

Detection of Mirror-Symmetric Image Patches

Viorica Pătrăucean
 École Polytechnique Palaiseau, France
 Military Technical Academy of Bucharest, Romania
 patraucean@lix.polytechnique.fr

Rafael Grompone von Gioi
 ENS Cachan, France
 grompone@cmla.ens-cachan.fr

Maks Ovsjanikov
 École Polytechnique Palaiseau, France
 maks@lix.polytechnique.fr

Abstract

We propose a novel approach for detecting partial reflectional symmetry in images. Our method consists of two principal stages: candidate selection and validation. In the first step, candidates for mirror-symmetric patches are identified using an existing heuristic procedure based on Hough voting. The candidates are then validated using a principled statistical procedure inspired from the a contrario theory, which minimizes the number of false positives. Our algorithm uses integral image properties to enhance the execution time.

1. Introduction

Symmetry cues play an important role in human perception for detecting and recognizing objects, as stressed by numerous psychophysical studies (e.g. [2]). Symmetry can provide global information about the structure of an object that is otherwise difficult to capture. Inspired by this finding, recent works in computer vision show considerable improvement when symmetry information is incorporated into automatic methods for object classification and recognition, by using appropriate symmetry descriptors [20, 11].

The problem of classifying various types of symmetries is well studied from a theoretical point of view [25], while the majority of methods for detection of symmetries in images have concentrated on mirror (reflectional), rotation and translation symmetries.

The most popular approaches for symmetry detection rely on voting procedures, e.g. Hough transform, carried out on symmetric pairs of feature points, like SIFT [14] or ASIFT [19] assuming either a (near-) orthogonal view of the symmetric objects [15], or trying to take into account the perspective skew induced in the image formation process [24, 4, 5].

A major weakness of these approaches is the absence of a validation procedure, as pointed out in [20], and highlighted by the results of CVPR11 contest on symmetry detection [1]. Most of the time, the candidates proposed by the voting procedure are accepted or rejected based on predefined detection thresholds. In the absence of a careful tuning, the result in terms of false detections (false positives and false negatives), is difficult to predict.

Choosing the right values for the detection thresholds is a generic issue in many computer vision detection tasks. Using hard thresholds is clearly not appropriate, since they cannot adapt to various noise levels, image resolution or content. Some works used different empirical approaches to compute threshold values that behave satisfactory in certain predefined scenarios [3, 17]. This improved the performance to some extent, but still one cannot predict threshold values for all possible image and object sizes.

A more principled approach to resolve this issue has been proposed by the *a contrario* theory [7], which formalises the *non-accidentalness principle* [13]. Informally, the principle states that there should be no detection in a random (white noise) image. Putting this very simple idea into an appropriate probabilistic framework has given uniformly good results in tasks including geometric feature detection [10, 21], or fundamental matrix estimation [18], as it allows to automatically compute self-tuning thresholds. The main goal of this paper is to give an *a contrario* formulation for the mirror-symmetry detection problem in images taken from orthogonal views.

Our algorithm focuses on the minimization of the number of false detections, and consists of two major stages. The first phase – *candidate selection* – can be carried out using any of the existing symmetry detectors, and the main concern is to avoid the false negatives. For efficiency reasons, we use the algorithm proposed by Loy and Eklundh [15], that we will denote by *baseline algorithm*. The second

step – *a contrario validation* – aims at reliably filtering out the false positives, without the need for parameter tuning.

The next two sections detail the underlying theory of the validation step and the main implementation aspects of the whole pipeline. Section 4 illustrates the advantages of the proposed validation procedure on several representative detections. Finally, section 5 outlines the directions of future improvement.

2. A *Contrario* Validation of Symmetric Image Patches

In the following, we denote by I a gray-scale image of $m \times n$ pixels. A symmetry axis is represented in polar coordinates by (ρ, θ) , with ρ being the distance between the axis and the origin, and θ the angle of the vector from the origin to the closest point on the axis. For simplicity, a mirror-symmetric patch is represented by its bounding box, defined by the limits of the symmetry axis and the width.

We first remind the generic *a contrario* framework, and then we adapt it to the symmetry detection problem.

Following the non-accidentalness principle, the *a contrario* validation starts by defining an unstructured (random) model, and then it accepts as valid detections only candidates that represent unexpected deviations (outliers) from this model.

The reasoning is similar to the multiple hypothesis testing formulation [8], where the unstructured model corresponds to the *null hypothesis*; we denote it by \mathcal{H}_0 . Let X be a random image drawn under \mathcal{H}_0 , having the same size as the analyzed image I . Given a candidate s , we define a function $k(s)$ to compute its *test statistic*, on which we evaluate the probability that s appeared in X . If the *p-value* of the test s is smaller than a predefined significance level α — *i.e.* the probability of observing in the random model a candidate with a test statistic at least as extreme as $k(s)$ is less than α — then there exists enough evidence to reject the null hypothesis and declare the candidate meaningful. With this choice, the probability of accepting one false positive is smaller than α . In other words, the number of false positives is controlled by α , which can be set as small as desired. Common values are $\alpha = 0.01$ or $\alpha = 0.05$. However, one needs to consider that in an image I there are in fact N_I candidates to test. In this case, the control of the number of false positives at the α level is ensured only if each test s_i rejects the null hypothesis whenever the p-value of the test is smaller than α/N_I , *i.e.* $\mathbb{P}_{\mathcal{H}_0}[k_X(s_i) \leq k_I(s_i)] \leq \alpha/N_I$ (Refer to [7, p. 71] or [22, p. 45] for the proof of this result).

By definition, the *number of false alarms (NFA)* of a candidate s_i is given by $NFA(s_i) := N_I \mathbb{P}_{\mathcal{H}_0}[k_X(s_i) \leq k_I(s_i)]$. The *NFA* of a candidate is a central notion in the *a contrario* theory, and can be interpreted as the overall expected number of false positives that we could observe when accepting candidates at least as extreme as s_i . With

this definition, the above inequality writes: $NFA(s_i) \leq \alpha$. Note that in this case, α up-bounds in fact an expectation, and its value is no longer restricted to $(0, 1)$. We denote by ε the significance level having this new meaning to avoid confusion. Thus a candidate is declared meaningful if its *NFA* is smaller than ε .

In most practical applications, the simple value $\varepsilon = 1$ appears to be suitable, so it can be set once and for all. With this choice, the expected number of false positives in a random image is guaranteed to be upper-bounded by 1. If the unstructured model is properly chosen, then this result, translated to the analyzed image I , ensures that the number of false positives, *i.e.* accidental detections reported in the background, is very small.

This framework can be applied to the symmetry detection problem by defining an appropriate unstructured model \mathcal{H}_0 , and a function $k(s)$ to compute the test statistic of a symmetric patch candidate s .

We suggest that a robust evaluation of the degree of symmetry that a patch s exhibits can be obtained by analyzing the gradient orientation errors of the pixels symmetrically positioned w.r.t. the symmetry axis supporting the patch. Let p_{1i} and p_{2i} be the pixels of the i -th symmetric pair, $i = \{1, \dots, N_p\}$, where N_p is the number of pairs contained in s . Then the orientation error of the pair is given by $\delta_i = |\text{Angle}(\nabla I(p_{1i}), \nabla_s I(p_{2i}))|$, where $\nabla I(p_{1i})$ is the gradient vector at p_{1i} and $\nabla_s I(p_{2i})$ is the gradient vector of the pixel p_{2i} , symmetrized w.r.t. the given symmetry axis (see figure 1 b). A similar idea was used in [11] to compute symmetry descriptors for image patches. Our goal is not only to assign a symmetry score to each patch, but to provide a decision-taking value about the meaningfulness of the patch symmetry.

With this choice, an appropriate (unstructured) null hypothesis \mathcal{H}_0 is a gradient field whose orientations are i.i.d. random variables, uniformly distributed over $[0, 2\pi]$. These properties hold in a Gaussian white noise model, under certain conditions of sub-sampling [7, p. 67]. From now on, X will denote a Gaussian white noise image, with the same size as the analyzed image I .

The function $k(s)$ can now be defined as the additive orientation error of the symmetric pairs of pixels in s . Let $\bar{\delta}_i$ be the normalized orientation error of a pair, *i.e.* $\bar{\delta}_i = \delta_i/\pi$. The test statistic of the patch s is given by $k(s) = \sum_{i=1}^{N_p} \bar{\delta}_i$. A perfect symmetric patch has all $\bar{\delta}_i = 0$, whilst the worst has all $\bar{\delta}_i = 1$.

To validate a given candidate s , we need to compute the probability of observing in the random image X , candidates with test statistics at least as extreme as $k_I(s)$, *i.e.* $\mathbb{P}_{\mathcal{H}_0}[k_X(s) \leq k_I(s)]$.

Generally speaking, we need to consider every configuration in which r independent uniformly distributed random variables taking values in $[0, 1]$ could sum to a value

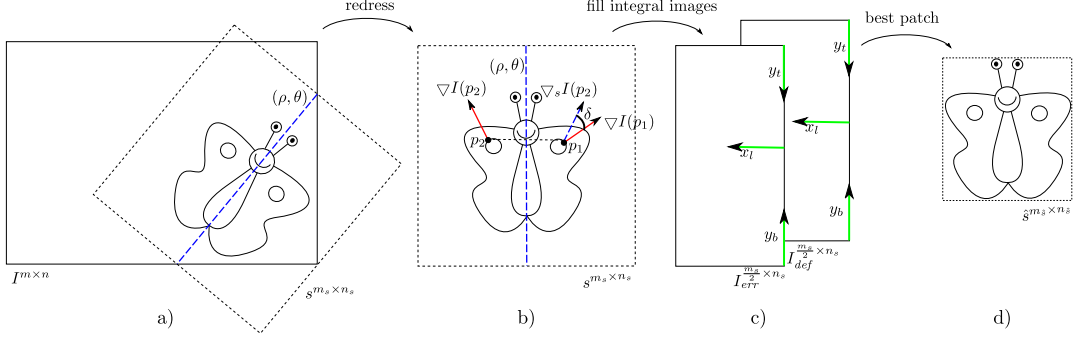


Figure 1. Validation pipeline: a)-b) Given a symmetry axis candidate (ρ, θ) , we extract and redress the maximal patch $s^{m_s \times n_s}$ along it; using the per-pair gradient orientation error δ , we fill the integral images $I_{err}^{\frac{m_s}{2} \times n_s}$ and $I_{def}^{\frac{m_s}{2} \times n_s}$, and then scan them (c) to locate the best patch (d).

smaller than a given $z \in \mathbb{R}$. Observe that these configurations are uniformly distributed on the unit hypercube. Let \mathcal{I}^r denote the r -dimensional unit hypercube, placed in the positive hyperoctant of \mathbb{R}^r , with one vertex at the origin. For any real z , let G_z^r denote the half-space in \mathbb{R}^r given by $G_z^r = \{\mathbf{x} = (x_1, \dots, x_r)^T \in \mathbb{R}^r \mid \sum_{i=1}^r x_i \leq z\}$. The cumulative distribution of the sum of r uniformly distributed random variables taking values in $[0, 1]$, is the volume of the slice obtained by intersecting the hypercube \mathcal{I}^r with the half-space G_z^r . Indeed, given r independent and uniformly distributed random variables $x_i \in [0, 1]$, with sum $S = \sum_{i=1}^r x_i$, then

$$\mathbb{P}[S \leq z] = \text{Vol}(\mathcal{I}^r \cap G_z^r) = \frac{1}{r!} \sum_{k=0}^{\lfloor z \rfloor} (-1)^k \binom{r}{k} (z - k)^r,$$

for $0 \leq z \leq r$, where $\lfloor z \rfloor$ is the largest integer not bigger than z . This formula is derived from a more general case studied in [16].

In our case, the sought probability is given by the volume of the slice obtained by intersecting the hypercube \mathcal{I}^{N_p} with the half-space $G_{k_I(s)}^{N_p}$. Moreover, it can be proven that the first term of the sum gives an upper bound of this probability. For computational reasons, we keep only this first term, as it is a sufficient approximation to evaluate the *NFA* test. Thence:

$$\mathbb{P}_{\mathcal{H}_0}[k_X(s) \leq k_I(s)] \leq \frac{[k_I(s)]^{N_p}}{N_p!}.$$

Finally, to complete the reasoning, we need to give an approximation for the number of tests N_I , *i.e.* the total number of patch candidates in image I . Assuming the representation of a patch by its bounding box, it follows that a patch has 5 d.o.f.: the coordinates of the symmetry axis limits (4 d.o.f.) and the width (1 d.o.f.). For a 1-pixel precision, the number of tests can then be approximated by $N_I = (mn)^{\frac{5}{2}}$. The number of tests ensures that the proposed validation

adapts to the image size, in order to keep under control the false positives for increasing image sizes.

To conclude, a symmetry candidate s , observed in an image I of size $m \times n$, containing N_p pairs of pixels, is accepted as valid detection if its *NFA* satisfies the simple test:

$$\frac{(mn)^{\frac{5}{2}}}{N_p!} \left[\sum_{i=1}^{N_p} \frac{|\text{Angle}(\nabla I(p_{1i}), \nabla_s I(p_{2i}))|}{\pi} \right]^{N_p} \leq 1. \quad (1)$$

As an aside, note that this is a *continuous* formulation of the *NFA*, similar to those proposed in [12] to merge segments in a segmentation application, or to compute detection thresholds in line segment detection problems [9]. Another possibility would be to use a *discrete* formulation similar to [7, 10, 21], in which for example, the test statistic $k(s)$ would denote the number of pairs whose orientation error is less than a precision threshold $\hat{\delta}$. However, we prefer the continuous formulation, which avoids this precision parameter.

3. Implementation Details

3.1. Candidate Selection

The detection pipeline starts by performing the candidate selection stage using the baseline algorithm [15]. Recall that this algorithm uses matched pairs of SIFT points to vote for salient symmetry axes. The votes are weighted by the gradient magnitude of the pixels in the pair, and by their orientation error, similar to Reisfeld transform [23].

3.2. Selection – Validation

This voting procedure is just the first part of the candidate selection. It provides the orientation and location (ρ, θ) within the image of possible symmetry axes. Optionally, the bounding box of the symmetric patch candidate can be recovered as the rectangle that covers the convex hull of the

pairs supporting the axis. However, this latter information may lack accuracy because of frequent errors in detecting and matching the SIFT feature points. Thus, we keep only the (ρ, θ) parameters of the axis. Afterwards, we use the properties of integral images to efficiently evaluate for symmetry all possible patches along the given symmetry axis. This is possible due to the additive nature of the $k(s)$ term involved in the *NFA* computation. Figure 1 shows the main steps of the proposed candidate validation procedure.

First, we extract and redress the maximal patch along the given axis (figure 1 a-b). The gradient of the patch is computed on 2×2 windows, and the gradient orientations in each symmetric pair are evaluated for symmetry with respect to the vertical middle axis. The error is then normalized as needed by the *NFA* test. Note that in a perfectly uniform image region, the gradient orientation is not defined. We choose to ignore pairs for which both pixels have undefined orientations. If only one of the pixels has undefined orientation, the pair gets maximum normalized error, *i.e.* 1.0.

The cumulative normalized error associated to the pairs in the patch is stored in an integral image I_{err} , which has the same number of lines n_s as the maximal patch, but halved number of columns $m_s/2$. Additionally, another integral image I_{def} of the same size as I_{err} , is used to keep track of the number of pixels that have defined orientations (figure 1 c). Then, any triplet (y_t, y_b, x_l) , where $n_s \geq y_t \geq 1$, and $0 \leq y_b < y_t$, and $\frac{m_s}{2} > x_l \geq 0$, defines a possible patch along the middle vertical axis, with size $(y_t - y_b, 2x_l)$. Thanks to integral image properties [6], by simply scanning the two integral images, we can extract the quantities $(k(s), N_p)$ needed to evaluate each patch, with only two additions and four subtractions, allowing to perform an exhaustive search for the best bounding-box in a reasonable time.

Note that in this case, the *NFA* score acts not only as a validation criterion, but also as a model selection criterion, guiding the search for the best meaningful symmetric patch. This dual usage has been proposed also in [7, 18, 21]. The exemplar detections given in section 4 show the advantage of this approach, as it allows to recover from imprecise bounding boxes given by the baseline algorithm as a result of accidental merging of pairs not belonging to the same object.

It is important to mention that natural images usually contain only approximate (partial) symmetries. Moreover, the gradient orientations are inherently affected by the image noise. Hence, a multi-scale evaluation of the *NFA* is necessary. We perform this by smoothing and successively sub-sampling the original image. This has the effect of concealing the noise and the imperfections of the symmetric patches, so we can identify meaningful detections.

Thus, to validate a candidate, one needs to rotate the

maximal patch along the axis, and then fill and scan the two integral images to obtain the best patch. On a regular computer (CPU i5 2.53GHz), this operation takes about 2s per candidate, for an 800×600 pixels image. For efficiency, one can choose to keep and validate only the first 5 or 10 candidates reported by the baseline algorithm which are more likely to correspond to real symmetries, so that the execution time of the validation stage is comparable to the execution time of the candidate selection stage. Note that if the precision of the symmetry axis candidates is poor, the whole validation pipeline can be repeated for small variations around the (ρ, θ) values proposed by the baseline algorithm, and keep eventually the most meaningful patch. However, this increases the execution time.

4. Exemplar Detections

This section illustrates in a qualitative manner the advantages of the proposed detection pipeline, and also the main directions of improving its performance.

We first carried out tests on random noise images, as a sanity check for the probabilistic model underlying the validation procedure (figure 2, 1st row). We used 15 random noise images of 512×512 pixels, generated with different variance values (figure 2, 1st row, 1st column gives a sample). The baseline algorithm reported on average 10 (false positive) detections per image (figure 2, 1st row, 2nd column). None of the candidates passed the validation test (1) (figure 2, 1st row, 3rd column).

The rows 2 – 5 in figure 2 illustrate the satisfactory behavior of the proposed pipeline on natural images containing single or multiple symmetries. It succeeds in pruning most of the false positives proposed by the baseline algorithm, while identifying with relatively good accuracy the expected bounding box of the symmetric patches.

The 6th row presents a typical false negative case: the candidates proposed by the baseline algorithm are too far from the real symmetry, so no detection is validated by our approach, even after repeating the validation as mentioned in the previous section. This underlines the need of a more precise candidate selection. The Hough voting used by the baseline algorithm has the advantage of reducing the number of detections by clustering the pairs that possibly support the same symmetry axis, but at the same time it can accidentally merge pairs belonging to different objects or to the background, resulting in low quality candidates.

Finally, the last row exemplifies two types of false positives reported by our procedure, and points to directions of future work. The detections reported on the background correspond to smooth, relatively uniform image regions. Indeed, such regions are symmetric according to our model, *i.e.* the gradient orientation error is smaller than the one expected in a Gaussian noise image. However, they should be pruned because they are not meaningful for the human

perception. The second type of false positives are the detections reported on actually symmetric patches which were not labeled by the human experts, because they do not belong to the foreground objects which usually capture experts' attention, *e.g.* the horizontal detection on the body of the zebra, or some of the detections on rows 2 – 5, unlabeled by experts. This also points out the difficulty of designing an accurate benchmark for symmetry detectors evaluation.

5. Conclusion

In this paper we proposed a novel approach for the detection of mirror-symmetric image patches, which uses a heuristic method to select potential symmetric patches, and a principled validation step, based on the *a contrario* theory, which accepts as valid detections only candidates that are not likely to appear in noise. However, this work is in progress, and several improvements are foreseen. The validation step could be improved through a principled method of weighting the gradient orientation errors according to the gradient magnitude. This would prevent the pixels with low gradient to contribute to the symmetry score, as their orientation can be strongly affected by noise, thus not reliable. The expected effect of this idea is the decrease of the number of false positives reported on smooth, uniform image regions, where pixels have low gradient magnitude.

The candidate selection step remains a challenging issue. Here we used an existing heuristic combined with integral image properties to recover efficiently the most salient symmetry candidates. However, this method still introduces false negatives either because of the absence of feature points on regions with low contrast, or because of feature points wrongly matched or point pairs wrongly merged by the Hough voting procedure. Another important direction is the handling of the perspective skew, hence we need to use feature points and descriptors invariant to affine transforms, like ASIFT [19].

Acknowledgments: This work is partially funded by a Qualcomm postdoctoral grant, École Polytechnique and Marie-Curie CIG grant 334283.

References

- [1] <http://vision.cse.psu.edu/research/symmcomp/index.shtml>.
- [2] H. B. Barlow and B. C. Reeves. The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision research*, 19(7):783–793, 1979.
- [3] A. Chia, S. Rahardja, D. Rajan, and M. K. Leung. A split and merge based ellipse detector with self-correcting capability. *IEEE T. Image Process.*, 20:1991–2006, 2011.
- [4] H. Cornelius and G. Loy. Detecting bilateral symmetry in perspective. In *Proc. of CVPR Workshop*, 2006.
- [5] H. Cornelius, M. Perdoch, J. Matas, and G. Loy. Efficient symmetry detection using local affine frames. In *Proc. of SCIA*, 2007.
- [6] F. Crow. Summed-area tables for texture mapping. In *Proc. of SIGGRAPH*, 1984.
- [7] A. Desolneux, L. Moisan, and J.-M. Morel. *From Gestalt Theory to Image Analysis: A Probabilistic Approach*. Springer, 2008.
- [8] A. Gordon, G. Glazko, X. Qiu, and A. Yakovlev. Control of the mean number of false discoveries, bonferroni and stability of multiple testing. *Ann. Appl. Stat.*, 1:179–190, 2007.
- [9] R. Grompone von Gioi and J. Jakubowicz. On computational gestalt detection thresholds. *J. Physiology-Paris*, 103:4–17, 2009.
- [10] R. Grompone von Gioi, J. Jakubowicz, J.-M. Morel, and G. Randall. Lsd: A fast line segment detector with a false detection control. *PAMI*, 32:722–732, 2010.
- [11] D. C. Hauagge and N. Snavely. Image matching using local symmetry features. In *Proc. of CVPR*, 2012.
- [12] L. Igual, J. Preciozzi, L. Garrido, A. Almansa, V. Caselles, and B. Rougé. Automatic low baseline stereo in urban areas. *Inverse Probl. and Imaging*, 1:319–348, 2007.
- [13] D. G. Lowe. *Perceptual Organization and Visual Recognition*. Kluwer Academic Publishers, 1985.
- [14] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.
- [15] G. Loy and J.-O. Eklundh. Detecting symmetry and symmetric constellations of features. In *Proc. of ECCV*, 2006.
- [16] J.-L. Marichal and M. J. Mossinghoff. Slices, slabs, and sections of the unit hypercube. *Online Journal of Analytic Combinatorics*, 2006.
- [17] J. Matas, C. Galambos, and J. Kittler. Progressive probabilistic hough transform. In *Proc. of BMVC*, 1998.
- [18] L. Moisan and B. Stival. A probabilistic criterion to detect rigid point matches between two images and estimate the fundamental matrix. *IJCV*, 57:201–218, 2004.
- [19] J. M. Morel and G. Yu. Asift, a new framework for fully affine invariant image comparison. *SIAM Journal on Imaging Sciences*, 2(2):438–469, 2009.
- [20] M. Park, S. Lee, P.-C. Chen, S. Kashyap, A. A. Butt, and Y. Liu. Performance evaluation of state-of-the-art discrete symmetry detection algorithms. In *Proc. of ECCV*, 2012.
- [21] V. Pătrăucean, P. Gurdjos, and R. Grompone von Gioi. A parameterless ellipse and line segment detector with enhanced ellipse fitting. In *Proc. of ECCV*, 2012.
- [22] V. Pătrăucean. *Detection and Identification of Elliptical Structure Arrangements in Images: Theory and Algorithms*. PhD thesis, Institut National Polytechnique de Toulouse, France, 2012.
- [23] D. Reissfeld, H. Wolfson, and Y. Yeshurun. Context free attentional operators: the generalized symmetry transform. *IJCV*, 14:119–130, 1995.
- [24] T. Tuytelaars, A. Turina, and L. van Gool. Noncombinatorial detection of regular repetitions under perspective skew. *PAMI*, 25(4), 2003.
- [25] H. Weyl. *Symmetry*. Princeton University Press, 1952.




Figure 2. Comparative qualitative results. 1st column: human-labeled ground truth, 2nd column: first 10 detections reported by the baseline algorithm, 3rd column: detections validated by the proposed method. Since we keep only the (ρ, θ) parameters of the candidates proposed by the baseline algorithm, the detections reported by our method may have different locations along the axes.