

Research Article

Features for Identification of Uppercase and Lowercase Letters

Daniel Fiset,¹ Caroline Blais,² Catherine Éthier-Majcher,² Martin Arguin,² Daniel Bub,¹ and Frédéric Gosselin²

¹Department of Psychology, University of Victoria, and ²Centre de Recherche en Neuropsychologie et Cognition, Département de Psychologie, Université de Montréal

ABSTRACT—*The determination of the visual features mediating letter identification has a long-standing history in cognitive science. Researchers have proposed many sets of letter features as important for letter identification, but no such sets have yet been derived directly from empirical data. In the study reported here, we applied the Bubbles technique to reveal directly which areas at five different spatial scales are efficient for the identification of lowercase and uppercase Arial letters. We provide the first empirical evidence that line terminations are the most important features for letter identification. We propose that these small features, represented at several spatial scales, help readers to discriminate among visually similar letters.*

Considerable evidence indicates that words are recognized by letters, rather than by the whole-word shape (Legge, Mansfield, & Chung, 2001; Paap, Newsome, & Noel, 1984; Peroa & Rosa, 1995; but see Allen & Emerson, 1991; Hadley & Healy, 1991). Most convincingly, Pelli, Farell, and Moore (2003) showed that a word is unreadable unless its letters are separately identifiable. Using a procedure with a noise mask, they demonstrated that the stimulus energy required for the recognition of a word increases linearly with the number of letters. Thus, whether a word can be identified is a straightforward function of the probability of identifying each of its individual letters. This is true even for the five most common three-letter words. Because these words have very familiar whole-word shapes, the authors concluded that global shape has little or no impact on the probability of recognizing a word correctly. Accordingly, a detailed knowledge of

the mechanisms involved in visual letter identification is fundamental to understanding reading.

Echoing this discussion about word recognition, a basic question about letter recognition is whether letters are perceived as global patterns or by features. A 45-year-old person who has read only 1 hr a day will have identified more than 1 billion letters (Pelli, Burns, Farell, & Moore-Page, 2006). Given this level of exposure, we would expect that literate human adults have a high level of efficiency in letter recognition. To reach high levels of efficiency, perceptual strategy matters. In fact, using an ideal-observer analysis, it is possible to demonstrate that the best solution—memory issues notwithstanding—to any object-identification task is template matching, rather than detection of individual features (Tjan, Braje, Legge, & Kersten, 1995). Does literate adults' extensive practice with letters lead them to use such an "optimal" strategy? Pelli et al. (2006) studied letter identification by measuring the contrast necessary for identifying a letter embedded in visual noise. To isolate the visual constraints on letter identification, they compared human performance across different fonts (e.g., Helvetica, Sloan, Kunstler) and alphabets (e.g., Arabic, Chinese). They found that human efficiency was mainly explained by Attneave and Arnoult's (1956) formula for perimetric complexity (i.e., sum of the inside and outside perimeters squared, divided by the "ink" area), with high complexity leading to low efficiency. This finding suggests that even though template matching may be the optimal solution in principle, normal human readers appear constrained to proceed by features, even for identifying simple and highly trained stimuli such as letters. The next question then becomes, what are the features upon which human letter recognition is based?

In the past three decades, insights about the features underlying letter recognition have come from experiments examining how recognition errors for each individual letter of the alphabet are distributed across the range of possible responses (i.e., the remainder of the alphabet; but see Petit & Grainger, 2002, for a

Address correspondence to Martin Arguin or to Frédéric Gosselin, Département de psychologie, Université de Montréal, C.P. 6128, Succ. Centre-ville, Montréal, Québec H3C 3J7, Canada, e-mail: martin.argin@umontreal.ca or frederic.gosselin@umontreal.ca.

different approach). Such confusability matrices, however, can be obtained only under very unusual conditions that elicit sufficiently high error rates (typically around 50%). For this reason, some researchers studied children who had not yet integrated the visual appearance of letters (E.J. Gibson, Gibson, Pick, & Osser, 1962). Others instead studied skilled readers, but with target letters exposed very briefly (Townsend, 1971) or with extremely low contrast (Geyer, 1977). Researchers have proposed sets of individual features to predict these letter-confusability matrices (Briggs & Hocevar, 1975; Geyer & DeWald, 1973; E.J. Gibson, 1969; E.J. Gibson et al., 1962; Laughery, 1969). In striking discrepancy with the findings of Pelli et al. (2006), however, template-matching models predict the outcome of letter-confusability experiments better than feature-based models do (Gervais, Harvey, & Roberts, 1984; Holbrook, 1975). An assumption that might resolve this lack of convergence is that the feature sets proposed by theories do not properly match those actually used by the human visual system. Indeed, such a mismatch would cause feature models to underperform in predicting human performance. This view highlights a fundamental weakness in using confusability matrices to indirectly reveal letter features: There is a vast gap between theories as to what constitutes a feature and the data used to test these theories. We propose an alternative approach for determining more directly which components of letters drive accurate recognition of them.

Recently, new tools referred to as classification image techniques have been developed to reveal aspects of a distal visual stimulus responsible for the measurable performance of observers in a specific task (e.g., Eckstein & Ahumada, 2002; Gosselin & Schyns, 2004). The underlying logic of these techniques is relatively simple. If specific visual information is necessary for the task at hand, depriving the observer of this information (using additive noise or a mask) will strongly impair his or her performance. In contrast, depriving the observer of nondiagnostic information will not substantially alter performance. For example, it seems obvious that for a reader to correctly discriminate between an *F* and an *E*, the bottom parts of the letters are necessary. Hiding these parts with a mask should have a great impact on discrimination between these letters. Classification image techniques, therefore, allow a direct empirical examination of the diagnostic features used by human observers in letter identification.

In the study we report here, skilled adult readers attempted to identify letters randomly sampled at different spatial scales (see Fig. 1). We performed multiple linear regression on the location of the samples and accuracy scores to establish which regions of the stimuli mediated letter recognition, and at which spatial scales. This method, called Bubbles (Gosselin & Schyns, 2001; e.g., Adolphs et al., 2005; B. Gibson, Lazareva, Gosselin, Schyns, & Wasserman, 2007; Nielsen, Logothetis, & Rainer, 2006; Smith, Gosselin, & Schyns, 2006), belongs to the general classification image approach. The results obtained using this

direct and unbiased method were compared with the various features that have been proposed to explain letter recognition.

METHOD

Participants

Six graduate students from the Université de Montréal took part in this experiment. All had normal or corrected-to-normal visual acuity.

Stimuli

The stimuli were the 26 letters of the Roman alphabet displayed in lowercase (152-point Arial font) and uppercase (117-point Arial font), and printed in dark gray (luminance of 2.1 cd/m²) against a light-gray background (luminance of 57.3 cd/m²). On average, both lowercase and uppercase letters subtended 1.35° of visual angle horizontally (64 pixels wide by 84 pixels high for uppercase letters, 64 pixels wide by 99 pixels high for lowercase letters). To reveal the visual features diagnostic for letter identification, we used Bubbles (Gosselin & Schyns, 2001). In a Bubbles experiment, stimulus information is randomly sampled, and multiple linear regression performed on the samples' locations and corresponding accuracy scores reveals which parts of the stimuli, on the dimensions that were sampled, are correlated with performance. We sampled letter stimuli in image space (*x*- and *y*-coordinates) and at varying spatial scales to uncover which letter parts are most correlated with letter identification.

The steps involved in creating an experimental stimulus were as follows: A letter stimulus was decomposed in five spatial-frequency bands (128–64, 64–32, 32–16, 16–8, and 8–4 cycles/image, or 32–16, 16–8, 8–4, 4–2, and 2–1 cycles/letter; the remaining bandwidth served as constant background; see Fig. 1, first row), using the Laplacian pyramid (Burt & Adelson, 1983). The letter information at the five scales was then sampled using an opaque mask punctured by randomly located Gaussian holes (henceforth called “bubbles”) to avoid introducing spatial-frequency artifacts. The size of the bubbles was adjusted according to frequency band so that each bubble revealed 1.5 cycles of spatial information (Fig. 1, second row). Because the size of the bubbles increased as the spatial scale became coarser, the number of bubbles differed across scales to keep the size of the sampled area constant across frequency bands. Finally, the information revealed by the bubbles was fused across the five frequency bands to produce an experimental stimulus (Fig. 1, third row). For each letter stimulus, the total number of bubbles was adjusted on a trial-by-trial basis to maintain a correct identification rate of 52% (approximately halfway between chance level—3.85%—and perfect identification, which is optimal for the multiple linear regression).

Procedure

Each participant completed 100 blocks of 260 trials for each letter case, for a total of 52,000 trials per participant. Such a




Fig. 1. Illustration of the stimulus-generation process. Each original letter (upper left) was first decomposed into five spatial-frequency bandwidths of one octave each (top row). Each bandwidth was then independently sampled with randomly positioned Gaussian windows (i.e., bubbles), so that sparse information was revealed (middle row). The information samples were summed across the five scales to produce an experimental stimulus (bottom row).

large number of trials was necessary to enable reliable statistical inference about the correlations between letter parts and accuracy. The uppercase and lowercase letters were presented in separate blocks, and each participant completed all the blocks in a particular case before moving to the other case. Half the participants began with the lowercase letters.

Each trial began with a fixation cross displayed at the center of the screen for 494 ms. The cross was then immediately replaced by a “bubbled” letter (see the stimulus at the bottom right of Fig 1 for an example), which remained on the screen for 200 ms. The participant had to identify the letter and then press the appropriate key (e.g., the “a” key when the letter was a or A) on a keyboard. No feedback was provided. The next trial started as soon as the software computed the next stimulus (approximately 500 ms after the participant’s response).

The stimuli were displayed on a 21-in. monitor set with a refresh rate of 75 Hz and calibrated to allow a linear manipulation of luminance. The experiment was run on a PC-Pentium IV computer. The experimental program was written in Matlab, using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). The viewing distance was maintained constant at 107 cm using a chin rest.

RESULTS AND DISCUSSION

To achieve 52% correct performance, participants needed, on average, 54.2 bubbles for lowercase letters and 30.9 bubbles for uppercase letters (see Table 1 for the number of bubbles required for each letter, on average). To pinpoint the features that different observers used to discriminate letters, we performed a least squares multiple linear regression on the bubble masks and

accuracy data. The plane of regression coefficients yielded by this operation is called a classification image. To obtain a classification image, we computed the *correct plane* by summing all the bubble masks (see Fig. 1, second row) that led to a correct answer and computed the *incorrect plane* by summing all the bubble masks that led to an incorrect answer; we then subtracted the incorrect plane from the correct plane. One such classification image was computed per letter, per case and per frequency band. To estimate the mean and the standard deviation of the distribution of the null hypothesis (i.e., no correlation between accuracies and sampled stimulus information), we repeated this procedure on permuted accuracies. This mean and this standard deviation were used to calculate the *Z* scores of the classification images. To determine the letter information significantly correlated with accuracy, we applied the pixel test to the *Z*-scored classification images ($p < .01$). The statistical threshold provided by this test corrects for multiple comparisons while taking the spatial correlation inherent to structured images into account (Chauvin, Worsley, Schyns, Arguin, & Gosselin, 2005). In Figure 2, this effective information—that is, the statistically thresholded classification image—is shown in red for each lowercase and uppercase letter; the classification images are superimposed on the corresponding letters, in gray, to help with interpretation.

A first glance at the left-most columns of Figure 2 may give the impression that the letters are almost completely revealed. However, this is far from the case. Only 32% and 24% of the ink area of the uppercase and lowercase letters, respectively, is depicted. The impression of completeness is due to the fact that what is revealed are the most informative regions of the letters. We believe that this is a rather impressive demonstration that the

TABLE 1

Average Number of Bubbles Participants Required to Maintain Performance at 52% Correct at the End of the Experiment

| Letter | Uppercase | Lowercase |
|--------|-----------|-----------|
| A | 16 | 45 |
| B | 38 | 44 |
| C | 42 | 84 |
| D | 33 | 35 |
| E | 37 | 54 |
| F | 41 | 55 |
| G | 31 | 37 |
| H | 33 | 55 |
| I | 21 | 133 |
| J | 32 | 45 |
| K | 24 | 29 |
| L | 33 | 128 |
| M | 22 | 18 |
| N | 30 | 69 |
| O | 50 | 100 |
| P | 45 | 44 |
| Q | 37 | 45 |
| R | 29 | 47 |
| S | 32 | 57 |
| T | 32 | 61 |
| U | 34 | 64 |
| V | 23 | 45 |
| W | 12 | 20 |
| X | 16 | 25 |
| Y | 27 | 31 |
| Z | 23 | 37 |

experiment indeed succeeded in revealing the letter regions that effectively drive recognition performance (see also Fig. 3).

A comparison of the number of pixels that were significantly useful for letter identification at each frequency band revealed a clear advantage for the information between 2 and 4 cycles per letter (see Fig. 4). This analysis was performed by calculating the proportion of significant pixels that fell on each frequency band. Note that only the pixels falling directly on letter ink were included in this analysis and all those that follow.

The classification images allowed us to reevaluate the various proposals regarding potent letter features. The letters were decomposed into the full complement of local features that have been proposed in the literature, except for global features such as symmetry, cyclic change, and parallelism, which we did not consider. We also included terminations, a feature that had not been considered previously. For each letter of the alphabet and each case, we created 213 masks containing the following 10 feature classes: vertical, horizontal, slant tilted left, slant tilted right, curves opened up, curves opened down, curves opened left, curves opened right, terminations, and intersections. For example, the uppercase letter **A** was decomposed in eight masks, that is, one slant tilted left, one slant tilted right, one horizontal, two terminations, and three intersections. The terminations and



Fig. 2. Classification images for the human observers in the experiment. Results for lowercase Arial letters are on the left, and results for uppercase Arial letters are on the right. For each letter, the overall classification image is shown on the left, and the next five columns display in red the significant pixels for each spatial-frequency bandwidth (from fine to coarse).

intersections were defined as letter ink within a radius of 13 pixels from the center of the feature, according to our own identification. To make sure that the masks for terminations and intersections were independent of those for the other features, we subtracted the area corresponding to the terminations and intersections from the other feature masks. For example, in the uppercase letter **A**, the two slants and the horizontal did not contain the pixels of the three intersections and the two terminations.



Fig. 3. Oscar Wilde's (1899) famous quote, "The truth is rarely pure and never simple" (p. 17) displayed with letter "ink" correlated (top; see Fig. 2) and uncorrelated (bottom) with accuracy. The same quantity of letter ink was used in both cases.

To examine directly the correspondence between the features and the classification images, we disposed of the spatial-frequency dimension by collapsing it prior to smoothing, thereby retaining only the x and y dimensions. From these bidimensional classification images, we kept only the portions of the stimuli corresponding to the highest 5% of the regression coefficients. For each image, we then calculated the proportion of the total number of these pixels that fell on each feature and divided this proportion by the total number of pixels in that feature, thus normalizing the proportion for feature size (not adjusting for feature size only amplified the effects reported). After conducting this analysis for each letter separately, we combined the results by summing the proportions calculated for each feature class across the 26 letters of each case, and subsequently dividing this sum by the number of occurrences of that feature class within the alphabet. To reveal the relative importance of the features for the correct identification of letters, we normalized the grand sum across feature classes to 1. Figure 5 summarizes the results thus obtained. Terminations were, by far, the most important features for both letter cases, with scores of .30 and .35 for upper- and lowercase letters, respectively. In fact, terminations were, respectively, 1.5 and 1.8 times more important than horizontals, which constituted the next most important feature class for both upper- and lowercase letters.

To compare human observers' use of features for letter identification with optimal use of diagnostic information, we built an ideal-observer model, which used all the visual information available for uppercase and lowercase letter identification. The ideal observer was submitted to the same experiment as the human participants and performed the same number of trials per letter as human participants. The number of bubbles for each letter was set to the average number of bubbles used by the human participants for that letter. An adjustable quantity of white Gaussian noise was added to the letters prior to sampling them with Gaussian apertures, in order to equate human and model performance (i.e., 52% accuracy for each letter). On each trial, the model determined the Pearson correlation between the sparse input (i.e., the noisy letter revealed by bubbles) and each of the 26 letters of the relevant case as revealed by the same bubble mask. The categorization response was the letter with the highest correlation with the stimulus.

The same feature analysis conducted with the human data was performed on the data for the ideal observer, and the results are presented in Figure 5. Note that the profile of usefulness of the different feature classes is flatter for the ideal observer than for humans. For the ideal observer, the difference between the first and second most important features is 2% for both lowercase and uppercase letters, whereas for humans, the difference is 16% for




Fig. 4. Relative use of the five sampled spatial-frequency bandwidths for recognition of uppercase and lowercase Arial letters.




Fig. 5. Relative use of 10 letter-feature classes for recognition of Arial letters. Results are shown for human observers and for the ideal observer. The letter-feature classes are ordered in terms of their usefulness to humans, with the least useful feature on the left; for each feature, results for lowercase letters are presented to the right of results for uppercase letters.

lowercase letters and 10% for uppercase letters. Most important, the usefulness of the terminations was much lower for the ideal observer (ranked 5th and 6th out of 10 for lowercase and uppercase letters, respectively) than for human observers. This indicates that the importance of the terminations for human participants is largely attributable to constraints imposed by properties of the human visual system, rather than to constraints that are exclusively determined by the stimulus set.

GENERAL DISCUSSION

We used Bubbles, a classification image technique, to reveal the letter areas responsible for accurate letter identification. Although this method has been applied successfully to face recognition, it had never been applied to letter identification. Figure 4 clearly shows that the frequency bandwidth from 2 to 4 cycles per letter conveys the most potent visual information for letter identification. This finding is congruent with results obtained by other researchers (Chung, Legge, & Tjan, 2002; Ginsburg, 1980; Legge, Pelli, Rubin, & Schleske, 1985; Majaj, Pelli, Kurshan, & Palomares, 2002; Parish & Sperling, 1991; Solomon & Pelli, 1994), who found that letters subtending 1.35°

of visual angle, as in the present study, are optimally masked by noise containing spatial frequencies around 3.1 cycles per letter. However, the 260 classification images of Figure 2 tell a more subtle and complete story. For example, the uppercase letters **C** and **G** share three curves, with convexities directed upward, to the left, and downward. The classification images show that one of these features (the curve opening on the right) is efficiently processed in low-spatial-frequency bandwidths. But identification of these letters also requires small features (the intersection and horizontal bar for **G** and the lower-part termination for **C**) represented by several spatial-frequency bandwidths. All spatial frequencies, including relatively high spatial frequencies, are necessary to resolve the edges of these small features. In fact, in determining the relative importance of letter features, we found that terminations, although relatively small, were the most distinctive features (Fig. 5). For example, the inferior termination of the uppercase letter **C** clearly allows the discrimination of this letter from the uppercase letters **G**, **Q**, and **O**, and is in fact sufficient for the correct identification of **C**.

In our feature analysis of the classification images of the ideal observer (see Fig. 5), terminations ranked 5th (lowercase letters) and 6th (uppercase letters) out of the 10 potential features

considered. These rankings are markedly different from those of the human observers. The stimulus areas effective for the ideal observer are those that are most diagnostic for the unique identification of the target letter—that is, for its discrimination against the remainder of the alphabet—in the absence of any other constraints. The divergence in the effective features for the ideal observer and humans is thus (at least partially; we discuss another potential determinant later) attributable to the fact that, in addition to needing to use features that can indeed discriminate among letters, humans are constrained by the constitution and organization of their visual system. This observation suggests, for instance, that the great importance of terminations for human participants results from an interaction between the relative diagnosticity of this feature and a disproportionately strong disposition of the human visual system to encode it (relative to other features). Some neurons in the primary visual cortex of monkeys respond to terminations. These are a subset of the cells with strongly end-stopped receptive fields (hypercomplex cells; Hubel & Wiesel, 1968). Moreover, it is believed that a number of these special cells converge to single neurons in V2 (von der Heydt & Peterhans, 1989). This might provide an early mechanism for letter identification. In contrast, humans' poor use of features such as verticals and curves opening up, which are highly effective for the ideal observer (see Fig. 5), suggests that the human visual system may be poorly equipped to process such features.

The feature analysis reported in this article is the first, to the best of our knowledge, to demonstrate the crucial importance of terminations for letter identification in humans. E.J. Gibson (1969) suggested that discontinuities were important for letter recognition, and terminations can be construed as discontinuities, but her proposal remained rather vague and was not substantiated empirically.

Why are small features such as terminations so important for letter identification in humans? Terminations are clear discontinuities in bars and curves and thus provide reliable information about the absence of intersections—also called coterminations in the object-recognition literature—in the target letter. An informal examination of other fonts suggests that the presence and the relative locations of terminations and intersections might be font-invariant properties of letters. In fact, it is relatively obvious that one can create an A without any slant. However, the two central intersections and the two terminations are critical properties of the letter A. This analysis suggests that a novel font would remain identifiable as long as a subset of these small features is available for visual extraction. It also points to another possible determinant of the divergence between humans' and the ideal observer's feature use, namely, the invariant diagnosticity of terminations across fonts, to which our ideal observer was obviously insensitive. Previous attempts to improve reading speed in individuals with low vision by filtering word images with a bandpass in the mid to high spatial-frequency range led to equivocal results (e.g., Fine & Peli, 1995). These failures may be attributable to the fact that the

diagnosticity (with respect to letter identity) of the visual features of letters in that spatial-frequency range was not improved by this manipulation. The results presented in this article should allow the creation of a font in which the diagnosticity of the features most effective for letter identification is enhanced. It remains to be seen whether such a font would lead to faster letter recognition and, in turn, to faster word recognition in normal readers and individuals with letter-by-letter dyslexia.

Acknowledgments—We thank the subjects who took part in this study. This research was supported by a grant from the Canadian Institute of Health Research (CIHR) to Martin Arguin, Frédéric Gosselin, and Daniel Bub; by grants from the Natural Sciences and Engineering Research Council of Canada (NSERC) to Martin Arguin and to Frédéric Gosselin; by a scholarship from the James S. McDonnell Foundation (Perceptual Expertise Network) and a postdoctoral scholarship from the Fonds Québécois de Recherche en Nature et Technologies (FQRNT) to Daniel Fiset; and by a FQRNT graduate scholarship to Caroline Blais.

REFERENCES

- Adolphs, R., Gosselin, F., Buchanan, T.W., Tranel, D., Schyns, P.G., & Damasio, A.R. (2005). A mechanism for impaired fear recognition after amygdala damage. *Nature*, *433*, 68–72.
- Allen, P.A., & Emerson, P.L. (1991). Holism revisited: Evidence for independent word-level and letter-level processors during word and letter processing. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 489–511.
- Attneave, F., & Arnoult, M.D. (1956). The quantitative study of shape and pattern perception. *Psychological Bulletin*, *53*, 452–471.
- Brainard, D.H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.
- Briggs, R., & Hocevar, D.J. (1975). A new distinctive theory for upper case letters. *The Journal of General Psychology*, *93*, 87–93.
- Burt, P., & Adelson, E. (1983). The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, *31*, 532–540.
- Chauvin, A., Worsley, K.J., Schyns, P.G., Arguin, M., & Gosselin, F. (2005). Accurate statistical tests for smooth classification images. *Journal of Vision*, *5*(9), Article 1. Available <http://journalofvision.org/5/9/1/>
- Chung, S.T.L., Legge, G.E., & Tjan, B.S. (2002). Spatial-frequency characteristics of letter identification in central and peripheral vision. *Vision Research*, *42*, 2137–2152.
- Eckstein, M.P., & Ahumada, A.J. (Eds.). (2002). Classification images: A tool to analyze visual strategies [Special issue]. *Journal of Vision*, *2*(1). Retrieved April 25, 2004, from <http://journalofvision.org/2/>
- Fine, E.M., & Peli, E. (1995). Enhancement of text for the visually impaired. *Journal of the Optical Society of America A*, *12*, 1439–1447.
- Gervais, M.J., Harvey, L.O., & Roberts, J.O. (1984). Identification confusions among letters of the alphabet. *Journal of Experimental Psychology: Human Perception and Performance*, *10*, 655–666.
- Geyer, L.H. (1977). Recognition and confusion of the lowercase alphabet. *Perception & Psychophysics*, *22*, 487–490.

- Geyer, L.H., & DeWald, C.G. (1973). Feature lists and confusion matrices. *Perception & Psychophysics*, *14*, 471–482.
- Gibson, B., Lazareva, O.F., Gosselin, F., Schyns, P.G., & Wasserman, E.A. (2007). Non-accidental properties underlie shape recognition in mammalian and non-mammalian vision. *Current Biology*, *17*, 336–340.
- Gibson, E.J. (1969). *Principles of perceptual learning and development*. New York: Meredith.
- Gibson, E.J., Gibson, J.J., Pick, A.D., & Osser, H.A. (1962). A developmental study of the discrimination of letter-like forms. *Journal of Comparative and Physiological Psychology*, *55*, 897–908.
- Ginsburg, A.P. (1980). Specifying relevant spatial information for image evaluation and display design: An explanation of how we see certain objects. *Proceedings of the SID*, *21*, 219–227.
- Gosselin, F., & Schyns, P.G. (2001). Bubbles: A technique to reveal the use of information in recognition. *Vision Research*, *41*, 2261–2271.
- Gosselin, F., & Schyns, P.G. (2004). A picture is worth thousands of trials: Rendering the use of visual information from spiking neurons to recognition. *Cognitive Science*, *28*, 141–146.
- Hadley, J.A., & Healy, A.F. (1991). When are reading units larger than the letter? Refinement of the Unitization Reading Model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 1062–1073.
- Holbrook, M.B. (1975). A comparison of methods for measuring the inter-letter similarity between capital letters. *Perception & Psychophysics*, *17*, 532–536.
- Hubel, D.H., & Wiesel, T.N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology (London)*, *195*, 215–243.
- Laughery, K.R. (1969). Computer simulation of short-term memory: A component decay model. In G.T. Bower & J.T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 3, pp. 135–200). New York: Academic Press.
- Legge, G.E., Mansfield, J.S., & Chung, S.T.L. (2001). Psychophysics of reading: XX. Linking letter recognition to reading speed in central and peripheral vision. *Vision Research*, *41*, 725–743.
- Legge, G.E., Pelli, D.G., Rubin, G.S., & Schleske, M.M. (1985). Psychophysics of reading: I. Normal vision. *Vision Research*, *25*, 239–252.
- Majaj, N.J., Pelli, D.G., Kurshan, P., & Palomares, M. (2002). The role of spatial frequency channels in letter identification. *Vision Research*, *42*, 1165–1184.
- Nielsen, K.J., Logothetis, N.K., & Rainer, G. (2006). Discrimination strategies of humans and rhesus monkeys for complex visual displays. *Current Biology*, *16*, 814–820.
- Paap, K.R., Newsome, S.L., & Noel, R.W. (1984). Word shape's in poor shape for the race to the lexicon. *Journal of Experimental Psychology: Human Perception and Performance*, *10*, 413–428.
- Parish, D.H., & Sperling, G. (1991). Object spatial frequencies, retinal spatial frequencies, noise, and the efficiency of letter discrimination. *Vision Research*, *31*, 1399–1415.
- Pelli, D.G. (1997). The Video Toolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Pelli, D.G., Burns, C.W., Farell, B., & Moore-Page, D.C. (2006). Feature detection and letter identification. *Vision Research*, *46*, 4646–4674.
- Pelli, D.G., Farell, B., & Moore, D.C. (2003). The remarkable inefficiency of word recognition. *Nature*, *423*, 752–756.
- Peroa, M., & Rosa, E. (1995). Does 'whole word shape' play a role in visual word recognition? *Perception & Psychophysics*, *65*, 784–795.
- Petit, J.P., & Grainger, J. (2002). Masked partial priming of letter perception. *Visual Cognition*, *9*, 337–354.
- Smith, M.L., Gosselin, F., & Schyns, P.G. (2006). Perceptual moments of conscious visual experience. *Proceedings of the National Academy of Sciences, USA*, *103*, 5626–5631.
- Solomon, J.A., & Pelli, D.G. (1994). The visual filter mediating letter identification. *Nature*, *369*, 395–397.
- Tjan, B.S., Braje, W.L., Legge, G.E., & Kersten, D. (1995). Human efficiency for recognizing 3-D objects in luminance noise. *Vision Research*, *35*, 3053–3069.
- Townsend, J.T. (1971). Alphabetic confusion: A test of models for individuals. *Perception & Psychophysics*, *9*, 449–454.
- von der Heydt, R., & Peterhans, E. (1989). Mechanisms of contour perception in monkey visual cortex. I. Lines of pattern discontinuity. *Journal of Neuroscience*, *9*, 1731–1748.
- Wilde, O. (1899). *The importance of being earnest*. London: Leonard Smithers and Co.

(RECEIVED 7/18/07; REVISION ACCEPTED 5/29/08)