The role of local and global symmetry in pleasure, interest, and complexity judgements of natural scenes

Claudia Damiano[1], John Wilder[2], Elizabeth Yue Zhou[2], Dirk B. Walther[2,3], & Johan Wagemans[1]

Affiliations:

(1) Laboratory of Experimental Psychology, Department of Brain and Cognition, KU Leuven, Belgium

(2) Department of Psychology, University of Toronto, Canada

(3) Samsung Artificial Intelligence Centre Toronto, Canada

Author's Note: The ideas and data appearing in this manuscript were presented at the Vision Science Society's annual conference (virtual conference; June 2020). The data, stimuli, and statistical analysis code are available on OSF at the following link: https://osf.io/zge9r/?view_only=b36958f6b67a4410818f62395511dc8b. Additionally, the code for computing local symmetry can be found here: https://github.com/mrezanejad/SalienceScoresForScene.
Correspondence concerning this article should be addressed to Claudia Damiano, Laboratory of Experimental Psychology, KU Leuven. Tiensestraat 102 - box 3711, BE-3000 Leuven, Belgium. Phone: +3216377786 E-mail: claudia.damiano@kuleuven.be.

Abstract

Symmetry generally makes stimuli less complex, and symmetric arrangements are also generally preferred to asymmetric ones. Here we investigated the roles of both local and global symmetry in subjective judgements of natural scenes. We collected ratings of complexity, aesthetic pleasure, and interest for 720 scene images and calculated average ratings for each image, as well as several measures of local and global symmetry. Global symmetry measures were calculated by creating an axis of symmetry at every column (vertical) and row (horizontal) of the grayscale image and correlating the rows or columns of pixels on either side of the symmetry axis, weighted by the proportion of pixels included in the correlation. Local symmetry measures were computed by converting each photograph into a line drawing and calculating the parallelism (ribbon symmetry) and distance (separation) between contours. To investigate the relationship between symmetry and participants' ratings, we ran a canonical correlation analysis using twelve symmetry measures as predictors of the three subjective rating measures. The analysis revealed two significant and interpretable canonical roots. In the first root, local symmetry and vertical global symmetry were negatively related to complexity, aesthetic pleasure, and interest (i.e., symmetry reduces complexity and renders the scene boring and unpleasant). Conversely, in the second root, local symmetry and horizontal global symmetry were positively related to pleasure and interest. Our work reveals the distinct roles of global and local symmetry in perceptual judgements, and lends further support to the pleasure-interest model of aesthetic liking.

Key words

Local symmetry, global symmetry, aesthetic pleasure, complexity, interest, pleasure-interest model of aesthetic liking

Introduction

From the early 20th century, Gestalt psychologists described ways in which parts of visual stimuli were grouped together by the human visual system to form coherent percepts (Koffka, 1922, 1935; Köhler, 1920; Wertheimer, 1922, 1923). These grouping principles have been studied extensively to uncover their roles in the perceptual phenomena of figure-ground segmentation and contour integration, among others (see Wagemans et al., 2012 for an extensive history and review of the Gestalt grouping principles). In the current study, we highlight one principle in particular – symmetry. In general, symmetry makes abstract shapes and patterns less complex (Bertamini et al., 2013; Day, 1968; Gartus & Leder, 2017) by repeating certain parts of the stimulus, thus reducing the amount of novel information to be processed. Notably, symmetry has also been studied in relation to the pleasantness of a stimulus, where symmetric shapes and patterns are often preferred compared to asymmetric ones (Bertamini et al., 2013; Chen et al., 2011; Day, 1968; Gartus & Leder, 2013), although art experts find symmetric patterns less beautiful (Leder et al., 2019). Here, we move beyond simplistic stimuli and examine these relationships between symmetry, complexity, and aesthetic pleasure in images of real-world natural scenes. Importantly, we also distinguish between local and global aspects of symmetry, allowing us to gain insight into the different facets of these relationships that may exist.

Global symmetry, as studied in visual perception, typically refers to mirror symmetry (i.e., symmetry with a reflection axis) across a central vertical, horizontal, or oblique axis, although vertical symmetry is the most salient one (Wagemans, 1997; Wagemans et al., 1992; Wagemans et al., 1993). Local symmetry, in contrast, refers to symmetry that occurs on a smaller and more local scale. For example, a person's body is typically (roughly) symmetric across the central vertical axis (i.e., globally symmetric), and each individual limb and digit is likewise symmetric on their own (i.e., locally symmetric). In addition to abstract patterns and shapes, as mentioned above, global mirror symmetry is also negatively related to complexity in higher-level stimuli, such as objects (Sukumar et al., 2008), scenes (Oliva et al., 2004), and even mobile app interface designs (Miniukovich & De Angeli, 2014). The study of local symmetry, on the other hand, has mostly been limited to dot patterns (Nucci & Wagemans, 2007) or simple object outlines (Panis & Wagemans, 2009), where it has also been found to generally increase processing fluency. How local symmetry relates to complexity within *complex visual scenes* is unclear. The role of local symmetry within scenes has only recently begun to be investigated, with the development of a novel method developed by Rezanejad and colleagues (2019) to obtain the shape skeleton of a line drawing of a real-world scene, which is then used to compute local symmetry at each contour. With this method, the researchers found that scene images that contained only the 50% most symmetric contours were significantly easier to categorize into their basic-level categories than scenes made up of the other 50% of contours (i.e., the least symmetric contours) (Rezanejad et al., 2019; Wilder et al., 2019). This superior performance afforded by the presence of local symmetry over asymmetry suggests that that local symmetry, like global symmetry, may also increase fluency by decreasing complexity in high-level stimuli. The current study formally tests this hypothesis by relating the presence of local symmetry to subjective ratings of perceived complexity in scenes.

Furthermore, it is well-known that complexity, processing fluency, and aesthetic pleasure are tightly linked (Van Geert & Wagemans, 2020), although the results are somewhat inconsistent. Sometimes a medium level of complexity results in the highest aesthetic rating (Berlyne, 1971; Güçlütürk et al., 2016), while other times, the most fluently processed stimuli are preferred (Reber et al., 2004). In any case, an excessive level of complexity seems to have a negative effect on aesthetic preference, which, along with the negative relationship between symmetry and complexity described above, suggests that symmetry should increase aesthetic pleasure by increasing processing fluency. Indeed, we know that people exhibit a preference for globally symmetric patterns over asymmetric ones (Bertamini et al., 2013; Chen et al., 2011; Gartus & Leder, 2013; Makin et al., 2012). However, local symmetry has not yet been explored in this regard. Additionally, while abstract patterns may allow for greater control of stimulus manipulations, they do not necessarily allow us to generalize the findings to higher-level stimuli, such as natural scenes. In terms of more ecologically-valid stimuli, the most one can say at the moment is that global mirror symmetry is associated with attractiveness and perceived health of faces (Little et al., 2011; Rhodes et al., 1998; Thornhill & Gangestad, 2006). However, the role of symmetry in aesthetic pleasure judgements of natural scenes has surprisingly only been investigated in one instance, as far as we know (Study 3 in Mayer & Landwehr, 2018).

The current study uses photographs of natural scenes to allow us to study how these relationships might act in stimuli with greater ecological validity. More specifically, we will explore these relationships in natural scenes, which is especially important given the findings on the cognitive and emotional benefits of interacting with nature (Berman et al., 2008; Bowler et al., 2010). For instance, the Attention Restoration Theory posits that natural environments promote a “soft fascination”, meaning that they are automatically and fluently processed and that they are interesting without being overwhelming (Kaplan, 2001; Schertz & Berman, 2019), unlike a bustling city street, for example, which would demand too many attentional resources. This soft fascination allows for the finite top-down attention network to restore its resources, improving cognitive abilities and decreasing stress and anxiety. The theory seems to suggest that nature is aesthetically pleasing because it is fluently processed, therefore implying that complexity is negatively related to aesthetic pleasure in natural scenes. In support of this hypothesis, Mayer and Landwehr (2018, Study 3) found that computationally derived simplicity and self-similarity positively predicted the number of views on images of landscapes, suggesting that people preferred to view scenes with low complexity. However, a positive linear relationship between complexity and aesthetic pleasure ratings in natural scenes has also been found (Biederman & Vessel, 2006; Yue et al., 2007), suggesting that complexity in natural scenes may actually be pleasurable in some cases. Thus, the link between pleasure and complexity in natural scenes is still an open question. Additionally, how a range of local and global symmetry measures may modulate the effects has yet to be examined. Note here that when we refer to aesthetic pleasure, we are referring to the joy one experiences when viewing particular scenes.

Due to these inconclusive results, and the scarcity of research on the role of symmetry in aesthetic experiences when viewing high-level stimuli, the current study uses an exploratory multivariate analysis approach – a canonical correlation – to understand the many relationships that may exist between pleasure, complexity, and interest in complex real-world scenes, and the roles that both global and local symmetry may play in those relationships. We find that different aspects of global and local symmetry play divergent roles. Vertical global symmetry reduces complexity, pleasure, and interest, while horizontal global symmetry is related to increased pleasure and interest ratings. Local symmetry is both positively and negatively related to the subjective ratings. We discuss our findings within the framework of the Pleasure-Interest model of Aesthetic liking (PIA; Graf & Landwehr, 2015), providing further evidence to the idea that both complexity and simplicity influence the aesthetic pleasure of natural scenes. Our work highlights the importance of considering both local and global aspects of symmetry when discussing symmetry's relationship with subjective complexity judgements and aesthetic experiences.

Methods

Stimuli

In order to obtain a stimulus-set which could vary on ratings of pleasure and complexity, we collected images spanning a range of categories, such as beaches, deserts, fields, caves, waterfalls, etc. through Google image searches. Key words used in the searches, along with the category label, included "beautiful", "ugly", "calm", "chaotic", "simple", "complex", "interesting", and "boring". The images depict natural outdoor landscapes, though sometimes include people or human-made objects (such as benches or fences). All images were cropped or resized to 500 (width) x 375 (height) pixels. If the quality of the image could not be retained during resizing, the image was discarded. We were left with a total of 720 images, which were used in the rating experiment.

Participants

A total of 189 people (age 19 to 70; mean age = 37.2) participated in the rating experiment on Amazon Mechanical Turk and were paid 1.50 USD for their participation. Due to a technical error, gender information was not saved for the majority of participants (N = 142). Of the other 47 participants, 55% (N = 26) were women. This study was approved by the University of Toronto Research Ethics Board (Protocol number 30999). We recruited MTurk workers with at least a 99% approval rating, had completed a minimum of 100 HITS, and were located only in the United States or Canada. To check whether participants were paying sufficient attention to the task, we calculated the variance of responses within a 15-response-long (i.e., 5 trial-long; 3 responses per trial) sliding window, as well as the mean reaction time across all trials, for each participant. Participants failed the vigilance check if their mean variance was less than 0.2. Responses were made on 5-point Likert scales, thus a low variance implied that participants pressed the same number for long stretches of time during the experiment. Participants also failed the vigilance check if their mean reaction time was less than 250ms, implying that they were rushing through the experiment and may not have been paying attention (for reference, the average of mean reaction times across participants was 1946ms). Data from four participants were excluded from further analyses due to failing the vigilance criteria.

Procedure

The experiment was run using Inquisit software (millisecond.com). There were two versions of the experiment, with differing numbers of trials. In one version, each participant (N = 142) viewed and rated a total of 244 images, chosen pseudo-randomly from the entire image set of 720 images. In another version, an additional 188 images were added to the total image set, and each participant (N = 47) now rated 310 images chosen from the new image set of 908 images. The additional images were added as a pilot for a different experiment, thus are not analyzed in this experiment. In total, we collected approximately 54 individual ratings (range = 49 to 58 ratings per image) for each of the original 720 images. The average time it took participants to complete the first version, with 244 photographs, was approximately 32 minutes, and the average time for the version with 310 photographs was approximately 40 minutes. These times included instructions and self-paced breaks.

In both versions, a trial began with a blank white screen for 500ms. Then, each image was displayed in the center of the screen on a white background for 500ms before the first rating scale would appear below the image. Participants then had to respond on a 5-point Likert scale, using their computer mouse, to the following question: “How much do you enjoy looking at this image?”. The response options were 1 = “not at all”, 2 = “barely enjoy”, 3 = “somewhat enjoy”, 4 = “enjoy”, and 5 = “enjoy very much”. Following the response, the first rating scale disappeared, the image remained on screen, and the second rating scale appeared: “How simple or complex is this image?”. The response options were 1 = “very simple”, 2 = “simple”, 3= “neutral”, 4 = “complex” and 5 = “very complex”. After the response, the scale again disappeared, and the third rating scale was displayed: “How boring or interesting is this image?”. The response options were 1 = “very boring”, 2 = “boring”, 3 = “neutral”, 4 = “interesting”, and 5 = “very interesting”. Participants had to respond to each question in turn before the next image would appear. There was no time limit for the responses. See Figure 1 for an example trial sequence.

Participants were given a break every 40-50 trials and had to press the spacebar to continue with the experiment. Following the main experiment, participants filled out demographic information such as age and gender, and responded to the following questions by typing their answer into the response box: 1) “How would you describe the relationship between pleasure (liking the image) and complexity? (e.g., Did you prefer images that were more complex? Less? In the middle?). Which factors affected this relationship?”; 2) “How would you describe the relationship between pleasure and interest? Which factors affected this relationship?”; and 3) “How would you describe the relationship between interest and complexity? Which factors affected this relationship?”. These questions were included in order to gain insight into participants’ thoughts on the relationships between pleasure, complexity, and interest. These responses are not analyzed here.

Subjective Ratings

Average ratings of complexity, pleasure, and interest were calculated for each image based on participants' ratings on the task described above. Firstly, the ratings were z-scored within each individual participant. Then, for each image, the z-scored ratings were averaged across all participants who saw that image, resulting in one rating per image for each of the three subjective measures (complexity, aesthetic pleasure, and interest).

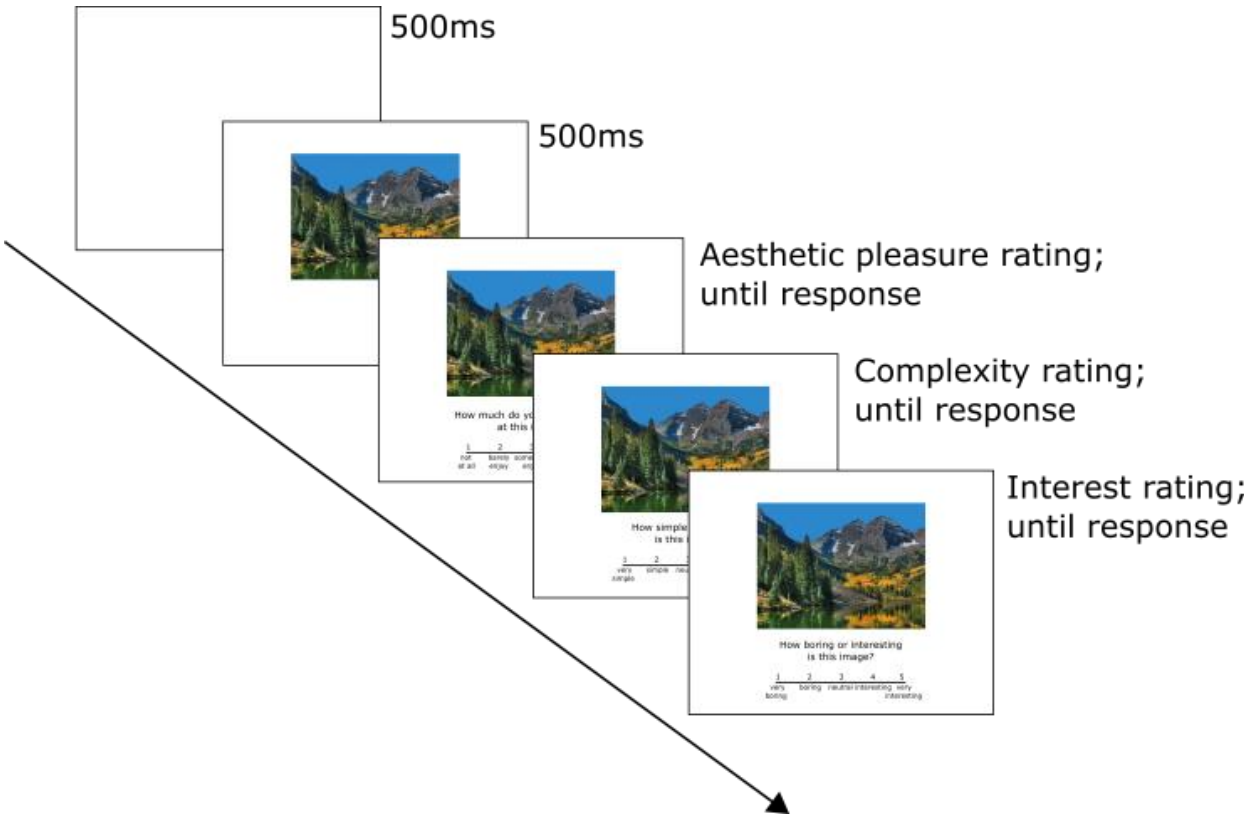

Figure 1: Example trial sequence.

Symmetry Analysis

In order to relate symmetry to subjective ratings of complexity, pleasure, and interest, several global and local symmetry features were calculated from the images.

Global Symmetry

The images were first transformed into grayscale using Matlab's rgb2gray function. Global symmetry measures were then obtained by creating an axis of symmetry at every column (vertical) and row (horizontal) of the image and correlating the rows or columns of pixel luminance values on one side of the axis of symmetry to an equal number of pixels on the other side. For instance, imagine the left-most column in an image. This column's global symmetry score would be the correlation of that column to an equal number of columns to its right (in this case, one column). This correlation is then multiplied by the proportion of the total pixels included in the correlation in order to downweigh the symmetry scores at the edges of the image. We do this because adjacent columns (or rows) of pixels tend to be similar to each other. Thus, on the edges of the image, when only a small portion of columns or rows are involved in the calculation, the symmetry score (i.e., the correlation) tends to be disproportionately high. We also do not believe that those high values truly capture high global symmetry at the edges, since calculating the symmetry over a small portion of the image is inherently not truly a "global" measure. We then move to the next column, this time correlating the two columns on the left of the axis to the two columns on the right (weighted by proportion of pixels), and so on.

In the end, this pixel correlation procedure along the length (and height) of the image results in a vector of vertical global symmetry scores and a vector of horizontal global symmetry scores for each image. See Figure 2 for example images displayed with their vertical and horizontal global symmetry vectors. The largest peak (i.e., global maximum) value of the vertical and horizontal global symmetry vectors for each image were recorded.

We also included two other measures of global symmetry, which represented variations of global symmetry within an image: the sum of the global symmetry range, and the number of peaks in the global symmetry vector. First, we calculated the sum of absolute differences in global symmetry values between every peak and trough of the global symmetry vector, separately for vertical and horizontal symmetry. In order to calculate the range, the vector needs to contain at least one peak and one trough. If the vector has only one peak without a trough, the symmetry range is 0. This global symmetry range score represented the degree to which there are large variations in luminance in the horizontal or vertical directions within an image. For example, the left image in Figure 2 has clear horizontal bands of dark-light-dark-light alternations, and its horizontal symmetry range score was large due to the large differences between each peak and trough in the horizontal symmetry vector. The final measure we included was the number of peaks in the global symmetry vectors, again separately for vertical and horizontal symmetry. The number of peaks represents the amount of unidirectional (i.e., vertical or horizontal) luminance changes, as described above, within an image.

To summarize, the global symmetry measures included in our analysis were the overall vertical and horizontal global symmetry of the image (Vertical Peak, Horizontal Peak), the global symmetry range (Vertical Range, Horizontal Range), and the volatility of global symmetry (# Peaks Vertical Symmetry, # Peaks Horizontal Symmetry).

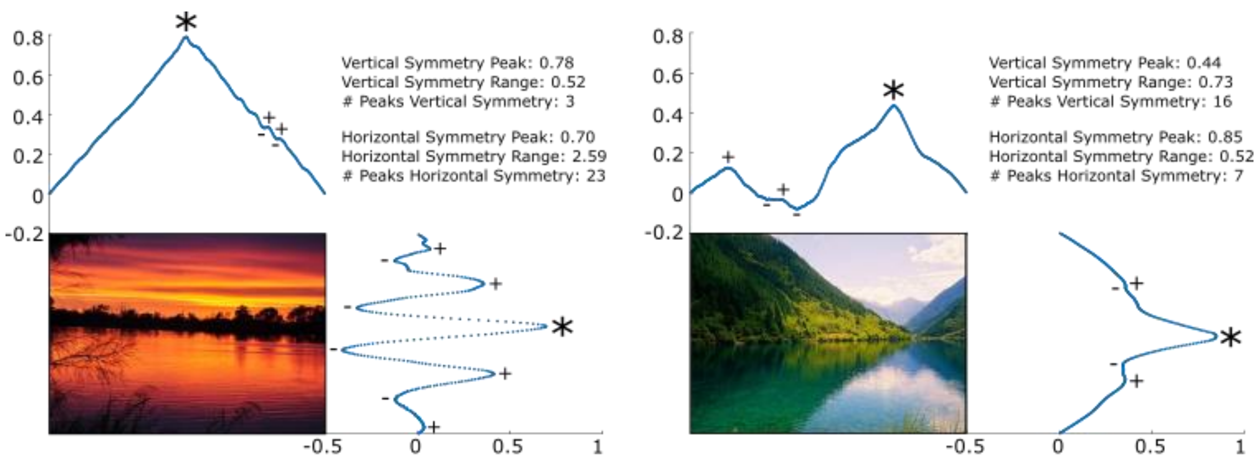

Figure 2: Vertical and horizontal global symmetry vectors for two example images. The global symmetry values corresponding to each image are displayed. The main global symmetry peak and several other obvious minor peaks and troughs have been marked along the global symmetry vectors using the * and $^{+}/^{-}$ symbols, respectively. Not all peaks and troughs have been marked, as they sometimes occur very close together along the symmetry vector.

Local Symmetry

Where global symmetry looks for a mirror, translational, or rotational symmetry in the image as a whole, local symmetry looks for a similar type of pattern in a local region of the image. For example, a vase is a symmetric object, but can be placed in a scene such that the scene as a whole is not symmetric. In this case the vase has local mirror symmetry. A garden hose, when laid straight has local mirror symmetry, but when coiled, or bent, the mirror symmetry is disrupted. However, there is still a constant distance from a central axis through the hose. This is another type of local symmetry that is referred to as local parallelism or local ribbon symmetry. Since local symmetry survives many changes in an image, such as those due to changes in the camera position, the relative positions of objects in scenes, and articulation of object parts, local symmetry is more common in scene images than global mirror symmetry, which requires the non-accidental alignment of objects in the scene and a particular positioning of the camera or observer. Indeed, local symmetry has been shown to be important for human perception (Kootstra & Schomaker, 2009; Nucci & Wagemans, 2007; Rezanejad et al., 2019; Wilder et al., 2019).

Local symmetry measures, such as ribbon symmetry, were obtained from line drawing versions of each image. In order to render all images as line drawings and calculate their local symmetry measures, we used the publicly available method developed and described by Rezanejad and colleagues (Rezanejad et al., 2019).

This method is used to score the local symmetry of contours or lines in the images, that is, the extent to which two nearby contours are locally symmetric with one another. As we are using colour photographs, we first perform edge detection on the colour image, using the publicly available structured edge detection toolbox as implemented in the LineDrawingExtraction package by Rezanejad et al. (2019; see https://github.com/mrezanejad). Once the edge map is obtained, the image is binarized using per image adaptive thresholding, resulting in a line drawing image with contour fragments that are 1 pixel wide. Finally, each contour fragment is spatially smoothed with a Gaussian kernel ($\sigma = 1$) to remove any discretization artifacts (i.e., to obtain smooth and continuous contours).

To calculate the symmetry scores on the line drawing (i.e., the contour image), we computed the medial axis of each region in the image (i.e., the regions between neighbouring pairs of contours) using the SalienceScoresForScene package (Rezanejad et al., 2019). It may be useful to think of the medial axis, sometimes called the shape skeleton, as being a stick figure for a shape. We wish to score the contours in our image based upon their behaviour around their medial axes. Traditionally the medial axis is defined for closed shapes, where it is the union of the centers of maximal inscribed discs within a shape (Blum, 1973). In our images, it is not easy to define what is the interior vs. exterior of shapes, therefore we obtained the medial axis for each region of the image. At each point on the medial axis, we know the radius of the disc centered at that point. For an entire medial axis, the set of radii are characterized by the radius function. Using the radius function, we calculated two types of symmetry scores: ribbon symmetry, and separation. Local parallelism (ribbon symmetry) is the extent to which the contours maintain a constant distance from the medial axis. A pair of contours with constant separation receives a high score. A pair of contours that have varying distances from one another will receive a low score. Local separation is the overall distance between two contours. More distal contours, yet still neighbouring contours (i.e. without any other contours between them), will receive a high score. Proximal neighbouring contours will receive a low score.

Concretely, for separation, the score assigned to each contour pixel is simply the radius itself. The larger the radius of the circle, the higher the separation score. Thus, contours that are far apart have high separation scores, while contours that are very close together have low separation scores. In the case that the contour was associated with two medial axes, the contour's separation score was the maximum of the two values.

For ribbon symmetry, the rate of change of the radius function was calculated along the length of each skeleton contour, within a threshold window with a length of $2K + 1$, where $K = 5$ pixels. The lower the rate of change, the higher the ribbon symmetry score, because this means that the distance between contours remains locally constant along the length of the contours. Again, the contour pixels were assigned the maximum of the ribbon symmetry values from the closest points along the medial axis on either side. See Figure 3a for an example image at each step of the local symmetry calculation process. See Figure 3b for simplified shape examples that illustrate how the local symmetry measures change depending on the relationships between contours.

To obtain the overall local ribbon symmetry within an image, the median value of all ribbon symmetry scores within an image was recorded. We also included two other measures of local symmetry for each image, which were chosen to explore the influence of the presence of *any* local symmetry or asymmetry within an image. For example, an image may have several parallel contours and several asymmetric contours (e.g., an image of a forest with parallel tree trunks and asymmetric groups of branches and leaves). This image would receive a medium overall symmetry score, but the mere presence of the locally symmetric or asymmetric objects may play a role in subjective judgements of complexity, pleasure, or interest. For this reason, we included a measure of the proportion of symmetric pixels (i.e., the proportion of pixels within the image with ribbon symmetry scores greater than 0.75) and the proportion of asymmetric pixels (the proportion of pixels with ribbon symmetry scores less than 0.25). The same measures were also included for the separation measure.

While ribbon symmetry is a clear measure of local symmetry, or parallelism, separation reflects local symmetry in a very different way. It measures relatively longer-range interactions between local parts of a scene, capturing the extent to which a medial axis extends its influence. However, just like images with symmetric contours are better categorized than images with asymmetric contours, images with the 50% of contours with the highest separation scores are categorized more easily than images with the 50% lowest separation contours (Rezanejad et al., 2019). This result suggests that the contours with greater separation also facilitate mid-level perceptual organization (e.g., figure-ground segregation, segmentation into salient regions) and may cue the layout of the scene, allowing for more efficient scene categorization. Thus, we chose to include both ribbon symmetry and separation in our analysis.

In sum, the local symmetry measures included in our analysis were the overall ribbon symmetry and separation scores of each image (Ribbon: median, Separation: median), the proportion of symmetric and far contours (Ribbon: proportion symmetric, Separation: proportion far), and the proportion of asymmetric and near contours (Ribbon: proportion asymmetric, Separation: proportion near).

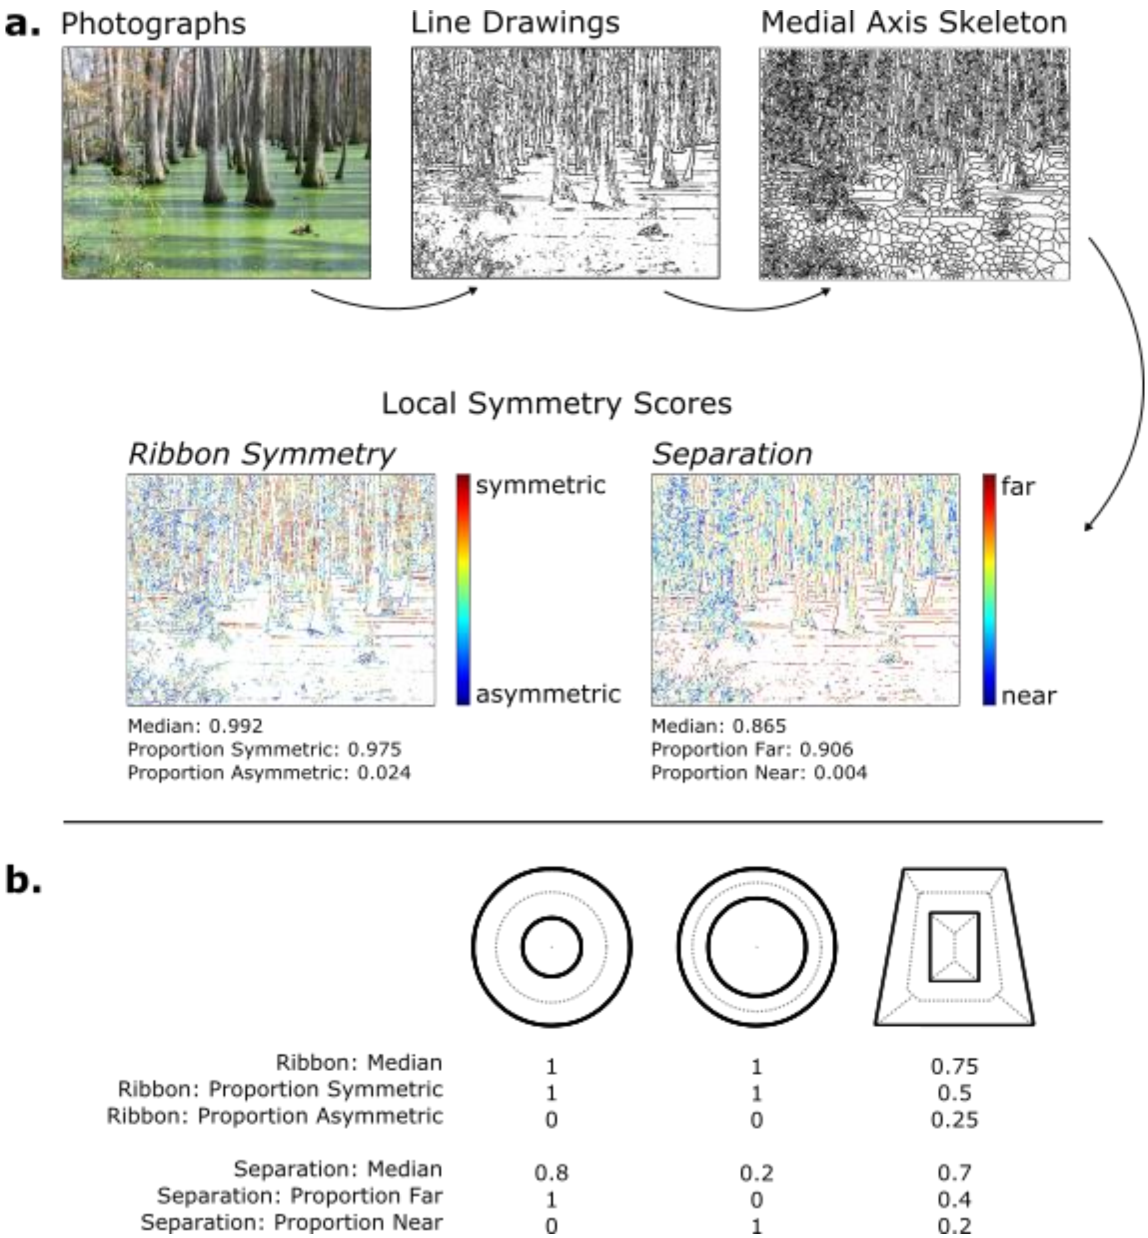

Figure 3: a. Calculation of local symmetry measures for one example image; The image is first transformed from a colour photograph to a line drawing. The medial axis transform is then computed, allowing for the symmetry scoring of each contour pixel. The local symmetry values corresponding to the image are displayed. **b.** Simplified examples of shapes with varying amounts of ribbon symmetry and separation, with their corresponding local symmetry values. The grey dotted lines represent the medial axis skeletons of the shapes. The circles' contours are perfectly parallel, therefore their ribbon symmetry is high, while their separation scores depend on the distance between contours. The rightmost shape contains sections that are not perfectly parallel, therefore its ribbon symmetry is relatively lower compared to the circles. Its separation scores also differ from those of the circles because this shape contains sections of contours that are close together as well as those that are further apart.

Canonical Correlation Analysis

A canonical correlation analysis was used in order to determine the relationships within and between two variable sets. One variable set included the symmetry features; 6 global symmetry variables, and 6 local symmetry variables. The other variable set included the three subjective ratings of complexity, aesthetic pleasure, and interest.

Due to the high correlations between the three subjective rating variables (see Appendix A), a multivariate statistical approach was preferred over separate univariate regressions in order to account for these relationships among the dependent variables. The additional advantage of using a canonical correlation analysis, rather than a multivariate regression analysis, for example, is that it further allows us to find more than one relationship between the dependent and independent variables, if more than one relationship exists. In this sense, a canonical correlation is equivalent to a multivariate regression followed by linear discriminant analysis (Lutz & Eckert, 1994). Due to the inconsistent findings in the literature regarding the relationships between aesthetic pleasure, complexity, and symmetry, we hypothesized that a singular regression outcome would not properly capture the nuances of these relationships. Thus, we believe that a canonical correlation analysis was the most appropriate way to explore our study question.

Concretely, the canonical correlation seeks to maximize the correlations between two variable sets (Sherry & Henson, 2005) – a set of dependent variables and a set of independent, or predictor, variables. It does so by fitting a linear equation to the observed predictor variables (in our case, the symmetry measures) to create a single synthetic predictor variable, and another linear equation to the observed dependent variables (i.e., the subjective rating variables) to create a synthetic criterion variable, and then maximizing the correlation between the two synthetic variables. The process repeats itself using the residual variance of the first maximal correlation. Each of these maximal correlations is called a canonical root. This analysis yields as many canonical roots as variables in the smaller variable set, and each root is orthogonal to the previous root(s). Thus, this approach allows us to assess the several different relationships that may exist within and between our variable sets (i.e., within and between different symmetry measures and subjective ratings of complexity, pleasure, and interest).

Results

In order to evaluate the multivariate shared relationship between symmetry and subjective ratings of complexity, pleasure, and interest, we conducted a canonical correlation analysis using the twelve symmetry measure variables as predictors of the three subjective rating variables. The analysis was run in R (R Core Team, 2018) using the "cc" function from the CCA package (González & Déjean, 2012). The analysis yielded three canonical roots with squared canonical correlations (R_c^2) of 0.28, 0.17, and 0.04 for each successive root. Collectively, the full model across all three roots was statistically significant, Wilks' $\lambda = 0.576$, $F(36, 2083.73) = 11.89$, $p < 0.0001$. Wilks' λ represents the variance that was not explained by the model, therefore $1 - \lambda$ yields the full model effect size. Thus, the full model explained a substantial portion, approximately 42.4%, of the variance shared between the variable sets.

Statistical significance of the canonical roots can be assessed with a dimension reduction approach, using the asymptotic test function ("p.asym") from the CCP package (Menzel, 2012) in R. As noted above, the full model (i.e., roots 1 to 3) was significant. Additionally, roots 2 to 3, and root 3 alone were also statistically significant, $F(22, 1412) = 7.83$, $p < 0.0001$, and $F(10, 707) = 2.85$, $p = 0.002$, respectively (see Table 1). Although this test shows that all three roots are statistically significant, only the first two roots were considered noteworthy in the context of this study. The first and second roots explained 27.5% and 12.6% of the model variance, respectively, while the third root only explained 2.3%, and thus was not examined further (see Table 2).

Table 1: Hierarchical statistical significance tests

Roots	*Wilks' λ*	*F*	*Hypothesis DF*	*Error DF*	*Significance of F*
1 to 3	0.576	11.893	36	2083.729	0.000
2 to 3	0.794	7.829	22	1412.000	0.000
3 to 3	0.961	2.846	10	707.000	0.002

Table 2: Canonical correlations for each root

Roots	*Canonical Correlations*	*Squared Correlation*	*% of shared variance*	*% of model variance*
1	0.524	0.275	27.5	27.5
2	0.417	0.174	17.4	12.6
3	0.197	0.039	3.9	2.3
			Total variance explained by the model:	***42.4%***

Table 3 presents the standardized canonical function coefficients and structure coefficients for roots 1 and 2. Canonical function coefficients are analogous to beta weights in a regression, as they represent the standardized coefficients used in the linear equations to combine the observed predictor variables (in our case, symmetry measures) and criterion variables (subjective ratings) into two synthetic variables. Structure coefficients represent the correlation between an observed variable and the synthetic variable created from the linear equation, and define which variables are useful in the model. The squared structure coefficients are also shown.

Finally, we include the communalities (h^2) across the two functions for each variable. Communalities are the sums of the squared structure coefficients, and represent the proportion of variance explained across the interpretable canonical roots for each variable. Across the two roots, the most useful variables (i.e., communality coefficients > 35%) for the canonical correlation solution were the three subjective rating variables (Pleasure, Complexity, and Interest), and several local and global symmetry variables: the overall local symmetry (Ribbon: median) and separation (Separation: Median) within an image, the proportion of separated contours (Separation: Proportion Far), as well as the overall global symmetry (Vertical Peak, Horizontal Peak), and the number of luminance changes within an image (# Peaks Vertical Symmetry, # Peaks Horizontal Symmetry).

We could also assess which variables were most relevant to each individual root according to their structure coefficients (i.e., structure coefficients greater than |0.45|). Firstly, in root 1, we found that all three subjective rating variables were relevant, especially Complexity, which has a very large structure coefficient. These rating variables are positively related to each other, as their structure coefficients have the same sign (i.e., positive).

Regarding the symmetry variable set in root 1, the primary contributors were local symmetry (Ribbon: Median) and separation (Separation: Median), the proportion of separation (Separation: Proportion Far), vertical global symmetry (Vertical Peak), and large luminance changes within an image (Vertical Range and Horizontal Range). With the exception of Vertical Range, all relevant symmetry variables had the same sign (i.e., negative), indicating that they were all positively related to each other, though they were negatively related to the subjective rating variables. Conversely, Vertical Range was negatively related to the other symmetry variables, but positively related to the rating variables.

Table 3: Canonical solution for features predicting ratings for root 1 and 2

		Root 1			*Root 2*		
Variable	*Coef*	r_s	r_s^2(%)	*Coef*	r_s	r_s^2(%)	h^2(%)
Ribbon: Median	-0.792	-0.614	37.66	-0.541	-0.523	27.31	64.97
Ribbon: Proportion Symmetric	16.194	-0.330	10.91	3.505	-0.393	15.47	26.38
Ribbon: Proportion Asymmetric	15.063	0.330	10.89	3.548	0.392	15.38	26.27
Separation: Median	-0.431	-0.531	28.15	0.097	-0.397	15.76	43.91
Separation: Proportion Far	-0.342	-0.507	25.66	0.177	-0.501	25.12	50.78
Separation: Proportion Near	0.282	-0.241	5.82	-0.200	-0.165	2.72	8.54
Vertical Symmetry Peak	-0.420	-0.607	36.87	0.587	0.079	0.63	37.50
Vertical Symmetry Range	0.369	0.564	31.86	-0.240	-0.037	0.14	32.00
# Peaks Vertical Symmetry	-0.099	0.385	14.84	0.762	0.537	28.86	43.70
Horizontal Symmetry Peak	-0.244	-0.195	3.79	-0.332	-0.643	41.28	45.07
Horizontal Symmetry Range	-0.071	-0.548	30.00	0.163	-0.214	4.56	32.84
# Peaks Horizontal Symmetry	-0.548	-0.083	0.69	0.353	0.705	49.64	50.33
Pleasure	0.287	0.659	43.41	-0.589	-0.748	56.00	99.41
Complexity	2.060	0.912	83.24	1.928	-0.396	15.70	98.94
Interest	-1.497	0.714	50.95	-1.921	-0.689	47.44	98.39

Note: Structure coefficients (r_s) greater than |0.45| are underlined. Communality coefficients (h^2) greater than 35% are underlined. Coef = standardized canonical function coefficient; r_s = structure coefficient; r_s^2 = squared structure coefficient; h^2 = communality coefficient.

These relevant variables also tended to have medium-to-large standardized canonical function coefficients, with the exception of Horizontal Range, which had a relatively small function coefficient but a large structure coefficient. This result is due to multicollinearity that this variable had with other symmetry variables. Multicollinearity within the context of a canonical correlation is not a problem, since the canonical correlation is an exploratory analysis to find relationships within the data, and the variables are not necessarily assumed to be independent. In any case, the multicollinearity still indicates that the Horizontal Range variable's specific relationship to the dependent variables cannot be estimated independently from other symmetry variables and should be interpreted with caution.

Conversely, the proportion symmetric and asymmetric variables (Ribbon: Proportion Symmetric, and Ribbon: Proportion Asymmetric) had very high function coefficients but small structure coefficients. This indicates that those variables are acting as suppressor variables, which are variables that do not necessarily relate to the dependent variable set on their own, but nonetheless may contribute to the model because their inclusion increases the predictive validity of other variables (MacKinnon et al., 2000; Nimon et al., 2015). In other words, the proportion symmetric and asymmetric variables may explain some of the variance in other variables, thus including them makes the effects of those other variables stronger. However, *excluding* them does not have a large effect on the model nor the structure coefficients of the symmetry and rating variables, thus we can safely interpret the proportion symmetric and proportion asymmetric variables as not contributing much to the canonical correlation model.

Overall, the pattern of results found in root 1 indicates that presence of symmetry decreases complexity, pleasure, and interest. This is the case for both the local symmetry measures (i.e., ribbon symmetry and separation), as well as vertical global symmetry. Large global symmetry variations in the horizontal and vertical directions are also related to the subjective ratings, with horizontal variations being related negatively to complexity, pleasure, and interest, while vertical variations are positively related to complexity, pleasure, and interest.

Moving on to root 2, the structure coefficients in Table 3 suggest that the rating variables of relevance for this root were Pleasure and Interest, but *not* Complexity. Recall that the relationships found in this root are orthogonal to those found in the first root, thus this root is explaining the variance in subjective pleasure and interest scores that are independent from their relationship to subjective complexity. Again, the relevant rating variables (i.e., Pleasure and Interest) had the same sign and thus were positively related to each other.

As for the symmetry variables, the variables related to the number of luminance changes within an image (# Peaks Vertical Symmetry, # Peaks Horizontal Symmetry) were dominant contributors to this root and were inversely related to Pleasure and Interest. The relevant local symmetry variables (i.e., Ribbon: Median, and Separation: Proportion Far), however, were this time *positively* related to Pleasure and Interest. Horizontal global symmetry (Horizontal Peak) also contributed significantly to this root and was positively related to the rating variables. Therefore, unlike root 1, in root 2, we found the overall pattern that symmetry increases pleasure and interest.

In a control analysis (see Appendix B) including measures of simplicity (i.e., ZIP compressibility of the image) and self-similarity, the symmetry measures no longer contribute to the model in the first root. Instead, simplicity is the only contributing variable. However, a multiple linear regression analysis revealed that the symmetry measures that are significant in our original canonical correlation analysis are highly predictive of the simplicity measure in this regression analysis, $F(12,707) = 98.91$, $p < 0.0001$, adjusted $R^2 = 0.62$, suggesting that symmetry is still indeed predictive of subjective complexity, pleasure, and interest (i.e., negatively correlated), but the contributions of the symmetry measures are captured within this broader simplicity measure (see Appendix B, Table B2). This makes sense, as the redundant information available in symmetric images makes the image more easily compressible.

In root 2, the pattern of results was not affected by the inclusion of the control variables. In this root, the symmetry measures still positively related to aesthetic pleasure and interest, and self-similarity also contributed positively to the model. The simplicity measure did not contribute to the model in root 2.

Taken together, these results are supportive of the theoretically expected relationships between (dis)fluency/complexity and pleasure and interest suggested by the Pleasure-Interest Model of Aesthetic Liking (PIA; Graf & Landwehr, 2015). The PIA model suggests that aesthetic preference is partially due to automatic and fluent processing (termed “aesthetic pleasure”), and partially due to the reduction of disfluency that is achieved through cognitive elaborations and reflections (termed “aesthetic interest”). If there is no disfluency to be reduced, the result is boredom and thus a negative aesthetic experience. Conversely, if disfluency exists and is sufficiently reduced, the result is interest and a positive aesthetic experience.

Since symmetry is known to increase fluency, and symmetry was negatively related to complexity, pleasure, and interest in the first root, we labeled root 1 as “boredom”. We labeled root 2 as “interest”, since symmetry here was positively related to pleasure and interest (see Discussion section for further elaboration).

Discussion

Using a canonical correlation analysis, we found two significant and interpretable canonical roots that revealed two distinct relationships between symmetry and subjective ratings of complexity, pleasure, and interest. The first root, which we termed “boredom”, involved many negative relationships between the symmetry measures and the subjective ratings. The second root, on the other hand, involved positive relationships between symmetry and pleasure and interest, thus we labeled it as “interest”. We will now delve into each root individually and describe the specific relationships we found, interpreting our results within the framework of the pleasure-interest model of aesthetic liking (PIA) (Graf & Landwehr, 2015).

Beginning with root 1, as expected, we found that vertical global symmetry relates negatively to complexity. The more globally symmetric the scene, the lower the complexity score it received. In addition to complexity, we also found that a similar relationship exists between vertical global symmetry and aesthetic pleasure and interest as well, and that all three subjective rating measures were positively related to each other. This implies that, as symmetry increases and makes a scene less complex, it also renders the scene less pleasing and interesting. Importantly, we have shown that this is not only the case for vertical global symmetry, but also for our *local* symmetry measures. Local symmetry measures had not previously been explored in relation to aesthetic pleasure, especially with natural scene stimuli. Using our novel method of extracting local symmetry directly from colour photographs of scenes (Rezanejad et al., 2019), we were able to show that local ribbon symmetry and separation related negatively to the subjective ratings in the first canonical root. In other words, local symmetry also decreases scene complexity and makes the image less pleasant and interesting. It is also important to note that local symmetry and vertical global symmetry were positively related to each other in root 1, meaning that images with high vertical global symmetry also tended to have higher local symmetry. Similar to how a combination of global and local symmetries results in the greatest detectability of regularities in dot patterns (Nucci & Wagemans, 2007), we suggest that the combination of vertical global symmetry and local symmetry in scenes results in a highly simple scene without much to be processed, rendering it boring and therefore unpleasant.

The other two measures involved in this root were the vertical and horizontal global symmetry range scores. More specifically, vertical global symmetry range was positively related to the subjective ratings, while the horizontal global symmetry range was negatively related to the subjective ratings. Though they were calculated directly from the global symmetry scores, the range measures are not so much measures of symmetry, but rather of horizontal or vertical regularities (i.e., bands of alternating high and low luminance values) that exist within an image. This is an interesting finding because it implies that vertical bands are considered to be relatively complex while the horizontal bands are seen as less complex. This may be due to people's general biases to scan their environments along the horizon, or in a horizontal fashion (Foulsham et al., 2008; Tatler & Vincent, 2008), even as infants (Van Renswoude et al., 2016), likely allowing for more fluent processing when visual features (e.g., luminance) remain relatively constant horizontally compared to when there are fluctuations (as would be the case in an image with high vertical range).

Interestingly, in a control analysis, we found that the symmetry measures which contribute to the first root are almost completely captured by a computational measure of simplicity (i.e., ZIP compression). This finding is particularly important because these types of objective measures of simplicity and symmetry are often treated as separate measures. However, we have shown that such a division is inappropriate when more specific aspects of symmetry, such as local symmetry and vertical or horizontal luminance regularities, are considered. Just as symmetry is thought to contribute to subjective simplicity (Bertamini et al., 2013; Day, 1968; Gartus & Leder, 2017), our findings suggest that symmetry measures may also contribute to more objective or computational measures of simplicity, such as ZIP compression. Thus, future studies that explore both symmetry and simplicity should consider this potential overlap before drawing strong conclusions regarding their relative contributions.

Overall, the pattern of results found in the first canonical root (labeled "boredom") seems to reflect the idea that when landscapes are too simple, due to the symmetry and horizontal regularities present in the scene, people find them uninteresting and not particularly pleasing. For this reason, we suggest that symmetry leads to increased boredom, within the framework of the PIA model (Graf & Landwehr, 2015). The model suggests that one aspect of aesthetic liking comes from a controlled contemplative process meant to decrease disfluency (e.g., coming to admire a piece of abstract art after reflecting on one's personal experience while viewing the art piece). Within our stimulus set of natural scene images, however, all scenes are relatively easy to understand, and therefore would not require much contemplation to begin with. Thus, if a scene was highly symmetric, this would leave even less room for contemplation or elaborated processing, rendering the experience dull and unpleasant.

The second canonical root reveals a contrasting pattern. In the second root, the contributing local symmetry measures are *positively* related to pleasure and interest. Additionally, horizontal global symmetry is also positively related to pleasure and interest. This finding now implies that combinations of local symmetry and *horizontal* global symmetry render the scene *more* pleasant and interesting. This is further supported by the negative relationship between the subjective ratings and the number of vertical and horizontal global symmetry variations (# peaks vertical symmetry and # peaks horizontal symmetry). A high number of peaks represents an unstable global symmetry vector, implying a lack of a single clear global symmetry axis. Thus, we suggest that it is the presence of a well-defined horizontal global symmetry axis, as well as local ribbon symmetry and a large proportion of far contours, that make a scene interesting and pleasant.

Again, within the framework of the PIA model, we believe that symmetry in the second canonical root (labeled "interest") relates to increased aesthetic pleasure and interest. A point of significance is that horizontal global symmetry relates positively to pleasure and interest here while vertical global symmetry relates negatively (in the first root). This is perhaps due to the fact that horizontal global symmetry is not as salient as vertical global symmetry (Wagemans et al., 1992; Wenderoth, 1994), thus requiring a slightly longer amount of time to process than vertical symmetry. This extra time may have helped to avoid processing the image too quickly or fluently. Indeed, a recent study found that vertical global symmetry, but not horizontal symmetry, was related to subjective fluency of artistic photographs (Vissers & Wagemans, *in press*). However, we argue that the presence of (horizontal) symmetry still allows for a reduction of *disfluency*, resulting in feelings of interest and pleasure. Similarly, in predictive coding terms, pleasure comes from being able to explain away prediction errors (Van de Cruys & Wagemans, 2011). Thus, it is also possible that the uniqueness of scenes with horizontal symmetry initially creates a prediction error which is then resolved by further processing of the scene.

Finally, our work also strongly highlights the importance of investigating the role of *local* symmetry within high-level stimuli. Local symmetry cues in real-world scenes are important for rapid scene categorization (Rezanejad et al., 2019; Wilder et al., 2019), and they attract category-specific eye movements (Damiano et al., 2019). Here, we have shown that they also relate to subjective aesthetic judgements, in both a negative (root 1) and positive (root 2) way. These results help to clear up the apparent discrepancies that exist within the literature regarding the link between complexity and aesthetic pleasure in natural scenes (i.e., fluency vs. complexity). In our stimulus set, complexity and pleasure were highly correlated, suggesting that complexity is indeed pleasurable in natural scenes. However, it is also the case that symmetry was, in part, positively related to pleasure, which implies that fluency, at least to the extent that symmetry can be related to fluent processing, also plays a role.

One previous study that also explored the link between aesthetic pleasure and visual features in landscape images found that symmetry did not significantly relate to the number of views on an image, which was used as a measure of aesthetic liking (Study 3 in Mayer & Landwehr, 2018). However, they found that two objective measures (i.e., ZIP compression and fractal self-similarity) did positively predict aesthetic liking. Thus, we ran a control analysis that included these two variables (see Appendix B). Interestingly, when we controlled for objective simplicity (i.e., ZIP compressibility) and self-similarity, we found that the symmetry measures no longer contributed to our model in the first canonical root. The symmetry measures were instead accounted for by the simplicity measure. This is not surprising, given the nature of ZIP compression. That is, because compressibility proceeds in a row-wise fashion, redundant information that occurs horizontally along the image (as is the case with vertical symmetry and the horizontal range measures) is directly related to its compressibility (i.e., the simplicity value). Thus, a possible reason that symmetry did not seem to relate to aesthetic judgements in Mayer and Landwehr's (2018) Study 3 is that its role may have been captured by the simplicity measure.

Another difference between this previous study and our current study is that we found a negative relationship between simplicity and aesthetic pleasure while they found a positive relationship. We speculate that this difference is due to the aesthetic liking measure used in their study vs. ours (i.e., number of image views vs. subjective ratings by participants, respectively). Their measure likely captures an automatic aesthetic experience because, while scrolling through a photography site (Flickr), people will only click on (i.e., view) images that automatically capture their attention. Thus, in terms of the PIA model, their measure better reflects the automatic portion of the model where aesthetic pleasure is based on fluent processing (i.e., the simpler the better), as opposed to the more contemplative portion of the model where the reduction of disfluency is more important to the aesthetic experience. Our subjective rating measure likely captures this contemplative portion more strongly, therefore simplicity in our case is actually negatively related to pleasure because it leads to boredom.

It is important to note that we cannot rule out the possibility that some of our results may be driven by the content of the images rather than the symmetry, per se. For example, many of the scenes with high horizontal global symmetry depict a mountain or forest scene reflected on a still body of water. Thus, the category of the scene may drive the aesthetic pleasure judgement, rather than the horizontal symmetry alone. To rule out this possibility, rather than complex natural scenes, future studies could explore these relationships in other high-level stimuli, such as collections of real-world objects that have been organized in unique and pleasant ways (i.e., neatly organized compositions) that also show a positive relationship between complexity and aesthetic preferences (Van Geert & Wagemans, 2019).

In summary, we found that symmetry is negatively related to complexity judgements and both positively and negatively related to aesthetic pleasure and interest judgements. We discussed our results within the framework of the Pleasure-Interest model of Aesthetic liking (PIA). The results found in the first canonical root suggest that scenes with combinations of high vertical global symmetry and local symmetry create a highly simplistic visual percept, resulting in boredom. Conversely, horizontal global symmetry along with local symmetry result in aesthetic interest. Within the model, these outcomes come about via the contemplative processes that occur once a stimulus has been processed and are the result of a lack of disfluency (i.e., no disfluency to be reduced), and disfluency reduction, respectively. Our experimental design did not allow for us to assess the role of symmetry in the ‘automatic processing’ portion of the PIA model, as participants had an unlimited amount of time to view the image before making their judgements. In the future, it would be fruitful to explore whether the relationships between symmetry, complexity, pleasure, and interest, change if the image can only be processed for a very short amount of time (i.e., a brief flash).

Overall, our results lend support to the idea that both complexity and simplicity influence the aesthetic pleasure of natural scenes. Our work highlights the importance of considering both local and global symmetry when exploring symmetry’s relationship with subjective judgements, as it is the combination of different aspects of local and global symmetry that reveals the nuances of the relationships that exist between symmetry perception and aesthetic experiences in natural scenes. By studying the fundamental principles of perceptual organization in relation to aesthetics, especially using an ecologically valid stimulus set, we move closer to the ultimate goal of understanding the elaborate links between human perception, cognition, and emotion.

References

Berlyne, D. E. (1971). *Aesthetics and psychobiology*. New York, NY: Appleton-Century-Crofts.

Berman, M. G., Jonides, J., & Kaplan, S. (2008). The Cognitive Benefits of Interacting With Nature. *Psychological Science*, *19*(12), 1207–1212. https://doi.org/10.1111/j.1467-9280.2008.02225.x

Bertamini, M., Makin, A., & Rampone, G. (2013). Implicit Association of Symmetry with Positive Valence, High Arousal and Simplicity. *I-Perception*, *4*(5), 317–327. https://doi.org/10.1068/i0601jw

Biederman, I., & Vessel, E. (2006). Perceptual Pleasure and the Brain: A novel theory explains why the brain craves information and seeks it through the senses. *American Scientist*, 8.

Bowler, D. E., Buyung-Ali, L. M., Knight, T. M., & Pullin, A. S. (2010). A systematic review of evidence for the added benefits to health of exposure to natural environments. *BMC Public Health*, *10*(1). https://doi.org/10.1186/1471-2458-10-456

Chen, C.-C., Wu, J.-H., & Wu, C.-C. (2011). Reduction of Image Complexity Explains Aesthetic Preference for Symmetry. *Symmetry*, *3*(3), 443–456. https://doi.org/10.3390/sym3030443

Damiano, C., Wilder, J., & Walther, D. B. (2019). Mid-level feature contributions to category-specific gaze guidance. *Attention, Perception, & Psychophysics*, *81*(1), 35–46. https://doi.org/10.3758/s13414-018-1594-8

Day, H. (1968). The importance of symmetry and complexity in the evaluation of complexity, interest and pleasingness. *Psychonomic Science*, *10*(10), 339–340. https://doi.org/10.3758/BF03331550

Foulsham, T., Kingstone, A., & Underwood, G. (2008). Turning the world around: Patterns in saccade direction vary with picture orientation. *Vision Research*, *48*(17), 1777–1790. https://doi.org/10.1016/j.visres.2008.05.018

Gartus, A., & Leder, H. (2013). The Small Step toward Asymmetry: Aesthetic Judgment of Broken Symmetries. *I-Perception*, *4*(5), 361–364. https://doi.org/10.1068/i0588sas

Gartus, A., & Leder, H. (2017). Predicting perceived visual complexity of abstract patterns using computational measures: The influence of mirror symmetry on complexity perception. *PLOS ONE*, *12*(11), e0185276. https://doi.org/10.1371/journal.pone.0185276

Graf, L. K. M., & Landwehr, J. R. (2015). A Dual-Process Perspective on Fluency-Based Aesthetics: The Pleasure-Interest Model of Aesthetic Liking. *Personality and Social Psychology Review*, *19*(4), 395–410. https://doi.org/10.1177/1088868315574978

Güçlütürk, Y., Jacobs, R. H. A. H., & Lier, R. van. (2016). Liking versus Complexity: Decomposing the Inverted U-curve. *Frontiers in Human Neuroscience*, *10*. https://doi.org/10.3389/fnhum.2016.00112

Kaplan, S. (2001). Meditation, Restoration, and the Management of Mental Fatigue. *Environment and Behavior*, *33*(4), 480–506. https://doi.org/10.1177/00139160121973106

Koffka, K. (1922). Perception: an introduction to the Gestalt-Theorie. *Psychological Bulletin*, *19*(10), 531.

Koffka, K. (1935). Principles of Gestalt Psychology (New York, 1935). *Chap. II.*

Köhler, W. (1920). Physical Gestalten. Reprinted in WD Ellis (Ed.)(1929), A Source Book of Gestalt Psychology.

Kootstra, G., & Schomaker, L. R. (2009). Prediction of human eye fixations using symmetry. In *The 31st Annual Conference of the Cognitive Science Society (CogSci09)* (pp. 56-61). Cognitive Science Society.

Leder, H., Tinio, P. P. L., Brieber, D., Kröner, T., Jacobsen, T., & Rosenberg, R. (2019). Symmetry Is Not a Universal Law of Beauty. *Empirical Studies of the Arts*, *37*(1), 104–114. https://doi.org/10.1177/0276237418777941

Little, A. C., Jones, B. C., & DeBruine, L. M. (2011). Facial attractiveness: Evolutionary based research. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1571), 1638–1659. https://doi.org/10.1098/rstb.2010.0404

Lutz, J. G., & Eckert, T. L. (1994). The relationship between canonical correlation analysis and multivariate multiple regression. *Educational and Psychological Measurement*, *54*(3), 666-675. https://doi.org/10.1177/0013164494054003009

MacKinnon, D. P., Krull, J. L., & Lockwood, C. M. (2000). *Equivalence of the Mediation, Confounding and Suppression Effect*. 9.

Makin, A. D. J., Pecchinenda, A., & Bertamini, M. (2012). Implicit affective evaluation of visual symmetry. *Emotion*, *12*(5), 1021–1030. https://doi.org/10.1037/a0026924

Mayer, S., & Landwehr, J. R. (2018). Quantifying visual aesthetics based on processing fluency theory: Four algorithmic measures for antecedents of aesthetic preferences. *Psychology of Aesthetics, Creativity, and the Arts*, *12*(4), 399. http://dx.doi.org/10.1037/aca0000187

Miniukovich, A., & De Angeli, A. (2014). Visual impressions of mobile app interfaces. *Proceedings of the 8th Nordic Conference on Human-Computer Interaction Fun, Fast, Foundational - NordiCHI ’14*, 31–40. https://doi.org/10.1145/2639189.2641219

Nimon, K. F., Zientek, L. R., & Thompson, B. (2015). Investigating bias in squared regression structure coefficients. *Frontiers in Psychology*, *6*. https://doi.org/10.3389/fpsyg.2015.00949

Nucci, M., & Wagemans, J. (2007). Goodness of Regularity in Dot Patterns: Global Symmetry, Local Symmetry, and Their Interactions. *Perception*, *36*(9), 1305–1319. https://doi.org/10.1068/p5794

Olivia, A., Mack, M. L., Shrestha, M., & Peeper, A. (2004). Identifying the perceptual dimensions of visual complexity of scenes. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 26, No. 26).

Panis, S., & Wagemans, J. (2009). Time-course contingencies in perceptual organization and identification of fragmented object outlines. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(3), 661–687. https://doi.org/10.1037/a0013547

Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing Fluency and Aesthetic Pleasure: Is Beauty in the Perceiver's Processing Experience? *Personality and Social Psychology Review*, *8*(4), 364–382. https://doi.org/10.1207/s15327957pspr0804_3

Rezanejad, M., Downs, G., Wilder, J., Walther, D. B., Jepson, A., Dickinson, S., & Siddiqi, K. (2019). Scene categorization from contours: Medial axis based salience measures. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 4116-4124).

Rhodes, G., Proffitt, F., Grady, J. M., & Sumich, A. (1998). Facial symmetry and the perception of beauty. *Psychonomic Bulletin & Review*, *5*(4), 659–669. https://doi.org/10.3758/BF03208842

Schertz, K. E., & Berman, M. G. (2019). Understanding Nature and Its Cognitive Benefits. *Current Directions in Psychological Science*, *28*(5), 496–502. https://doi.org/10.1177/0963721419854100

Sherry, A., & Henson, R. K. (2005). Conducting and Interpreting Canonical Correlation Analysis in Personality Research: A User-Friendly Primer. *Journal of Personality Assessment*, *84*(1), 37–48. https://doi.org/10.1207/s15327752jpa8401_09

Sukumar, S. R., Page, D. L., Koschan, A. F., & Abidi, M. A. (2008). Towards understanding what makes 3D objects appear simple or complex. *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 1–8. https://doi.org/10.1109/CVPRW.2008.4562975

Tatler, B. W., & Vincent, B. T. (2008). Systematic tendencies in scene viewing. *Journal of Eye Movement Research*. https://doi.org/10.16910/JEMR.2.2.5

Thornhill, R., & Gangestad, S. W. (2006). Facial sexual dimorphism, developmental stability, and susceptibility to disease in men and women. *Evolution and Human Behavior*, *27*(2), 131–144. https://doi.org/10.1016/j.evolhumbehav.2005.06.001

Van de Cruys, S., & Wagemans, J. (2011). Putting Reward in Art: A Tentative Prediction Error Account of Visual Art. *I-Perception*, *2*(9), 1035–1062. https://doi.org/10.1068/i0466aap

Van Geert, E., & Wagemans, J. (2019). Order, complexity, and aesthetic preferences for neatly organized compositions. *Psychology of Aesthetics, Creativity, and the Arts*. https://doi.org/10.1037/aca0000276

Van Geert, E., & Wagemans, J. (2020). Order, complexity, and aesthetic appreciation. *Psychology of Aesthetics, Creativity, and the Arts*, *14*(2), 135–154. https://doi.org/10.1037/aca0000224

Van Renswoude, D. R., Johnson, S. P., Raijmakers, M. E. J., & Visser, I. (2016). Do infants have the horizontal bias? *Infant Behavior and Development*, *44*, 38–48. https://doi.org/10.1016/j.infbeh.2016.05.005

Vissers, N., & Wagemans, J. (*in press*). Fluency, Processing Style and Aesthetic Response to Artistic Photographs. *Psychology of Aesthetics, Creativity, and the Arts.*

Wagemans, J. (1997). Characteristics and models of human symmetry detection. *Trends in Cognitive Sciences*, *1*(9), 346–352. https://doi.org/10.1016/S1364-6613(97)01105-4

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & von der Heydt, R. (2012). A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure–ground organization. *Psychological Bulletin*, *138*(6), 1172–1217. https://doi.org/10.1037/a0029333

Wagemans, J., Van Gool, L., & d'Ydewalle, G. (1992). Orientational Effects and Component Processes in Symmetry Detection. *The Quarterly Journal of Experimental Psychology Section A*, *44*(3), 475–508. https://doi.org/10.1080/14640749208401295

Wagemans, J., Van Gool, L., Swinnen, V., & Van Horebeek, J. (1993). Higher-order structure in regularity detection. *Vision Research, 33,* 1067-1088. https://doi.org/10.1016/0042-6989(93)90241-N

Wenderoth, P. (1994). The Salience of Vertical Symmetry. *Perception*, *23*(2), 221–236. https://doi.org/10.1068/p230221

Wertheimer, M. (1923). Laws of organization in perceptual forms. *A source book of Gestalt Psychology.*

Wertheimer, M. (1922). Untersuchungen zur Lehre von der Gestalt. *Psychologische forschung, 1*(1), 47-58.

Wilder, J., Rezanejad, M., Dickinson, S., Siddiqi, K., Jepson, A., & Walther, D. B. (2019). Local contour symmetry facilitates scene categorization. *Cognition*, *182*, 307–317. https://doi.org/10.1016/j.cognition.2018.09.014

Yue, X., Vessel, E. A., & Biederman, I. (2007). The neural basis of scene preferences: *NeuroReport*, *18*(6), 525–529. https://doi.org/10.1097/WNR.0b013e328091c1f9

Appendix A

Univariate Analyses

Here, we include three separate univariate regressions, predicting each of the subjective rating variables from the local and global symmetry measures (see Tables A1, A2, and A3). While all regressions are significant, we opted for a multivariate statistical analysis (i.e., canonical correlation) as our main approach, due to the high correlations between the three subjective rating variables (see Table A4).

Table A1: Regression model predicting subjective complexity from symmetry measures. $F(12,707) = 20.36$, $p < 0.0001$, adjusted R^2 of 0.24.

	coefficient	*std. error*	*t*	*p*
(Intercept)	-302.0	202.0	-1.49	0.14
Ribbon: Median	-50.74	12.97	-3.91	<0.0001***
Ribbon: Proportion Symmetric	360.4	203.2	1.77	0.08
Ribbon: Proportion Asymmetric	333.8	205.0	1.63	0.10
Separation: Median	-5.27	2.18	-2.42	<0.05*
Separation: Proportion Far	-2.37	2.18	-1.10	0.27
Separation: Proportion Near	2.41	10.30	2.34	<0.05*
Vertical Symmetry Peak	-0.83	0.15	-5.63	<0.0001***
Vertical Symmetry Range	0.38	0.008	4.67	<0.0001***
# Peaks Vertical Symmetry	-0.004	0.002	-2.60	<0.001**
Horizontal Symmetry Peak	-0.28	0.03	-1.71	0.09
Horizontal Symmetry Range	-0.13	0.08	-1.72	0.09
# Peaks Horizontal Symmetry	-0.001	0.002	-7.07	<0.0001***

Table A2: Regression model predicting subjective aesthetic pleasure from symmetry measures. $F(12,707) = 16.32$, $p < 0.0001$, adjusted R^2 of 0.20.

	coefficient	*std. error*	*t*	*p*
(Intercept)	-191.3	218.7	-0.88	0.38
Ribbon: Median	-17.28	14.02	-1.23	0.21
Ribbon: Proportion Symmetric	215.2	219.7	0.98	0.32
Ribbon: Proportion Asymmetric	198.1	221.7	0.89	0.37
Separation: Median	-4.21	2.36	-1.79	0.07
Separation: Proportion Far	-1.94	2.32	-0.84	0.40
Separation: Proportion Near	18.43	11.14	1.66	0.09
Vertical Symmetry Peak	-0.98	0.16	-6.16	<0.0001***
Vertical Symmetry Range	0.39	0.09	4.51	<0.0001***
# Peaks Vertical Symmetry	-0.008	0.002	-4.33	<0.0001***
Horizontal Symmetry Peak	-0.10	0.17	0.57	0.57
Horizontal Symmetry Range	-0.12	0.08	-1.48	0.14
# Peaks Horizontal Symmetry	-0.01	0.002	-6.36	<0.0001***

Table A3: Regression model predicting subjective interest from symmetry measures. $F(12,707) = 16.93$, $p < 0.0001$, adjusted R^2 of 0.21.

	coefficient	*std. error*	*t*	*p*
(Intercept)	-240.7	217.0	-1.11	0.27
Ribbon: Median	-27.34	13.91	-1.97	<0.05*
Ribbon: Proportion Symmetric	275.6	218.0	1.26	0.21
Ribbon: Proportion Asymmetric	252.0	219.9	1.15	0.25
Separation: Median	-4.77	2.34	-2.04	<0.05*
Separation: Proportion Far	-2.35	2.30	-1.02	0.31
Separation: Proportion Near	25.34	11.05	2.29	<0.05*
Vertical Symmetry Peak	-0.95	15.73	-6.01	<0.0001***
Vertical Symmetry Range	0.38	0.09	4.33	<0.0001***
# Peaks Vertical Symmetry	-0.007	0.002	-3.79	<0.001**
Horizontal Symmetry Peak	-0.009	0.17	-0.05	0.96
Horizontal Symmetry Range	-0.17	0.08	-2.04	<0.05*
# Peaks Horizontal Symmetry	-0.01	0.002	-6.63	<0.0001***

Table A4: Bivariate correlation matrix.

		1.	2.	3.	4.	5.	6.	7.	8.	9.	10.	11.	12.	13.	14.	15.	16.	17.
Subjective Ratings	1. Pleasure	–																
	2. Complexity	**.89**	–															
	3. Interest	**.98**	**.94**	–														
Local Symmetry	4. Ribbon: Median	-.05	**-.20**	-.07	–													
	5. Ribbon: Proportion Symmetric	.002	**-.08**	.001	**.80**	–												
	6. Ribbon: Proportion Asymmetric	-.001	**.08**	-.001	**-.79**	**-.99**	–											
	7. Separation: Median	-.07	**-.18**	**-.07**	**.69**	**.72**	**-.71**	–										
	8. Separation: Proportion Far	-.03	**-.15**	-.04	**.85**	**.89**	**-.89**	**.90**	–									
	9. Separation: Proportion Near	-.03	**-.09**	-.04	-.05	**-.45**	**.45**	**.15**	**-.06**	–								
Global Symmetry	10. Vertical Symmetry Peak	**-.19**	**-.28**	**-.20**	**.34**	**.34**	**-.34**	**.40**	**.44**	**.10**	–							
	11. Vertical Symmetry Range	**.21**	**.27**	**.22**	**-.29**	**-.29**	**.29**	**-.32**	**-.36**	**-.10**	**-.58**	–						
	12. # Peaks Vertical Symmetry	-.03	**.09**	-.01	**-.44**	**-.49**	**.48**	**-.51**	**-.59**	**-.08**	**-.75**	**.57**	–					
	13. Horizontal Symmetry Peak	**.14**	.008	**.11**	**.26**	**.21**	**-.21**	**.21**	**.27**	**.08**	**.23**	-.01	**-.31**	–				
	14. Horizontal Symmetry Range	**-.12**	**-.23**	**-.15**	**.47**	**.32**	**-.31**	**.37**	**.43**	**.14**	**.45**	**-.48**	**-.47**	**.30**	–			
	15. # Peaks Horizontal Symmetry	**-.25**	**-.16**	**-.23**	**-.30**	**-.34**	**.34**	**-.38**	**-.45**	**-.11**	**-.38**	**.27**	**.60**	**-.47**	**-.34**	–		
Controls	16. Simplicity (ZIP compression)	**-.33**	**-.48**	**-.36**	**.57**	**.41**	**-.41**	**.55**	**.59**	**.27**	**.60**	**-.53**	**-.58**	**.28**	**.54**	**-.37**	–	
	17. Self-similarity	**-.25**	**-.15**	**-.23**	**-.51**	**-.51**	**.51**	**-.57**	**-.61**	**-.07**	**-.19**	**.14**	**.48**	**-.27**	**-.31**	**.47**	**-.38**	–

Note: all significant correlations ($p < 0.05$) are bolded.

Appendix B

Control Analysis

Previous work investigating the influence of visual features on aesthetic pleasure of landscape scenes found that computational simplicity (i.e., ZIP compression) and fractal self-similarity, but not symmetry, significantly predicted aesthetic pleasure (Mayer & Landwehr, 2018). Thus, we performed a control analysis, including simplicity and self-similarity in our model to see how they relate to our symmetry measures and the subjective ratings.

We first extracted simplicity and self-similarity using the image-fluency package in R. We then ran another canonical correlation analysis, this time including simplicity and self-similarity in the predictor variable set along with our symmetry measures. The control model was significant, Wilks' $\lambda = 0.428$, $F(42, 2086.20) = 16.44$, $p < 0.0001$, explaining approximately 57.2% of the variance shared between variable sets. The critical question is, what happened to the relationship between the symmetry variables and the subjective ratings when the control variables were added.

Table B1 shows the variables that significantly contribute to the model in root 1 and root 2. We can see that, in root 1, the simplicity measure is the only independent variable that contributes to the model. The effect of the symmetry variables has disappeared. This result suggests that the variance explained by the symmetry variables in the original model is contained in the variance explained by the simplicity variable. To test this, we ran a multiple linear regression analysis, predicting the simplicity measure on each image from the symmetry values. The multiple linear regression analysis yielded a significant regression equation, $F(12,707) = 98.91$, $p < 0.0001$, with an adjusted R^2 of 0.62. The significant predictors of simplicity were the symmetry values that had significantly contributed to the original canonical correlation model.

Taken together, these results show that the simplicity measure can be explained in large part by the symmetry measures, indicating that symmetry is indeed significantly related to subjective ratings of complexity, aesthetic pleasure, and interest.

Table B1: Canonical solution of control analysis for features predicting ratings for root 1 and 2.

	Root 1			*Root 2*			
Variable	*Coef*	r_s	$r_s^2(\%)$	*Coef*	r_s	$r_s^2(\%)$	$h^2(\%)$
Ribbon: Median	-0.244	-0.370	13.71	-0.814	-0.691	47.73	61.44
Ribbon: Proportion Symmetric	9.308	-0.182	3.30	9.930	-0.473	22.40	25.70
Ribbon: Proportion Asymmetric	8.924	0.182	3.30	9.455	0.472	22.29	25.59
Separation: Median	-0.465	-0.331	10.93	-0.014	-0.545	29.75	40.68
Separation: Proportion Far	0.176	-0.295	8.68	-0.035	-0.633	40.12	48.80
Separation: Proportion Near	0.268	-0.152	2.31	-0.068	-0.235	5.51	7.82
Vertical Symmetry Peak	-0.019	-0.442	19.50	0.287	-0.280	7.85	27.35
Vertical Symmetry Range	0.065	0.431	18.56	-0.033	0.157	2.50	21.06
# Peaks Vertical Symmetry	-0.026	0.195	3.82	0.588	0.628	39.42	43.24
Horizontal Symmetry Peak	-0.054	-0.032	0.10	-0.405	-0.662	43.88	43.98
Horizontal Symmetry Range	-0.028	-0.371	13.77	0.124	-0.387	14.99	28.76
# Peaks Horizontal Symmetry	-0.385	-0.185	3.41	0.105	0.622	38.78	42.19
Simplicity (ZIP compression)	-0.956	-0.767	58.90	0.169	-0.446	19.87	78.77
Self-similarity	-0.425	-0.163	2.67	0.197	0.678	46.00	48.67
Pleasure	0.539	0.844	71.25	-0.561	-0.534	28.49	99.47
Complexity	1.457	0.984	96.96	2.411	-0.122	1.48	98.44
Interest	-1.016	0.876	76.72	-2.170	-0.458	20.97	97.69

Note: Structure coefficients (r_s) greater than |0.45| are underlined. Communality coefficients (h^2) greater than 35% are underlined. Coef = standardized canonical function coefficient; r_s = structure coefficient; r_s^2 = squared structure coefficient; h^2 = communality coefficient.

Table B2: Regression model predicting simplicity from symmetry measures.

	coefficient	*std. error*	*t*	*p*
(Intercept)	-17.93	42.0	-0.43	0.67
Ribbon: Median	14.05	2.7	5.22	<0.0001***
Ribbon: Proportion Symmetric	2.89	42.1	0.07	0.95
Ribbon: Proportion Asymmetric	8.14	42.5	0.19	0.85
Separation: Median	-0.33	0.45	-0.73	0.47
Separation: Proportion Far	1.65	0.45	3.71	<0.0001***
Separation: Proportion Near	3.26	2.14	1.53	0.13
Vertical Symmetry Peak	0.22	0.03	7.09	<0.0001***
Vertical Symmetry Range	-0.10	0.02	-6.04	<0.0001***
# Peaks Vertical Symmetry	-0.00003	0.0003	-0.10	0.92
Horizontal Symmetry Peak	0.08	0.03	2.53	<0.05*
Horizontal Symmetry Range	0.05	0.02	2.97	<0.01**
# Peaks Horizontal Symmetry	0.0003	0.0004	0.74	0.46