Extinguished stimuli activated striate and early extrastriate cortex in the damaged right hemisphere – a pattern of activation no different from that elicited by left-side stimuli that were consciously perceived. This activation of early visual cortex occurred regardless of whether the patient was aware of the stimulus, suggesting that these areas are not sufficient for conscious awareness. More importantly, a region of interest analysis revealed low-threshold, category-specific activation in the right FFA in association with extinguished face stimuli, suggesting that the extinguished face was processed by the same regions as for consciously perceived faces (for reviews of evidence for preserved activation in response to unreported stimuli, see Driver & Vuilleumier, 2001; Driver, Vuilleumier, Eimer, & Rees, 2001).

This basic pattern was replicated in a similar experiment using both fMRI and ERPs (Vuilleumier et al., 2001). An extinction patient with right-lateralized posterior inferior parietal damage indicated on each trial whether or not a face was presented. Stimuli (i.e., schematic faces and shapes) were presented unilaterally in the right or left hemifield or bilaterally. Again, extinguished faces activated right striate cortex as well as an area of inferior temporal cortex just lateral to the FFA, although the level of activation was much reduced relative to that for visible face stimuli. Furthermore, ERPs revealed a right-lateralized negativity over posterior temporal regions approximately 170–180 ms after a face was presented in the left hemifield. This N170, a component known to be face-selective, was evident regardless of whether the face was perceptible or not. Interestingly, this experiment varied the duration of the bilateral presentations in order to vary whether or not extinction occurred. Awareness of the left visual field stimulus evoked activation of striate cortex and fusiform gyrus coupled with increased activation of a network of frontal and parietal brain regions, reflecting the sorts of long-range associations or widespread activation thought to accompany consciousness (Baars, 1988). Thus, differences in activation strength and functional connectivity distinguish conscious from unconscious perception.

Evidence from patients with bilateral amygdala damage (e.g., Adolphs, Tranel, Damasio, & Damasio, 1995) and neuroimaging of intact individuals (Breiter et al., 1996; Morris et al., 1996) support a role for the amygdala in processing fear-related stimuli, such as fearful faces. At least one theory suggests that the human homologue of a direct short-latency pathway between the thalamus and the amygdala might underlie the processing of emotional stimuli even in the absence of awareness (Le Doux, 1996). In one fMRI study (Whalen et al., 1998), fearful and happy faces were presented for 33 ms followed immediately by a neutral-face mask. Based on previous behavioral studies, the 33-ms masked presentation was assumed to be below the threshold for awareness. Post-study questioning found that eight of ten subjects denied having seen emotional faces and did not select these faces as having been in the stimulus set. Under these conditions, the unnoticed fearful faces did elicit a relatively circumscribed increase in amygdala activation relative to masked happy faces and a fixation baseline. This amygdala activation was attenuated with repeated exposure to masked fearful faces, a finding consistently observed with visible faces as well. Further, increased activation in response to both masked fearful and happy faces extended into the adjacent sublenticular substantia innominata of the basal forebrain, a region thought to be involved in more general processing of emotional stimuli and arousal (although activation was more pronounced for fearful faces). This pattern of results is consistent with the notion that the amygdala is selectively recruited when subliminal fear stimuli are presented (for additional evidence of early affective word processing in the absence of awareness see Bernat, Bunce, & Shevrin, 2001a), but as for similar behavioral results, such dissociations must be interpreted with caution because of methodological shortcomings in the assessment of awareness. For example, awareness was not measured directly on each trial – doing so might change the subject's strategy, and most similar fMRI studies examine the pattern of activation with passive viewing rather than active search (Whalen et al., 1998).

In another study, relative to neutral faces, fearful faces were associated with significant activation of the left amygdala, left fusiform, lateral orbitofrontal, and right intraparietal cortex (Vuilleumier et al., 2002). Activation of fusiform gyrus in response to extinguished faces was much reduced relative to activation for visible faces though, and the activation evident in association with extinguished stimuli may be a consequence of feedback from the amygdala. Together, these findings suggest that emotional stimuli can receive substantial processing even if they fail to reach awareness. Emotional stimuli are among the most promising approaches to the study of implicit processing, precisely because of the hypothesized existence of a direct, perhaps more primitive neural pathway that bypasses higher cognitive areas.

These studies provide interesting, suggestive support for the hypothesized short-latency pathway originating in the thalamus (Le Doux, 1996). Such a pathway might reasonably allow for processing even in the absence of more complex cognitive processes, and by inference without awareness. More importantly, amygdala activation was not significantly modulated by awareness (Vuilleumier et al., 2002), suggesting that processing of extinguished stimuli extends beyond early visual processing areas and that activation need not be less robust in the absence of conscious detection. Similar

P1: JzG
0521857430c09 CUFX049/Zelazo 0 521 85743 0 printer: cupusbw November 6, 2006 15:48

approaches have been taken in the study of implicit processing of unnoticed, emotionally arousing stimuli (see Lane & Nadel, 2000 for an overview of work on the cognitive neuroscience of emotion).

Another neuroimaging-based approach to studying processing in the absence of awareness relies on the phenomenon of binocular rivalry. When two patterns are presented simultaneously, one to each eye, the contents of conscious awareness spontaneously alternate between one monocular percept and the other over time. The visual percepts compete for awareness such that only one image is consciously perceived, and the other is suppressed (Levelt, 1965; Wheatstone, 1838). The oscillation of perceptual awareness between two simultaneously presented stimuli provides a useful tool to identify the neural correlates of conscious awareness (for reviews, see Rees, Kreiman, & Koch, 2002; Tong, 2001, 2003).

A growing number of investigations have been conducted using fMRI to address, in particular, the contributions of specific brain regions to perceptual awareness of rivalrous stimuli. One recent study using fMRI (Tong, Nakayama, Vaughan, & Kanwisher, 1998) presented face and house images separately to each eye and measured the neural activity in two predefined regions of interest: the FFA, which responds preferentially to faces (Kanwisher et al., 1997; McCarthy et al., 1997), and a parahippocampal region that responds most strongly to places and less so to faces (the Parahippocampal Place Area, or PPA; Epstein & Kanwisher, 1998). During imaging, participants continuously reported whether they saw a face or a house, and the pattern of neural activity extracted from a region of interest analysis was time-locked to these conscious perceptual experiences. Interestingly, neural activation corresponded to the conscious perceptual experience, even though the stimulus pair was invariant within a trial; FFA activation increased when participants reported perception of a face stimulus, and PPA activation increased when they reported a house. Critically, the pattern of activation when subjects consciously perceived a face or a house when both were present (in the rivalrous stimulus) was no different than when the face or house was presented alone, suggesting that the competitive neural interactions responsible for rivalry are largely resolved before conscious perception occurs.

This finding might suggest that activation in the FFA or the PPA produces visual awareness of the presence of a face or a house. However, the FFA also is active when faces are not consciously reported (Rees et al., 2000), suggesting that reliable FFA activation is not sufficient for conscious perception of a face. This discrepancy might result from different degrees of activation, though. If neural activity is graded with respect to the level of perceptual awareness (i.e., low-level activity reflects low-level awareness) or if activity must surpass some threshold before conscious awareness occurs, then it is entirely possible that sufficient activation of the FFA or PPA does correspond to conscious awareness of a face or house, respectively. Stricter criteria for measuring conscious awareness are needed to determine whether activation in these specialized processing regions is sufficient for conscious perception.

Implicit Processing of Words and Numbers

Just as consciously perceived emotional stimuli activate the amygdala, read words tend to activate a prescribed set of brain regions more than do other stimuli (e.g., left-lateralized extrastriate cortex, fusiform gyrus, and precentral sulcus). Therefore, studies of implicit word perception can use neuroimaging evidence to determine whether words activate a similar set of regions without awareness. Such studies first assess the visibility of the critical words using behavioral measures. In one study (Dehaene et al., 2001), masked words were presented such that they were detected only 0.7% of the time (a rate slightly higher than the false alarm rate of 0.2% for trials in which no word was presented) and almost never named successfully (see also Rees, 2001). Moreover, recognition tests after the imaging portion of the study revealed no memory for the masked words. Of course, subjects might adopt a conservative criterion for indicating whether or not a word was present if they knew they would then be asked to name it. If so, then the task might not exhaustively measure conscious awareness, raising the possibility that the masked words were at least temporarily available to consciousness.

Assuming that low detection rates and failed recognition performance imply the absence of conscious awareness of the presence of the masked words, and if neural activity is consistent with reading, then perception presumably occurred implicitly. Interestingly, when compared to control conditions that mimicked the masking conditions of the critical trials but without any masked words, the unseen stimuli activated the previously mentioned set of brain regions known to be associated with reading (Dehaene et al., 2001). This pattern is consistent with the idea that the unseen stimuli were processed similarly to visible words. However, the pattern of neural activity evoked by the masked words in the ventral visual pathway was less widely distributed and of smaller magnitude than that obtained with consciously perceived words. The discrepancy was increasingly evident from posterior to anterior brain regions, suggesting that visual masking begins to suppress neural activity early in the visual processing stream, rendering later stages of visual processing less likely. Furthermore, visible words elicited neural activity in parietal, prefrontal, and cingulate cortices, but corresponding activation was not evident when the words were not available to conscious awareness. Finally, increased correlated activity among the ventral visual stream, parietal, and prefrontal areas was evident only when the words were visible. Some of these differences might well result from the naming task used in the study rather than from the perceptibility of the stimuli. Visible words could be named, but the masked words were not. However, it cannot be determined on the basis of these results whether some of the activity associated with visible words is a consequence of the naming task. In sum, masking resulted in less robust neural activation, but also in reduced correlated neural activity that might contribute to conscious awareness.

Similar patterns have emerged in neuroimaging studies of the perception of numerical stimuli (Naccache & Dehaene, 2001a). Neuroimaging and lesion data suggest a role for the parietal lobe (and particularly the intraparietal sulcus) in the mental representation and understanding of the quantity meaning of numbers (for a review, see Dehaene, Piazza, Pinel, & Cohen, 2003). Can implicit stimuli lead to similar patterns of activation? A recent paper (Naccache & Dehaene, 2001a) reanalyzed earlier neuroimaging data (Dehaene et al., 1998) and addressed this issue by using the phenomenon of repetition suppression. A number of imaging studies have shown that when a stimulus is repeated, localized neural activity associated with processing of that stimulus or its attributes typically decreases (Schacter, Alpert, Savage, Rauch, & Albert, 1996; Schacter & Buckner, 1998; Squire et al., 1992). Whole brain analysis of fMRI data revealed two isolated brain regions with reduced activity when the target repeated the prime relative to an otherwise categorically congruent prime: the left and right intraparietal sulci (Naccache & Dehaene, 2001a). The priming effect was not influenced by the use of different notations for the prime and target (1 vs. one), suggesting that the intraparietal sulcus encodes numbers in a more abstract format. Assuming that the prime stimuli were not consciously perceived, these effects indicate that repetition suppression can occur even when observers are unaware of the repetition. Presumably, this effect reflects the fairly extensive processing of an implicitly perceived stimulus. Additionally, ERP studies of the number response compatibility effect reveal covert activation from an incongruent prime – a lateralized readiness potential (LRP) on the incorrect side of response – presumably because the incongruent prime activates the incorrect motor response (Dehaene et al., 1998). fMRI data revealed greater overall activation in right motor cortex when both the prime and target were consistent with a left hand response (and vice versa), providing additional evidence for processing of the prime stimulus without awareness (Dehaene et al., 1998). In all of these studies, perception of the prime in the absence of awareness was not limited to sensory mechanisms alone, but also influenced higher-level processing.

ERPs in Priming Experiments

The influence of an unseen prime stimulus has been explored by examining general changes in ERPs to a target as a result of the presence of a prime. These studies measure the influence of an implicit prime indirectly, looking for changed neural processing of the target rather than activation directly in response to the prime stimulus. Studies of priming by masked stimuli represent the paradigmatic application of the dissociation paradigm, and the use of ERPs in conjunction with this approach may well contribute to a more complete assessment of the processing of an unseen stimulus. To the extent that semantic processing of a prime takes place, it should lead to modulation of the N400 (i.e., a negative-going ERP component sensitive to manipulations of semantic relatedness). Experiments with supraliminal words and sentences consistently find larger deflections in N400 amplitude for incongruent than for congruent targets (Kutas & Hillyard, 1980). For example, the N400 generated in response to the word "lemon" would likely be more negative when preceded by the unrelated prime "chair" than when preceded by the related prime "citrus."

Unfortunately, studies of N400 modulation by semantically related, unseen primes have produced mixed results (for a review, see Deacon & Shelley-Tremblay, 2000). For instance, in one experiment, masked primes led to faster responses to semantically related targets, but modulation of N400 was evident only when primes were completely visible (Brown & Hagoort, 1993). This finding implied that the N400 might constitute an electrophysiological marker of conscious semantic processes. Yet, other experiments that account for potential methodological shortcomings of this experiment induce modulation of the N400 even when the primes were consciously inaccessible (e.g., Deacon, Hewitt, Yang, & Nagata, 2000; Kiefer, 2002). Moreover, the effects of a prime on the N400 are qualitatively different for visible and masked primes. Masked primes modulate the N400 with a short SOA between the prime and target, but not with a longer SOA. In contrast, for visible primes, the modulation of the N400 increases as the SOA increases (Kiefer & Spitzer, 2000). This qualitative difference suggests that implicit and explicit perception of prime stimuli might rely on different processing mechanisms.

Taken together, these studies provide support for N400 activation in response to an unseen prime. However, they are subject to many of the critiques leveled at the dissociation paradigm (Holender, 1986). For example, visibility of the prime on some trials might well contribute to the observed effects – the measure of awareness might not have been exclusive. One recent ERP study made a valiant effort to address many of the requirements of the dissociation paradigm (Stenberg, Lindgren, Johansson, Olsson, & Rosen, 2000). Most dissociation paradigm studies attempt to render the prime invisible using a masking procedure, assuming that the prime is invisible to all subjects on all trials. An alternative approach is to vary the visibility of the target itself and to measure the ERP response to an unseen target stimulus (Stenberg et al., 2000). This approach has the advantage of allowing the trial-by-trial measurement of the target.

In these experiments (Stenberg et al., 2000), a visible category name (the prime) was followed by a word that either was from the primed category or from a different category. Target perceptibility was varied across blocks so that individual subjects could successfully name the target on 50% of trials. Because this subjective naming task leaves criterion setting in the hands of the subject and does not sample conscious awareness exhaustively, several measures of conscious awareness were also administered at the end of each trial: (a) subjects indicated whether or not the word had been a member of the prime category, (b) they named the word (guessing if necessary), and (c) they attempted to select the target word from either a 2- or 6-alternative forced-choice test. The 6-alternative test was considered the most sensitive, hence the most exhaustive measure of awareness. Interestingly, the semantic priming effect (i.e., N400) distinguished between categorically consistent and categorically inconsistent words, irrespective of visibility. Although modulation of the N400 was less pronounced when the words could not be explicitly identified, the topographical pattern of activation did not differ across conditions. Qualitative differences in hemispheric lateralization were evident in an extended positive-going complex that typically accompanies cognitive tasks like the one employed in these experiments. This ERP component remained consistent irrespective of categorical classification, but had a different topography depending upon whether or not targets were explicitly identified. Consciously reported targets were associated with left- lateralized activity, whereas implicitly perceived targets elicited more distributed or right-lateralized activity, suggesting that different neural populations were recruited under these circumstances (Stenberg et al., 2000).

Together, the consistency of the N400 irrespective of visibility and differences in lateralization of raw amplitudes for visible and implicit targets strengthen claims for semantic processing of words that are not readily identified. When the criterion for conscious awareness was based on the more conservative 6-alternative forced-choice test, 30% of the words that could not be named were correctly identified and dropped from subsequent analyses. The binary categorization responses collected on the remaining trials were used to calculate d′, which was not different from 0 – providing even stronger evidence that the remaining target words were not available to conscious awareness. Despite using a more stringent objective criterion, modulation of the N400 remained intact (Stenberg et al., 2000). In fact, when a regression analysis was conducted to determine whether the N400 was more sensitive to categorical deviations than the binary-choice discrimination task, the intercept was reliably greater than 0. This experiment adopts most of the controls needed to make clear inferences from behavioral studies using the dissociation paradigm, but also adds a more sensitive neuroimaging measure to provide additional evidence for both quantitative and qualitative differences in the processing of consciously perceived and implicitly perceived stimuli.

Additional evidence for a change in the ERP pattern in response to an unseen stimulus comes from studies of the P300, a component typically occurring 260–500 ms after exposure to a relatively rare stimulus. In this case, the "rarity" of the target stimulus depends on its relation to other stimuli presented in the study. Would the target stimulus reveal this rarity response if the other stimuli were not consciously perceived? A number of studies have explored this question (e.g., Brazdil, Rektor, Dufek, Jurak, & Daniel, 1998; Devrim, Demiralp, & Kurt, 1997), but most are subject to the critique that subjects were aware of the regular or frequent stimuli and that they had a strong response bias when awareness was assessed in a separate block of trials (Bernat, Shevrin, & Snodgrass, 2001b).

One more recent experiment (Bernat et al., 2001b) showed modulation of the P300 to a rare target word even when more rigorous criteria for measuring conscious awareness were applied to make sure that the frequent words were not consciously detected (for a review, see Shevrin, 2001). The words LEFT and RIGHT were presented tachistoscopically in an oddball design with an 80:20 frequent-to-rare ratio. Frequent stimuli were made subliminal by presenting them for only 1 ms, and subjects were given a forced-choice detection block after the experiment. Collapsed across subjects, d′ did not differ from 0, but not all subjects showed a d′ of 0. Consequently, the effect could be driven by a few subjects who showed awareness on some trials. However, the correlation between the d′ score for a given subject and their P300 was negative, suggesting that more awareness of the frequent stimuli actually diminished the P300 amplitude. Moreover, a regression of P300 magnitude against d′ revealed a significant P300 effect even when d′ was extrapolated to 0.

Summary

One recurrent theme in this overview of the neuroimaging of implicit perception is that, when stimuli are not consciously perceptible, activation is often reduced relative to when they are consciously perceived. Importantly, activation in response to an unseen stimulus is not limited to early sensory processing and often activates brain regions associated with processing that particular type of stimulus. These findings suggest that implicit perception might be a weaker version of the same processes occurring for explicit perception. As for most studies of implicit perception, neuroimaging studies rely almost exclusively on the dissociation paradigm, attempting to eliminate explicit awareness and then attributing the residual effects to implicit perception. To the extent that these studies fail to meet the exhaustiveness assumption of the dissociation paradigm, they are subject to the same critiques often leveled at behavioral studies (Hannula et al., 2005). The strength of the evidence for implicit perception based on neuroimaging approaches depends on the extent to which the studies successfully demonstrate that processing has really occurred in the absence of awareness.

A Believer's Interpretation

Although some of the experiments fail to address the exhaustiveness assumption sufficiently, others provide more convincing tests of explicit awareness. Few individual studies provide unequivocal evidence for the effects of an unseen stimulus on brain activity; however, when considered holistically, the literature provides strong converging evidence. Some experiments provide evidence that processing was implicit and simultaneously demonstrate neural consequences of implicit perception. The strongest evidence comes from studies of differences in N400 amplitude in response to an implicitly perceived stimulus (Stenberg et al., 2000). Reliable differences in N400 amplitude were evident even when a fairly conservative 6-alternative forced-choice task was used to rule out explicit awareness on a trial-by-trial basis. By probing for awareness of the critical stimulus immediately after presentation, this study reduced concerns about fleeting conscious perception of the stimuli (i.e., memory failure following conscious perception). Further, the study adopted the regression technique (Greenwald et al., 1995) to show that N400 patterns persisted even when the d′ measure was extrapolated to null sensitivity. Finally, and perhaps most importantly, the patterns of neural activity elicited by implicitly and explicitly visible stimuli were qualitatively different, suggesting different neural mechanisms for the processing of implicit and explicit stimuli. Each piece of evidence can be criticized if considered in isolation, but taken together they provide one of the most complete and convincing demonstrations of implicit perception.

Further valuable evidence for implicit perception comes from fMRI studies of emotionally valenced faces (Whalen et al., 1998). Implicitly perceived fearful faces produce amygdala activation, and a subtraction analysis revealed no additional activation of visual cortex relative to happy faces, suggesting the possibility that fearful faces are processed automatically via a non-cortical route. Of course, this subtraction does not eliminate the possibility of cortical activation; both happy and fearful faces could produce visual cortex activation, and the subtraction just reveals the lack of additional cortical processing of fearful faces. Even so, the fact that amygdala activation was greater for fearful faces in the absence of greater activation of visual cortex is suggestive of an alternative, non-cortical source of the activation. Together these results provide converging support implicit perception.

A Skeptic's Interpretation

Although these investigations provide some of the strongest evidence for implicit perception, and despite the advantages of using sensitive neuroimaging measures, all of them adopt the dissociation paradigm without fully meeting the exhaustiveness assumption for each subject (Hannula et al., 2005). Most of these studies find diminished responses to less visible stimuli, raising the possibility that the effect results from residual explicit processing rather than from a different mechanism altogether. That is, these findings are consistent with a failure to meet the exhaustiveness assumption. Moreover, neural activation might not be as sensitive a measure as we assume. Perhaps a sizable amount of conscious processing is necessary to produce robust neural activation and to produce the distributed processing that is typically attributed to consciousness. If so, then "implicit" stimuli may have been fleetingly or weakly perceived, and the amount of conscious information available might not be enough to drive robust neural activation. This distinction might account for qualitative differences in the pattern of activation for identified and unidentified words (Stenberg et al., 2000). The unidentified words might have received an insufficient amount of conscious processing to produce the pattern typically associated with full awareness; however, that qualitative difference does not imply the absence of explicit processing. Implicitly and explicitly perceived stimuli may produce qualitatively different patterns of activation only because the implicit stimuli received less conscious processing (not no conscious processing).

The strongest evidence reviewed here is the N400 effect for unseen stimuli. This series of studies represents the most careful and systematic exploration of implicit perception that we are aware of in any of the studies discussed in this chapter. The studies carefully segregated aware and unaware trials on the basis of both subjective (word identification) and objective (6-alternative forced-choice [6AFC] decisions) measures and then examined the N400 for both correct and incorrect/absent responses. Although explicit sensitivity (measured using d′ for a binary category decision task) was effectively nil for mistaken responses in the 6AFC task, it was reliably above chance for the word identification task (in Experiments 2 and 3). Thus, the identification task clearly was not exhaustive. The 6AFC task comes closer, but a skeptic could quibble with several of the procedures in this study. First, the mean d′ was often greater than 0, and some subjects had d′ values greater than 0.5 on the binary choice. Although the regression method revealed an intercept significantly greater than 0, suggesting implicit processing even when d′ was extrapolated to 0, the fact that many observers had greater than nil sensitivity raises concerns that a few of the subjects might partially drive the results. A better approach would be to set the stimulus characteristics separately for each subject such that d′ is as close as possible to 0 on the explicit task. Another concern is that the task used to measure d′ was a binary category judgment (in the category vs. not in the category). This task might not be as sensitive as a presence/absence judgment, raising the possibility that a more sensitive measure might reveal some explicit processing even when observers show no sensitivity in the category judgment. These critiques aside, this study represents one of the greatest challenges to a skeptic because it uses multiple explicit measures and a sensitive imaging measure to examine implicit processing.

Evidence for Implicit Perception – Patient Data

Studies of brain-damaged patients provide some of the most compelling evidence for implicit perception. In fact, some have noted the surprising acceptance of evidence for implicit perception in brain-damaged subjects even by researchers who reject similar methods in the study of unimpaired subjects (Merikle & Reingold, 1992). In part, this

P1: JzG
0521857430c09 CUFX049/Zelazo 0 521 85743 0 printer: cupusbw November 6, 2006 15:48

acceptance of evidence from patient populations derives from the belief that brain damage can entirely disrupt some aspects of conscious perception or memory. If so, then the brain damage may provide the most effective elimination of explicit perception, much more so than simply reducing the visibility of a stimulus via masking. In unimpaired populations, the mechanisms for conscious perception are potentially available, leaving the persistent concern that any evidence for implicit perception might derive from explicit contamination. However, if the mechanisms themselves are eliminated by brain damage, then any residual processing must be attributed to implicit processes.

The challenge for researchers wishing to provide evidence for implicit perception is different for patient studies. Rather than trying to show that a particular task rules out the use of explicit perception, researchers must demonstrate that the patient entirely lacks the capacity for explicit processing in any task. Given that most such "natural experiments" are inherently messy, with some spared abilities intermixed with impairments, conclusions from patient studies depend on a systematic exploration of the nature and extent of the deficit in processing. In many cases, such studies require the same level of empirical precision necessary in behavioral studies, but they are further hampered by the limited subject population.

In this section, we consider three different sorts of evidence for a distinction between implicit and explicit processing. In two of these cases, conclusions rely heavily on the data of a relatively small number of patients. First, we consider the implications of studies of DF who is a visual form agnosic. We then consider two different classes of brain-damage phenomena, each of which has led to striking findings of preserved processing in the absence of awareness: blindsight and visual neglect.

DF and the Two Pathways Argument

The visual form agnosic patient DF (Goodale & Milner, 1992; Milner & Goodale, 1995) acquired her deficit from bilateral damage to portions of extrastriate visual cortex in the ventral visual processing stream. Although she can perceive and discriminate surface features such as color and texture, she shows a strikingly impaired ability to visually discriminate figural properties of objects, such as form, size, and orientation. Her preserved haptic and auditory discrimination of objects reveals preserved general knowledge and object recognition abilities; her deficit is one of visual object perception. Despite her inability to recognize objects visually, she can use the visual structure of objects to guide her motor responses. For example, she shows normal performance when trying to insert a slate into a slot, using the proper orientation and directed movement even though she cannot report the orientation of the slate in the absence of a motor interaction (Goodale, Milner, Jakobson, & Carey, 1991). Furthermore, she cannot report the orientations of blocks placed on tables, but can still reach out and pick up the blocks with appropriate grip aperture and limb movements (Jakobson & Goodale, 1991).

These results countermand the intuition that perception produces a unitary representation of the world, that interactions with the visual world should rely on the same representations and mechanisms as visual interpretation of the world. The dissociation in DF's ability to interpret and act on the world provides evidence for two distinct mechanisms to process visual information. One system involves the phenomenal recognition of parts of the visual world, and the other, operating without our awareness of the identities of objects, allows us to act on the world. In other words, this system seems to allow guided motor responses to objects even if we are unaware of what those objects might be.

The case of DF does not provide evidence for implicit perception in the same sense discussed throughout the rest of this chapter; at some level, DF is aware of the existence of the object even if she cannot name it. However, the case has some interesting parallels, and it reveals the importance of looking for qualitative differences in performance. One

P1: JzG
0521857430c09 CUFX049/Zelazo 0 521 85743 0 printer: cupusbw November 6, 2006 15:48

obvious parallel is that a visual stimulus can elicit an appropriate action or response even if some aspects of it are unavailable to consciousness. Visual analysis of an object does not guarantee conscious perception of its properties. More importantly, some aspects of visual processing occur outside of what can be consciously reported. The case of DF differs from other studies of implicit perception in that her spared abilities do not involve the processing of symbolic representations outside of awareness. Rather, she can engage in actions toward objects without needing to use a symbolic representation or any recognition-based processes. Most studies of implicit perception focus on whether or not implicit symbol manipulation or representation is possible (Dulany, 2004).

Blindsight

Neurologists had long speculated that some visual functioning might persist even in patients blinded by cortical damage (see Teuber, Battersby, & Bender, 1960 for a review), but the phenomenology of "blindsight" was not convincingly demonstrated until the 1970s (Pöppel, Held, & Frost, 1973; Weiskrantz, Warrington, Sanders, & Marshal, 1974). Patients suffering damage to primary visual cortex (V1) experience a visual scotoma; they fail to consciously perceive objects that fall into the affected portion of their visual field. They do not perceive a black hole or an empty space. Rather, the missing region of the visual field simply does not reach awareness, much as neurologically intact individuals do not normally notice their blind spot when one eye is closed. Blindsight refers to the finding that some cortically blind patients show evidence of perception in their damaged field in the absence of awareness. In essence, such patient evidence constitutes an application of the dissociation logic; the patient reports no awareness of the stimulus but still shows some effect of it. In a classic study of blindsight (Weiskrantz et al., 1974), lights were flashed in the damaged visual field of patient DB. Although DB reported no awareness of the lights, he could point out the location of the light more accurately than would be expected by chance. This finding suggests that V1 contributes to visual awareness, because in its absence, patients do not consciously experience visual stimuli. Perhaps the most established explanation for blindsight posits two routes to visual perception: (a) a pathway via V1 that leads to conscious awareness and (b) a more primitive pathway bypassing V1, perhaps via the superior colliculus. The latter route presumably allows perception in the absence of awareness. Indeed, in animals, cells in MT specialized for the detection of motion continue to respond normally to moving stimuli in the scotoma (Rosa, Tweedale, & Elston, 2000).

The two-routes hypothesis provides a strong claim about the nature of implicit perception, with one route operating outside awareness and the other generating awareness. Over the past 20 years, this hypothesis has faced a number of challenges designed to undermine the claim that conscious perception is entirely absent in blindsight. In other words, these alternative explanations question the exhaustiveness of the measure of conscious awareness, which in this case is the subjective report of the subject. For example, damage to V1 might be incomplete, with islands of spared cortex that function normally, thereby allowing degraded visual experience in small portions of the scotoma region (Fendrich, Wessinger, & Gazzaniga, 1992, 1993; Gazzaniga, Fendrich, & Wessinger, 1994; Wessinger, Fendrich, & Gazzaniga, 1997). Brain imaging of blindsight patients has returned mixed results: at least one patient (CLT) showed a small region of metabolically active visual cortex (Fendrich et al., 1993), whereas other researchers found no evidence for intact visual cortex in structural scans of other blindsight patients (e.g., Trevethan & Sahraie, 2003; Weiskrantz, 2002). Moreover, lesions of V1 in animals produce blindsight-like behavior even though these controlled lesions likely are complete (e.g., Cowey & Stoerig, 1995). Another alternative is that neurologically spared regions surrounding the scotoma receive differential sensory input as a result of the presence of an item in the blind region, thereby allowing better-than-chance guessing (Campion, Latto, & Smith, 1983). For example, a light source in the blind field might also generate some visual input for regions outside the blind field via scattering of light, thereby indicating the presence of something unseen (see the commentary in Campion et al., 1983 for a discussion). A final challenge comes from the argument that blindsight itself might indicate a change in response criterion rather than a change in awareness or sensitivity per se (Azzopardi & Cowey, 1998). This challenge is based on the idea that subjective reports on single trials do not fully measure awareness and that a signal detection approach is needed to verify that the response criterion cannot entirely account for the results. We address this final alternative in more detail here because it is the most theoretically relevant to the topic of this chapter.

Most evidence for blindsight comes from a comparison of performance on two tasks: a presence/absence judgment (direct measure of awareness) and a forced- choice task (indirect measure of perception), and most data are reported in terms of percentage correct (Azzopardi & Cowey, 1998). Yet, the use of percent correct to compare performance in these two tasks could well lead to spurious dissociations between implicit and explicit perception because percent correct measures are affected by response biases (Campion et al., 1983). For example, subjects tend to adopt a fairly conservative response criterion (responding only when certain) when asked to make a presence/absence judgment about a near-threshold stimulus. Furthermore, subjects may well vary their criterion from trial to trial. In contrast, when subjects are forced to choose between two alternative stimuli or to pick which temporal interval contained a stimulus, response bias is less of a concern; subjects have to choose one of the two stimuli. Thus, direct comparisons of forced choice and presence/absence judgments pit a potentially biased measure against an unbiased one.

To examine the possibility of bias, a frequently tested blindsight patient's (GY) sensitivity to the presence of stimuli was measured with d′ (or da where appropriate) along with his response criterion for a variety of tasks often used to study blindsight (Azzopardi & Cowey, 1998). As expected, responding was unbiased in a forced-choice task. In contrast, response criterion in a presence/absence judgment was fairly conservative ($c = 1.867$), and interestingly, it was substantially reduced by instructing GY to guess when unsure ($c = .228$). These findings reveal the danger of relying on percent correct as a primary measure of blindsight; with sensitivity set to $d' = 1.5$, these levels of bias elicit 75% correct responding for a forced-choice task, but 55% performance for a presence/absence judgment. In fact, any $d' > 1$ would lead to an apparent dissociation in percentage correct, but the result could entirely be attributed to response criterion rather than differential sensitivity. Using this signal detection approach, GY showed greater sensitivity for static displays in forced-choice responses than in presence/absence responses, but the same did not hold for moving displays. Thus, evidence for "blindsight" to motion stimuli in which patients report no awareness (presence/absence) but still show accurate forced-choice performance might result entirely from shifts in response criterion. These results underscore the danger of relying on percent correct scores in investigations employing blindsight patients and highlight the benefits of using bias-free tasks. To date, relatively few investigations of blindsight have adopted these important methodological changes, despite an active literature that possesses surprising scope given the apparent rarity of blindsight patients.

Inferences from one recent study are less subject to the criterion problem (Marcel, 1998). Two patients (TP and GY) completed a series of tasks that required forced-choice judgments, and only some showed evidence of implicit perception. For example, neither showed much priming from single letters presented to their blind field when the task was to pick the matching letter. Also, neither was more likely to select a synonym of a word presented to the blind field. However, when defining a polysemous

P1: JzG
0521857430c09 CUFX049/Zelazo 0 521 85743 0 printer: cupusbw November 6, 2006 15:48

word, both showed priming in their choice of a definition when a word presented to the blind field disambiguated the meaning. Given that these tasks all involve forced-choice decisions, differences between them are unlikely to result from response biases. Interestingly, the finding that the least direct measure shows an effect implies that semantic concepts are activated without activating the representation of the word itself.

One intriguing finding is that some blindsight patients apparently consciously experience afterimages of stimuli presented to their blind field (Marcel, 1998; Weiskrantz, 2002). Such afterimages, if frequently experienced by blindsight patients, might explain some residual perception in the damaged field. Interestingly, the afterimages can arise after information from the blind and sighted fields have been combined. When different colored filters were used for the blind and sighted field, patient DB experienced an afterimage that was specific to the combination of those two colors, suggesting that information from the blind field was processed beyond the point required to resolve binocular differences (Weiskrantz, 2002).

The phenomenon of blindsight represents one of the most striking demonstrations of non-conscious perception. It provides potentially important insights into the need for V1 in order to consciously perceive our environment. However, the approaches typically used to study blindsight are subject to methodological critiques because they often do not account for response biases in the measurement of awareness. Perhaps more importantly, most such studies are couched in the dissociation framework, inferring implicit perception based on the absence of direct evidence for conscious perception. Consequently, blindsight findings are subject to many of the same objections raised for behavioral work on implicit perception.

Parietal Neglect

Visual neglect involves deficient awareness of objects in the contralesional visual field, typically resulting from damage to the posterior inferior parietal lobe in the right hemisphere, secondary to middle cerebral artery infarction. Although both blindsight and neglect are associated with spared processing in the absence of awareness, neglect is characterized as an attentional (rather than sensory) deficit, and commonly occurs in the absence of a visual scotoma (or blind spot). The damage in neglect occurs later in the perceptual processing stream than it does in cases of blindsight, raising the possibility that neglected stimuli might be processed semantically to a greater extent as well (see Driver & Mattingley, 1998 for a review). In many patients, the failure to notice or attend to contralesional stimuli is exacerbated when stimuli are presented simultaneously to both left and right visual fields, presumably because these stimuli compete for attention (visual extinction, as described earlier). For example, neglect patients might fail to eat food on the left side of their plate. Some patients fail to dress the left side of their body or to brush the left side of their hair. Although neglect can affect other processing modalities (e.g., haptic and auditory processing), we limit our discussion to visual neglect.

Evidence for preserved processing of neglected visual stimuli takes several forms: (a) successful same/different discrimination of bilaterally presented stimuli despite a failure to report the contralesional stimulus, (b) intact lexical and semantic processing of extinguished stimuli, and (c) activation of responses consistent with an extinguished prime. Here we review evidence from each of these areas, and we also describe experiments designed to test the claim that extinguished stimuli are not consciously perceived.

Many studies demonstrate the preserved ability to discriminate identical pairs of items from those that are physically or categorically dissimilar (Berti et al., 1992; Verfaellie, Milberg, McGlinchey-Berroth, & Grande, 1995; Volpe, Ledoux, & Gazzaniga, 1979). Typically, two pictures or words are briefly presented to the right and left of fixation, and patients judge whether they are the same (have the same name) or are different. Both patients and intact control subjects perform this task better than chance even when the object orientations differ or when they are two different exemplars of the same category (e.g., two different cameras). Further, patients can reliably indicate that physically similar (and semantically related) items are in fact different from one another. In all cases, patients report little or no awareness of the stimulus in the contralesional visual field.

These findings suggest that neglect patients can process extinguished stimuli semantically and that their representation of these unseen stimuli is fairly complex and complete. The dissociation between the naming and matching indicates that visual processing of extinguished stimuli proceeds relatively normally despite the absence of awareness (but see Farah, Monheit, & Wallace, 1991). However, more concrete evidence for the absence of explicit awareness of extinguished stimuli is required for a clear conclusion in favor of implicit perception. This approach is logically equivalent to the dissociation paradigm; demonstrate that subjects cannot perceive a stimulus and then look for residual effects on performance. Subjects claim no conscious experience of extinguished stimuli but still are able to perform a fairly complex discrimination on the basis of the stimulus presentation. In fact, when patients were required to name both stimuli, they frequently named only the ipsilesional item. However, some of them felt that something had appeared on the contralesional side. None of these studies demonstrate that sensitivity to the presence of the extinguished stimulus is objectively no better than chance.

An alternative to examining preserved judgments about extinguished stimuli is to explore whether such stimuli produce lexical or semantic priming. To the extent that neglect is only a partial disruption of the ability to form representations, the inability to name extinguished stimuli might result from a failure to access existing representations rather than a failure to form a representation (McGlinchey-Berroth, Milberg, Verfaellie, Alexander, & Kilduff, 1993). Repetition priming tasks have been used to examine lexical, orthographical, or phonological priming by neglected stimuli (Schweinberger & Stief, 2001), and semantic priming studies have explored whether neglected primes receive more extensive cognitive processing (e.g., Ladavas, Paladini, & Cubelli, 1993; McGlinchey-Berroth et al., 1993). Primes are presented briefly in either the contralesional or ipsilesional visual field followed by a visible target stimulus, and priming is reflected in faster processing of the target stimulus. Lexical priming apparently survives visual neglect: Priming was evident for both patients and normal controls only when word stimuli were repeated and not when non-word stimuli were repeated, suggesting that the neglected word activated an existing representation (Schweinberger & Stief, 2001). Furthermore, the magnitude of priming was comparable in the contralesional and ipsilesional visual fields. In fact, left visual field priming was actually greater than that of normal controls for patients who neglected their left visual field. This counter-intuitive finding may result from a center-surround mechanism that increases activation for weakly accessible or subconscious visual stimulus while simultaneously inhibiting activation of other related items (see Carr & Dagenbach, 1990).

Similar claims have been made with respect to higher-level semantic processing of neglected visual stimuli. In one of these experiments (McGlinchey-Berroth et al., 1993), pictures, used as prime stimuli, were presented peripherally in the left or the right visual field, and filler items (a meaningless visual stimulus made up of components of the target items) were presented on the side opposite. After 200 ms, the pictures were replaced by a central target letter string, and subjects indicated whether or not the string was a word. Semantic priming should lead to faster lexical decisions if the prime pictures were related to the word. Although patients responded more slowly than controls, they showed semantic priming even though they could not identify the prime pictures in a 2-alternative forced-choice task (McGlinchey-Berroth et al., 1993). Other semantic priming tasks have found faster processing of a right-lateralized word following a left-lateralized prime word (Ladavas et al., 1993). Given that the patient in this study was unable to read a single word presented in the left visual field, and performed no better than chance with lexical decision, semantic discrimination, and stimulus detection tasks even without a bilateral presentation, the prime presumably was not consciously perceived.

However, none of the studies discussed thus far provided an exhaustive test for conscious awareness of left-lateralized stimuli, leaving open the possibility that residual awareness of the "neglected" stimulus might account for preserved priming effects. More generally, the use of different paradigms or stimuli in tests of awareness and measures of priming does not allow a full assessment of awareness during the priming task; measures of awareness may not generalize to the experiment itself. Many of the priming studies also introduce a delay interval between the prime and target, leaving open the possibility that patients shift their attention to the extinguished stimulus in advance of target presentation (see Cohen, Ivry, Rafal, & Kohn, 1995).

Other studies of implicit perception in visual neglect have adopted a response compatibility approach in which a central target item is flanked by an irrelevant item on either the left or right side of fixation (Cohen et al., 1995). These flanker items were either compatible, incompatible, or neutral with respect to the response required for the target. Interestingly, responses to the target were slower when the flanker was incompatible, even when it was presented to the contralesional visual field. In a control experiment using the same materials, the patients were asked to respond to the flankers and were given unlimited time to respond. Responses were reliably slower and more error prone for stimuli presented to the contralesional visual field. This finding confirms the impairment of processing of stimuli in the contralesional visual field, but it also undermines the response compatibility results as a demonstration of implicit perception. That subjects could, when instructed, direct their attention to the "neglected" contralesional stimulus implies that they might have had some residual awareness of flankers presented to the contralesional visual field. The patients in this experiment also had more diffuse damage than is typical in neglect experiments, and one had left-lateralized damage. The diffuse damage might affect performance on the flanker task for reasons other than hemispatial neglect.

Most studies of implicit perception by neglect patients have focused on determining the richness of processing of neglected stimuli, but relatively few studies have focused on producing convincing demonstrations that neglected stimuli truly escape conscious awareness. One recent study adopted the process dissociation procedure in an attempt to provide a more thorough demonstration that processing of neglected stimuli is truly implicit (Esterman et al., 2002). A critical picture appeared in one visual field and a meaningless filler picture appeared in the other. After a 400-ms delay, a two-letter word stem appeared in the center of the screen, and subjects were either instructed to complete the stem with the name of the critical picture (inclusion) or to complete it with any word other than the picture name (exclusion). Relative to normal control subjects, patients were less likely to complete word stems with picture names in the inclusion task, particularly when the picture was presented in the neglected visual field. In contrast, patients were more likely than controls to complete the stems with picture names in the exclusion task when the picture was presented in the neglected field. Moreover, they completed such stems with the picture name more frequently than in a baseline condition with no pictures. If patients had explicitly perceived the stimulus, they would not have used it to complete the word stem. However, they still processed it enough that it influenced their stem completion. Although this study provides clearer evidence for implicit perception of neglected stimuli, the methods are subject to the same critiques discussed in our review of behavioral

P1: JzG
0521857430c09 CUFX049/Zelazo 0 521 85743 0 printer: cupusbw November 6, 2006 15:48

evidence using the exclusion paradigm. Clearly, more systematic assessments of explicit perception of neglected words are needed before unequivocal claims about implicit perception in neglect are possible.

A Believer's Interpretation

Perhaps more interesting than the evidence itself is the face validity of evidence for implicit perception in patient populations. Neglect and blindsight illustrate the serious behavioral ramifications of the absence of awareness, and their normal behaviors are in essence a constant, real-world version of the process dissociation paradigm. Such patients' daily actions reflect their lack of awareness of some aspects of their visual world, and if they had awareness of those aspects, they would perform differently. Evidence for perception despite the absence of awareness in these patients is particularly convincing because the absence of explicit awareness is their primary deficit. In combination with behavioral and neuroimaging evidence, these data confirm that implicit perception is possible in the absence of awareness.

A Skeptic's Interpretation

Studies of patient populations rely extensively on the logic of the dissociation paradigm; patients lack awareness of parts of their visual world, so any residual processing of information in those areas must reflect implicit perception. Unfortunately, few studies exhaustively eliminate the possibility of partial explicit processing of visual information in the face of these deficits. Nobody doubts that explicit awareness is affected in both blindsight and neglect. These deficits of awareness have clear behavioral consequences. However, impaired awareness does not mean absent awareness. One of the few studies of blindsight to measure awareness using signal detection theory found that many of the most robust findings supporting implicit perception could be attributed to bias rather than residual sensitivity (Azzopardi & Cowey, 1998). This study reveals the danger in relying solely on subjective reports of awareness rather than on systematic measurement of performance. Most patient studies use subjects' ability to report their visual experience as the primary measure of explicit awareness, with implicit perception inferred from any spared processing in the "blind" field. Such subjective reports in the context of the dissociation paradigm do not provide an adequately exhaustive measure of explicit perception. Consequently, performance on indirect measures might reflect residual explicit perception rather than implicit perception.

What do Dissociations in Perception Mean?

Despite protestations to the contrary, the century-old debate over the mere existence of implicit perception continues to this day. The techniques and tools have improved, but the theoretical arguments are surprisingly consistent. In essence, believers argue that the converging evidence provides overwhelming support for the existence of implicit perception, whereas skeptics argue that almost all findings of implicit perception fail to provide adequate controls for explicit contamination. As with most such debates, different conclusions can be drawn from the same data. Believers can point to improved methodologies that provide more sensitive measures of implicit processing or that more effectively control for explicit processing. Skeptics can point to the fact that none of these controls are airtight and that the effects, when present, tend to be small. Believers point to converging evidence from patients and from imaging studies with neurologically intact individuals, and skeptics point to the even greater inadequacies of the controls for explicit processing in those domains. The conclusions drawn from these data are colored by assumptions about the parsimony of each conclusion. Believers find the conclusion in favor of implicit processing more parsimonious because a variety of critiques, some fairly convoluted, are needed to account for all of the converging support for implicit processing. Skeptics find conclusions in favor of implicit processing unappetizing because they often posit the existence of additional mechanisms when all of the data can potentially be explained using solely explicit processing.

More recent behavioral techniques have made progress toward eliminating the more obvious objections of skeptics. Qualitative differences, signal detection measures of sensitivity, and regression techniques are appropriate first steps toward overcoming the critiques of the dissociation paradigm, although a staunch critic might never be satisfied. From this possibly irresolvable debate, perhaps some additional insights can be gleaned. Regardless of whether or not implicit perception exists, what can we learn about perceptual processing from attempts to reveal implicit perception? What do dissociations in perception, whether between implicit and explicit or entirely within explicit processing, tell us about the nature of awareness and about the nature of perception? What do these dissociations mean for our understanding of perception?

For the moment, let's assume that implicit perception exists. If it does exist, what does it do? Does it play a functional role in the survival of the perceiver? Our perceptual systems exist to extract information from the world to allow effective behavior. The world, then, presents a challenge to a perceptual system: The information available far exceeds our ability to consciously encode and retain it. Our perceptual systems evolved to extract order and systematicity from the available data, to encode those aspects of the world that are relevant for the behavioral demands of our ecological niche (Gibson, 1966). Our perceptual systems adapted to extract signal from the noise, allowing us to survive regardless of whether we are aware of the variables that influence our behavior.

We are only aware of a subset of our world at any time, and we are, of course, unaware of those aspects that fail to reach awareness. Just as a refrigerator light is always on whenever we look inside, any time we examine the outputs of perceptual processing, we are aware of those outputs. Someone unfamiliar with the workings of a refrigerator might assume that the light is on when the door is closed. Similarly, given that the only information available to consciousness is that information that has reached awareness, an intuitive inference would be to assume that all processing involves awareness – we never "see" evidence of processing without awareness. Claims that mental processes happen outside of consciously mediated operations run counter to this intuitive belief. That same belief might also underlie the willingness to accept subjective reports as adequate measures of conscious processing (see the subjective vs. objective threshold discussion above). The goal of the implicit perception literature is to determine whether or not perceptual processing occurs with the metaphorical refrigerator light out.

A fundamental issue in the implicit perception literature concerns the similarities of the types of processing attributed to implicit and explicit mechanisms. Is there a commonality to those operations that apply with and without awareness? If implicit perception exists, do the implicit mechanisms apply to everything outside of awareness equally, or is there some selectivity? Does the spotlight of consciousness perpetually move about, randomly illuminating implicitly processed information, thereby bringing it to awareness? Or are there fundamental differences between implicit and explicit processes? This question is, in many respects, more interesting and important than the question of whether or not implicit perception exists at all. Implicit perception would lack its popular appeal and broad implications if it produced nothing more than a weak version of the same processing that would occur with awareness. The broad popular appeal (or fear) of the notion of implicit processing is that it could, under the right circumstances, lead to behaviors different from those we would choose with awareness.

One way to conceptualize the implicit/explicit distinction is to map it onto the intentional/automatic dichotomy. Consciousness presumably underlies intentional actions (Searle, 1992), those in which perceivers can perform new operations, computations, or symbol manipulations on the information in the world. Automatic behaviors, in contrast, reproduce old operations, computations, or symbol manipulations, repeating processes that were effective in the past in the absence of intentional control. Automatic computations occur in a data-driven, possibly encapsulated fashion (Fodor, 1986). Much of the evidence for perception without awareness is based on this sort of data-driven processing that could potentially affect explicit processing, but that occurs entirely without conscious control. Previous exposure to a stimulus might lead to more automatic processing of it the next time. Or, if the implicit and explicit processing mechanisms overlap substantially, a prior exposure might provide metaphorical grease for the gears, increasing the likelihood that bottom-up processing will lead to explicit awareness.

The vast majority of evidence for implicit perception takes this form, producing behavioral responses or neural activation patterns that mirror those we might expect from explicit processing. Behavioral response compatibility findings are a nice example of this form of implicit perception: The interference shown in response to an implicitly perceived prime is comparable to that we might expect from a consciously perceived prime (except for NCE effects). Similarly, much of the fMRI evidence for implicit perception shows that activation from unseen stimuli mirrors the pattern of activation that would occur with awareness.

These findings are not as theoretically interesting as cases in which the outcome of implicit perception differs from what we would expect with conscious perception – cases in which implicit and explicit processes lead to qualitatively different outcomes. When the results are the same for implicit and explicit processing, the standard skeptical critiques weigh heavily. The failure to meet the assumptions of the dissociation paradigm adequately leaves open the possibility that "implicit" behaviors result from explicit processing. In contrast, qualitative differences are theoretically significant regardless of whether or not they reflect a difference between implicit and explicit processing. Take, for example, the case of DF (Goodale & Milner, 1992). She can accurately put an oriented card through a slot, but lacks conscious access to the orientation of the card. Although she cannot subjectively perceive the orientation, other mechanisms allow her to access that information and to use it in behavior. In other words, the dissociation implies the operation of two different processes. Moreover, the finding is significant even if her accurate behavior involves some degree of explicit awareness. The dissociation reveals the operation of two different processes and different uses of the same visual information. If the behavior happened to result entirely from implicit perception, that would be interesting as well, but it is not as important as the finding that two different processes are involved. Similarly, evidence for amygdala activation from unreported fearful faces is interesting not because the faces cannot be reported but because it suggests a possible alternative route from visual information to neural activation (Le Doux, 1996). Subsequent control experiments might show that the unreported faces were explicitly perceivable. However, the more interesting question is whether a subcortical route exists, not whether that subcortical route operates entirely without awareness. Of course, to the extent that inferences about alternative processing mechanisms depend on the complete absence of awareness, these findings will always be open to critique. Although evidence from the process dissociation paradigm is subject to shifts in bias and motivation (Visser & Merikle, 1999), the underlying goal of that paradigm has at its base the demonstration of a qualitative difference in performance. Whether or not this difference reflects the distinction between implicit and explicit processing or between two forms of explicit processing is of secondary importance.

In sum, the evidence for implicit perception continues to be mixed and likely will remain that way in spite of improved tools and methods. A diehard skeptic likely will be able to generate some alternative, however implausible, in which explicit processing alone can explain a dissociation between implicit and explicit perception. Similarly, believers are unlikely to accept the skeptic's discomfort with individual results, relying instead on the convergence of a large body of evidence. Methodological improvements might well force the skeptic to adopt more convoluted explanations for the effects, but are unlikely to eliminate those explanations altogether.

Here we propose a somewhat different focus for efforts to explore perception with and without awareness. Rather than trying to eliminate all aspects of explicit perception, research should focus instead on demonstrating differences in the perceptual mechanisms that vary as a function of manipulating awareness. Qualitative differences in perception are interesting regardless of whether they reflect purely implicit perception. Differences between performance when most explicit processes are eliminated and performance when all explicit processes can be brought to bear are interesting in their own right and worthy of further study.

References

Abrams, R. L., & Greenwald, A. G. (2000). Parts outweigh the whole (word) in unconscious analysis of meaning. *Psychological Science*, *11*(2), 118–124.

Abrams, R. L., Klinger, M. R., & Greenwald, A. G. (2002). Subliminal words activate semantic categories (not automated motor responses). *Psychonomic Bulletin & Review*, 9(1), 100–106.

Adolphs, R., Tranel, D., Damasio, H., & Damasio, A. R. (1995). Fear and the human amygdala. *Journal of Neuroscience*, *15*, 5879–5891.

Azzopardi, P., & Cowey, A. (1998). Blindsight and visual awareness. *Consciousness & Cognition*, 7(3), 292–311.

Baars, B. J. (1988). *A cognitive theory of consciousness*. New York: Cambridge University Press.

Bar, M., & Biederman, I. (1998). Subliminal visual priming. *Psychological Science*, 9(6), 464–469.

Bernat, E., Bunce, S., & Shevrin, H. (2001). Event-related brain potentials differentiate positive and negative mood adjectives during both supraliminal and subliminal visual processing. *International Journal of Psychophysiology*, *42*, 11–34.

Bernat, E., Shevrin, H., & Snodgrass, M. (2001). Subliminal visual oddball stimuli evoke a P300 component. *Clinical Neurophysiology*, *112*(1), 159–171.

Berti, A., Allport, A., Driver, J., Dienes, Z., Oxbury, J., & Oxbury, S. (1992). Levels of processing for visual stimuli in an "extinguished" field. *Neuropsychologia*, *30*, 403–415.

Brazdil, M., Rektor, I., Dufek, M., Jurak, P., & Daniel, P. (1998). Effect of subthreshold target stimuli on event-related potentials. *Electroencephalography & Clinical Neurophysiology*, *107*, 64–68.

Breiter, H. C., Etcoff, N. L., Whalen, P. J., Kennedy, W. A., Rauch, S. L., Buckner, R. L., et al. (1996). Response and habituation of the human amygdala during visual processing of facial expression. *Neuron*, *17*, 875–887.

Brown, C., & Hagoort, P. (1993). The processing nature of the N400: Evidence from masked priming. *Journal of Cognitive Neuroscience*, *5*, 34–44.

Brown, M., & Besner, D. (2002). Semantic priming: On the role of awareness in visual word recognition in the absence of an expectancy. *Consciousness and Cognition*, *11*(3), 402–422.

Campion, J., Latto, R., & Smith, Y. M. (1983). Is blindsight an effect of scattered light, spared cortex, and near-threshold vision? *Behavioral & Brain Sciences*, 6(3), 423–486.

Carr, T. H., & Dagenbach, D. (1990). Semantic priming and repetition priming from masked words: Evidence for a center-surround attentional mechanism in perceptual recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *16*, 341–350.

Cheesman, J., & Merikle, P. M. (1984). Priming with and without awareness. *Perception & Psychophysics*, 36(4), 387–395.

Cherry, E. C. (1953). Some experiments upon the recognition of speech, with one and with two ears. *Journal of the Acoustical Society of America*, 25, 975–979.

Cohen, A., Ivry, R. B., Rafal, R. D., & Kohn, C. (1995). Activating response codes by stimuli in the neglected visual field. *Neuropsychology, 9*, 165–173.

Coles, M. G., Gratton, G., & Donchin, E. (1988). Detecting early communication: Using measures of movement-related potentials to illuminate human information processing. *Biological Psychology, 26*, 69–89.

Corteen, R. S., & Dunn, D. (1974). Shock-associated words in a nonattended message: A test for momentary awareness. *Journal of Experimental Psychology, 102*(6), 1143–1144.

Corteen, R. S., & Wood, B. (1972). Autonomic responses to shock-associated words in an unattended channel. *Journal of Experimental Psychology, 94*(3), 308–313.

Cowey, A., & Stoerig, P. (1995). Blindsight in monkeys. *Nature, 373*, 247–249.

Damian, M. (2001). Congruity effects evoked by subliminally presented primes: Automaticity rather than semantic processing. *Journal of Experimental Psychology: Human Perception and Performance, 27*(1), 154–165.

Daza, M. T., Ortells, J. J., & Fox, E. (2002). Perception without awareness: Further evidence from a Stroop priming task. *Perception & Psychophysics, 64*(8), 1316–1324.

Deacon, D., Hewitt, S., Yang, C., & Nagata, M. (2000). Event-related potential indices of semantic priming using masked and unmasked words: Evidence that N400 does not reflect a post-lexical process. *Cognitive Brain Research, 9*, 137–146.

Deacon, D., & Shelley-Tremblay, J. (2000). How automatically is meaning accessed: A review of the effects of attention on semantic processing. *Frontiers in Bioscience, 5*, E82–94.

Debner, J., & Jacoby, L. (1994). Unconscious perception: Attention, awareness and control. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*(2), 304–317.

Dehaene, S., Naccache, L., Cohen, L., Le, B., D., Mangin, J., Poline, J., et al. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience, 4*(7), 752–758.

Dehaene, S., Naccache, L., Le Clech, G., Koechlin, E., Mueller, M., Dehaene-Lambertz, G., et al. (1998). Imaging unconscious semantic priming. *Nature, 395*(6702), 597–600.

Dehaene, S., Piazza, M., Pinel, P., & Cohen, L. (2003). Three parietal circuits for number processing. *Cognitive Neuropsychology, 20*(3–6), 487–506.

Dell'Acqua, R., & Grainger, J. (1999). Unconscious semantic priming from pictures. *Cognition, 73*(1), B1–B15.

Deutsch, J. A., & Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review, 70*(1), 80–90.

Devrim, M., Demiralp, T., & Kurt, A. (1997). The effects of subthreshold visual stimulation on P300 response. *Neuroreport, 8*, 3113–3117.

Dosher, B. A. (1998). The response-window regression method – some problematic assumptions – comment on Draine and Greenwald (1998). *Journal of Experimental Psychology: General, 127*(3), 311–317.

Draine, S. C., & Greenwald, A. G. (1998). Replicable unconscious semantic priming. *Journal of Experimental Psychology: General, 127*(3), 286–303.

Driver, J., & Mattingley, J. B. (1998). Parietal neglect and visual awareness. *Nature Neuroscience, 1*, 17–22.

Driver, J., & Vuilleumier, P. (2001). Perceptual awareness and its loss in unilateral neglect and extinction. *Cognition, 79*(1–2), 39–88.

Driver, J., Vuilleumier, P., Eimer, M., & Rees, G. (2001). Functional magnetic resonance imaging and evoked potential correlates of conscious and unconscious vision in parietal extinction patients. *Neuroimage, 14*, S68–75.

Dulany, D. E. (2004). Higher order representation in a mentalistic metatheory. In R. J. Gennaro (Ed.), *Higher order theories of consciousness*. Amsterdam: John Benjamins.

Eimer, M. (1999). Facilitatory and inhibitory effects of masked prime stimuli on motor activation and behavioural performance. *Acta Psychologica, 101*(2–3), 293–313.

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

Eimer, M., & Schlaghecken, F. (2002). Links between conscious awareness and response inhibition: Evidence from masked priming. *Psychonomic Bulletin & Review, 9*(3), 514–520.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature, 392*, 598–601.

Esterman, M., McGlinchey-Berroth, R., Verfaellie, M., Grande, L., Kilduff, P., & Milberg, W.

(2002). Aware and unaware perception in hemispatial neglect: Evidence from a stem completion priming task. *Cortex*, *38*(2), 233–246.

Farah, M. J., Monheit, M. A., & Wallace, M. A. (1991). Unconscious perception of "extinguished" visual stimuli: Reassessing the evidence. *Neuropsychologia*, *29*, 949–958.

Fendrich, R., Wessinger, C. M., & Gazzaniga, M. S. (1992). Residual vision in a scotoma: Implications for blindsight. *Science*, *258*, 1489–1491.

Fendrich, R., Wessinger, C. M., & Gazzaniga, M. S. (1993). Sources of blindsight – Reply to Stoerig and Weiskrantz. *Science*, *261*, 493–495.

Fodor, J. (1986). Modularity of mind. In Z. W. P. a. W. Demopoulos (Ed.), *Meaning and cognitive structure: Issues in the computational theory of mind* (pp. 129–167). Norwood, NJ: Ablex.

Freud, S. (1966). *Introductory lectures on psychoanalysis* (J. Strachey, Trans.). New York: W. W. Norton.

Gazzaniga, M. S., Fendrich, R., & Wessinger, C. M. (1994). Blindsight reconsidered. *Current Directions in Psychological Science*, *3*(3), 93–96.

Gibson, J. J. (1966). *The senses considered as perceptual systems*. Boston: Houghton Mifflin.

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, *15*(1), 20–25.

Goodale, M. A., Milner, A. D., Jakobson, L. S., & Carey, D. P. (1991). A neurological dissociation between perceiving objects and grasping them. *Nature*, *349*, 154–156.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.

Greenwald, A. G. (1992). New Look 3: Unconscious cognition reclaimed. *American Psychologist*, *47*, 766–779.

Greenwald, A., Abrams, R., Naccache, L., & Dehaene, S. (2003). Long-term semantic memory versus contextual memory in unconscious number processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(2), 235–247.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, *102*(4–27).

Greenwald, A. G., Draine, S. C., & Abrams, R. L. (1996). Three cognitive markers of unconscious semantic activation. *Science*, *273*(5282), 1699–1702.

Greenwald, A. G., Schuh, E. S., & Klinger, M. R. (1995). Activation by marginally perceptible (subliminal) stimuli: Dissociation of unconscious from conscious cognition. *Journal of Experimental Psychology: General*, *124*(1), 22–42.

Greenwald, A. G., Spangenberg, E. R., Pratkanis, A. R., & Eskenazi, J. (1991). Double-blind tests of subliminal self-help audiotapes. *Psychological Science*, *2*(119–122).

Hannula, D. E., Simons, D. J., & Cohen, N. J. (2005). Imaging implicit perception: Promise and pitfalls. *Nature Reviews Neuroscience*, *6*, 247–255.

Holender, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *Behavioral and Brain Sciences*, *9*, 1–66.

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic and intentional uses of memory. *Journal of Memory & Language*, *30*, 513–541.

Jakobson, L. S., & Goodale, M. A. (1991). Factors affecting higher-order movement planning: A kinematic analysis of human prehension. *Experimental Brain Research*, *8*, 199–208.

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.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, *17*, 4302–4311.

Kiefer, M. (2002). The N400 is modulated by unconsciously perceived masked words: Further evidence for an automatic spreading activation account of N400 priming effects. *Cognitive Brain Research*, *13*(1), 27–39.

Kiefer, M., & Spitzer, M. (2000). Time course of conscious and unconscious semantic brain activation. *Neuroreport*, *11*(11), 2401–2407.

Klapp, S., & Hinkley, L. (2002). The negative compatibility effect: Unconscious inhibition influences reaction time and response selection. *Journal of Experimental Psychology: General*, *131*(2), 255–269.

Klinger, M., Burton, P., & Pitts, G. (2000). Mechanisms of unconscious priming: I. Response competition, not spreading activation. *Journal* *of Experimental Psychology: Learning, Memory, and Cognition*, 26(2), 441–455.

Koechlin, E., Naccache, L., Block, E., & Dehaene, S. (1999). Primed numbers: Exploring the modularity of numerical representations with masked and unmasked semantic priming. *Journal of Experimental Psychology: Human Perception & Performance*, 25(6), 1882–1905.

Kunst-Wilson, W. R., & Zajonc, R. B. (1980). Affective discrimination of stimuli that cannot be recognized. *Science*, 207, 557–558.

Kutas, M., & Hillyard, S. A. (1980). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science*, 207, 203–205.

Ladavas, E., Paladini, R., & Cubelli, R. (1993). Implicit associative priming in a patient with left visual neglect. *Neuropsychologia*, 31, 1307–1320.

Lane, R. D., & Nadel, L. (2000). *Cognitive neuroscience of emotion*. New York: Oxford University Press.

Le Doux, J. E. (1996). *The emotional brain*. New York: Simon & Shuster.

Levelt, W. J. M. (1965). *On binocular rivalry*. Assen, The Netherlands: Royal VanGorcum.

Levin, D. T., Momen, N., Drivdahl, S. B., & Simons, D. J. (2000). Change blindness blindness: The metacognitive error of overestimating change-detection ability. *Visual Cognition*, 7, 397–412.

Mack, A., & Rock, I. (1998). *Inattentional blindness*. Cambridge, MA: MIT Press.

Macmillan, N. A. (1986). The psychophysics of subliminal perception. *Behavioral and Brain Sciences*, 9, 38–39.

Marcel, A. J. (1983a). Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes. *Cognitive Psychology*, 15, 238–300.

Marcel, A. J. (1983b). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology*, 15, 197–237.

Marcel, A. J. (1998). Blindsight and shape perception: Deficit of visual consciousness or of visual function? *Brain*, 121(8), 1565–1588.

McCarthy, J. C., Puce, A., Gore, J. C., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. *Journal of Cognitive Neuroscience*, 9, 604–609.

McCormick, P. (1997). Orienting attention without awareness. *Journal of Experimental Psychology: Human Perception & Performance*, 23(1), 168–180.

McGlinchey-Berroth, R., Milberg, W. P., Verfaellie, M., Alexander, M., & Kilduff, P. T. (1993). Semantic processing in the neglected visual field: Evidence from a lexical decision task. *Cognitive Neuropsychology*, 10, 79–108.

Merikle, P. (1992). Perception without awareness. Critical issues. *American Psychologist*, 47(6), 766–779.

Merikle, P. M. (1994). On the futility of attempting to demonstrate null awareness. *Behavioral & Brain Sciences*, 17(3), 412.

Merikle, P., & Joordens, S. (1997a). Parallels between perception without attention and perception without awareness. *Consciousness & Cognition*, 6(2–3), 219–236.

Merikle, P. M., & Joordens, S. (1997b). Measuring unconscious influences. In J. W. Schooler (Ed.), *Scientific approaches to consciousness* (pp. 109–123). Mahwah, NJ: Erlbaum.

Merikle, P. M., & Reingold, E. M. (1991). Comparing direct (explicit) and indirect (implicit) measures to study unconscious memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(2), 224–233.

Merikle, P. M., & Reingold, E. M. (1992). Measuring unconscious perceptual processes. In T. S. Pittman (Ed.), *Perception without awareness: Cognitive, clinical, and social perspectives* (pp. 55–80). New York: Guilford Press.

Merikle, P., & Reingold, E. (1998). On demonstrating unconscious perception – Comment on Draine and Greenwald (1998). *Journal of Experimental Psychology: General*, 127(3), 304–310.

Merikle, P., Smilek, D., & Eastwood, J. (2001). Perception without awareness: Perspectives from cognitive psychology. *Cognition*, 79(1–2), 115–134.

Milner, A. D., & Goodale, M. A. (1995). *The visual brain in action*. New York: Oxford University Press.

Mitroff, S. R., Simons, D. J., & Franconeri, S. L. (2002). The siren song of implicit change detection. *Journal of Experimental Psychology: Human Perception & Performance*, 28(4), 798–815.

Moray, N. (1959). Attention in dichotic listening: Affective cues and the influence of instructions. *Quarterly Journal of Experimental Psychology*, 11, 56–60.

Morris, J. S., Frith, C. D., Perrett, D. I., Rowland, D., Young, A. W., Calder, A. J., et al. (1996). A differential neural response in the human amygdala to fearful and happy facial expressions. *Nature*, 383, 812–815.

Most, S. B., Simons, D. J., Scholl, B. J., Jimenez, R., Clifford, E., & Chabris, C. F. (2001). How not to be seen: The contribution of similarity and selective ignoring to sustained inattentional blindness. *Psychological Science*, *12* (1), 9–17.

Naccache, L., Blandin, E., & Dehaene, S. (2002). Unconscious masked priming depends on temporal attention. *Psychological Science*, *13* (5), 416–424.

Naccache, L., & Dehaene, S. (2001a). The priming method: Imaging unconscious repetition priming reveals an abstract representation of number in the parietal lobes. *Cerebral Cortex*, *11*(10), 966–974.

Naccache, L., & Dehaene, S. (2001b). Unconscious semantic priming extends to novel unseen stimuli. *Cognition*, *80*, 223–237.

Neumann, O., & Klotz, W. (1994). Motor responses to nonreportable, masked stimuli: Where is the limit of direct paramater specification? In M. Moscovitch (Ed.), *Attention and performance* (Vol. XV, pp. 123–150). Cambridge, MA: MIT Press.

Pöppel, E., Held, R., & Frost, D. (1973). Residual visual function after brain wounds involving the central visual pathways in man. *Nature*, 243, 295–296.

Pratkanis, A. R. (1992). The cargo-cult science of subliminal persuasion. *Skeptical Inquirer*, *16*, 260–272.

Pratkanis, A., Eskenazi, J., & Greenwald, A. (1994). What you expect is what you believe (but not necessarily what you get) – A test of the effectiveness of subliminal self-help audiotapes. *Basic & Applied Social Psychology*, *15* (3), 251–276.

Rees, G. (2001). Seeing is not perceiving. *Nature Neuroscience*, 4(7), 678–680.

Rees, G., Kreiman, G., & Koch, C. (2002). Neural correlates of consciousness in humans. *Nature Reviews Neuroscience*, 3, 261–270.

Rees, G., Wojciulik, E., Clarke, K., Husain, M., Frith, C., & Driver, J. (2000). Unconscious activation of visual cortex in the damaged right hemisphere of a parietal patient with extinction. *Brain*, *123* (8), 1624–1633.

Reingold, E., & Merikle, P. (1988). Using direct and indirect measures to study perception without awareness. *Perception & Psychophysics*, *44*(6), 563–757.

Rosa, M. G. P., Tweedale, R., & Elston, G. N. (2000). Visual responses of neurons in the middle temporal area of new world monkeys after lesions of striate cortex. *Journal of Neuroscience*, 20(14), 5552–5563.

Schacter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 501–518.

Schacter, D. L., Alpert, N. M., Savage, C. R., Rauch, S. L., & Albert, M. S. (1996). Conscious recollection and the human hippocampal formation: Evidence from positron emission tomography. *Proceedings of the National Academy of Sciences*, 93 (1), 321–325.

Schacter, D. L., & Buckner, R. L. (1998). On the relations among priming, conscious recollection, and intentional retrieval: Evidence from neuroimaging research. *Neurobiology of Learning & Memory*, 70(1–2), 284–303.

Schweinberger, S., & Stief, V. (2001). Implicit perception in patients with visual neglect: Lexical specificity in repetition priming. *Neuropsychologia*, 39(4), 420–429.

Searle, J. R. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT Press.

Shevrin, H. (2001). Event-related markers of unconscious processes. *International Journal of Psychophysiology*, 42, 209–218.

Sidis, B. (1898). *The psychology of suggestion*. New York: D. Appleton.

Simons, D. J. (2000). Attentional capture and inattentional blindness. *Trends in Cognitive Sciences*, 4(4), 147–155.

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28, 1059–1074.

Squire, L. R., Ojemann, J. G., Miezin, F. M., Petersen, S. E., Videen, T. O., & Raichle, M. E. (1992). Activation of the hippocampus in normal humans: A functional anatomical study of memory. *Proceedings of the National Academy of Sciences*, 89(5), 1837–1841.

Stenberg, G., Lindgren, M., Johansson, M., Olsson, A., & Rosen, I. (2000). Semantic processing without conscious identification: Evidence from event-related potentials. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 26(4), 973–1004.

Teuber, H. L., Battersby, W. S., & Bender, M. B. (1960). *Visual field defects after penetrating* *missile wounds of the brain*. Cambridge, MA: Harvard University Press.

Tong, F. (2001). Competing theories of binocular rivalry: A possible resolution. *Brain and Mind*, *2*, 55–83.

Tong, F. (2003). Primary visual cortex and visual awareness. *Nature Reviews Neuroscience*, *4*, 219–229.

Tong, F., Nakayama, K., Vaughan, J. T., & Kanwisher, N. (1998). Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron*, *21*, 753–759.

Treisman, A. M. (1960). Contextual cues in selective listening. *Quarterly Journal of Experimental Psychology*, *12*, 242–248.

Treisman, A. (1964). Monitoring and storage of irrelevant messages in selective attention. *Journal of Verbal Learning and Verbal Behavior*, *3*, 449–459.

Trevethan, C. T., & Sahraie, A. (2003). Spatial and temporal processing in a subject with cortical blindness following occipital surgery. *Neuropsychologia*, *41*(10), 1296–1306.

Verfaellie, M., Milberg, W. P., McGlinchey-Berroth, R., & Grande, L. (1995). Comparison of cross-field matching and forced-choice identification in hemispatial neglect. *Neuropsychology*, 9, 427–434.

Visser, T., & Merikle, P. (1999). Conscious and unconscious processes: The effects of motivation. *Consciousness & Cognition*, 8(1), 94–113.

Volpe, B. T., Ledoux, J. E., & Gazzaniga, M. S. (1979). Information processing of visual stimuli in an "extinguished" field. *Nature*, 282, 722–724.

Vuilleumier, P., Armony, J., Clarke, K., Husain, M., Driver, J., & Dolan, R. (2002). Neural response to emotional faces with and without awareness: Event-related fMRI in a parietal patient with visual extinction and spatial neglect. *Neuropsychologia*, *40*(12), 2156–2166.

Vuilleumier, P., Sagiv, N., Hazeltine, E., Poldrack, R. A., Swick, D., Rafal, R. D., et al. (2001). Neural fate of seen and unseen faces in visuospatial neglect: A combined event-related functional MRI and event-related potential study. *Proceedings of the National Academy of Sciences*, 98(6), 3495–3500.

Watanabe, T., Nanez, J., & Sasaki, Y. (2001). Perceptual learning without perception. *Nature*, *413*(6858), 844–848.

Weiskrantz, L. (2002). Prime-sight and blindsight. *Consciousness & Cognition*, *11*(4), 568–581.

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

Wessinger, C. M., Fendrich, R., & Gazzaniga, M. S. (1997). Islands of residual vision in hemianopic patients. *Journal of Cognitive Neuroscience*, 9(2), 203–221.

Whalen, P. J., Rauch, S. L., Etcoff, N. L., McInerney, S. C., Lee, M. B., & Jenike, M. A. (1998). Masked presentations of emotional facial expressions modulate amygdala activity without explicit awareness. *Journal of Neuroscience*, *18*, 411–418.

Wheatstone, C. (1838). Contributions to the physiology of vision – Part the first. On some remarkable and hitherto unobserved phenomena of binocular vision. *Philosophical Transactions of the Royal Society of London*, *128*, 371–394.

Wolfe, J. M. (1999). Inattentional amnesia. In V. Coltheart (Ed.), *Fleeting memories: Cognition of brief visual stimuli* (pp. 71–94). Cambridge, MA: MIT Press.