Image Symmetries: The Right Balance between Evenness and Perception

Fabrizio Guerrini*, Alessandro Gnutti*, Riccardo Leonardi*

*Department of Information Engineering, Univeristy of Brescia, Italy

{*fabrizio.guerrini, a.gnutti006, riccardo.leonardi*}*@unibs.it*

Abstract—A recent and fascinating interest in computational symmetry for computer vision and computer graphics applications has led to a remarkable realization of new symmetry detection algorithms. Such a concern is culminated in a symmetry detection competition as a workshop affiliated with the 2011 and 2013 CVPR Conferences. In this paper, we propose a method based on the computation of the symmetry level associated to each pixel. Such a value is determined through the energy balance of the even/odd decomposition of a patch with respect to a central axis (which is equivalent to estimate the middle point of a row-wise convolution). Peaks localization along the perpendicular direction of each angle allows to identify possible symmetry axes. The evaluation of a feature based on gradient information allows to establish a classification confidence for each detected axis. By adopting the aforementioned rigorous validation framework, the proposed method indicates significant performance increase.

Index Terms—Symmetry detection, even-odd decomposition, gradient image analysis, object detection.

I. INTRODUCTION

Local and global symmetries in the digital world represent an important characteristic in computer vision and machine intelligence, since they can simplify tasks such as object detection and visual search [1]. They could also represent an unexplored relevant feature for image compression, where most recent interest has been devoted to manage increased resolution and dynamic range [2] [3]. Indeed, symmetry characteristics has already been suggested to minimize information redundancy [4], or to improve source modeling [5].

Symmetry detection has been mostly focused on bilateral reflection [4]. Early methods were simple but sensitive to noise [6], [7]. In 2011 and 2013 the IEEE CVPR conference proposed a symmetry detection competition on two scenarios: single and multiple symmetry axes. In the most recent instance [8], two methods stood out. In [9], which obtained the best overall performance, SIFT-like descriptors on constellations of interest points were used, while [10] employed SIFT descriptors with gradient-based weighting and a principled statistical procedure, achieving marginally better performance for high recall. In [11], however, it is argued that those results are still not yet satisfactory. In [12] a robust and efficient method has been proposed, but it requires to know the number of symmetry axes.

This paper describes a new algorithm for the detection of symmetric objects in natural images. When symmetric objects detected by an algorithm are compared to the answers given by a human being, it is imperative to consider how the human brain works. The performance of the aforementioned algorithms can be hampered by not taking into account human perception, since it is impossible to neglect those perception mechanisms employed by the human brain when engaging in object recognition tasks [13] [14]. In particular, symmetry detection is just one of the many tool employed during object recognition by humans [15] [16], although it is still unclear how and when symmetry comes into play [17]. Of course, symmetry is a very simple stimulus and it appears that a subconscious process is able to directly detect symmetric patterns in retinal frontoparallel images using some kind of 2D correlation [18]. In addition, 2D symmetric planar views help the human brain to learn 3D shapes [19] and conversely 3D contextual information allows to compensate for the viewing angle before actually detecting 2D frontoparallel symmetry [16]. Also, there is evidence that the brain may also perform residual symmetry detection in oblique views as well [20]. Furthermore, color seems to slow down symmetry detection because color channels are likely analyzed separately even when symmetry is present [13] [20]. In the end, there is the distinct need to complement a mathematic tool dealing with symmetry detection in the raw data with processing akin to what the brain cortex or frontoparallel retinal images do to bridge the semantic gap between mathematical symmetry and perceived symmetry.

As argued above, using just raw symmetry can be tricky to identify symmetric objects, either because a perceived symmetric object may not possess a real symmetry in the data (due to illumination changes, partial occlusions, and noise of other types) or because very symmetric data may correspond to almost uniform background patterns. Trying to mimic basic processing happening in the human brain, spatial correlation and relevant gradient information are thus exploited to help raw data-driven symmetry detection. Gradient information is surely processed by the brain to detect object edges [21] and as such has been largely proposed in the past for appearance-based methods for image segmentation [22]. Here, the gradient image is used to complement raw data symmetry detection, greatly boosting its performance.

The rest of the paper is organized as follows. In Section II the problem of detecting reflection symmetries in a 1D digital sequence is described by decomposing it using the even/odd decomposition and then by comparing the energies of the output sequences. This is partly inspired by [23] that concerned the search for hierarchies of symmetries in 1D raw

data. The method is then expanded to take into account the 2D, natural images domain. In Section III, the details of the symmetry detection algorithm, which also incorporates the gradient information, are given. In Section IV the algorithm is benchmarked using the 2013 CVPR competition dataset, and it is shown that it outperforms all of those reported in [8]. The paper concludes with Section V.

II. REFLECTION SYMMETRY DETECTION IN NATURAL IMAGES

The even/odd decomposition of a real, 1D discrete sequence $x[n]$ around a candidate symmetry point m is given by:

$$x_e[n; m] = \frac{x[n] + x[2m - n]}{2} \quad (1)$$

$$x_o[n; m] = \frac{x[n] - x[2m - n]}{2} \quad (2)$$

$$x[n] = x_e[n; m] + x_o[n; m] \quad (3)$$

As detailed in [23], the energy E_e of the even part w.r.t. the energy E of $x[n]$ can be computed as:

$$E_e(m) = \frac{1}{2}E + (x * x)[2m] \quad (4)$$

where $(x * x)$ represents the convolution of the discrete sequence with itself, *i.e.* the auto-convolution of $x[n]$. When $x[n]$ is prevalently reflection symmetric w.r.t. m (namely, even), $E_e(m)$ is more than half of E. Therefore, local maxima of the auto-convolution, which are also maxima for E_e discounting for the factor 2, correspond to the best candidate reflection symmetry points m w.r.t. neighboring positions.

However, consider the case where $x[n]$ possesses a clear local symmetry of modest energy whereas outside the symmetry support the signal is non-symmetric and of somewhat higher energy. In this scenario, the non-symmetric part affects the position of the (globally) optimal symmetry point, whose desired position is in the center of the local symmetry. This happens because the computation of the energy of the even part, namely the auto-convolution, also takes into account the non-symmetric signal with non negligible energy and the local symmetry becomes thus non detectable.

The proposed solution comprises two refinements to the strategy above. The first is to confine the convolution computation on a window W of size $2n_p + 1$ centered around the candidate symmetry point. The windowed auto-convolution is no more dependent on the parts of the signal outside the window support and hence this allows to precisely locate the position of the local symmetry if the window support is (at least approximately) correct. On the other hand, non-symmetric portions of the signal may still attain a high auto-convolution value due to the sheer energy of the windowed signal. Therefore, the second refinement is to normalize the computed auto-convolution by the energy of the windowed sequence. This way, it is possible to compare found symmetries in windowed signals with different energy.

In the end, the best local symmetries in the 1D sequence is found by sliding the window W at every position and then compute a measure of symmetry S associated to the center of the window, as follows:

$$S(W) = \frac{\sum_{n=-n_p}^{n_p} x[n] \cdot x[-n]}{\sum_{n=-n_p}^{n_p} |x[n]|^2} \quad (5)$$

The measure in Eq. (5) is the normalized, windowed auto-convolution computed on the window W centered around a candidate symmetry point, which is in effect equivalent to the normalized inner product between the windowed sequence and its flipped version.

Moving on to the problem of detecting symmetric object in 2D images, some additional care is needed to apply the insights given by the measure $S(W)$. For now, let us consider the case of an horizontally symmetric object such that the sought symmetry axis is vertical. If $S(W)$ is applied row-wise, the noisy image data can delete or misplace the 1D local symmetry in some rows and therefore it can happen that in such rows there is no local maximum correspondent to the vertical symmetry axis, making it impossible to simply connect the peaks in the vertical direction to reveal the axis. The human brain, though, still extends the symmetry axis along the perceived symmetry independently of the absence of the 1D symmetry in some rows because it uses the global symmetry information.

Consequently, instead of computing $S(W)$ separately for each row, it is better to compute the row-wise, windowed and normalized auto-convolution over a 2D square patch P all at once, as follows:

$$S_2(P) = \frac{\sum_{m=-m_p}^{m_p} \sum_{n=-n_p}^{n_p} x[m, n] \cdot x[m, -n]}{\sum_{m=-n_p}^{n_p} \sum_{n=-n_p}^{n_p} |x[m, n]|^2} \quad (6)$$

Using Eq. (6), the 2D spatial correlation information is properly used to smooth displaced symmetries in individual rows. Computing the symmetry on the position at the center of P allows to extend the symmetry axis all the way even in the presence of noisy data, such as possible variations in the object texture. As detailed in Section III, by connecting local maxima in Eq. (6), candidate symmetry axes will be derived.

The human brain capabilities for detecting symmetries go beyond the use of a 2D context, as introduced in Section I. To at least partially mimic the mechanisms underlying object recognition, the proposed algorithm also incorporate the information provided by the gradient image to segment the symmetric object out of the background and distinguish it from non-symmetric objects. In particular, in Section III it is proposed to compare the evenness of the gradient magnitude in the same 2D patch P used above to detect the presence of a symmetric object. Even if the image segmentation process performed by the brain is much more complex, the proposed simple approximation still leads to a significant performance increase and create correspondences between detected and perceived symmetry.

III. PROPOSED ALGORITHM DETAILS

This Section describes the algorithm proposed for the symmetric object detection, including only those details that we

deem essential to properly understand the argument that we have put forward in previous Sections to improve the correct localization and verification of detected symmetry segments.

The objective of the first part of the algorithm, as described in the previous Section, is to attribute a symmetry matrix whose values are associated to the pixels of the image I, for any given angle. We store that information in a 3-D stack. As shown in Fig. 1, such a process is composed of three main tasks. In the first one, we rotate I by n different angles α_i that are referred to as I_{α_i}. The symmetry value is calculated separately for each angle, so that all possible symmetries can be captured independently of their tilt. I_{α_i} will return symmetries in the α_i direction ($0 \leq \alpha_i < 180$, with $\alpha_i = 180°i/n$).

As described in Section II, the auto-convolution operation applied to a sequence allows to locate the position of the optimal symmetry axis. To improve the location up to half pixel accuracy, a columns interpolation by a factor of 2 is performed on I_{α_i}. So, at the end of the first block of Fig. 1, n images are formed, that we refer to as $\tilde{I}_{\alpha_i}$, representing rotated and interpolated versions of the original image I.

In the map computation process, the level of symmetry of each pixel p of $\tilde{I}_{\alpha_i}$ is computed. First, a 2-D patch P of fixed size $(2n_p + 1)$ and centered in p is extracted. Then, the symmetry measure $S_2(P)$ is computed according to Eq. (6), that takes values in the $[-1, 1]$ interval. This is strictly related to the ratio between the energy of the signal in the patch and its even-odd decomposition around its middle vertical axis. In particular:

$$S_2(P) = 2 \cdot \frac{E_e}{E} - 1 = 1 - 2 \cdot \frac{E_o}{E} \qquad (7)$$

where E is the energy of the signal on the patch, E_e is the energy of the even part of the patch and E_o is the energy of the odd part of the patch.

The map computation is performed for all images $\tilde{I}_{\alpha_i}$, so that n maps M_{α_i} are generated. Fig. 2b shows the computed map for $\alpha = 0°$ (vertical symmetry axis) superimposed on the test image shown in Fig. 2a (we scaled down the map to match the horizontal sizes). The borders in Fig. 2a are discarded in Fig. 2b since the patches fall outside of the image boundaries. Observe how, in addition to two clear vertical axes where the objects are, the uniform background also exhibits a significant symmetry response. Last, a 3-D stack of n maps, in which the
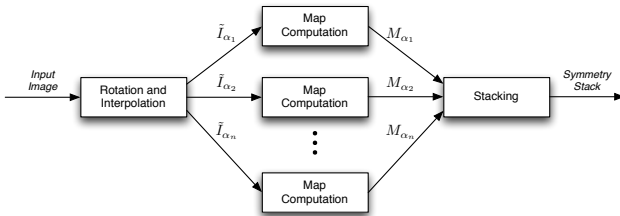
i-th slice is the map M_{α_i}, is constructed. We refer to such a stack as "symmetry stack" since it represents the local (pixel) information about the level of symmetry.

At this point all symmetry axes present in the image I can be identified by processing the symmetry stack. In the end, all candidate symmetry axes are extracted and mapped back to the original image domain. The process is constituted by four distinct operations (see Fig. 3). First, a straightforward (half-wave) rectification is performed, by setting to 0 all negative values of the stack (we did that in Fig. 2 too). Following the discussion we have carried so far, one could assume that symmetry measures close to 1 identify stronger symmetries and should be thus sufficient to identify symmetric objects. On the contrary, we observed experimentally that this is not necessarily true. For example, large values can be associated to uniform background regions that correspond to highly symmetric structures. Moreover, objects that are perceived by humans as clearly symmetric could have lower values due to shadows, illumination changes, low resolution, etc. The use of 2-D patches may only alleviate the problem.

Instead of just taking the symmetry measure, a key observation is that searching for the peaks (local maxima) of the symmetry map is more relevant. So, the absolute value of the coefficient is not as important as its relationship with respect to its neighbors. All local maxima along the rows of every symmetry map can be associated to a specific direction α_i. A connectivity analysis between such maxima can be performed by means of a flood-fill algorithm. Consequently, all possible reflection symmetries define a series of segments linking connected local maxima existing in each row.

In order to project back the symmetry information onto the original image coordinate system, each map M_{α_i} is first horizontally scaled down by a factor 2 to be consistent with the size of the input image I and then rotated by $-\alpha_i$.

Even with the expedients described so far to improve the detection accuracy of data symmetries, such as the search for

(a) The test image, with two clearly symmetric objects.

(b) The map $M_{0°}$, associated to vertical symmetry axes.

Fig. 2: The symmetry computation stage for a test image for $\alpha = 0°$.

Fig. 3: Main blocks of the candidate axes identification process, starting from the symmetry stack.

Fig. 1: Main processing stages involved in the symmetry stack computation.

peaks in the orthogonal direction, the candidate symmetry axes still may correspond to a lot of false positives, in the sense that the detected symmetries are not necessarily associated to symmetric objects. In an effort to correct this, we employ a kind of processing similar to what our brain does when it performs object recognition.

The idea is to favor those axis whose edge information around them is as specular as possible, hinting at the fact that the axis is running through the center of a symmetric object. Fig. 4 shows a simple example with just two of the candidate symmetry axes. As it turns out, these are the first two answers in terms of $S_2(P)$, both axes pertaining to highly symmetric data, but of course the red one is perceived as corresponding to a symmetric object, while the blue one is detected because of the particular texture in that image area. To favor the former, first we compute the gradient image using a standard Sobel operator and take the gradient magnitude. Then, we use a threshold M on the percentage of most significant gradient magnitude values applied on the whole image that we retain whereas the least significant are set to 0. By setting M we obtain the desired gradient magnitude image $G[m, n]$.

For each axis, a patch P_A is extracted that is as wide as the patches P used for the symmetry measure computation (namely n_p) but in this case it is extended vertically to cover all the axis locations (recall that for each axis the image is rotated so as to have the considered axis in the vertical direction). On P_A, we again perform the even-odd decomposition and evaluate:

$$G_2(P_A) = \frac{\sum_{m=m_1}^{m_2} \sum_{n=-n_p}^{n_p} G[m, n_1 + n] \cdot G[m, n_1 - n]}{\sum_{m=m_1}^{m_2} \sum_{n=-n_p}^{n_p} |G[m, n_1 + n]|^2} \quad (8)$$

where the axis is detected in the n_1-th column and in the $[m_1, m_2]$ rows interval. If the gradient magnitude is approximately specular around the axis, $G_2(P)$ takes on values close to 1, whereas if no gradient is significant its value is around 0. So, the candidate axes are finally sorted according to their G_2 values.

IV. EXPERIMENTAL RESULTS

In this Section, we report the results of the algorithm we presented using the dataset of the 2013 CVPR competition on symmetry detection [8]. To obtain these results, we proceeded as follows. First, we computed a set of candidate symmetry axes using the symmetry measure $S_2(P)$, averaged on all the positions belonging to the axis, as the discriminant. In particular, we kept as candidate symmetry axes only those whose mean symmetry measure $S_2(P)$ is above a threshold $T = 0.2$ (remember that the symmetry measure is in the $[-1, 1]$ interval, with value 1 identifying perfectly symmetric data). Then, these candidate axes are sorted according to their G_2 values. The threshold M is set to 0.3.

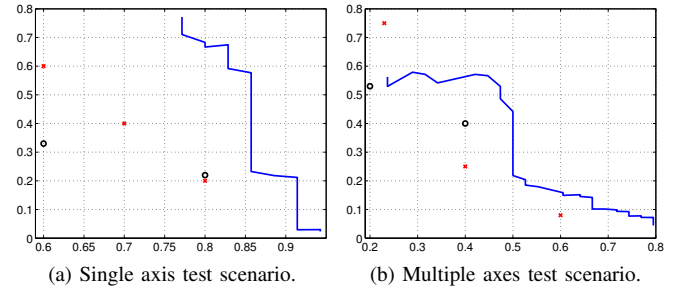(a) Single axis test scenario. (b) Multiple axes test scenario.

Fig. 5: Precision/recall curve on the dataset (blue line). We have also drawn selected values of the other considered techniques: [9] with red crosses and [10] with black circles.
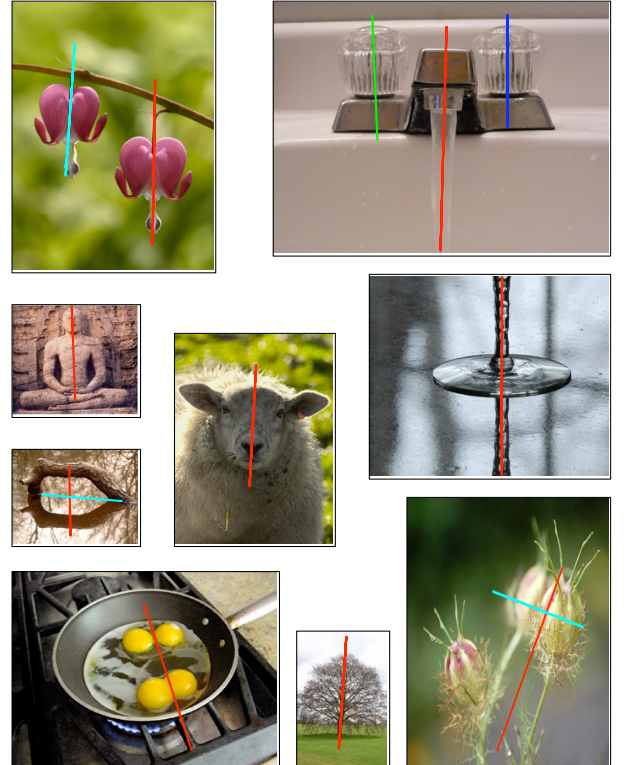

Fig. 6: Collage of visual results on selected images taken from the 2013 CVPR competition dataset.

(a) The original example image, with two candidate symmetry axes shown as well.

(b) The gradient magnitude of the image in Fig. 4a. The threshold in this case is $M = 0.3$.

Fig. 4: A visual example of the gradient image processing that we employed in this work.

Fig. 5 depicts the precision/recall obtained on the provided dataset, for both the single axis and multiple axes detection scenarios. Such a curve shows very good performance. To put our results in perspective, we have also drawn some selected precision/recall values obtained from [8] for the best two techniques previously described [9] [10]. In particular it is clear how the proposed algorithm outperforms the other baseline methods.

Last, Fig. 6 shows a collage of detected symmetries, considering $G_2 > 0.85$, superimposed on the actual images, taken from the provided dataset, that we think best represent the success and the limits of our algorithm. Most of them show how accurately existing symmetries are detected. We also show a particularly interesting example, the bottom right image of Fig. 6, where our algorithm mis-detect the symmetric objects twice. For the red axis, the problem is that the background divides two identical objects. The difference with what our brain does here is the fact that we perceive as objects the peripheral symmetric area and not the central one. Edge information is in this case insufficient to avoid this occurrence. For the blue axis instead, the detection mixes foreground and background objects into a single object. Again, our brain is efficiently able to extract the foreground using 3-D processing, something that is not consider yet in the method described in this paper.

V. Conclusions

In this paper we presented a new method for reflection symmetric objects detection. The proposed algorithm complements data analysis aimed at the search for structural symmetries with arguments taken from object recognition processing done by the human brain. First, a symmetry value for each pixel and each angle is calculated to form a stack of symmetry maps, through the computation of the energy-normalized, 2D row-wise auto-convolution of the pixels in a 2D patch. The patch is centered on the considered pixel of the suitably rotated image. The candidate symmetry points are found by looking for local maxima in the orthogonal direction and not by simply thresholding the values in the symmetry maps. Such candidate points are then connected through a morphological operator, constructing candidate symmetry axes.

Then, these axes are validated using the magnitude of the image gradient. In a manner similar to what our brain unconsciously does, the specular character of the gradient magnitude in the vicinity of the axis is evaluated, indicating whether the considered symmetry axis runs through the center of a symmetric object. When the gradient information is taken into account to validate the axes derived from the purely mathematical symmetry analysis, the correspondence between the outcomes of our algorithm and a human-built ground truth is significantly improved.

The method has been experimentally tested using the framework of the 2013 CVPR symmetry detection competition, and we have shown how our algorithm outperform all of those reported in [8].

References

[1] N. Mitra, M. Pauly, M. Wand, and D. Ceylan, "Symmetry in 3d geometry: Extraction and applications," in *EUROGRAPHICS State-of-the-art Report*, 2012.

[2] V. Sze, M. Budagavi, and G. Sullivan, "High efficiency video coding (hevc)," *Integrated Circuit and Systems, Algorithms and Architectures. Springer*, pp. 1–375, 2014.

[3] A. Boschetti, N. Adami, R. Leonardi, and M. Okuda, "Flexible and effective high dynamic range image coding," in *Image Processing (ICIP), 2010 17th IEEE International Conference on*. IEEE, 2010, pp. 3145–3148.

[4] Y. Liu, H. Hel-Or, C. Kaplan, and L. Gool, "Computational symmetry in computer vision and computer graphics.," *Foundations and Trends in Computer Graphics and Vision*, vol. 5, no. 1-2, pp. 1–195, 2010.

[5] V. Sanchez, R. Abugharbieh, and P. Nasiopoulos, "Symmetry-based scalable lossless compression of 3d medical image data," *IEEE Transactions on Medical Imaging*, vol. 28, no. 7, pp. 1062–1072, 2009.

[6] M. Atallah, "On symmetry detection," *IEEE Trans. Computers*, vol. 34, no. 7, pp. 663–666, 1985.

[7] H. Blum and R. Nagel, "Shape description using weighted symmetric axis features," *Pattern Recognition*, vol. 10, no. 3, pp. 167–180, 1978.

[8] J. Liu, G. Slota, G. Zheng, Z. Wu, M. Park, S. Lee, I. Rauschert, and Y. Liu, "Symmetry detection from realworld images competition 2013: Summary and results," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2013, pp. 200–205.

[9] G. Loy and J. Eklundh, "Detecting symmetry and symmetric constellations of features.," in *ECCV (2)*, Ales Leonardis, Horst Bischof, and Axel Pinz, Eds. 2006, vol. 3952 of *Lecture Notes in Computer Science*, pp. 508–521, Springer.

[10] V. Patraucean, R. Grompone von Gioi, and M. Ovsjanikov, "Detection of mirror-symmetric image patches," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2013, pp. 211–216.

[11] C. Funk and Y. Liu, "Symmetry recaptcha," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 5165–5174.

[12] I. Atadjanov and L. Seungkyu, "Bilateral symmetry detection based on scale invariant structure feature," *Image Processing (ICIP), 2015 IEEE International Conference on*, pp. 3447 – 3451, September 2015.

[13] J. DiCarlo, D. Zoccolan, and N. Rust, "How does the brain solve visual object recognition?," *Neuron*, vol. 73, no. 3, pp. 415–434, 2012.

[14] K. Grill-Spector, Z. Kourtzi, and N. Kanwisher, "The lateral occipital complex and its role in object recognition," *Vision research*, vol. 41, no. 10, pp. 1409–1422, 2001.

[15] F. Schmidt and T. Schmidt, "Rapid processing of closure and viewpoint-invariant symmetry: behavioral criteria for feedforward processing," *Psychological research*, vol. 78, no. 1, pp. 37–54, 2014.

[16] P. Van der Helm, *Simplicity in vision: A multidisciplinary account of perceptual organization*, Cambridge University Press, 2014.

[17] S. Palmer, J. Brooks, and R. Nelson, "When does grouping happen?," *Acta psychologica*, vol. 114, no. 3, pp. 311–330, 2003.

[18] G. van der Vloed, A. Csatho, and P. van der Helm, "Symmetry and repetition in perspective," *Acta psychologica*, vol. 120, no. 1, pp. 74–92, 2005.

[19] T. Vetter, T. Poggio, and H. Bulthoff, "The importance of symmetry and virtual views in three-dimensional object recognition," *Current Biology*, vol. 4, no. 1, pp. 18–23, 1994.

[20] A. Makin, G. Rampone, and M. Bertamini, "Conditions for view invariance in the neural response to visual symmetry," *Psychophysiology*, vol. 52, no. 4, pp. 532–543, 2015.

[21] D. Marr and E. Hildreth, "Theory of edge detection," *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 207, no. 1167, pp. 187–217, 1980.

[22] J.-M. Morel and S. Solimini, *Variational methods in image segmentation: with seven image processing experiments*, vol. 14, Springer Science & Business Media, 2012.

[23] A. Gnutti, F. Guerrini, and R. Leonardi, "Representation of signals by local symmetry decomposition," in *Proc. of the 23rd European Signal Processing Conference (EUSIPCO '15)*, 2015, pp. 983–987.