

Consensus Graphs for Symmetry Plane Estimation

Andrea Albarelli¹, Marcello Pelillo¹, and Sebastiano Viviani²

¹ Dipartimento di Informatica, Universita' Ca' Foscari di Venezia

² Evolvenda s.r.l.

Abstract. Finding the most relevant symmetry planes for an object is a key step in many computer vision and object recognition tasks. In fact such information can be effectively used as a starting point for object segmentation, noise reduction, alignment and recognition. Some of these applications are strongly affected by the accuracy of symmetry planes estimation, thus the use of a technique that is both accurate and robust to noise is critical. In this paper we introduce a new weighted association graph which relates the main symmetry planes of 3D objects to large sets of tightly coupled vertices. This technique allows us to cast symmetry detection to a classical pairwise clustering problem, which we solve using the very effective Dominant Sets framework. The improvement of our approach over other well known techniques is shown with several tests over both synthetic data and sampled point clouds.

Keywords: symmetry, dominant sets, shape registration, clustering.

1 Introduction

Among all the features that can be extracted from a 2D image or 3D object symmetry is one of the most important. For instance symmetry axes can be used to set a canonical orientation, which in turn can be useful for matching different objects, computing partial registration of surfaces or simply to offer a comparable point of view in visual application. Other applications exploit symmetries as a descriptor for searching in large databases or as a starting point for locating other features. In addition partial local symmetries can be very effective as a cue in segmentation. Finally, due to its abundance in both natural and crafted object, symmetries play an important role in visual perception and Gestalt theory. Because of this wide range of applications, symmetry estimation has been a popular research field for long time. Most seminal papers appeared in the eighties and were focused on perfect symmetries, mainly because this kind of analysis required less computational power. Atallah[1] and Wolter[2] solved the problem very efficiently by casting it into substring matching. Some more recent approaches use orientation histograms[3], singular value decomposition[4], extended Gaussian image[5] and generalized complex moments[6]. Also principal component analysis[7] can be used with good results when dealing with strongly symmetric objects. All these approaches rely on the existence of a perfect plane

of symmetry, thus they are very sensitive to noise and cannot be applied to objects where partial or local symmetries are present.

In order to deal with partial or imperfect symmetries some measures of the distance between a given object and the nearest perfectly symmetric shape have been proposed[8,9]. Such a measure could be used to search in the full plane space for the solution that minimizes the error, however since it can be applied to just one plane at a time it would not be an efficient solution.

A more practical approach, that has been proved effective in a wide range of contexts, is to use some kind of transform. That is a mapping from a function defined over the space of the object data to a function defined over the space of the feasible solutions, where local and global maximum are usually associated to local and global solutions of the problem itself. The most well known transform in computer vision is the Hough transform, which maps points on a plane to the space of parametrized straight lines. Over the last decade several symmetry transform have been proposed[10,11,12,13,14]. Recently the Planar-Reflective Symmetry Transform (PRST) has been presented in[15,16] and extended in[17]. This is a transform that maps the points of three dimensional objects to a real-valued function in the three dimensional space of parametric planes. As for any transform technique in practical implementations it is computed in a discrete form: that is, the function is evaluated only on a large set of randomly selected points and the results are used to vote for the best solutions in a discrete set of bin that are uniformly distributed over the space of planes. This approach leads to coarse errors, due to the discrete size of the bins. However a continuous refinement technique is proposed that greatly enhances the solution by searching for a more accurate symmetry plane inside the winning bin by using an error minimization technique inspired by the Iterative Closest Point algorithm[18].

In this paper we propose a novel approach which casts the problem to the search for pairwise clusters in a newly defined weighted association graph that we named Symmetry Consensus Graph. In such a graph each vertex is associated to a candidate symmetry plane and the similarity between vertices expresses the consensus between planes. This way each large cluster in the graph captures a set of highly similar plane candidates and leaves out all the unwanted outliers, according to the high internal homogeneity and external inhomogeneity properties of clusters. In section 2 we define the Symmetry Consensus Graph in both a binary and a weighted form. In section 3 we suggest a very effective framework in order to find large cluster in a Consensus Graph. Finally, in sections 4 and 5 we compare our approach with PCA and the very accurate PRST technique.

2 The Symmetry Consensus Graph

Given a large set of 3D points describing an object our goal is to find the largest subsets that share a common symmetry plane. The most widely used approach is to select a large number of point pairs and thus collect votes for the plane between them in an appropriate parameters space. As for any other voting technique, concrete implementations of this approach need to choose a quantization of the

parameters space into discrete bins. In many scenarios the choice of the size of the bins can be very critical as a too coarse quantization can lead to a poor precision in the estimation of the symmetry, while a too fine one could suffer of vote dispersion due to the dimensionality curse. In [16] the authors suggest to adopt a quite coarse parameters space quantization and thus apply a continuous refinement step in order to minimize the local symmetry error among points that voted for the winning bin. This step enhances the quality of the solution found, but still is quite sensible to the presence of local outliers in the rather large winning bin. In order to obtain the best results an optimal bin size should be used, that is a quantization that allows to collect just the most relevant votes, leaving out local outliers. Unfortunately it is very hard to know in advance which size to use as it depends on many factors such as noise and quality of the symmetry. By contrast, the main idea which underpin our method is to avoid the quantization by using a more global approach able to weight all the voting point pairs at once.

2.1 Binary Symmetry Consensus Graph

In this section we focus our attention on objects that, although not being symmetric as a whole, contain one or more subset of points that are perfectly specular by a plane. Obviously this scenario is not very realistic from a practical point of view, nevertheless it permits us to define an useful structure that we will expand in section 2.2 to be more suitable for real world applications:

Definition 1. *Given a set of points in space $P = \{p_1, p_2 \dots p_n\}$, their Binary Symmetry Consensus Graph (BSCG) is the undirected graph $G_B = (V_B, E_B)$ where $V_B = \{(p_a, p_b) \in P | a < b\}$ is the set of pairs of points in P and where $((p_a, p_b), (p_c, p_d)) \in E_B$ if and only if the plane between p_a and p_b is exactly the same plane between p_c and p_d .*

Each vertex in such a graph represents a pair of points in the original object and is attributed with the plane that separates such points in a symmetric way. The key property of this graph is that two vertices are connected if and only if they share the same plane. As this is an equivalence relation the transitive property holds and it is easy to see that each set of vertices that share the same symmetry plane with each other forms a complete subgraph (clique). By building the BSCG over a set of points we can thus cast the search for the most relevant symmetry planes into the search for the largest cliques over the Consensus Graph. Although this is known to be a NP-complete problem many powerful heuristics have been suggested for it.

In figure 1 we illustrate a simple example of a BSCG. On the left an object with two partial perfect symmetry planes is shown. On the right the BSCG is built over all the pairs of point of the object. Only two cliques are found in the BSCG, respectively between pairs ag, bf, ce , which are symmetric with respect to the vertical axis, and between pairs bc, ef , which are symmetric with respect to the horizontal axis. All other pairs do not share their symmetry plane with any other pair in the graph, thus all of them are singleton.




Fig. 1. A very simple object with perfect partial symmetries (left) and its Binary Symmetry Consensus Graph (right). Only pairs of vertices that share the same symmetry plane are connected by an edge. For practical reasons only a part of the graph is shown.

2.2 Weighted Symmetry Consensus Graph

The assumptions taken in section 2.1 are very strict in practical scenarios. In fact noise, numerical errors and not perfect symmetries in the object itself can all make very difficult for many pairs of points to share exactly the same symmetry plane. One possible solution to this problem could be the use of a threshold in order to state that two plane are equals. Unfortunately this approach introduces two more hurdles: the first is the correct choice for such threshold, which is not different from the choice of the correct bin size, the other is that this way the transitive property would be no more valid, thus breaking the correspondence between cliques in the Consensus Graph and symmetries.

Our proposal is to extend the BSCG presented in 2.1 to a weighted graph where each edge is weighted according to the similarity between the planes associated to the vertices connected:

Definition 2. Given a set of points in space $P = \{p_1, p_2 \dots p_n\}$, their *Weighted Symmetry Consensus Graph (WSCG)* is the edge weighted undirected graph $G_W = (V_W, E_W, \omega)$ where $V_W = \{(p_a, p_b) \in P | a < b\}$ is the set of pairs of points in P , $E_W = V_W \times V_W$ is the set of edges and $\omega : E_W \rightarrow \mathbb{R}^+$ is a weight function that assigns to each edge a positive real value $\omega((p_a, p_b), (p_c, p_d))$ proportional to the similarity between the symmetry plane of p_a, p_b and the symmetry plane of p_c, p_d .

From a topological point of view a WSCG is by construction a complete graph, so it would not be interesting to search for cliques in it as we made for the BSCG. Instead we are interested in finding large sets of vertices that are all associated to very similar symmetry planes. In other words, since the edges reflect the similarity between planes we are interested in finding sets of vertices that present a high pairwise similarity among them and a low pairwise similarity with respect to vertices external to the set itself. These two properties (high internal similarity and low external similarity) are commonly used to define the notion of a cluster, which indeed generalizes to the weighted graphs context the notion of clique in binary graphs.

In figure 2 we illustrate a simple example of a WSCG. The object shown is a slightly perturbed version of the one shown in figure 1. In this case there are




Fig. 2. A very simple object without perfect symmetries (left) and its Weighted Symmetry Consensus Graph (right). The weight of the edges is proportional to their thickness. For practical reasons negligible edges are not drawn and only a part of the graph is shown.

no subset of points that exhibit a perfect symmetry, nevertheless the separation plane between a, g is very similar to the one between b, f , which in turn is just a little less similar to the ones between c, d and c, e , and so on (note that only the edges with relevant weight are drawn). The two largest clusters correspond to the two most relevant symmetry planes in the model. It is worth noting than a cluster could assign to the same point in the model two or more different specular points, as in the case of point c which is symmetric to both e and d . This many-to-many property of the assignment is critical when dealing with data coming from sampling, where artifacts and random noise make very unlikely a direct correspondence between each point and its specular instance, even in the presence of perfect symmetry in the sampled object.

3 Dominant Sets Cluster Extraction

In section 2.2 we shown how the problem of finding optimal symmetry planes can be casted into the search for large clusters in a set of vertices over which we defined a pairwise similarity measure ω . In order to perform this search we propose to use the Dominant Set framework [19] which is a very effective pairwise clustering technique.

Given an edge weighted undirected graph $G = (V, E, \omega)$, a subset of vertices $S \subseteq V$ and two vertices $i \in S$ and $j \notin S$ the following function measures the coherence between nodes j and i , with respect to the average coherence between node i and its neighbors in S

$$\phi_S(i, j) = \omega(i, j) - \frac{1}{|S|} \sum_{k \in S} \omega(i, k) \tag{1}$$

And overall weighted coherence between i and other nodes in S is defined as:

$$w_S(i) = \begin{cases} 1 & \text{if } |S| = 1 \\ \sum_{j \in S \setminus \{i\}} \phi_{S \setminus \{i\}}(i, j) w_{S \setminus \{i\}}(j) & \text{otherwise} \end{cases} \tag{2}$$

Intuitively, $w_S(i)$ will be high if i is highly coherent with vertices in S . Given this measure $S \subseteq V$ is said to be *dominant* if $\sum_{i \in T} w_T(i) > 0$ for any nonempty $T \subseteq S$ and the following conditions hold:

$$w_S(i) > 0, \forall i \in S \quad (3)$$

$$w_{S \cup \{i\}}(i) < 0, \forall i \notin S \quad (4)$$

The conditions above correspond to the two main properties of a cluster: namely internal homogeneity and external inhomogeneity, which are exactly the properties that we expect to hold for the sets we are searching for.

Another compelling reason to use the Dominant Set framework is that, despite the complexity of their recursive definition, they are very easy to find in practice: in fact a single set (corresponding to a symmetry plane) can be found in polynomial time and a convenient technique exists in order to enumerate all the Dominant Sets in a graph.

Pavan and Pelillo [19] have shown that dominant sets correspond to local maximizer over the standard simplex (i.e. the sets of vectors with components that sum up to 1) of the quadratic function

$$f(\mathbf{x}) = \mathbf{x}^t A \mathbf{x} \quad (5)$$

where A is the weighted adjacency matrix of the graph (thus $A_{ij} = \omega(i, j)$).

These maximums can be found by exploiting the convergence properties of the payoff monotonic replicator dynamic

$$x_i(t+1) = \frac{(Ax(t))_i}{x(t)^t Ax(t)} \quad (6)$$

which is guaranteed to converge to a local maximum when the association graph is undirected and, thus, the matrix A is symmetric[20]. At convergence the value of the function f is a measure of the global coherence of the extracted set, which in turn express the strength of the symmetry found.

Since an object could express more than a single symmetry plane it would be useful to be able to enumerate at least the most relevant Dominant Sets present in the graph. In addition, although replicator dynamics are guaranteed to converge it does not need to stop on the largest Dominant Set, thus by enumerating them we could choose always the most relevant ones. To obtain such enumeration we used the technique proposed in [21].

4 Experimental Results

In order to show the effectiveness of the proposed approach we performed both qualitative tests on a database of real world objects and quantitative validation over a larger set of synthetic data generated under different conditions of noise and occlusion. In both cases our technique has been compared with the PRST[15] algorithm which we implemented with 64 bins for each dimension of the space of planes parametrization and with ISP refinement. As an optimization meant

to keep the size of the Symmetry Consensus Graph reasonable we filtered the vertex set V_W by sampling a fixed number of random vertices.

The qualitative tests have been made against a database of about 40 3D scans of sunglasses models acquired as point clouds with a structured light scanner. Each model counts between 10.000 and 20.000 points which have been sampled with a metric error which is guarantee to be less than 10 micrometers. Our goal is to find for each model the main symmetry plane, that is, in this context, the plane that split each sunglasses pair in two specular halves. Despite the high quality of the sampling process it is difficult to obtain an accurate result because of slight differences in the two halves of the model, incomplete sampling and parts of the object that are not fixed and can move during the sample process (for instance nose pads). These hurdles strongly limit the use of traditional global techniques, such as PCA, and the presence of non perfect symmetry can easily lead to some slight imprecision also in more refined methods, such as PRST.

In figure 3 we show the performance of these latter methods in comparison to WSCG. The central blue plane is the symmetry plane found, the two lateral yellow planes are added reference planes that are parallel and equidistant to it. The rows in the figure show respectively the plane obtained with PCA, PRST and WSCG. The plane found by PCA is clearly skewed because of the asymmetric sampling of the object during range scanning. The result obtained by the PRST algorithm is far better than PCA, as expected, while WSCG gives another improvement over it.

Qualitative results offer a visual hint about the kind of improvement obtained with the WSCG technique, but as PRST is in itself very accurate this enhancement could appear minimal. In order to give a more quantitative insight we




Fig. 3. Qualitative results on 3D scans. In the first row the symmetry plane obtained with PCA is shown from different angles and details. The second row contains the results of PRST and the third the results of WSCG.

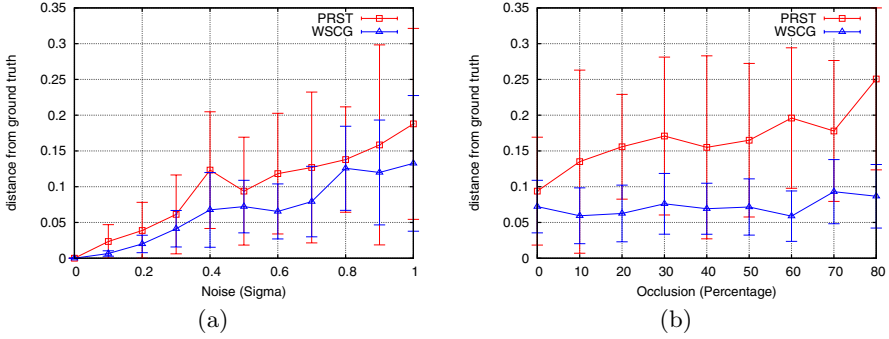


Fig. 4. Accuracy of PRST and WSCG with different level of noise and occlusion

performed a set of experiments with synthetic data which allowed us to compare the symmetry planes obtained with respect to an accurate ground truth plane. We generated a random distribution of 2000 points within a range of 400 millimeters perfectly symmetric with respect to a known plane. To evaluate the performance of WSCG in different conditions we added gaussian random noise along all three axes with a value of the parameter σ ranging from 0.1 to 1 millimeters. Then we repeated the experiments by setting $\sigma = 0.5$ and simulating occlusion with the deletion of a growing percentage of random points from one side of the symmetry plane. Each experimental condition has been tested with 20 independent runs. In each run the two algorithms are applied and the distance between the obtained principal symmetry plane and the ground truth is measured. While several symmetry measures have been proposed [8,9], we are not really interested in knowing how much the found symmetry diverge from a perfect one, as we already know that our data are noisy: rather we are interested in a plane distance measure. Each plane has been parametrized by the vector P perpendicular to it and the distance between two planes is measured as the norm of the difference between their respective associated vectors P_1 and P_2 :

$$d = \| P_1 - P_2 \| \quad (7)$$

It must be noted that this simple measure mixes up translational and rotational errors, which is reasonable in this context because errors are always very small for both PRST and WSCG. The measure of equation 7 is also used to define the function ω that weights the edges of the WSCG in definition 2.2. If P_{ab} is the vector perpendicular to the separation plane between points p_a, p_b and P_{cd} the one associated to the plane between p_c, p_d then we measure the similarity between vertices (p_a, p_b) and (p_c, p_d) of the WSCG as:

$$\omega((p_a, p_b), (p_c, p_d)) = e^{-\|P_{ab} - P_{cd}\|} \quad (8)$$

In figure 4(a) the performance of WSCG is shown for several levels of gaussian noise added to the test points and in figure 4(b) the robustness with respect




Fig. 5. Running time and accuracy of WSCG with respect to the number of points

to occlusion is shown for different percentages of sample points removal. The comparison with PRST shows that both techniques offer comparable very good performance and that WSCG is more accurate on average. In particular WSCG seems to be more tolerant to occlusion, which is expected as PRST relies on a voting algorithm. In figure 5(a) we show the convergence time of the WSCG algorithm with respect to the number of vertices that are kept in the Consensus Graph. Each test has been performed 20 times with gaussian noise of parameter $\sigma = 0.5$ and an occlusion of 50% and has been executed on a 2Ghz Athlon CPU. These result show that by keeping the number of sampled points low, WSCG can be executed in less than a second. This optimization does not imply a lower quality in the symmetry plane estimation, in fact figure 5(b) shows that the error becomes very small when keeping as few as 300 vertices.

5 Conclusions

The Weighted Symmetry Consensus Graph is a new approach to planar symmetry detection that exploits pairwise clustering to find a solution that is maximally consistent from a global perspective. The main features of WSCG are that it does not require any quantization of the solutions space and is adaptive with respect to outliers. Experimental validation shows that this approach gives better results than another state-of-the-art algorithm. Finally WSCG is quite general as the compatibility measure needs not to be related to planar symmetry. In fact future work will aim to extend the approach to radial symmetry and other features that can be estimated by exploiting a global consensus.

Acknowledgments. We thank Luxottica Group S.p.A. for funding this research and providing the experimental data. We also acknowledge the financial support of the 7th Framework Programme, under FET-Open project SIMBAD grant no. 213250.

References

1. Atallah, M.J.: On symmetry detection. *IEEE Trans. Computers* 34(7), 663–666 (1985)
2. Wolter, J.D., Woo, T.C., Volz, R.A.: Optimal algorithms for symmetry detection in two and three dimensions. *The Visual Computer* 1(1), 37–48 (1985)
3. Sun, C.: Symmetry detection using gradient information. *Pattern Recogn. Lett.* 16(9), 987–996 (1995)
4. Shah, M.I., Sorensen, D.C.: A symmetry preserving singular value decomposition. *SIAM J. Matrix Anal. Appl.* 28(3), 749–769 (2006)
5. Sun, C., Sherrah, J.: 3d symmetry detection using the extended gaussian image. *IEEE Trans. Pattern Anal. Mach. Intell.* 19(2), 164–168 (1997)
6. Shen, D., Ip, H.H.S., Cheung, K.K.T., Teoh, E.K.: Symmetry detection by generalized complex (GC) moments: A close-form solution. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21(5), 466–476 (1999)
7. Lu, C., Zhang, C., Wen, F., Yan, P.: Principle component analysis-based symmetry detection. *Tien Tzu Hsueh Pao/Acta Electronica Sinica* 27(5), 25–28 (1999); Cited By (since 1996): 4
8. Zabrodsky, H., Peleg, S., Avnir, D.: Symmetry as a continuous feature. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 17(12), 1154–1166 (1995)
9. Kulkarni, P., Dutta, D.: An investigation of techniques for asymmetry rectification. *Journal of Mechanical Design* 117(4), 620–626 (1995)
10. Reisfeld, D., Wolfson, H., Yeshurun, Y.: Context free attentional operators: the generalized symmetry transform. *Int. J. of Computer Vision, Special Issue on Qualitative Vision* (1994)
11. Choi, I., Chien, S.I.: A generalized symmetry transform with selective attention capability for specific corner angles. *Signal Processing Letters, IEEE* 11(2), 255–257 (2004)
12. Bigün, J.: Pattern recognition in images by symmetries and coordinate transformations. *Computer Vision and Image Understanding* 68(3), 290–307 (1997)
13. Gesú, V.D., Valenti, C., Strinati, L.: Local operators to detect regions of interest. *Pattern Recogn. Lett.* 18(11-13), 1077–1081 (1997)
14. Loy, G., Zelinsky, A.: Fast radial symmetry for detecting points of interest. *IEEE Trans. Pattern Anal. Mach. Intell.* 25(8), 959–973 (2003)
15. Podolak, J., Shilane, P., Golovinskiy, A., Rusinkiewicz, S., Funkhouser, T.: A planar-reflective symmetry transform for 3d shapes. In: *SIGGRAPH 2006: ACM SIGGRAPH 2006 Papers*, pp. 549–559. ACM, New York (2006)
16. Podolak, J., Shilane, P., Golovinskiy, A., Rusinkiewicz, S., Funkhouser, T.: A planar-reflective symmetry transform for 3d shapes. *ACM Trans. Graph.* 25(3), 549–559 (2006)
17. Rustamov, R.M.: Augmented symmetry transforms. In: *SMI 2007: Proceedings of the IEEE International Conference on Shape Modeling and Applications 2007*, Washington, DC, USA, pp. 13–20. IEEE Computer Society, Los Alamitos (2007)
18. Besl, P.J., McKay, N.D.: A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.* 14(2), 239–256 (1992)
19. Pavan, M., Pelillo, M.: Dominant sets and pairwise clustering. *IEEE Trans. Pattern Anal. Mach. Intell.* 29(1), 167–172 (2007)
20. Pelillo, M.: Replicator equations, maximal cliques, and graph isomorphism. *Neural Comput.* 11(8), 1933–1955 (1999)
21. Rota Bulò, S., Torsello, A., Pelillo, M.: A continuous-based approach for partial clique enumeration. In: Escolano, F., Vento, M. (eds.) *GbRPR 2007*. LNCS, vol. 4538, pp. 61–70. Springer, Heidelberg (2007)