

Local Image Features Resulting from 3-Dimensional Geometric Features, Illumination, and Movement: II*

James Damon[†], Peter Giblin[‡], and Gareth Haslinger[‡]

Abstract. This is the second part of an investigation into the visual clues in illuminated scenes, in terms of the interactions between apparent contours, shade and cast shadow curves, boundary edges, and markings on smooth surfaces, pairs of surfaces meeting in a crease, and triples of surfaces meeting in a corner. We consider both “stable” and “codimension 1” cases, the latter meaning that we list events which occur in a generic “fly-past” of the scene. We assume that there is a single principal source of light. The first part of this work is [*Internat. J. Comput. Vision*, 82 (2009), pp. 25–47]; in this second part we give details of the cases which involve creases and corners, and the “multilocal” cases where two surfaces separated in space interact via occlusion or cast shadows. We also give some details of the mathematical background to our work; the full mathematical treatment will appear in [Damon, Giblin, and Haslinger, *Characterizing Stable Local Features of Illuminated Surfaces and Their Generic Transitions from Viewer Movement*, in preparation].

Key words. illuminated surface, piecewise smooth surface, geometric features, shade, shadow, viewer movement, stable configuration, generic transition, singularity theory

AMS subject classifications. 58K05, 58K15, 68T45

DOI. 10.1137/090771053

1. Introduction. For a scene illuminated by a single light source, a number of visual clues about the shapes and positions of objects in the image are provided by the interaction of the geometric features of the objects (F), the curves bounding the shade/shadow regions on the objects (S), and the (apparent) contours (C) resulting from the viewing direction. (We use orthogonal or parallel projection here.) Typically the clues are obtained from the local configurations which result from the interaction of one or more of these ingredients. These visual clues then allow us to differentiate between objects and determine their shapes and positions. In [8], we gave the first part of the classification for fixed light source of the configurations resulting from the stable interactions of any subset of these ingredients (such configurations persist under small movements of the viewing direction). We also gave the first part of the classification of the generic changes which can occur as a result of movement of the viewpoint. This included the generic transitions of apparent contours (C), and interaction of apparent contours with either shade/shadow curves (SC) or with geometric features (FC), with the one exception of the transitions resulting from apparent contours and corners.

*Received by the editors September 14, 2009; accepted for publication (in revised form) January 7, 2011; published electronically March 29, 2011. This work was partially supported by an Insight 2+ grant from the European Commission.

<http://www.siam.org/journals/siims/4-1/77105.html>

[†]Department of Mathematics, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-3250 (jndamon@math.unc.edu). The work of this author was partially supported by National Science Foundation grants DMS-0405947 and DMS-0706941.

[‡]Department of Mathematical Sciences, The University of Liverpool, Liverpool L69 7ZL, England (pjgiblin@liv.ac.uk).

In this paper we complete the classification of generic codimension 1 transitions. These are the transitions in the configuration of features, shade curves, or contours which one would expect to see when the view direction changes continuously, moving in a general direction. The additional codimension 1 cases covered here include the interaction of apparent contours with corners (FC) (section 3.2); the interaction of all three ingredients (SFC) (section 4.1, further detailed in sections 4.2 and 4.3); and the interactions for the multilocal case (section 5.1), which involves interactions via cast shadows or occlusions. The complete list of codimension 1 transitions is given in section 2.1. As one consequence of the classification of transitions involving apparent contours and corners, we will see that optical illusions involving corners are resolved from the transitions arising from movement of viewpoint (see, e.g., parts (i) and (ii) of Figure 6).

Because the viewsphere is 2-dimensional, at isolated viewpoint positions special transitions can occur which are difficult to observe. These are called “codimension 2 transitions.” To understand all the nearby configurations near such points, the viewpoint must move in a small circle about such a special view position. In section 6, we briefly discuss the semiswallowtail transition from Table 3 of [8]. The qualitative changes can be a subtle succession of generic transitions, and for this reason we do not attempt to give the full analyses for cases of codimension 2 which we have listed.

Also, to be precise, when we speak of classifying either the stable configurations or the generic codimension 1 changes under viewer movement, we mean allowing “equivalence up to applying local diffeomorphisms” which preserve the geometric features and the shade/shadow curves. We explain in section 7 how the methods of singularity theory may be used to carry out the classifications when the local diffeomorphisms preserve both geometric features and shade/shadow curves while capturing the viewing direction.

As for the earlier parts of the classification given in [8], these results are consequences of mathematical theorems proven in [9].

In [8] and the present article we concentrate on describing the classification of stable phenomena and also those “codimension 1 phenomena” which occur at isolated moments during a “fly-past,” when the view direction changes smoothly and in a generic way. Our initial motivation has been to correct and augment previous classifications, using rigorous mathematical methods. Here we give some preliminary thoughts on how our results might be used. It is likely that the stable phenomena will be easier to apply than those of codimension 1.

The “edges” which we include in our classification are of several kinds, some much easier to detect in images than others. Apparent contours, cast shadows, surface markings, creases, and surface boundary edges (which are a mathematical abstraction from sharply curving thin surfaces producing a sharp apparent contour) are relatively well defined, but shade curves, where the main illumination grazes the surface, producing a gradual diminution of brightness, are more problematic.¹

Junctions between these various features are notoriously difficult to localize, and edge

¹We understand that Dvir Haviv, a student of Yosef Yomdin, has experimented with new methods of detection [13], and Amir Tamrakar, a student of Benjamin Kimia, has investigated the luminance profiles characteristic of the difference edge types (personal communication).

detectors commonly trace continuations of edges through a junction incorrectly, matching, say, an apparent contour to a cast shadow. By listing all the valid junction types and how they can evolve during a fly-past, our classification reduces the number of possibilities and thereby reduces the number of errors in identifying continuations of detected edges through a junction. Having identified the likely combination of edges meeting at a junction, we can then either identify the type of junction or at least list the few possibilities for the type of junction, that is, describe the type of feature point that it is in the image and what it represents in the original scene.

2. Codimension 1 generic transitions. In part I we gave the classification for both stable configurations of geometric features (F), shade/shadow curves (S), and apparent contours (C) and their interactions (FC), (SC), (SF), and (SFC). We also gave the codimension 1 transitions for (F), (SC), and (FC) (excluding the (FC) corner transitions). Such codimension 1 transitions are the “generic transitions” for those cases. This means that these are the transitions we expect to see if we move our viewing position and direction along a curve in 3-space. Here we complete this list by including the remaining generic transitions for each of the remaining four types of configurations already listed.

In contrast with the earlier cases in [8], for a number of these remaining cases the unstable configurations strictly have codimension greater than 1 when we consider the configurations and the way that they change up to a local diffeomorphism of the image. However, the *qualitative* behavior is simpler and is captured by a single parameter, as with a viewpoint moving on a curve in the viewsphere. Hence, we still refer to them as codimension 1 generic transitions.

2.1. List of codimension 1 generic transitions. The codimension 1 generic transitions include the following classes given below: (1)–(6) already given in [8], together with additional codimension 1 generic transitions (FC) for corners (7); the cases (8)–(10), which are generic (SFC) transitions; and two general classes of multilocal transitions, (11)–(12).

Transitions involving (SC) and (FC).

- (1) Semicusp: (SC) and (FC); see Statements 5.1 and 6.1 in [8].
- (2) Semi lips/beaks: (SC) and (FC), *ibid.*
- (3) Boundary cusp: (SC) and (FC), *ibid.*
- (4) Light direction cusp—view fold (SC); see Statement 6.3 and Table 4 of [8].
- (5) Nontransverse semifold (FC) (type iii) in Statement 6.6(2) of [8].
- (6) Fold apparent contour passing over an isolated stable geometric feature point (FC): in Statement 6.6 of [8].
- (7) Corner transitions (FC) (see section 3.2): This collection of generic transitions occur when a fold contour generator curve passes through the corner point. The classification is based on the classification of transitions for crease curves/contours (which initially ignores occlusions) given in section 3.1. This classification is then expanded by taking into account which possibilities occur for each of the four corner types, and then including visibility and illumination.

(SFC) transitions.

- (8) Notch or saddle corner with shadow transitions (SFC) (see type 4 of section 4.1, and section 4.3): Transitions occur when a fold contour generator curve passes a (notch

or saddle) corner which has a cast shadow curve. The classification is derived from that for corners (section 3.2), taking into account the extra cases resulting from the presence of the cast shadow curve. The transitions have higher codimension; however, the qualitative transition behavior is still derived from that in section 3.2.

- (9) Apparent contour passing isolated stable (SF) point (see types (1)–(3) of sections 4.1 and 4.2): Such a point will be where a shade curve meets a marking curve, crease, or edge curve. The transition occurs when a fold contour generator curve passes over the isolated point in a generic way. This means that the contour generator is not tangent to the crease/edge/marking curve, and also the latter curve is smooth in all the views near the transition moment.
- (10) Shade or cast shadow curve from an edge or crease (see sections 4.1 and 4.2): There can be some combination of a shade curve, a cast shadow curve from the shade curve, or cast shadow curve from the edge or crease which end at the edge or crease. A fold apparent contour generator passes the meeting point of these curves at the edge or crease.

Multilocal transitions.

- (11) Multilocal case for cast shadows (see class (1) of section 5.1): An isolated point results from a stable cast shadow for either a V-point or from a cast shadow curve transversally meeting an edge, marking curve, ridge crease, or shade curve. The transition involves a fold contour curve moving over the isolated point.
- (12) Multilocal case for occlusions (see class (2) of section 5.1). There are two general types: (i) the curve of an occluding object (edge, crease, or apparent contour) passes over a stable isolated point for (SF)); (ii) the occluding curve becomes tangent to a shade or feature curve, and the transition occurs as it passes through the tangent point.

We shall make occasional reference to codimension 2 in what follows, and one example is touched on in section 6. However, from now on, unless explicitly stated otherwise, the term *generic transition* will always refer to *codimension 1 generic transition*.

3. Classification of generic transitions for corners and contours (FC). First, we complete the list of generic (FC) transitions by considering the case where the features are the three creases meeting at a corner. Then, we are considering the interaction of these creases with an apparent contour on one of the three sheets of the corner. We note that while apparent contours can occur on more than one sheet which meet in a corner, it is a consequence of the analysis that when a generic transition occurs, the apparent contour on only one of the sheets interacts with the corner.

Call the three smooth surface sheets which meet at a corner P , Q , and R . They meet at a common point, the corner point, and P , Q , R meet pairwise in three crease curves which themselves meet at the corner. As explained in [8, section 1], there are four basic *types of corners*: concave, convex, saddle, and notch, as illustrated in Figure 1. We will denote them respectively by Cc, Cv, S, and N.

The surface sheets divide space around the corner point into two regions, one thought of as a region occupied by an object and the other as “empty space.” A convex or notch corner has this empty space occupying *more* than a hemisphere of solid angle around the corner point, whereas a concave or saddle corner has it occupying *less* than a hemisphere.

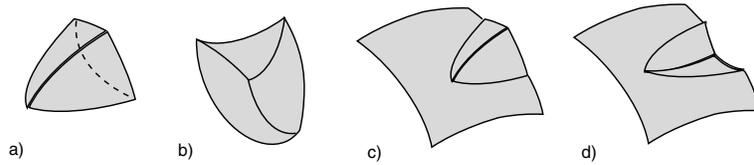


Figure 1. General (curved version) corner types: (a) is a convex corner (Cv), (b) is a concave corner (Cc), (c) is a saddle corner (S), and (d) is a notch corner (N).

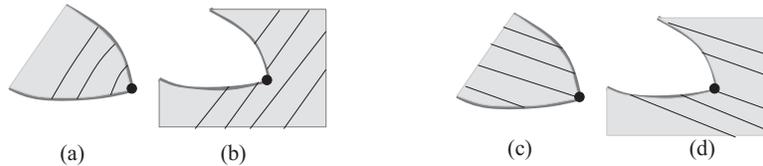


Figure 2. A schematic diagram of one of the three sheets, R say, making up a corner. The dark grey lines are the crease curves where the sheet R meets the other two sheets, P and Q , and the corner point where all three sheets meet is marked with a dot. The thin black lines are representations of contour generators on R which slide across R as the viewpoint moves. In (a) and (b) the contour generators meet both creases and then neither crease. Each intersection of crease and contour generator produces a crease semifold, in the image. In (c) and (d), on the other hand, the contour generators meet one crease at a time throughout the viewer movement. The pair represents (a) and (c) a concave or convex corner and (b) and (d) a saddle or notch corner. (a) and (c) are referred to as “one-quarter sheet” and (b) and (d) as “three-quarter sheet.”

Also, we first consider the case where the illumination is assumed to be uniform on each sheet, which means that a sheet is either entirely lit or entirely in shadow; this is the case of (FC). (Later in section 4 we consider the transitions where there is a cast shadow from one of the creases onto one of the sheets.) By the *illumination of a corner* we mean the assignment to each sheet of being entirely lit or in shadow. For example, the possible stable configurations of corners with illuminations are given in Figures 13, 14, and 15 of [8].

Third, there is the issue of *visibility*. Suppose the object is viewed so that we can see the corner; that is, we can see the point where the three creases meet (even though some of the creases and sheets themselves may not be visible). If the object were transparent, then we would see all of the crease curves and apparent contours. We refer to this as the *crease curve/contour configuration*. Changing the direction of view to the opposite direction may alter the visibility of some of the creases and sheets, even though the crease curve/contour configuration would remain unchanged.

Fourth, for the classification of codimension 1 transitions for corners (FC) we need consider only the case where *one* of the three surface sheets, say R , has an apparent contour in the direction of view (this is the case of a “C-semi-fold” in [19, 20]). The properties of this apparent contour can change as we move the viewpoint along a curve in space. A fourth fundamental distinction among these transitions can be explained with reference to the surface R and its two crease curves (where it meets P and Q ; see Figure 2). We can extend R to be part of a smooth surface, and the contour generators on R can be extended to curves in the smooth surface which also projects in the viewing direction with contour generators given by the extended curves. We obtain a different corner by replacing R by its complement in the smooth surface,

as illustrated in Figure 2. The complement then has its contour generator curves, which are the complements of those in the initial R . We shall refer to this process as taking the *complementary contour curves* in the complementary sheet. The contour generators can meet the crease curves of R in two different ways, as illustrated in Figure 2. Together with the complementary contours, this leads to one of the four situations illustrated schematically in Figure 2.

Hence, for a given crease curve/contour configuration for a corner, there are four factors which will alter the possible image of the corner: corner type, illumination, visibility, and complementary contours. The procedure for classifying transitions for corners in the case (FC) is to first classify the possible transitions for the crease curve/contour configurations, and then to further refine the classification by taking into account the four additional characteristics: corner type, illumination, visibility, and complementary contours.

Classification of corner transitions (FC) via crease curve/contour configurations. We separate (FC) transitions for corners into eight distinct cases by means of the following three properties.

Contour/crease number. In Figure 2 we show abstractly the two different ways in which a contour generator can move across a surface which forms one of the three sheets of a corner. We observe that the contour generators can meet the crease curves in either two points or zero points in cases (a) and (b), and then during the transition this number changes, or in exactly one point as in (c) and (d), and this number does not change during the transition. Also these numbers do not change if we replace R by its complementary sheet. We refer to the maximum number of meeting points as the *contour/crease number* for the transition. This number will be 1 or 2.

Crease direction. In the *image*, two of the three creases (and also the apparent contour) will become tangent at the moment of transition. This is because the view direction lies in the tangent plane to the sheet R , and therefore this plane is viewed as a single line in the image. Any curve in that tangent plane through the corner point will therefore be seen as tangent to this line. This includes the two crease curves which lie in the sheet R . The two creases in the image approach the corner point from either the *same* direction or the *opposite* direction. The *crease direction* is then denoted respectively by s or o .

Apparent contour position. Consider the property of whether the three creases in the image create a “reflex angle” (> 180 degrees). When two creases as above are tangent from the same direction, this property will hold throughout the transition; when they are tangent from opposite directions, the property will hold only on one side of the transition. Now follow this region through the transition and ask: is the apparent contour in the image entirely in this region throughout the transition? If “yes,” the *apparent contour position* will be denoted by y , and if “no,” by n .

Together, these three invariants provide a triple of values $(1/2, s/o, y/n)$ with eight distinct possibilities. In Figure 3 we give eight “basic” configurations which correspond to all eight possible triples. For any triple (a, b, c) , we then have an associated configuration, which we denote by $(a, b, c)^*$, obtained using the complementary contour curve configuration to that associated to (a, b, c) . This complementary configuration will have the same invariants. Then, these configurations yield a complete classification of the crease/contour configurations.

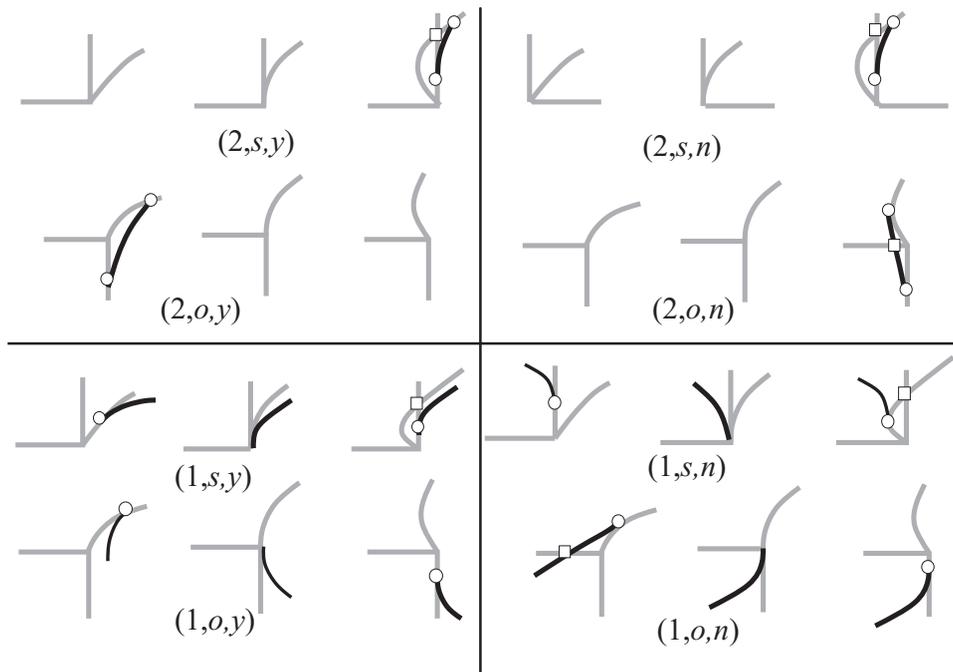


Figure 3. The eight basic crease curve/contour configurations for corner transitions. Creases are represented by grey curves, and apparent contours by black curves corresponding to the triples given in section 3.1. Open circles represent tangential contact, and squares represent transverse crossings. Transitions occur in each horizontal sequence. We do not take account here of corner type, illumination, or visibility, and each case has a “complementary” version, where the apparent contour is replaced by the complementary contour curve. One case, $(2, s, y)$, is given in full detail in Figures 4 and 5. The complete set is available at [4].

3.1. Classification of crease curve/contour configurations for (FC) corner transitions.

The classification of the 14 crease curve/contour configurations for generic (FC) corner transitions is given as follows:

- (1) The classification of “basic” crease curve/contour configurations for corner transitions (FC) corresponds exactly to the eight possible combinations of the triple of invariants $(1/2, s/o, y/n)$. They are illustrated in Figure 3.
- (2) The four complementary contour configurations corresponding to basic ones with contour/crease number 2 are the configurations which have apparent contours with two components (see, e.g., Figures 4, 5).
- (3) For those basic configurations with contour/crease number 1 and the same crease direction $(1, s, y/n)$ (which is the same direction as the shown apparent contours), the complementary contour configuration will have apparent contours with opposite direction from the creases.
- (4) For those basic configurations with contour/crease number 1 and the opposite crease direction $(1, o, y/n)$, the complementary contour configurations are equivalent to the original configurations.

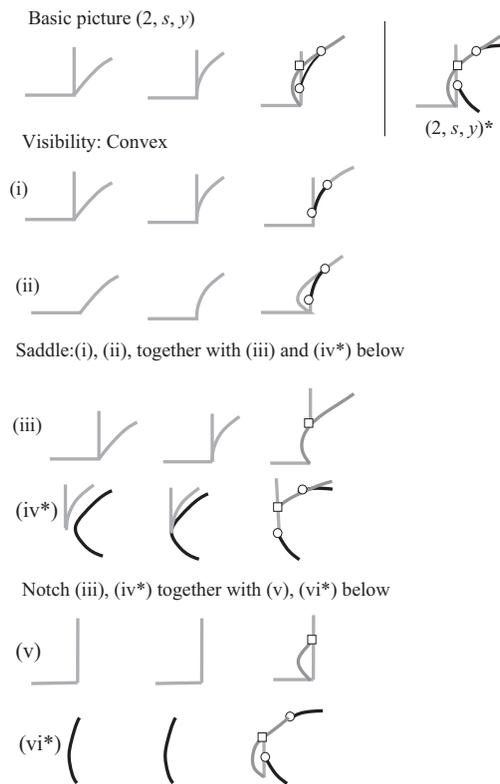


Figure 4. Complete list of transitions on corners of type $(2, s, y)$, showing visibility. The grey lines are the creases, and the black line is the apparent contour. The asterisk $*$ indicates that the apparent contour is of the “complementary” type.

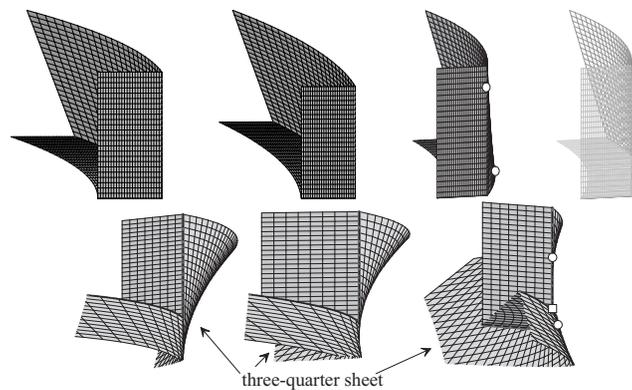


Figure 5. Top row: A transition on a convex corner of type $(2, s, y)$ (ii) (see Figure 4). Note that it is the arrangement of crease edges and contours which is important, not their shapes, when comparing actual examples with the schematic diagrams in Figure 4. The rightmost figure is a wireframe view of the figure to its left, showing the occluded self-intersection of creases in the image. Bottom row: Transition on a notch corner of type $(2, s, y)$ (iv*). See the accompanying animations [77105_01.gif](#) [1.3MB] and [77105_02.gif](#) [1.4MB].

Table 3.1

Corner types having transitions corresponding to crease curve/contour configurations.

Crease curve/contour configuration	Corner type	Illustrated
$(2, s, y)$	Cv, S, N	Cv Fig. 5
$(2, s, n)$	Cv, S, N	
$(2, o, y)$	Cv, S, N	Cv Fig. 7
$(2, o, n)$	Cv, Cc, N	Cc Fig. 6
$(2, s, y)^*$	S, N	N Fig. 5
$(2, s, n)^*$	S, N	
$(2, o, y)^*$	N	N Fig. 7
$(2, o, n)^*$	S, N	N Fig. 6
$(1, s, y)$	Cv, S, N	
$(1, s, n)$	Cv, S, N	
$(1, o, y) = (1, o, y)^*$	Cv, S, N	
$(1, o, n) = (1, o, n)^*$	Cc, Cv, S, N	
$(1, s, y)^*$	N	
$(1, s, n)^*$	S	

We can then combine the classification of the crease curve/contour configurations for corner transitions given in (3.1) with the additional three characterizing properties of corners, namely, corner type, illumination, and visibility, to give a classification of the generic transitions for corners and contours (FC).

3.2. Generic transitions for corners and apparent contours (FC). The generic transitions for local configurations involving corners and apparent contours (FC) can be classified as follows:

- (1) For each of the 14 crease curve/contour configurations, the corner types having transitions with the configuration are given in Table 3.1.
- (2) *Convex corner types* occur for all of the eight basic crease curve/contour configurations; however, they do not occur for the complementary contour configurations (except the self-complementary $(1, o, y/n)$).
- (3) *Concave corner types* occur only for the crease curve/contour configurations $(2, o, n)$ and $(1, o, n)$.
- (4) *Saddle corner types* occur for all eight configurations, though in the case of type $(2, o, n)$ only for the complementary configuration $(2, o, n)^*$, and in types $(2, o, y)$ and $(1, s, y)$ only for the standard configuration.
- (5) *Notch corner types* occur for all eight configurations, though in the case of type $(1, s, n)$ only for the standard configuration, not for $(1, s, n)^*$.
- (6) The number of cases is large, and we have not attempted to illustrate all of them here. However, a table of all cases of visibility, together with more examples of actual surfaces, is available in [4]. The full list of visibility cases for type $(2, s, y)$ appears in Figure 4 with illustration in Figure 5; other cases are illustrated more briefly in Figures 6 and 7. See Table 3.1.

See section 7 for an indication of how in [9] we arrive at this exhaustive list by means of realization of abstract forms of singularities.

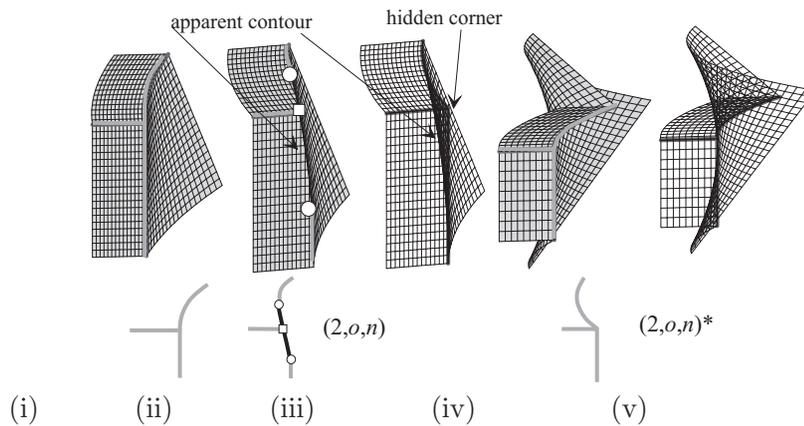


Figure 6. (i)–(ii): A concave corner transition with crease curve/contour configuration $(2, o, n)$, together with the corresponding schematic diagram from Figure 3, but with visible parts only indicated. Notice that it is not possible to determine from (i) whether the corner is concave or convex, but (ii) settles that it must be concave. (iii) A wireframe view of (ii) showing the hidden corner. (iv) A notch corner of type $(2, o, n)^*$ at the transition moment, obtained by replacing the curved “quarter-sheet” in (i) with a “three-quarter sheet.” Note that all creases are visible, but both parts of the broken apparent contour are occluded. (v) A wireframe view of (iv). See the accompanying animations [77105_03.gif](#) [2.0MB] and [77105_04.gif](#) [1.3MB].

Remark 3.1. The cases of concave corners were originally classified by Tari [19, 20], who concentrated on analyzing one type of transition (for the case of crease/contour number 1) and gave the normal forms for the equations. That classification extends here to the four types of corners, yielding the classification given in section 3.2.

4. Classifications of generic transitions for triple interactions (SFC). Next, we give in subsection 4.1 the local classification for the interactions of all three ingredients. This will be further divided into subclassifications in terms of distinct geometric features.

4.1. Generic transitions for configurations involving all three geometric features, shade or shadow curves, and apparent contours (SFC). First, the configurations of geometric features and shade/shadow curves are stable and are given by Figures 9, 10, 12, 15, and 16 of [8]. In addition, there is one case of a shade curve and a cast shadow from a boundary edge which cannot be simultaneously seen as a stable view. Second, if we ignore the shade/shadow curves, the underlying interaction of geometric features and apparent contours belongs to the classification in Statement 6.6 of [8]. Third, this classification is then refined by taking into account the shade/shadow curve configuration. The classification of generic transitions for local configurations involving all three geometric features, shade/shadow curves, and apparent contours is given as follows.

- (1) **Marking curve:** The marking curve stably intersects a shade curve transversely (i.e., nontangentially). The view projection is a fold mapping at the intersection point, whose tangential and kernel directions are distinct from the tangent lines for the other two curves. The generic transition corresponds to movement of the fold curve (apparent contour) from the intersection point. See Figure 8.

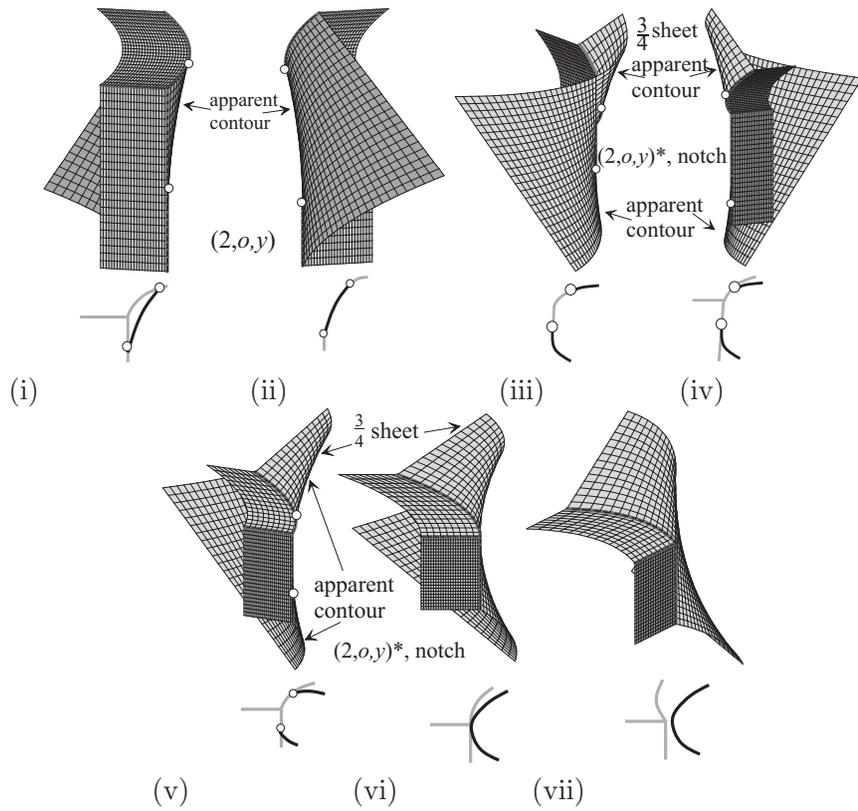


Figure 7. The schematic diagrams, from Figure 3, show only the visible creases and contour. Parts (i) and (ii) are opposite views of the same convex corner involved in the transition for $(2, 0, y)$; this transition cannot occur for concave corners. Parts (iii) and (iv), which are also opposite views, show a notch corner of type $(2, 0, y)^*$ (saddle corners cannot be of this type). They are obtained from (i) and (ii), respectively, by replacing the one-quarter curved sheet with a three-quarter sheet. Parts (v)–(vii) show a transition on a notch of type $(2, 0, y)^*$; note that (v) is qualitatively the same as (iv), with all creases and both parts of the broken contour visible. See the accompanying animations [77105_05.gif](#) [1.4MB], [77105_06.gif](#) [1.2MB], and [77105_07.gif](#) [1.2MB].



Figure 8. SFC transition involving the interaction of a marking and a shade curve with an apparent contour. The marking curve (M) is grey, the contour (C) is black, and the shade curve (S) is also black. Dashed curves denote occluded curves.

- (2) Edge curve: The stable configurations involve a shade/cast shadow curve meeting the edge curve transversely (this is illustrated in Figure 10(g) and (h) of [8]). A codimension 1 transition occurs when a fold contour generator curve moves over the meeting point of the edge with the shade/shadow curve. This transition is analogous to that occurring on one sheet of a crease (see the analogous point (3) below and section 4.2). Figure 9 shows the four possible configurations in which this can happen for a shade curve, without regard to visibility, and Figure 10 takes visibility into account.

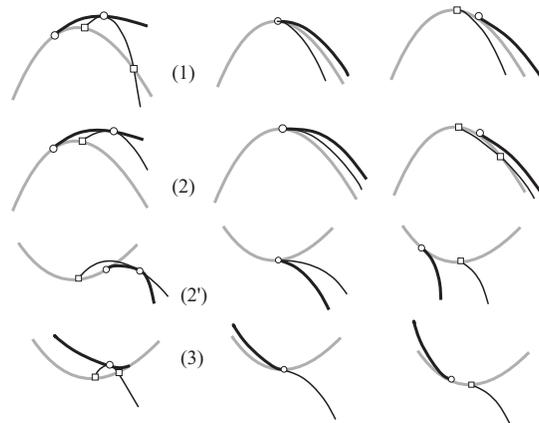


Figure 9. SFC transitions involving the interaction of an edge curve and a shade curve with an apparent contour. These are the four basic cases, not taking account of visibility. They will be referred to as (1), (2), (2'), and (3), as here. See also Figure 10, which allows for visibility. The edge curve is grey, the contour is a thick black line, while the shade curve is a thin black line. These also apply to a crease where the shade and the contour are on the same sheet.

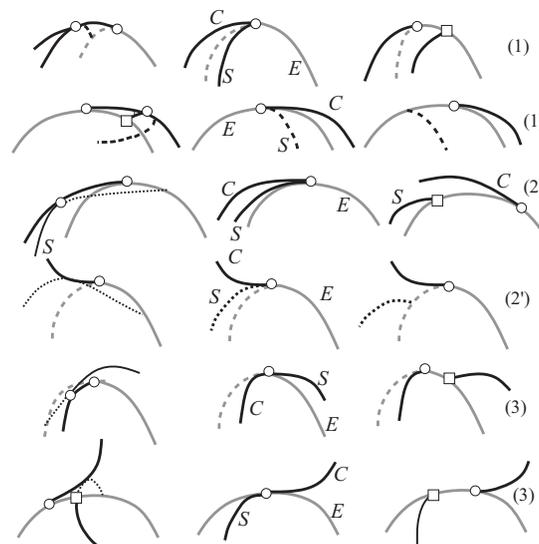


Figure 10. SFC transitions involving the interaction of an edge curve and a shade curve with an apparent contour, showing visibility. The edge curve (*E*) is grey, the contour (*C*) is black, and the shade curve (*S*) is also black. Dashed curves denote occluded curves. The numbers in brackets refer to the type, as in Figure 9.

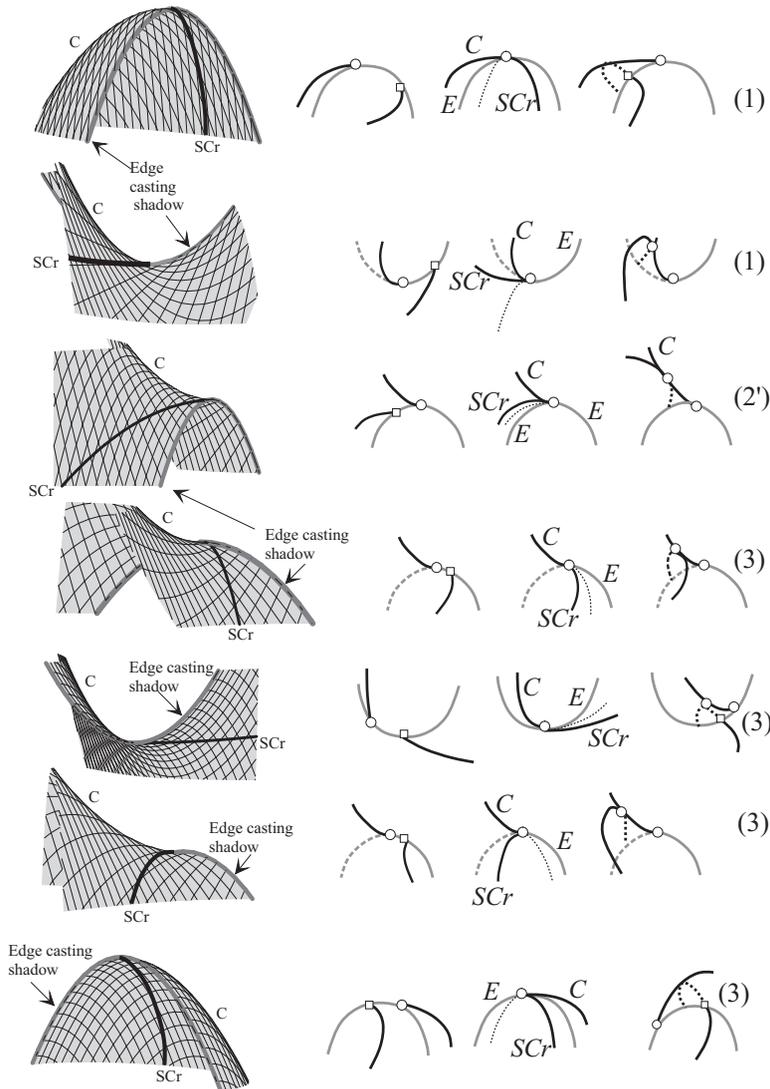


Figure 11. *SFC transitions involving the interaction of an edge curve and its cast shadow with an apparent contour. The cast shadow is marked SCr (“shadow of crease”) since this also applies to a crease surface when the second sheet does not occlude the features shown. The thin dashed line in the middle “transition” drawing of each row is the shade curve, which is always occluded. In the first, third, fifth, and seventh rows a very small portion of the shade curve becomes visible when the view changes. This is illustrated for the seventh row in Figure 12. See also Figure 13.*

The cases where the edge throws a cast shadow are illustrated in Figure 11; see also Figure 12. There is a further exceptional case where a shade and a cast shadow from the edge occur on opposite sides of the same sheet, meeting at a single point on the edge. Both curves meeting at the common edge point are not stably visible and only become so at a transition, where the point also is a point on the contour generator.

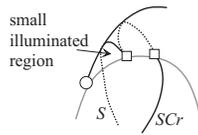


Figure 12. In the bottom row of Figure 11 a small portion of the shade curve (S) becomes visible when the view is moved, and a small illuminated region appears. This also happens for the first, third, and fifth rows of the same figure.

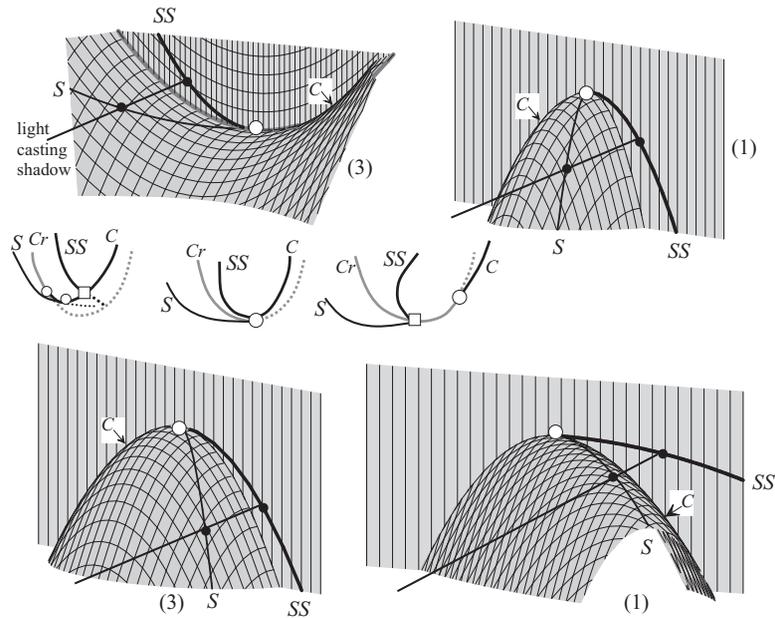


Figure 13. Each surface with boundary edge having an SFC interaction of edge, apparent contour C , and shade curve S can be augmented by a second sheet to make a crease, and this falls within our classification, provided that the second sheet has uniform lighting (i.e., no shade curve) and no apparent contour. In most cases the addition of a second sheet merely occludes features such as those in Figure 11 in an obvious way. There are a few cases, exemplified here, where the shade curve on the first sheet can throw a cast shadow SS on the introduced second sheet. The number in brackets is the type, as in Figure 9. The transition sketch is for the top left case; the others are similar.

- (3) Crease curve meeting a shade/shadow curve: The stable (SF) interactions involving creases are given in Figure 12 of [8]. In each of these cases there is a distinguished point on the crease where either a shade curve (possibly with a cast shadow curve from the shade curve) or a cast shadow from the crease meets the crease—in (d) of that Figure 12 both shade curve and its cast shadow occur. The generic transitions for the shade curve on one sheet correspond to the possible cases where a fold apparent contour generator on either sheet passes the distinguished point. (The other cases, listed in Statement 6.6(2) of [8], will have higher codimension in the presence of the shade/shadow curve.) These possible transitions are given in section 4.2; the illustrations for the case when the contour and shade curve lie on the same sheet are Figures 9–13; in particular, Figure 11 shows the cases where there are shade curve, cast shadow from the crease,

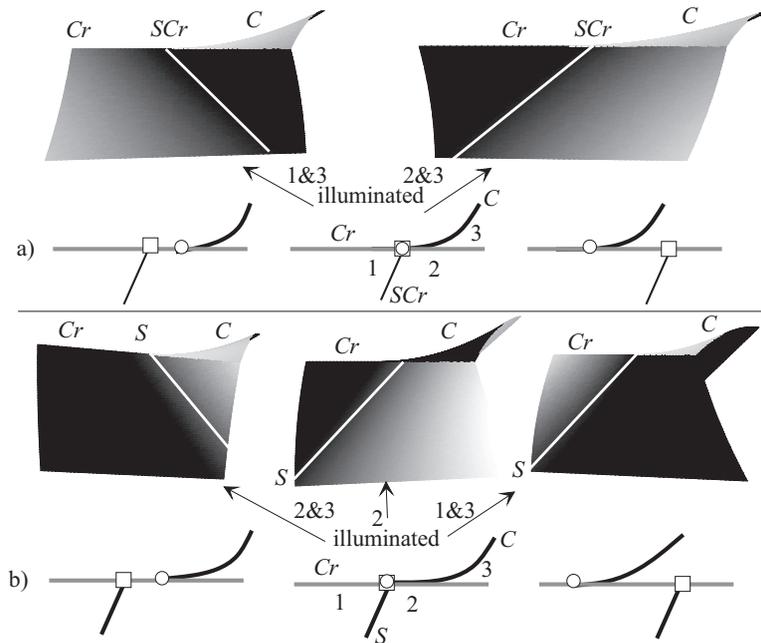


Figure 14. *SFC generic transitions for a ridge crease, cases (a) and (b). (See section 4.2.) The schematic diagrams of each case illustrate the transition, as the cast shadow of the crease (SCr, case (a)) or the shade curve (S, case (b)) moves across the point where the crease (Cr) meets the apparent contour (C). The shaded figures illustrate, for the transitional moment, the possible illuminations of the regions on the crease. The cast shadow and shade curve are artificially emphasized in white in these shaded figures.*

and contour all on the same sheet. When shade/shadow and contour lie on different sheets the relevant images are in Figures 14, 15, and 16.

- (4) **Corner:** The stable (SF) cases for corners occur when for a notch or saddle corner one of the crease curves casts a shadow on one of the sheets meeting at the corner. These include the cases for saddle corners, Figure 15(e), (f), and (g) in [8] and for notch corners, Figure 16(i), (k), and (l) there. The generic transitions occur for these cases when a fold contour generator on one of the sheets moves across the corner point. The classification of these possibilities is given in section 4.3 and involves a refinement of the classification of corner transitions involving notch or saddle corners given in section 3.2 and Table 3.1. See Figures 17 and 18. A complete list of the visibility possibilities for the notch and saddle cases is available in [4].

4.2. SFC generic transitions involving crease curves. The generic (SFC) transitions involving a crease curve and shade/cast shadow curve are obtained from the stable (SF) interactions involving creases given in Figure 12 of [8]. The generic transitions correspond to the possible cases where a contour generator on one of the crease sheets, which gives a fold contour in the image, moves over the distinguished point on the crease where a shade/cast shadow curve meets the crease. The cases are distinguished by whether there is a shade curve or cast shadow curve, or both, and whether the contour generator is on the same sheet as the shade/shadow curve. The cases are then given as follows:

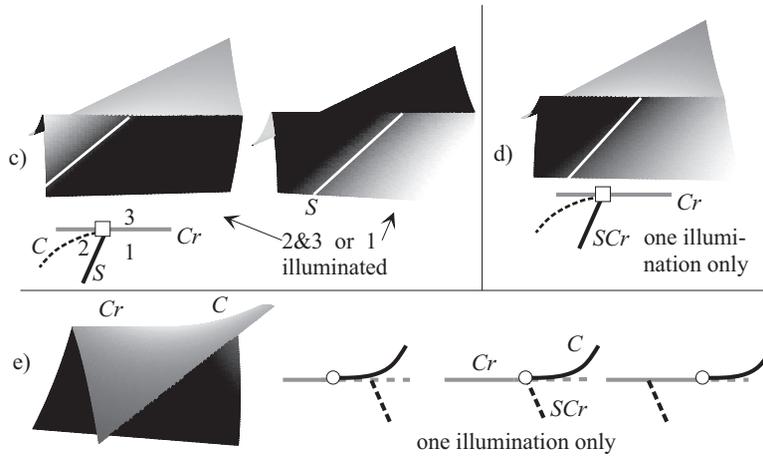


Figure 15. SFC generic transitions for a ridge crease, cases (c)–(e). (See section 4.2.) In these cases there is no visible change since either the contour is occluded (cases (c) and (d)) or the cast shadow of the crease is occluded (case (e)). For the notation, see Figure 14.

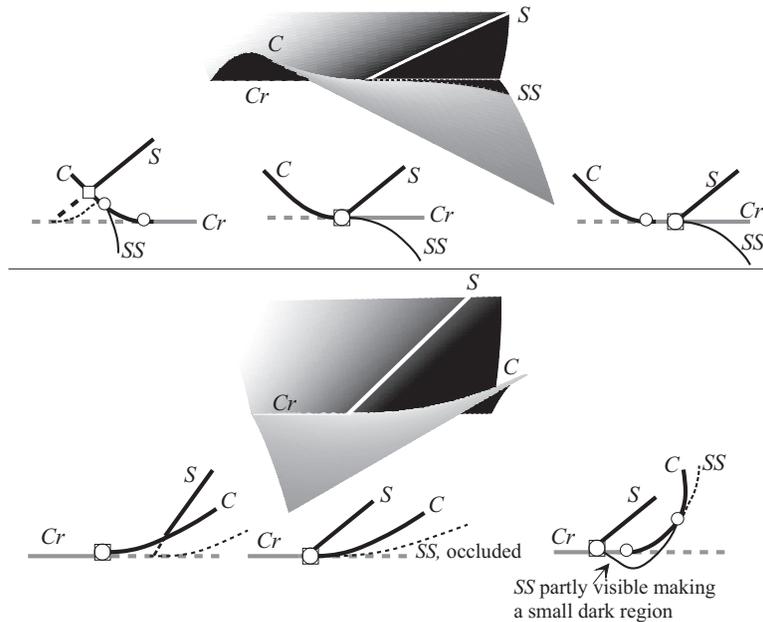


Figure 16. SFC generic transitions for a valley crease. The shade curve (S) (here artificially emphasized in white) on one sheet of the valley crease casts a shadow curve (SS) on the other sheet. The crease (Cr) is again the grey curve. Again the transition results from the movement of the apparent contour passing the meeting point of (S) and (SS) with (Cr).

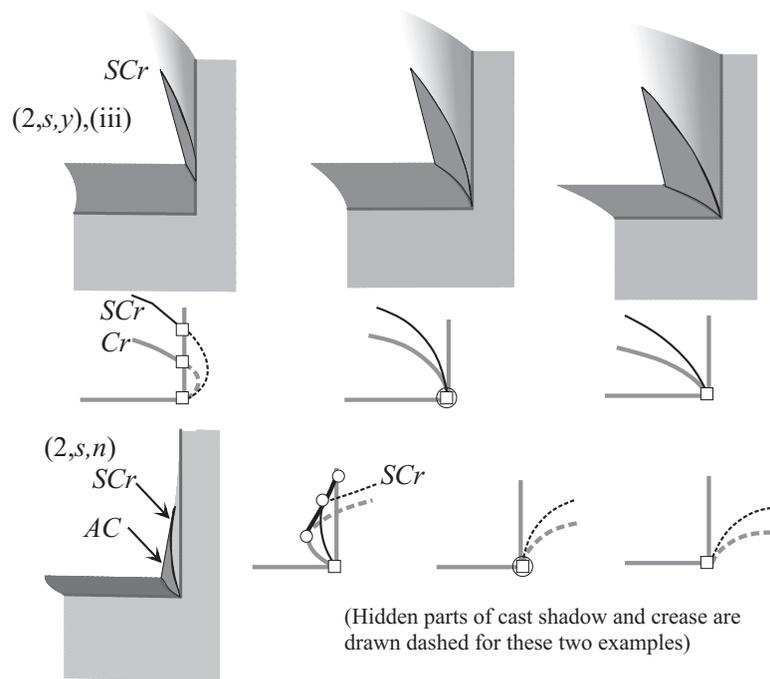


Figure 17. Two examples of notch corners with cast shadow transitions, of types $(2, s, y)$ and $(2, s, n)$. SCr = cast shadow of a crease, AC = apparent contour, Cr = Crease. The label $(2, s, y)(iii)$ refers to Figure 4. A complete list with visibility indicated is available at [4].

- (1) Contour generator and shade curve on different sheets (ridge crease): For ridge creases [8, Fig. 12 a) and b)], there are two cases for each where the contour generator is on the sheet without the shade curve. For one of these the apparent contour is visible, and for the other it is not visible (see cases (b) and (c), respectively, in Figures 14 and 15).
- (2) Contour generator and shade curve on same sheet: There are a number of cases here; for the case where there is also a cast shadow from the crease falling on the same sheet, see the next item. The second sheet plays no role except possibly to occlude the transition, so this is covered by the corresponding edge case of section 4.1(2). The transition occurs when the contour generator moves past the distinguished point where the edge meets the shade curve. Figure 9 show the basic situation without regard to visibility, while visibility is allowed for in Figure 10.
- (3) The same as (2), but with a cast shadow from the edge (or crease) on the same sheet: The transition occurs when the distinguished meeting point of cast shadow and shade curve is passed by the contour. However, cast shadow and shade curve are not simultaneously visible, except in certain cases on one side of the transition. See Figures 11 and 12.
- (4) The same as (2), but with a shade curve on one sheet throwing a cast shadow on the second sheet of a valley crease. See Figures 13 and 16.

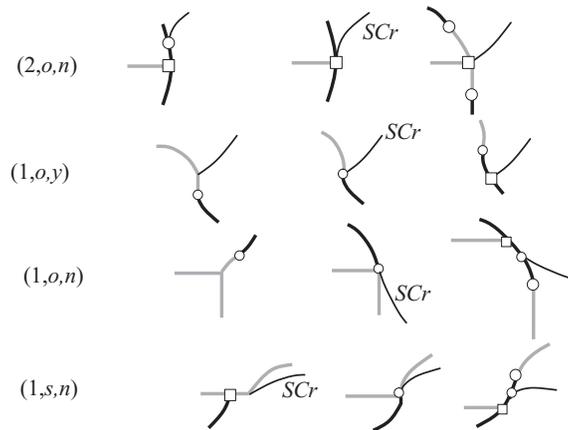


Figure 18. Schematic pictures of all the saddle corners with cast shadow transitions, showing visibility. SCr = cast shadow of a crease, and the creases are grey, the apparent contour black. The labels refer to Figure 3.

- (5) Contour generator and cast shadow curve on different sheets (ridge crease): For ridge crease [8, Fig. 12 e)], the crease casts a shadow on one of the sheets. The contour generator can occur on the sheet without the cast shadow. The contour generator can be either visible for case (e) or not visible for case (f); see (a) and (d) in Figures 14 and 15; also the cast shadow can be invisible as in (e) of Figure 15.

Next we give the corresponding generic (SFC) transitions for corners.

4.3. SFC generic transitions involving corners. The stable (SF) cases for corners occur when, for a notch or saddle corner which has a cast shadow on one of the sheets, an apparent contour generator on one of the sheets moves over the corner point, as stated in section 4.1. For a notch or saddle corner the types $(2, o, n)$, $(1, o, n)$, $(1, s, n)$, and $(1, o, y)$ can all give rise to this transition. The visibility diagrams for the saddle case are in Figure 18, and illustrations of three notch cases are in Figure 17. Full visibility diagrams for the saddle and notch cases are available in [4].

5. Classifications of generic transitions: The multilocal case.

Stable multilocal classification. Before considering the generic transitions in the multilocal cases, we first refer to the classification of stable multilocal configurations. These were given in [8, Fig. 20] to provide a complete list of stable configurations and were briefly discussed in section 2.5 of that paper. These classifications arise from two distinct occurrences: the occlusion of one geometric feature or apparent contour by one from a distant object (or part of the same object) and the intersection of a cast shadow (from a distant object or part of the same object) with a geometric feature or apparent contour.

The occlusion results from the partial occlusion of a marking curve, edge curve, crease curve, fold apparent contour, or shade/shadow curve by a region of an object bounded at the occlusion point by either an edge, ridge crease, or apparent contour. Genericity implies that the occluding and the occluded curves meet nontangentially. This is the “Hard T” in [8, Fig. 20] and is traditionally referred to as a “T-junction.”

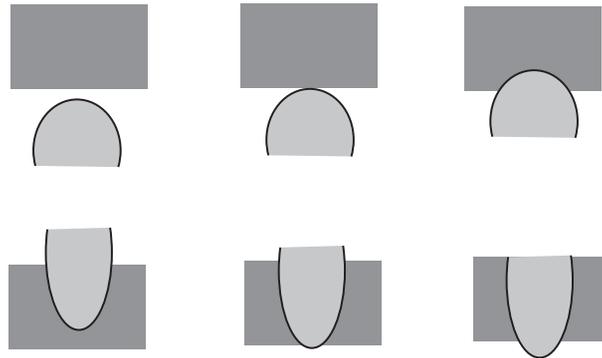


Figure 19. Abstract representation of the tangency transitions 1 and 4 in section 5.1, where for 4 the configuration curve which is behind would denote the cast shadow curve. The transitions can occur in either direction. They are analogues of the lip-beak transitions given in [8].

Because the light source and object are fixed, we assume that the cast shadow from a distance meets any other geometric feature generically. There are two general classes of possibilities. One is that the cast shadow is a smooth curve which cuts nontangentially across a marking curve (the “Hard-Soft X” [8, Fig. 20]), both surfaces meeting in a crease curve (the “Hard-Soft Broken X” [8, Fig. 20]), edge curve, one surface of a crease curve, or apparent fold contour (all “Soft T” [8, Fig. 20]). Also, generically the cast shadow curve does not intersect isolated points such as corner points. The second possibility is that the shadow is cast by a geometric feature such as a V -point from a corner or cast shadow curve intersecting a shade/shadow curve (“Soft V” [8, Fig. 20]). The cast shadow of the vertex of the V lies in a smooth part of a surface, that is, not on a crease, edge curve, or marking curve, with the shaded region filling the interior or exterior of the V . These possibilities lead to the classification given in [8, Fig. 20].

Generic multilocal transitions. With the knowledge of the stable multilocal configurations, we then complete the catalogue of generic transitions by giving those for the multilocal cases. These arise as transitions for the two distinct classes of occurrences involving occlusion or cast shadows. These are given by the following classification.

5.1. Generic transitions for multilocal configurations. The generic transitions for multilocal configurations involving some combination of geometric features, shade/shadow curves, and apparent contours are given as follows:

Occlusion.

- (1) Configuration curves meeting tangentially: The image of a smooth configuration curve, namely, an edge curve, ridge crease curve with one sheet visible, or apparent contour curve, for an object meets tangentially the image of another smooth configuration which may also be a marking curve. As the tangency disappears, one curve occludes a segment of the other, as in Figure 19 (this is analogous to the semilips transitions for the local case in Table 3 of [8]).
- (2) Configuration curve moving across an isolated configuration point: A configuration curve which is an edge curve, ridge crease curve with one sheet visible, or apparent

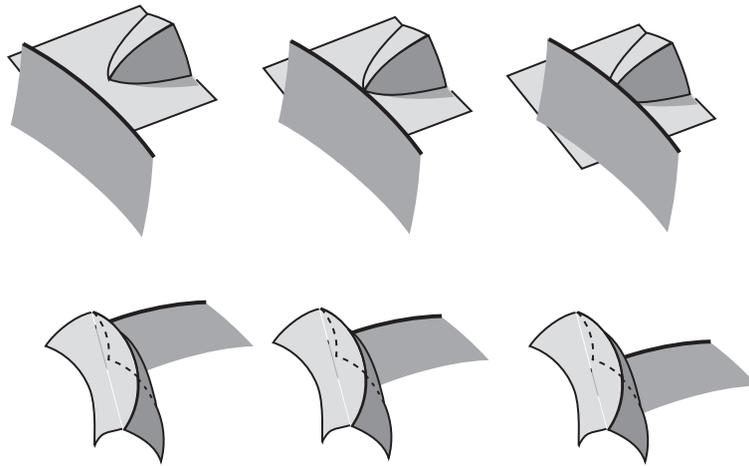


Figure 20. Two examples of the transitions involving a configuration curve meeting and moving past and occluding an isolated configuration point (2) given by (a) or an isolated configuration point moving across a configuration curve (3) given by (b) in section 5.1. The transitions can occur in either direction.

contour curve moves over (and in front of) an isolated configuration point which can be any isolated type of point in [8, Fig. 3] (this excludes only the “separating curves”). At the point where it meets the isolated point its tangent line is distinct from the tangent lines of any of the curves of the isolated point. For each type of isolated point, there are distinct transition types depending on the relative positions of the tangent lines and the side of the configuration curve on which the front object lies. This is illustrated in Figure 20(a).

- (3) Isolated configuration point moving across configuration curve: An isolated configuration point on an object which does not fill a neighborhood of the image of the point can be any of the stable types: (SC) semifold; (SF) edge curve of type (c)–(h) in [8, Fig. 10]; (SF) ridge crease of type (c) in [8, Fig. 12]; (FC) creases of type (a), (a’), (d), or (d’) in [8, Fig. 13]; (F) convex corner of type (e)–(i) in [8, Fig. 14]; (SF) notch corner of types (g), (h), (l) in [8, Fig. 16]; or (FC) or (F) marking curve of types (a), (c), (d), (e) in [8, Fig. 17]. Such a point moves past (and in front of) a smooth configuration curve of any type. Again there are different transitions depending on the relative positions of the tangent lines, as illustrated in (b) of Figure 20.

Cast shadow from a distance.

- (4) Configuration curve tangent to cast shadow curve: A configuration curve which is edge curve, ridge crease curve, or fold apparent contour for one object becomes tangent to the image of a smooth cast shadow curve on another object (or distant part of the same object). As the configuration curve moves past the tangent point, either the cast shadow breaks into two components or two components join together (see Figure 19).
- (5) Configuration curve moving across a V-point: A configuration curve which is an edge curve, ridge crease curve, or fold apparent contour for one object moves across

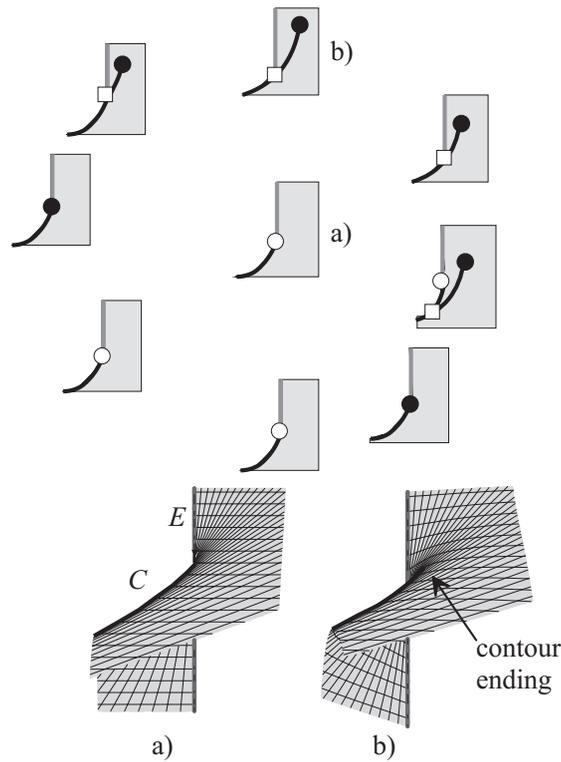


Figure 21. (a) A surface with a boundary edge E and apparent contour C viewed in the special semiswallowtail direction. Geometrically this means that the view direction is asymptotic and has higher (“A-point”) contact with the surface at the point where E and C meet. (b) The view is slightly moved to reveal a contour ending (cusp on the apparent contour). The clock diagram above illustrates the way in which the configuration of boundary edge and apparent contour changes as the viewpoint is moved around the swallowtail direction in the center. The views (a) and (b) are illustrated schematically at the center and one point of the clock.

the image of a V-point, formed as a cast shadow so the V-point becomes occluded or unoccluded.

6. A codimension 2 example. Codimension 2 phenomena are those which are visible only from special isolated directions: in a general fly-past an observer will not see these phenomena since a general path in the viewsphere will miss the isolated points. Such phenomena are therefore hard to realize, and for this reason we do not provide an extensive classification here. The usual way to describe them is to “circle round” one of the special directions and observe all the codimension 1 transitions which occur during the circuit, showing the result on a “clock diagram.”

Here is one brief example to illustrate the ideas. The “semi-swallowtail” is included in Table 3 of [8] for surfaces with boundary edges (among other configurations). Figure 21 shows two views of a surface with boundary edge, one of them with the view in the special semiswallowtail position and the other slightly moved. Other slight movements around the initial direction will reveal other configurations of apparent contour and boundary edge.

7. Explaining how singularity theory yields the classifications. We explain in this section how we apply the methods of singularity theory to obtain the classifications of both the stable views and the generic transitions occurring for configurations of geometric features, shade/shadow curves, and apparent contours.

We assume that for a fixed light source the shade/shadow curves form a stable configuration with any geometric features of M (without involving the viewpoint). First, we carry out a classification of the stable configurations that are possible. Then, to classify the interactions of apparent contours with these stable configurations, we introduce the equivalence relation for view projections, which will allow local changes of coordinates on the viewplane and on M which preserve the stable shade/shadow–geometric feature configuration.

Reduction to an abstract classification of local mappings. We let M denote a surface in 3-space \mathbb{R}^3 which at a point \mathbf{p} has one of the geometric features we have already introduced. The geometric features consist of sheets of the surface, their intersection along crease curves or corner points, or edge curves, and surface marking curves. These make up a configuration which we shall denote by X_g . See [8, Table 1]. The shade and shadow curves (if present) make up a configuration on M which we shall denote by C_g .

To model the directions for light and viewpoint, we consider two projection mappings near $\mathbf{p} \in M$ as in (7.1). Here, φ is the orthogonal projection of M in the view direction (“towards the viewer”), and ψ is the projection of M in the direction “towards the light.” By translation we may assume that \mathbf{p} projects to the origin 0 in both the light and view directions.

$$(7.1) \quad \begin{array}{ccc} \mathbb{R}^3, \mathbf{p} \supset M, \mathbf{p} & \xrightarrow{\psi} & \mathbb{R}^2, 0 \\ \varphi \downarrow & & \\ \mathbb{R}^2, 0 & & \end{array}$$

One way to understand the geometric features of (7.1) is to classify such diagrams of mappings allowing nonlinear change of coordinates for M and each of the \mathbb{R}^2 representing planes perpendicular to the light and view directions. This was the approach proposed by [14] and used by Donati [10] and Donati and Stolfi [11]. There is a fundamental problem with this approach if we hope to use all of the tools of singularity theory; namely, such a diagram is an example of a “divergent diagram of mappings,” and a basic theorem needed for singularity theory does not apply (see DuFour [12]). We take an alternate approach which extracts the essential features of the stable interactions of shade/shadow with geometric features. This yields either a subspace X_g or pair of subspaces (X_g, C_g) of M . Then, we introduce an equivalence among local mappings on M where we allow a local change of coordinates for the viewplane \mathbb{R}^2 and also for M , except we require that we preserve the configuration X_g or (X_g, C_g) , respectively. We call this \mathcal{S} -equivalence (which differs from the more restrictive notion of \mathcal{S} -equivalence used in [14] and [10]).

In order to make these concepts precise and to carry out the calculations, we need to parametrize the surface M locally, say by $\chi : \mathbb{R}^r, 0 \rightarrow \mathbb{R}^3, \mathbf{p}$. For a single smooth surface $r = 2$, and \mathbb{R}^2 is the “parameter plane” of the surface. For surfaces with creases or corners we use $r = 3$, and the local models explained in section 1 of [8] would provide a parametrization. For example, for a corner, the three sheets are parametrized by appropriate parts of the three

coordinate planes in \mathbb{R}^3 ; if the corner is convex, the region bounded by M is parametrized by the first octant where all coordinates are ≥ 0 . We then define a configuration C in \mathbb{R}^r corresponding to C_g via χ . In general for $r = 3$ we have a pair (X, C) , where X , which corresponds to X_g via χ , consists of the appropriate parts of the coordinate planes in \mathbb{R}^3 , whose images under χ make up the sheets of M , together with curves corresponding to any surface markings.

Because we are concerned only with local classifications of view projections near \mathbf{p} , we classify local mappings $f_0 : \mathbb{R}^r, 0 \rightarrow \mathbb{R}^2, 0$ under arbitrary local diffeomorphisms in the target and local diffeomorphisms in the source preserving C for $r = 2$ or the pair (X, C) for $r = 3$. This equivalence is called $C\mathcal{A}$ - or $X, C\mathcal{A}$ -equivalence. It is a specialized form of the \mathcal{A} -equivalence used for apparent contours. These equivalences can be viewed in terms of the groups of diffeomorphisms which give the equivalences, which are denoted by the same symbols.

If we now move our viewer direction, we can locally identify the new viewer plane with the original y - z plane, but the local mapping will have changed and, depending on the surface, the point \mathbf{p} need no longer go to 0 under the identification of planes. Because we can move viewer direction in two independent directions $\mathbf{u} = (u_1, u_2)$ orthogonal to the x axis, we obtain a family $F_1(\mathbf{x}, \mathbf{u}) : \mathbb{R}^{r+2}, 0 \rightarrow \mathbb{R}^2, 0$ such that when $\mathbf{u} = 0$, then we recover f_0 . Such a family is called a (2-parameter) unfolding of f_0 . Along with determining f_0 we also wish to determine the form of these unfoldings. We give further details of all of these reductions in [9].

Methods of singularity theory. To classify both the local mappings and their unfoldings we now will employ methods from singularity theory. Corresponding to the equivalences are analogous equivalences of the unfoldings. These equivalences of both local mappings and unfoldings are what singularity theory allows us to analyze. We think of these in terms of groups \mathcal{G} applied to a space of local mappings \mathcal{F} and their unfoldings. These groups and spaces can be thought of as infinite-dimensional manifolds, and there is a specific way to determine their tangent spaces. Provided that their tangent spaces have a certain special algebraic structure and the groups and unfoldings satisfy several natural conditions, then the basic theorems of singularity theory apply for the equivalence defined by them; see, e.g., [5] (or the more expository (but still very mathematical) [6]). Equivalence groups which satisfy these conditions are usually referred to as *geometric subgroups of \mathcal{A} or \mathcal{K}* . In our case, for stable configurations of geometric features with shade/shadow, the C or (X, C) forms a *special semianalytic set* or *special semianalytic pair*, as explained in [9]. This provides the conditions needed to ensure that both $C\mathcal{A}$ and $X, C\mathcal{A}$ are geometric subgroups of \mathcal{A} so that the basic theorems of singularity theory are valid.

What are these theorems and how do they allow us to carry out the classifications which we have given? From the tangent spaces to the groups we can compute the tangent spaces to the subspace formed from the local mappings which are equivalent to each other (these form “orbits of the group action”). Although both the space of local mappings and the equivalence class are each infinite-dimensional, it is possible to determine a dimension measuring by how much the full space is larger than the subspace formed by the equivalence class. This number is called the \mathcal{G} -codimension and is the same for any local mapping in the equivalence class. (Here \mathcal{G} denotes one of our equivalence groups.) The first major theorem of singularity theory,

the *finite determinacy theorem*, asserts that if f_0 has finite \mathcal{G} -codimension, then it is equivalent under \mathcal{G} -equivalence to a finite part of its Taylor expansion. The original finite determinacy theorem for the groups \mathcal{A} (and several others) was due to Mather [16], and a form which is applicable for geometric subgroups is given in [5].

This theorem is what allows us to give, as in [8, Tables 2 and 4], local models which are polynomial but represent an entire equivalence class of local mappings. Moreover, together with a further argument of Mather [17], it allows us to replace the problem of classifying local mappings within the infinite-dimensional space by that of classifying finite parts of Taylor expansions, reducing to problems involving finite-dimensional Lie groups. This approach allows the further introduction of other extended Lie group methods in [3] and [2], which allow symbolic computations to be carried out on a computer.

Second, if f_0 has finite \mathcal{G} -codimension, then it is possible to determine all possible ways that f_0 can be locally deformed. There is a special unfolding, the (uni)versal unfolding $F(\mathbf{x}, \mathbf{v})$ of f_0 , which has the property that all other unfoldings $F_1(\mathbf{x}, \mathbf{u})$ of f_0 are obtained from F by a mapping of the unfolding parameters $\mathbf{v} = \psi(\mathbf{u})$. A second major theorem of singularity theory, the *(uni)versal unfolding theorem*, gives a sufficient condition in terms of tangent spaces and the infinitesimal deformations of F in the parameter directions to ensure that F is a versal unfolding. It gives a specific method for constructing versal unfoldings of a finite codimension local mapping f_0 , and shows that the codimension specifies how many unfolding parameters are needed. Again the versal unfolding depends on the equivalence group \mathcal{G} . This result also originated with Mather, who implicitly used it in [17] and explicitly stated it for one equivalence group in [18]; it was given for \mathcal{A} and another equivalence group by Martinet [15], and again is generally valid for geometric subgroups as shown in [5].

The above ideas are applied in our situation for the germs of codimension ≤ 2 , which we classify for the various configurations X and (X, C) . (However, we have only touched on codimension 2 here, in section 6.) Once we verify that the local mappings occur as the result of projection from a surface with geometric features, we may verify the criterion and apply the versal unfolding theorem to conclude that movement by viewer direction gives a versal unfolding of the local mapping. Because versal unfoldings for equivalent local mappings are themselves equivalent as unfoldings, we have completely determined the local transition behavior as we change view direction. Furthermore, the unfolding theorem implies that local mappings of codimension 0 are stable under viewer movement. This also allows us to identify the stable configurations.

The same theory applies to the multilocal case, providing a classification of multilocal mappings, yielding the stable multilocal mappings and the versal unfoldings for multilocal mappings. These provide the corresponding stable multilocal configurations and the generic transitions

There is one point which distinguishes certain parts of the classifications we obtain from the earlier classification of apparent contours given in [8, Table 2]. In this table, there is a finite list of equivalence classes represented by the polynomials. By contrast, for certain stable configurations of geometric features and shade/shadow curves, the list of local mappings of codimension ≤ 2 is infinite. There are families of equivalence classes described by parameters, called *moduli*, which have the property that continuously changing the parameter continuously changes the equivalence classes of local mappings. To overcome this problem and reduce to a

finite classification, we have to replace the equivalence by a corresponding *topological equivalence*, where diffeomorphisms are replaced by (piecewise differentiable) homeomorphisms. Now the equivalence captures the qualitative properties of the configuration, which is what our visual recognition really captures as well. The *topological classification* allows us to reduce to a finite number of representative parameter values in order to obtain an inclusive classification. This also reduces the codimension by the number of moduli. For example, the topological classification already appears for “lips/beaks on the boundary” and the “double cusp” in [8, Table 3] and the “semi-swallowtail” in [8, Table 4]. In this paper it appears in many cases, where the model mappings have higher codimension, but all but one of the parameters in the versal unfolding do not alter the topological behavior. Hence, there is only one interesting transition for the model, and this is the generic transition.

There are analogues given in [7] for topological equivalence of the main theorems of singularity theory for the finite determinacy, classification, and versality theorems. These allow us to carry out the analogous steps of the classification in these cases. This additional feature illustrates that there are many important but subtle points in the application of singularity theory that we have had to gloss over in this brief explanation. A complete and thorough treatment is carried out in [9].

Carrying out the classification. To carry out the classification, we must first classify the stable configurations of geometric features with shade/shadow curves. We do this by applying the classification of Tari [19, 20], taking into account visibility and allowing multiple types of corners. Then for each stable configuration we obtain, we must carry out the classification of the abstract local mappings by the equivalence which preserves the individual stable configurations. This is potentially an incredibly lengthy process. Fortunately, it is considerably simplified because, as we have mentioned, there already exist several classifications for abstract local mappings preserving a marking curve, by Bruce and Giblin [1], and creases and (convex) corners, by Tari [19, 20]. Because of results in [9], these classifications also apply to other configurations and, in addition, help in the classifications for still more complicated configurations.

We may then apply the abstract classifications for specific stable configurations as given in [8, Table 3].

Abstract mappings versus realization. In order to apply the abstract classifications to the situation of illuminated surfaces in 3-space we need to take into account the special geometry of our situation. Consider the (SC) case of a single smooth surface M without surface marking, and suppose that both the light and view directions lie in the tangent plane at the point $\mathbf{p} \in M$. Then there is a shade curve passing through \mathbf{p} , and also \mathbf{p} lies on the contour generator, that is, the curve which projects in the view direction to the apparent contour. Shade curves and contour generators are not arbitrary curves on M : they arise as critical sets of projection maps. This has a significant consequence: setting aside the nongeneric case where the view and light directions coincide, these curves can be *tangent* only when \mathbf{p} is parabolic and their tangents are both in the unique asymptotic direction at \mathbf{p} . This immediately gives a restriction on any singularity from the abstract list which requires that the critical sets on the two projections be tangent: such a singularity must occur at a parabolic point. In fact there are two cases in which the geometrical restrictions prevent abstract singularities from

being realized at all. (These appear as “N” in the right-hand column of Table 3 of [8].) It is also possible to use geometrical arguments to show that singularities are not versally unfolded by moving the view direction (such a case appears as “NV” in the same table).

Thus there is an important step after obtaining the abstract classifications of maps which preserve the geometric features and shade/shadow curves: we need to consult the geometry of the situation to discover exactly which abstract singularities are realized and versally unfolded. Again, details are in [9].

Having obtained a realization, there remains the consideration of *visibility*: for example, the view direction can be reversed, changing the occlusions of one sheet of a crease or corner by another. In the case of corners, visibility considerations are particularly complex, as discussed above in sections 3 and 4.

8. Comments and summary. We have completed in this paper the explanation of how the complex interactions of geometric features, light, and viewer movement can be analyzed using the methods of singularity theory to yield a classification of both expected local features of images and their generic transitions under viewer movement. Together with the results of part I [8], these provide a concise alphabet of local curve configurations that we expect to see in images, along with the possible geometric properties that accompany them. Further, we provide a specific classification of the generic transitions which occur in these configurations under viewer movement. These results provide a catalogue which subsumes and considerably refines the earlier work of a number of groups on special aspects of images.

Acknowledgments. The first two authors wish to express their sincere gratitude (i) to the IMA in Minneapolis, and its Director Doug Arnold, and (ii) to Stanisław Janeczko and the Center for Advanced Studies of the Warsaw University of Technology. The generous hospitality afforded the authors during their visits enabled them to complete some parts of the work presented here.

REFERENCES

- [1] J. W. BRUCE AND P. J. GIBLIN, *Projections of surfaces with boundary*, Proc. London Math Soc. (3), 60 (1990), pp. 392–416.
- [2] J. W. BRUCE, N. P. KIRK, AND A. A. DU PLESSIS, *Complete transversals and the classification of singularities*, Nonlinearity, 10 (1997), pp. 253–275.
- [3] J. W. BRUCE, A. A. DU PLESSIS, AND C. T. C. WALL, *Determinacy and unipotency*, Invent. Math., 88 (1987), pp. 521–554.
- [4] *Visibility diagrams for corner transitions*, pdf file available as supplementary material for this article or, with updates, at <http://www.liv.ac.uk/~pjgiblin> and <http://www.math.unc.edu/Faculty/jndamon>.
- [5] J. DAMON, *The unfolding and determinacy theorems for subgroups of \mathcal{A} and \mathcal{K}* , Mem. Amer Math. Soc., 50 (1984), no. 306.
- [6] J. DAMON, *The unfolding and determinacy theorems for subgroups of \mathcal{A} and \mathcal{K}* , Proc. Sympos. Pure Math., 40 (1983), pp. 233–254.
- [7] J. DAMON, *Topological triviality and versality for subgroups of \mathcal{A} and \mathcal{K} II: Sufficient conditions and applications*, Nonlinearity, 5 (1992), pp. 373–412.
- [8] J. DAMON, P. GIBLIN, AND G. HASLINGER, *Local image features resulting from 3-dimensional geometric features, illumination, and movement: I*, Internat. J. Comput. Vision, 82 (2009), pp. 25–47.
- [9] J. DAMON, P. GIBLIN, AND G. HASLINGER, *Characterizing Stable Local Features of Illuminated Surfaces and Their Generic Transitions from Viewer Movement*, in preparation.

- [10] L. DONATI, *Singularités des vues des surfaces éclairées*, Ph.D. thesis, Université de Nice, Sophia Antipolis, France, 1995.
- [11] L. DONATI AND N. STOLFI, *Shade singularities*, Math. Ann., 308 (1997), pp. 649–672.
- [12] J. P. DUFOUR, *Sur la stabilité diagrammes d'applications différentiables*, Ann. Sci. École Norm. Sup. (4), 10 (1977), pp. 153–174.
- [13] D. HAVIV, *Representation and Processing of Images based on Geometric Models*, Ph.D. thesis, Weizmann Institute, Rehovot, Israel, 2010.
- [14] J.-P. HENRY AND M. MERLE, *Shade, shadow and shape*, in Computational Algebraic Geometry (Nice, 1992), Progr. Math. 109, Birkhäuser Boston, Boston, MA, 1993, pp. 105–128.
- [15] J. MARTINET, *Déploiements versels des applications différentiables et classification des applications stables*, in Singularités d'Applications Différentiables (Sém. Plans-Sur-Bex, 1975), Lecture Notes in Math. 535, Springer, Berlin, 1976, pp. 1–44.
- [16] J. N. MATHER, *Stability of C^∞ mappings III: Finitely determined map-germs*, Publ. Math. IHES, 35 (1969), pp. 127–156.
- [17] J. N. MATHER, *Stability of C^∞ mappings IV: Classification of stable germs by \mathbf{R} -algebras*, Publ. Math. IHES, 37 (1970), pp. 223–248.
- [18] J. N. MATHER, *Notes on Right Equivalence*, manuscript, 1970.
- [19] F. TARI, *Some Applications of Singularity Theory to the Geometry of Curves and Surfaces*, Ph.D. thesis, University of Liverpool, Liverpool, UK, 1990.
- [20] F. TARI, *Projections of piecewise-smooth surfaces*, J. London Math. Soc. (2), 44 (1991), pp. 152–172.