Journal
xxx

Learning an Artist's Style: Just What *Does* a Pigeon See in a Picasso?

John R. Vokey
Department of Psychology
University of Lethbridge

Jason M. Tangen
School of Psychology
University of Queensland

Judgements of style in art, music, and literature are commonplace, although the mechanisms providing for this structural sensitivity are not well understood. Watanabe, Sakamoto, and Wakita (1995) showed that pigeons trained to discriminate colour slides of paintings of Picasso from those of Monet could generalise this discrimination not only to new paintings of Picasso and Monet, but also to paintings of other Cubist and Impressionist painters. This tacit sensitivity to artistic style is explored in terms of a simple PCA network model applied to pixel-maps of the paintings. The eigenvectors obtained from the singular value decomposition of sets of these pixel-maps provide for descriptions of the stimuli in terms of visual "macro-features". These macro-features provide a simple basis for the successful discrimination of novel paintings into various style categories. The results suggest that the eigen-decomposition is a necessary first-step, and that the bases for judgments of style may indeed be quite simple.

From the initial demonstration by Herrnstein and Loveland (1964), we have seen over 40 years of relatively—and sometimes, remarkably—*sophisticated* discriminative judgements by pigeons. It is apparent that pigeons can be trained to discriminate complex categorical structure depicted in visual images of natural and artificial scenes. Herrnstein and Loveland (1964) found that laboratory pigeons could be trained to discriminate correctly photographic slides depicting humans (or parts thereof) from otherwise similar slides not depicting humans. More important, they could generalise this ability to previously unseen slides.

This work was supported by an operating grant from the Natural Sciences and Engineering Research Council of Canada (NSERC) to the first author, and a postgraduate NSERC fellowship to the second author. This work was presented as part of a larger project at the annual meeting of the Canadian Society for Brain, Behaviour, and Cognitive Science (BBCS), in Edmonton, Alberta, June 20, 1999, the joint meeting of BBCS and the Experimental Psychology Society of Britain (EPS) in Cambridge, UK, July 21, 2000, and the annual meeting of the Canadian Society for Brain, Behaviour, and Cognitive Science (BBCS) at Université Laval, Québec, June 25, 2001.

We thank Hervé Abdi, Alice O'Toole, Dominique Valentin, and Betty Edelman whose earlier work with autoassociative neural network models of face perception served as a model for the current research. We also thank Shigeru Watanabe who responded with both patience and generosity to our many inquiries about the materials and procedures of his experiments on the learning of artistic style in pigeons. We beg the indulgence of Monen, Brenner, and Reynarts (1998), from whom we stole the subtitle.

Reprint requests should be sent to Dr. John R. Vokey, Behaviour and Evolution Research Group, Microcognition Laboratory, Department of Psychology, University of Lethbridge, Lethbridge, Alberta, Canada, T1K 3M4. Electronic mail should be sent to: vokey@uleth.ca

What is less clear is how the pigeons accomplish the task. The Herrnstein and Loveland (1964) slides were of a wide variety of indoor and outdoor scenes, and differed widely in the persons shown and in the way they were depicted. Because of this variety and the fact that the person slides did not otherwise appear to the experimenters to differ in any other systematic way from the non-person slides, Herrnstein and Loveland (1964) concluded that the pigeons were discriminating the concept of "person" as opposed to having learned to respond to discriminative differences along some simple, physical dimension of the slides that happened to be correlated with the person/non-person distinction among them (see Siegel & Honig, 1970, for related findings).

Malott and Siddall (1972) trained pigeons to make this same person/non-person discrimination with as few as a dozen training exemplars rather than the hundreds used by Herrnstein and Loveland (1964) and yet their pigeons still appeared able to generalise the discrimination to an apparently unlimited set of test stimuli. Herrnstein and Loveland (1964) had argued that the ease with which their pigeons learned to make the discrimination suggested that the training served not so much to teach the pigeons a new concept as to train them to use their pre-existing concept of "person" to behave discriminatively in the task. That is, in their view, what the pigeons acquired in the task was not a new conceptualisation (i.e., perceptual organisation and representation) of two-dimensional scenes, but a simple mapping of a pre-existing conceptualisation to the reward demands of the task. The minimal training results of Malott and Siddall (1972) were subsequently taken by Herrnstein, Loveland, and Cable (1976) and Herrnstein and de Villiers (1980) to reinforce that view. The issue of whether the pigeons should be seen as *acquiring* new concepts (i.e., perceptual organisations) in these experiments or simply *mapping* pre-existing ones, if any, in response to reward contingencies is one we return to in part later.

Herrnstein et al. (1976) followed up the Herrnstein and Loveland (1964) study with a new set of slides and discriminations of what they referred to as "natural concepts in pigeons". Herrnstein et al. (1976) collected a set of slides consisting of hundreds of colour photographs of indoor and outdoor scenes of Cambridge (USA) taken over a period of a year around 1970. In three separate experiments, pigeons were trained to discriminate either slides containing trees from those containing no trees (see also Herrnstein, 1979; Vaughan & Herrnstein, 1987), slides depicting water from those with no water, or slides displaying a particular female person from slides not showing her, including those depicting other persons. As in Herrnstein and Loveland (1964), not only could the pigeons learn to make these discriminations with respect to particular training slides, they also were able to generalise this ability to previously unseen test slides. Also as in Herrnstein and Loveland (1964), Herrnstein et al. (1976) could find no common elements that they believed would enable such discrimination at a level below that of or different from the complex categories of "tree", "water", or the particular person. Such categories are difficult, if not impossible, to specify in simple stimulus terms. In acknowledging this difficulty, Herrnstein et al. (1976) wrote:

> {H}aving looked at the hundreds of instances used here or even at the two positive instances {shown in a Figure in the original paper} (let alone the tens of thousands involved in real-life discriminations), we cannot begin to draw up a list of common elements. To recognise a tree, the pigeons did not require that it be green, leafy, vertical, woody, branching, and so on (overlooking the problem of common elements nested within terms like leafy, vertical, woody, and so on). Moreover, to be recognizable as a non-tree, a picture did not have to omit greenness, woodiness, branchiness, verticality, and so on. Neither could we identify common elements in the other two experiments. (pp. 297-298)

Other research has replicated and extended these basic findings. For example, Poole and Lander (1971) trained pigeons to discriminate colour slides of different breeds of pigeons from those of other species of birds, animals, or objects. Ceralla (1979) found that pigeons could be trained to discriminate silhouettes of oak leaves from those of other deciduous trees, even when the training consisted of only a single positive exemplar. Interestingly, the pigeons could *not* be trained to discriminate this same exemplar from other examples of oak leaves. Lubow (1974) trained pigeons to discriminate between photographs of aerial views of artefacts and similar photographs containing no artefacts, and as in the other studies, to generalise successfully this discriminative ability to previously unseen aerial views. Even more surprising for a nominally land-based, cliff-dwelling species, Herrnstein and de Villiers (1980) showed that pigeons could learn to discriminate (and subsequently generalise from) colour photographs of underwater scenes containing fish from those without fish.

Moving to perhaps less natural categorical distinctions, Morgan, Fitch, Holman, and Lea (1976) taught pigeons to discriminate between slides depicting the letter 'A' and the numeral '2' across a wide set of type-faces, and found that the pigeons could generalize this ability to previously unseen type-faces, and even to previously unseen, hand-written versions of the characters. And Ceralla (1980) showed that pigeons could be trained to discriminate between cartoon panels from the *Peanuts* comic strip containing the Charlie Brown character from those depicting other characters (e.g., Lucy, Linus, Snoopy) see Huber (2001, for further examples and discussion).

Several recent experiments have set out to pinpoint the particular aspects of the photographs that pigeons rely on in making these discriminations. Troje, Huber, Loidolt, Aust, and Fieder (1999) independently manipulated the shape and texture of photographs of synthetic male and female faces, and demonstrated that pigeons could correctly discriminate the gender of the faces depicted in the photos, and suggested that "texture" information (i.e., colour, intensity, and other first-order statistics) was particularly relevant. Gibson, Wasserman, Gosselin, and Schyns (2005) used the "Bubbles" procedure (Gosselin & Schyns, 2001) to isolate the features of similar synthetic face stimuli that control pigeons' discriminative response. They found that the pigeons were particularly sensitive to the area around the mouth to discriminate photos of happy and neutral facial expressions, and around the eyes and chin to discriminate male and female face images. Aust and Huber (2001) replicated the original Herrnstein and Loveland (1964) human/non-human discrimination task and scrambled the stimuli separating the spatial arrangement of the photographs into fragments. They found that the pigeons were still largely capable of discriminating the coloured photos despite the fragmentation (see Aust & Huber, 2003, for a similar approach using scrambled human figures within the intact photograph). Aust and Huber (2002) presented pigeons with photographs showing humans in the distance, in silhouette, isolated parts of human figures (e.g., hands, feet, heads), and "pseudohuman stimuli" (e.g., dolls, scarecrows, snowmen), mammals and birds, and photos of isolated clothing and clothes worn by animals. They found that the pigeons regularly classified novel photos depicting humans in the distance and in silhouette as instances of the people-absent category. They tended to classify photos containing hands as members of the people-present category, whereas feet and skin were assigned to the people-absent category. The pigeons' responses to the pseudohuman stimuli, non-humans and clothing were highly variable and inconclusive.

A Matter of Style: Watanabe et al. (1995)

In a paper that forms the focus of the current research, Watanabe et al. (1995) showed that pigeons could learn to discriminate colour slides of paintings by Picasso from his Cubist period from Impressionist paintings by Monet, and then generalise this discrimination not only to previously unseen Cubist and Impressionist paintings by Picasso and Monet, but to paintings by other Cubist (e.g., Braque) and Impressionist (e.g., Cezanne) artists (see also Watanabe, 2001, for similar results with paintings by Van Gogh and Chagall). These results are especially noteworthy because, unlike most of the previous stimuli and discriminations nominally required of the pigeons—such as the "tree", "water" and particular person discriminations in Herrnstein et al. (1976)—it is clearly the individual stimuli *as a whole* and not specific aspects of their content that are to be discriminated into categories.

In the Watanabe et al. (1995) experiments, the pigeons were required to discriminate art on the basis of overall *style*: Cubist paintings by Picasso are "Cubist Picassos" by virtue of the fact they were painted by a particular artist who painted in a particular style. Picasso, for example, painted in multiple styles, such as his earlier "Blue" and "Rose" periods and his subsequent "Cubist" period. Cubist painters are so-called at least in part because they painted with very similar styles to one another. Indeed, the Cubist paintings of Picasso and Braque, for example, are so remarkably similar that it is a significant challenge for the typical artistic layperson—us, at any rate—to discriminate one from the other. It is important to note, however, that visual commonality does not *define* whether or not a painting is, say, Cubist. That is a matter of history and a particular artistic philosophy. Instead, the resulting common visual style found among Cubist paintings should be considered to be a consequence of the artistic period and philosophy that constrain the paintings to have a certain "look" (i.e., a highly-correlated, but not defining property). Monen et al. (1998) used precisely this point to argue against the conclusion from Watanabe et al. (1995) that the pigeons could sensibly be said to have acquired the concepts of "Picasso" or "Cubist".

In contrast, in the experiments of, say, Herrnstein et al. (1976), what the birds learned was said to rely on the "objects" depicted: the overall style of the photograph (e.g., background, other objects, colouring, shading, etc.) was nominally irrelevant. Indeed, it is precisely this nominal irrelevance (and wide variability over slides) that led Herrnstein et al., among others, to conclude that it was indeed the *concepts* of "tree", "water", "person", and so on that the pigeons were discriminating in these tasks (see also Herrnstein, 1990; Herrnstein & de Villiers, 1980).

But was it? By "style" we don't mean to imply anything fundamentally different from the visual manifestation of the (usually) natural language-based concepts of Wittgenstein (1968) or Ryle (1951). Indeed, many of the visual classes used in these studies with pigeons, such as "tree", "water", and "person" are probably best characterised as "polymorphous" (Ryle, 1951) categories held together by a visual form of "family-resemblance" (Wittgenstein, 1968). Herrnstein et al. (1976), for example, explicitly cited these ideas in their discussion of what the pigeons were discriminating (see also Bhatt, Wasserman, Reynolds, Jr., & Knauss, 1988; Lea & Harrison, 1978; Roberts & Mazmanian, 1988; Wasserman, Kiedinger, & Bhatt, 1988, for further examples and discussion). Our point in using "style" is to emphasise the distinction revealed by contrasting the photos of embedded objects used by Herrnstein et al. (1976) with the paintings used by Watanabe et al. (1995). The birds may not have learned to discriminate the polymorphous natural language concepts of the embedded objects in particular, but rather, as with the Cubist and Impressionist paintings, the overall "style" of the photographs. Because the photographs used by Herrnstein et al. (1976) were required to depict or not depict particular embedded objects (e.g., "tree"), they may have been similarly constrained to have a certain "look"(and, thereby, equally, a polymorphous category based on family-resemblance).

It is for this reason that we have taken some care in our summaries of the earlier experiments—such as those of Herrnstein et al. (1976)—to refer to the discriminations as one of discriminating *slides* or *photographs* of, say, trees from those of non-trees, rather than one of discriminating trees directly; we really have no firm evidence that pigeons can discriminate trees, water, etc. *per se* in photographs. Of course, the argument cuts both ways: possibly the pigeons in Watanabe et al. (1995) relied on embedded objects that were unique to Cubist or Impressionist paintings, and it was the presence of these depicted objects that supported the discrimination. After all, some of Picasso's Cubist paintings did tend to favour a mandolin, and Monet was terribly fond of water lilies. Inspection of the paintings used by Watanabe et al. (1995) reveals no such obvious commonalities, but, as with the earlier quotation from Herrnstein et al. (1976), our failure to detect such is not compelling evidence that none exist.

If photographs of trees or water or a particular person tend to have a certain "look" because of the general constraints on the photograph imposed by the presence of the embedded object—or, if Cubist paintings tend to have similarly constrained "look" imposed by painting in that style, how could it be captured for analysis? Our approach to this issue is to analyse the stimuli—in this case the representations of the paintings used by Watanabe et al. (1995), at their lowest level of representation, below that of the natural language, perceptual categories and features favoured by (explicit) human perception, that of pixel-maps. We present a statistical analysis of the colour values of the individual pixels ("picture elements") that comprise each of the paintings. We do so by systematically simulating the results of each of the Watanabe et al. (1995) experiments using a simple, artificial neural network, an autoassociator, applied to the pixel-maps of the images (see Huber, Troje, Loidolt, Aust, & Grass, 2000; Troje et al., 1999, for a related experiments).[1]

Simulating Watanabe et al. (1995)

The first experiment reported by Watanabe et al. (1995) used two sets of training stimuli, referred to as "set A" and "set B". Each set contained a unique set of 10 paintings by Monet, and 10 paintings by Picasso from his Cubist period,

[1] We have conducted similar simulations with some of the photographic stimuli and categories of Herrnstein et al. (1976). The results of these simulations have been reported at the joint meeting of the British Experimental Psychology Society and the Canadian Society for Brain, Behaviour & Cognitive Science, Cambridge, UK, July 21, 2000, and at the annual meeting of the Canadian Society for Brain, Behaviour, & Cognitive Science, Edmonton, Alberta, June 19, 1999.

thought to be typical of each artist.[2] For presentation to the pigeons, the set A paintings were reproduced as colour slides, and set B paintings were converted to video. Because we were unable to locate copies of all of the set B paintings, only the set A stimuli were used in the simulations reported here. The complete set of stimuli used in the simulations is shown in Table 1.

Method

Stimuli. The set A paintings varied in both size and shape from one another; for conversion to slides, Watanabe et al. (1995) re-scaled each painting so that the smaller of its height or width just filled the corresponding slide dimension; each image was then cropped, centred on the other dimension.[3] The same approach was used in the simulations reported here. Photographic slide dimensions are in the ratio of 5:7. We arbitrarily designated the longer slide dimension to correspond to the height of the images (i.e., as if taller than wide paintings had been photographed with the camera rotated by 90 degrees from the camera's "normal" landscape orientation). Each of the paintings used by Watanabe et al. (1995) was scanned from various art books and reproduced as colour, computer graphic images. Each image was then re-scaled and cropped as just described so that the height was 175 pixels, and the width was 125 pixels. The colour of the paintings was represented by coding each image as three separate Red, Green, and Blue (RGB), 8-bit sub-images—the intensity of the pixel in each sub-image corresponding to the intensity (on a scale of 0-255) of the corresponding colour. Thus, each painting was represented as a cropped computer graphic image consisting of $125 \times 175 \times 3 = 65{,}625$ pixels.

Simulated Subjects. In Watanabe et al. (1995) one-half the pigeons receiving set A were trained to discriminate the Monet paintings from the Picasso paintings with Monet as the positive (rewarded) category, and the remainder with Picasso as the positive category. The training criterion for each pigeon was a *discrimination ratio* exceeding 90% over two successive sessions—where a session was a random ordering of all 20 training stimuli for that pigeon. The discrimination ratio was calculated as the ratio of pecks-to-targets to pecks-to-targets *plus* pecks-to-foils, and the pigeons took from 6 to 22 sessions to reach the criterion.

To simulate this variability both between and within pigeons, each trial a given simulated pigeon had with a given stimulus consisted of a random block of the pixels from the stimulus. On each trial, a random square of 100×100 pixels was extracted from the 125×175 pixel stimulus, and stored, coded for colour, as a vector $100 \times 100 \times 3 = 30,000$ pixels in length, representing the simulated pigeon's visual experience of the stimulus for that trial. The idea here was to capture the notion that no stimulus is ever encountered in exactly the same way twice. Thus, as there are $26 \times 76 = 1{,}976$ different samples possible for each image, both within and between the simulated pigeons, a trial with a given stimulus is unlikely to be literally identical with any other.[4] The variability also provides a basis for statistical analyses. Each of 20 simulated pigeons, one-half trained with Monet as the positive category and the remainder with Picasso as the positive category, was given 10 such trials with each of the 20 stimuli from set A of Watanabe et al. (1995).

Procedure. For each simulated pigeon, the 200 samples (10 trials $\times$ 10 stimuli $\times$ 2 artists) were learned as a linear autoassociative memory using the Widrow-Hoff learning algorithm. This memory—the 30,000 pixels $\times$ 30,000 pixels weight matrix, $\mathbf{W}$, relating the connection value between each pixel and every other pixel over the 200 samples—can be computed via the singular-value-decomposition (SVD) of the 30,000 pixels $\times$ 200 matrix, $\mathbf{X}$, of training stimuli. The SVD of a rectangular matrix, $\mathbf{X}$, is expressed as $\mathbf{X} = \mathbf{U}\boldsymbol{\Delta}\mathbf{V}^{\mathrm{T}}$, for which $\mathbf{U}$ is the matrix of eigenvectors of $\mathbf{X}\mathbf{X}^{\mathrm{T}}$, $\mathbf{V}$ is the matrix of eigenvectors of $\mathbf{X}^{\mathrm{T}}\mathbf{X}$, and $\boldsymbol{\Delta}$ is the diagonal matrix of singular values—the square-root of the eigenvalues of either $\mathbf{X}\mathbf{X}^{\mathrm{T}}$ or $\mathbf{X}^{\mathrm{T}}\mathbf{X}$ (as they are the same).[5] In statistics, the related eigendecomposition of the data matrix is called principal components analysis (PCA), and so such linear autoassociators are often referred to as PCA neural networks (see Abdi, Valentin, & Edelman, 1999). From this perspective, $\mathbf{W}$ can be represented in terms of the eigenvectors, $\mathbf{U}$, of the pixels $\times$ pixels cross-products matrix (see Abdi et al., 1999):

$$\mathbf{W} = \boldsymbol{\delta}\mathbf{U}\mathbf{U}^{\mathrm{T}}$$

where $\boldsymbol{\delta}$ corresponds to the eigenvalues. The effect of Widrow-Hoff learning is to *spherise* the weight matrix, i.e., render all of the resultant eigenvectors equally important in reconstructing the stimuli (Abdi et al., 1999), yielding:

$$\mathbf{W} = \mathbf{U}\mathbf{U}^{\mathrm{T}}$$

Retrieval of an item from this memory, $\hat{\mathbf{x}}_i$, is computed as

$$\begin{aligned} \hat{\mathbf{x}}_i &= \mathbf{W}\mathbf{x}_i && (1) \\ &= \mathbf{U}_{l:m}(\mathbf{U}_{l:m}^{\mathrm{T}}\mathbf{x}_i) && (2) \end{aligned}$$

where the subscript, *l:m*, denotes the range of eigenvectors used to reconstruct the item. For our purposes, the eigenvectors are ordered in terms of the magnitude of the associated eigenvalues (i.e., proportion of variance accounted for), from most to least. As only the eigenvectors with associated eigenvalues greater than zero are retained, there are at most as

[2] The names of the paintings originally used are given in Table 1 of Watanabe et al. (1995). Unfortunately, this table is not complete: one painting by Monet from each set is not listed. The missing Monet paintings are "An Impression: Sun Rise" for set A, and "Dinner of Sisley" for set B (S. Watanabe, personal communication, March 4, 2001).

[3] S. Watanabe, personal communication, March 25, 1999.

[4] Admittedly, rendering these experiences as completely random with respect to one another, constrained only by the stimulus itself, so that between and within simulated pigeon variances on the same stimulus are the same, is a bit unrealistic. However, including the more realistic encoding correlations and biases (e.g., a pigeon rewarded for looking at the bottom-left of an image would be more likely to do so again) seemed like overkill, and, at any rate was not necessary for successful simulations.

[5] $\mathbf{X}^{\mathrm{T}}$ denotes the transposition of matrix $\mathbf{X}$.

Table 1
Paintings used as the training set in the simulation of Watanabe et al. (1995)

Paintings by Artists	Year
Monet	
Terrace at Saint-Adresse	1866
Dinner of Sisley	1869
Mrs. Monet	1870
Impression, Sunrise	1872
Red poppy	1873
Capucines boulevard	1873
Pears and grapes	1880
Poplars of Giverny	1888
Pond of water lily	1899
Palazzo da mula in Venezia	1908
Picasso	
Girls in Avignon	1907
Man with a violin	1911
Girl with a ring	1919
Guitar, cup and fruits	1924
Model and artist	1928
Woman playing with a ball on the beach	1932
Nude woman with a comb	1940
Still life with a pan	1945
Women of Algeria	1955
Nude woman under the pine tree	1959

many eigenvectors as there are items in the training set. The expression in parentheses of Equation 2 can be interpreted as the *projection*, $\mathbf{p}_{i|l:m}$, of the item into the space defined by the eigenvectors,

$$\mathbf{p}_{i|l:m} = \mathbf{U}_{l:m}^{\mathrm{T}}\mathbf{x}_i$$

where the values of $\mathbf{p}_{i|l:m}$ are the *weights* on each eigenvector used to reconstruct the item from the linear combination of eigenvectors:

$$\hat{\mathbf{x}}_i = \mathbf{U}_{l:m}\mathbf{p}_{i|l:m}$$

Thus, given the eigenvectors of the set as a whole, each item can be represented in a very reduced form as its 200 projection weights on the eigenvectors. It is in this sense that the eigenvectors can be seen as the "macrofeatures" of the items (see, e.g., Abdi, Valentin, Edelman, & O'Toole, 1995; Huber et al., 2000; O'Toole, Abdi, Deffenbacher, & Valentin, 1993; O'Toole, Deffenbacher, Valentin, & Abdi, 1994; Turk & Pentland, 1991; Valentin, Abdi, & O'Toole, 1994; Valentin, Abdi, O'Toole, & Cottrell, 1994, for similar analyses of photographs of faces).

The learning of the labels (Monet/Picasso) associated with the stimuli was simulated by training a simple classifier, a variant of a *perceptron* known as an "adaline" (see, e.g., Anderson, 1995). The adaline is a simple linear heteroassociator with Widrow-Hoff error-correction, composed of a multiple-unit input layer and one binary output unit. In statistical terms, it is a simple linear discriminant function analysis of the inputs to predict the binary classification of the items (see, e.g., Abdi et al., 1995).[6] The inputs to the classifier were the projection weights on the eigenvectors for each trial item to produce a final set of discriminative weights to predict the artist category, in the form of a simple linear equation, from the projection weights for any given input item. This approach is equivalent to fitting a hyper-plane to the projections of the items that best (in the sense of the least-squares criterion) separates the Monet training inputs from the Picasso training inputs—perfectly, if all eigenvectors are used (a consequence of the Widrow-Hoff error-correction), or maximally if some sub-set of eigenvectors is used. These prediction weights were then frozen for test, and used to predict the artist category from the projection weights of the test stimuli.

The first question is how well what the simulated pigeons learned from their 200 training experiences generalised to new experiences with the same training images. To answer that question, each simulated pigeon was presented with one more randomly-sampled block of each training image, and was asked to classify it based on its earlier experiences. No feedback was given; that is, none of the learning from the previous experiences was adjusted, nor were these test experiences stored in the autoassociative memory. Thus, each new sample of each of the training stimuli was projected into the space defined by the eigenvectors of the 200 training experiences for a given simulated pigeon; that is, the new sample image was *encoded* in terms of the original 200 training experiences.

[6] Given the binary dependent variable, it is in this case the mathematical equivalent of a simple, multiple linear regression of the projection weights on the eigenvectors as the independent variables.

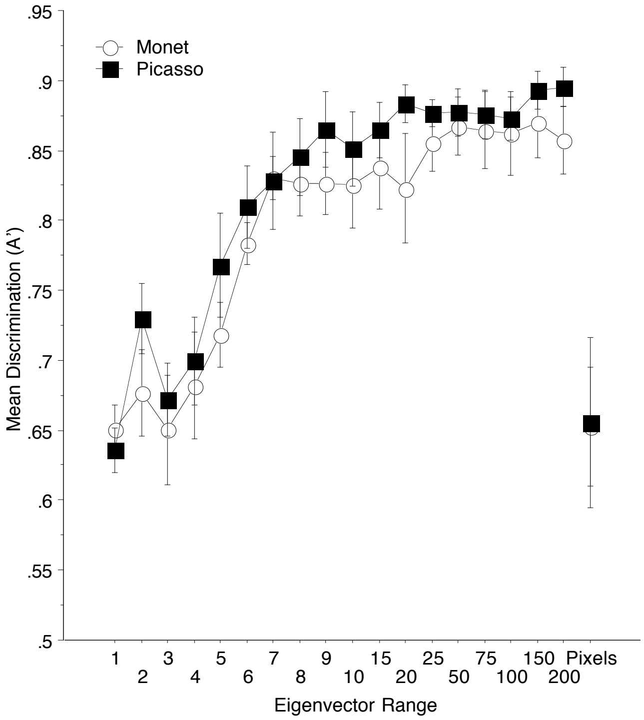

Figure 1. Mean training discrimination (A') of the Monet from the Picasso training stimuli as a function of "macrofeature" or eigenvector range for the simulated pigeons that were trained with colour stimuli from either Monet (open circles) or Picasso (filled squares) as the positive category. Error-bars are within-cell, standard errors.

The projection weights of this encoding were then passed to the classifier to predict the artist category of the new sample image. This test was intended to correspond to the last session of training trials of the pigeons in Watanabe et al. (1995), and provides a measure of the maximum performance possible for the simulated pigeons when given subsequent tests in which the test images are manipulated in various ways.

Results and Discussion

For each simulated pigeon, the 20 classification responses for different ranges of eigenvectors (i.e., representing the use of more and more of the "macrofeatures" of higher dimensionality) of the 10 Monet and 10 Picasso test experiences were scored as hits and false-alarms (for the respective positive category it had been trained with), and then converted to a non-parametric, signal-detection measure of discrimination, A'. Values of A' vary between 0.00 and 1.00, and approximate the results of a two-alternative forced-choice task with the same discriminative stimuli; a value of $A' = 0.50$ indicates "chance" discrimination; values of A' greater than 0.50 indicate increasingly successful levels of discrimination. This discrimination index was computed for each simulated pigeon for classification based on just the first "macrofeature" or eigenvector, the first 2, the first 3, . . . , 10, 15, 20, 25, 50, 75, 100, 150, and all 200 eigenvectors.

The data were subjected to a 2 (positive category) × 18 (eigenvector range) mixed ANOVA with simulated pigeons crossing eigenvector range, but nested within positive category, as the random variate. The results are shown in Figure 1. Clearly, substantial levels of discrimination are possible with these simulated pigeons. As can be seen in Figure 1, discrimination increased as more eigenvectors were used to make the discrimination [$F(17, 306) = 33.42$, $MSE = 0.004$], although the effect appeared to asymptote once the first 8 or so eigenvectors were included. Whether Monet or Picasso served as the positive training category had no effect [$F(1, 306) = 1.04$, $MSE = 0.051$], nor did it interact significantly with eigenvector range ($F < 1$).

Is the eigen-decomposition necessary?. It is possible that the high levels of successful discrimination achieved here reflect *not* the eigendecomposition of the items, but rather the robusteness of the discriminant function analysis inherent in the perceptron classifier. That is, is there any advantage to representing the items in terms of eigenvectors or "macrofeatures"? For example, it is possible that the Monet items are generally darker than the Picasso items, or contain more blue, and, hence, may be discriminated directly in terms of these mean rather than covariant differences in the pixel values themselves. To assess this possibility, the classifier was applied directly to the pixel-maps of the 200 training images for each simulated pigeon, and then the obtained weights applied directly to the pixels of the test samples of the training items to predict their classification. As with the earlier analysis, the resulting hits and false-alarms were converted to values of the discrimination index A'.

These data were subjected to a 1-way (positive category) between-subjects ANOVA. The results are also shown in Figure 1 as the two isolated points on the right-hand side of the figure above the abscissa label "Pixels". Discrimination A' did not differ significantly as a function of positive category ($F < 1$), but was significantly different from the nominal chance level of 0.50 for both positive training categories, $t(9) = 3.57$ for Monet as the positive training category and $t(9) = 2.57$ for Picasso. Clearly, it is possible to achieve above-chance levels of discrimination without encoding the items in terms of the eigenvectors of the set. But, as is also clearly evident in Figure 1, much higher levels of discrimination are possible if the eigen-decomposition is performed as the first step, especially if more than just the first few eigenvectors are used. This observation was confirmed in a 2 (positive training category) × 2 (eigenvectors vs. pixel-maps) mixed ANOVA with simulated pigeons crossing the latter factor, but nested within positive category, as the random variate. On average (i.e., collapsed over eigenvectors), encoding the simuli in terms of the eigenvectors resulted in significantly greater discrimination than did applying the perceptron directly to the pixel-maps [$F(1, 18) = 27.53$, $MSE = 0.010$], and this effect was the same in both positive training conditions [$F(1, 18) = 1.16$, $MSE = 0.010$], which did not differ significantly from one another ($F < 1$).

The role of colour, sharpness, and orientation. In Watanabe et al. (1995), training was followed by three tests, designed to probe various potential differences between the Monet and Picasso training stimuli as the source of the pigeons' successful discrimination of them. They focused on three such differences: the colours favoured by the two artists (Picasso preferred bright, saturated colours to Monet's muted, pastel tones), the marked differences in image sharpness (edges in the paintings by Monet are blurred relative to the sharp edges in the Picasso paintings), and the degree of orientation dependence (Picasso's relatively abstract representations contrast with the more naturalistic representations of Monet). To test whether the successful discrimination of the training stimuli was a function of any of these differences, following training to criterion, Watanabe et al. (1995) gave each pigeon one session each of (a) the training stimuli presented in monochrome (i.e., gray-scale) (b) blurred (to soften the sharp edges of the Picasso paintings) (c) and flopped (left-right reversal) or inverted (up-down) orientation. In each case, Watanabe et al. (1995) found that their pigeons were able to maintain the discrimination with the modified training stimuli, suggesting that such differences did not play a major role in their pigeons' successful discrimination of Monet from Picasso training slides.

The same tests were simulated here. To create the monochrome versions of the training stimuli, each of the original training images for each simulated pigeon was converted to an 8-bit grey-scale image, and the grey-scale was then propagated to each of the RGB sub-images. The test images were then sampled from these grey-scale images as described for the test of training.

Watanabe et al. (1995) blurred their stimuli by presenting the training test images out of focus. The criterion for the degree of blur was that two, 0.43 mm wide lines separated by 0.43 mm fused when viewed on the screen. We simulated that degree of blur by independently convolving each of the three colour separation sub-images of the (newly) sampled images of the training stimuli with a 3×3 filter consisting of a uniform distribution of coefficients of 1.0. On such small (100×100 pixels) images, this degree of smoothing is most likely even more extreme than that used by Watanabe et al. (1995).

For the orientation tests, the original training images were flopped on the horizontal or the vertical axis, and then test images sampled from them. Otherwise, these test simulations were identical to the original training test.

The results are shown in Figure 2. The associated inferential statistics essentially replicated those of the training test, and will not be repeated here. As is evident in Figure 2, in each case, substantial levels of discrimination were still possible with these simulated pigeons despite the removal or manipulation of the relevant information at test, replicating the results of the three tests with real pigeons of Watanabe et al. (1995).

For each test, discrimination increased significantly as more eigenvectors were used to make the discrimination, although, as with the training test analysis, for each test, the effect appeared to asymptote once the first 8 or so eigenvectors

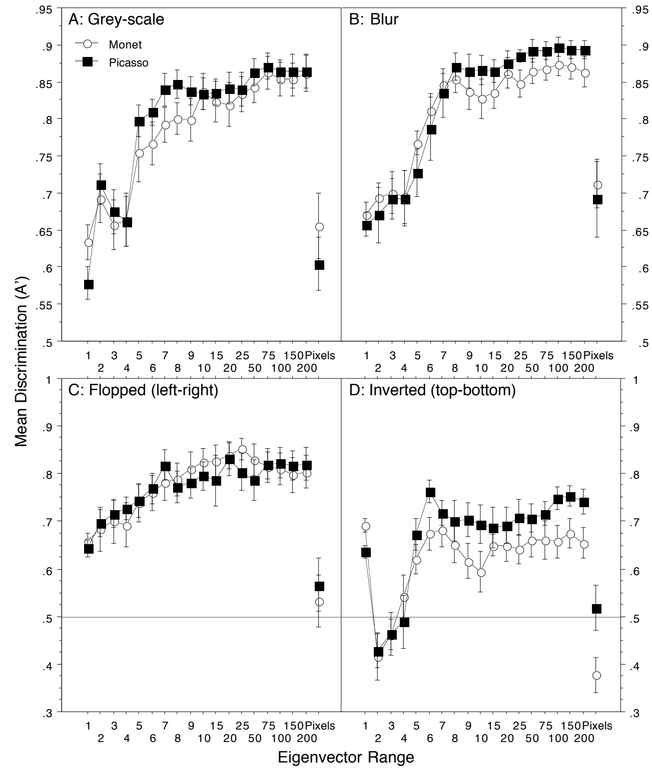

Figure 2. Mean discrimination (A') of Monet from Picasso training stimuli in the tests of grey-scale (Figure 2A), blur (Figure 2B), and flopped (left-right: Figure 2C) and inverted (top-bottom: Figure 2D) orientation as a function of eigenvector range for the simulated pigeons that were trained with colour stimuli from either Monet (open circles) or Picasso (filled squares) as the positive category. Error-bars are within-cell, standard errors.

were included. Positive training category did not significantly influence the results, nor did it interact significantly with eigenvector range for any of the tests.

As with the original training test, eigendecomposition was not strictly necessary to achieve above-chance discrimination, at least for the colour and blur tests. However, for all tests, substantially better performance was obtained by encoding the stimuli in terms of the eigenvectors as compared with applying the perceptron classifier directly to the pixel-maps. Discussion of some of the differences in performance among the grey-scale, blur, flopped and inverted conditions is deferred to the General Discussion.

Generalisation. The fourth test in Experiment 1 of Watanabe et al. (1995) was a *generalisation* test. It is the results of this test that have been the source of most of the interest in the Watanabe et al. (1995) paper because the results implied that not only could the pigeons learn the style of a particular artist, they could also generalise that learning to other artists who painted in the same general artistic styles, such as Impressionist and Cubist. Following training with the set A stimuli, the pigeons of Watanabe et al. (1995) were tested with three novel paintings each by Monet and

Table 2
Paintings used as the generalisation set in the simulation of Watanabe et al. (1995)

Paintings by Artists	Year
New Monet	
La Grenouillere	1869
Lady with a parasole	1869
Water lily	1870
New Picasso	
Dance	1925
Woman looking at the glass	1937
Still life with an ox head	1942
Cezanne	
Sitting man	1898
Still life with onions	1895
Big water bathing	1898
Braque	
Female musician	1917
Still life with "le Jour"	1929
An easel and a woman	1936
Delacroix	
Still life with a lobster	1827
July 28th	1830
Atelier	1830

Picasso, and three paintings each by Cezanne, Braque, and Delacroix. Cezanne was chosen to represent an artist who painted in the Impressionist style of Monet. Braque was similarly chosen as an artist who painted in the Cubist style of Picasso. And Delacroix was chosen as a representative of artists who painted in neither the Impressionist nor Cubist styles. The names of these 15 paintings may be found in Table 2. Watanabe et al. (1995) reported that pigeons trained with Monet as the positive category pecked preferentially at the novel Monet paintings and the paintings of Cezanne, and least to the Braque and novel Picasso paintings. Those pigeons trained with Picasso as the positive category, showed the reverse pattern. In both cases, paintings by Delacroix received an intermediate number of pecks (i.e., neither preferred nor rejected).

To investigate whether our simulated pigeons could evince a similar pattern of generalisation, the 15 new paintings were scanned and converted to pixel-maps as outlined earlier for the training stimuli. Each simulated pigeon was presented with a randomly-sampled block of each of the new test images, and was asked to classify it based on its memory for the training stimuli. Thus, each new sample of each of the test stimuli was projected into the space defined by the eigenvectors of the 200 original training experiences for a given simulated pigeon, and the projection weights were then passed to the classifier to produce a response to the image of either "Monet" or "Picasso". For simulated pigeons trained with Monet as the positive category, positive responses signify "Monet", whereas they signify "Picasso" for those trained with Picasso

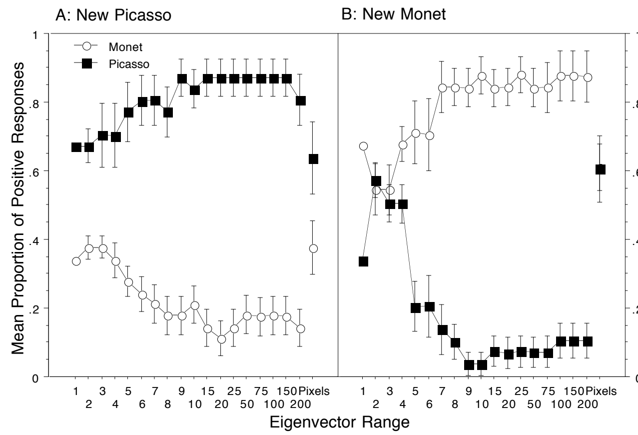

Figure 3. Mean proportion of new Picasso (Figure 3A) or new Monet (Figure 3B) images labelled either "Monet" or "Picasso" in the generalisation tests as a function of eigenvector range for the simulated pigeons that were trained with coloured stimuli from either Monet (open circles) or Picasso (filled squares) as the positive category. Error-bars are within-cell standard errors.

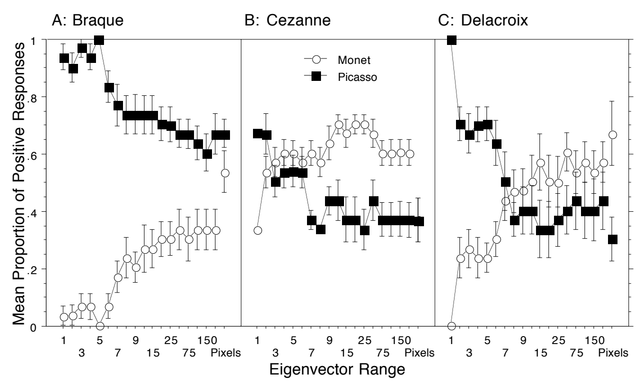

Figure 4. Mean proportion of new Braque (Figure 4A), new Cezanne (Figure 4B), and new Delacroix (Figure 4C) images labelled either "Monet" or "Picasso" in the generalisation tests as a function of eigenvector range for the simulated pigeons that were trained with coloured stimuli from either Monet (open circles) or Picasso (filled squares) as the positive category. Error-bars are within-cell standard errors.

as the positive category. The results are shown in Figures 3 and 4.

Figure 3 depicts the results for the generalisation to new Picasso and Monet images. For simulated pigeons trained with Picasso as the positive category, positive responses increased as a function of eigenvector range for new Picasso images (Figure 3A), and decreased as a function of eigenvector range for new Monet images (Figure 3B); simulated pigeons trained with Monet as the positive category evinced the opposite pattern of responding. Although the new Picasso images could apparently be discriminated with no more than the first eigenvector (and to a slight degree even by the perceptron classifier applied directly to the pixel-maps), the new Monet images required a minimum of the first 5 eigenvectors. However,

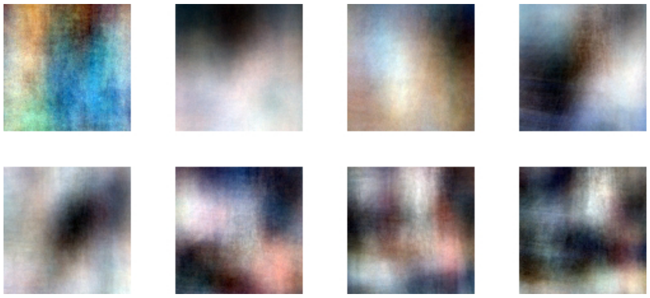

Figure 5. Images of the first eight (from top-left to bottom-right) eigenvectors of the complete set of 2000 training stimuli used in the simulations. Animations or "eigenmovies" of these eigenvectors may be viewed on the web at <http://people.uleth.ca/~vokey/movies>.

it also is clear that at the point at which performance in the various tests with the training stimuli appeared to asymptote (i.e., 8 or so eigenvectors), both the new Picasso and the new Monet images were predominantly labelled correctly. Thus, as with the pigeons of Watanabe et al. (1995), the simulated pigeons were able to transfer the training discrimination to the new Picasso and Monet images.

Figure 4 depicts the results for the generalisation to the images from the Cubist painter Braque, the Impressionist painter Cezanne, and the non-Cubist, non-Impressionist painter Delacroix. Responding to the Braque images (see Figure 4A), simulated pigeons trained with Picasso as the positive category clearly responded positively, whereas those trained with Monet as the positive category, clearly did not, although this pattern decreased as a function of eigenvector range. For the Cezanne images (see Figure 4B), the pattern was reversed; these images were preferred by the simulated pigeons trained with Monet as the positive category, although this result was predominantly the case for larger ranges of the eigenvectors; again, though, if the first 8 or so eigenvectors were used, there is a clear bias to label the Cezanne stimuli as Monet. Hence, the Braque images were responded to as if they were Picasso images, especially if only the first few eigenvectors were used, and the Cezanne images were responded to as if they were Monet images, especially as more and more eigenvectors were used. The Delacroix images (see Figure 4C), in contrast, were responded to as Picasso images for the first few eigenvectors, but were favoured as neither artist type after the first 8 or so eigenvectors were used, replicating the results of Watanabe et al. (1995) with these images.

General Discussion

Using a simple PCA neural network to encode sub-samples of the pixel-maps of images coupled with an even simpler perceptron classifier to discriminate them, we were able to simulate the results of Experiment 1 of Watanabe et al. (1995) on the same set A images they used. The match with the pattern of their results was most acute if the classifier was restricted to encodings of test items based on the first 8 or so eigenvectors or macro-features. Given the similarity in the pattern of results, it seems reasonable to suggest that the pigeons in Watanabe et al. (1995) were also responding to visual aspects of the stimuli either highly correlated with or identical to these first 8 components.

Accordingly, these visual macro-features are deserving of closer scrutiny. Depicted in Figure 5 are the first eight eigenvectors or "macro-features" of the complete set of 2000 sub-sampled images used as the training stimuli.

The issues associated with the analysis of such macro-features appear to be three-fold. First, what aspects of the variation in the visual world do each of these macro-features appear to capture? Second, how stable are these macro-features as dimensions of variation? Are they likely to vary widely from stimulus-set to stimulus-set or sample to sample? And third, what are the advantages (and disadvantages) in taking this approach to categorical discrimination?

Eigenvector interpretation (including the distinction between early and late vectors)—the four principle reasons to look at pixel-maps

Lubow

Stability

Relation to Huber et al.

How is all that Style?

Scale (selling PCA of images).

What are the advantages of the approach of PCA of pixel-maps?

1. Avoids the potential problem of arbitrary or a priori codings (lower than)

2. A more complete description of the potential sources of variance in your stimuli than the usual strategy of picking one or few theoretically motivated descriptions because they can be at most complete and most likely incomplete. (Hence, control groups).

3. It allows for description of the stimuli along independent dimensions for which we may not have good labels.

4. Generality of mechanism (can be applied to pictures of all types)

Advantages 1 and 2 relate to the distinction between the *analytic* and *synthetic* approaches to investigating visual category learning (von Fersen & Lea, 1990).

References

Abdi, H., Valentin, D., & Edelman, B. (1999). *Neural networks*. Thousand Oaks, CA: Sage Publications.

Abdi, H., Valentin, D., Edelman, B., & O'Toole, A. (1995). More about the difference between men and women: Evidence from linear neural netweorks and the principal component approach. *Perception*, *24*, 539–562.

Anderson, J. A. (1995). *An introduction to neural networks*. Cambridge, Massachusetts: The MIT Press.

Aust, U., & Huber, L. (2001). The role of item-and category-specific information in the discrimination of people versus nonpeople images by pigeons. *Animal Learning & Behavior*, *29*(2), 107–119.

Aust, U., & Huber, L. (2002). Target-defining features in a "people-present/people-absent" discrimination task by pigeons. *Animal Learning & Behavior*, *30*(2), 165–176.

Aust, U., & Huber, L. (2003). Elemental versus configural perception in a people-present/people-absent discrimination task by pigeons. *Learning & Behavior*, *31*(3), 213–224.

Bhatt, R. S., Wasserman, E. A., Reynolds, Jr., W. F., & Knauss, K. S. (1988). Conceptual behavior in pigeons: Categorization of both familiar and novel examples from four classes of natural and artificial stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, *14*, 219–234.

Ceralla, J. (1979). Visual classes and natural categories in the pigeon. *Journal of Experimental Psychology: Human Perception and Performance*, *5*, 68–77.

Ceralla, J. (1980). The pigeon's analysis of pictures. *Pattern Recognition*, *12*, 1–6.

Gibson, B., Wasserman, E., Gosselin, F., & Schyns, P. (2005). Applying bubbles to localize features that control pigeons' visual discrimination behavior. *Journal of Experimental Psychology: Animal Behavior Processes*, *31*, 376–382.

Gosselin, F., & Schyns, P. (2001). Bubbles: A technique to reveal the use of information in recognition tasks. *Vision Research*, *41*(17), 2261–2271.

Herrnstein, R. J. (1979). Acquisition, generalization, and discrimination reversal of a natural concept. *Journal of Experimental Psychology: Animal Behavior Processes*, *5*, 116–129.

Herrnstein, R. J. (1990). Levels of stimulus control: A functional approach. *Cognition*, *37*, 133–166.

Herrnstein, R. J., & de Villiers, P. A. (1980). Fish as a natural category for people and pigeons. *The Psychology of Learning and Motivation.*, *14*, 59–95.

Herrnstein, R. J., & Loveland, D. H. (1964). Complex visual concept in the pigeon. *Science*, *146*, 549–551.

Herrnstein, R. J., Loveland, D. H., & Cable, C. (1976). Natural concepts in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *2*, 285–302.

Huber, L. (2001). *Avian visual cognition.* Comparative Cognition Press. Retrieved February 16, 2004, from http://www.pigeon.psy.tufts.edu/avc/huber. Available from `http://www.pigeon.psy.tufts.edu/avc/huber`

Huber, L., Troje, N. K., Loidolt, M., Aust, U., & Grass, D. (2000). Natural categorization through multiple feature learning in pigeons. *The Quarterly Journal of Experimental Psychology*, *53B*(4), 341–357.

Lea, S. E. G., & Harrison, S. N. (1978). Discrimination of polymorphous stimulus sets by pigeons. *Quarterly Journal of Experimental Psychology*, *30*, 521–537.

Lubow, R. E. (1974). High-order concept formation in the pigeon. *Journal of the Experimental Analysis of Behavior*, *21*, 475–483.

Malott, R. W., & Siddall, J. W. (1972). Acquisition of the people concept in pigeons. *Psychological Reports*, *31*, 3–13.

Monen, J., Brenner, E., & Reynarts, J. (1998). What does a pigeon see in a Picasso? *Journal of the Experimental Analysis of Behavior.*, *69*, 223–226.

Morgan, M. J., Fitch, M. D., Holman, J. G., & Lea, S. E. (1976). Pigeons learn the concept of an "A". *Perception*, *5*, 57–66.

O'Toole, A. J., Abdi, H., Deffenbacher, K. A., & Valentin, D. (1993). A low-dimensional representation of faces in the higher dimensions of the space. *Journal of the Optical Society of America A*, *10*, 405–411.

O'Toole, A. J., Deffenbacher, K. A., Valentin, D., & Abdi, H. (1994). Structural aspects of face recognition and the other race effect. *Memory and Cognition*, *22*, 208–224.

Poole, J., & Lander, D. G. (1971). The pigeon's concept of pigeon. *Psychonomic Science*, *25*, 157–158.

Roberts, W. A., & Mazmanian, D. S. (1988). Concept learning at different levels of abstraction by pigeons, monkeys, and people. *Journal of Experimental Psychology: Animal Behavior Processes*, *14*, 247–260.

Ryle, G. (1951). Thinking and language. *Proceedings of the Aristotelian Society, Supplement, 25*, 65–82.

Siegel, R. K., & Honig, W. K. (1970). Pigeon concept formation: Successive and simultaneous acquisition. *Journal of the Experimental Analysis of Behavior*, *13*, 385–390.

Troje, N. F., Huber, L., Loidolt, M., Aust, U., & Fieder, M. (1999). Categorical learning in pigeons: The role of texture and shape in complex static stimuli. *Vision Research*, *39*, 353–366. Available from `http://www.sciencedirect.com/science?`$_ob = MImg_imagekey = B6T0W - 3WP2K56 - F - 1_cdi = 4873_orig = browse_coverDate = 01dGLbVlz - zSkzV_acct = C000051241_version = 1_userid = 1067350md5 = 268e412106b0e82517824a38d5db6233ie = f.pdf$

Turk, M., & Pentland, A. (1991). Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, *3*, 71–86.

Valentin, D., Abdi, H., & O'Toole, A. J. (1994). Categorization and identification of human face images by by neural networks: A review of the linear autoassociative and principal component approaches. *Journal of Biological Systems*, *2*, 413–429.

Valentin, D., Abdi, H., O'Toole, A. J., & Cottrell, G. W. (1994). Connectionist models of face processing: A survey. *Pattern Recognition*, *27*, 1209–1230.

Vaughan, W., & Herrnstein, R. J. (1987). Choosing among natural stimuli. *Journal of the Experimental Analysis of Behavior.*, *47*, 5–16.

von Fersen, L., & Lea, S. E. G. (1990). Category discrimination by pigeons using five polymorphous features. *Journal of the Experimental Analysis of Behavior*, *54*, 69–84.

Wasserman, E. A., Kiedinger, R. E., & Bhatt, R. S. (1988). Conceptual behavior in pigeons: Categories, subcategories, and pseudocategories. *Journal of Experimental Psychology: Animal Behavior Processes*, *14*, 235–246.

Watanabe, S. (2001). Van Gogh, Chagall and pigeons: Picture discrimination in pigeons and humans. *Animal Cognition*, *4*, 147–151.

Watanabe, S., Sakamoto, J., & Wakita, M. (1995). Pigeon's discrimination of paintings by Monet and Picasso. *Journal of the Experimental Analysis of Behavior.*, *63*, 165–174.

Wittgenstein, L. (1968). *Philosophical investigations* (3rd ed.). Oxford: Blackwell.