

A crowding model of visual clutter

Ronald van den Berg

Institute of Mathematics and Computing Science and School
of Behavioral and Cognitive Neurosciences,
University of Groningen, Groningen,
The Netherlands



Frans W. Cornelissen

Laboratory for Experimental Ophthalmology,
School of Behavioral and Cognitive Neurosciences,
University Medical Center Groningen,
University of Groningen, Groningen,
The Netherlands



Jos B. T. M. Roerdink

Institute of Mathematics and Computing Science and
School of Behavioral and Cognitive Neurosciences,
University of Groningen, Groningen,
The Netherlands



Visual information is difficult to search and interpret when the density of the displayed information is high or the layout is chaotic. Visual information that exhibits such properties is generally referred to as being “cluttered.” Clutter should be avoided in information visualizations and interface design in general because it can severely degrade task performance. Although previous studies have identified computable correlates of clutter (such as local feature variance and edge density), understanding of why humans perceive some scenes as being more cluttered than others remains limited. Here, we explore an account of clutter that is inspired by findings from visual perception studies. Specifically, we test the hypothesis that the so-called “crowding” phenomenon is an important constituent of clutter. We constructed an algorithm to predict visual clutter in arbitrary images by estimating the perceptual impairment due to crowding. After verifying that this model can reproduce crowding data we tested whether it can also predict clutter. We found that its predictions correlate well with both subjective clutter assessments and search performance in cluttered scenes. These results suggest that crowding and clutter may indeed be closely related concepts and suggest avenues for further research.

Keywords: visual clutter, crowding, visual search, visualization

Citation: van den Berg, R., Cornelissen, F. W., & Roerdink, J. B. T. M. (2009). A crowding model of visual clutter. *Journal of Vision*, 9(4):24, 1–11, <http://journalofvision.org/9/4/24/>, doi:10.1167/9.4.24.

Introduction

The main purpose of information visualization and graphical design in general is to present information in a form that facilitates understanding and improves task performance. A pressing problem today is that while data sets continue to grow in size and complexity, computer displays are limited in their capacity to show visual information. At the same time, the human visual system is limited with respect to its capacity to process incoming visual information. Advances in visualization research have provided a variety of techniques to deal with this problem of “information overload,” such as “filtering,” “zooming,” and “focus + context” (Shneiderman, 1996). What all these methods seem to aim at is to reduce “clutter” without hindering task performance.

While most of us have an implicit sense of what it means for a display to be cluttered, it is not at all obvious

how to make this explicit, let alone how to quantify and predict it. Clutter can be defined in various ways. First, it can refer to the subjective impression of “visual chaos.” However, in order to study clutter, it is useful to have an operational definition. Rosenholtz, Li, Mansfield, and Jin (2005) therefore proposed to define clutter as “the state in which excess items, or their representation or organization, lead to a degradation of performance at some task.”

Based on the operational definition of clutter, we can identify two factors that appear to play an important role in clutter: information density and information layout. This implies that there are also two ways to deal with clutter, viz., reducing the information density and changing the layout.

Previous studies that have addressed the issue of information density in relation to clutter include Woodruff, Landay, and Stonebraker (1998), who developed a system to keep information density constant in interactive displays, and Yang-Peláez and Flowers (2000),

who proposed an information content measure of visual displays based on Shannon’s information criterion. In addition, Oliva, Mack, Shrestha, and Peepers (2004) studied how visual complexity for real-world images is represented by a cognitive system. Although they did not identify a single perceptual dimension that fully accounts for visual complexity, they did find that subjects reported variety and quantity of objects and colors, and their spatial arrangement (thus, “clutter”) as the most important factors. Furthermore, in a recent paper by Baldassi, Megna, and Burr (2006), it was shown that perceptual clutter not only leads to increases in (orientation) judgment errors but also in perceived signal strength and confidence in erroneous judgments. An implication of these results is that an increase in the amount of displayed information not only leads to more error-prone judgments but, paradoxically, also to more confidence in erroneous decisions. Finally, the most comprehensive studies of visual clutter in information displays that we know of are those carried out by Rosenholtz, Li, Mansfield, and Jin (2005) and Rosenholtz, Li, and Nakano (2007). These authors hypothesized that clutter is inversely related to saliency, which had earlier been shown to relate to local feature variance (Rosenholtz, 2001). They proposed a model that estimates clutter by measuring local variance in several visual feature channels. Their experimental data showed that there is indeed a strong correlation between local feature variance and subjective clutter assessments of images. However, the question *why* feature variance correlates with clutter remained unanswered.

The present work is motivated by the expectation that clutter can be measured and controlled more adequately when we have an understanding of its roots. We hypothesize that clutter has its basis in visual “crowding,” that is, the (extensively studied) phenomenon that closely spaced objects hinder each other’s recognition, most notably in the periphery of the visual field. This hypothesis is based on a number of conspicuous similarities between both phenomena. First, both crowding and clutter increase with information density. Second, both phenomena are most prominent in the periphery of the visual field, yet cannot be (fully) explained by acuity loss. Third, one of the defining aspects of clutter is that it degrades performance on visual tasks (Baldassi et al., 2006; Beck, Lohrenz, Trafton, & Gendron, 2008; Bravo & Farid, 2008; Rosenholtz et al., 2005). The same is true for crowding. Significant decreases in search performance can be observed as a result of increased numbers of fixations, increased fixation durations, and increased saccade amplitudes in crowded search tasks (Vlaskamp & Hooge, 2005).

Although the neural basis of crowding is not yet understood, evidence is accumulating that it involves feature integration occurring over inappropriately large areas (Levi, 2008; Pelli & Tillman, 2008). We hypothesize that it is this integration—which will often result in information loss—that underlies both the performance degradation and the feeling of “confusion” that is

characteristically experienced when viewing cluttered displays. In order to test this hypothesis, we developed an algorithm that estimates how much information in an image is lost due to crowding (Model description section) and we evaluated the predictions of this model against subjective clutter assessments and search performance in cluttered scenes (Simulations and results section).

Background: Crowding

A peripherally viewed object that is easy to recognize when shown in isolation is much harder to identify when surrounded by other objects, especially when object spacing is small (Figure 1). This effect was first described in the 1920s, when Korte (1923) discovered that flanking a letter by other letters makes it more difficult to recognize. This phenomenon is now popularly known as “crowding” (Stuart & Burian, 1962). The crowding effect has since been studied extensively (reviewed in Levi, 2008 and Pelli & Tillman, 2008) and has led to the view that vision is usually limited by object spacing rather than size. Much of the literature concentrates on studying the spatial extent over which crowding acts, which is commonly referred to as the “critical spacing” and considered by many the defining property of crowding. Time and again, researchers found that the critical spacing for letter and shape recognition scales with eccentricity in the visual field. This fundamental crowding property is now sometimes referred to as “Bouma’s law” (Pelli & Tillman, 2008), after its original discoverer (Bouma, 1970). Recent studies suggest that Bouma’s law is a universal property of vision. It has been demonstrated not to be confined to letter and shape recognition but to hold for a wide range of stimuli and tasks, including the identification of orientation (Wilkinson, Wilson, & Ellemberg, 1997), object size, hue, saturation of colors (van den Berg, Roerdink, & Cornelissen, 2007), and face recognition (Pelli et al., 2007).

Several theories have been proposed to explain crowding. While these proposals vary widely in detail and scope, there seems to be a growing consensus toward a *two-stage* model, consisting of a feature detection stage followed by an integration stage. Proponents of this theory argue that whereas feature detection remains unaffected in

A B A + ABA

Figure 1. An example of crowding. The two B’s are at equal distance from the fixation cross. On the left, the spacing between the letters is approximately 0.5 times the eccentricity of the B. On the right, letter spacing is approximately 0.2 times the eccentricity of the B. While the central item on the left can easily be recognized when fixating the cross, the central item on the right cannot and appears to be jumbled with its neighbors.

crowding, integration happens over inappropriately large areas, sometimes referred to as “integration fields” (Pelli, Palomares, & Majaj, 2004). Because of Bouma’s rule, these putative integration fields should have a size that equals roughly 0.4 times the eccentricity of its center position. Furthermore, the relation between object spacing and crowding magnitude (e.g., Pelli et al., 2004; van den Berg et al., 2007) suggests a *weighted* form of integration over these fields, i.e., non-target objects in the center of a field contribute more than objects near the border of a field.

Model description

From a computational standpoint, crowding appears to be the result of feature pooling, carried out by (weighted)

integration fields with sizes proportional to retinal eccentricity (Pelli et al., 2004). This inevitably results in a loss of perceived detail of objects, in particular in the periphery, where integration fields are large. We conjecture that at a subjective level this loss of information is responsible for the feeling of “confusion” that people experience when viewing a cluttered scene. At a more objective level, we suspect that it is also the reason for the elevated recognition thresholds, longer inspection times, and increased number of fixations. If this is true, then the information loss due to crowding should be an apt indicator of visual clutter.

The amount of information loss can be estimated by simulating the putative integration fields and comparing the information content before and after integration. We implemented this idea in a model that consists of the following steps (see Figure 2 for a schematic description; each step is explained in more detail later in this section):

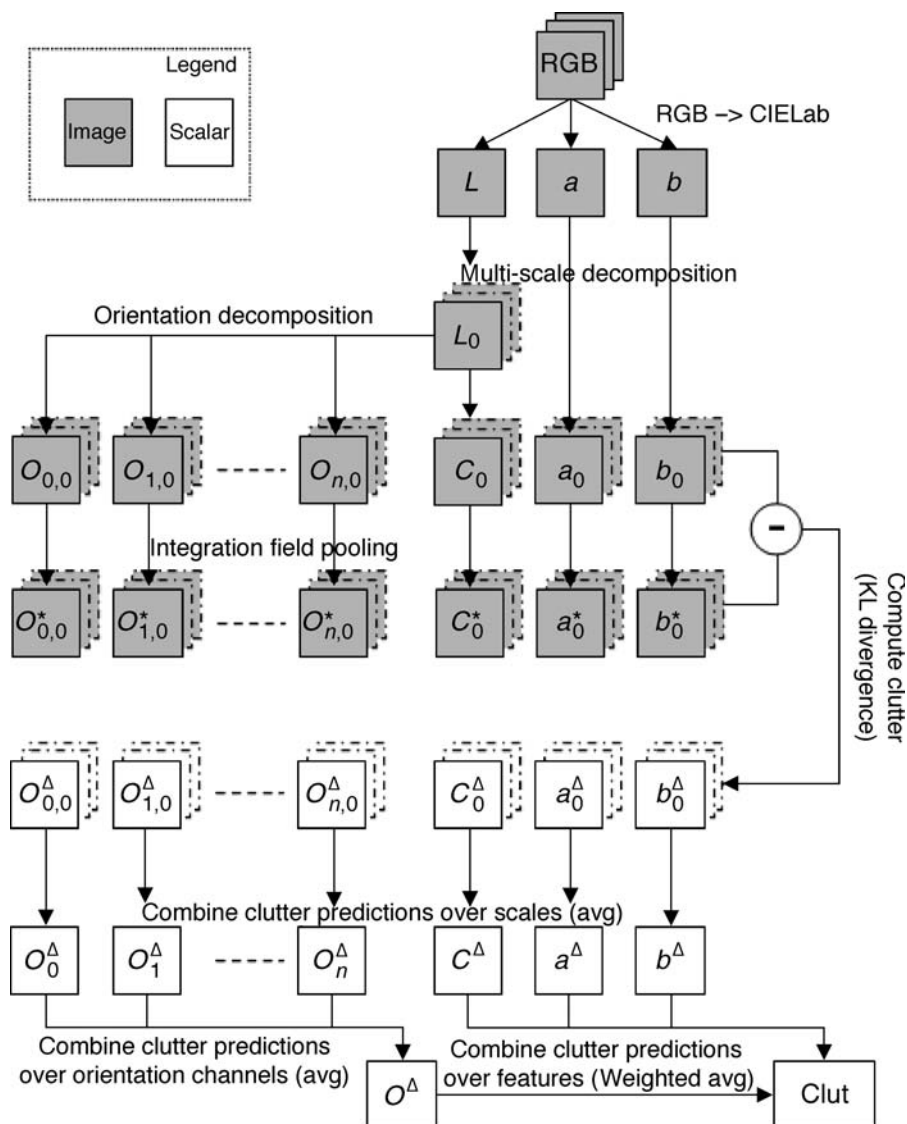


Figure 2. Schematic illustration of the crowding-based clutter measurement algorithm.

1. Convert the input (an sRGB image) to CIELab space. Output: a luminance image L_0 , a red/green image a_0 , and a blue/yellow image b_0 .
2. Perform a multi-scale decomposition of L_0 , a_0 , and b_0 (N scales). Output: a set of luminance images L_i , a set of red/green images a_i , and a set of blue/yellow images b_i , $i = 0 \dots N - 1$.
3. Perform an orientation decomposition of $L_0 \dots L_{N-1}$ (M orientations). Output: a set of orientation images $O_{i,j}$, $i = 0 \dots N - 1$, $j = 0 \dots M - 1$.
4. Perform contrast filtering of $L_0 \dots L_{N-1}$. Output: a set of contrast images C_i , $i = 0 \dots N - 1$.
5. Simulate crowding (integration fields), by performing local averaging of C_i , a_i , b_i , $O_{i,j}$. Output: images C_i^* , a_i^* , b_i^* , $O_{i,j}^*$, $i = 0 \dots N - 1$, $j = 0 \dots M - 1$.
6. Estimate for each image the amount of information loss in step 5. Output: clutter estimates $C\Delta_i$, $a\Delta_i$, $b\Delta_i$, $O\Delta_{i,j}$, $i = 0 \dots N - 1$, $j = 0 \dots M - 1$ (scalars).
7. Pool over scales and features. Output: image clutter prediction $CLUT$ (a scalar).

Step 1: RGB to CIELab conversion

The first step consists of decomposing the input RGB image into a set of feature channels that reflect the decomposition as it occurs in the human visual system, viz., into a luminance channel and two color channels (red/green and blue/yellow). The RGB to CIELab conversion gives a luminance (L) component and two color (a , b) components. This conversion is carried out in two steps; we first convert the RGB image to an XYZ image, which subsequently is converted to CIELab.

Step 2: Multi-scale decomposition

Next, the images are analyzed on multiple scales. For this purpose, N -level Gaussian pyramids for the L -, a -, and b -images are created (Burt & Adelson, 1983). In the experiments described below, the number of levels of the Gaussian pyramid was set to 3.

Step 3: Orientation decomposition

From the luminance images a number of orientation images are constructed. This is done by filtering the luminance images with oriented Gabor filters, with equally spaced orientations in the range $[0,180)$. We chose to use biologically motivated filters with non-classical receptive fields with lateral inhibition, as described in Grigorescu, Petkov, and Westenberg (2004). Briefly summarized, these center-surround filters are of the form $f = H(E - \alpha I)$, where E is the Gabor energy response to the center, I is the Gabor energy response to the

surround, α is a factor controlling the inhibition strength, and H is a non-linearity that clips negative values to zero (for details, consult the cited paper).

Prior to the orientation decomposition we filter the luminance image with a sigmoid kernel (with $\mu = \text{mean luminance of } L$, and $\sigma = \mu/10$). This reduces contrast differences across the image and, therefore, decorrelates the luminance and orientation channels. We checked the effect of this non-linearity by computing the mean correlation between the contrast and orientation channel (at scale 1) for the 25 images from the map sorting task (see Figure 7 below). It appeared that without applying the non-linearity, the correlation was 0.49, while the non-linear operation reduced it to 0.32.

In the experiments described below, we used a decomposition into 6 orientation images (0, 30, 60, 90, 120, 150 deg) and the inhibition factor α was set to 1.

Step 4: Contrast filtering

Using a Difference-of-Gaussians filter ($\sigma_1 = 2$; $\sigma_2 = 6$), the luminance images are converted into contrast images (negative values are clipped to zero).

Step 5: Pooling

The next step consists of carrying out, for all images created in steps 1–4, the feature integration that occurs in the putative integration fields. We chose to implement this step as a weighted averaging operation, in accordance with an earlier finding that crowded orientation signals are perceived as being averaged (Parkes, Lund, Angelucci, Solomon, & Morgan, 2001). The images were filtered with Gaussian kernels, so that the kernel width controls the size of the integration field. In the experiments below, the width (sigma) was set to 1/16th of the eccentricity of the integration field center (see Figure 3 for an example of the effect of this step).

With regard to the orientation domain, we note that averaging takes place within subbands and not over the entire orientation domain. As a consequence, predicted clutter will be higher for similar orientations than for dissimilar orientations. This is in line with the “feature similarity” effect reported in the crowding literature. The more similar two different objects are, the stronger they will crowd each other. In addition, orientation averaging only occurs when tilt differences are relatively small (hence, presenting patches with -45 and $+45$ deg tilt clearly does not result in observing 0 deg tilted patches).

As the filter kernel size scales with eccentricity, this step requires that we know the eccentricity of each integration field, i.e., it requires that a fixation location is defined. In order to obtain a clutter estimate that is relatively independent of where one is looking, we can

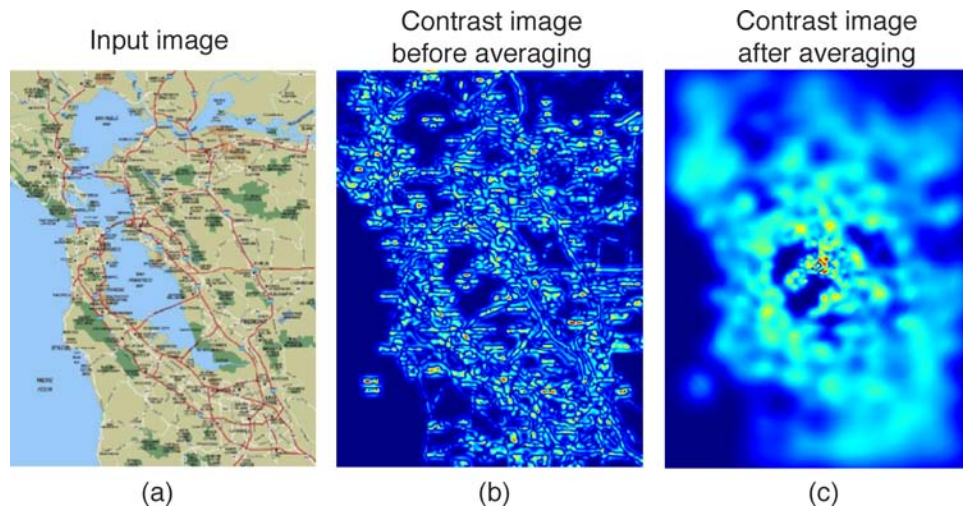


Figure 3. Example showing the effect of local feature averaging (with fixation set to the center of the image): (a) Input image, (b) contrast image before pooling, (c) contrast image after pooling.

repeat this step and all subsequent ones several times, with fixation set to different locations in the image, and then average the results. To assess to what extent the simulation results depend on the number of fixations chosen, we performed the following experiment. We let the model compute clutter values for 25 images. Based on these values, we ranked the images from least to most cluttered. We performed this with 1, 2, 4, 8, and 16 randomly chosen points of fixation. It appeared that the rankings produced for these different numbers of fixations were highly correlated (mean pairwise Spearman correlation was 0.91), indicating that the (ranking) results of our model only weakly depend on the number of fixations chosen, at least for the images used in the evaluation experiments that are described in the next section; this means that the amount of clutter in these images is apparently rather uniform over space. In the experiments reported here, we chose to use only a single fixation point, set to the center of the image.

Step 6: Determine information loss

As a measure for the amount of information loss in the integration step, we use a sliding window to locally compute the Kullback–Leibler (KL) divergence (Kullback & Leibler, 1951) between the input and output of the previous step. The KL divergence is a measure of the difference between two probability distributions P and Q and is computed as follows:

$$D_{KL}(P\|Q) = \sum_i P(i) \log \frac{P(i)}{Q(i)} \quad (1)$$

for an input region P (consisting of a set of pixels $P(i)$) and an output region Q . To obtain a global clutter value for an image, we average local KL divergence values over all image regions.

Step 7: Pool over scales and features

The last step consists of combining the clutter values for orientations, scales, and features, in order to obtain a global clutter estimate of the input image. We first combine orientations and scales, by averaging over orientation channels and, subsequently, over scales. After this step, we have one clutter value per feature channel. Since it is known from previous research that crowding does not affect all feature channels equally, we assign different weights to the features when combining them, thus computing a weighted average.

Simulations and results

Crowding

The defining property of crowding is that object recognition thresholds decrease with object spacing. The smallest spacing at which objects do not affect recognition of a target is called the “critical spacing” and is usually found to equal approximately 0.4 times the eccentricity of the target (Pelli & Tillman, 2008). To verify whether our model can reproduce this key property of crowding, we ran the following simulation (with parameters set to the values reported in the previous section). Stimuli consisted of images with 25 objects from Bravo and Farid’s (2004) study described below, with a size of approximately

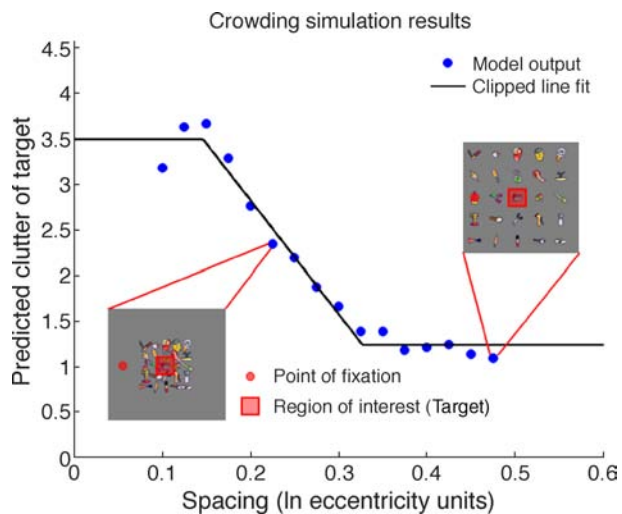


Figure 4. Effect of element spacing on predicted clutter (in the region of interest). Predicted clutter decreases with increased spacing, in a way that is very similar to the crowding effect (compare for example with Pelli et al., 2004 and van den Berg et al., 2007). These results demonstrate that our clutter-model computation gives output comparable to that occurring in crowding.

30 × 30 pixels each, organized in a regular 5 × 5 grid (Figure 4). We varied the spacing between objects from 20 to 120 pixels. With a fixation point set to 200 pixels away from the target image’s center, we computed local clutter in a 25 × 25 pixel region of interest located at the image center (the center object was thus defined as target object).

The results show that predicted clutter decreases with spacing in a similar way as found in crowding studies, and up to a (critical) spacing of about 0.33 times the eccentricity. We therefore conclude that our clutter model indeed demonstrates behavior akin to crowding.

Comparison with subjective clutter judgments

In order to evaluate how well the model performs in predicting clutter, and to compare its performance with the feature congestion model, we partly repeated the experiment from Rosenholtz et al. (2005). In that experiment twenty subjects were asked to sort 25 US maps (Figure 7) according to how cluttered they were perceived to be. Based on the obtained rankings, an average subjective ranking was computed and compared with the clutter ranking as produced by their feature congestion model.

Rosenholtz et al. found a significant (Spearman’s rank) correlation of 0.83 ($p < 0.001$) between subjective and model ranking. This was comparable to the correlation between subjects (which, on average, was 0.70 between every pair of subjects). This indicates that their local feature variance measure is a good indicator for perceived clutter and performs as well as is possible given the between-subject variance.

We used the same set of images as input to our model. All model parameters were fixed to the values reported in the Model description section. If we set the weights in step 6 equal for all channels we find a correlation of 0.82 ($p < 0.00001$) between the ranking produced by our model and the average subjective ranking. This is comparable to the correlation reported by Rosenholtz et al. (Figure 5a).

We obtain a slightly stronger correlation ($\rho = 0.84$, $p < 0.00001$) if we assign the color channels about half the weight of those of the orientation and contrast channels. This is in line with our earlier finding that crowding is stronger in the orientation channel than the color channel (van den Berg et al., 2007).

To compare the predictions of our crowding-based model and the feature congestion model, we computed

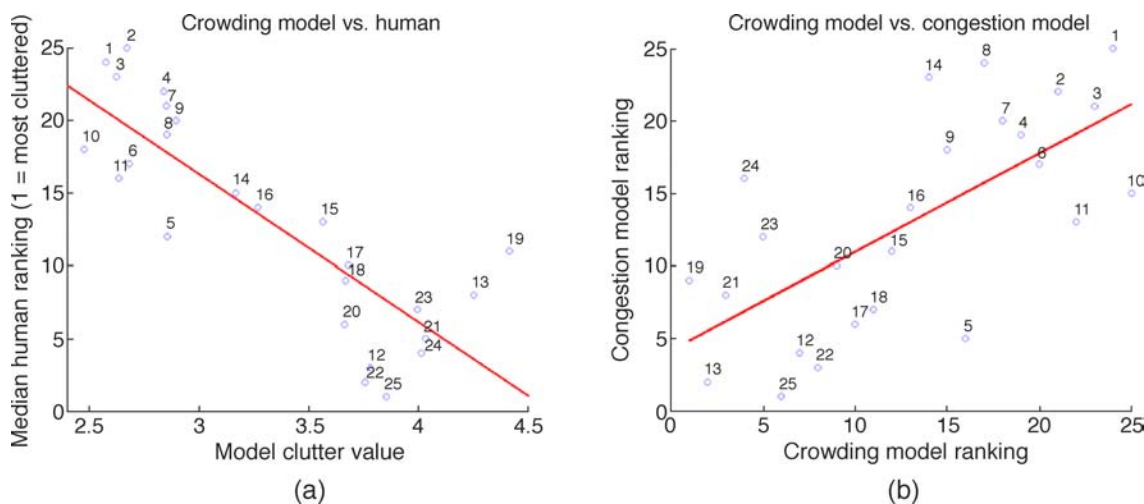


Figure 5. (a) Median subject ranking as a function of clutter as estimated by our crowding-based model (cf. Figure 2 from Rosenholtz et al., 2005). (b) Clutter rank order as predicted by our crowding-based model vs. rank order as predicted by the feature congestion model.

the correlation between their rankings. It is 0.68 ($p < 0.001$; Figure 5b). Clearly, even though the measures used by both models correlate, they obviously differ in their predictions (we will elaborate on this in the Discussion section).

Comparison with visual search performance in cluttered images

Bravo and Farid (2004) studied how clutter affects visual search. They performed a target present/absent search experiment with images that varied in terms of number of objects ($N = 6, 12, 24$), spatial arrangement (sparse versus cluttered layout), and distractor type (simple versus complex). Two sample images are shown in Figure 6.

Their main findings (Figure 6c) were that:


- i. search times are longer for cluttered layouts compared to sparse layouts;
- ii. search times increase faster (as a function of N) for cluttered layouts compared to sparse layouts;
- iii. search times for a cluttered layout are longer for images with complex distractors compared to images with simple distractors.

We used the full set of 960 images of Bravo and Farid's study as input to our model. Model parameters were set to the same values as in the two simulations that were described above. There was one difference however. Unlike the map images of Rosenholtz et al.'s study, Bravo and Farid's images have a clear figure/background separation. Since crowding is an adverse interaction between objects, and not objects and their background, we decided to ignore all background pixels in the averaging step.

The results are shown in Figure 6d. The predicted clutter curves are similar to the search time curves from Bravo and Farid's study (Figure 6c) in the following respects:


- i. images with cluttered layout are predicted to be more cluttered than images with sparse layout;
- ii. predicted clutter increases with the number of objects N in an image;
- iii. the dependence of predicted clutter on N is stronger for cluttered images than for sparse images (slope is about twice as large).

For the case of sparse layout our model predicts higher clutter for images with simple objects compared to images with complex objects. We were not able to identify the source of this result.

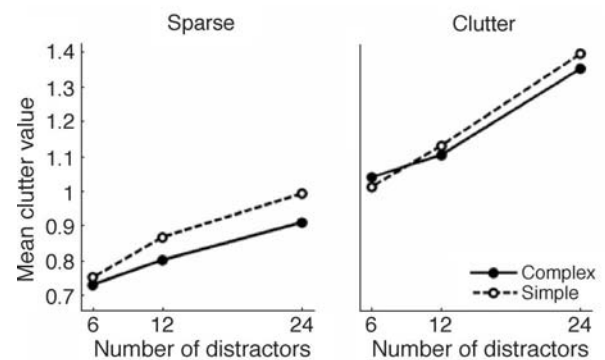


(a)

(b)



(c)



(d)

Figure 6. (a) Sample image from Bravo and Farid's (2004) study: $N = 6$, sparse arrangement, simple objects. (b) Another example: $N = 12$, cluttered arrangement, complex objects. (c) Human subject search time results from Bravo and Farid's study (data from Bravo & Farid, 2004). (d) Prediction results from our crowding-based clutter model.

Altogether, our model performs quite well on these data. In our view, this suggests that Bravo and Farid's manipulation of clutter (by varying layout, complexity, and number of distractors in their images) appears to have been largely the result of influencing crowding.

Discussion

The main aim of this study was to examine the hypothesis that crowding is an important, if not the main constituent of clutter. To do so, we constructed a model

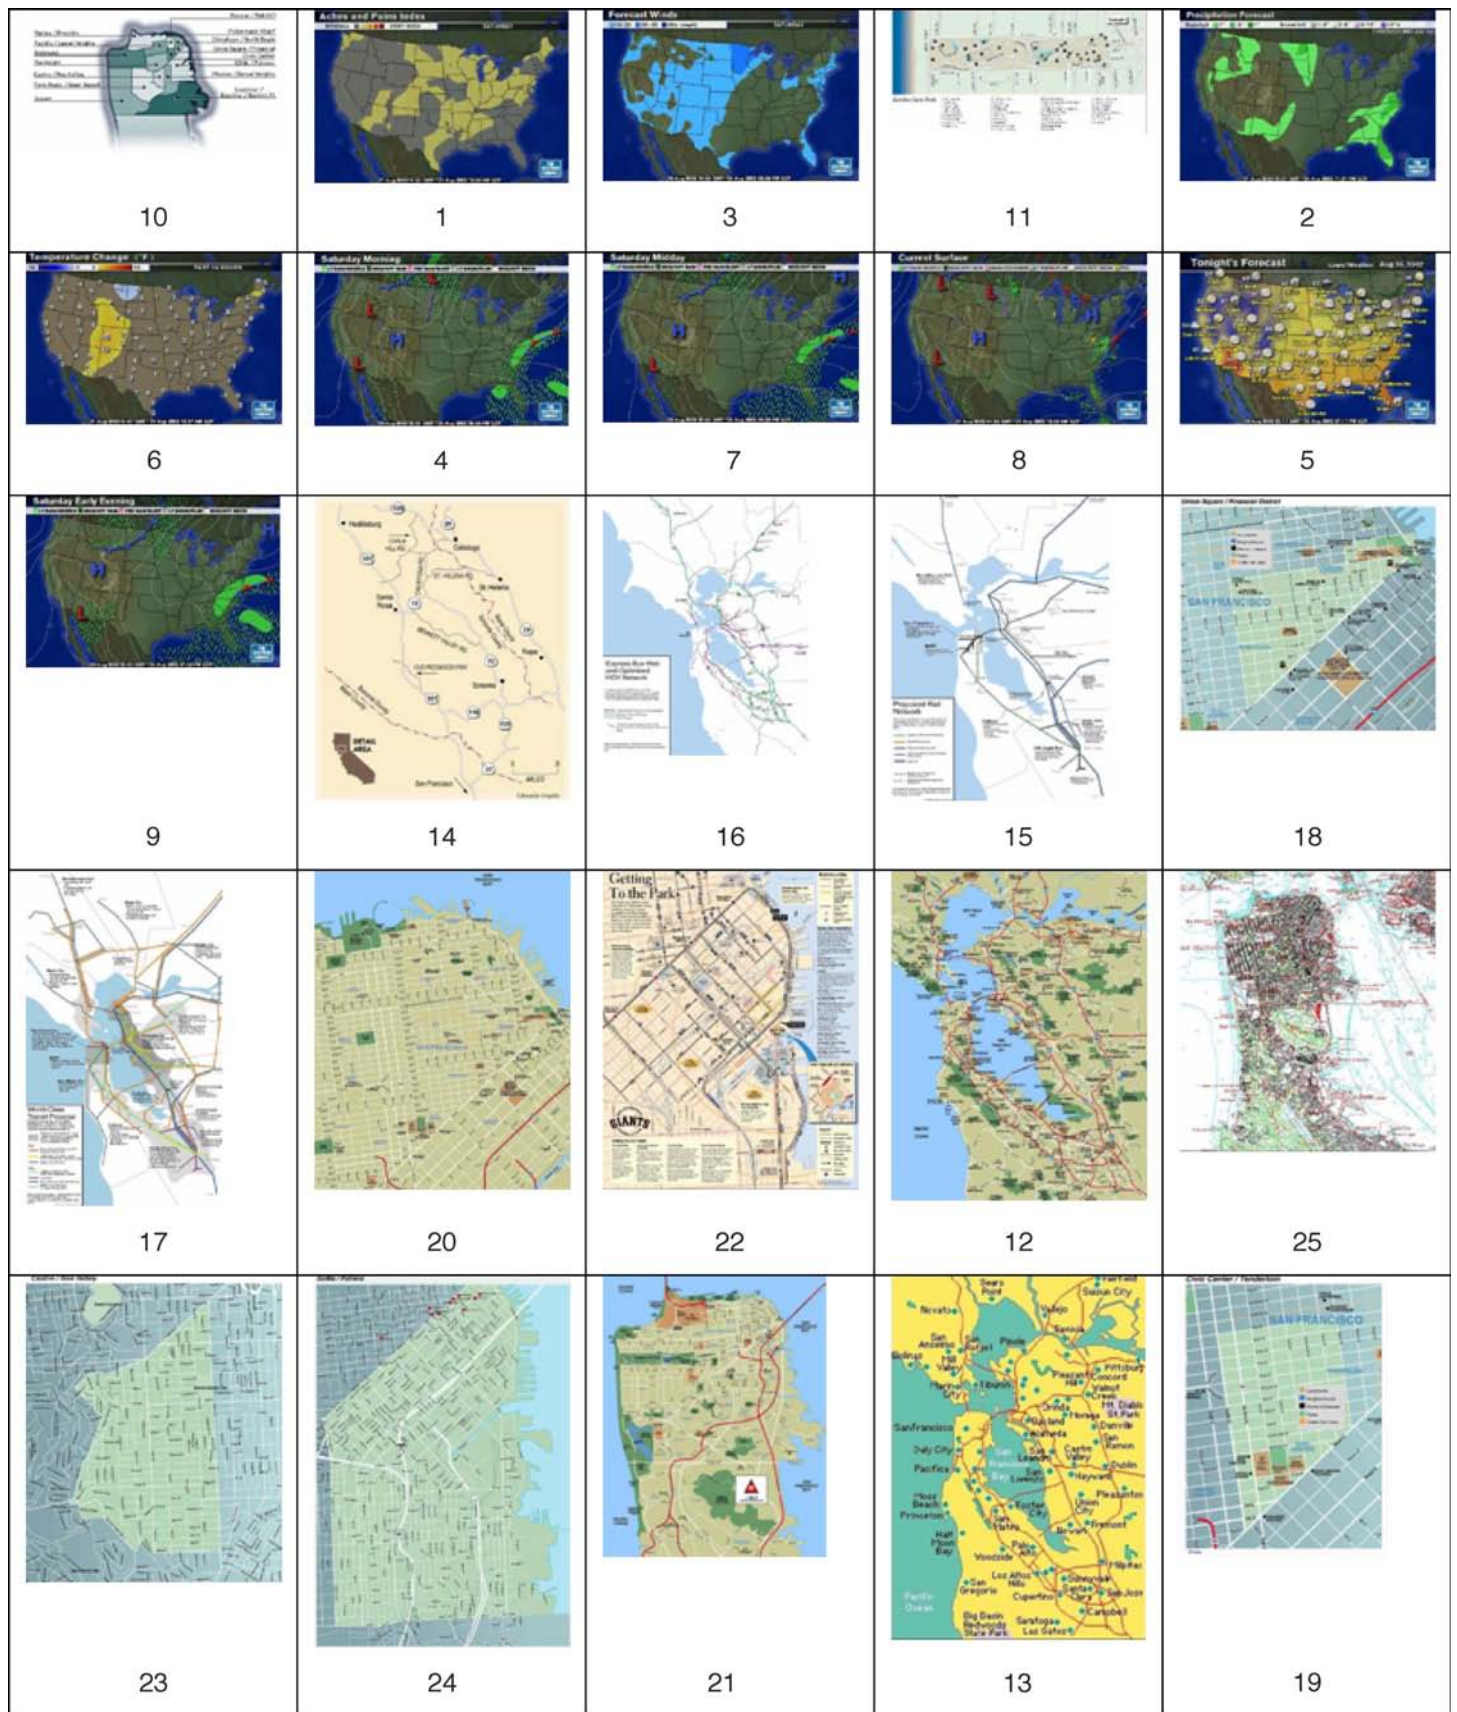


Figure 7. The maps used in the evaluation experiment, sorted from least cluttered (top left) to most cluttered (bottom right) as estimated by the crowding-based clutter model.

that mimics crowding. We found that such a model can also capture many findings reported in relation to clutter.

Comparison with other models

The model that we presented in this paper is not the first one to predict clutter. Rosenholtz et al. (2005) proposed a measure that relates clutter to local feature variance. Bravo and Farid (2008) found that clutter correlates with the number of regions in an image. In our view, the important question is not so much which of these measures is the “correct” one, but rather what is the common aspect that makes them successful in predicting clutter? It seems that all three clutter measures either explicitly (as in our present model) or implicitly (as in the other models) compute how much information is lost in peripheral vision: the higher the local feature variance (Rosenholtz) or the more “regions” an image consists of (Bravo & Farid), the greater the loss of information when information is compressed (as in peripheral vision, where sampling density is lower).

While the predictions of each model correlate strongly with perceived clutter, the correlation between the predictions of both models is much lower (see also Figure 5b and Figure 7). This suggests that the predictions of the models are partly based on a common factor determining clutter and partly on independent factors. It would be interesting to disentangle these effects. Varying feature variability could be a first manipulation, as this is where both models appear to make deviating predictions. However, an issue that immediately arises is that even if local variance is high throughout an image, there still may be higher order structure. Configural effects have been found for crowding (Livne & Sagi, 2007), but how they affect perceived clutter, search performance, and feature congestion is not known. Disentangling the effects of feature variance and crowding on clutter will thus require careful experiments that should also take configural effects into account.

Practical implications

The information visualization field is currently lacking a clear underlying theory (Purchase, Andrienko, Jankun-Kelly, & Ward, 2008); we believe that theoretical understanding of clutter should be part of such a theory.

Having established a link between clutter and crowding, a number of interesting consequences follow for the field of information visualization. Most importantly, this link suggests that the subjective concept of clutter has roughly the same properties as the much better understood crowding effect. In other words, precise predictions can be made about how clutter depends on (and can be controlled by) manipulation of object spacing and object similarity, among other things.

Another interesting question related to visualization is how clutter and crowding relate to texture perception. There is some evidence that crowding blocks access to local feature estimates, while access to global statistics is preserved. Based on this, some authors have proposed that crowding facilitates texture perception (Balas, Nakano, & Rosenholtz, submitted for publication; Parkes et al., 2001; Pelli & Tillman, 2008). If this is true, we should expect that the use of texture in visualizations can significantly reduce clutter and, therefore, improve their effectiveness. Although textures are already used in some visualization techniques (e.g., Healey & Enns, 1998; Kanatani & Chou, 1989), a perceptual theory explaining their effectiveness is lacking to date.

Directions for further research on clutter

The results of our study suggest that crowding is an important constituent and (thus) apt predictor of visual clutter. Although these results should not be interpreted as a definitive proof that clutter is “just a matter of crowding”, they do go a long way in suggesting these concepts are closely related. Therefore, further research in this direction is warranted.

Psychophysical experiments can be used to verify whether effect of things as object spacing and feature variability is the same for clutter and crowding. Furthermore, there are several ways in which our model could be improved. Theoretical knowledge about the mechanisms behind crowding is still quite limited. Hence, the model presented in this paper does probably not capture all details about the crowding effect. As better models of crowding become available, it should also be possible to make more accurate predictions of visual clutter.

One may argue that visual search is an even better candidate for modeling clutter. Clutter and search performance clearly correlate and several long-standing models exist for visual search (e.g., Treisman & Gelade, 1980; Wolfe, 2007), which might thus be used to predict clutter. However, while crowding has been shown to affect search (e.g., Vlaskamp & Hooge, 2005), to our knowledge, there are currently no models of visual search that take these effects into account. Although it will be interesting to study how well visual search models can predict clutter, we believe that crowding should be accounted for also in these models if they are to make accurate predictions of clutter.

Acknowledgments

We thank three *anonymous reviewers* for their helpful comments. This work was partly funded by the European

Commission under Grant No. 043157, project SynTex. It reflects only the authors' views.

Commercial relationships: none.

Corresponding author: R. van den Berg.

Email: r.van.den.berg@rug.nl.

Address: Laboratory for Experimental Ophthalmology, University Medical Center Groningen, University of Groningen, Antonius Deusinglaan 2, 9713 AW Groningen, Netherlands.

References

- Balas, B., Nakano, J., & Rosenholtz, R. (submitted for publication). A statistical model of peripheral vision explains visual crowding.
- Baldassi, S., Megna, N., & Burr, D. C. (2006). Visual clutter causes high-magnitude errors. *PLoS Biology*, *4*, e56. [PubMed] [Article]
- Beck, M., Lohrenz, M., Trafletton, J. G., & Gendron, M. (2008). The role of local and global clutter in visual search [Abstract]. *Journal of Vision*, *8*(6):1071, 1071a, <http://journalofvision.org/8/6/1071/>, doi:10.1167/8.6.1071.
- Bouma, H. (1970). Interaction effects in parafoveal letter recognition. *Nature*, *226*, 177–178. [PubMed]
- Bravo, M. J., & Farid, H. (2004). Search for a category target in clutter. *Perception*, *33*, 643–652. [PubMed]
- Bravo, M. J., & Farid, H. (2008). A scale invariant measure of clutter. *Journal of Vision*, *8*(1):23, 1–9, <http://journalofvision.org/8/1/23/>, doi:10.1167/8.1.23. [PubMed] [Article]
- Burt, P., & Adelson, E. H. (1983). The Laplacian pyramid as compact image code. *IEEE Transactions on Communication, COM-31*, 532–540.
- Grigorescu, C., Petkov, N., & Westenberg, M. A. (2004). Contour and boundary detection improved by surround suppression of texture edges. *Image and Vision Computing*, *22*, 609–622.
- Healey, G. H., & Enns, J. T. (1998). Building perceptual textures to visualize multidimensional datasets. In *IEEE Visualization. Proceedings of the Conference on Visualization '98* (pp. 111–118). Los Alamitos, CA, USA: IEEE Computer Society Press.
- Kanatani, K., & Chou, T. C. (1989). Shape from texture: General principle. *Artificial Intelligence*, *38*, 1–48.
- Korte, W. (1923). Über die Gestaltauffassung im indirekten Sehen. *Zeitschrift für Psychologie*, *93*, 17–82.
- Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. *Annals of Mathematical Statistics*, *22*, 79–86.
- Levi, D. M. (2008). Crowding—An essential bottleneck for object recognition: A mini-review. *Vision Research*, *48*, 635–654. [PubMed] [Article]
- Livne, T., & Sagi, D. (2007). Configuration influence on crowding. *Journal of Vision*, *7*(2):4, 1–12, <http://journalofvision.org/7/2/4/>, doi:10.1167/7.2.4. [PubMed] [Article]
- Oliva, A., Mack, M. L., Shrestha, M., & Peeper, A. (2004). Identifying the perceptual dimensions of visual complexity of scenes. In *Proceedings of the 26th Annual Meeting of the Cognitive Science Society Meeting*. Chicago.
- Parkes, L., Lund, J., Angelucci, A., Solomon, J. A., & Morgan, M. (2001). Compulsory averaging of crowded orientation signals in human vision. *Nature Neuroscience*, *4*, 739–744. [PubMed]
- Pelli, D. G., Palomares, M., & Majaj, N. J. (2004). Crowding is unlike ordinary masking: Distinguishing feature integration from detection. *Journal of Vision*, *4*(12):12, 1136–1169, <http://journalofvision.org/4/12/12/>, doi:10.1167/4.12.12. [PubMed] [Article]
- Pelli, D. G., & Tillman, K. A. (2008). The uncrowded window of object recognition. *Nature Neuroscience*, *11*, 1129–1135. [PubMed]
- Pelli, D. G., Tillman, K. A., Freeman, J., Su, M., Berger, T. D., & Majaj, N. J. (2007). Crowding and eccentricity determine reading rate. *Journal of Vision*, *7*(2):20, 1–36, <http://journalofvision.org/7/2/20/>, doi:10.1167/7.2.20. [PubMed] [Article]
- Purchase, C. P., Andrienko, N., Jankun-Kelly, T. J., & Ward, M. (2008). Theoretical foundations of information visualization. In A. Kerren, J. T. Stasko, J.-D. Fekete, & C. Chris North (Eds.), *Information visualization: Human-centered issues and perspectives. Vol. 4950 of LNCS State-of-the-Art Survey* (pp. 46–64). Heidelberg: Springer Berlin.
- Rosenholtz, R. (2001). Search asymmetries? What search asymmetries? *Perception & Psychophysics*, *63*, 476–489. [PubMed]
- Rosenholtz, R., Li, Y., Mansfield, J., & Jin, Z. (2005). Feature congestion, a measure of display clutter. In *SIGCHI* (pp. 761–770). NY, USA: ACM New York.
- Rosenholtz, R., Li, Y., & Nakano, L. (2007). Measuring visual clutter. *Journal of Vision*, *7*(2):17, 1–22, <http://journalofvision.org/7/2/17/>, doi:10.1167/7.2.17. [PubMed] [Article]
- Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the 1996 IEEE Symposium on Visual Languages* (p. 336). Washington, DC, USA: IEEE Computer Society.

- Stuart, J. A., & Burian, H. M. (1962). A study of separation difficulty. Its relationship to visual acuity in normal and amblyopic eyes. *American Journal of Ophthalmology*, *53*, 471–477. [[PubMed](#)]
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*, 97–136. [[PubMed](#)]
- van den Berg, R., Roerdink, J. B., & Cornelissen, F. W. (2007). On the generality of crowding: Visual crowding in size, saturation, and hue compared to orientation. *Journal of Vision*, *7*(2):14, 1–11, <http://journalofvision.org/7/2/14/>, doi:10.1167/7.2.14. [[PubMed](#)] [[Article](#)]
- Vlaskamp, B. N. S., & Hooge, I. T. C. (2005). Crowding degrades saccadic search performance [[Abstract](#)]. *Journal of Vision*, *5*(8):956, 956a, <http://journalofvision.org/5/8/956/>, doi:10.1167/5.8.956.
- Wilkinson, F., Wilson, H. R., & Ellemberg, D. (1997). Lateral interactions in peripherally viewed texture arrays. *Journal of the Optical Society of America A, Optics, Image Science, and Vision*, *14*, 2057–2068. [[PubMed](#)]
- Wolfe, J. M. (2007). Guided Search 4.0: Current progress with a model of visual search. In W. Gray (Ed.), *Integrated models of cognitive systems* (pp. 99–119). New York: Oxford.
- Woodruff, A., Landay, J., & Stonebraker, M. (1998). Constant information density in zoomable interfaces. In *Proceedings of Advanced Visual Interfaces* (pp. 57–65). NY, USA: ACM New York.
- Yang-Peláez, J., & Flowers, W. C. (2000). Information content measures of visual displays. In *Proceedings of the IEEE Symposium on Information Visualization* (p. 99). Washington, DC, USA: IEEE Computer Society.