

Figure 7.8 shows the cross-sections so recovered for the connection of the SHGC of Figure 6.8.b and the termination of the SHGC of Figure 6.8.d.



Figure 7.8 Axis based cross-section recovery for previous SHGCs.

For discontinuous connections where there are no limb patches on either side of a connection, no cross-sections can be recovered. This leaves *holes* in the final description (see top right SHGC in Figure 7.12.b and c for which discontinuity is caused by self occlusion). In the case of LSHGCs, however, the recovery is straightforward as the limbs are known on both sides.

7.5.2 Limb reconstruction method

Cross-sections recovered by the previous method for a connection or a termination can be used to infer the missing limb boundary. The limb reconstruction method consists of finding a point on each of the recovered cross-sections that is a limb point (in the projection of an SHGC, limbs and internal cross-sections are tangential to each other).

One possible way of finding these points would be first to find the intersection point P_t of the tangent line at each point P_u with the axis. Then the corresponding limb point would be the point on the recovered cross-section at P_u whose tangent line intersects the axis at P_t (i.e. property P3). However, we have already mentioned that this property is sensitive to noise and distortions. Instead, we use a method that finds the *tangential envelope* of the set of recovered cross-sections. Given a starting point, call it P_0 as in Figure 7.9, taken to be an extremity at the *open side* of the connection (or termination), the method consists of finding a point P_1 on the first recovered cross-section whose tangent line passes through P_0 . P_1 is marked as a limb point. The process is then repeated for P_1 and the next cross-section until the limb point on the last cross-section is so determined. Since the axis-based cross-section recovery produces a dense set of very close cross-sections,

7.5.1 Axis-based cross-section recovery method

Given the axis of the global SHGC and a reference cross-section (of any of the local patches), this method recovers cross-sections for unmatched limb patches (at connections or terminations). First, for each point P_u of a given unmatched limb, its corresponding point R_u on the reference cross-section is found (they have parallel tangents⁸). See Figure 7.7. The reference cross-section is translated so that P_u and R_u coincide. The scale of the recovered cross-section relatively to the reference one is determined as follows:

In the case of an LSHGC (Figure 7.7.a), the corresponding point P_c is simply the intersection of the line from P_u parallel to the limb correspondence line of the reference cross-section (line R_u-R_c in the figure) and the other straight limb of the LSHGC. The scale is given by the ratio $dist(P_u, P_c) / dist(R_u, R_c)$.

In the case of a non-linear SHGC (Figure 7.7.b), first the intersection point P_x of the line connecting P_u to R_u and the axis is determined⁹. Then, by the property of linear parallel symmetry between the cross-sections, it can be shown that the scale is given by the ratio $dist(P_x, P_u) / dist(P_x, R_u)$.




Figure 7.7 Axis based cross-section recovery method. a. for LSHGCs. b. for nonlinear SHGCs

For a termination, the method is applied until the end of the unmatched limb is reached or there is overlap between the recovered cross-section and the *bottom end* cross-section of the SHGC. In doing so, we obtain a more accurate segmentation of limbs and cross-section.

8. there may be many points having the same tangent but only one is the corresponding point. It is selected based on the continuity of the cross-section sweep and of the cross-section function.

9. the reference cross-section is chosen so that the correspondence line is not parallel to the axis.

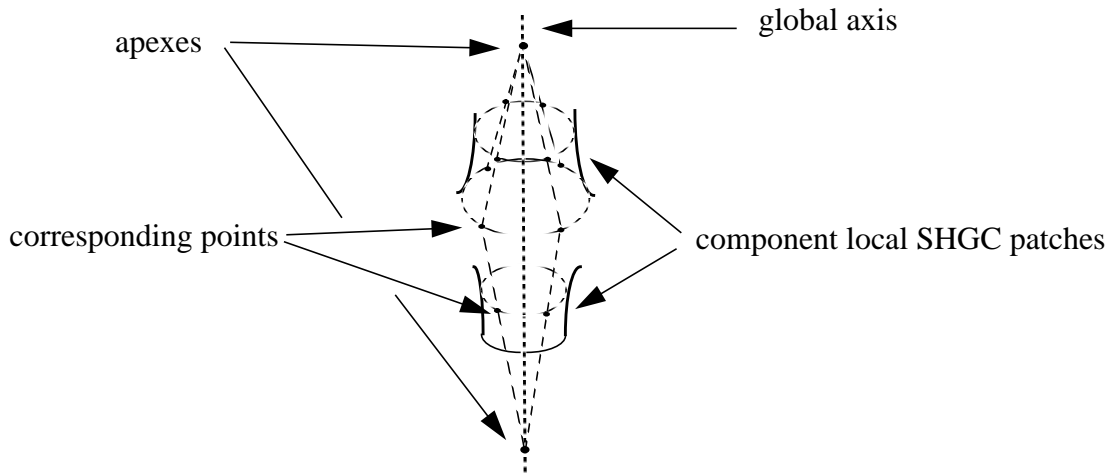


Figure 7.6 Global axis detection

Verification between the global axis and the component local ones uses the rules described in 7.2. Successfully verified groupings form a global SHGC with a more accurate estimate of the axis.

7.5 Boundary completion

Global SHGC patches formed in the previous step consist only of aggregates of local patches believed to make up a single global surface. The descriptions of those global surfaces may thus be discontinuous if they are occluded or simply bounded by broken contours. This can be seen in the case of the occluded vase of Figure 6.8d, for which the surface boundary does not terminate due to that occlusion (other examples include the objects of Figure 1.1 and Figure 7.5c and d). However, our (human) perception of a surface is clear there. In fact, we can even guess the shape of the hidden boundary due to the symmetric nature of the shape. We show that projective invariants can also be used for completion of surface descriptions. In this step, gaps in descriptions of verified global SHGCs are completed. Boundary completion is done for connections between adjacent local SHGC patches and *terminations* where a local SHGC patch, at an extremity of the global SHGC, has an incomplete limb correspondence due, say, to occlusion as is the case for the lower part of the occluded vase in Figure 6.8.d. The method consists of, first, recovering cross-sections at all points of unmatched limb patches, using a method we call *axis-based cross-section recovery method*, then finding the (missing) corresponding limb patch using a method we call *limb reconstruction method*. We discuss the two methods separately.

7.3 Selection of hypotheses

Because of the highly constrained nature of the compatibility measures, conflicting hypotheses are rare at this level. When they do happen, they involve alternative connections at the same extremity of some local SHGC patch. The only filtering done at this step consists of preferring continuous connections over discontinuous ones. Among the remaining hypotheses, the one involving the closest connection is selected. This method has been sufficient in all tested examples where very few conflicts have arisen.

7.4 Verification

In this step, grouping hypotheses (candidate global SHGCs) generated in the previous steps are checked for global consistency. It is possible that some connected set of local SHGC patches are locally consistent but globally inconsistent. For example, cylindrical patches with small local variations in the direction of the axis will not correspond to a global SHGC if the variation between, say the first and last patches, is too large. Global consistency is checked by first determining the global axis of each set of connected local patches, then verifying the compatibility of that axis with each of the local ones.

The global axis detection depends on the nature of the local patches. If all the local patches are cylindrical with mutually colinear limbs, then they form a global cylindrical LSHGC with an axis direction determined from the direction of the global limbs. Similarly, if all the local patches are conical with colinear limbs, then they form a global conical LSHGC whose apex is the intersection of the global limbs. Otherwise, the global axis (line) is detected. One way of detecting that axis would be to exploit the types of the component local SHGC patches. For example, the apexes of two conical patches (with non colinear limbs) theoretically determine the axis. Similarly, the apex of a local cone and the direction of a local cylinder also determine the axis. However, the axes so detected are sometimes not accurate (relying on few points), particularly in the presence of short limbs. Instead of this method, we use the more expensive but more accurate one of combining the recovered cross-sections of *all* component local SHGC patches in the same way discussed for the non-linear patches in 7.1 (i.e. using corollary P4) and fitting a line to the obtained axis points. Figure 7.6 illustrates this procedure.

- *conical* and *cylindrical*: a line is generated (apex and direction) ⁶
- *cylindrical* and *cylindrical*: the limbs must be colinear (for the same LSHGC), otherwise the directions must be parallel

Figure 7.5 shows some examples of geometrically compatible local SHGC patches.

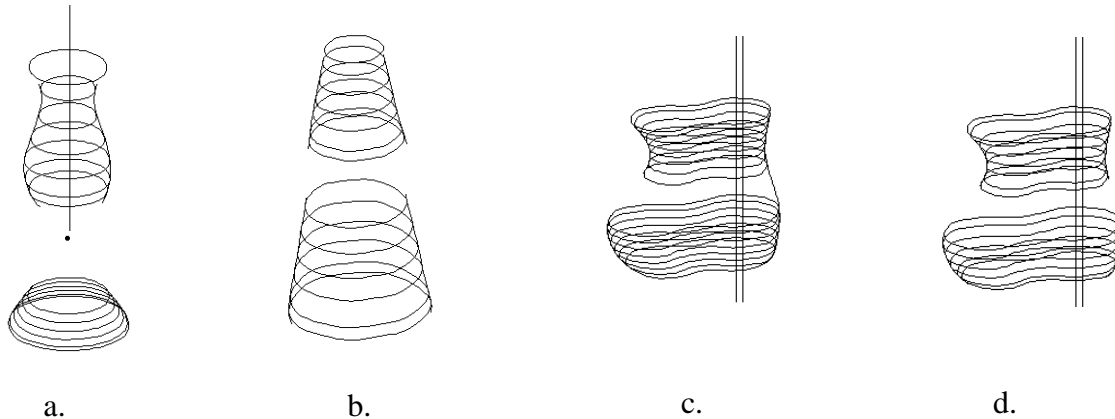


Figure 7.5 Examples of geometrically compatible local SHGC patches.

Structural compatibility involves measures of proximity and continuity between SHGC patches. We distinguish two cases: *continuous connection* where the patches share a limb curve segment as for the SHGC in Figure 7.5.c, and *discontinuous connection* where there is no common limb as in Figure 7.5.d. A connection hypothesis is generated between an SHGC patch and a geometrically compatible neighbor if the connection is continuous or is discontinuous with the constraint that the limb extremities are co-curvilinear *or* form self-occlusion ⁷. The co-curvilinearity measure uses looser thresholds than at the curve level since more information about the contours is given at this more global level.

Note finally, that a grouping of local SHGC patches implies a grouping of their limb curves. Therefore, gaps that have not been bridged in the curve level may be bridged at this more global level. For example, in Figure 7.5.b, grouping of the two patches of the cone connects the limbs, on the right side, which were judged unrelated in the curve level as could be noticed in Figure 5.2.b.

6. same remark for this line or equivalently the apex should belong to the global axis and the direction parallel to it.

7. self-occlusion can be detected by T-junctions and discontinuities in cross-section scaling between local SHGC patches.




Figure 7.4 Examples of hypothesized local SHGC patches (and their corresponding image contours) detected from the contours of Figure 1.1. a. the right hypotheses . b. examples of wrong hypotheses. The total number is 94 hypotheses, only those shown in a (4 patches) are the right ones.

- *non-linear and non-linear*: the axes must be (almost) colinear (Figure 7.5.c and d)
- *non-linear and conical*: the cone apex must lie on the axis (up to some error; Figure 7.5.a)
- *non-linear and cylindrical*: the direction of the cylinder must be (almost) parallel to the axis
- *conical and conical*: either the limbs are colinear (same apex as in Figure 7.5.b), or a line is generated (between to the apexes) ⁵

5. in this case, this line will be constrained to be colinear to the global SHGC axis or equivalently the apexes belong to it.

tions, a tangent line direction is less reliable than the direction of a line joining two separate (distant) points. Further, our method can be applied to $O(n^2)$ cross-sections (all combinations), providing more points for the voting process than the $O(n)$ property P3 would. Of course, corollary P4 can be applied only when the cross-section is detectable (though not necessarily complete) which we are assuming to be the case in this work.

In summary, hypothesized local SHGC patches are classified into three types each giving a local estimate of the axis:

- *cylindrical*: giving an estimate of the *direction* of the axis
- *conical*: giving the cone *apex* (belonging to the axis)
- *non-linear*: giving the projection of the *axis*

7.2 Grouping of Local SHGC Patches

The previous step may generate sparse local patches not all corresponding to perceived objects. Figure 7.4 shows some local SHGC patches detected from the contours of Figure 1.1 (the total number of such hypotheses is 94, only 4 of them correspond to real objects). It is impossible to identify, at the local patch detection step, which of those patches correspond to scene objects (*good* hypotheses as in Figure 7.4.a) and which are simply produced by irrelevant (or non-related) contours (*wrong* hypotheses such as the three patches of Figure 7.4.b). Further, due to breaks and occlusion, several local patches can be obtained for the same (global) SHGC. Therefore, selection of the good hypotheses requires grouping of relevant and related ones. In this step, grouping hypotheses are generated so as to form global patches providing global SHGC descriptions of real scene objects. A combination of local geometric and structural compatibility constraints is used for generating such grouping hypotheses.

Geometric compatibility is based on the projective properties of SHGCs (mainly corollary P4 and property P2). More precisely, it involves, between two local SHGC patches, compatibility of their estimated axes (the patches must refer to the same geometric description of an object). Several cases can occur depending on the types of the SHGC patches:



Figure 7.3 Finding local SHGC patches.

Pairs of candidate limbs having such a correspondence are checked whether they form local SHGC patches. This is done as follows.

- If the limb patches are straight, then a local LSHGC patch is hypothesized. The patch is further classified as being *cylindrical* if the two limb segments are parallel, or otherwise *conical*. A cylindrical patch, does not fix the axis uniquely. However, the parallel limbs give an estimate of its direction (corollary P4). Similarly, for a conical patch, the lines supporting the limbs intersect at the cone apex which belongs to the axis (also corollary P4). Cross-section recovery in this case is simple since all limb correspondence segments are parallel (property P5).
- If the limb patches are not both straight then corresponding points between the limb patches are identified using the limb-based recovery method. This yields a set of recovered cross-sections. Between *each pair* of such cross-sections having different scales, the intersection point of lines of symmetry is determined⁴ (Figure 7.2.c). A local SHGC patch is hypothesized if the locus of such points is a straight line (using fitting criteria). This line is a local estimate of the projection of the axis (corollary P4). We call this patch a *non-linear* SHGC patch.

We believe this method of finding axis points to be more robust and accurate than the method, used by Ponce et. al [1989], based on tangent lines (property P3), as the latter is sensitive to distortions of the limbs. In a sense, their method (property P3) can be thought of as the limit case of ours (corollary P4) where the two cross-sections get arbitrarily close to each other. Since the error in the slope of a line of correspondence is inversely proportional to the distance between the cross-

4. lines of symmetry are easily determined since the relative scaling and correspondence between the two cross-sections have already been determined.

7.1 Detection of local SHGC patches

Before describing the method, we define local SHGC patches.

Definition 3: A *local SHGC patch* is given by a hypothesized closed cross-section and a pair of corresponding limb curves (*limb patches*) satisfying the projective properties P2 or P4; i.e. the limb patches are either straight (for a local LSHGC) or have the property that lines of symmetry between any pair of projected cross-sections intersect on a straight line (projection of the axis). Figure 7.2 shows sample local SHGC patches.

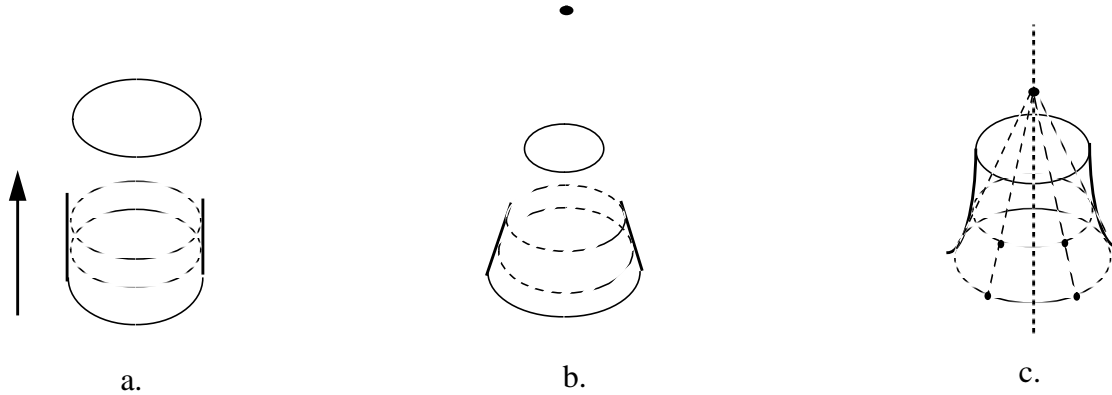


Figure 7.2 Sample local SHGC patches. a. Cylindrical patch. b. Conical patch. c. Non-linear patch

For each hypothesized cross-section, the method consists of finding limb curves having such a correspondence. Curve segments lying between the two curves of a parallel symmetry (involving the hypothesized cross-section) are considered as potential limbs. These latter are further classified in two sets lying on the two sides of the parallel symmetry. They can be thought of as a “left” and a “right” side. For each pair of such candidate limbs we check whether they form corresponding limb patches. Using the hypothesized cross-section, the correspondence can be found using a method that minimizes the scale of the cross-section³ joining corresponding points [Ulupinar 1990a] (call it *limb-based recovery method*). This correspondence is continuous and monotonic, thus invertible. Detection of candidate limb patches proceeds first by finding pairs of potential limb curves having an invertible correspondence at each of the mapped extremities (see Figure 7.3).

3. the scale is with respect to the hypothesized (completed) cross-section (which we will henceforth call “top” cross-section)

features are themselves difficult to detect and sensitive to image imperfections. However, even under heavy occlusion, local surface patches, or surface segments, can still be detected. A fragmented surface can then be recovered by *grouping* those surface patches whenever there is evidence that they project from the same global surface. In this level, we use a hypothesize-verify process of several steps in order to detect *local SHGC patches* and generate grouping hypotheses of these patches. The constraints used in this process are based on projective invariant properties discussed in section 3. Figure 7.1 shows the block diagram of this level. The steps are discussed in detail in the following sections.

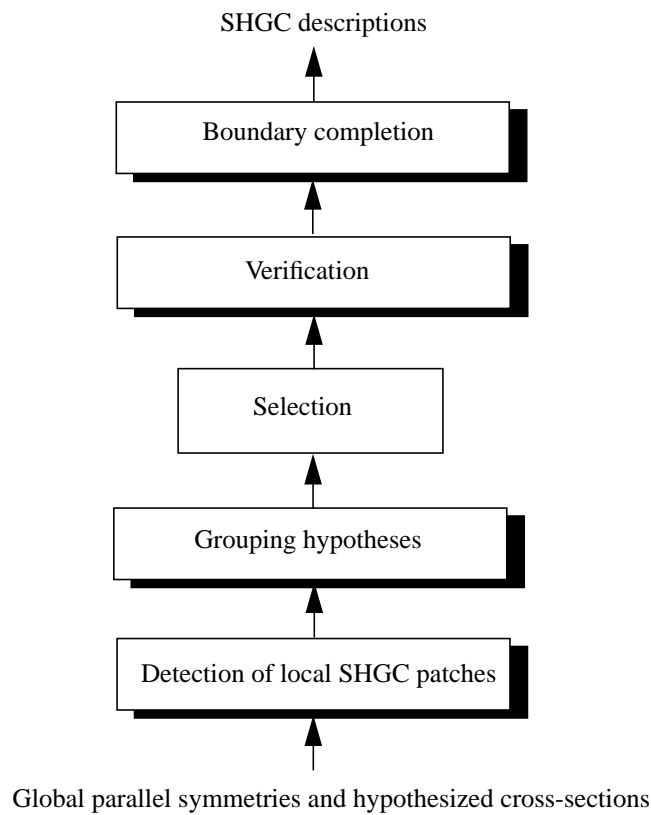


Figure 7.1 Block diagram of the SHGC patch level