A User's Guide to *Bubbles*

Frédéric Gosselin

Philippe G. Schyns

Correspondence concerning this article should be addressed to Frédéric Gosselin, Département de psychologie, Université de Montréal, C.P. 6128, Succursale centre-ville, Montréal, Québec, Canada, H3C 3J7; or to Philippe Schyns, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow, Scotland, United Kingdom, G12 8QB. Electronic mail may be sent via the Internet to frederic.gosselin@umontreal.ca or to philippe@psy.gla.ac.uk.

Abstract

This article provides a user's guide to *Bubbles*, a technique that reveals the information that drives a measurable response. We illustrate the technique with a complete example: the Face Inversion Effect and discuss the six basic decisions that must be made to set up a *Bubbles* experiment (i.e., the stimulus set, the generation space, the "bubbles", the task, the group of observers, and the response). We describe methods to analyze the data and provide practical advice for the researcher intending to use the technique.

A User's Guide to *Bubbles*

The herring gull chick begs for food by pecking at its mother's beak. In a seminal experiment, Nobel-prize-winning ethologist Nikko Tinbergen and co-worker (Tinbergen & Perdeck, 1950) sought to discover the stimulus that maximized this response. This enterprise led to the remarkable discovery of the *super-stimulus*: An artificial stimulus that evokes a stronger response than the original, natural stimulus. For example, a white stick with three red annuli moving up and down produces a stronger pecking response than the head of the herring gull mother's.

At an abstract level, the search for the super-stimulus can be framed as a generic search problem. Given a measurable dependent variable (e.g. the pecking rate response), the problem is to find the specific parameters of the independent variable(s) (e.g. the characteristics of the mother's head) that optimize the dependent variable. Obviously, this approach is not limited to ethology. An approach similar in spirit is that of Nobel-prize-winners Hubel and Wiesel who searched for the stimulus that optimizes the response of cells in the primary visual cortex (see Hubel, 1988, for a review). Much to their surprise, they discovered that small spots of light, which are so effective in the retina and Lateral Geniculate Nucleus (LGN) were much less effective in visual cortex. Instead, simple cells in primary visual cortex responded optimally to inputs with linear properties, such as a line with a specific width and orientation in the plane. At the next level of cortical integration, optimal inputs become more complicated. For example, complex cells tend to respond better to a stimulus with a critical orientation, but also a characteristic speed and direction, adding to the width and orientation search space a third dimension. Further up the integration ladder, cells in temporal cortex respond to complex object properties (Kobatabe & Tanaka, 1994) such as orientation in depth (Perrett, Mistlin & Chitty, 1987), object similarities (Vogels, 1999), and the information responsible for visual object categorization (Sigala & Logothetis, 2002; Vogels, 1999).

However, even though IT cells are just “a few synaptic connections away” from primary visual cortex, their optimal stimuli are hidden in a much more complex search space: the physical world. With its many faces, objects and scenes, this space does not comprise just the few degrees of freedom required to represent the little spots of light positioned within the visual field, or the moving orientated bars. Instead, IT cells respond to structured information that varies in 2D retinal position, 2D rotation in the image plane, 3D rotation in depth, illumination, articulation and so forth. Amongst these multiple degrees of freedom, different subspaces of parameters represent the effective stimuli of IT cells. The challenge is to understand what these subspaces are.

And it is still one of the greatest methodological challenges in Cognitive Science: when dealing with complex visual stimuli, how can a brain event (an ERP response, or an fMRI measurement) or a human behavior (e.g. a categorization response) be attributed to a specific object category (e.g. a beach scene), a specific object (e.g. a deck chair), a specific feature (e.g., the texture of the deck chair) or a specific function (e.g. a beach, deck chair, or texture detector)? In the absence of a principled method, the specificity of the response (e.g., to the beach) is determined by contrast with responses from carefully chosen contrast categories (e.g., roads, cities, mountains, fields and so forth), and informal hypotheses tested. Unfortunately, a dense correlative structure exists in the low-level visual properties of category members (e.g., luminance energy, main directions of orientation, spatial frequency composition and so forth), only a small subset of which can be controlled with a finite number of carefully chosen contrast categories. Consequently, the specificity of the brain or behavioral response might be due to incidental input statistics and not to the category *per se* (Schyns, Jentzsch, Johnson, Schweinberger & Gosselin, 2002).

In this article, we present *Bubbles* (Gosselin & Schyns, 2001), a method designed to solve the problem of finding the effective stimulus in the complex search spaces that are characteristic of visual categorization. From the outset, it is important to stress that the method can be scaled down and applied to simpler search spaces. However, originality of *Bubbles* is that it can handle search spaces that have so far proven to be elusive (e.g., the information responsible for face recognition, Gosselin & Schyns 2001, Schyns, Bonnar & Gosselin, 2002; scene recognition, Nielsen, Rainer, Brucklacher & Logothetis, 2002 ; or the perception of complex figures, Bonnar, Gosselin & Schyns, 2002). The article is organized as a user's guide. First, we introduce a typical research problem never before addressed with *Bubbles*: the Face Inversion Effect. We then discuss the six main decisions that must be made to set up a *Bubbles* experiment, discussing critical issues with examples from our own research.

The problem: The Face Inversion Effect

In a seminal article, Yin (1969) reported that the recognition of face pictures was disproportionately affected by a 180 deg rotation in the image plane from the normal, upright viewing condition. This phenomenon is now commonly called the Face Inversion Effect (FIE). Since then, the FIE has been replicated in multiple experimental situations (e.g. Carey, Diamond & Woods, 1980; Philips & Rawles, 1979; Scapinello & Yarmey, 1970 ; Carey & Diamond, 1977; Diamond & Carey, 1986; Freire et al., 2000;

Leder & Bruce, 2000; Scapinello & Yarmey, 1970; Tanaka & Farah, 1993; Valentine & Bruce, 1986; Yarmey, 1971).

There is now agreement amongst most face recognition researchers that the FIE does not arise from long-term memory interferences, but instead from a greater difficulty to perceptually encode inverted face information (e.g. Farah et al., 1998; Moscovitch, Berhmann & Winocur,1997; Phelps & Roberts, 1994; Searcy & Bartlett, 1996, Freire et al., 2000). Therefore, recent studies have examined more closely the encoding differences that occur when experiencing an upright or an inverted face. However, the specification of these differences has so far remained largely unknown (Rossion & Gauthier, in press).

To address the FIE with *Bubbles*, we need to make six basic decisions (1) what is the stimulus set, (2) in which space will stimuli be generated, (3) what is the "bubble," (4) what is the observer's task, (5) what are the observer's possible responses (6) is the analysis per observer, or per group of observers. In resolving all of these, we will set up a search space and vary the parameters of the independent variables (upright and inverted face information) that determine the measurable dependent variable (the observer's response). The *Bubbles* solution should specify the difference between the information subspaces driving the processing of upright and inverted faces.

(1) Stimulus set. In a Bubbles experiment, the stimulus set is crucial because it critically bounds what will be tested. Here, we used a total of 10 greyscale faces (5 males, 5 females), each one of which displaying 3 different expressions (neutral, angry and happy). Hairstyle was normalized, and so were global orientation, and the location of the light source. Stimuli could be upright or inverted, but when inverted, we flipped the image so as to keep the light source to the right of the face.

Generally speaking, the larger the stimuli set, the better the *Bubbles* solution should be. A large stimulus set will tend to prevent observers from adopting strategies atypical of natural processing. In the FIE example, the stimulus set restricts the search space for differences in upright and inverted face encodings to a few males and females with a limited set of expressions, in highly restricted conditions of presentation (only one light source, two poses, and static images). Although this also applies in most face recognition experiments, it is important to point out that the *Bubbles* solution will be tied to these limitations. In our research we have already used faces in other experiments with human participants (Gosselin & Schyns, 2001; Schyns, Gosselin & Bonnar, 2002; Schyns, Jentzsch, Schweinberger, Johnson & Gosselin, 2002), but also in animals experiments (Gibson, Wasserman, Gosselin & Schyns, 2002). Other stimuli used ranged from 3D models of Gouraud shaded animals (Schyns & Gosselin, 2002) to a painting of Dali (Bonnar, Gosselin & Schyns, 2002). Other researchers have also applied *Bubbles* to natural scenes (Nielsen, Rainer, Brucklacher, & Logothetis, 2002). Although these applications only involved visual stimuli, the technique should straightforwardly generalize to auditory and tactile stimulus sets, or to cross-modal combinations of these.

(2) Stimulus Generation Space. The choice of a proper stimulus generation space is one of the most important decisions when setting up a *Bubbles* experiment. Remember that we are searching for the parameters of the independent variables (upright and inverted face information) that determine the FIE. Each independent variable considered constitutes one independent dimension whose parametric values will be searched. To illustrate, our face stimuli are 2D pictures. The axes of the 2D plane could be searched to find the (x,y) coordinates of face information that determine upright, vs. inverted

performance. The stimulus generation space would then be two-dimensional, and the solution would be a subset of the plane.

However, there is evidence that early vision does analyze the input at multiple spatial scales (or spatial frequencies, see de Valois & de Valois, 1990, for a review), and that mechanisms of face recognition rely on this input (see Morrison & Schyns, 2001, for a review). Thus, a better space to search for the determinants of FIE could include a third dimension of spatial scales. Specifically, we segmented the third dimension into 5 independent bands of fine to coarse spatial frequencies of one octave each—with cutoffs at 90, 45, 22.5, 11.25, 5.62, and 2.81 cycles per face. The solution subspace becomes an interaction between the 2 dimensions of face feature location, and the third dimension of spatial scale.

In setting up this search space, we are making a number of assumptions that are worth pointing out. We are assuming that the face pictures are normalized for the position of their main features (i.e., the *x,y* locations of the eyes, nose, mouth, chin, cheeks and forehead are roughly similar across face pictures). This is necessary because the selected search space is not invariant to translation in the image plane. Similarly, we are assuming that the faces in the pictures will have the same size, because the search space is not invariant to scale changes. Note that these constraints on the search space are not constraints on the technique itself. It is possible to set up translation invariant search spaces (see Schyns & Gosselin, 2002, for a Fourier implementation), and it is also possible to set up scale invariant search spaces. However, the experimental question (the nature of face features that determine FIE) suggested a search space where the location of face features would be known.

In our research, we have used a variety of stimulus generation space, ranging from the 2D image plane (Gosselin & Schyns, 2001; Gibson, Gosselin, Wasserman & Schyns, 2002; Schyns, Jentzsch, Schweinberger, Johnson and Gosselin, 2002; O'Donnell, Gosselin & Schyns, 2002), a 3D plane identical to the one used here (2D image x spatial scales, Bonnar, Gosselin & Schyns, 2002; Gosselin & Schyns, 2001; Schyns, Bonnar & Gosselin, 2002), a translation invariant 1D scale space (Schyns & Gosselin, 2002), and a 3D space comprising the standard 2D image plane and time (Vinette & Gosselin, 2002).

From the discussion above, it should be clear that the number of dimensions making up the stimulus generation space is critical to the number of trials required to reach a stable *Bubbles* solution. Generally speaking, to visit each point of a search space, there is a combinatorial explosion of steps with the increasing number of dimensions. Note, however, that if the dimensions of the search can be collapsed for the analyses, then the search space itself can be large. For example, one could decide that spatial scales are, after all, not that important for FIE, collapse the data along this dimension and analyze feature use in the 2D image plane.

(3) The bubbles. At this stage, two important decisions have been made and the search can almost begin. In the search, information is sampled from the set up space, and the next decision to make concerns the unit of sampling. This unit depends on a number of factors, including the stimuli, the nature of the search space and the task to be performed.

To bring the observer away from ceiling, relevant information must sometimes be sampled, but sometimes not sampled. The parameters of the sampling unit must be adjusted to ensure this modulation of performance. A first parameter is the geometry of the sampling unit. An information sample is effectively a cut in the search space. Sampling unit with different "punch-hole" geometries will change the information sampled and displayed to the observer. Our research has mostly used a Gaussian shaped geometries, either in 2D (Gibson, Gosselin, Wasserman & Schyns, 2002; Gosselin & Schyns, 2001; O'Donnell, Gosselin & Schyns, 2002; Nielsen, Rainer, Brucklacher & Logothetis, 2002 ; Schyns, Jentzsch, Schweinberger, Johnson & Gosselin, 2002) or in 3D (Bonnar, Gosselin & Schyns, 2002; Gosselin & Schyns, 2001; Schyns, Bonnar & Gosselin, 2002). This choice was motivated by two main factors: Gaussians functions produce a *smooth* cut (producing a sample that does not introduce hard edge artifacts), without orientation biases of the sampled information (i.e. a Gaussian is *circularly symmetric*).

A different search space could require geometries other than Gaussians. For example, if orientation information was searched as an independent dimension, the sampling unit would need to introduce orientation biases. For example, a Gabor function could be designed to sample information at several orientations (e.g., 0, 45, 90 and 135 deg). More abstract geometries can also be used, when the search space is itself abstract. For example, in Schyns and Gosselin (2002), the bubble was a dot sampling Fourier coefficients in a Fourier Transform search space.

Another important parameter of the sampling unit is its resolution. The resolution is largely determined by considering the scale of the stimulus and the expected resolution of the relevant information for the task at hand. To illustrate, we know that the eyes, the mouth and the nose are the most useful features to make face decisions. It would therefore be advisable that the resolution of the sampling unit in a FIE task be lower (i.e.

smaller) than the resolution of the important features. A very low-resolution sampling unit (e.g. the pixel of an image) provides a precise sample of the search space, but many trials are required to converge on a solution of the search. Clearly, the resolution of the sampling unit must be chosen with a priori considerations of the expected scale of the solution.

For the reasons just discussed, the bubble of our FIE example has a Gaussian geometry. The scale of the bubble was chosen to sample three cycles per face (i.e. stds of .13, .27, .54, 1.08, and 2.15 deg of visual angle, from fine to coarse scales, see Figure 1). On any given trial, information is sampled from the search space by a number of bubbles. The sampling is typically performed randomly and is thus non-biased. Figure 1 a-e illustrates the sampling procedure. In (b) the face shown in (a) is decomposed into five independent scales. In (c) bubbles with a Gaussian geometry sample the information space at random locations (overlap is permitted). In (d) the bubbles (c) are applied to the appropriate scales in (b). Finally, in (e) the pictures of (d) are added together to produce a sub-sample of the face information in (a).

Insert Figure 1 about here

One important point about bubbles: their number. It can either be adjusted on-line to maintain performance at a given level (e.g., Gosselin & Schyns, 2001; Schyns, Bonnar & Gosselin, 2002), or be kept constant throughout the experiment (i.e., Gibson, Gosselin, Wasserman & Schyns, 2002; Jentzsch, Gosselin, Schweinberger & Schyns, 2002). The technique will work so long as performance is between floor and ceiling. The advantage of adjusting bubble numbers to equate performance is that *Bubbles* solutions are comparable. In the FIE example, we maintained categorization of sampled face information at about 75% correct by adjusting the number of bubbles using a gradient descent algorithm on a trial per trial basis. The initial bubble number resulted from an informed guess (i.e., between 50 and 60 bubbles for a first session), and we let the gradient descent algorithm take over and adjust the bubble number to maintain performance at 75%.

(4) The task. At this stage, the sampling procedure has been fully specified. The final decision is that of the task. We have explored a variety of face categorizations in humans and animals (Gibson, Gosselin, Wasserman & Schyns, 2002; Gosselin & Schyns, 2001; Jentzsch, Gosselin, Schweinberger and Schyns, 2002; O'Donnell, Gosselin & Schyns, 2002; Schyns, Bonnar & Gosselin, 2002; Vinette & Gosselin, 2002), basic and subordinate categorizations of models of animals (Schyns & Gosselin, 2002), and disciminations of an ambiguous painting by Dali (Bonnar, Gosselin & Schyns, 2002). In the FIE example, observers will identify the faces in the upright and inverted conditions.

(5) Observers. Depending on the objectives of the research, different types of observers can interact with the Bubbles algorithm. For example, we have applied the technique to groups of human observers (Gosselin & Schyns, 2001; Schyns, Bonnar & Gosselin, 2002; Bonnar, Gosselin & Schyns, 2002), individual observers to track down effects of expertise acquisition (Jentzsch, Gosselin, Schweinberger and Schyns, 2002; Gosselin & Vinette, 2002; O'Donnell, Gosselin & Schyns, 2002), infants to tackle issues in development (Humphreys, Gosselin, Schyns, Kaufman & Johnson, 2002), pigeons (Gibson, Gosselin, Wasserman & Schyns, 2002), and ideal observers which are models

providing a benchmark of the information available in a task (Gosselin & Schyns, 2001). We have several on-going research projects involving brain-damaged patients (patients suffering from prosopagnosia and hemi-neglect).

(6) Response. The response is an interesting parameter of a *Bubbles* experiment because the technique is in principle sensitive to any measurable dependent variable. Here, observers pressed labeled keys corresponding to the names of 10 individuals. We have used such key-press responses to derive correct and incorrect responses (Gibson et al., 2002; Gosselin & Schyns, 2001; Gosselin & Vinette, 2002; O'Donnell, Gosselin & Schyns, 2002; Schyns et al., 2002; Bonnar, Gosselin & Schyns, 2002) and response latencies (Schyns et al., 2002). In addition, we also used preferential looking (Humphreys, Gosselin, Schyns, Kaufman & Johnson, 2002) and N170 amplitudes (Schyns et al., 2002). Other responses could be the firing rate of single cells, fMRI, galvanic skin response, pleismograph, eye movements, and so forth. To the extent that *Bubbles* is essentially an empirical tool, it is useful to record as many different responses as possible (e.g., correct/incorrect, latencies and N170 in a face recognition experiment). It is difficult to predict before the experiment how responses will correlate with the parameters of the search space.

Analyses

Now that the search has been run, the data are collected, and the analyses can be performed. Remember that the goal of the search is to isolate a subspace of information that determines the measured response(s). Technically, a multiple linear regression on the samples (explanatory variable) and the responses (predictive variable) provides this solution. This reduces to summing all the bubble masks in different response bins, where the number of responses is a function of the nature of the response itself[1]. For example, two bins are sufficient to tease apart correct and correct responses, but more bins are necessary to cover the range of electric activity (or cell firing rate) of a brain response. To reveal the most important information, we can perform a linear operation on the bins (e.g. subtracting the wrong response from the correct response bin; divide the correct response bin by the sum of the correct and incorrect response bins). The result of this operation is usually transformed into Z-scores, and thresholded (e.g., at 1.65, $p < .05$, or 2.33, $p < .01$). The outcome of this analysis is the product of *Bubbles*, revealing the effective subspace of input information. In the visual domain, this outcome is a filtering function that can be applied on the original stimulus to reveal the information that drives the task.

In the FIE example, three observers learned to criterion (perfect identification of all faces twice in a row) the name attached to each of the 10 faces from printed pictures with corresponding name at the bottom. During the experiment, observers had to determine the identity of each sparse face (from 10 possibilities). The experiment comprised six sessions of 780 trials (i.e., 13 presentations of the 30 faces upright and inverted), but we only used the data of the last five sessions (for a total of 3900 trials per subject), when observers were really familiar with the faces and experimental procedure. In a trial, one sparse face computed as described earlier appeared on the screen either upright or inverted. To respond, observers pressed labeled computer-keyboard keys (self-paced, and with correct vs. incorrect feedback). A chin-rest maintained subjects at a 100 cm

[1] However, it is also useful to keep all the sampled information of each trial, to be able to do more detailed analysis such as the conjunctive use of information (see Schyns et al., 2002).

constant viewing distance. Stimuli subtended 5.72 x 5.72 deg of visual angle on the screen.

On average, observers required an average of 46 and 126 bubbles to reach the 75% performance criterion in upright and inverted conditions, respectively. The number of bubbles (between 197 and 30 bubbles, depending on observers and condition) and average performance (between 86% and 75%) did vary across the six experimental sessions, to stabilize in the last session. In this session, observers in upright and inverted respectively required an average of 30 and 65 bubbles for performance levels 75% and 76%. The comparatively higher number of bubbles in the inverted condition suggests a higher requirement of visual information, suggesting a more difficult inverted condition, diagnosing a FIE.

We can now turn to a comparison of the required information in each condition to attain the same level of performance. To this end, we first perform a linear multiple regression. Practically, for each spatial scale, we computed two independent sums: we added together all the information samples leading to correct responses in one sum, and all the information samples leading to incorrect responses in another sum. At each spatial scale, we then subtracted these two sums to construct an image that discriminates the information leading to correct and incorrect responses (see the first row of Figure 2 for the discrimination images at each scale). If all regions of the search space were equally effective at determining the response, the image would be a uniform gray. To pull out the most effective region, we computed Z-scores for each discrimination image, and indicated in red the regions that are 1.65 std away from their mean (corresponding to a $p < .05$). These regions circumscribe the subspace driving upright and inverted face classification responses. If we project the original face in Figure 1 a into this diagnostic subspace, we obtain the effective stimuli displayed the extreme right of the rows in Figure 2. Technically, each effective stimulus is obtained by multiplying the face information at each scale in Figure 1b with the corresponding thresholded coefficients in the rows of Figure 2.

Insert Figure 2 about here

For upright faces, the eyes are the most important local features. The only scales with diagnostic information are the second and the third. This is consistent with the results of Gosselin and Schyns (2001) and Schyns, Bonnar and Gosselin (2002). However, these experiments did not include an inverted condition. Observers saw two, not three expressions, and they were less familiar with the faces (i.e., 1,000 rather than 3,900 trials).

For inverted faces, observers do not seem to rely on any specific features to perform the task. They all seem equally good, or bad. This is reflected in the second row of Figure 2. As the number of bubbles was unequal in upright and inverted, we computed the discrimination image on the normalized Z-score images. Here, the first three scales contain diagnostic information. The eyes, the nose and the right part of the mouth are the most important local features that explain the difference between inverted and upright information use in face processing.

Discussion

We started this problem with a generic methodological question: Given a dependent measurable response (behavioral, electrophysiological, or other) of an organism, how can we determine the optimal subset of parameters from the independent variables that determine the response? With simple stimuli (e.g. Gabor functions, or sinewaves), this is not much of a challenge, because they can only vary along a few degrees of freedom, limiting the complexity of the task. With the stimuli that are typical of realistic face, object and scene categorizations the task had proven so far intractable.

Bubbles is a technique that can resolve the credit assignment problem of attributing the determinants of a response to the parametric subspace of a carefully specified information search space. Using the Face Inversion Effect as an example, we reviewed the six basic decisions to be made to set up a *Bubbles* experiment. In order, deciding the stimulus set, the dimensions within which to search for information, the geometry of the unit to sample information, the task and the response(s) to measure. Applying these to the FIE, we revealed that differences in local face information use (the eyes, nose, and mouth) represented at the scales between 90 and 22.5 cycles per face, determined the effect. This subspace, at least in the visual realm, takes the shape of a diagnostic filtering function that can be applied to render the effective stimulus of the task.

There are obvious shortcomings to *Bubbles*. The first main shortcoming is the combinatorial explosion in number of trials that are required to exhaustively explore the search space. Practically, *Bubbles* is tractable within low-dimensional search spaces, and users are advised to restrict the dimensionality of the search space to be as low as possible. High-dimensional problems are made tractable with heuristics that guide the search towards regions of more promising solutions, leaving aside less promising regions of the space. Heuristic search can be performed with adaptive statistical sampling, or their implementations. However, any heuristic search introduces biases, resulting in a trade-off between speed and suboptimal solutions—i.e. local minima. In any case, the number of trials in *Bubbles* will need to be reduced to apply the techniques to brain damaged patients, children, or when learning is itself a factor of the experiment. To illustrate, with a Gaussian of sigma = 10 pixels and a 256 x 256 pixels image, the search space comprises minimally 25 x 25 different locations to visit, and the solution should converge within less than 500 trials. If 5 scales are added as a third dimension, the solution converges within about 5000 trials. There is little doubt that significant learning occurs during these 5000 trials. We are currently developing several heuristic searches to minimize these numbers (e.g., Leclerc & Gosselin, 2002).

A second shortcoming of *Bubbles* the relationship between the scale and geometry of the sample, and the scale of the solution. The scale and geometry of the sample impose biases on the search space. If the scale is too small with respect to that of the solution, important information will not be revealed within one sample, and the same situation occurs when the geometry of the punch-hole sampling unit does not fit that of the solution. Remember that the *Bubbles* algorithm adaptively adjusts the number of bubbles to maintain the observer at a given performance criterion (e.g. 75% correct). Thus, a higher sampling density, leading to sample overlap, can partially overcome the problems just discussed. However, there is always the possibility that observers will adopt strategies that enable performance to criterion, but are nevertheless atypical of natural human categorization.

At a more theoretical level, one could ask the question of: "What is the information revealed by the *Bubbles* algorithm?" The safe response is "the information required to drive a response at a given performance level." However, to the extent that this information is processed somewhere between the input and the response, it has interesting implications for psychological processing. To illustrate, consider the high-level task of face categorization, and its underlying face features. One interesting property of *Bubbles* is that the researcher can set up a search space that subsumes that of the assumed categorization features (e.g., the eyes, the nose and the mouth). For example, the 3D search space discussed earlier (2D image x spatial scales) formally represents any face feature as a linear combination of information from the scales. Consequently, task-specific face features can emerge in the *Bubbles* solution from the use of information at one, or several of these scales (see Gosselin & Schyns, 2001; Schyns et al., 2002, for examples). Thus, while not applying the search directly to the features, but in a space that represents these features, *Bubbles* can reveal the subspace in which important features are actually represented. It is in such spaces that *Bubbles* solutions tend to be most interesting.

To illustrate, some would argue that the subspace in which important features are represented is in fact the information subspace to which attention is allocated. Earlier, we argued that this diagnostic subspace could be used as a diagnostic filtering function to reveal the effective stimulus. Future research will need to characterize this filtering function, enable bridges to be erected between cognition, attention and perception.

References

Ahumada, A. J., Jr. (1996). Perceptual classification images from Vernier acuity masked by noise, *Perception (Supplement), 26*, 18.

Ahumada, A. J. and Lovell, J. (1971). Stimulus features in signal detection. *Journal of the Acoustical Society of America, 49,* 1751-1756.

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.

Carey, S. and Diamond, R. (1977). From piecemeal to configurational representation of faces. Science, 195, 312-314.

Carey, S., Diamond, R. and Woods, B. (1980). The development of race recognition : a maturational component, Developmental Psychology, 16, 257-269.

Diamond, R. and Carey, S. (1986). Why Faces Are and Are Not Special: An Effect of Expertise. Journal of expérimental Psychology: General, 115, 107-117.

Farah, M.J., Wilson, K.D., Drain, M. and Tanaka, J.N. (1998). What is "special" about face perception? Psychological Review, 105, 482-498.

Freire A., Lee, K. and Symons, L.A. (2000). The face-inversion effect as a deficit in the encoding of configural information: Direct evidence. Perception, 29,159-170.

Gibson, B. M., Gosselin, F., Wasserman, E. A., and Schyns, P. G. (2002). Pigeons use specific and consistent features to discriminate human faces. Submitted.

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

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

Hubel, D. H. (1988). *Eye, Brain and Vision*. New York: Scientific American Library.

Humphreys, K., Gosselin, F., Schyns, P. G., Kaufman, J. and 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.

Jentzsch, I., Gosselin, F., Schweinberger, S. R. and Schyns, P. G.. Applying “Bubbles” to Understand the Face Information Driving Event-Related Brain Potentials. *ECVP*.

Kobatabe, E. & Tanaka, K. (1994). Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex. *Journal of Neurophysiology*, **71**, 856-867.

Leclerc, J. and Gosselin, F. (2002). Adaptative Bubbles. *ECVP.*

Leder, H. and Bruce, V. (2000). When inverse faces are recognized: The role of configural information in face recognition. Quarterly Journal of Experimental Psychology A, 53, 513-536.

Moscovitch, M. Behrmann, M. and Winocur, G. (1997). What is special about face recognition? Nineteen experiments on a person withvisual object agnosia and dyslexia but normal face recognition. Journal of Cognitive Neuroscience, 9, 555-604.

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

Perrett, D. I., Mistlin, A. J., & Chitty, A. J. (1987). Visual neurons responsive to faces.

Trends in Neuroscience, **10**, 358-364.

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

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

Phelps, M.T. and Roberts, W.A. (1994). Memory for pictures of upright and inverted primate faces in humans squirrel-monkeys, and pigeons. Journal of Comparative Psychology, 108, 114-125.

Philips, R.J. and Rawles, R.E. (1979). Recognition of upright and inverted faces: a correlational study. Perception, 43, 39-56.

Rossion, B. and Gauthier, I. (in press). How does the brain process upright and inverted faces? Behavioral and Cognitive Neuroscience Reviews.

Scapinello, K.F. and Yarmey, A.D. (1970). The role of familiarity and orientation in immediate and delayed recognition of pictorial stimuli. Psychonomic Science, 21, 329-331. 18, 555-586.

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

Schyns, P. G. and 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. (2002). A principled method for determining the functionality of brain responses. Submitted.

Searcy, J.H.& Bartlett, J.C. (1996). Inversion and processing of component and spatial-relational information in faces. Journal of experimental Psychology: Human Perception and Performance, 22, 904-915.

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

Simoncelli, E. P. (1999). Image and Multi-scale Pyramid Tools [Computer software]. NewYork: Author

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

Tinbergen, N. and 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.

Valentine, T. and Bruce, V. (1986). Recognizing familiar faces: the role of distinctiveness and familiarity. Canadian Journal of Psychology, 40, 300-305.

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

Vogels, R. (1999). Categorization of complex visual images by rhesus monkeys. Part 2: single-cell study. *European Journal of Neuroscience*. 11, 1239-1255.

Yarmey, A.D. (1971). Recognition memory for familiar “ public ”faces : effects of orientation and delay. Psychonomic Science, 24(6), 286-288

Yin, R.K. (1969). Looking at upside-down faces. Journal of experimental Psychology, 81, 41-145.

Figure captions

Figure 1. Pictures in (b) decomposes (a) in five scales; (c) illustrate the "bubbles" applied to each scale; (d) are the revealed information of (b) by the bubbles of (c). Note that on this particular trial there is no revealed information at the fifth scale. By integrating the pictures in (d) we obtain (e), a sample stimulus (Gosselin & Schyns, 2001; Schyns, Bonnar & Gosselin, 2002). Picture (f) is (a) sampled in the image plane (Gosselin & Schyns, 2001; Gibson, Gosselin, Wasserman & Schyns, 2002; Jentzsch, Gosselin, Schweinberger & Schyns, 2002; O'Donnell, Schyns & Gosselin, 2002). Picture (g) is a 3D shape model of a dog sampled in phase space (Schyns & Gosselin, 2002). Finally, picture (h) is the ambiguous area of a Dali painting sampled in the same generation space as (e) (Bonnar, Gosselin & Schyns, 2002).

Figure 2. The first row of this figure applies to the upright condition. It gives the five classification images at each scale, from finest to coarsest. The red areas revealing a face are significantly above chance ($p < .05$). The rightmost picture is the effective face. The second and third rows are the same the first one, except that they apply to the inverted condition and the difference between the upright and inverted conditions, respectively.

Figure 1

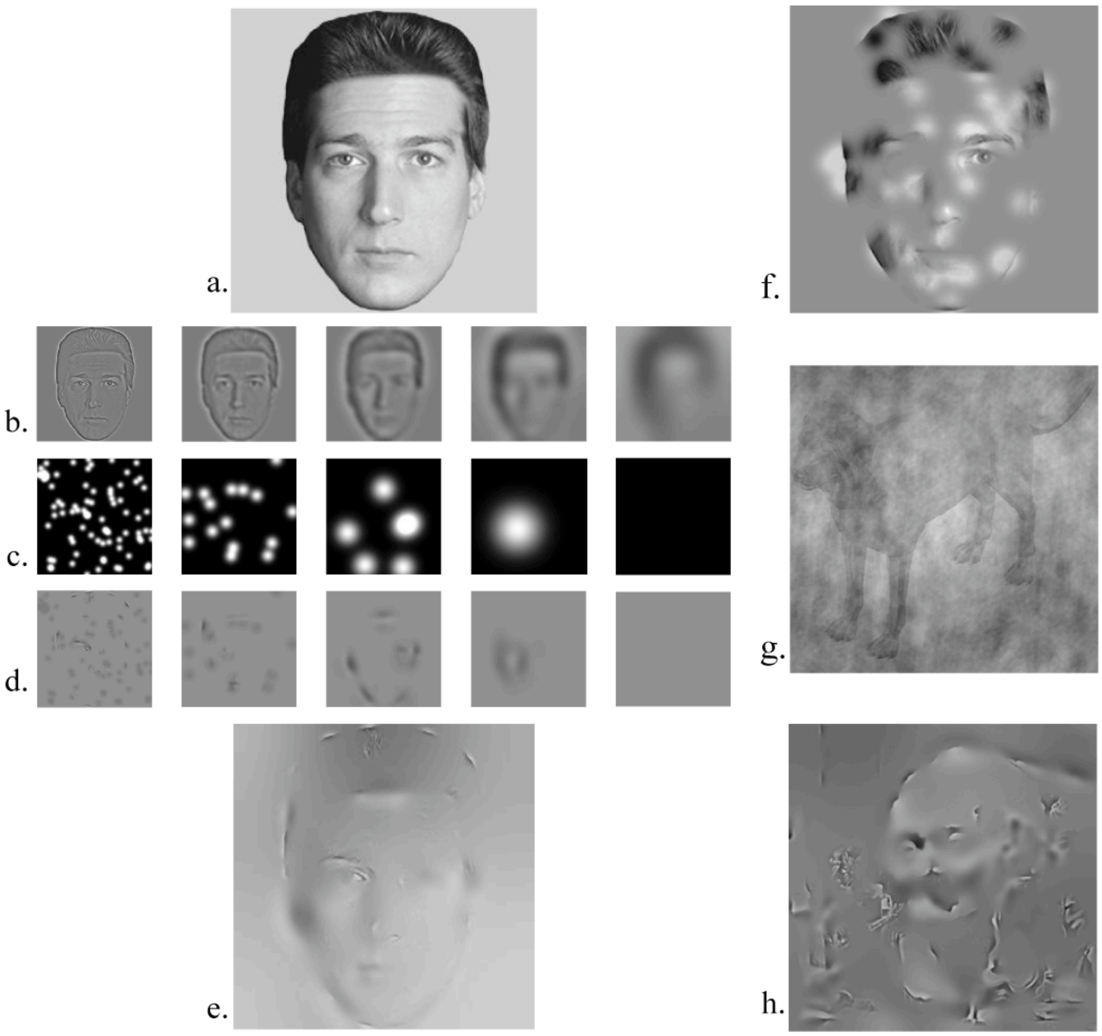

Figure 2

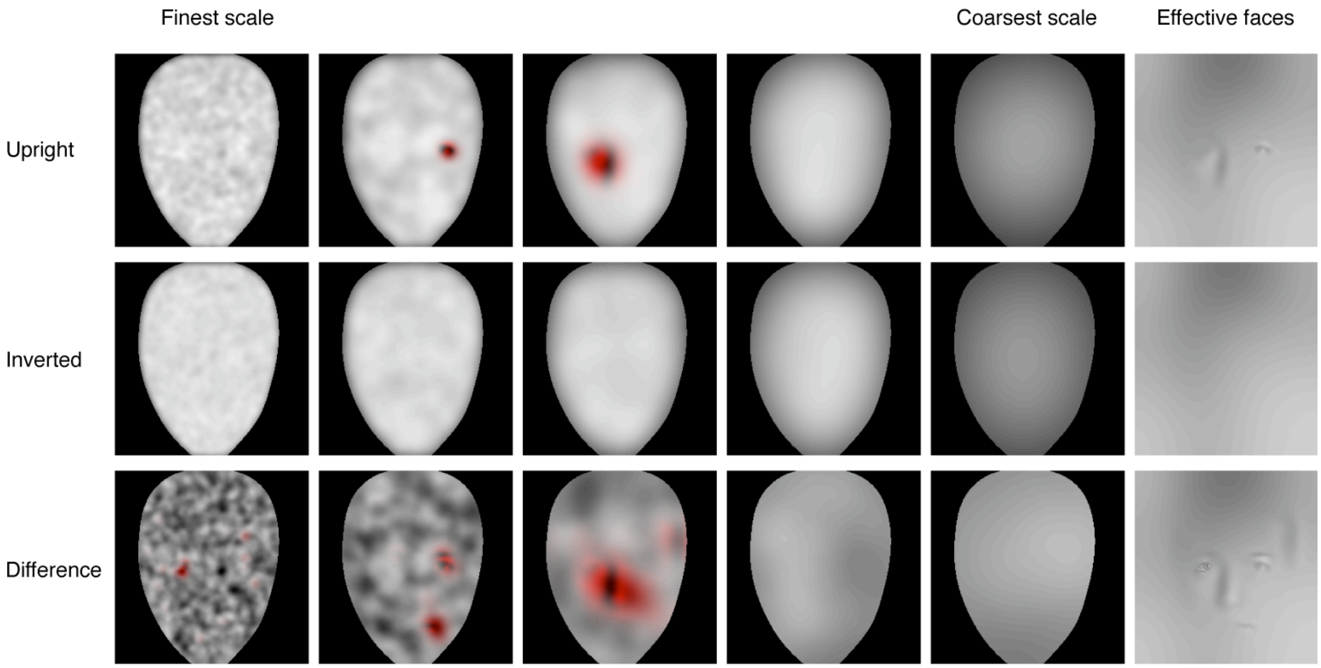
Finest scale
Coarsest scale
Effective faces
Upright
Inverted
Difference