Constraints on the Interaction Between Context and Stimulus Information

Michael E. J. Masson

Department of Psychology
University of Victoria
Victoria, British Columbia
Canada V8W 3P5
bitnet: mmasson@uvunix.ca

Abstract

A central issue in the development of models of context effects, such as the interactive activation model, concerns the relationship between contextual and stimulus information. Empirical evidence regarding perception of spoken and printed words indicates that context and stimulus information make independent contributions to perceptual identification. Recent research on visual object recognition, however, suggests that context may have a direct influence on the rate or accuracy of visual analysis. These results imply that contextual influences in language comprehension and object recognition may operate in fundamentally different ways. A series of experiments is described that lead to a reinterpretation of the object recognition results. It is concluded that contextual information contributes to the interpretation of stimulus input without altering its form or the rate of its acquisition.

Introduction

The influence of context on the identification of stimuli is very powerful and has been demonstrated in a number of domains including auditory and visual word identification, letter identification, and picture identification (e.g., Biederman, Mezzanotte, & Rabinowitz 1982; McClelland & Rumelhart 1981; Stanovich & West 1983; Connine 1987). Two major theoretical views of how context influences stimulus identification have emerged. The *interactionist* view emphasizes the interactive nature of cognitive and perceptual processes and suggests that contextual knowledge directly enhances the speed or accuracy with which visual perception operates, thereby increasing the discriminability of alternatives (McClelland & Elman 1986; Rumelhart 1977). The *classic* view adheres to a modular approach in which it is assumed that processes devoted to visual and contextual analysis operate autonomously and that evidence independently obtained from each process is integrated to complete stimulus identification (Fodor 1983; Massaro 1989; Oden & Massaro 1978; Rueckl & Oden 1986). According to the classic view, context does not directly influence featural analysis of a stimulus or discriminability of alternatives, but instead affects the criterion applied when selecting among stimulus alternatives.

Massaro (1989) challenged the class of interactive activation models, and in particular the TRACE model of speech perception (McClelland & Elman 1986), on the grounds that such models predict a contextual influence on acquisition of stimulus information, whereas data from a speech perception paradigm clearly indicated that asymptotic performance reflected independent contributions of contextual and stimulus information. Contrary to the suggestion that these results rendered interactive models invalid as descriptions of perceptual identification, a revision to the interactive activation model proposed by McClelland (1991) reproduced the independence effect obtained by Massaro. A critical alteration to the original interactive activation model was an architectural constraint that prevented direct interactions between contextual units and stimulus input units. Interactive processing was retained in the sense that units representing alternatives among which a choice must be made (e.g., letters in a letter identification task) could influence activation among lower level feature units. Contextual information (e.g., the word in which a target letter is embedded) could indirectly affect stimulus information, through its effect on units representing response alternatives.

The question addressed here is whether, in the early stages of perceptual identification, context can influence the rate or accuracy of stimulus analysis, even indirectly as suggested by McClelland's (1991) revised approach. An example of such an effect offered by McClelland is the finding that forced-choice identification of a letter is enhanced when the target letter is presented in the context of a word or pronounceable pseudoword (Reicher 1969). This discriminability effect, however, is not at the level of letter features, but at the level of letter identity which under even the revised interactive model has bidirectional links to the word level. Furthermore, tests of a possible interaction between sentence context and word identification have shown that sentence context serves to influence response criterion, not discriminability among alternatives (Masson 1988). In a related study Kroll (1990) demonstrated that sentence context has similar effects on identification latencies for both printed words and drawings of objects. The cross-modality nature of this result suggests that sentence

This project was supported by grants from the Natural Sciences and Engineering Research Council of Canada and from the University of Victoria.

context interacted with stimulus processing at the level of meaning rather than featural analysis. Moreover, the influence of sentence context on auditory word perception appears to operate post-perceptually, rather than on the encoding of the auditory stimulus (Connine, 1987).

More convincing evidence for an interaction between context and stimulus processing has been obtained with visual object identification tasks. In general, observers are more accurate at identifying objects when presented as part of or following a relevant context as opposed to an irrelevant context (e.g., Biederman 1972; Biederman et al. 1982; Boyce, Pollatsek, & Rayner 1989; Reinitz, Wright, & Loftus 1989). Boyce et al. and Reinitz et al. used methods to correct for decision bias effects and still obtained an advantage for objects presented with relevant contextual information.

In this paper I report a series of experiments designed to explore the source of the context effects obtained in the Reinitz et al. (1989) and Boyce et al. (1989) studies. The general issue was whether these context effects arise from the influence of context on the rate or accuracy of stimulus processing. The alternative hypothesis underlying these new experiments was that enhanced object identification is a result of using context to interpret visual information without actually modifying the amount or accuracy of that information. In the case of the Reinitz et al. study, the method of assessing decision bias effects was not deemed appropriate, and in the Boyce et al. experiments it was not clear whether contextually enhanced identification was accompanied by improved acquisition of visual details. These two issues are addressed in the experiments reported here.

Experiment 1

The purpose of the first experiment was to replicate the Reinitz et al. (1989) results using a set of materials that could be altered for use in a paradigm amenable to signal detection analysis. A set of line drawings of objects served as the stimuli rather than photographs of real objects as in the Reinitz et al. experiments. This format allowed objects to be redrawn with alterations for use in Experiment 2, which examined the possible influence of context on detection of visual detail.

Method

The subjects were 24 university students. A set of 72 objects were selected from the Snodgrass and Vanderwart (1980) norms and an artist reproduced a line drawing of each one. Eight of the objects were used as practice items and the remaining 64 were treated as critical items. In addition, each object's name served as its appropriate context word. Half of the critical and practice items were presented with their appropriate context word and half with an inappropriate context word. Inappropriate context words were obtained by randomly reassigning words to objects with the constraint that the resulting pairs were unrelated. Assignment of critical items to context conditions was counterbalanced across subjects.

Stimulus materials were shown using slide projectors and tachistoscopic shutters controlled by an Apple II microcomputer. On each trial the subject was shown a context word for 1 s followed by the target object which was in view for 75 ms before being replaced by a pattern mask. The subject's task was to name the briefly presented object. Subjects were informed that sometimes the word would name the object and sometimes it would not. Eight practice trials were presented first, followed by a randomly ordered presentation of 64 critical trials.

Results and Discussion

The mean proportions of correctly identified critical objects were .64 in the appropriate context condition and .36 in the inappropriate context condition. These means were significantly different, $t(23) = 7.45$, $p < .001$. The probability of subjects offering the context word as their response in the inappropriate condition was .08. When this value was used to correct for guessing in the appropriate condition (as did Reinitz et al. 1989) the corrected probability was .59, which was significantly different from the inappropriate condition, $t(23) = 5.57$, $p < .001$. The results of Experiment 1 successfully replicated the Reinitz et al. (1989) finding of an advantage of appropriate context even when a correction for guessing was applied.

Experiment 2

To test the hypothesis that the context advantage observed in Experiment 1 stemmed from more efficient visual analysis in the appropriate context condition, a second experiment was conducted. In this experiment the displays were very similar to those of Experiment 1, except that the task was to examine a clearly visible comparison drawing to determine whether it was identical to the briefly displayed object or whether it was an altered version of the object. If context serves to enhance the accuracy or rate of visual analysis, as claimed by Reinitz et al. (1989), subjects should be more accurate in this task when objects are presented following an appropriate context word.

Method

The subjects were 64 university students. The materials and procedure were the same as in Experiment 1 with the following exceptions. For each of the objects a second version was drawn with a visual detail altered (see Figure 1 for an example). On each trial the pattern mask was followed by a comparison item that represented the same object as the one that had appeared before the mask. On half of the trials the comparison item was physically identical to the target object and on

Figure 1. An example of an object in its original (left) and altered (right) form.

half of the trials the altered version of the target object was shown. The task was to decide whether the comparison item was identical to the target. For half of the trials of each type the context word was appropriate for the target object and for half of the trials it was inappropriate. In order to bring performance on this discrimination task above chance, it was necessary to raise the exposure duration. Two different durations were used, 180 ms and 230 ms, with half of the subjects tested at each duration.

Results and Discussion

For the purposes of signal detection analysis a hit was defined as correctly classifying an identical comparison item as being the same as the target object, and a false alarm was defined as classifying an altered comparison item as the same as the target object. Performance did not vary significantly as a function of exposure duration, so the data presented here were collapsed across that factor. The mean hit and false alarm rates for the appropriate and inappropriate context conditions are shown in Table 1. A nonparametric, two-high threshold signal detection analysis was also applied to the data. This analysis was chosen because it maintains independence between measures of sensitivity and bias and does not rely on assumptions regarding normal distributions with equal variance (Snodgrass & Corwin 1988). The sensitivity measure, P_r, is computed as hit rate minus false alarm rate, with a value of 1.0 representing perfect discrimination and 0.0 representing chance. The bias measure, B_r, represents the probability of a "same" response when the subject is in a state of uncertainty. The mean P_r and B_r values are shown in Table 1.

Analyses of the sensitivity and bias scores indicated that there was no significant effect of context for either measure, $t(62) = 1.34$, $p > .15$, and $t < 1$, for sensitivity and bias, respectively. Given the failure to find a reliable effect of context on sensitivity, an analysis of the power of the study to find such an effect was carried out. The size of the context effect in Experiment 1 was approximately .8, a large effect size (Cohen 1977). It was assumed that the effect of context on detection of a visual alteration likely would be smaller, so the power analysis was carried out assuming a medium effect size of .5. The power of the study to detect an effect of this size was approximately .86.

Table 1
Mean Values of Signal Detection Measures as a Function of Context in Experiment 1

Context	Hit rate	FA rate	P_r	B_r
Appropriate	.74	.40	.34	.61
Inappropriate	.72	.42	.30	.61

Despite adequate power in Experiment 2 to detect a medium effect of context on a visual discrimination task, no reliable evidence for such an effect was obtained. It is conceivable that a smaller effect, not detectable by this experiment, exists. Even if such an effect exists, however, it would appear to be so small as to be incapable of accounting for the large effect of context observed in Experiment 1. By this reasoning, it seems plausible that a different process is responsible for generating much if not all of the context effect found in Experiment 1. In particular, it is possible that the context effect arises not from increased efficiency in visual processes, but from improvement in the accuracy of the interpretation of the results of visual analysis.

Experiment 3

Experiment 3 was designed to test the hypothesis that the effect of a context word on object identification is due to the availability of an appropriate interpretation or schema. Further, if contextual and stimulus information are integrated in some way *after* both are constructed, it should not matter whether the contextual information is supplied before or after the object is presented. The procedure for this experiment was similar to that of Experiment 1, except that the test format was a two-alternative forced choice. It was predicted that if the main function of the context word was to supply an interpretive schema, rather than a means of enhancing visual processing, then the availability of an appropriate schema in the form of a test alternative might serve a similar function. Consequently, the forced choice format was expected to eliminate the context effect observed in Experiment 1.

Method

A group of 24 university students was tested. The same materials and procedure as those in Experiment 1 were used for Experiment 3, with one exception. On each trial after the mask was removed, the subject was shown two alternative object names and was asked to select the one that corresponded to the briefly presented object.

Results and Discussion

The mean proportion of correct responses in the appropriate and inappropriate context conditions were

.80 and .79, respectively. These means were not significantly different, $t < 1$. Given that the design of Experiment 3 was identical to that of Experiment 1 it would be expected that similar context effects should have been found had the the context word served to increase the rate of visual information acquisition. The power of Experiment 3 to detect a context effect of the size observed in Experiment 1 was estimated at .98.

The outcome of Experiment 3 confirmed the hypothesis that the advantage provided by an appropriate context word that precedes the target object can be virtually eliminated by using a forced choice testing format. The presence of a correct alternative following the object on every trial created a level playing field for both context conditions, apparently by always providing a valid interpretive schema for the object. The results of Experiments 1-3 generally support the conclusion that the effect of a preceding context word on object identification is not to increase the rate of visual processing. Rather the influence of context appears to be on the interpretation of the results of visual analysis, as suggested by information integration models (e.g., Massaro 1989).

Experiment 4

An alternative method of providing context for object identification is to embed an object in an appropriate scene. The Boyce et al. (1989) experiments clearly demonstrated an improvement in object discriminability (sensitivity in signal detection analyses) under these conditions. The critical issue examined in the remaining experiments described here was the source of this enhanced identification accuracy. Experiment 4 was designed to produce a replication of the effect of context on identification accuracy and Experiment 5 tested a hypothesis regarding the source of that enhancement.

Method

The subjects were 24 university students. The 64 critical objects used in the earlier experiments were arranged as pairs and a background scene relevant to each object was created with the constraint that an object could be replaced in its scene by its partner to create a clearly inappropriate combination of object and context. An example of an object in an appropriate and an inappropriate scene is shown in Figure 2. For each scene a relevant and an irrelevant object name representing objects not in the scene were chosen for use as foils on the identification test. The 32 pairs of objects were arranged as four sets of eight pairs. Two sets were assigned to a condition in which the scene contained an appropriate target object, and two were assigned to a condition in which the scene contained an inappropriate target object. In each of these conditions, one set of scenes was probed with the name of the target object it contained and one set of scenes was probed with with the name of a foil. In the case of foils, the foil appropriate to the scene was used for scenes that contained an appropriate target and the inappropriate foil was used for scenes that contained an inappropriate target (see Figure 2). Assignment of the sets of scenes to context conditions was counterbalanced across subjects. Scenes and probes were also created for the eight practice items.

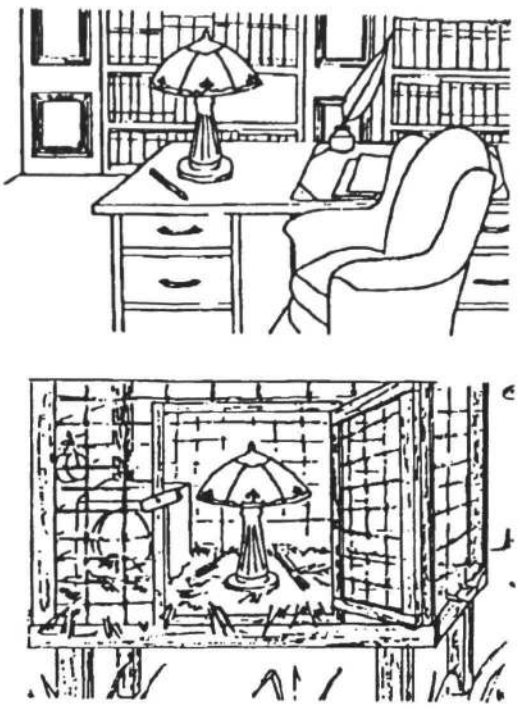

Figure 2. An example of a target object (lamp) embedded in an appropriate (upper) and an inappropriate (lower) scene used in Experiment 4. An appropriate scene was probed with the name of the target or the name of a foil object appropriate to the scene. An inappropriate scene was probed with the name of the target or the name of a foil object inappropriate to the scene.

The same apparatus was used as in earlier experiments. On each trial the subject was shown a scene for 130 ms followed by a pattern mask. The mask was followed by a probe consisting of the name of an object. The subject's task was decide whether the object named by the word was present in the scene. A subject was first shown a series of eight practice trials then 64 trials involving critical items presented in a random order.

Results and Discussion

A hit was defined as correctly claiming that a probed object was present in the scene and a false alarm was defined as making such a claim when a foil probe was presented. The mean hit and false alarm rates for the appropriate and inappropriate context conditions are shown in Table 2, along with the results of a two-high threshold signal detection analysis. Significance tests

Table 2
Mean Values of Signal Detection Measures as a Function of Context in Experiment 4

Context	Hit rate	FA rate	P_r	B_r
Appropriate	.80	.43	.36	.66
Inappropriate	.52	.31	.21	.38

were applied to the sensitivity and bias scores and it was found that there was higher sensitivity in the appropriate context condition, $t(23) = 2.52$, $p < .02$. In addition, differences in the bias scores indicated that there was a reliably greater tendency for subjects to claim a probe object was present when it was appropriate to the scene, $t(23) = 5.05$, $p < .001$.

The results of Experiment 4 constitute a replication of the Boyce et al. (1989) findings regarding the contribution of a relevant contextual scene to the identification of constituent objects, and also provide evidence for a strong bias on the part of subjects to claim an object was present if it were relevant to the scene. It is reasonable to conclude that these two effects of context, sensitivity and bias, occurred independently because the signal detection analysis applied to the data allows the separation of these two influences. One might argue, then, that the bias effect reflects the influence of context on interpretation processes and the sensitivity effect is evidence for the role of context in enhancing the acquisition of visual information. Before accepting this conclusion, however, an alternative explanation for the source of the sensitivity effect must be considered. It is conceivable that context serves to enable the identification of an object with *less* visual analysis than might be required for an object placed in an inappropriate scene. In this view, visual analysis may proceed at the same rate when viewing objects regardless of the scene in which they are placed, but less visual information may be needed in order to identify an object appropriate to the scene. When an inappropriate object is encountered more visual analysis may be needed in order to identify it. If viewing time is restricted, an observer would be less likely to identify an object that is not appropriate to the scene.

A counterintuitive prediction follows from this alternative view. If subjects are required to report on visual details of objects presented in appropriate versus inappropriate scenes, they may have a tendency to engage in less visual analysis of objects in the former case because they fit so well into the interpretive schema instantiated by the scene. Consequently, subjects might be expected to show *lower* sensitivity to objects appearing in appropriate scenes when the task requires discrimination of visual details of objects appearing in appropriate as compared to inappropriate scenes.

Experiment 5

Method

A group of 48 university students served as subjects in the experiment. The materials consisted of the scenes from Experiment 4 and the drawings of the target objects in their original and altered forms that were used in Experiment 2. The design and procedure was the same as in Experiment 4, with target objects appearing in appropriate or inappropriate scenes, except that the probe item consisted of a drawing of the target object in its original or altered form. The subject's task was to decide whether the drawing was visually identical to the version of the object that had appeared in the preceding scene. In order to bring performance above chance on this task it was necessary to increase the exposure duration of the scenes to 280 ms and to dispense with presentation of the pattern mask following the scene.

Results and Discussion

A hit occurred when the probe drawing was identical to the target that appeared in the scene and the subject claimed that it was identical. A false alarm occurred when the same claim was made about a probe that was an altered version of the target. The mean hit and false alarm rates, and mean sensitivity and bias scores are shown in Table 3. An analysis of the sensitivity scores revealed a marginally reliable difference in favor of the inappropriate context condition, $t(47) = 1.69$, $p < .10$, and no reliable difference on the bias measure, $t < 1$.

The results of Experiment 5 provided weak support for the prediction that subjects would tend to apply more visual analysis to objects that do not meaningfully fit in a scene, and clearly contradicted the notion that context enabled faster or more efficient acquisition of visual information. Given this outcome, the favored view of the influence of contextual scenes on object identification appears to be one in which context is assumed to enable earlier identification of objects that fit the schema invoked by the context. Rather than directly affecting the rate of visual information acquisition, then, context appears to reduce the amount of such information that is necessary for an accurate identification response to be produced.

Table 3
Mean Values of Signal Detection Measures as a Function of Context in Experiment 5

Context	Hit rate	FA rate	P_r	B_r
Appropriate	.70	.53	.17	.65
Inappropriate	.73	.48	.24	.64

Conclusion

The experiments reported here examined two very different kinds of context, both of which were shown to have large effects on the ability of observers to identify target objects. One type of context consisted of an object name, presented asynchronously with a target object. The other type of context comprised a scene presented contemporaneously with the target object. There was no evidence, however, to support the claim that presentation of appropriate contextual information enhanced the rate or accuracy of visual information acquisition. Moreover, it was shown that (a) when an object name served as context, its contribution was equally powerful when presented before or after visual processing had taken place, and (b) the contribution of appropriate contextual scenes may consist of enabling observers to identify objects with less visual information rather than increasing the rate or accuracy of visual processing. All of these effects were obtained under very brief exposure durations, enabling snapshots of processing as it occurs at very early stages of perceptual identification.

The present results converge with recent work involving context effects on visual and spoken word identification in localizing the influence of context at a level removed from direct perceptual analysis of a stimulus. In general, this research supports the view that context serves to influence the interpretation of stimulus information rather than to affect the rate at which stimulus processing goes forward or the accuracy with which it is conducted.

The empirical evidence regarding early stages of perceptual identification is consistent with a view of a connectionist architecture that includes bidirectional movement of activation between "intermediate level" processing units that represent words, letters, or objects and the units representing contexts in which they occur, but that constrains the connections between units representing featural details and units at the intermediate level to a bottom-up movement of activation. The similarity of the results across both linguistic units and objects suggests that this constraint may be a fundamental aspect of the cognitive system.

References

Biederman, I. 1972. Perceiving real-world scenes. *Science* 177: 77-80.

Biederman, I., Mezzanotte, R. J., and Rabinowitz, J. C. 1982. Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology* 14: 143-177.

Boyce, S. J., Pollatsek, A., and Rayner, K. 1989. Effect of background information on object identification. *Journal of Experimental Psychology: Human Perception and Performance* 15: 556-566.

Cohen, J. 1977. *Statistical power analysis for the behavioral sciences.* New York: Academic Press.

Connine, C. M. 1987. Constraints on interactive processes in auditory word recognition: The role of sentence context. *Journal of Memory and Language* 26: 527-538.

Fodor, J. A. 1983. *The modularity of mind.* Cambridge, MA: MIT Press.

Kroll, J. F. 1990. Recognizing words and pictures in sentence contexts: A test of lexical modularity. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 16: 747-759.

Massaro, D. W. 1989. Testing between the TRACE model and the fuzzy logical model of speech perception. *Cognitive Psychology* 21: 398-421.

Masson, M. E. J. 1988. The interaction of sentence context and perceptual analysis in word identification. *Memory and Cognition* 16: 489-496.

McClelland, J. L. 1991. Stochastic interactive processes and the effect of context on perception. *Cognitive Psychology* 23: 1-44.

McClelland, J. L., and Elman, J. L. 1986. The TRACE model of speech perception. *Cognitive Psychology* 18: 1-86.

McClelland, J. L., and Rumelhart, D. E. 1981. An interactive activation model of context effects in letter perception. Part I: An account of basic findings. *Psychological Review* 88: 375-407.

Oden, G. C., and Massaro, D. W. 1978. Integration of featural information in speech perception. *Psychological Review* 85: 172-191.

Reicher, G. M. 1969. Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology* 81: 274-280.

Reinitz, M. T., Wright, E, and Loftus, G. R. 1989. Effects of semantic priming on visual encoding of pictures. *Journal of Experimental Psychology: General* 118: 280-297.

Rueckl, J. G., and Oden, G. C. 1986. The integration of contextual and featural information during word identification. *Journal of Memory and Language* 25: 445-460.

Rumelhart, D. E. 1977. Toward an interactive model of reading. In S. Dornic (ed.), *Attention and performance VI* (pp. 573-603). Hillsdale, NJ: Erlbaum.

Snodgrass, J. G., and Corwin, J. 1988. Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General* 117: 34-50.

Snodgrass, J. G., and Vanderwart, M. 1980. A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory* 6: 174-215.

Stanovich, K. E., and West, R. F. 1983. On priming by a sentence context. *Journal of Experimental Psychology: General* 112: 1-36.