The University of Lethbridge

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

John R. Vokey
Department of Psychology and Neuroscience
University of Lethbridge

Jason M. Tangen
Department of Psychology
McMaster University

ABSTRACT

Judgments 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. These results suggest that the bases for such judgments of artistic style may be simpler than normally thought. 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 not only for recognising previously-experienced paintings, but for the successful discrimination of novel paintings into various style categories. A summary of simulations of the performance of Watanabe et al.’s pigeons using precisely the same stimuli and tasks is provided. 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.

INTRODUCTION

Judgments of style in art, music, architecture, and literature are ubiquitous, as are similar if more mundane judgments about everything from music videos, television commercials, and sports’ performances, to automobile driving, academic lectures, petty argumentation, and handwriting. To say that some event or performance is in the *style* of some other event or performance or *category* of events or performances is to make a *polymorphous* (Ryle, 1951) categorical distinction, typically characterised as one held together by a form of *family-resemblance* (Wittgenstein, 1968). That is, however effortless and intuitive the judgment of style is to make (for one so trained or inclined), the basis for the judgment is often tacit, and typically resistant to an analytic reduction to a small set of simple, discriminative “features” (that are apparent even to the uninitiated once drawn to their attention). This resistance to analysis is often taken as evidence for an extreme intellectual “depth” and sophistication of the concept underlying the categorical distinction–often with the accompanying implied extension of these same qualities to the expert making the style judgment. But it is possible that however truly deep and sophisticated the *intellectual* concept underlying the categorical distinction is, the immediate bases for the judgment may be more superficial.

The appeal to human vanity of these claims for the intellectual depth apparently required to make such judgments of style notwithstanding, regardless of what we may commonly believe about the intellectual sophistication required for our sensitivity to the “deep” structure of many of these domains, we explore here the possibility that it may often be much simpler. Our focus is that organisms no more complicated than pigeons can apparently acquire a similar sensitivity to the structure of these domains, and, thus, appear capable of rendering similar “style” judgments. Since the seminal research of Herrnstein and Loveland (1964) showing that pigeons could be trained to discriminate and to generalise to new stimuli polymorphous categorical distinctions such as person (or parts thereof) vs. non-person from photographs of natural scenes, we have seen 35 years of demonstrations of relatively–and sometimes, remarkably–*sophisticated* categorical judgments by pigeons that appear to track closely human categorical distinctions or *concepts*. Although few would argue that pigeons have the *intellectual* capacity to comprehend in any sense, say, the philosophical and historical bases for the artistic movements of cubism and impressionism (see Monen, et al., 1998), they do nonetheless appear capable of learning to discriminate visually the paintings of cubists from those of impressionists from training experience with just a few exemplars of each artistic school.

Watanabe, et al. (1995) found that pigeons could learn to discriminate colour slides of paintings by Picasso from those of Monet, and then to generalise this discrimination not only to slides of previously unseen paintings by Picasso and Monet, but to paintings by other cubist (e.g., Braque) and impressionist (e.g., Cezanne) painters. The Watanabe, et al. (1995) results with slides of paintings from Picasso and Monet 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 from photographs of natural scenes in Herrnstein et al. (1976)–it is clearly the individual stimuli *as a whole*–that is, their *style*–and not specific aspects of their content that are to be discriminated into discrete categories: Paintings by Picasso are “Picassos” by virtue of the fact they were painted by a particular artist who painted in a particular style. In contrast, in the experiments of, say, Herrnstein et al. (1976), the pigeons were required to discriminate on the basis of the “objects” depicted: the overall style of the slide or photograph (e.g., background, other objects depicted, colouring, shading, etc.) in these cases was nominally irrelevant. Indeed, it is precisely this nominal irrelevance (and wide variability in it over slides) that led Herrnstein et al. (1976), among others, to conclude that it was indeed the concepts of “person”, “tree”, “water”, and so on that the pigeons were discriminating in these tasks.

In what follows, we attempted to simulate the results of Watanabe et al. (1995) using a simple neural network model applied to pixel-maps of the paintings used by Watanabe et al. (1995) as a way of investigating these ideas. Watanabe et al. (1995) trained each of their pigeons to discriminate 10 paintings by Monet from 10 from Picasso to a criterion of a response ratio of 90% over two successive trials with the 20 training stimuli. The pigeons required between 6 and 22 trials to reach criterion. Training was followed by 3 transfer tests to various manipulations of the training stimuli, and one transfer test to new stimuli. To test whether possible colours differences between the paintings of Monet and Picasso were the source of the discrimination, Watanabe et al. (1995) presented the training stimuli with colour removed (i.e., in grey-scale), and found that the discrimination was maintained. As paintings by Picasso tend to have sharper edges than those of Monet, the pigeons were tested with the training stimuli blurred, and again Watanabe et al. (1995) found that the discrimination was maintained. The pigeons also were tested with some of the training stimuli mirror-reversed or turned upside-down. Again, the discrimination was maintained, although at a lower level, particularly for the upside-down stimuli. Finally, the pigeons were tested with new paintings by Monet and Picasso, as well as paintings by other cubists (e.g., Braque) and other impressionists (e.g., Cezanne), and some from neither category (e.g., Delacroix). The pigeons responded as trained to the new Monet and Picasso paintings, and tended to respond to those from the same artistic schools as the training stimuli in the appropriate way (e.g., if trained with Monet as S+, paintings of other impressionists were responded to more than those from other cubists).

FIGURE 1

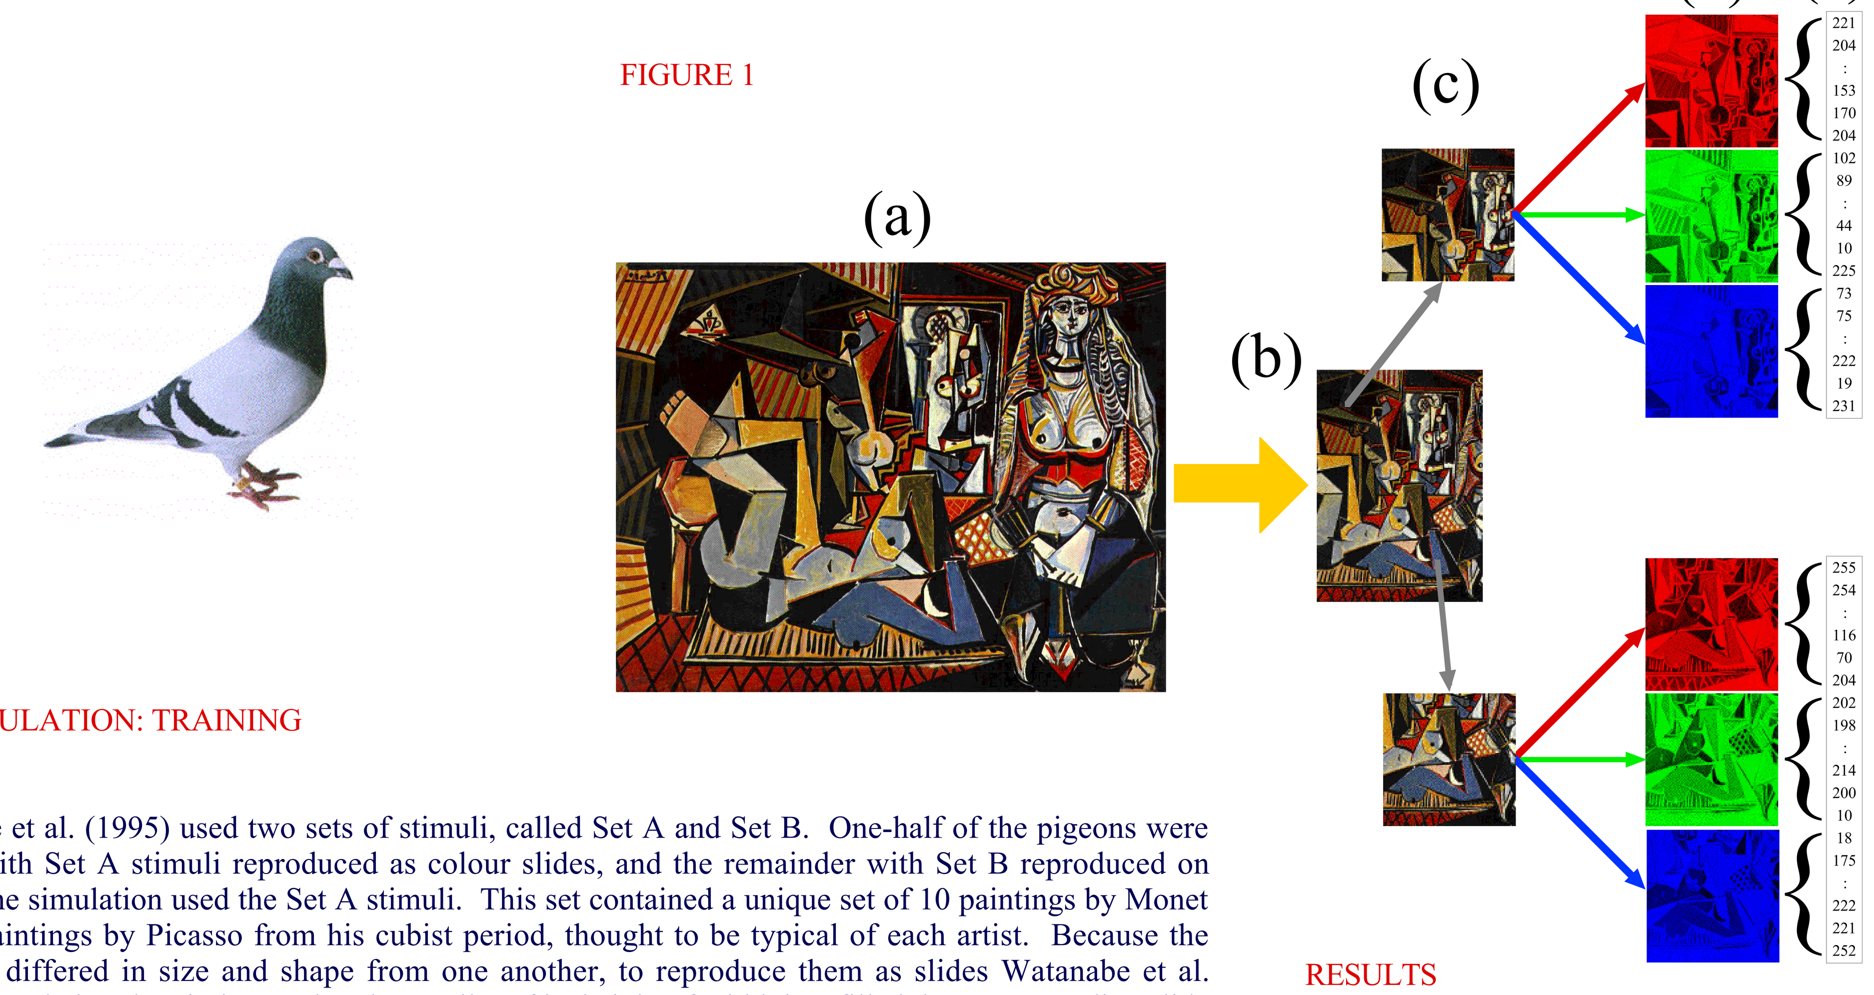

THE SIMULATION: TRAINING

STIMULI

Watanabe et al. (1995) used two sets of stimuli, called Set A and Set B. One-half of the pigeons were trained with Set A stimuli reproduced as colour slides, and the remainder with Set B reproduced on video. The simulation used the Set A stimuli. This set contained a unique set of 10 paintings by Monet and 10 paintings by Picasso from his cubist period, thought to be typical of each artist. Because the paintings differed in size and shape from one another, to reproduce them as slides Watanabe et al. (1995) re-scaled each painting so that the smaller of its height of width just filled the corresponding slide dimension, and then cropped it centred on the other dimension (S. Watanabe, personal communication, 1999). The same approach was used in the simulation. Photo slide dimensions are in the ratio of 5:7. We designated the longer slide dimension to correspond to the height. The paintings in Set A were scanned from various art books and reproduced as colour, computer graphic images. An example of one of the paintings by Picasso is shown in Figure 1a, which also shows how the images were re-scaled and cropped to have a height of 175 pixels and a width of 125 pixels (Figure 1b).

SIMULATED SUBJECTS

The variability both between and within pigeons was simulated by having each trial a given simulated pigeon had with a given stimulus consist of a random block of the pixels from the cropped stimulus. On each trial, a random square of 100 x 100 pixels was extracted from the 125 x 175 pixel stimulus, and stored, coded for colour (as RGB sub-images), as a vector 100 x 100 x 3 = 30,000 pixels in length, representing the simulated pigeon’s visual experience of the stimulus for that trial. Figures 1c–1e depict the extraction of two such samples. The idea here was to capture the notion that no stimulus is ever encountered in exactly the same way twice. Thus, as there 26 x 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. The sampling variability also provided 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 [roughly midway between the 6 and 22 trials of Watanabe et al.’s (1995) pigeons] such trials with each of the 20 stimuli.

PROCEDURE

For each simulated pigeon, the 200 samples (10 trials x 10 stimuli x 2 artists) were learned as a linear autoassociative memory using the Widrow-Hoff learning algorithm. This memory–the 30,000 pixels x 30,000 pixels weight matrix, **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 x 200 matrix of training stimuli, and represented in terms of the eigenvectors, **U**, of the pixels x pixels cross-products matrix (see Abdi et al., 1999): $\mathbf{W} = \mathbf{U}\mathbf{U}'$. Retrieval of an item (a vector) from this memory, $\hat{\mathbf{X}}_i$, is computed as: $\hat{\mathbf{X}}_i = \mathbf{W}\mathbf{X}_i = \mathbf{U}_{1:m}(\mathbf{U}'_{1:m}\mathbf{X}_i)$, 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 many eigenvectors as there are items in the training set. The expression in parentheses can be interpreted as the *projection*, $\mathbf{P}_{i|1:m}$, of the item into the space defined by the eigenvectors, $\mathbf{P}_{i|1:m} = \mathbf{U}'_{1:m}\mathbf{X}_i$, where the values of $\mathbf{P}_{i|1:m}$ are the *weights* on each eigenvector used to reconstruct the item from the linear combination of the eigenvectors: $\hat{\mathbf{X}}_i = \mathbf{U}_{1:m}\mathbf{P}_{i|1:m}$. 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 “macro-features” of the items (see, e.g., Abdi et al., 1999).

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”–a simple linear heteroassociator with Widrow-Hoff error-correction, composed of a multiple-unit input layer and one binary output unit–in essence, in statistical terms, a simple linear discriminant function analysis of the inputs to predict the binary classification of the items (see, e.g., Abdi et al., 1999, 1995). 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.

RESULTS

To investigate how well what the simulated pigeons learned from their 200 training experiences generalised to new experiences with the same training images, 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. 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.

For each simulated pigeon, the 20 classification responses for different ranges of eigenvectors (i.e., representing the use of more and more of the “macro-features” 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; an A’ value of 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 “macro-feature” or eigenvector, the first 2, 3, ..., 10, 15, 20, 25, 50, 75, 100, 150, and all 200 eigenvectors. The results for training are shown in Figure 2. Clearly, substantial levels of discrimination are possible with these simulated pigeons. As can be seen in Figure 2, discrimination increased as more and more eigenvectors were used to make the discrimination, although the effect appeared to asymptote once the first 8 or so eigenvectors were included.

FIGURE 2

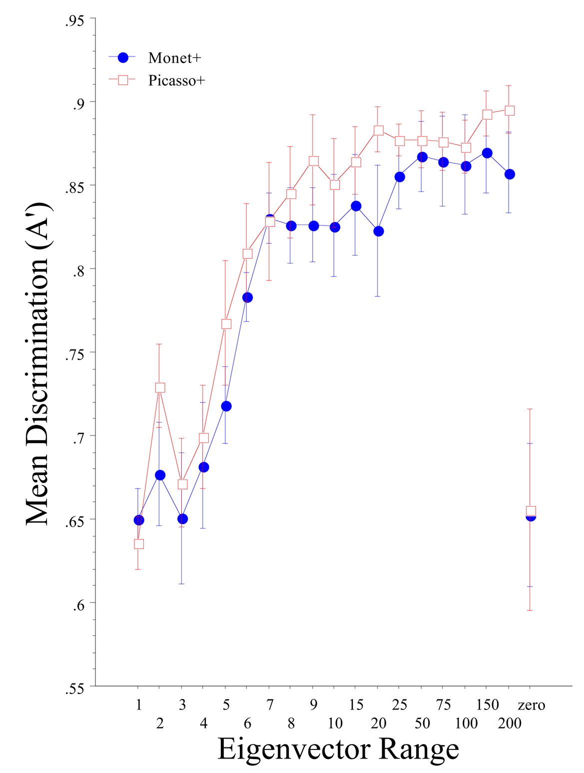

FIGURE 3

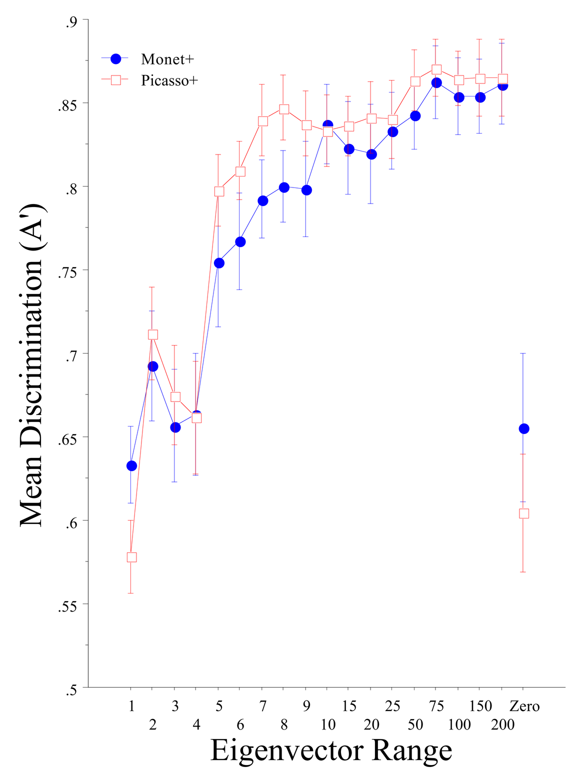

FIGURE 4

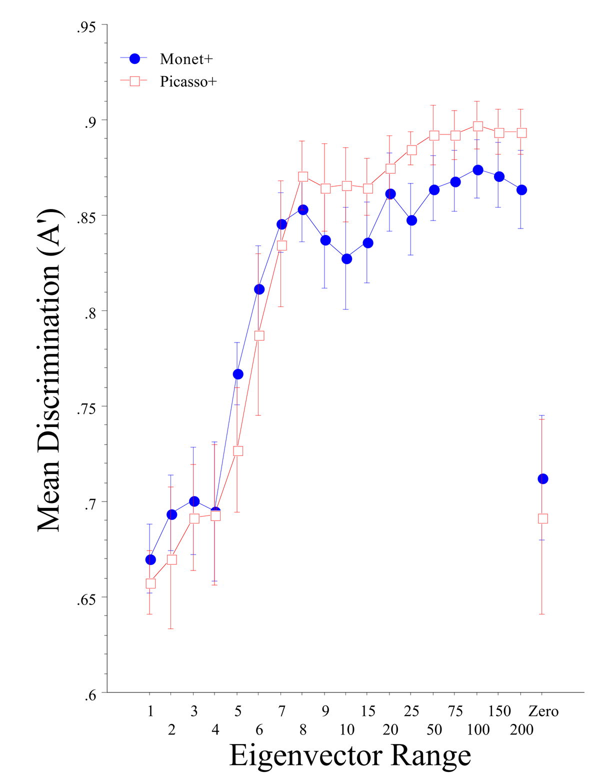

It is possible that the high levels of successful discrimination achieved here reflect *not* the eigen-decomposition of the items, but rather the power of the discriminant function analysis inherent in the classifier. That is, is there any advantage to representing the items in terms of eigenvectors or “macro-features”? 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 A’. These values are shown as “Zero” eigenvectors in each of the Figures. As can be seen, although applying the perceptron directly to the pixel-maps does result in better than chance performance, substantially better performance is possible by coding the items in terms of the macro-features.

As shown in Figures 3, 4, and 5, similar results were found for the grey-scale, blur, and mirror-reversal and upside-down tests, respectively, replicating each of these results with Watanabe, et al.’s (1995) pigeons.

Figure 6 depicts the rate of responding “Picasso” and “Monet”, respectively, for the generalisation test. Again, the results replicate those of Watanabe, et al. (1995). Simulated pigeons trained with Monet as S+, for example, responded preferentially to new paintings by Monet and Cezanne, and rejected new paintings by Picasso and Braque, with Delacroix somewhere in the middle. As with the other simulations, the effect is found principally with the first 8 or so eigenvectors or macro-features.

Figure 7 depicts the first 8 eigenvectors for one of the simulated pigeons that appear to be the bases for the discriminations. The corresponding eigenvectors for the other simulated pigeons are very similar. That is, it appears that these eigenvectors or macro-features are all that a pigeon need “see” in a Picasso to produce successful judgments of an artist’s style.

FIGURE 5

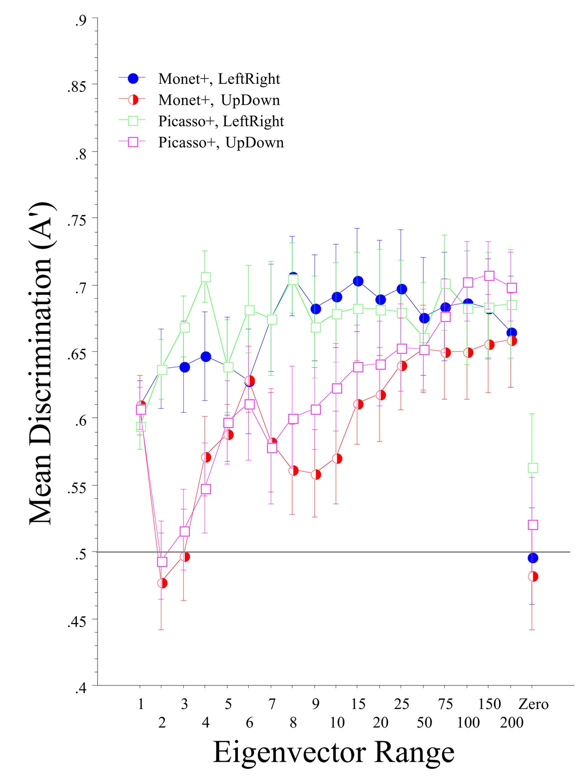

FIGURE 6

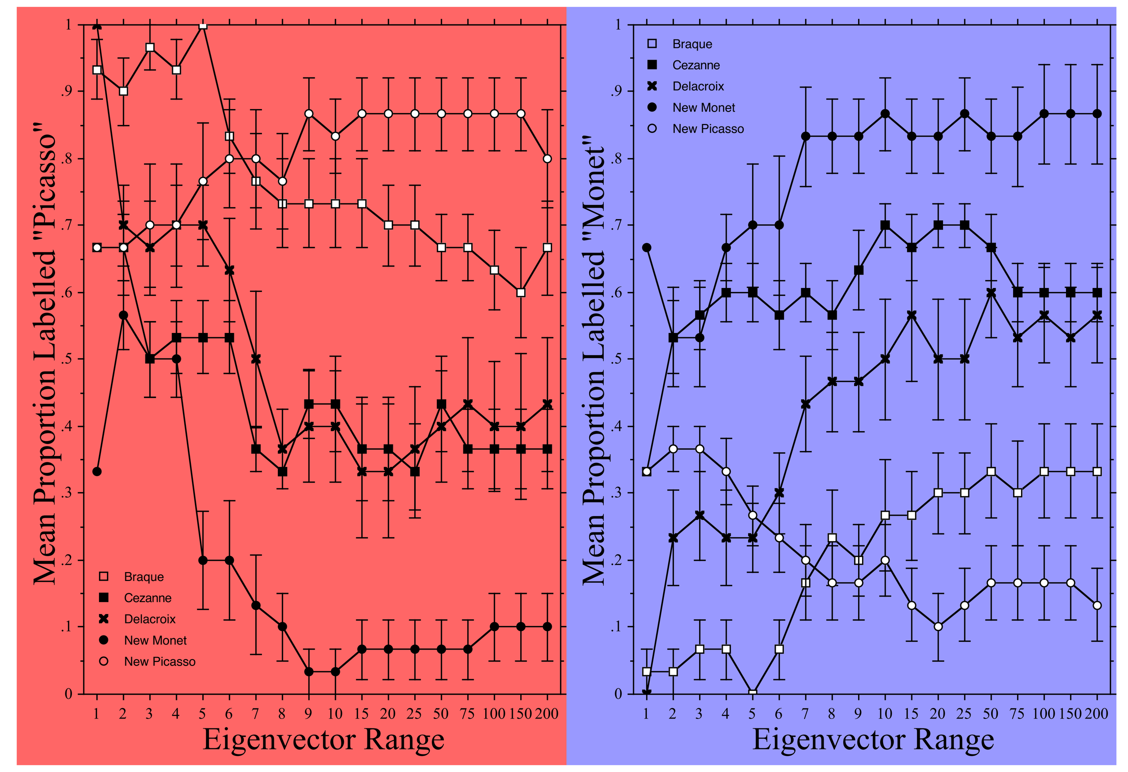

FIGURE 7

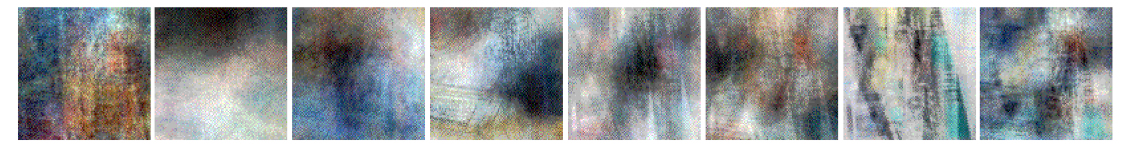

REFERENCES

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 mean and women: Evidence from linear neural networks and the principal component approach. *Perception, 24,* 539–562.

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.

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.

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

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.