

Judging Words at Face Value: Interference in a Word Processing Task Reveals Automatic Processing of Affective Facial Expressions

Georg Stenberg, Susanne Wiking, and Mats Dahl

University of Lund, Sweden

Earlier research has indicated that some characteristics of facial expressions may be automatically processed. This study investigated automaticity as evidenced by involuntary interference in a word evaluation task. Compound stimuli, consisting of words superimposed on pictures of affective faces, were presented to subjects who were given the task of evaluating the affective valence of the words while disregarding the faces. Results of three experiments showed that word evaluation was influenced by the concurrently shown affective faces. Overall, negative words were found to require longer latencies, indicating that more processing resources are invested in negative than in positive stimuli. This speed advantage for positive words was modified by the faces. Negative words were facilitated, relative to positive ones, when shown with a negative expression (e.g. a sad face). Correspondingly, negative words were inhibited, relative to positive ones, when shown with a positive expression (e.g. a happy face). The results are consistent with automatic, involuntary semantic processing of affective facial expressions.

INTRODUCTION

The face is an extraordinary visual stimulus, paramount in its social importance. Psychological theories from Darwin onward have given the processing of facial expressions a special place, supported by a genetically prepared capacity. Recent research has shown unusually rapid and efficient detection of some expressions, such as angry faces, with obvious relevance for fear and threat (Hansen & Hansen, 1988, 1994). Öhman (1986) has

Requests for reprints should be sent to Dr Georg Stenberg, Department of Psychology, Paradisgatan 5 P, S-223 50 Lund, Sweden; e-mail: Georg.Stenberg@psykol.lu.se.

This study was financed by a grant to Georg Stenberg from the Swedish Council for Research in the Humanities and the Social Sciences (HSFR). We are grateful for helpful comments from Craig A. Smith, Arne Öhman, and two anonymous reviewers on an earlier version. Peter Meyerhöffer helped run Experiment 1.

suggested that the processing of affective expressions, especially threatening ones, is preattentive, and in support of this idea facial expressions have been shown to have effects even in the absence of the subject's awareness of the eliciting stimulus (Esteves, Dimberg, & Öhman, 1994a; Esteves, Parra, Dimberg, & Öhman, 1994b; Murphy, Monahan, & Zajonc, 1995; Murphy & Zajonc, 1993). Findings such as these strongly suggest that affective facial expressions enjoy an unusual advantage, compared to other visual stimuli of similar complexity, in gaining access to automatic processing. The present study aims to examine further the extent to which affective facial expressions are automatically processed. It does so by examining involuntary interference effects of face processing on other ongoing cognitive activity, such as semantic processing of words.

Face Processing

The case for automaticity in the processing of some facial expressions, but not others, was made by Hansen and Hansen in an influential study (1988), which investigated a perceptual pop-out effect for deviant expressions in arrays of faces. Angry faces were found among distractors with an ease suggesting parallel search of the stimulus array, whereas the time to find a happy face was a function of array size, indicative of serial and controlled search. The authors have later (Hansen & Hansen, 1994) pointed out methodological problems in the earlier study, and now support their position with other data, especially the power of angry faces to attract attention quickly, as indicated by EOG saccade latency.

Further evidence for automatic processing of facial expressions, and a special status for angry faces, comes from a classical conditioning paradigm (Esteves et al., 1994a, b). Effective aversive conditioning can take place with angry (but not happy) faces as conditioned stimuli, even though the faces are exposed very briefly and prevented from reaching conscious awareness by backward masking (Esteves et al., 1994b). Similarly, after learning, conditioned skin conductance responses can be elicited without awareness by masked angry faces (Esteves et al., 1994a).

The benefits of preattentive detection of threat have been discussed in an evolutionary context by Öhman (1993). A low threshold for discovering threat may have been established over the course of evolution, where environmental pressures have made false negatives carry more devastating consequences than false positives. In addition, threat stimuli may have acquired the capacity to gain privileged access to further strategic processing after initial detection.

Automaticity in the access to facial expressions is also suggested by the way in which the perceiver mimics the perceived expression, typically without intention and awareness (Dimberg & Öhman, 1996). The very

short latency with which this muscle activity is initiated (300–400msec) leaves little room for controlled processing.

The idea that semantic characteristics of facial expressions are made available to the observer automatically and preattentively was carried further by White (1995), who examined valence (happy or sad) as a characteristic of schematic faces. Effects of valence on the intercept of search functions (i.e. on overall differences in the speed of processing happy and sad expressions irrespective of array size), suggested that faces were effortlessly categorised as to valence in an early stage of processing.

Affective priming with facial expressions has been demonstrated (Murphy et al., 1995; Murphy & Zajonc, 1993). In these studies, very brief (4msec) exposures were found to affect ratings of Chinese ideographs shown immediately after the primes. Longer exposures (1sec), on the other hand, had no effect. The authors interpret these findings as a result of the instantaneous accessibility of affective valence, which reigns unopposed in early processing, but may be overruled by slower cognitive assessments in later stages.

Automaticity and Interference Tasks

The notion of automatic processing has several connotations, perhaps the most basic of which is involuntariness. Thus, an automatic process runs its course to completion once it is started, in much the same sense as a ballistic projectile in mid-flight is beyond modification of its course. Several other meanings of automaticity, which may or may not be co-extensive with the involuntariness criterion, have attracted interest in studies of controlled versus automatic processing (discussed in Kahnemann & Treisman, 1984). Among them are the absence of demands on central, limited resources and the absence of interference with other concurrent processes. However, the sense in which we intend to investigate the automaticity of affective face processing is the sense of involuntariness. In this respect, we follow a tradition from explorations of the Stroop effect and the Eriksen's flanker task (Eriksen & Eriksen, 1974; Stroop, 1935). In both of these, diversionary stimuli are presented (words in the Stroop task, flanking letters in the Eriksen task) which distract the subject from the task at hand (naming the colour, reacting to the central letter). It is in the best interest of the subject to avoid dealing with the distractors. Failures to reject the distractors are nevertheless consistently observed by way of their interference effects on reaction times. Automaticity is deduced from the involuntary way in which task-irrelevant stimuli are processed to the detriment of performance on the given task.

In analogy to the flanker and Stroop tasks, this study uses pictures of faces, each surrounding a centrally placed word, as distracting stimuli. In

each trial, one face is shown, with a word superimposed on the lower nose area. The task is to respond to the word only. We purport to show that the face is given automatic processing on a semantic level, as evidenced by its side-effects on the processing of the words. By the same token as in the Stroop and flanker tasks, we expect relative facilitation if the distracting stimulus leads to the same response as the task-relevant one, and we expect relative inhibition if it leads to a competing response.

Other studies have examined interference phenomena between pictures and words. In his review of the Stroop effect, MacLeod (1991) characterised the picture-word task as an analogue of the Stroop task. In it, a picture, spatially adjacent to a word, is to be named, and inhibition (or facilitation) effects on naming latency are usually found to indicate that the word is involuntarily processed. In Smith and Magee's study (1980), both semantic classification and naming tasks were given, and the pattern of results indicated that word reading was immune to interference from pictures, but semantic classification of the words was not. Conversely, semantic classification of pictures was not disturbed by words, but picture naming showed interference. A parsimonious interpretation of these effects involves the assumptions that pictures have privileged access to semantic memory, whereas words have privileged access to a lexicon, devoid of semantic features. This interpretation has been developed and tested by Glaser and Glaser (1989).

The Glaser and Glaser model has been extended by De Houwer and Hermans (1994), who applied it to such semantic attributes as affective valence. They examined the task of speeded good/bad evaluation, performed on pictures with word distractors and words with picture distractors. In accordance with their expectations, derived from the Glaser and Glaser model, pictures interfered with word evaluation, but not vice versa.

Attitude Activation

The view that evaluations are associatively connected with object representations in memory has been proposed by Fazio and his colleagues and corroborated in a series of investigations. If associative strength is high, attitudes can be activated automatically. Fazio, Sanbonmatsu, Powell, and Kardes (1986) found that words connected with highly accessible attitudes could prime the evaluation of ensuing words. Priming was automatic, because it obtained only when stimulus onset asynchrony between prime and target was short (300msec), and not when it was long enough (1000msec) to allow conscious expectancies to develop. Similarly, Roskos-Ewoldsen and Fazio (1992) found that objects toward which subjects held highly accessible attitudes attracted attention in a visual display. This allocation of attention was against the best interest of the subject in

performing the task, and the activation of the attitude can therefore be said to be automatic by the criterion of involuntariness. Although Fazio and co-workers reserved the label of automaticity for attitudes which are strongly held and primarily idiosyncratic, others (Bargh, Chaiken, Gollwitzer, & Pratto, 1992) have argued that automatic attitude activation has a broader range of applicability, excluding only the most weakly held attitudes.

Other evidence indicates that there may be an asymmetry between positive and negative attitude objects in their power to attract and hold attention automatically. Pratto and John (1991) used positively and negatively valenced trait adjectives in a Stroop task, and found that the negative words delayed colour naming more than positive ones. The result is compatible with the hypothesis of automatic vigilance (i.e. the view that attention is drawn preferentially to negative information). Asymmetries between positive and negative events have been reviewed by Taylor (1991), who concluded, based on a wide range of data, that negative events tend to elicit more cognitive (and other) activity in an initial phase, although the long-term effects are damped down in a characteristic minimisation pattern. In individual-differences studies of the attentional probe paradigm, anxious subjects have been found to give more extensive processing and a larger share of attentional resources to negative stimuli (MacLeod, Mathews, & Tara, 1986; Mathews & MacLeod, 1994).

Thus, an objective for the present investigation will be to assess possible differences in the speed of processing of differently valenced objects. In terms of experimental design, our primary hypothesis has been cast as an interaction effect between word valence and facial affect. Additionally, we will examine the main effects for possible valence differences.

In so doing, we need to take into account a methodological caveat, which concerns reaction time differences between positive and negative words. Word frequency exerts an obvious influence on latencies, in for example, the lexical decision task, more frequent words being processed more quickly. The frequency effect can be confounded with valence, because there is a pervasive tendency for positive words to be used more frequently than negative ones, as described by Zajonc (1968). We have tried to eliminate this confound by statistical means (Experiments 1 and 2) and by matching the selections of positive and negative words for frequency (Experiment 3).

Objectives

To summarise, we intend to present compound stimuli, composed of words superimposed on an affective face. The task will be to categorise the word as good or bad, while ostensibly disregarding the face. This task was chosen because the evaluative dimension has been found to be a

fundamental one in factor analyses of the semantic content of words (Osgood, Suci, & Tannenbaum, 1957). Also, the evaluative dimension is one which can easily be applied to affective facial expressions.¹

Our expectations, based on our own pilot work and on findings in the literature (Pratto & John, 1991; Taylor, 1991), are to find latency differences between positive and negative words in this task. We hypothesise that positive words will be at an advantage with respect to speed, because negative words will tend to evoke more extended processing, leading to slower categorisation. This Positive Valence Advantage (PVA) is expected to figure as a main effect of word valence in the analyses, but we also anticipate that it will be importantly modified by concurrently shown faces. If a happy expression is shown, then the processing of positive words will be facilitated, even more so than in a neutral condition, and the Positive Valence Advantage will be enlarged. Correspondingly, negatively valenced expressions (e.g. sadness), will tend to diminish or reverse the PVA. The way these effects are achieved is by adding activation to—or lowering thresholds for—processing in the pathways congruent with the facial expression. This is in essence the way semantic priming is often construed, although in contrast to typical priming studies, we show prime and target together.

Our principal dependent variable will be reaction time, but we will also examine the pattern of errors. Although the task is not difficult, occasional errors can be expected to be dictated by the valence of the face, when this is in conflict with the valence of the word. Conceptualised in the terms of Signal Detection Theory, this corresponds to an adjustment of the criterion. As the parameter to measure this adjustment, both bias, beta, and criterion, *C*, could in principle have been used, but Snodgrass and Corwin (1988) have given convincing reasons for preferring *C*, and we have followed their recommendation.

EXPERIMENT 1

In the first experiment, we examined the effects of angry, happy, and neutral facial expressions on performance in the word evaluation task. A happy expression is arguably the only example of a facial gesture at once elementary and unequivocally positive. Several expressions are, on the other hand, negative. We chose anger to represent the negative pole, because much important theorising has concerned this emotion and its

¹ Furthermore, it is a task which can be accomplished with reasonable speed. In an experiment with a more difficult task (i.e. deciding whether a word was emotional in meaning or neutral), we found reaction times to be slow and highly variable. The expected effects on reaction times were therefore not obtained, although there was a predicted effect on errors.

place in evolution (Dimberg & Öhman, 1996; Hansen & Hansen, 1994). We expected to see angry faces facilitate negative judgements relative to positive ones, and happy faces facilitate positive judgements relative to negative ones. We also expected these modulations to be superimposed upon an overall speed advantage for positive judgements.

Method

Subjects. Twenty women and thirty men, all students at the University of Lund, participated on a volunteer basis without any compensation. The median age was 26, with a range of 20–34.

Materials and Procedure. The pictures of facial expressions were taken from Ekman and Friesen's *Unmasking the face* (1975). Five pictures showing an angry expression, five with a happy expression, and three with a neutral expression were selected for the experiment. They were digitised and saved as grey-scale picture files.

In the experiment, the pictures were shown on a monitor in 4-bit monochrome mode (i.e. set to display 16 shades of grey). The pictures occupied an area 6.3cm wide and 9.0cm high, centred on the screen. The words were displayed in black 24-point text on a grey background in a rectangle, the size of which was adjusted for the length of the word, typically about 2×1 cm. The positioning of the word rectangle was such as to occupy an area across the lower part of the nose, chosen to be central in the face without hiding any of the facial features essential for the affective expression. Instructions were to judge and react on the displayed word only and to disregard the face.

The 120 words selected for the experiment were mainly, but not exclusively, Swedish translations of words from the Toglia–Battig corpus (Toglia & Battig, 1978). This material contains a large number of words classified along several semantic dimensions, among them the pleasant-unpleasant dimension. Words were chosen from the word clusters at the extreme poles of this dimension. The frequencies of the words in written Swedish were checked in Allén (1971). Mean frequencies were 60.7 per million (SD 70.6) for positive words, and 22.4 per million (SD 38.7) for negative words. Because the frequency variable is typically very skewed, a logarithmic transformation [$\log(\text{frequency}+1)$], is used later in computations where word frequency serves as a covariate.

Three lists were constructed by randomly pairing words with pictures, subject to the constraint that the six combinations of word and picture valences each be represented by 20 items in the list. For each subject, one of the three lists was randomly selected at the time of the experiment.

Subjects were tested individually in a small room, where a Macintosh LC II computer presented the stimuli and recorded the responses. The subject was instructed to classify each word as positive or negative in meaning by pressing one of two response keys on the keyboard. The pictures were to be ignored, and the responses were to be given as rapidly as possible without sacrificing accuracy. A practice sequence of 12 trials preceded the main sequence of 120 trials.

In each trial, the picture was first transferred to the screen, and, after a delay of 250msec, the word, on its grey background rectangle, was displayed across the lower nose area. The compound picture-word stimulus remained on screen for 1500msec. Then the screen was erased to a middle-grey background colour, and remained blank for another 1500msec. Responses were recorded in the 2sec interval following the presentation of the word.

Results

Reaction Times. Mean RTs for correct responses to the six picture-word combinations are shown in Fig. 1. They were entered into a two-way repeated measures ANOVA, using Word valence (2 levels) and Facial expression (3 levels) as factors. There was a marked effect of word valence, reflecting the fact that positive words were processed faster than negative ones [$F(1,49) = 34.31$; $MS_e = 4049.38$; $P < .001$]. The Positive Valence Advantage was sizeable: 719msec for positive words versus 762msec for negative ones. There was also a main effect of face type. On average, angry faces induced slower responding (753msec) than neutral and happy faces (733 and 734msec) [$F(2,98) = 5.70$; $MS_e = 2050.48$; $P = .005$]. The critical prediction concerned the interaction of the factors. It showed a reliable effect [$F(2,98) = 4.93$; $MS_e = 1971.5$; $P = .012$]. (Reported probabilities, here and in the following, have been adjusted according to the Huynh-Feldt procedure to compensate for deviations from the sphericity assumption.)

A set of planned contrasts was computed to elucidate the nature of the interaction. These were directed at the question whether the difference between RTs for negative and positive words (the PVA) was altered by the affective conditions relative to the neutral one. Also, a comparison was made of the size of the PVA between the two affective conditions. In total, three contrasts were computed, each having (1,98) degrees of freedom and $MS_e = 1971.5$.

First, a comparison between the Happy and the Neutral condition showed that happy faces increased the PVA from 21 to 60msec ($F = 9.49$; $P = .004$). As can be seen in Fig. 1, the results for the Happy facial

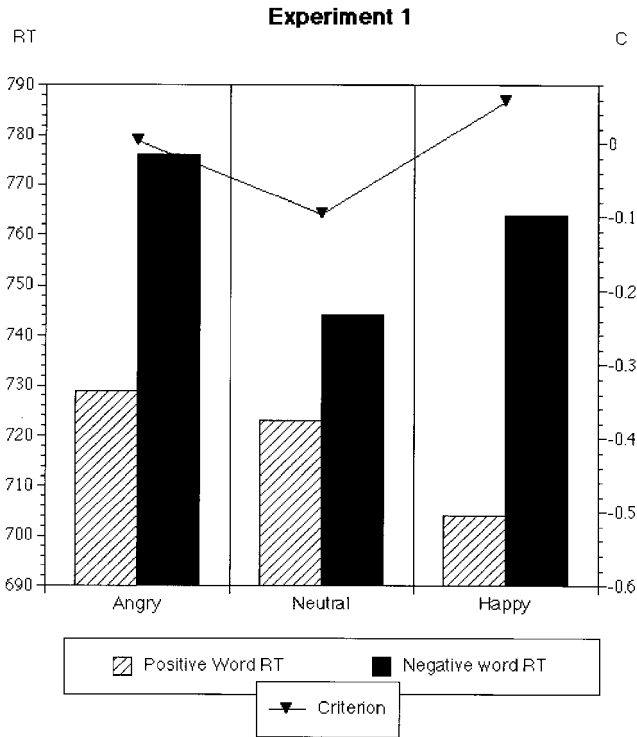


FIG. 1. Reaction times for positive and negative words, shown with angry, neutral, and happy faces in Experiment 1. The superimposed line diagram shows the criterion, *C* (scale on right *y*-axis), induced by the three types of faces. Positive values of *C* indicate a bias towards evaluating words as positive, negative values indicate the opposite bias.

expression are in line with predictions. The processing of positive words was speeded up, and the processing of negative words was slowed down in comparison to the neutral condition, the combined effect being a substantial augmentation of the PVA.

For Angry facial expressions, the results ran counter to our predictions. The PVA was altered by angry faces in relation to neutral ones ($F = 4.27$; $P = 0.046$), but not in the predicted direction. Instead, the PVA increased to 47msec from the neutral 21. As a result, the contrast between the two affective conditions did not reveal a reliable difference ($F = 1.03$; n.s.).

The word valence difference was followed up by a separate, item-wise test to determine whether it was caused by word frequency differences between the two word populations. In this item-based analysis, the average reaction time for each word was computed across all subjects and across all facial-affect combinations. In the one-way ANOVA of these RTs (two

groups of 60 words each), word frequency was entered as a covariate to assess whether there was an effect of valence independently of word frequency. The ANCOVA results showed that both effects were reliable [frequency: $F(1,117) = 8.97$; $MS_e = 2,887.0$; $P = .003$; valence: $F(1,117) = 6.85$; $MS_e = 2,887.0$; $P = .010$]. Thus, although positive words were more frequent than negative ones, and high-frequency words generally had shorter latencies than low-frequency ones, the fact that positive words had shorter latencies than negative ones could not be reduced to a side-effect of their high frequencies. There was an independent effect of valence on RT, and controlling for frequency only led to a reduction of the difference from 43 to 29msec.

Error Rates and Decision Criteria. The sometimes idiosyncratic nature of value judgements made it necessary to check whether there were large deviations from the intended, normative value categorisation of the words. The error rates indicated that the deviations were small (average error rate across all stimulus types: 3.4%). To assess whether, despite the overall conformity with normative evaluations, any stray items were particularly liable to be misclassified, an item-based analysis was performed. It showed only three items, out of 120, with error rates in excess of 12%, and none above 30%.

Based on the pattern of errors, the signal detection parameter C was computed as a measure of the criterion applied by the subject in making the response decision. A negative evaluation of the word was arbitrarily designated as a "no" response in signal detection terms, making negative values of C (a liberal criterion) denote a tendency toward negative evaluations, and positive values of C (a conservative criterion) denote a tendency toward positive evaluations. The mean C values for the three different facial expressions are shown in Fig. 1.

The visual impression gleaned from Fig. 1 is that the result pattern for C parallels that of the PVA, and statistical analysis bears this out. A one-way repeated-measures ANOVA showed a reliable difference in criterion placement as a function of facial expression [$F(2,98) = 6.18$; $MS_e = 0.050$; $P = .003$]. As in the RT analysis, a set of planned contrasts compared the two affective conditions with the neutral one, and with each other (for each contrast, $df = 1,98$ and $MS_e = 0.050$). As was the case for the RT effects, the Happy expression had an effect in accordance with our predictions ($F = 11.98$; $P < .001$). It shifted the criterion significantly toward a positive response bias. Whereas the Angry expression also had a significant effect ($F = 5.13$; $P = .026$), it ran counter to our prediction by also shifting the criterion toward a positive response bias. The Angry and Happy conditions did not reliably differ ($F = 1.43$; n.s.).

Discussion

In a general sense, support was provided for the hypothesis of task-irrelevant facial affect displays exerting an influence on a concurrently performed word evaluation task. The more specific hypotheses were, however, only partly confirmed. As to the effect of a happy expression, our expectations were fully supported. Happy faces slowed negative word processing, speeded positive word processing, and consequently increased the Positive Valence Advantage. Additionally, it shifted the decision criterion between the two in a positive direction. For angry expressions, on the other hand, the effect challenged our assumptions. Reaction times were overall slower, but especially so for negative words, resulting in an increased PVA. Moreover, the criterion was moved in a positive direction.

A factor of possible importance is the special signal value of angry facial expressions, discussed by both Öhman and Hansen and Hansen. Angry faces may tend always to attract attention, and hold it, once caught, which could explain why angry faces cause prolonged latencies for both word valences (Hansen & Hansen, 1994, p. 233). "Automatic vigilance may introduce effects at both the automatic and controlled level; once attention is attracted to a threatening face, it may be difficult to suppress its hold". However, to explain the present result, one would have to assume that this propensity for extended controlled processing interacts with the type of word being processed, resulting in an increased PVA.

Niedenthal, Setterlund, and Jones (1994) make a different, relevant point. They point out that it may be a mistake to arrange emotions along a single dimension of valence. In particular, the putatively negative emotions are very different from each other, and they may not all serve equally well as an evaluative opposite to happiness. For this reason, we repeated the experiment, but now with a different negative facial expression, namely sadness.

EXPERIMENT 2

This experiment aimed at reproducing the evaluation task of Experiment 1, with the exception that the facial expression Sad was substituted for the Angry expression of Experiment 1. We wanted to examine the possibility that sadness is closer to being the opposite of happiness in evaluative terms than is anger.

An additional test has been appended to the procedure, presented below as Experiment 2B. We test the memory retention of the stimuli used in the word evaluation task. Assuming that word-picture combinations are processed as a whole, incidental memory retention should show an advantage for words that are presented in the same facial-affect combination at test,

as compared to words presented in a changed face context. Thus, we expect a face-context effect to influence performance in the word recognition task, similarly to a context effect found by Niedenthal (1990). She used cartoon characters as the targets in a study-recognition paradigm, and preceded presentations of them with subliminally presented human face primes. Recognition was better when the targets were primed with the same facial expression at test as at study.

Furthermore, we expect words encountered in incongruent word-face combinations in the initial task to have been more elaborately processed (as evidenced by longer latencies), leading to better memory retention in accordance with the levels-of-processing findings in the memory literature. Therefore, we have a second hypothesis for the memory test, namely, better word recognition for words initially presented in incongruent word-face combinations.

Experiment 2A: Method

Materials and Procedure. The face pictures were again taken from the Ekman–Friesen set. The expressions represented were Happiness and Sadness, with six pictures selected from each category, using equal representation of male and female models. As neutral baseline pictures, we did not, however, reuse the (few) neutral expression pictures in the set. Instead, we constructed a set of six pseudoface pictures to provide us with a baseline condition. They were produced by a digital distortion of some of the original pictures, which preserved average brightness and general shape, but smeared features and texture into an unrecognisable blur, creating the impression of a very unfocused picture.

A set of words was constructed and pre-tested for use in this experiment. An original list of 288 words was designed to contain evaluatively positive, negative, and neutral words. The list was evaluated by 52 subjects, mostly psychology students, on a 7-point scale, where +3 meant very positive, and -3 very negative connotations. From this normatively evaluated corpus, a set of 200 words was selected, representing the 100 most positively and the 100 most negatively evaluated words. Out of this basis set, two different random selections were made, each of 60 positive and 60 negative words, to serve in the evaluation task. Due to partial overlap, the two selections together comprised 173 different words. The negative words had an average frequency of 18.23 per million (SD 31.46), and an average word length of 6.11 (SD 1.38) letters. For the positive words, average frequency was 34.19 per million (SD 53.93), and the average word length 6.79 (SD 1.29) letters. Ratings for the 88 negative words ranged from -2.98 to -1.35, and for the 85 positive words from 1.23 to 2.75. Six lists were prepared, with the

two selected sets in three different random orders. For each subject, one of these six lists was chosen. In each case, it comprised six combinations of word valence (2 levels) with face type (3 levels), each combination being represented by 20 stimuli, resulting in a list length of 120 trials.

Compared to the earlier experiment, the display procedure was altered in some respects, designed mainly to make it less likely that subjects were to analyse the picture-word combination into its separate parts. First, a random component was introduced in the positioning of the picture-word combination on the screen, making it impossible to fixate a certain position on the screen knowing that the word would be displayed there. A rectangularly distributed random term ranging from -64 to $+64$ pixels ($\pm 2.26\text{cm}$) was added to both vertical and horizontal coordinates before each stimulus display. The position of the word relative to the face was unaltered. Second, the display period was shortened to 750msec. Also, the lead time of 250msec, by which the completion of the picture drawing operation preceded the onset of the word in the earlier experiments, was now eliminated.

The period during which responses were accepted was 2000msec after word onset. A 1000msec display of a fixation cross on the centre of the screen preceded each stimulus. Assignment of "yes" or "no" responses to left- or right-arrow-keys was randomised individually. A practice block with auditory feedback for incorrect responses was given before the experiment.

Subjects. A total of 25 subjects, 12 men and 13 women, all students or employees at the University of Lund, participated without monetary reward. They had a median age of 22 years (range: 19–30).

Results

Reaction Times. Response latencies were analysed in a two-way repeated measures ANOVA, with Face type (3 levels: Neutral, Happy, and Sad), and Word valence (2 levels: Negative and Positive) as factors. The central hypothesis concerned the interaction, which proved to be reliable [$F(2,48) = 13.17$; $MS_e = 1,447.19$; $P < .001$]. As shown in Fig. 2, positive words were processed faster than negative ones in the Neutral condition (by 22msec), and this PVA was increased to 54msec when Happy faces were shown. The contrast was reliable ($F = 4.43$; $P = .041$) (this and the other contrasts within this interaction have $df = 1,48$ and $MS_e = 1,447.19$). With Sad faces, the PVA was reversed to -23msec . The contrast between Sad and Neutral was reliable ($F = 9.01$; $P = 0.004$), and so was the contrast between Sad and Happy ($F = 26.07$; $P < .001$).

Experiment 2

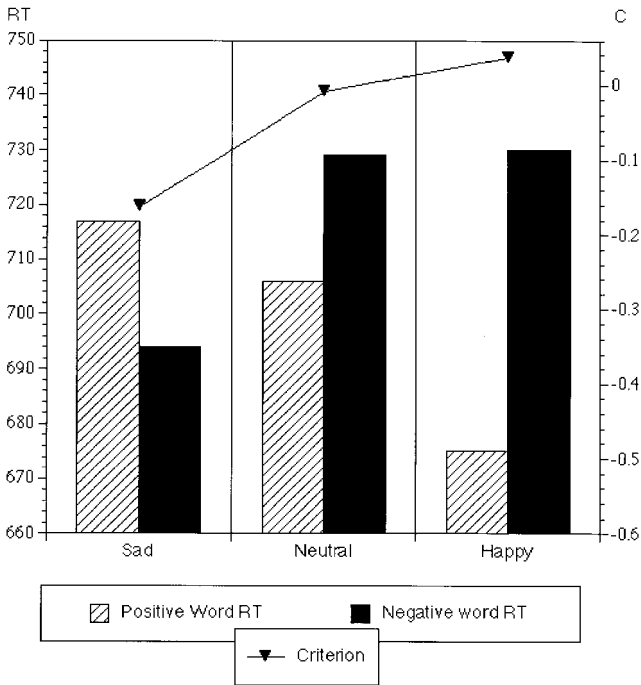


FIG. 2. Reaction times for positive and negative words, shown with sad, neutral, and happy faces in Experiment 2. Criterion values shown in the superimposed line diagram as in Fig. 1.

The main effects were not reliable (both P s > .10). However, to elucidate our earlier findings concerning the effect of Word valence on RT, an item-based analysis was also performed. The 173 words were subjected to an ANCOVA, using Word valence as the categorical and log frequency as the continuous variable. Because the Word valence effect turned out to be reliable (and frequency not reliable), Word rating was entered as a further factor to clarify the word value effect.

Depending on the *degree* of positive or negative valence, the words were classified into 3×2 groups, comprising very high, high, or moderate degrees of positive or negative valence. Cut-off points were ± 2.2 between very high and high, and ± 1.9 between high and moderate. (Recall that all words had been rated on a scale from -3 to $+3$ by an independent group of subjects.) The analysis of these variables showed a marginal effect of Word frequency [$F(1,166) = 3.67$; $MS_e = 4149.06$; $P = .057$], but reliable effects of both Valence [$F(1,166) = 7.25$; $MS_e = 4149.06$; $P = .008$] and Degree of

Valence [$F(2,166) = 5.72$; $MS_e = 4149.06$; $P = .004$]. The RT data are shown in Fig. 3.

Decision Criteria. Criteria are shown in Fig. 2. The analysis of C in a one-way ANOVA showed a significant effect of Face type (3 levels: Neutral, Happy, and Sad): [$F(2,48) = 4.635$; $MS_e = 0.059$; $P = .015$]. While the criterion in the Pseudoface condition was unbiased ($C = 0.009$), it was liberal ($C = -0.162$) for designating words as negative in the Sad condition. There was an opposite tendency for happy faces to make subjects more inclined to call words positive ($C = 0.038$). Contrasts showed significant differences between the Sad and Neutral conditions [$F(1,48) = 4.97$; $P = 0.031$], and between the Sad and Happy conditions ($F = 8.47$; $P = .006$), but no significant difference between the Happy and Neutral conditions ($F < 1$).

Discussion

The results of Experiment 2A fully support our hypothesis concerning automatic processing of facial expressions and extend the findings of Experiment 1, which it replicates with some variations. The most important of these alterations is the substitution of sad faces for angry faces. The outcome of the experiment gives unequivocal support to the notion that sad expressions function as the evaluative opposites of happy faces. Happy

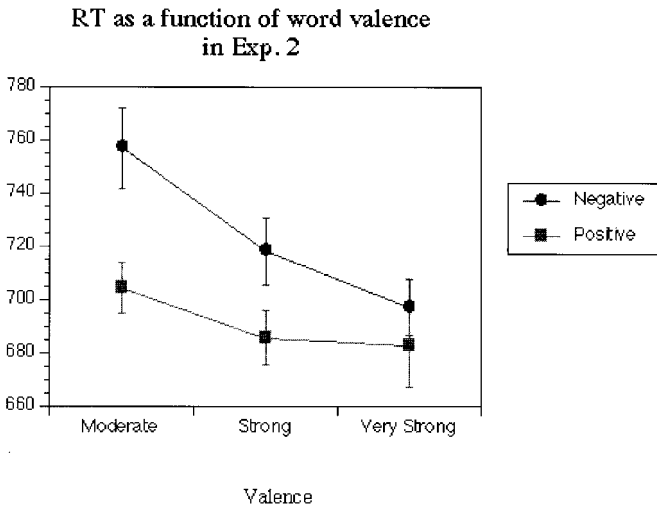


FIG. 3. Reaction times for positive and negative words of very strong, strong, and moderate degrees of valence (averaged across different face backgrounds). Error bars indicate SEM.

faces substantially increased the Positive Valence Advantage, and sad faces reversed it. The errors indicated a positive bias with happy faces, and a negative bias with sad ones.

Thus, the findings for happy faces were closely replicated, but the outcome for angry and sad faces were quite different. It remains to be determined whether this should be attributed to fundamental dissimilarities between different negative emotions (Niedenthal et al., 1994), to special attention-catching qualities in angry faces (Hansen & Hansen, 1994; Öhman, 1993), or to some subtle methodological variation between our two experiments. This issue will be addressed in Experiment 3.

Our analysis of reaction time differences between positive and negative words showed that positive words were judged more quickly, even when allowance was made for the fact that they were also more frequent in written Swedish. However, the extreme cases of negative valence were reacted to almost as quickly as the extremely positive ones. More doubtful cases were judged more slowly, as could be predicted from the attitude accessibility literature (Fazio et al., 1986). Although this accessibility effect applied to both types of words, moderately negative words were judged much more slowly than moderately positive ones. Thus, valence *per se* seemed to be less important than the added cognitive processing needed for the doubtful cases. Candidates for a negative verdict seemed to receive more extended scrutiny, in line with findings of more extensive cognitive processing of negative events (Taylor, 1991).

Experiment 2B: Method

Materials and Procedure. The memory test followed the evaluation task after a brief (1 minute) interval and without prior notification. On-screen instructions informed the subjects that a recognition test of the words from the previous test would now ensue. For each word, the subject was to decide whether it had been shown earlier and press a key on the keyboard (left- or right-arrow) accordingly.

All words that had been shown with an affective face (i.e. 80 out of the 120), were presented as old words in the memory test, intermingled with 80 new words. All word stimuli were presented against the background of an affective face, using the same timing and stimulus parameters as before. For half the old words, the affective face previously shown was retained, whereas for the other half, a picture of the opposite affect (happy changed to sad and vice versa), posed by the *same* person, was shown. This was to ensure that other aspects of the pictorial context were kept constant, while only the affect was being manipulated. The new words were taken from the same corpus and comprised positive and

negative words in equal proportions. They were also presented against a backdrop of an affective face, happy and sad expressions being presented with equal probability and completely crossed with the valence of the word.

Results

As a measure of recognition accuracy, we used the Signal Detection parameter d' (Snodgrass & Corwin, 1988). Values of d' were entered into a three-way ANOVA with the factors Word Valence (2 levels: Positive and Negative), Test Facial Expression (2 levels: Happy and Sad), and Study Facial Expression (2 levels: Happy and Sad).

The first hypothesis to be considered was that of a context influence on memory accuracy. Specifically, we expected recognition to be better of those words shown with the same facial expression during test and during study. This translates into an interaction term between Test Facial Expression and Study Facial Expression. The interaction revealed a tendency in the expected direction, which did not reach the conventional significance level [$F(1,24) = 3.20$; $P = .086$; $MS_e = 0.12$]. Contrasts showed that, for words first encountered with a Sad face, it made a difference whether it was tested with the same background ($d' = 1.39$), or not ($d' = 1.20$); ($F = 7.53$; $P = .011$). For words studied with a Happy face background, the test background did not matter (d' for same = 1.19, d' for different = 1.21; $F < 1$, n.s.).

Our second hypothesis was that words encountered in an incongruent context during study would be more elaborately encoded, and consequently better retained. The test of this hypothesis was the interaction term Word Valence \times Study Facial Expression. It showed a reliable effect [$F(1,24) = 11.83$; $P = .002$; $MS_e = 0.19$]. Positive words first encountered with a Sad face were much better remembered than those encountered with a valence-congruent Happy face [($d' = 1.48$ vs. 1.17), $F = 12.54$; $P = .002$]. For negative words, valence-incongruence conferred only a smaller and nonsignificant advantage [($d' = 1.23$ vs. 1.11), $F = 1.75$; $P = .198$].

Because we had given speeded instructions for the memory test, we also examined recognition latency as a dependent variable. Interestingly, the latency data showed the same Word Valence \times Study Facial Expression² interaction as did the accuracy data [$F(2,48) = 10.37$; $P = .002$; $MS_e = .19$]. Positive words studied with an incongruent, sad face were recognised much faster (890msec) than positive words studied with a congruent, happy face (955msec), ($F = 17.99$; $P < .001$). For negative words, the

² The Study Facial Expression in this analysis has three levels: Sad, Happy, and None, because new words were included in the analysis.

background face of the study episode did not matter (1002 vs. 984msec), ($F = 1.42$, n.s.).

Discussion

These memory results suggest that affective faces are involuntarily processed for affective valence, and that the outcome of this evaluation may affect the episodic memory trace of concurrent controlled processing. Incongruent combinations, such as a positive word shown with a sad face, require more extensive processing than a congruent one (positive word with a happy face), which leads to slower responding in the immediate task, but in compensation better subsequent memory retention. However, the results will have to be regarded as suggestive rather than conclusive, because the expected effect was obtained only for positive words, not for negative ones.

There was also a tendency for preserved facial expressions (from study to test) to improve memory of the words. As a test of the effect of context on memory, this was a very strict one, because in each case, pictures of the same person provided the facial background for the word, and only the expression changed.

EXPERIMENT 3

The purpose of the third experiment was to extend the previous ones by clarifying the relation between the negative emotions. Up to this point, our findings have been consistent regarding the automatic evaluation of happy faces, but the effects of angry and sad faces have been divergent. Only a test using the same subjects and precisely the same experimental parameters for both can ascertain whether the differences between anger and sadness are real. If they are, they may be attributable to special, attention-attracting properties of angry faces, to which our evolutionary past has attuned us (Hansen & Hansen, 1994; Öhman, 1986, 1993). But a hitherto unmentioned consideration may also play a part. Some facial expressions are more frequent than others, and it is natural to assume that the ones seen seldom are more prone to catch attention. We venture the generalisation that angry faces are probably less frequent than sad ones, and certainly less frequent than neutral and happy ones, in cultures we know, although we have seen no empirical data to bolster this assertion. If frequency plays a part, it may be appropriate to include an even more infrequent expression than anger in our design, because any argument based on frequency would apply *a fortiori* to it. Disgust is such an emotional expression, which is visually very distinctive, yet typically rare in any real-life environment. Our design therefore includes three negative

conditions: sadness, anger, and disgust, along with a neutral and a happy condition.

Method

Materials and Procedure. Six pictures for each of the four affects, and three neutral ones, were selected from the Ekman–Friesen set. The 10 cells of the factorial design (5 affective conditions by 2 word valences) were each represented by 30 trials, resulting in 300 trials for the whole experiment. Within these constraints, the assignment of words to affective faces was randomly determined for each subject at run time, as was the order of the stimuli. Thus, no two subjects saw the same list of stimuli.

A new selection of words was made for this experiment, with the special purpose of matching the positive and negative words for frequency. Recent word frequencies for newspaper Swedish from the Press95 database³ were used to select 150 positive and 150 negative words. Average frequency was 32.47 per million (SD 70.34) for negative words, and 30.24 per million (SD 46.94) for positive words. Average word length was 5.65 characters (SD 1.06) for negative words, and 5.98 characters (SD 0.99) for positive ones.

The display parameters were altered somewhat, with the purpose of encouraging subjects to respond quickly. Each trial began with a fixation cross, displayed at the centre of the screen for 1 second. The picture of the face was then displayed at a randomly selected location at a distance of 100 pixels (3.5cm) from the fixation cross. Thus, by selecting random positions along the periphery of a circle, positional uncertainty was created, while keeping saccade length constant. After 300msec, the word was superimposed on the face at its customary place across the lower nose area. The face-word compound was displayed for 500msec. Responses given within a 2000msec period from the word onset were accepted.

Because the experimental session was longer than in the earlier experiments (about 20 minutes), five short breaks with feedback about performance (average RT and error rate) were inserted to maintain motivation throughout. The experimental trials were preceded by practice trials with auditory feedback for errors. The assignment of “positive” and “negative” responses to the left- and right-arrow keys was balanced over subjects.

Subjects. A total of 51 subjects, 24 men and 27 women, students at the University of Lund, participated. They had a median age of 23 years (range: 19–34).

³ Courtesy of the University of Gothenburg, Sweden, the database is available on the Internet at <http://logos.svenska.gu.se/>.

Results

Reaction Times. Figure 4 illustrates the reaction times for the 5 affective conditions by 2 word valences. The critical interaction of Facial Expression \times Word Valence was reliable [$F(4,200) = 12.26$; $P < .001$; $MS_e = 663.28$]. Planned contrasts revealed the nature of the interaction. (For each contrast, degrees of freedom are 1,200 and $MS_e = 663.28$.) Happy faces increased the Positive Valence Advantage to 53.7msec, compared to 31.4msec in the Neutral condition. The difference was reliable ($F = 9.50$; $P = .003$).

The negatively valenced facial expressions all decreased the PVA: Angry 16.1, Disgusted 9.3, and Sad 15.6msec. They were all reliably different from the Neutral condition (Angry: $F = 4.49$, $P = .036$; Disgusted: $F = 9.41$, $P = .003$; Sad: $F = 4.80$, $P = .031$). They were also all different from the Happy condition (Angry: $F = 27.04$, $P < .001$; Disgusted: $F = 37.82$, $P < 0.001$; Sad: $F = 27.80$, $P < .001$).

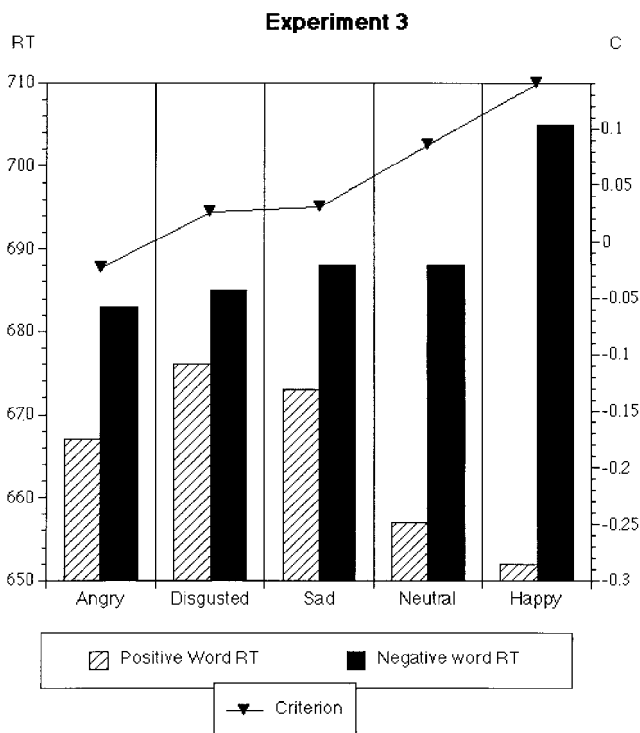


FIG. 4. Reaction times for positive and negative words, shown with angry, disgusted, sad, neutral, and happy faces in Experiment 3. Criterion values shown in the superimposed line diagram as in Figs. 1 and 2.

There was a main effect of Word Valence, because positive words were responded to faster than negative words across all conditions [$F(1,50) = 34.06$, $P < .001$, $MS_e = 2381.95$].

Errors and Criteria. As before, the average error rate was relatively low: 7.6%. Criteria were significantly different between conditions [$F(4,200) = 5.19$; $MS_e = 0.038$; $P < 0.001$]. Planned contrasts showed that the Angry condition was reliably different from the Neutral one [$F = 7.80$, $P = .006$]. (For this and the following contrasts, degrees of freedom are 1,200, and $MS_e = 0.038$.) Other emotional conditions were not different from the Neutral one [Happy: $F = 1.98$, Disgusted: $F = 2.32$, Sad: $F = 2.02$; all P s $> .05$].

Opposite-valenced emotional conditions were reliably different from each other. The Angry-Happy contrast amounted to [$F = 17.63$, $P < .001$; Disgusted-Happy: $F = 8.59$, $P = .004$; Sad-Happy: $F = 8.00$, $P = .005$].

Discussion

The findings were in line with predictions generated by the view that all affective expressions are automatically processed for valence, and all affects can be placed along a single dimension of valence. All negative affects acted as opposites to the one positive affect in their effects on the PVA and the criterion, and the neutral condition fell expectedly between. The negative affects differed only negligibly from each other, and no affect, frequent or infrequent, had any generally retarding effect on reaction times (i.e. one acting on both positive and negative words).

As before, negative words were processed more slowly than positive ones, and concerns about a possible confound with word frequency could be allayed. The more powerful experimental control of word frequency gave the same outcome as the statistical control in Experiments 1 and 2 (i.e. there was an effect of valence on evaluation latency independently of word frequency).

The PVA is a net effect, potentially consisting of both facilitation and inhibition effects in relation to a neutral baseline. Decomposition shows that inhibition generally exceeded facilitation in this experiment. The data for all three experiments are shown in Table 1. In contrast to the present case, the data for Experiment 2 show a predominance of facilitation effects. The difference lies in the choice of neutral baselines. The unusual, artificial pseudofaces of Experiment 2 seem to have caught some visual attention and retarded processing across both positive and negative words. The neutral faces were apparently processed more quickly. In part, this may result from the fact that there were fewer individual faces among them than

TABLE 1
Differences Between Reaction Times for Affective-face Conditions and Neutral Conditions in the Three Experiments

	<i>Exp 1</i>		<i>Exp 2</i>		<i>Exp 3</i>	
	Positive	Negative	Positive	Negative	Positive	Negative
Angry	6.1 (n.s.)	32.0 (.001)	–	–	9.7 (.06)	–5.6 (n.s.)
Disgusted	–	–	–	–	19.2 (<i>< .001</i>)	–3.0 (n.s.)
Sad	–	–	11.0 (n.s.)	–34.7 (.002)	15.5 (.003)	–0.3 (n.s.)
Happy	–18.7 (.04)	20.0 (.03)	–31.0 (.006)	1.0 (n.s.)	–5.0 (n.s.)	17.2 (.001)

Note: Positive figures indicate inhibition in relation to neutral, negative figures indicate facilitation. Figures in parentheses show significance level for the contrast between the affective and the neutral condition. See further comments in the Discussion of Experiment 3.

among the affective expressions (3 neutral vs. 6 of each affect). The more frequently repeated faces in all likelihood demanded progressively less time for visual encoding. In sum, the vagaries of different neutral pictures demonstrate the difficulty of finding a truly neutral baseline, and justify the use of a baseline-independent variable, such as the PVA, as an indicator of the processing of valence.

GENERAL DISCUSSION

The three experiments reported in this article are consistent in showing that there is a speed advantage for positively valenced words in the word evaluation task. Our findings are summarised in Table 2. The Positive Valence Advantage was markedly and predictably modified by concurrently shown faces, such that the size of the PVA, which was 21–31msec when the words were shown with a neutral background, more than doubled, to 54–60msec, when a happy face was shown. Analogously, when faces expressing negative emotions were shown, the PVA was drastically reduced by about 50% in the third experiment (to 9–16msec). In the second experiment, the reduction induced by sad faces was even more radical, leading to a reversal of the PVA into an advantage for negative words by 23msec. The only exception to this generalisation was the results for angry faces in Experiment 1. The increase in PVA from neutral to angry faces was against our expectations, and inconsistent with the findings we obtained with the same faces in Experiment 3.

TABLE 2

The "Positive Valence Advantage" (PVA) for a Given Face Display Type is the Difference between RT for Correct Negative Responses and RT for Correct Positive Responses. C is the Criterion

		<i>Negative Valence</i>			<i>Neutral Valence</i>	<i>Positive Valence</i>
		<i>Angry</i>	<i>Sad</i>	<i>Disgusted</i>	<i>Neutral</i>	<i>Happy</i>
PVA	Exp. 1	47.45 ^a	—	—	21.49	60.18 ^a
	Exp. 2	—	-23.39 ^{a,b}	—	22.29	54.31 ^{a,b}
	Exp. 3	16.14 ^{a,b}	15.62 ^{a,b}	9.29 ^{a,b}	31.42	53.65 ^{a,b}
C	Exp. 1	.005 ^a	—	—	-0.096	.058 ^a
	Exp. 2	—	-.162 ^{a,b}	—	-.009	.038 ^b
	Exp. 3	-.023 ^{a,b}	.030 ^b	.026 ^b	.085	.140 ^b

Note: Within each experiment, superscript a denotes a significant difference (at $P < .05$) from the neutral condition, and superscript b denotes a significant difference from the opposite-valence condition (e.g. Happy vs. Sad).

Murphy and Zajonc (1993) have also found signs of automatic evaluation of affective faces, manifested by priming effects on the evaluation of subsequently presented stimuli (Chinese ideographs). However, in their findings, affective priming was effective only with faces shown at very brief (4msec) durations, not at 1sec exposures. In contrast, our priming effects were obtained with the stimuli in full view. Reaction time may be more sensitive as a dependent variable than the rating scales used by Murphy and Zajonc, and may therefore capture subtle priming effects otherwise counteracted by controlled processing.

The pattern offered by our error data are consonant with the reaction time data. Although the proportion of errors was small on the whole, the distribution of the types of errors supported the conclusion that subjects were more willing to call a word positive in meaning when it was shown with a happy face, and more willing to call it negative when shown with a negatively valenced facial expression. Analysed in signal detection terms, the criterion for detecting negative word targets was raised by positive faces, and lowered by negative ones. The fact that there was a parallelism between RT and criterion data—although RT proved the more sensitive of the two—is useful in excluding trade-offs between speed and accuracy as explanations of the findings. It extended across all instances, including our one anomalous finding (i.e. the results for angry faces in Experiment 1).

There is no readily apparent explanation for the anger anomaly in the first experiment. The possibility that angry faces attract attention and involuntarily instigate controlled processing was not supported by our third experiment. There were, however, some procedural differences between our experiments which somewhat strain comparisons. The main

difference was the display duration for the stimuli. In Experiment 1 we let the word-face compounds dwell on screen for 1500msec, but in Experiments 2 and 3 we introduced both positional uncertainty and briefer displays, 750 and 500msec respectively. The purpose of these changes was to induce subjects to respond both quickly and holistically. The changes were effective in this respect, because average reaction times (across all stimuli) declined from 740 to 709 to 677msec across the three experiments, as a monotonic function of display duration. The longer inspections awarded to the stimuli in Experiment 1 may have allowed more controlled processing to develop, this preferential treatment being disproportionately often given to angry face stimuli. However, to yield the results we obtained, this tendency would have to interact with word valence, such that negative words pre-empted the extra controlled processing capacity more often than did positive words. All of these assumptions would need independent verification.

A different line of reasoning involves the stimulus material. The stock of words varied between experiments. Out of the total of 298 negative word tokens used—60 in Experiment 1, 88 in Experiment 2, and 150 in Experiment 3—only 64 participated in more than one experiment; 14 of these were common to all three. Thus, the composition of the word samples showed little overlap.

Because negative emotions are different from each other, it is possible that a selection of negative words may match the semantic characteristics of one emotion better than another (Niedenthal et al., 1994). Specifically, our selections of negative words could have varied in their proportion of anger-related meanings, creating varying degrees of match with the angry faces. To investigate this possibility, we made an effort to classify the emotional content of the words. Only 13% could be classified as directly related to anger, rage, or deliberate physical violence. There was little difference in this proportion between experiments (15, 17, and 13%). Therefore, it is not likely that word sampling variation could have created the different results for angry faces in Experiments 1 and 3, and the inconsistency will essentially have to remain unexplained.

Generality of Involuntary Evaluation

The interference pattern we obtained suggests that good-bad evaluation is often included as a feature in the mental representation of affective expressions. An issue worth considering is the generality of this finding. Does it apply to a broader class of images than just affective faces? Whereas other images have been found to interfere when words are to be categorised (Smith & Magee, 1980), this has concerned categorisation tasks where the semantic relation is very obvious, such as applying category label to

category members—"clothing" to "shirt", "hat", etc. Although evaluation features are far less obviously a part of the mental representation, findings indicate that attitudes, at least if they are strongly held and easily accessible, are automatically evoked on presentation of the object (Fazio et al., 1986). With pictures of animals, De Houwers and Hermans (1994) found results similar to ours (i.e. pictures involuntarily evoking good-bad evaluations, and therefore interfering with the oral positive/negative response given to a concurrently presented word). Their set of stimuli included six animals—snake, spider, beetle, in the negative category and duck, rabbit, bird in the positive one. Snakes and spiders are frequently objects of phobias, and they have presumably figured as natural enemies in our evolutionary history. It has therefore been suggested (Öhman, 1986, 1993) that humans are genetically prepared for preattentive processing of their visual appearance. The class of automatically processed objects may of course be larger than those we have been genetically prepared for. Learning plays a role in acquiring long-term automaticity of processing (Logan, 1988; Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977), and temporary accessibility of attitudes can be enhanced by repetition (Roskos-Ewoldsen & Fazio, 1992).

This study suggests that valence features are included in the mental representations of affective facial expressions and are made automatically accessible. Facial expressions therefore belong to a larger class of affectively tagged objects, whose evaluations are automatically evoked on mere presentation in a large majority of subjects, regardless of individual experience. Other members of this class are, in all likelihood, snakes, spiders, and other prominent enemies from our evolutionary history, as well as other fear-relevant and phobia-inducing objects. On an individual and idiosyncratic basis, a much larger class of objects may be able to evoke automatic evaluation, as shown by Fazio and co-workers (1986).

Valence and Latency

Our second main question concerned possible differences between positive and negative objects in the speed with which they were processed. The findings in this regard were:

1. Decision latencies in all experiments were shorter for positive words than for negative ones.
2. Word frequency influenced decision latency as expected, and positive words had, on the average, higher frequency in written Swedish than negative ones, paralleling Zajonc's (1968) observation for American English. However, the valence effect was independent of the frequency effect, as ascertained by both statistical control (Experiments 1 and 2) and stimulus selection control (Experiment 3).

3. As could be expected from earlier results (Fazio et al., 1986) on attitude accessibility, objects with more extreme normative evaluations, as indicated by the ratings of our pilot study group, were valence-categorised more rapidly than moderately valenced objects (Experiment 2). The increase in processing time for the doubtful exemplars was particularly noteworthy among negative words (Fig. 3). Thus, the extra processing demanded for the difficult decisions took a heavier toll on speed for negative than for positive words.

The task of incorporating these findings into the extant literature is at first glance made more difficult by the apparently contradictory findings concerning valence-induced latency differences of earlier studies. On the one hand, detection of negative objects has been found to be more rapid (Hansen & Hansen, 1988). On the other, latencies for negative words are longer in the affective Stroop task, where affective words serve as stimuli. For anxious subjects, detection of probe stimuli is more rapid in locations to which attention has recently been drawn by negative words (MacLeod et al., 1986). More generally, negative events have been found to be subjected to more elaborate and protracted processing than positive ones (Taylor, 1991).

Apparently, the task makes a large difference for the kind of latency differences, if any, that negative and positive objects will show. As a generalisation, it may be said that negative stimuli are detected more rapidly than positive ones, and if detection is the only processing the task demands, latencies will be shorter for negative stimuli. If on the other hand, some cognitive processing is called upon, negative stimuli will occupy cognitive resources for a longer time and to a greater extent (Taylor, 1991), leading to longer latencies. Thus, detection tasks such as that of Hansen and Hansen result in speeded reactions for negative faces. Tasks in which involuntary processing is invoked by background stimuli do so to a greater extent if the stimuli are negative (Pratto & John, 1991; White, 1996).

The principle of protracted cognitive processing of negative stimuli applies well to the present findings. Reaction times for negative words were longer, and as Fig. 3 shows, the difference seems to be attributable above all to more extended processing of doubtful cases.

Conclusion

This study has shown that some semantic qualities, particularly affective valence, are stored in a highly accessible form with mental representations of affective facial expressions. They are automatically evoked by mere presentation of affective faces and may therefore interfere with other

ongoing processing. In the face-to-face encounters of daily social life, we should perhaps keep in mind that the words we say may thus be—in an unexpected sense—taken at face value.

Manuscript received 19 November 1996

Revised manuscript received 6 March 1998

REFERENCES

- Allén, S. (1971). *Nusvensk frekvensordbok* (Vol. 2). Stockholm: Almqvist & Wiksell.
- Bargh, J.A., Chaiken, S., Govender, R., & Pratto, F. (1992). The generality of the automatic attitude activation effect. *Journal of Personality and Social Psychology*, *62*, 893–912.
- De Houwer, J., & Hermans, D. (1994). Differences in the affective processing of words and pictures. *Cognition and Emotion*, *8*, 1–20.
- Dimberg, U., & Öhman, A. (1996). Behold the wrath: Psychophysiological responses to facial stimuli. *Motivation and Emotion*, *20*, 149–182.
- Ekman, P., & Friesen, W. (1975). *Unmasking the face: A guide to recognizing emotions from facial clues*. Englewood Cliffs, NJ: Prentice-Hall.
- Eriksen, B.A., & Eriksen, C.W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception and Psychophysics*, *16*, 143–149.
- Esteves, F., Dimberg, U., & Öhman, A. (1994a). Automatically elicited fear: Conditioned skin conductance responses to masked facial expressions. *Cognition and Emotion*, *8*, 393–413.
- Esteves, F., Parra, C., Dimberg, U., & Öhman, A. (1994b). Nonconscious associative learning: Pavlovian conditioning of skin conductance responses to masked fear relevant facial stimuli. *Psychophysiology*, *31*, 375–385.
- Fazio, R.H., Sanbonmatsu, D.M., Powell, M.C., & Kardes, F.R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, *50*, 229–238.
- Glaser, W.R., & Glaser, M.O. (1989). Context effects in Stroop-like word and picture processing. *Journal of Experimental Psychology: General*, *118*, 13–42.
- Hansen, C.H., & Hansen, R.D. (1988). Finding the face in the crowd: An anger superiority effect. *Journal of Personality and Social Psychology*, *54*, 917–924.
- Hansen, C.H., & Hansen, R.D. (1994). Automatic emotion: Attention and facial efference. In P.M. Niedenthal & S. Kitayama (Eds.), *The heart's eye: Emotional influences in perception and attention* (pp. 217–243). San Diego, CA: Academic Press.
- Kahnemann, D., & Treisman, A. (1984). Changing views of attention and automaticity. In R. Parasuraman & D.R. Davies (Eds.), *Varieties of attention* (pp. 29–61). New York: Academic Press.
- Logan, G.D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492–527.
- MacLeod, C., Mathews, A., & Tara, P. (1986). Attentional bias in emotional disorders. *Journal of Abnormal Psychology*, *95*, 15–20.
- MacLeod, C.M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, *109*, 163–203.
- Mathews, A., & MacLeod, C. (1994). Cognitive approaches to emotion and emotional disorders. *Annual Review of Psychology*, *45*, 25–50.
- Murphy, S.T., Monahan, J.L., & Zajonc, R.B. (1995). Additivity of nonconscious affect: Combined effects of priming and exposure. *Journal of Personality and Social Psychology*, *69*, 589–602.

- Murphy, S.T., & Zajonc, R.B. (1993). Affect, cognition and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology*, *64*, 723–739.
- Niedenthal, P.M. (1990). Implicit perception of affective information. *Journal of Experimental Social Psychology*, *26*, 505–527.
- Niedenthal, P.M., Setterlund, M.B., & Jones, D.E. (1994). Emotional organization of perceptual memory. In P.M. Niedenthal & S. Kitayama (Eds.), *The heart's eye: Emotional influences in perception and attention* (pp. 87–113). San Diego, CA: Academic Press.
- Öhman, A. (1986). Face the beast and fear the face: Animal and social fears as prototypes for evolutionary analyses of emotion. *Psychophysiology*, *23*, 123–145.
- Öhman, A. (1993). Fear and anxiety as emotional phenomena: Clinical phenomenology, evolutionary perspectives, and information-processing mechanisms. In M. Lewis & J.M. Haviland (Eds.), *Handbook of emotions* (pp. 511–536). New York: Guilford Press.
- Osgood, C.E., Suci, G.J., & Tannenbaum, P.H. (1957). *The measurement of meaning*. Urbana, IL: University of Illinois Press.
- Pratto, F., & John, O.P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Social and Personality Psychology*, *61*, 380–391.
- Roskos-Ewoldsen, D.R., & Fazio, R.H. (1992). On the orienting value of attitudes: Attitude accessibility as a determinant of an object's attraction of visual attention. *Journal of Personality and Social Psychology*, *63*, 198–211.
- Schneider, W., & Shiffrin, R.M. (1977). Controlled and automatic human information processing: I. Detection, search and attention. *Psychological Review*, *84*, 1–66.
- Shiffrin, R.M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, *84*, 127–190.
- Smith, M.C., & Magee, L.E. (1980). Tracing the time course of picture-word processing. *Journal of Experimental Psychology: General*, *109*, 373–392.
- Snodgrass, J.G., & Corwin, J. (1988). Pragmatics of measuring recognition memory—applications to dementia and amnesia. *Journal of Experimental Psychology: General*, *117*, 34–50.
- Stroop, J.R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*, 643–662.
- Taylor, S.E. (1991). Asymmetrical effects of positive and negative events: The mobilization-minimization hypothesis. *Psychological Bulletin*, *110*, 67–85.
- Toglia, M.P., & Battig, W.F. (1978). *Handbook of semantic word norms*. Hillsdale, NJ: Erlbaum.
- White, M. (1995). Preattentive analysis of facial expressions of emotion. *Cognition and Emotion*, *9*, 439–460.
- White, M. (1996). Automatic affective appraisal of words. *Cognition and Emotion*, *10*, 199–211.
- Zajonc, R.B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, *9*, 1–27.