Greater cross-viewer similarity of semantic associations for representational than for abstract artworks

Astrid Schepman*

Department of Psychology, University of Chester, Chester, United Kingdom

Paul Rodway*

Department of Psychology, University of Chester, Chester, United Kingdom

Sarah J. Pullen

Department of Psychology, University of Chester, Chester, United Kingdom

It has been shown previously that liking and valence of associations in response to artworks show greater convergence across viewers for representational than for abstract artwork. The current research explored whether the same applies to the semantic content of the associations. We used data gained with an adapted unique corporate association valence measure, which invited 24 participants to give short verbal responses to 11 abstract and 11 representational artworks. We paired the responses randomly to responses given to the same artwork and computed semantic similarity scores using UMBC Ebiquity software. This showed significantly greater semantic similarity scores for representational than for abstract art. A control analysis in which responses were randomly paired with responses from the same category (abstract, representational) showed no significant results, ruling out a baseline effect. For both abstract and representational artworks, randomly paired responses resembled each other less than responses from the same artworks, but the effect was much larger for representational artworks. Our work shows that individuals share semantic associations with other viewers in response to artworks to a greater extent when the artwork is representational than when it is abstract. Our novel method shows potential utility for many areas of psychology that aim to understand the semantic convergence of people's verbal responses, not least aesthetic psychology.

Introduction

Aesthetic appreciation of visual art involves multiple complex processes, including visual, cognitive, emotional, social, and semantic processes (see, e.g., Jacobsen, 2010; Leder, 2013; Leder, Belke, Oeberst, & Augustin, 2004; Lindell & Mueller, 2011; Palmer, Schloss, & Sammartino, 2013). While responses to artwork may be subjective, there are some properties of artwork that predictably influence aesthetic appreciation across individuals. The property of interest in this article is the representational content of the art. We contrast representational art, which depicts the physical visual world, usually in a nondistorted way, with abstract art, which does not contain recognizable objects but instead features shapes, patterns, forms, or color compositions. It has been found by a number of researchers that viewers prefer representational art to abstract art, and it has been proposed that this may be because they find it more difficult to find meaning in abstract than in representational art, especially if they lack art expertise (see, e.g., Gordon, 1952; Hekkert & van Wieringen, 1996; Landau, Greenberg, Solomon, Pyszczynski, & Martens, 2006; Leder, Carbon, & Ripsas, 2006; Martindale, 1984; Mastandrea, Bartoli, & Carrus, 2011; Vartanian & Goel, 2004; Winston & Cupchik, 1992).

In addition to a global preference for representational art, particularly by naïve viewers, there is evidence that viewers agree more with other viewers in their preferences for representational than for abstract images. Vessel and Rubin (2010) argued that this is because representational images are likely to generate associations that are shared by other viewers, which also have similar emotional connotations (e.g., pleasant, unpleasant), while responses to abstract images may be more idiosyncratic. Schepman, Rodway, Pullen, and Kirkham (2015) provided support for Vessel and Rubin's (2010) claim that the shared liking was due to a greater level of shared valence of semantic

associations for representational art by asking participants to generate semantic associations and to provide valence ratings for these associations. Schepman et al. (2015) found, using this method, that representational artworks generated semantic associations that shared valence (positive, negative) with those of other viewers to a greater extent than was the case for abstract artwork. What Schepman et al. (2015) were not able to probe directly, and what was also not the empirical focus of Vessel and Rubin's (2010) work, was the semantic content of the associations generated by viewers. For Vessel and Rubin's claim to be fully supported, the semantic associations generated by viewers should overlap in meaning to a greater extent when they relate to representational artwork than when they relate to abstract artwork. Testing this hypothesis was the aim of the current study, which follows on from Schepman et al. (2015).

Method

Data collection

We analyzed a previously unanalyzed part of the data set generated by Schepman et al. (2015, experiment 2), briefly summarized here so that the study can be understood independently of the cited source. Twenty-four adults who were not art experts provided short verbal responses to 22 artworks (11 representational, 11 abstract). We classified artworks as representational if they resembled the ordinary shapes and colors of the entities represented (without major distortions in, e.g., color or shape), while abstract artworks contained no recognizable objects but could include shapes. A full description of the artworks is provided in Schepman et al. (2015), with a list appearing in its supplementary materials (http://jov.arvojournals.org/Article.aspx?articleid=2278788). In summary, a range of artworks with a variety of styles, colors, subjects, and visual appearances was chosen. Works by nonfamous artists were used to minimize the probability that participants had seen the works before or had been exposed to others' opinions or interpretations of the works. Works were presented in a printed booklet with blocks of abstract and representational artworks in a random order; blocks were counterbalanced across participants. Participants also rated the images on rating scales (see Schepman et al., 2015), but rating data are not featured in this article, which focuses on verbal responses elicited by the task. These verbal responses were elicited in writing using an adaptation of the unique corporate association valence measure (Spears, Brown, & Dacin, 2006). The instructions (also reported in Schepman et al., 2015) were as

follows: "Please write a word or short description in the boxes below of any thoughts that the work of art brought to mind. Please try to complete a minimum of three boxes and then please circle how positive, neutral, or negative the description is." Participants could complete a maximum of five response boxes. The circled ratings of the descriptions have been reported in Schepman et al. (2015) as measures of the valence of the associations and are not featured here. Instead, we concentrate on a semantic similarity analysis of the verbal responses. Participants generated responses consisting of an average of 6.61 words per representational artwork and 5.33 words per abstract artwork. We entered these responses for further semantic similarity analysis.

Analysis method

Building on Vessel and Rubin (2010) and Schepman et al. (2015), our hypothesis was that verbal responses to representational artworks would show greater semantic similarity across viewers than verbal responses to abstract artworks. To operationalize the analysis, we identified semantic similarity analysis software that could accommodate the types of responses that had been elicited and that could compute a numeric semantic similarity score for pairs of these responses for further statistical analysis. Based on our constraints, we chose UMBC Ebiquity (Han, Kashyap, Finin, Mayfield, & Weese, 2013; http://swoogle.umbc.edu/SimService/index.html). This software uses a hybrid approach to computing semantic similarity, namely distributional similarity and latent semantic analysis, supplemented with a thesaurus method using WordNet (see Han et al., 2013). Of the three variants of the software available, we chose Semantic Textual Similarity (http://swoogle.umbc.edu/StsService/index.html) because it is able to cope with the full range of responses (e.g., words, short phrases, sentences). For each pair presented, this software yields a score between 0 and 1. A score of 0 means no similarity at all, or it can indicate that a word is not in its dictionary, while a score of 1 is a perfect match. To illustrate, the words *ocean* and *sea* yield a score of 1, the phrases *old acquaintances* and *absent friends* yield a score of 0.369, and the sentences "The farm was located in a mountainous region" and "He read five books in two days" yield a score of 0. Note that these examples are not from our corpus but rather have been created by us specifically to illustrate the output from the semantic similarity software. As described more fully in Han et al. (2013), the software has a multilayered set of routines to optimize the accuracy of the semantic similarity scores and performs well against other similar software.

For each artwork, the 24 participants were asked to provide a minimum of three and a maximum of five short verbal responses in the description boxes. We randomly paired these verbal responses with other verbal responses using random numbers generated by a sequence generator (www.random.org) in one of two ways—experimental and control pairings—which are discussed in turn.

For the experimental pairings, for each artwork we paired the responses given by the 24 participants in one description box randomly with one of the responses from that set. We did this by description box to avoid the possibility that the response from a participant in one description box would be paired with his or her own response in a different description box, as that could inflate the similarity scores. For the first three description boxes, which yielded full data sets (bar very rare missing data), we did not constrain for the probability that the response would be randomly matched with itself, as this probability was deemed stable across the two conditions (abstract and representational). As participants had been asked to provide three to five responses, this process was repeated for all boxes and all artworks separately. Boxes 4 and 5 (which were optional) had fewer responses per artwork, but the same process was used except in cases with very few responses, when any matches to the response itself were rerandomized and any single responses by only one person to a particular artwork were deleted from the analysis. This process yielded 1,729 pairs, of which 842 were responses to abstract artwork and 887 were responses to representational artwork.

In addition to running the within-artwork and within-description box pairings, control pairings were created for a key control analysis. This was partly because it had been observed that more words were produced in response to representational than to abstract artworks. It was felt that this may introduce inflation of similarity scores in the representational artworks. In addition, other general aspects of the text may have led to higher similarity scores for representational artwork than for abstract artwork without these necessarily being attributable to the specific artworks. Thus, pairings were created in which all the responses within a category (abstract and representational) were randomly paired with other responses from across all artworks and description boxes of that category. These pairings were not subject to any constraints. It was hypothesized that if this analysis revealed a significant difference in similarity scores between abstract and representational artworks, then any significant difference in the experimental analysis would likely be a baseline effect. On the other hand, a nonsignificant result in the control comparison could be argued to rule out this baseline effect.

Custom-written Javascript code sent all experimental and control word pairs through the UMBC Ebiquity Semantic Textual Similarity service and stored the resulting output in an Excel spreadsheet. The semantic similarity scores yielded by this process were used to test the experimental hypothesis and the control hypotheses.

Results

Sample pairings and output

To illustrate the data, we report a sample of experimental response pairs and their similarity scores. A sample abstract artwork featuring white protruding forms with black and blue line shapes on a beige–gray background (Pol Ledent, Abstract 882140; http://c300221.r21.cf1.rackcdn.com/abstract-882140-1335238270_b.jpg) gave rise to experimental response pairs including "earthy tones" paired with "puzzled" (semantic similarity score: 0), "dark" paired with "mystery" (0.13), "messy and random" paired with "complicated" (0.31), "hidden meaning" paired with "abstract" (0.15), "ship in storm" paired with "cold flames" (0.13), "nature" paired with "anger" (0), "paranoid" paired with "cotton wool" (0), and "snow" paired with "interesting colors" (0). A sample representational artwork featuring a woman standing by a wall laughing (Jean Smith, Laugher #4; http://jeansmithartist.com/wp-content/gallery/laughter-project/laughter4.jpg) gave rise to experimental response pairs including "positive and happy" paired with "fun" (0.20), "shadow" paired with "happy" (0), "I want to meet this lady; she looks fun" paired with "yellow" (0), "amusing" paired with "I would love to know why she is laughing" (0.45), "colorful" paired with "I love the contrast between the background and the woman" (0.16), "good color choice" paired with "embarrassment" (0), "funny" paired with "snapshot" (0), and "good times" paired with "happy" (0.15). As can be seen from the sample response pairs, responses were quite varied for both types of art. However, in this small illustrative sample, it seems that the semantic content of the responses to the representational artwork may overlap to a greater extent and that the responses to the abstract art may be more varied. Our statistical analyses, set out in the next sections, aimed to put this notion to the test.

Experimental pairings

Normality tests and distribution plots (see top panels of Figure 1) showed a nonnormal distribution for both
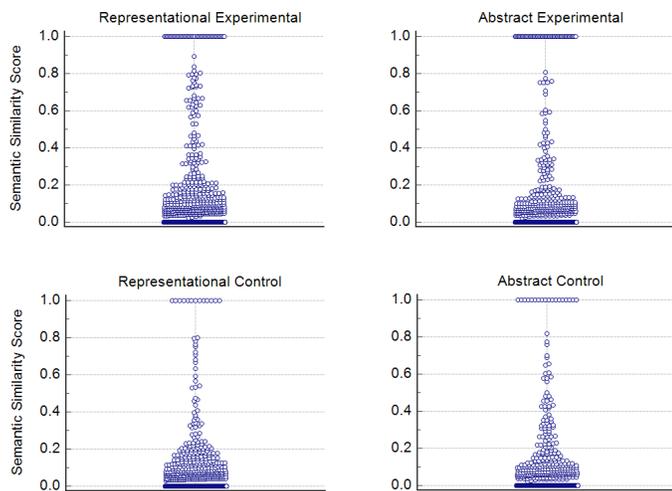
Figure 1. Dot plots of the distributions of semantic similarity scores for the representational and abstract artworks in the experimental and control pairings.

categories. Therefore, statistical analysis was carried out using the nonparametric Mann–Whitney test for two independent samples, which tested for differences in ranks. The mean semantic similarity score for abstract artworks was 0.1141 ($SD = 0.257$), while the score for representational artworks was higher at 0.1298 ($SD = 0.251$); the similarity scores differed significantly when comparing the two types of artwork ($Z = -3.622$, $p < 0.001$). The abstract set contained 504 scores of 0 (59.9%) and 53 scores of 1 (6.3%). The representational set contained 455 scores of 0 (51.3%) and 44 scores of 1 (5.0%).

Control pairings

This control analysis yielded a mean of 0.0681 ($SD = 0.159$) for representational artworks and a slightly higher mean of 0.0726 ($SD = 0.178$) for abstract artworks. The difference between conditions was not significant ($Z = -1.166$, $p = 0.244$). The abstract set contained 543 scores of 0 (64.5%) and 18 scores of 1 (2.1%), while the representational set contained 541 scores of 0 (61.0%) and 13 scores of 1 (1.5%). The distribution of scores for these two data sets can be seen in the lower panels of Figure 1.

Experimental versus random pairings

Given the patterns reported above we felt it would be useful to run a third analysis, which explored whether, for both abstract and representational artworks, the cross-viewer similarity of the responses given by participants to specific artworks significantly exceeded the similarity scores observed in the random control

pairings. The main focal points in the analysis were whether abstract artwork showed some convergence compared with a random baseline and, if so, on what order of magnitude the effect size may be different from the equivalent comparison in the representational artworks. This analysis was done using a pairwise nonparametric test, namely Wilcoxon's signed-ranks test. This showed that for representational artworks, semantic similarity of the experimental pairs exceeded that of the random control pairings significantly ($Z = -7.010$, $p < 0.0001$). Crucially, this also applied to the abstract artworks, but with a smaller effect size ($Z = -3.928$, $p < 0.001$).

Discussion

Our current work shows, for the first time, that there is a greater overlap in the semantic associations elicited by representational artworks than by abstract artworks. This finding directly supports Vessel and Rubin's (2010) associationist explanation of the greater consistency in preferences for representational versus abstract artworks. Although it is somewhat difficult to translate the software's semantic similarity value into real-world semantic overlap, the semantic similarity scores for both types of artworks in the experimental pairings were relatively low within the range of 0 to 1, which suggests that a large proportion of the responses were individual. Nevertheless, to the extent that responses are shared between viewers, the responses generated by representational artwork showed a greater similarity across viewers than those generated by abstract artwork. It could be argued that this is to be expected because representational art features obvious semantic referents in the physical entities depicted, while abstract art does not, and thus representational art may generate some description-based associations that are not available for abstract artworks. On the other hand, the quantitative data and the sample response pairs show that representational artworks generate considerably varied responses. Thus, the finding that these responses overlap is not likely to be solely due to basic object naming but rather seems more likely to be associated with higher level interpretative processes.

Our control analysis shows that our findings cannot be attributed to baseline aspects of the text. One may have expected, for example, that simply producing a higher number of words may lead to higher similarity scores, or one might expect that there may be a higher level of specificity in the responses to representational art than to abstract art, giving rise to higher similarity scores without this being connected to the specific artwork. However, the control analysis, which used a

different randomization than the experimental analysis, showed that this was not the case. In fact, numerically, the scores for abstract artworks were somewhat higher than those for representational artworks in this analysis, although not significantly so.

Our other key comparison showed that for both abstract and representational artworks, the semantic similarity of randomly paired responses is exceeded amply and significantly by those of the experimental pairings, though the effect size for this observation was much larger in representational than in abstract artworks. This suggests that, even in abstract artworks, there is some overlap between viewers' responses and that their responses are not purely idiosyncratic. The overlap is stronger in representational artworks, but, based on our data, the difference is one of degree and not of kind. This leaves interesting possibilities for future research, which could aim to examine the overlap in abstract artworks; this could serve to help understand the communication between artist and viewer of abstract entities.

Our work substantially extends Vessel and Rubin's (2010) and Schepman et al.'s (2015) empirical support for the idea that representational artworks generate internal states in viewers that resemble those of other viewers to a greater extent than abstract artworks because the entities depicted in representational art create associations that show greater semantic similarity with those of other viewers. This takes this evidence beyond that provided by Vessel and Rubin (2010), who provided evidence of similarity in preference and inferred that internal states were responsible. It also takes the evidence beyond that of Schepman et al. (2015), who found that the valence of the semantic associations overlapped across viewers to a greater extent in response to representational than to abstract art but who were not able to show actual semantic overlap.

In addition to providing evidence on this specific point, we feel that, more generally, using this method opens the door to many other interesting studies that could examine how viewers process the meaning of art and a multitude of other objects. It is particularly useful to extend the methods by which this can be studied because traditionally it is relatively difficult to study meaning empirically, particularly using quantitative statistical methods. This is especially important because meaning has been deemed a key factor in the appreciation of art (see, e.g., Martindale, 1984). While the process of generating meaning may be a crucial process in art viewers, this may be the case more strongly in expert than in naïve viewers. Thus, it would be interesting in the future to carry out the same experiment with art experts, who may show differences compared with the nonexpert viewers who took part in our study.

Conclusions

Our data show that responses to representational art show a greater semantic overlap across viewers compared with responses to abstract art. This bolsters the theoretical view that shared liking is associated with shared semantic representations of art. It also provides novel and original evidence that suggests that meaning plays an important role in the complex processes that lead to aesthetic appreciation.

Keywords: aesthetic psychology, semantic association, social psychology, art, computational linguistics

Acknowledgments

*AS and PR contributed equally to this article.
Commercial relationships: none.
Corresponding author: Astrid Schepman.
Email: a.schepman@chester.ac.uk.
Address: Department of Psychology, University of Chester, Chester, United Kingdom.

References

Gordon, D. A. (1952). Methodology in the study of art evaluation.*The Journal of Aesthetics and Art Criticism, 10*(4), 338–352.

Han, L., Kashyap, A., Finin, T., Mayfield, J., & Weese, J. (2013). *UMBC EBIQUITY-CORE: Semantic textual similarity systems*. Retrieved from http://ebiquity.umbc.edu/paper/html/id/621

Hekkert, P., & van Wieringen, P. C. W. (1996). Beauty in the eye of expert and nonexpert beholders: A study in the appraisal of art. *American Journal of Psychology, 109,* 389–407.

Jacobsen, T. (2010). Beauty and the brain: Culture, history and individual differences in aesthetic appreciation. *Journal of Anatomy, 216,* 184–191.

Landau, M. J., Greenberg, J., Solomon, S., Pyszczynski, T., & Martens, A. (2006). Windows into nothingness: Terror management, meaninglessness, and negative reactions to modern art. *Journal of Personality and Social Psychology, 90,* 879–892.

Leder, H. (2013). Next steps in neuroaesthetics: Which processes and processing stages to study? *Psychology of Aesthetics, Creativity, and the Arts, 7,* 27–37.

Leder, H., Belke, B., Oeberst, A., & Augustin, D. (2004). A model of aesthetic appreciation and aesthetic judgments. *British Journal of Psychology, 95,* 489–508.

Leder, H., Carbon, C. C., & Ripsas, A. (2006). Entitling arts: Influence of title information on understanding and appreciation of paintings. *Acta Psychologica, 121,* 176–198.

Lindell, A. K., & Mueller, J. (2011). Can science account for taste? Psychological insights into art appreciation. *Journal of Cognitive Psychology, 23,* 453–475.

Martindale, C. (1984). The pleasure of thought: A theory of cognitive hedonics. *Journal of Mind and Behavior, 5,* 49–80.

Mastandrea, S., Bartoli, G., & Carrus, G. (2011). The automatic aesthetic evaluation of different art and architectural styles. *Psychology of Aesthetics, Creativity, and the Arts, 5,* 126–134.

Palmer, S. E., Schloss, K. B., & Sammartino, J. (2013). Visual aesthetics and human preference. *Annual Review of Psychology, 64,* 77–107.

Schepman, A., Rodway, P., Pullen, S. J., & Kirkham, J. (2015). Shared liking and association valence for representational art but not abstract art. *Journal of Vision, 15*(5):11, 1–10, doi:10.1167/15.5.11. [PubMed] [Article]

Spears, N., Brown, T. J., & Dacin, P. A. (2006). Assessing the corporate brand: The unique corporate association valence (UCAV) approach. *Journal of Brand Management, 14,* 5–19.

Vartanian, O., & Goel, V. (2004). Neuroanatomical correlates of aesthetic preference for paintings. *Neuroreport, 15,* 893–897.

Vessel, E. A., & Rubin, N. (2010). Beauty and the beholder: Highly individual taste for abstract, but not real-world images. *Journal of Vision, 10*(2):18, 1–14, doi:10.1167/10.2.18. [PubMed] [Article]

Winston, A. S., & Cupchik, G. C. (1992). The evaluation of high art and popular art by naive and experienced viewers. *Visual Arts Research, 18,* 1–14.