No Troubles with Bubbles: a Reply to Murray and Gold

Frédéric Gosselin

Philippe G. Schyns*

Please send all correspondence either to:
Frédéric Gosselin
Département de psychologie
Université de Montréal
C.P. 6128, succursale Centre-ville
Montréal Québec Canada H3C 3J7
frederic.gosselin@umontreal.ca,
or to:
Philippe G. Schyns
Department of Psychology
University of Glasgow
56, Hillhead Street
Glasgow G12 8QB UK
philippe@psy.gla.ac.uk

* The authors are listed in alphabetical order. They are jointly responsible for the content of the article.

Abstract

Murray and Gold (in this issue of *Vision Research*) discuss two “shortcomings” of the *Bubbles* method (Gosselin & Schyns, 2001a). The first one is theoretical: *Bubbles* would not fully characterize the LAM (Linear Amplifier Model) observer, whereas reverse correlation would. The second “shortcoming” is practical: the apertures that partly reveal information in a typical *Bubbles* experiment would induce atypical strategies in human observers, whereas the additive Gaussian white noise used by Murray and Gold (and others) in conjunction with reverse correlation would not. Here, we show that these claims are unfounded.

In the first part of their article, Murray and Gold formalize Gosselin and Schyns' (2002a) analysis of the relationship between the three main types of information involved in visual tasks, the $R \quad A \approx P$ framework. We suggested that Potent information (P, rendered by *Bubbles*) is the result of an interaction (denoted by) of Represented information (R, measured by reverse correlation) with Available input information (A). Murray and Gold essentially confirmed that this analysis is valid in the context of the LAM observer and for a particular version of the *Bubbles* method.

We subscribe to the idea that without an analysis based on a precise observer model (such as LAM), it is difficult to know exactly what *Bubbles* and reverse correlation measure. Although we welcome Murray and Gold's formal effort, we do not believe that "Tout est pour le mieux dans le meilleur des mondes."[1], as Voltaire's *Candide* famously pointed out. Therefore, before describing the LAM formalization and evaluating its real implications for *Bubbles*, we will first discuss its general limitations.

Beyond Photometric Space: General Problems with the LAM Observer

Experiments with reverse correlation and *Bubbles* estimate the visual information used in a task by testing how information from an image generation space modulates response. This space should be chosen with great care, because its structure will constrain the information-use estimates. Factors such as the nature of the task to be resolved and the nature of the stimuli should guide the choice. To appreciate the range and diversity of possible search spaces, consider the problem of estimating the features of face recognition when the stimuli are a few static 2D images. If the faces are normalized for the positions of their main features (e.g., eyes, nose, mouth, and silhouette) then the 2D photometric image space can be fruitfully searched, but the information estimates will be restricted in scope – essentially revealing attention to features in the 2D image. Suppose now that stimuli are 3D-rendered, laser-scanned faces. A linear parametric space might be used (e.g., morphing), with loss of linearity in the 2D image projections (e.g., Leopold, O'Toole, Vetter & Blanz, 2001), but with greater scope for the information use estimates – essentially transformations of the 3D shape of faces.

Image generation spaces constitute a pressing research issue because the structure of visual information will tend to change with tasks (e.g., identification, gender, expression, age, and so forth) and object classes (e.g., static 2D pictures, static 3D-rendered objects, dynamic objects and so forth, see

---

[1] "All is for the better in the best of worlds."

Schyns, 1998). In our view (and that of others, e.g., Mangini & Biederman, 2003; Ollman & Kersten, 2003; Sadr & Sinha, 2003), recognition researchers should not restrict the scope of their research to the linear photometric space of the LAM observer, for a number of straightforward reasons discussed below.

Reverse correlation of the observer's responses onto Gaussian white noise of the type advocated by Murray and Gold does *not* provide a truthful representation of the information that the observer could use in the task. Instead, this version of reverse correlation projects the used information onto a linear space, as a weighted sum of Gaussian white noise pixel intensities. Such linear estimates have paid off particularly well in neurophysiological and psychophysical studies of lower-level vision, where intensity-based receptive fields are responsive to noise induced modulations of input contrast (Ringach & Shapley, 2003, for a review). The correlation between added input noise and response magnitude guarantees a good linear regression.

The success of models of low-level vision does not necessarily entail success of the same models at higher levels of vision. The neural systems of face and object categorization do not respond with the same magnitude to simple modulations of contrast like low-level ones do (Avidan, Harel, Hendler, Ben-Bashat, Zohary & Malach, 2001), but rather they respond to scale, higher-level features and shape information (Haxby, Hoffman & Gobbini, 2000; Gauthier, Tarr, Anderson, Skudlarski & Gore, 1999; Kanwisher, McDermott & Chun, 1997; Op de Beeck, Wagemans & Vogels, 2001; Sigala & Logothetis, 2002).

In sum, LAM is not sufficiently general in scope to impose any prescriptive standards on the conduct of research in visual categorization (see Gosselin & Schyns, in press). Thus, Murray and Gold's critique of *Bubbles* should be taken with a grain of salt. Not only is it limited because it assumes that humans are LAM observers – which they are not –, but also because it assumes a particular version of *Bubbles*. This last point is the object of the next section.

Within Photometric Space: Limitations in Scope of the Murray and Gold Analysis

Bubbles is an information sampling technique in which a "bubble" carves out an information sample from the generation space. With the exception of Schyns and Gosselin (2002), we have so far restricted the application of *Bubbles* to photometric spaces, not all of which have a 2D structure[2]. However, *Bubbles* is a generic method that can be applied to more complex parametric spaces than the photometric space (see Gosselin & Schyns, in press, for discussion), although the same also applies to reverse correlation (see Olman & Kersten, 2003). Murray and Gold's analysis applies only to the low-passed white noise version of *Bubbles* in photometric space. We will now show this analysis with limited scope is inconsequential for the practice of everyday research.

Still RAPing

The key formal development that Murray and Gold provide is reproduced below (their Equation 5):

$$B \approx b * b * (T \text{ o } (I_X - I_Y)).$$

For a LAM observer, T (or R, in our *RAP* framework) is an internal template, $(I_X - I_Y)$ (or A) is the ideal template, B (or P) the *Bubbles* information estimate, b the Gaussian bubble used to low-pass the white noise window, o a pointwise product, and * a convolution. This formulation implements $R \quad A \approx P$ in the context of a LAM observer: (1) the fuzzy operator becomes o, and (2) $(b * b)$, is introduced. The latter is a "double-blurring" term that quantifies the limit of the spatial resolution that can be achieved with a particular Gaussian b, given □. It is easy to show that $b(\square = s) * b(\square = s)$ reduces to $b'(\square = \sqrt{2}\, s)$, a larger single Gaussian bubble.

In *Bubbles*, □ is the only parameter that must be adjusted to sample information in a 2D image space. We have already shown that standard deviations of different sizes can be simultaneously used to search the image (e.g. at different spatial resolutions, Schyns et al., 2002, but also simultaneously to approximate the scale of diagnostic features in the 2D image). The choice of an appropriate □ is subject to a number of parameters (not all independent) ranging from the expected scale of visual information, the required smoothness of the solution, the number of parameters to estimate, and the required rate of convergence. *Bubbles* solutions can range from coarse (i.e., with a large Gaussian bubble, few parameters to estimate, and typically fast convergence), to fine (i.e. with many parameters to estimate and typically

[2] A list of generation spaces explored with *Bubbles*: the standard 2D image plane (Gosselin & Schyns, 2001a; Gibson, Gosselin, Wasserman & Schyns, 2002; Jentzsch, Gosselin, Schweinberger and Schyns, 2002; O'Donnell, Gosselin & Schyns, 2002), spatial frequency (Schyns & Gosselin, 2002), spatial frequency x 2D image (Bonnar, Gosselin & Schyns, 2002; Gosselin & Schyns, 2001a; Schyns, Bonnar & Gosselin, 2002), and 2D image plane x time (Vinette & Gosselin, 2002).

slow convergence). At the limit, sigma becomes a dot in discrete space, and $b * b$ vanishes from Murray and Gold's equation, which then reduces to $R \circ A \approx P$, and the spatial resolution of *Bubbles* is equivalent to that which is typical of white noise reverse correlation.

No Troubles in Theory: LAMe attack

The Murray and Gold formalization led them to conclude: "[…] if the LAM is a valid model of the system under study, then […] equation (5) shows that from a classification image [that is R] we can easily determine the result of any bubbles experiment. […] Thus in cases where the LAM is correct, the bubbles method is *superfluous* […]" (p. 9, italic added). Later on (on p. 19), they modulate this statement with: "[…] if one is interested only in what stimulus locations help an observer give a correct response [that is *P*], then the bubbles method is perfectly adequate [….]".

We have shown that Murray and Gold's equation (5) reduces to $R \circ A \approx P$ when □ is so small that the Gaussian bubble reduces to a dot. Whenever $R \circ A \approx P$ and $R \approx P / A$ (a pointwise division) are defined, a complete *RAP* characterization of the observer can be obtained either from *A* and *R* (estimated from reverse correlation) or from *A* and *P* (estimated from *Bubbles*). The theoretical superiority of reverse correlation over *Bubbles* rests solely on the domain of definition: $R \circ A$ is defined over all real numbers, whereas P / A is undefined whenever $A(x, y) = 0$, for any x and y. To make it absolutely clear, reverse correlation is superior to *Bubbles* whenever there is no information available at a given image location to resolve a task – whenever $A(x, y) = 0$.

In our own research, we used "superstitious" to refer to situations in which $A(x, y) = 0$ and $R(x, y) \neq 0$. For example, we induced three observers to "superstitiously" see an 'S' letter ($R(x, y) =$ 'S') in pure bit noise ($A(x, y) = 0$) by artificially restricting (via instructions) the number of candidate representations to be matched against the input noise. We then applied reverse correlation to depict the observer's share, i.e. $R(x, y)$. We are thus well aware that reverse correlation can – and that *Bubbles* cannot – be applied in such "superstitious" situations. We made this point explicit a number of times (Gosselin & Schyns, 2002a, 2002b, in press; Gosselin, Bacon & Mamassian, submitted).

Unless artificially created[3], however, $A(x, y) = 0$ is the exception. In most of the work by Murray

[3] Murray and Gold (p. 17) "[…] chose the fat-thin task […] in order to show how the bubbles method would affect strategies in a task that had actually been discussed in the literature, rather than a task that was

and Gold and colleagues, two-image experiments are designed. To derive *A* in such conditions, we computed the differences between all possible pairs of the 32 face images (grayscale 8-bits) used in Gosselin and Schyns (2001a, Experiment 1). $A(x, y) = 0$ occurs with a probability of .022, always within the forehead region[4]. In most recognition experiments (and real life), observers compare pairs of categories of images, not pairs of individual images. In LAM, the information available between pairs of categories is the sum of the images in one category minus the sum of the images in the other category. To set the stage for the experiment to come, Figure 1 presents the information available in the gender (sum(10 male faces) – sum(10 female faces)) and expressive or not categorizations (sum(10 smiling faces) – sum(10 neutral faces)) computed from 20 8-bit grayscale images of Schyns et al. (2002). The likelihood of $A(x, y) = 0$ is .00034 in GENDER and .000057 in EXNEX. As human observers tend to maximize within-category similarity and between-category dissimilarity (Rosch, 1978; Gosselin & Schyns, 2001b), the closer experimental stimuli get to real-world conditions of stimulation (in the example, hairstyle and lighting conditions were normalized), the smaller these probabilities will be.

---

designed to maximize the disruptive effect of showing only small fragments of a signal." Unfortunately, in this fat-thin task, *A* has many more zero information pixels than non-zero ones – non-zero information pixels of *A* are concentrated in two sectors of 3.5° at the palates of each "pacman". As this is a most unfortunate choice of experiment to compare *Bubbles* with reverse correlation we will not discuss Murray & Gold's experiment any further.

[4] We did the same for the 10 face images (grayscale 8-bits) of Gold, Bennett and Sekuler (1999a, 1999b). A(x, y) = 0 never occurs.

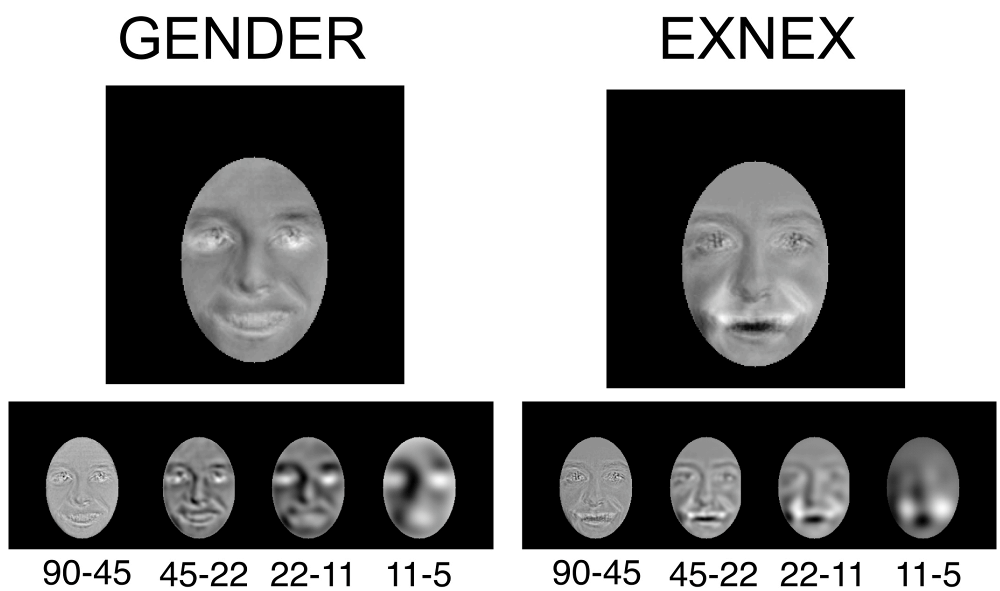

Figure 1. This figure depicts the linear information available (i.e., A) to resolve the GENDER and EXNEX categorization tasks. It also shows the breakdown of this information at four of the scales sampled in Schyns, Bonnar & Gosselin (2002) (11.25-5.65, 22.5-11.25, 45-22.5, 90-45 cycles per face, from fine to coarse).

If $A(x, y) = 0$, and an estimate of $R(x, y)$ is required from $P(x, y) / A(x, y)$, then 0 can be assigned to $R(x, y)$. If there is no information available for this pixel, it is reasonable to assume that the observer's representation will not comprise information for this pixel either.

To summarize, the fact that $P(x, y) / A(x, y)$ is not defined when $A(x, y) = 0$ gives little support to the claim that the *Bubbles* method is "superfluous" (Murray & Gold, p. 9 and p. 19). We have argued that $A(x, y) = 0$ is an outlier event in the real-world, that it only occurs in artificial laboratory situations. For most practical purposes, then *Bubbles* and reverse correlation are complementary techniques (see Gosselin & Schyns, 2002a).

No Troubles in Practice: A Fair Comparison of Bubbles and Reverse Correlation

To compare the outcomes of reverse correlation and *Bubbles*, we applied reverse correlation in two categorization tasks (GENDER, and expressive or not, EXNEX) in experimental conditions identical to the *Bubbles* experiment of Schyns et al. (2002). Fifteen University of Glasgow paid observers were assigned to each categorization task (a total of 30 observers). A trial consisted in the presentation of a randomly chosen stimulus to which Gaussian white noise was added. The sigma of the noise distribution was adjusted to maintain observers' performance at 75% correct. The stimuli were presented on a calibrated high-resolution Sony monitor, with a refresh rate of 85 Hz. The *Bubbles* version of the experiment differed only in how the facial information was sampled. For the reverse correlation experiment, results were derived from the pooled data of 15 observers each resolving 1000 trials, separately for the GENDER and EXNEX conditions. For the *Bubbles* experiment, we reanalyzed the data of Schyns et al. (2002), pooling this time 1000 trials per observer, separately for GENDER and EXNEX (the analyses reported in Schyns et al., 2002 only concerned the last 500 trials of each observer).

The raw *Bubbles* estimates, at different scales, in the GENDER and EXNEX tasks are shown in Figure 2 (see Schyns et al., 2002, for details). The raw classification images are presented in the leftmost pictures of Figure 2. The steps to compute these images are as follows: We summed together the Gaussian noise fields according to the categorization response they elicited (e.g. male_sum = hits(male) + false_alarms(male); female_sum = hits(female) + false_alarms(female)). We then subtract the summed image corresponding to each categorization (e.g., male_sum – female_sum) to derive the classification images.

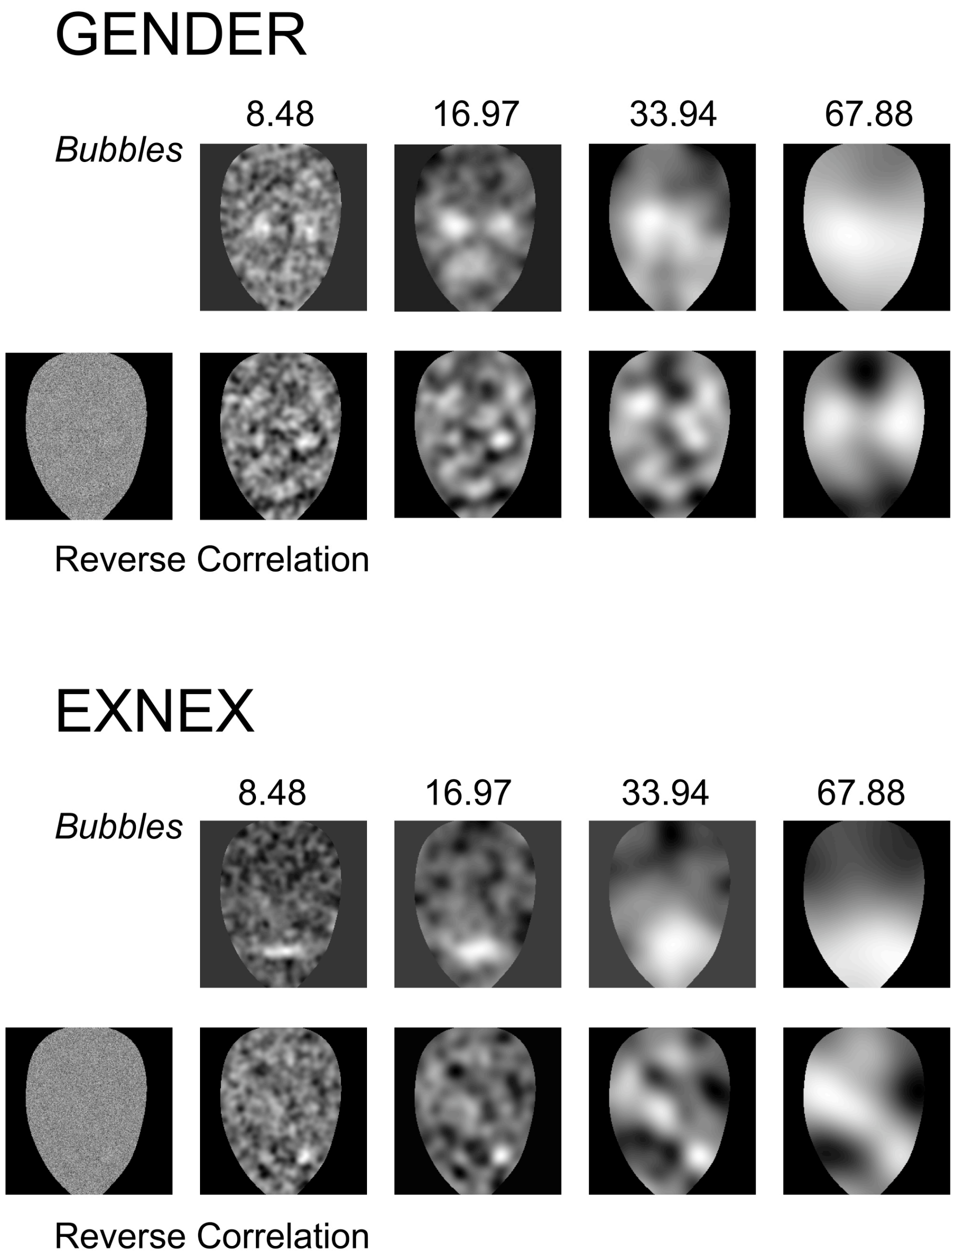

Figure 2. This figure depicts the *Bubbles* and reverse correlation visual information estimates (i.e., *P* and *R*, respectively), at different scales, in the GENDER and EXNEX tasks. (The raw reverse correlation classification image is shown in the leftmost picture). To compare the reverse correlation solution to the *Bubbles* solution we convolved the former with Gaussians of standard deviations 8.48, 16.97, 33.94 and 67.88 pixels, from fine to coarse (this takes into account the “double-blurring” discussed by Murray & Gold). This analysis reveals that the *Bubbles* and reverse correlation solutions represent similar information across scales.

To compare the reverse correlation solution to the *Bubbles* solution we convolved the former with Gaussians of standard deviations 8.48, 16.97, 33.94 and 67.88 pixels, from fine to coarse (this takes into account the "double-blurring" discussed by Murray & Gold). The outcome of the *Bubbles* and reverse correlation solutions is similar at the two finest scales (see Figure 2). It reveals the use of the eyes in GENDER, and the use of the mouth in EXNEX. While the *Bubbles* outcome remains consistent at coarser scales, the corresponding reverse correlation solutions become noisier and more difficult to interpret.

In sum, it does not appear that "…the *Bubbles* method drastically changes human observers' strategies," (Murray & Gold, p. 4) at least compared with Gaussian noise reverse correlation. Furthermore, there is no evidence that the value of the information extracted from the *Bubbles* experiment is reduced compared to that extracted from the reverse correlation experiment.

No Troubles Whatsoever in Practice

Even though we did not find empirical evidence for this, we are still left with Murray and Gold's firm belief that the type of windowing used in *Bubbles* "drastically" disrupts the observers' strategy, whereas additive Gaussian white noise does not (Murray & Gold, Abstract, p. 4, p. 12, p. 17, p. 19, and p. 20)[5]. They develop three arguments to support their belief. None of them survives closer scrutiny.

Murray and Gold wrote (p. 16): "First, it is intuitively clear why windowing stimuli through bubbles might change observers' strategies: when only small parts of a stimulus are shown on any given trial, observers may be forced to use stimulus features that they would not use if the whole stimulus was presented."

To the extent that added Gaussian white noise will mask certain diagnostic regions more than others, observers that would have used the former will have to resort to using the latter. The probability that a diagnostic region will be entirely destroyed by additive noise might appear small in comparison to the probability that it is not revealed through Gaussian bubbles. However, this will happen only if the sigma of the Gaussian bubbles is too large for the task at hand (see discussion in the *Still RAPping* section above).

[5] As demonstrated by Murray & Gold, a LAM observer uses the same strategy when presented with bubbled stimuli than when presented with stimuli plus Gaussian white noise. Therefore any difference between the strategies induced by windowing with Gaussian bubbles and additive noise on human observers refutes LAM as an adequate model. The success of Murray & Gold's practical argument entails the demise of their theoretical argument, and vice-versa.

This does not demonstrate the superiority of additive noise over windowing.

"[Second,] a great deal of psychophysical and physiological evidence shows that even under noiseless viewing conditions, observers' performance in threshold tasks is limited by internal noise, so by adding external noise we are probably not presenting observers with a task that is qualitatively different from a noiseless threshold task [….]" (Murray & Gold, p. 16-17)

This could be the case, but the critical part of this argument (i.e., that observers are not performing a task qualitatively different from an everyday visual task) applies to windowing as well. Aren't the eyes a window on the world? And aren't objects almost never seen in their entirety? As Murray, Sekuler and Bennett (2001) put it: "One of the challenges to object recognition is the fact that sensory information reaching the eyes is often incomplete: Objects occlude parts of neighboring objects and parts of themselves. Even though we constantly perceive partly occluded objects, we rarely notice that the visual information we receive is incomplete." (p. 1) Thus, windowing is certainly a type of noise we are accustomed to dealing with in the real world.

"Third, and most convincingly, observers' contrast energy thresholds have been found to be an approximately linear function of external noise power in practically every task in which this relationship has been tested, including discrimination of fat vs. thin Kanizsa squares, and this is strong evidence that observers use the same strategy at all levels of external noise, from negligible levels to high levels of noise." (Murray & Gold, p. 17)

An analoguous approximate linear relationship was found between contrast energy threshold and the area revealed by Gaussian bubbles. Two experienced psychophysical observers from the Université de Montréal with normal, or corrected to normal vision, identified either the gender (EM), or the expression (IF) of the 32 faces used in Gosselin and Schyns (2001, Experiment 1). Observers had previously participated in several other *Bubbles* experiments with the same face set and with the same tasks; their performance had thus stabilized at the time of this investigation. The experiment ran on a Macintosh G4 computer using a program written with the Psychophysics Toolbox for Matlab (Brainard, 1997; Pelli, 1997). The stimuli were presented on a calibrated high-resolution Sony monitor, with a refresh rate of 85 Hz. Contrast energy thresholds for a 75% correct response rate were measured at four levels of windowing (400, 500, 600, and 700 bubbles with a standard deviation equal to 3 pixels, or .067 deg of visual angle, for stimuli spanning 256 x 256 pixels, or 5.72 x 5.72 deg of visual angle) using the method of constant stimuli (we employed five levels of contrast energy for each level of windowing – a total of 96 * 5 levels of contrast energy * 4 levels of windowing = 1,920 trials per observer; R^2 for the bestfitted cumulative Gaussians ranged from .692 to .947, with an average of .803). We found approximate linear relationships between contrast energy thresholds and level of windowing (IF: R^2 = .947; EM: R^2 = 0.738).

In sum, the Murray and Gold claim that use of *Bubbles* "drastically" modifies observers' strategies and so reduces the value of the technique compared with additive Gaussian white noise reverse correlation is unfounded. Our face recognition comparison did not lend support to this claim. In addition, *none* of the arguments put forth by Murray and Gold claiming that Gaussian white noise is preferable to Gaussian bubbles stands up to closer scrutiny.

Conlcusions

Murray and Gold caimed that there are some "shortcomings" with *Bubbles*. We have addressed their criticisms and shown: (1) that their formal analysis is restricted in scope; (2) that the argument that *Bubbles* would not fully characterize the LAM observer is inconsequential; (3) that in a fair comparison, *Bubbles* and reverse correlation reveal a simlar use of information in human observers; and (4) that none of the arguments put forth for prefering additive Gaussian white noise over windowing with Gaussian bubbles survives closer scrutiny. Thus, there are no troubles with *Bubbles*.

References

Avidan, G., Harel, M., Ben-Bashat, D., Zohary, E. & Malach, R. (2001). Contrast sensitivity in human visual areas and its relationship to object recognition. *Journal of Neurophysiology, 87*, 3112-3116.

Bonnar, L., Gosselin, F. and Schyns, P. G. (2002). Understanding Dali's Slave Market with the Disappearing Bust of Voltaire: A case study in the scale information driving perception. *Perception, 31*, p. 683-691.

Gauthier I., Tarr M. J., Anderson A. W., Skudlarski P., & Gore J. C. (1999). Activation of the middle fusiform 'face area' increases with expertise in recognizing novel objects. *Nature Neuroscience, 2*, 568-573.

Gibson, B. M., Gosselin, F., Wasserman, E. A., & Schyns, P. G. (2002). Pigeons and people categorize human faces using similar features. Submitted.

Gold, J. Bennett, P.J. & Sekuler, A. B. (1999a). Signal but not noise changes with perceptual learning. *Nature*, 402, 176-178.

Gold J, Bennett P.J., & Sekuler AB (1999b). Identification of band-pass filtered faces and letters by human and ideal observers. *Vision Research*, 39(21), 3537-3560.

Gosselin, F. & Schyns, P. G. (2001a). Bubbles: A technique to reveal the use of information in recognition. *Vision Research, 41*, p. 2261-2271.

Gosselin, F. & Schyns, P. G. (2001b). Why do we SLIP to the basic-level? Computational constraints and their implementation. *Psychological Review, 108*, 735-758.

Gosselin, F. & Schyns, P. G. (2002a). RAP: a new framework for visual categorization. *Trends in Cognitive Science, 6*, p. 70-77.

Gosselin, F. & Schyns, P. G. (2002b). You are about to see picturial representations! [Commentary on Z. Pylyshyn's target article] *Behavioral & Brain Research, 25*, 191-192. 2002.

Gosselin, F., & Schyns, P. G. (in press). A User's Guide to Bubbles. To appear in L. Gershkoff-Stowe and D. Rakison, Building categories in developmental time. Hilldsale, NJ: Erlbaum..

Gosselin, F. & Schyns, P. G. (in press). Superstitious perceptions reveal properties of memory representations. *Psychological Science*.

Gosselin, F., Bacon, B. A., & Mamassian, P. (in press). Internal surface representations approximated by multiple regression. *Vision Research.*

Haxby J. V., Hoffman E. A., & Gobbini M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences, 4*, 223-233.

Humphreys, K., Gosselin, F., Schyns, P. G., Kaufman, J., & Johnson, M. H. Do 7-month-old infants use the same information as adults to process facial identity? *International Conference on Infant Studies*, Toronto, Canada, 18-21 avril, 2002.

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

Mangini, M. & Biederman, I. (2003). Making the ineffable explicit: Estimating the information employed for face classifications.

Murray, R. F., Sekuler, A. B., & Bennet, P. J. (2001). Time course of amodal completion revealed by a shape discrimination task. *Psychonomic Bulletin & Review, 8*, 713-720.

Nielsen, K. J., Rainer, G., Brucklacher, V., & Logothetis, N. K. (2002). Studying the representation of natural images with the use of behavioural reverse correlation, *ECVP.*

O'Donnell, C., Schyns, P. G., & Gosselin, F. (2002). The acquisition of facial expertise and how that mediates the information utilized to recognize the face. *Vision Sciences.*

Ollman, C. A. & Kersten, D. (2003). Classification objects, ideal observers and generative models. Submitted for publication.

Op de Beeck, H., Wagemans, J., & Vogels, R. (2001). Inferotemporal neurons represent low-dimensional configurations of parametrized shapes. *Nature Neuroscience,* **4**, 1244-1252 (2001).

Oshawa, I., DeAngelis, G. C., & Freeman, R. D. (1990). Stereoscopic depth discrimination in the visual cortex: neurons ideally suited as disparity detectors. *Science*, *249* (4972), 1037-1041.

Ringach, D. & Shapley, R. (2003). Reverse correlation in neurophysiology. Submitted.

Rosch, E. (1978). Principles of categorisation. In E. Rosch & B. B. Lloyd (Eds.), *Semantic factors in cognition* (pp. 137-168). Hillsdale, NJ: Erlbaum.

Sadr, J. & Sinha, P. (2003). Object recognition and random structure evolution. Submitted for publication.

Schyns, P. G. (1998). Diagnostic recognition: Task constraints, object information and their interactions. *Cognition*, 67, 147-179.

Schyns, P. G., Bonnar, L., & Gosselin, F. (2002). Show me the features! Understanding recognition from the use of visual information. *Psychological Science, 13*, 402-409.

Schyns, P. G. & Gosselin, F. (2002). Use of frequency informaion at the basic and the subordinate level in animal categorization. *Vision Sciences.*

Schyns, P. G., Jentzsch, I., Johnson, M., Schweinberger, S. R., & Gosselin, F. (in press). A principled method for determining the functionality of brain responses. *NeuroReport*.

Sigala, N. & Logothetis, N. (2002). Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature*, **415**, 318-320.

Tinbergen, N. & Perdeck, A.C. (1950). On the stimulus situation releasing the begging response in the newly hatched herring gull chick (Larus argentatus Pont.). *Behaviour, 3*, 1-39.

Tanaka, J.W. & Farah, M.J. (1993). Parts and wholes in face recognition. Quarterly Journal of Experimental Psychology, 46A, 225-245.

Vinette, C. & Gosselin, F. (2002). Spatio-temporal use of information in face recognition. *Vision Sciences.*

Vuilleumier P., Henson R. N., Driver J., & Dolan R. J. (2002). Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nature Neuroscience, 5*, 491-499.

Aknowledgments

This research was supported by a NSERC (R0010085) and a NATEQ (R0010287) grant awarded to Frédéric Gosselin; and by an ESRC grant (R000237901) awarded to Philippe Schyns. We thank Marie Smith for running the first experiment, and Éric McCabe and Isabelle Fortin for their help with the second experiment. Thanks to Marie Smith and Asifa Majid for their comments on the article.