



Amodal volume completion: 3D visual completion

Toby P. Breckon *, Robert B. Fisher

School of Informatics, University of Edinburgh, Edinburgh, Scotland, UK

Received 4 December 2004; accepted 17 May 2005

Available online 27 July 2005

Abstract

This work considers the common problem of completing partially visible artifacts within a 3D scene. Human vision abilities to complete such artifacts are well studied within the realms of perceptual psychology. However, the psychological explanations for completion have received only limited application in the domain of 3D computer vision. Here, we examine prior work in this area of computer vision with reference to psychological accounts of completion and identify remaining challenges for future work.

© 2005 Elsevier Inc. All rights reserved.

Keywords: Volume completion; 3D Vision; Perceptual psychology; Range data; 3D Surface completion

1. Introduction

Common 3D acquisition techniques in computer vision are realistically only $2\frac{1}{2}$ D in nature—such that the backs and occluded portions of objects cannot be realised from a single view.¹ However, the same physical limitation is also true of human stereopsis where visual completion, as an aspect of our extending visual reasoning system, allows

* Corresponding author.

E-mail address: toby.breckon@ed.ac.uk (T.P. Breckon).

¹ Here, we are considering common surface based capture techniques such as stereo vision and laser range capture rather than the volumetric approaches, such as MRI, more common to medical vision applications.



Fig. 1. Candidates for volume completion.

us to perceive the completion of occluded objects and reason how an object may appear from an alternative viewpoint (e.g., an opposing backward view).

Common 3D capture techniques result in models which are inherently $2\frac{1}{2}$ D in nature (Fig. 1) or 3D models which are constructed of multiple differently orientated views. In the latter case, the additional capture, view planning, and subsequent combination, via registration, are both time consuming and computationally expensive [18]. Additionally ensuring all faces and crevices of a target have been captured at least once is not an easy task and is often unapparent until final registration [11]. Ideally a method of visual completion is required to facilitate the completion of a 3D model, from a $2\frac{1}{2}$ D or incomplete 3D representation, akin to that present in human visual reasoning—a method of completing unknown 3D scene enclosures (i.e., volumes).

In incomplete scenes (Fig. 1) it is often perceivable to a human viewer how the missing portions should appear (i.e., what we *expect* to see) based on the ability of our visual system to compose these areas using supposed surface, contour, and pattern completion within the scene. The salient point, however, is that we do not perceive exactly what exists there but what we could reasonably expect to see. In this aspect, working with our limited knowledge of the scene, we trade off accuracy for plausibility—*what actually exists against what we expect exists*. Our expectation for these invisible scene portions can be thought of as being governed by two classes of knowledge—*Scene Evidence*, visual cues specific to the scene that trigger certain completion expectations, and *World Knowledge*, general a priori completion rules based on prior visual experience.

Based on this analysis of visual completion an underlying principle of our completion abilities can be stated as follows: as humans, we perpetually complete the portions of a scene we cannot physically visualise based on those portions which we can. This reasoning process is itself governed by our generic and specific knowledge of the world which we are perceiving in conjunction with relevant scene evidence.

This concept, of completion through controlled scene continuation, can be thought of as *visual propagation* and embodies an approach to *completion through generalisation* rather than the more specific *completion through recognition* [3,30,22]².

² That is, recognise an article from a partial view and hence determine its unseen portions from a specific a priori model. This is the common “recognition problem” of computer vision.

In this paper, we review prior psychological theory and computer vision to identify the supporting aspects for this view of completion with specific reference to the completion of unknown 3D volumes—considered here as the completion of the exterior bounding surface(s) for an unknown (i.e., unbounded) enclosure within a given 3D scene. Notably, we see that 3D completion through visual propagation is specifically supported by the recent psychological account of Volume Completion [72], itself surpassing earlier work in this field, although computer vision work in considering such aspects remains limited.

2. Completion in perceptual psychology

The phenomenon of human visual completion has been of interest to psychologists for a considerable period, ranging from the early Gestalt principles [41] to more modern interpretations [37,72].

Much of the early work centres around the study of various, now famous illusory examples (Fig. 2) and the cues and process that cause certain perceptions for these cases [34–36]. This work concentrates on the visual perception of 2D structures—to date an on-going aspect of psychological research (e.g. [88,23,54,87]).

At the top level work in this area concerns itself with the two primary phenomenon of visual completion:

- *Modal completion*: the illusory perception of a foreground object/shape, even though it is not explicitly present, due to the perceived occlusion of background objects/shapes (e.g., the perception of a triangle or rectangle in Fig. 2).
- *Amodal completion*: the perception of a background object/shape continuing behind a foreground occluder despite the fact that one cannot see it directly. The occluded portion is invisible yet perceived by the viewer (e.g., the perception of the small disks being occluded and thus continuing behind the illusory triangle/rectangle in Fig. 2).

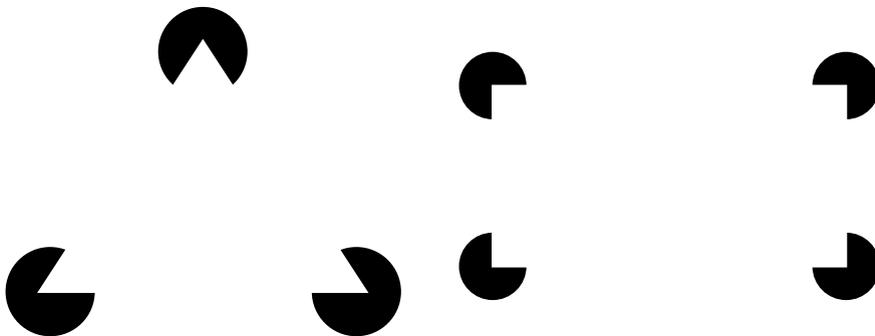


Fig. 2. Famous illusory figures.

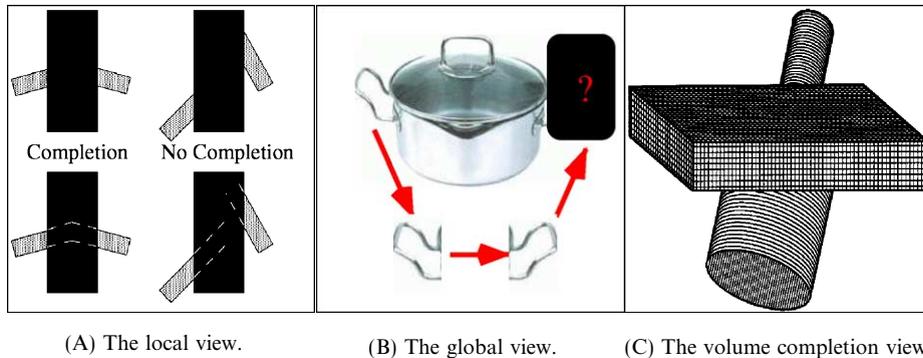


Fig. 3. Amodal completion examples: a clash of theories [72].

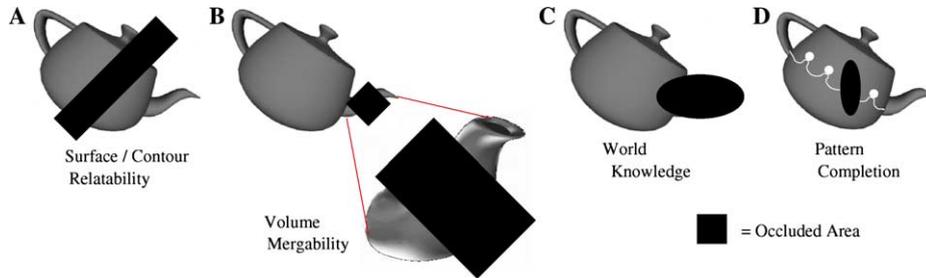
In reality, amodal completion can be thought of as occlusion among opaque surfaces (either partial or complete—i.e., hidden) whilst modal completion can be thought of as perfect camouflage [72]. Clearly amodal completion is the more commonly occurring instance with relation to the problems encountered in 3D computer vision and hence we are interested in investigating further.³

We are specifically interested in the aspects of amodal completion as they pertain to 3D completion. Theory in this area is still a source of debate but, following the review by Tse in [71], there are generally three types of theory in this area:

1. Completion occurs when certain conditions are satisfied among local scene cues (e.g., contours and junctions)—the *local* view (Fig. 3A).
2. Completion occurs when certain conditions are satisfied among global scene cues (e.g., symmetry, regularity, and simplicity of form or pattern)—the *global* view (Fig. 3B).
3. Completion occurs when certain conditions are satisfied among representations internal to the scene (e.g., contour, surfaces, and volumes)—the *volume based* view (Fig. 3C).

Theories following the local view (type 1) generally support completion through ‘good continuation’ (Kellman and Shipley [37]) and intersection of scene contours (and associated surfaces) without consideration for the wider structure within the scene (e.g. [37,83,62], Fig. 3A). In contrast, the global view (type 2) challenges this assumption of locally dependent completion to suggest the importance of global regularities, for instance symmetry, in completion patterns (e.g. [76,77]). For example, it is clear that Fig. 3B could not be completed based on a purely localised view of completion.

³ Although, admittedly, modal completion itself poses a similarly interesting future problem for computer vision research [88,87].



In each case the teapot is completed by different aspects of visual completion.

Fig. 4. Examples of visual completion aspects.

The more recent type 3 theories have evolved from the traditional views in this area, namely that completion is dependent on surface completion on a common depth plane (e.g. [53]), towards the concept of a volume completion based approach [72,75]. Fig. 3C shows an example where amodal completion would not occur through a simple local or global view as shown in Figs. 3A and B—yet the viewer amodally completes the occluded “wine bottle” volume to form a cohesive object. Here (i.e., Fig. 3C), completion occurs at a higher level based on perceived volumetric relationships within the scene—*volume based completion*.

2.1. Volume completion concept

Volume completion, suggested by Tse [72], argues that contour and surface based amodal completion theories are not alone sufficient to account for all completion phenomenon (see Figs. 3C and 6). In an extension to previous consideration of visual completion (types 1 and 2), volume completion considers the way the human visual system is able to complete entire 3D enclosures, which are not explicitly visible within a scene.

It is argued that volume completion can be achieved through a combined approach utilising:

1. *contour/surface relatability*—the inter-relationships in the completion surfaces and contours within the scene (Kellman and Shipley [37]). (see Fig. 4A). For formal definition see [72].
2. *volume mergability*—the relationship between unbounded volumes within the scene. “Mergability” refers to how the inside of one volume, unbounded due to occlusion, can join with a similarly unbounded volume to create a larger volume within the scene [72] (see Fig. 4B).
3. *world knowledge*—generic and specific knowledge relating to the world and to instances within the scene that is being perceived [70] (see Fig. 4C).
4. *pattern completion*—the presence of regularity or semi-regularity within the scene that could be extrapolated as being present in unknown scene areas [71] (see Fig. 4D).

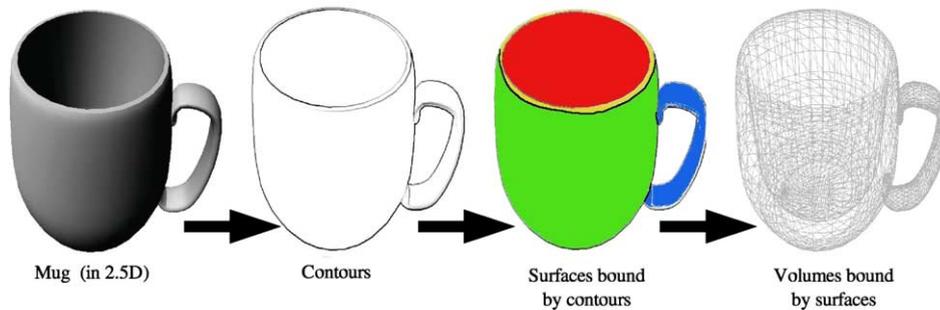


Fig. 5. Essence of volume completion.

2.2. Volume vs. surface completion

Q: So how does volume completion differ from earlier surface or contour based completion?

When performing any visual completion in 3D we are essentially performing the completion of a volume. Consider a $2\frac{1}{2}$ D representation of a mug (i.e., front faces only). First, assume we can identify the contours of the mug within our representation. From this, we can identify the surfaces of the mug as the regions bounded by these contours. In turn, the 3D volume of the mug is itself bound by the concurrent 3D completion of these surfaces. This is a realistic assertion as we do not expect the mug to extend beyond the bounds of the visible surfaces we perceive. Hence realistically we are completing a 3D structure enclosed by a set of visible and invisible surfaces—we are essentially completing a volume through the combined completion of bounding surfaces (Fig. 5).

Alternatively volume completion can be thought of as concurrently completing a set of surfaces to form a complete 3D representation over the 3D space the surfaces potentially enclose (i.e., the volume). In the case where we only have a single surface (e.g., a sphere) then we are completing the volume bounded solely by that singular surface. When dealing with the degenerate case of completing a single planar surface, a volume-less surface, we are essentially completing an undefined volume as the associated volume is essentially unbounded. This occurrence is unrealistic as every real world surface has an associated volume [72]—even the surface of a sheet of paper has an associated volume, albeit very small.

Formalising this view, the $2\frac{1}{2}$ D scene completion problem can be stated as follows:

- Let there be a set of natural contours C visible in the scene, $C = \{\text{fold/blade edges of scene}\}$.
- Let the set of visible surfaces S be defined such that:

$$\forall S_i \in S, S_i \text{ is bounded by } \{C'\} \subseteq C.$$

- There exists a completion of the scene, V , such that:

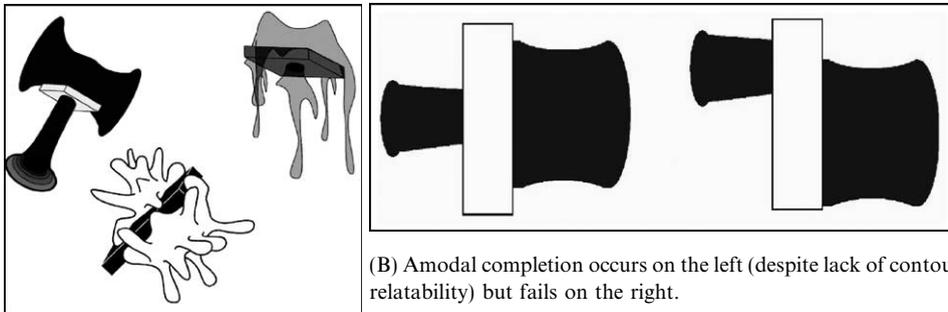
$$V = \sum_0^i \{S_i^C\}, \text{ where } S_i^C \text{ is the completed surface of } S_i.$$

The solution to the volume completion problem is to find the V that minimises the completion error with respect to constraints on completion, namely the need for the aspects of good surface/contour relatability, volume mergability, pattern completion, and plausibility (i.e., world knowledge constraints) as discussed previously.

Based on this premise that *volumes are bounded by surfaces, which are themselves bounded by contours*, the volume completion view subsumes surface completion as a special case by asserting that in fact every scene surface has an associated volume [72]. Through this paradigm volume completion operates by merging volumes, those enclosures bound by surfaces and contours in the scene, to achieve visual completion. This completion process is in turn governed by the constraints of world knowledge to assess plausibility and pattern completion to enforce regularity. Thus, we see our *visual propagation* principle underpinning the theory in this area.

Additionally, it has to be noted that the volume completion view is not in itself separate from the local and global views of completion (types 1 and 2). Volume completion supports the view that whilst contour relationships (type 1) and global regularities (type 2) are not solely responsible for completion they, together with other visual cues, can be used to infer edge, surface, and volume relationships. Completion then occurs at a higher level through the analysis of these relationships [71]. In fact, Tse [72] goes on to suggest that the process of volume completion may in fact “*involve a relaxation into a best-fit given these local constraints.*”

Examples of the concept of volume completion as a combined theory of completion can be seen by considering Figs. 3, 4, and 6 together. Each of the completion theories of Fig. 3 (local, global, and volume) can be attributed to the completion



(A) Amodal completion of visually disjoint volumes.

(B) Amodal completion occurs on the left (despite lack of contour relatability) but fails on the right.

Fig. 6. Amodal “volume based” completion examples where other theories fail [72].

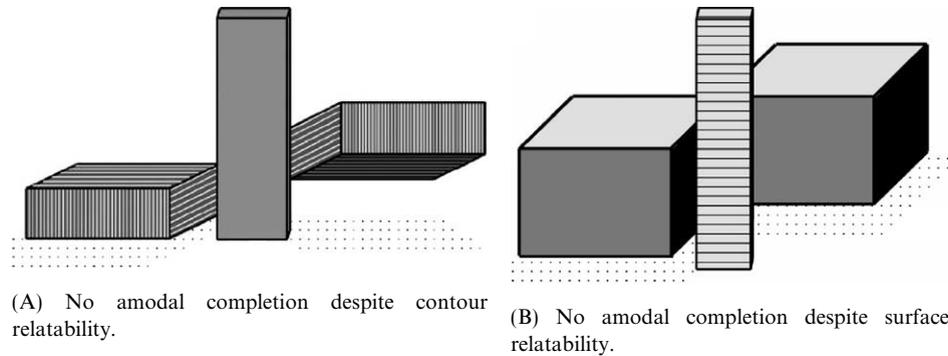


Fig. 7. Contour reliability—a counter example [72].

cases of Fig. 4. For instance, the localised completion of Fig. 3A can be attributed to the contour/surface reliability of Fig. 4A and the global completion of Fig. 3B to the world knowledge and pattern completion depicted through Figs. 4C and D. The more complex completion case of the “wine bottle” object in Fig. 3C shows a good example of volume based completion—here completion succeeds despite the lack of directly relatable contours or surfaces and in the absence of a global pattern (e.g., symmetry, regularity) to allow us to deduce the occluded portion. In this case, we see that the completion occurs primarily through volume mergability, as depicted in Fig. 4B, although the influence of world knowledge has also to be acknowledged.⁴ These examples support the case made by Tse [72] that amodal completion occurs through a combination of factors that all influence 3D completion and not simply through contour and surface reliability as suggested earlier by Kellman and Shipley [37]. The examples in Figs. 3B and C could not be completed based on the contour/surface “good continuation” theory of [37] whilst in Fig. 3A, the “good continuation” of contours and surfaces acts as a rough guide through which completion is achieved. It has to be noted, however, that volume mergability and/or world knowledge could also provide completion in this case using the contour reliability as completion constraints. Further, examples that rely on volume mergability for amodal completion, from [72], are given in Fig. 6 and develop an argument [69] that amodal volume completion occurs in the absence of specific contour based cues.

Support for a volume based completion account is further provided through Fig. 7, which shows that amodal completion in the human visual system can still fail despite the presence of surface and contour reliability. In this example, the viewer fails to complete the two cubic objects, despite the presence of reliability, because of the overriding perception of disjoint 3D enclosures within the scene (albeit through any combination of the volume completion aspects previously mentioned). Here, we see volume based completion overriding the ideas of completion by contour and surface reliability supporting Tse’s view [72] that amodal completion cannot be simply

⁴ That is, from the abundance of bottle shaped objects we may see in a lifetime!

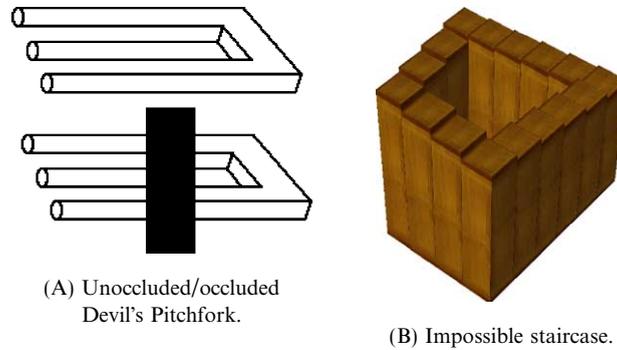


Fig. 8. Famous impossible works.

attributed to the “good continuation” of contours and surfaces of Kellman and Shipley [37].

This idea of volume completion overriding contour/surface relatable completion is also evident in the famous examples of Fig. 8 and commonly in the artistic works of E.C. Escher [46] (e.g., “Relativity”: E.C. Escher). Here, we perceive “impossible” objects despite the relatability of their contours and surfaces. It is our realisation of the object in 3D, considering the set of contours and surfaces together as a 3D volumetric shape, that allows us to identify the spatial contradictions in the scene. This issue was explored by Huffman [29] by likening scene interpretation to understanding sentences and hence impossible objects to nonsensical sentences.

2.3. Further work in volume completion

The volume completion theory is backed up by experimental evidence in Tse’s later work [71], which also further supported the idea of considering a global view of completion rather than completion based on an isolated ‘per surface’ approach. This work also gives a fuller discussion of the concept of ‘mergability’ underlying volume completion in terms of the concept “good volume continuation” similar to that of Kellman and Shipley [37] in relation to surfaces. Overall, it concludes that completion processes are “*inherently probabilistic in nature*” with multiple influences including proximity (‘localness’), global context, pattern, similarity of material, and form, orientation, and volume mergability. With relation to computer vision this can be likened to a co-operative algorithm such as Marr’s seminal work on stereo matching [50].

The issue of global completion, raised in [72], has been considered by recent work looking both experimentally and theoretically at global completion in non-regular shapes (Tessa and Van Lier et al. [66] and Van Lier [75]). Both conclude that evidence for globally influenced completion methods is strong with [75] going on to discuss further the concept of ‘fuzzy completion’—the concept of being able to derive a plausible completion to an irregular object based on analysis of the

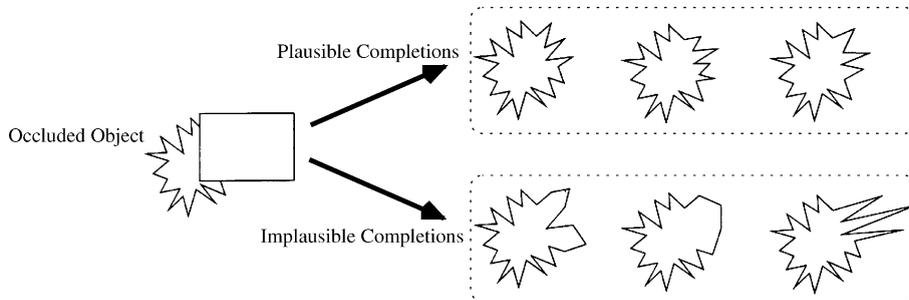


Fig. 9. Fuzzy completion: an experimental example [75].

visible portion. It is suggested that this allows us to plausibly complete irregular shapes by analysis of the general fuzzy regularities we perceive in the visible portion (see Fig. 9). This suggestion is backed up by experimental evidence through which the author concludes global completion offers the best account for the fuzzy and 3D completion phenomenon observed. This is further supported by [78] that more generally asserts, via experimentation, that well-established local and global completion tendencies for 2D surfaces appear valid for extension into 3D. Related work in extending 2D *modal* completion to 3D similarly supports the concept of volumetric completion from a combination of local and global occlusion cues [88,87].

Additionally orientation in 3D completion has been considered by [1] which showed experimentally that orientation and surface adjacency can affect the perception of a 3D volume. This shows that completion is affected by scene context and further supports the view of global constraints within 3D completion.

Contour propagation (and by extension surface propagation) is also considered by Tse [73], with regard to the visual interpretation of 2D silhouettes as 3D volume enclosing surfaces. This work strongly supports the concept of visual propagation through the propagation of single and multiple contours for the reconstruction of 3D volumes. It considers the identification of “propagatable segments” from occluding contours to propagate information about the underlying 3D form from the occluding contours to areas where 3D form is ambiguous.

Contour propagation [73] can be considered a specific subset of our earlier *visual propagation* proposal. Generalised visual propagation encompasses completion through the propagation of all visual stimuli—contours, surfaces, shape, texture, colour, reflectance, etc. Here, we see a specific case of visual propagation based on existing psychological evidence for human completion abilities.

The work [73] proposes a theoretical algorithm for generating 3D shape from 2D object silhouettes: given an identified propagatable segment (e.g., occluding edge), propagate it into the interior of the silhouette (in a given direction of propagation) and then scale (expand/contact) it until it touches the bounds of the silhouette at this point. Repeated propagation at different distances into the interior thus builds a $2\frac{1}{2}$ D contour map of the original silhouette.



Fig. 10. Tse's "seamonster" amodal volume completion case [72].

Although several practical issues remain in addressing the underpinning assumptions a future extension into computer vision work is suggested [73] and future extension to $2\frac{1}{2}$ D to 3D propagation, to facilitate true volume completion appears viable.

2.4. Volume completion summary

Ultimately the traditional theories of completion (types 1 and 2) can be challenged by considering Tse's "seamonster" example from [72] (see Fig. 10). Here, despite the lack of local cues to suggest occlusion, the lack of "good continuation" in terms of contours and surfaces and even the lack of any explicit occluding surface the viewer sees a coherent, complete "seamonster."

Tse suggests that the unbounded black portions are processed, as potential volumes, through volume completion to derive the completed "seamonster" concept that we perceive. However, the bearing of "world knowledge" within this volume completion instance, from the countless similar {sea|worm|snake|Loch Ness} monster images to which we are exposed, has also to be considered.

Overall, it is clear that there is no simplistic explanation for our visual completion abilities. A number of suggested theories fail to explain all of the completion capabilities of the human visual system. The most recent and advanced such theory, Volume Completion [72], offers a combined and somewhat open-ended approach that considers various visual influences in a "best-fit" optimisation based approach that subsumes the abilities of earlier theories. It would seem apparent that volume completion offers both a theoretical explanation of human visual completion abilities and a useful framework for implementing artificial visual completion with heavy support for our desired principle of *completion through visual propagation*. Notably, some recent theoretical work [73] supports the realisation of volume completion through a subset of the *visual propagation* paradigm.

3. Completion in computer vision

Work in visual completion in computer vision can be thought of in two distinct but related camps. First, there are those inspired by the 2D contour and surface completion theories of [41] interested in completion of 2D synthetic or photographic images. The majority of this work concerns itself with contour or pattern identification realised through a variety of techniques (e.g. [5,24,82,23,44,45,54]) to explore both the possibilities for automated 2D completion and to some extent to test psychological visual completion theories. Here, however, we are primarily concerned with the other camp in visual completion—3D visual completion—and thus our analysis of completion work in computer vision concentrates itself here.

Work in 3D visual completion is reasonably limited within the computer vision literature. However, a number of existing techniques support the concept of visual propagation and to a limited extent some of the previous psychologically based completion theories.

3.1. Occlusion resolution

In terms of tackling the limitations of $2\frac{1}{2}D$, there has been some initial work in the resolution of partial occlusions in range images by Fisher et al. [20,59,7,12]. Overall three types of occlusion are identified and considered: *occlusion preserving surfaces*, *occlusion breaking surfaces*, and *occlusion breaking boundaries* (see Fig. 11). This work completed occluded surfaces through the use of contour and surface relatability together with a limited aspect of “world knowledge” to constrain the overall completion (i.e., definition of ground and wall planes at scene limits). The technique operates by extracting contours and surfaces using segmentation and fitting techniques, extending the contours into the occluded area and then interpolating the surface between these now extended contours.

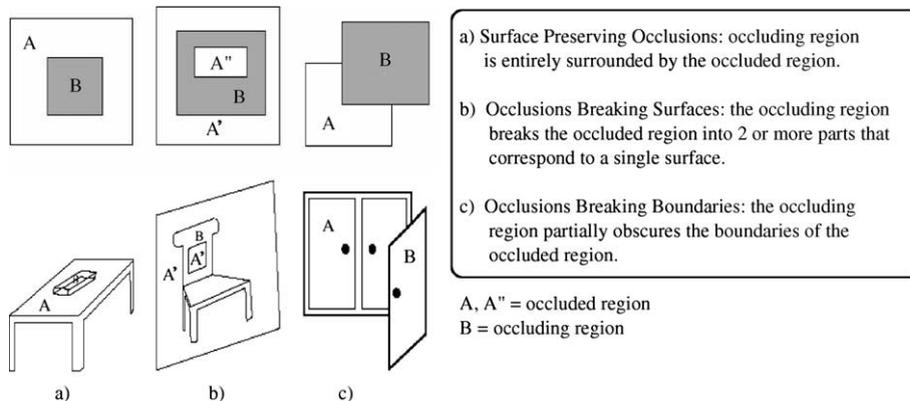


Fig. 11. Types of occlusion identified and resolved in [7].

In terms of visual completion theory this approach fits best with the completion paradigm put forward by Kellman and Shiplet [37], itself forming a subset of the later theories on volume completion [72]. The more advanced of this occlusion resolution work [7], considering surface completion through the use of contour relatability, falls short of considering true surface relatability in terms of multiple surface interactions and concentrates on the completion, through continuation, of uniform surface types. The limitations of [7] in completing certain surfaces (due to lack of constraining information) and in fully completing scenes (i.e., from $2\frac{1}{2}$ D to complete 3D) is acknowledged by the authors. Additionally this work considers the completion of individual surfaces in isolation rather than completing a set of multiple (possibly interacting) scene surfaces in combination. Possible constraints introduced by surface interaction are thus not available to aid completion.

In contrast to this data driven approach, related work has also been carried out using parametric shape fitting to facilitate the 3D completion of partially visible $2\frac{1}{2}$ D cylinders and cuboids. This model based approach is utilised in [51,19] with specific reference to building column completion in architectural reconstruction. The limitations in this work are that it considers the columns abstractly from the scene for the purposes of completion and completion is itself based purely on parametric completion of regular 3D shapes—a good example of generic “world knowledge” in terms of the volume completion paradigm.

3.2. $2\frac{1}{2}$ D to 3D reconstruction

Completion work in architectural reconstruction has also been carried out using structure and motion as a source for the initial $2\frac{1}{2}$ D representation [15,14]. This work uses a model fitting based approach together with a high level of “world knowledge,” specific both to architecture in general and to the classical/Gothic architectural styles considered. Both synthetic buildings and building completions are constructed through the utilisation of Bayesian priors for architectural scenes and a model fit of the structure from the motion image sequence.

The probability priors for model fitting to the video sequence fall into four categories: primitive usage (e.g., window, door frequency), shape (e.g., specific architectural styles such as narrow Gothic windows), texture parameters (e.g., windows are often dark with intersecting vertical/horizontal bars), and image likelihood (i.e., the likelihood of the images given a complete specification of the model) [15]. The primitive and shape priors are based upon both architectural rules, observations and practical considerations whilst the texture priors are learnt from a set of example images. Model fitting is achieved by obtaining a *maximum a posteriori* (MAP) set of model parameters. In the latter work [14], this is augmented with spatial organisation prior to allow the generation of building completions and synthetic buildings using a Markov Chain Monte Carlo (MCMC) algorithm with the initial building model as a seed. An example of this approach applied to the Downing Library, Cambridge, is shown in Fig. 12.

The overall result is a textured geometric model based representation of a given building at the level of defined geometric and architectural primitives (e.g., walls, doors, windows, columns, etc.). In terms of perceptual completion, this work relies

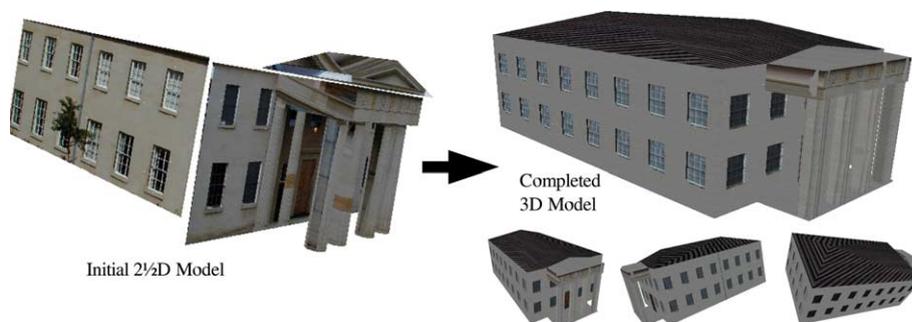


Fig. 12. $2\frac{1}{2}$ D to 3D reconstruction using MCMC.

heavily on the concept of “world knowledge” for success, both general (i.e., geometric models) and specific (i.e., architectural), but does follow the arguments put forward for visual propagation and plausible completion—in this case both embodied in a probability distribution.

Overall this work offers a successful model based, heavily knowledge-dependant approach to completion specific to the domain of architecture with structure from motion as its primary source of 3D information. It does, however, also support the idea of probability based plausible completion where a number of different completions are possible (cf. fuzzy completion [75]).

3.3. $2\frac{1}{2}$ D completion

In a variant on other work in this field [68] considers the completion of $2\frac{1}{2}$ D depth maps from sparse range data based on corresponding intensity image similarity. This relatively successful technique relies on having valid intensity data for the scene regions where depth is not explicitly known and thus uses a method akin to the texture synthesis of Efros [16] to infer depth based on the intensity similarity of localised regions to those where both intensity and depth are known. Markov random fields are used to capture the relationship between intensity and depth over localised regions neighbouring an unknown depth value and the unknown depth value then is inferred from this model in a deterministic fashion. Depth values are synthesised one value at a time (following the texture based method of Efros [16]) and are constrained, by neighbourhood localisation, to be similar to some region close to their location.

An example of this technique is shown in Fig. 13 where from left to right we see complete intensity knowledge of the scene, partial depth knowledge, and the resulting completion of the depth map compared against the ground truth. Apart from a few minor errors around the top of the main foreground object we see a successfully completed depth map based on intensity similarity.

In relation to volume completion, this work follows the concept of visual propagation but relies on at least partial knowledge of the entire scene (i.e., intensity) from which only $2\frac{1}{2}$ D rather than 3D scene knowledge is derived. By contrast the problem posed by volume completion contains no explicit knowledge of the areas of the scene

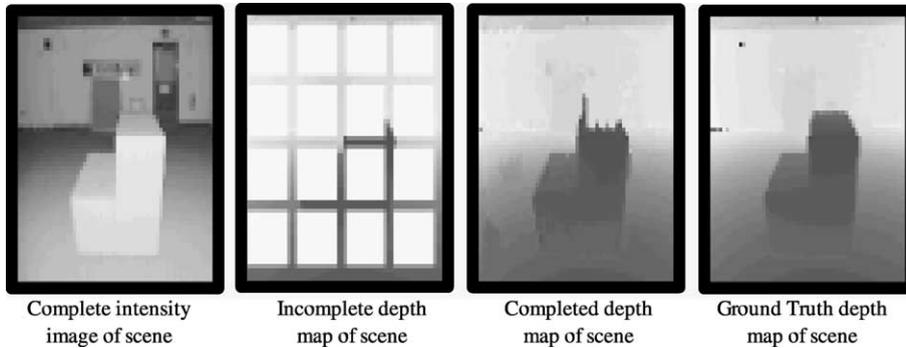


Fig. 13. $2\frac{1}{2}$ D completion of sparse depth maps [68].

to be completed (or synthesised) but instead relies on knowledge that can be inferred from the visible scene portions (i.e., implicit knowledge).

Additional work in $2\frac{1}{2}$ D completion has also been considered in the construction of solid models from 2D sketches [80]. In this work, the ‘frontal geometry of the 2D sketch’ (i.e., $2\frac{1}{2}$ D contour sketch) is initially derived from 2D using the techniques of [79]. Completion is then performed using geometric face and vertex addition through a method of greedy hypothesis construction and selection. In addition, special-case symmetry based completion based on [79] is also considered. This work is, however, limited in its purely geometric view of completion and currently only offers completion of simple 2D geometric sketches consisting of straight lines with right-angled intersections (i.e., block-type objects). Additionally, it assumes that the object has been sketched from the “most informative viewpoint” (i.e., no hidden portions exist that cannot be geometrically inferred from the visible portion) and that the sketch represents a singularly connected object (i.e., no frontal occlusion or multiple object consideration).

In terms of volume completion it most closely follows the surface and contour based completion paradigm of [37] although also exhibits the use of “world knowledge” through its use of geometric constraint. To some limited extent an awareness of “geometric volume” is also shown in terms of inferring hidden vertex junctions and faces to create a fully constrained solid object. However, in summary this work only offers very limited aspects of generalised volume completion due its reliance on the geometric properties, and hence the possible geometric ‘good continuation,’ of the objects it considers.

3.4. 3D surface filling

The topic of 3D surface filling from within the 3D modelling domain is another area related to visual completion. Here, the problem of hole filling in automated model acquisition relates directly to the idea of plausible completion.

Work in this area is fairly advanced and was notably used on the Michelangelo project—a complete 3D reconstruction of Digital Michelangelo’s David [11]. This

work concentrates on the merging of surfaces through diffusion between opposing edges of a surface hole or missing volume. The zero set, with respect to a pre-determined distance field, then forms a surface approximation over the diffused region. This work follows on from earlier mesh based reconstruction techniques [10] that used a similar method of blending distance functions together with space-carving based on a priori knowledge for the line of sight between the laser scanner and visible surfaces. Additionally there is a considerable body of work using regular mesh construction or point cloud interpolation for surface filling (e.g., [6,2,13]). These techniques primarily concentrate on the surface reconstruction and perform surface filling almost as an implicit by-product of this process—“*Reconstruct this surface—if a section is missing then just ‘fill over it.’*” By contrast, Curless/Davis [10,11] and Sogawa [55] treat surface filling as an explicit post-process to initial surface reconstruction—“*This section can be reconstructed but this section cannot therefore re-process it separately.*” This distinction is important in terms of considering scene evidence or world knowledge for surface completion although whether visual completion naturally occurs as part of a general scene interpretation or as a separate visual process remains an open question in the study of human visual completion. In general, a good overview of previous work in this area is given in the review by Davis et al. [11].

More recent work in this area [55] extends [11] by using a pre-processing technique to correct inconsistencies in the signed distance field used for surface reconstruction, resulting in hole filling success in previously difficult cases. This method performs correction by propagating a localised consensus from neighbouring voxels into inconsistent regions giving support to our concept of visual propagation from work in this area. Subsequent work on this area, specifically on polygon models [43,65,33], facilitates smooth surface completion via similar propagation based techniques.

Whilst this set of techniques is well suited to the continuous completion of holes within existing surfaces, it is less suited to the completion of large 3D enclosures because it inherently pursues a smooth surface completion over the void within the surface regardless of localised surface texture or surrounding contour completion constraints. As such it is not suited to instances where plausible completion may be bounded by existing contour or similarity constraints and surface data may be very sparse as shown in Figs. 14(a) and (b). In fact, the authors of [11] cite greater control over surface shape with respect to constraining factors (user specified or otherwise) as an area for future work.

In general, this work supports generalised visual propagation in terms of “smooth” surface shape completion through a propagation process akin to surface relatability (Kellman Shipley [37]) as a subset of wider volume completion (Tse [72]). However, as shown in Fig. 14 it falls short of the desired *plausible completion* criteria previously identified.

3.5. Completion using symmetry

Symmetry occurs frequently within both natural and built environments or objects and is thought to feature heavily in human visual recognition [81]. It thus appears a natural choice for use in the visual completion of occluded objects [85].

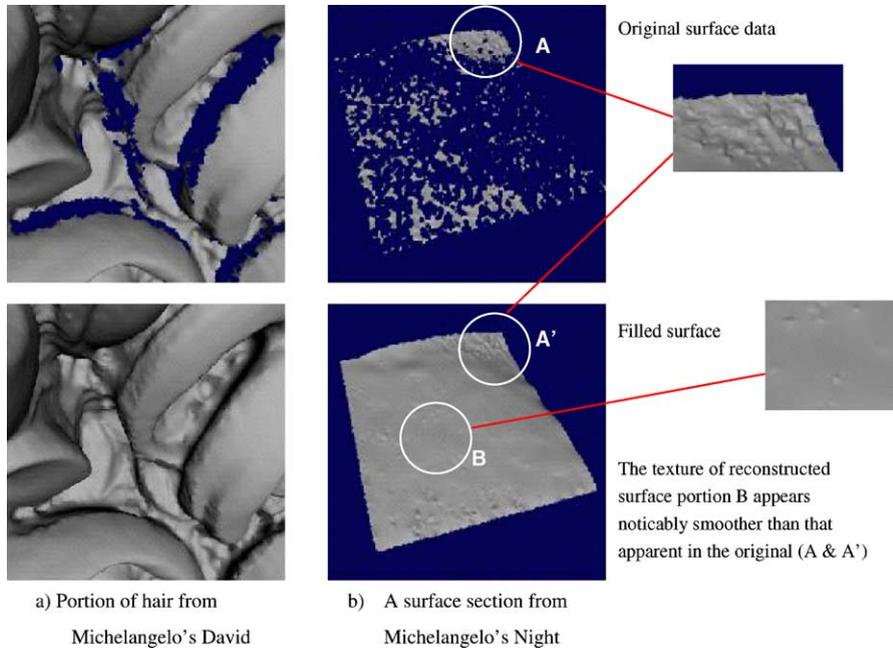


Fig. 14. Examples of 3D surface filling from [11].

However, whilst the symmetry properties of isolated 3D objects and 2D shapes are relatively easy to obtain and reason with (e.g. [85,67,60]), the identification of consistent sets of localised symmetries in a complex 3D scene remains a difficult problem [22] (e.g., finding the symmetrical components in Fig. 1B). This problem is made increasingly difficult when dealing in $2\frac{1}{2}$ D to 3D completion where the goal is to identify potential axes of symmetry based on limited knowledge of the object (i.e., potentially only one side of the symmetrical axis in some cases).

Work in detecting symmetry in isolated 2D objects is well established [85] and has thus been used for visual completion where necessary based on the idea of a continuous symmetry measure [84,86]. This extends the idea of symmetry from a discrete feature (i.e., *Is symmetric? yes/no*) to consider symmetry as a continuous measure of the minimum effort required to turn a shape into a symmetric shape.

Some work in this area has been extended to consider detecting 3D symmetry from simple 2D geometric sketches [79]. This work has been further extended to consider 3D completion of these objects [80], but suffers from a number of limitations— notably its reliance on the simple geometric nature of the sketches it successfully completes and possible ambiguity in interpreting the initial 2D representation (cf., $2\frac{1}{2}$ D completion). Furthermore, this work is also limited by its consideration of objects that are noiseless, isolated, and unoccluded.

From the literature, work on symmetry in real 3D data appears limited to identify symmetry axes and the recognition of objects based on their symmetric properties

[21,67]. This work is again limited to isolated objects rather than the more general case of identifying local object symmetries in a larger scene. A body of work is however present in identifying localised symmetry in isolated objects following from similar work in 2D [64]. Work is also available for detecting symmetries in multiple 2D object scenes based on edge detection and matching [52]. Again, this work is limited to simple scenes without consideration of the general effects of noise or occlusion on the process.

A potential solution to this difficult problem, posed by $2\frac{1}{2}$ D localised symmetry identification, could involve either the segmentation of individual scene objects prior to symmetry analysis, or some kind of (possibly exhaustive) search for subsets of data satisfying local symmetry constraints for a given axis. It would appear this problem remains largely unsolved for more complex 3D scenes such as those of Fig. 1B. However, an extension to the approaches of [64,52], perhaps utilising the continuous symmetry measure of [86] together with common 3D feature extraction techniques may offer a viable future solution in this area.

Despite these problems the use of symmetry, within the context of volume completion, can be seen most clearly as an example of completion through world knowledge (e.g., Fig. 3B) and as such would probably only form a small part, possibly through constraint, to a more general solution in this area. After all, few objects or scenes (especially in the natural world) are perfectly symmetric.

3.6. *Completing surfaces from contours*

Given that, based on earlier completion work [7], contours can be interpolated into occluded scene regions the derivation of surfaces from these contours needs to be itself considered as an aspect of volume completion.

A considerable portion of work in this area has been carried out with regard to 2D imaging and the 3D interpretation of 2D line drawings [79] or line structure extracted from a 2D images (cf. early work [3]). More recent work in this area uses contour identification and matching/tracking over a set of 2D images to recreate the underlying 3D structure (e.g. [9,74,89,61,56]). Similarly, shape from silhouette work has also been carried out by a number of researchers (e.g. [42,17]), where a series of silhouettes of an object, as opposed to image contours, are used to estimate 3D surface information. Generally, both techniques operate by considering the edge contours of an object over multiple views (image contours or silhouettes) from differing object rotations or camera positions. The inter-view relationships (a priori or derived) together with the relative feature placement in each view allows the derivation of 3D surface shape. Overall work in this area is now a well-established branch of 2D to 3D vision and a subset of the more general idea of tracking generic image features (be they lines, corners, blobs, etc.) over a set of images to recover 3D structural information [26]. The problem solved here is, however, subtly different from that posed in completing 3D surfaces from contours.

In our completion scenario, we have explicit knowledge of the $2\frac{1}{2}$ D contours from which likely completions into unknown scene portions can be estimated—unlike approaches based on 2D line drawings. Thus, we have explicit 3D knowledge of

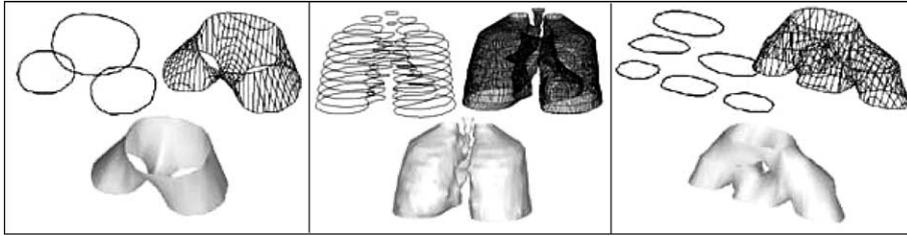


Fig. 15. Branching in surface from contour reconstruction.

our contours but not the surfaces which they bound. This is quite a unique reconstruction situation. If we were reconstructing a scene from a range scanner or multiple view vision, we would have both evidence of the contours and similarly of the surfaces they bound as both would be present in the intensity/range image(s)⁵ [22]. Here, we have a set of contours, or possible contours, evidence of surfaces in the visible portion of the scene (for possible continuation criterion, cf. Kellman and Shipley [37]) and no explicit knowledge of surfaces in the scene portion we wish to complete. Work in reconstruction problems of this kind is present in two different sub-fields of 3D vision—reconstructing surfaces from contour slices on an object (common to 3D medical MRI type vision) and in reconstructing surfaces from sparse 3D structure (common in feature based sparse stereo techniques).

3.6.1. Surfaces from contour slices

Looking first at reconstructing surfaces from contour slices, or more formally a set of regularly spaced concentric surface contours, a body of early work exists based on triangulation of surfaces between contours by Keppel [38]. A good overview of this work, including later extensions to consider the problems of branching between contour layers (see Fig. 15), is given in [31]. This work concentrates on finding the minimum cost triangulation between two contours with respect to a defined cost metric (e.g., minimum surface area or enclosed volume) using triangulated tiling. Essentially tackling this problem involves two important and inter-related sub-problems:

- *The correspondence problem:* common to many aspects of computer vision problems but here concerned with which points on each contour should be connected as triangular mesh vertices to provide an error free and realistic surface representation.
- *The branching problem:* should slices contain several contours (e.g., Fig. 15), how to solve the M to N contour correspondence problem to provide a realistic surface representation ($M \neq N$; $N, M > N$).

⁵ Surface presence is implied by the presence of surface range data for range images whilst the absence of detected contours together with possible intensity similarities or gradients imply the same for intensity images.

Unfortunately these sub-problems cannot be solved uniquely and consequently several methods of solution have been proposed with a varying degrees of success (for an overview see [31,48]). Although well suited to surface reconstruction between similar contours many approaches begin to fail when contour shape or branching varies greatly between slices resulting in erroneous surface triangulation [48]. However, some recent improvements to triangulation based techniques have attempted to tackle these problems by assuming the underlying shape follows an assumed definition of natural topology (cf. *world knowledge*) to successfully overcome the unnatural branching constructions produced from earlier work in this area [8].

Some more recent work in this area considers the use of Euclidean distance fields as an alternative method of solution [32,40]. Work by Jones [32] encloses the contour slices within a volume and uses distance fields to calculate the distance from every point within the volume (voxel) to the closest contour point. From this an isosurfacing approach can be used to hypothesize a surface while still maintaining the original contour slices. This technique operates by circumnavigating the explicit correspondence and branching problems by considering the surface space as a distance field. It is shown to work well on both practical medical imaging problems and on known ‘difficult’ problems, from literature in this area, beyond those of earlier approaches in the field [32]. However, as the author states this technique is currently limited to use on closed contours and non-intersecting contours (within the slices).

Related and later work by Klein et al. [57,39,40] also made use of distance fields as a basis for a similar solution to this problem. In this work [40], the distance field is used to find a medial contour, a half-way contour between two slices, based on summing the adjacent distance fields at every point and locating the zero crossings (where the dominance of a given distance field changes). Once this medial contour between two contour slices has been found, the approach can then be applied recursively considering this previously derived medial contour as a new contour slice itself. This can be repeated until the desired level of contour density has been achieved—up to that of a dense range image if required. From this now simplified problem, in terms of the correspondence and branching issues, a polygonised mesh can be constructed using any of the wealth of suitable techniques available for mesh construction over a dense point cloud [47,18,27,2,13]. However, this technique can result in an overly complex triangulation and [40] simultaneously presents a fast simplification algorithm based on the calculated distance fields used for surface construction as a solution to this problem. In addition, numerous techniques for general mesh simplification exist within the literature based on augmentations to common mesh construction algorithms [47,18,27].

Overall this technique appears to work well on example problems and, although not explicitly mentioned, appears not to suffer from the complete contour constraint of the earlier work by Jones and Chen [32]. As such it offers a viable method for surface reconstruction from contours for use as an aspect of volume completion. Augmentations to the process may also allow the reconstruction to be tailored to incorporate other volume completion aspects as detailed previously.

3.6.2. Surfaces from sparse data

Surface reconstruction from contour slices can be thought of as a specific case of the more general problem of reconstruction from disparate points [32]—reconstruction from sparse data. Work in this area is primarily motivated from the interpretation of data from sparse stereo capture and generally involves 3D reconstruction from a set of sparse features (e.g., points, lines, and corners) whose presence and relative position are scene dependent [18].

A number of approaches in this area concern themselves with surface fitting to sparse data points using methods analogous to those used in range data segmentation to fit surfaces and feature descriptors to sparse 3D data (e.g. [28,25,63]). In [25], constraints are used to extend the common surface fitting approach to impose good surface continuity in the resulting surface model. Such constraints can be considered similar to those proposed by Kellman and Shipley [37] with respect to visual completion and hence an extension of this work into the area of visual completion may be possible. However, as later stated by one of the authors [63], this work is limited to fitting surfaces and identifying features in 3D data in which “*surface and curvature discontinuities are explicitly preserved.*” Thus although dealing with sparse range data this technique is performing a task closer to advanced surface interpolation, through the use of additional constraints, rather than that of 3D completion as we are considering here. As such this work is functionally limited to $2\frac{1}{2}$ D surface interpolation—a common paradigm for realising complete $2\frac{1}{2}$ D scenes from sparse stereo techniques [18].

Another approach is that of 3D model construction from sparse stereo capture. Recent work in constructing 3D models from sparse 3D data [49] has utilised triangulation on sparse data sets obtained using structure-from-motion over a sequence of 2D images. In this work, sparse corner and edge features were used as primes for the construction of a triangular mesh to describe the surfaces of the scene. Visibility constraints, encoding the visibility of features from a given viewpoints, are used to constrain the construction of the mesh towards a model globally consistent with structural data gathered across all viewpoints. This technique relies on useful constraining features being present in multiple images (i.e., not consistently occluded or hidden) and being present at the surface intersections and edges in order to guide the triangulation.

In terms of completing entirely unknown 3D enclosures, through volume completion, these luxuries of constraining knowledge are not explicitly available and may have to be inferred from visible scene portions prior to application of a technique such as this. An alternative would be to extend the principle of this technique, through iterative constraint optimisation, to consider more general constraints. Probabilistic inferences rather than absolute positional constraints could possibly be used to guide the reconstruction towards the most probable representation of an unknown scene portion (cf. fuzzy completion [75]).

Overall, despite the clear parallels to general 3D completion, work on both surface from contour and sparse data completion techniques relies on the presence of constraining information not present in our stated $2\frac{1}{2}$ D to 3D completion problem. As such they focus on completion through good continuation of contours and

surfaces—hence following the local view of completion shown by Tse [72] to be inadequate to many 3D completion cases. However, despite this limitation, this approach may offer a viable completion aspect into the wider volume completion framework from which the additional constraint, required here for success, can otherwise be derived.

4. Summary

Our examination and contrast of 3D completion in both perceptual psychology and computer vision show us that completion is approached very differently in each discipline but still a vein of commonality, in terms of usable visual attributes, can be seen running through both.

Perceptual psychology offers us several explanations for human visual completion. The most recent and advanced in terms of 3D completion is that of Volume Completion proposed by Tse [72]. Volume Completion is a combined and open-ended approach to complete 3D enclosures utilising contour and surface relatability, volume mergability, world knowledge, and pattern completion. The volume completion concept asserts that every surface has an associated volume and hence volume completion subsumes earlier surface based completion (Kellman and Shipley [37]) as a special case. Tse goes on to suggest that the volume completion process is a relaxation into a best-fit between given constraints and that completion processes are in general inherently probabilistic.

Tse's work is supported by both experimental evidence [71] and work by others in the field—notably in the work of van Lier et al. [75] with his suggestion, and supporting experimental evidence, for plausible completion through probabilistic ‘fuzzy completion.’ Further work by Tse [73] goes on to highlight the importance of contours and surfaces in volume completion and proposes a visual propagation technique as an explanatory theory for this area.

We have examined primarily 3D completion work in computer vision although a body of work in 2D similarly exists. Work explicitly in $2\frac{1}{2}$ D to 3D completion of range images is limited to completion of partially occluded surfaces and this work is itself limited to well-posed and isolated instances of this completion task [12,7,19]. Similar work exists in 3D building reconstruction and completion using $2\frac{1}{2}$ D structure from motion data and a probabilistic model of building features and textures [14]. This work is itself limited to the domain of architecture and both rely heavily upon aspects of “world knowledge” either specific to this domain [14] or more generally to geometric shape conformity [12,7,51]. Related work has also been carried out on $2\frac{1}{2}$ D depth map completion based on having at least partial knowledge (i.e., intensity) of areas to be completed [68] and similarly in completing $2\frac{1}{2}$ D sketches of simple geometric constructions [80]. Here again we see a reliance upon knowledge not readily available within the confines of the generalised volume completion problem although both offer a useful insight into achievable goals when we allow the propagation process to be closely governed by either specific scene knowledge [68] or geometric constraint [80].

By contrast, the relevant work in 3D surface filling [11,43,33,65] offers what can be considered a more generalised propagation based approach to isolated surface completion which follows the conceptual surface continuation idea of Kellman and Shipley [37]. However, here the primary consideration is surface completion, through merging over a hole, rather than completing 3D volumes with regard to other governing constraints. This is notable when this work is considered within the wider aspect of producing locally plausible completion in addition to filling a surface hole that is consistent with more global (i.e., higher level) aspects of the scene.

A less well-developed area is that of symmetry for 3D visual completion. Although work on symmetry identification has been successfully carried out in 2D [85] together with successful work on symmetry identification in 3D data sets [67,64] it would appear that identifying potential axes of localised symmetry in a $2\frac{1}{2}$ D scene remains an open, and potentially difficult problem. It is not clear how helpful advances in this area would be to the more general problem of 3D completion. Although symmetry naturally appeals to any proposed completion approach its application is limited to symmetrically regular objects/scenes—a very limited subset of general completion instances.

In our final area of computer vision in consideration, surface from contour completion [32,25,40,49], we again see the requirement for additional knowledge moving us away from the generalised volume completion problem. Here, however, unlike earlier knowledge dependent completion approaches we foresee potential for using surface/contour relatability as a sub-part of generalised volume completion where the required constraining knowledge is itself derived from the other identified aspects of volume mergability, world knowledge, and pattern completion.

Overall, the problem of 3D completion is certainly difficult—both in terms of understanding the psychological phenomenon and in terms of replicating this artificially in computer vision.

Many aspects of the perceptual psychology work highlighted here remain unexplored in the computer vision literature. From our examination, despite the many successes we have detailed, computer vision work in 3D completion is limited to considering, somewhat disjointly, a subset of the aspects attributed to visual completion by perceptual psychology and only mildly approaches our concept of visual propagation and that of Tse's volume completion derived from the psychological literature. From our earlier discussion of aspects attributable to volume completion it is notable that work in 3D pattern completion, volume mergability, and, to a lesser extent, surface/contour relatability is somewhat lacking. Instead, we see a variety of successful “world knowledge” dependent techniques themselves knowledge-limited to specific visual domains.

Work on unified completion approaches is also similarly limited. Here, we show an example of work in progress [4] utilising combined aspects of world knowledge, contour/surface relatability, and 3D pattern completion (Fig. 16). Utilising geometric shape fitting and completion [7,12,19] to achieve a constrained “good surface continuation” of the underlying shape for a scale model of Pisa Tower we use a novel 3D pattern completion technique based on [16] to achieve a realistic and plausible

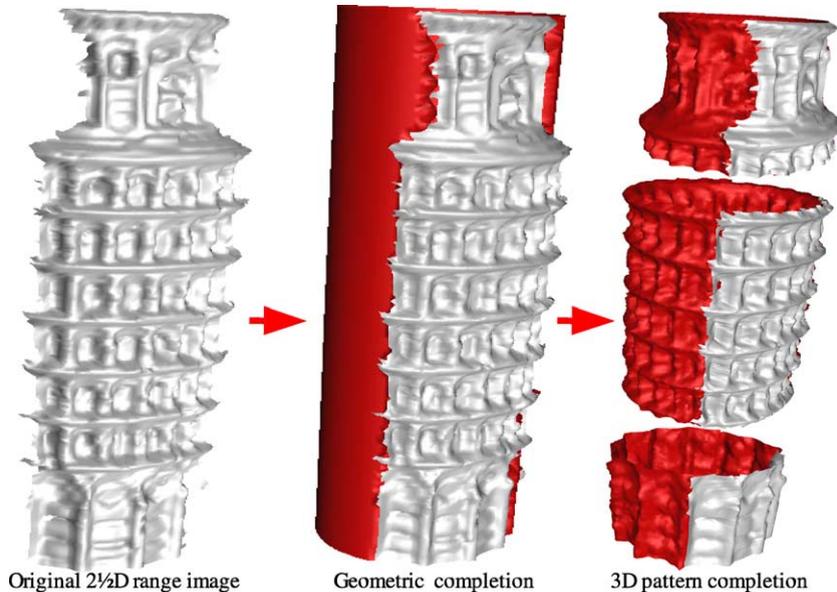


Fig. 16. 2½D to 3D completion of Pisa Tower.

completion of this 2½D range image (Fig. 16)—details in [4]. Concurrent work [58] has considered a similar approach in contrast with [4]. Despite the success of this work in progress, the technique still falls some way short of generalised volume completion in the computer vision arena.

Perhaps future computer vision work in this area, building upon the psychological accounts highlighted here, should further explore the available psychological hypotheses. Overall a realistic and plausible 3D completion process remains an open goal in computer vision research. Notably, several promising general areas in 3D volume completion that have to date received little attention in computer vision are:

- *Volume mergability*: the completion of objects through the identification and correspondence of disjoint volumes within the scene. Proposed by Tse [72] with strong visual supporting evidence (Fig. 10) and theoretical proposals [73].
- *3D pattern completion*: one of the main contributing aspects to the identified phenomenon of volume completion. Although a wide body of work exists in 2D image analysis and synthesis (e.g. [16]), similar work in 3D remains in its infancy [4,58].
- *Co-operative completion approaches*: multi-modal approaches utilising a breadth of available scene evidence and knowledge from all of the aspects of volume completion discussed here in an approach driven by the principle of completion via *visual propagation*. Akin to algorithms proposed in other computer vision domains [50] and representing the “holy grail” of work in this area.

Acknowledgments

The authors thank the following people for their kind assistance: Marc Pollefeys for permission to use the Arenberg Castle data set (Fig. 1). Craig Robertson for permission to use the toy data set (Fig. 1). Peter Tse for permission to reproduce Figs. 3, 6, 7, and 10 in part or in full from [72]. Rob van Lier to reproduce Fig. 9 from [75]. Antony Dick for permission to use the Downing Library models in Fig. 12. Abril Torres-Mendez for permission to reproduce Fig. 13 from [68]. James Davis for permission to reproduce Fig. 14 from his work related to [11]. Mark Jones for permission to reproduce Fig. 15 from his work related to [32].

This work was supported by the UK EPSRC and QinetiQ PLC (www.qinetiq.com).

References

- [1] M.K. Albert, P.U. Tse, The role of surface attraction in perceiving volumatic shape, *Perception* 29 (2000) 409–420.
- [2] N. Amenta, S. Choi, R.K. Kolluri, The power crust, in: *Proc. Sixth ACM Symposium on Solid Modeling and Applications*, ACM Press, New York, 2001, pp. 249–266.
- [3] D.H. Ballard, C.M. Brown, *Computer Vision*, Prentice Hall, New Jersey, 1982.
- [4] T.P. Breckon, R.B. Fisher, Non-parametric 3D surface completion, in: *Proc. of the 5th International Conference on 3D Digital. Imaging and Modelling (3DIM)*. IEEE Press, 2005, pp. 573–580.
- [5] J. Bruske, E. von Puttkamer, U.R. Zimmer, Spin-nfds: learning and preset knowledge for surface fusion—a neural fuzzy decision system, in: *Proc. ANZIIS 1993*, Perth, Western Australia, 1993.
- [6] J.C. Carr, R.K. Beatson, J.B. Cherrie, T.J. Mitchell, W.R. Fright, B.C. McCallum, T.R. Evans, Reconstruction and representation of 3D objects with radial basis functions, in: *Proc. 28th SIGGRAPH*, ACM Press, New York, 2001, pp. 67–76.
- [7] U. Castellani, S. Livatino, R.B. Fisher, Improving environment modelling by edge occlusion surface completion, in: *Internat. Symp. on 3D Data Processing Visualization and Transmission (3DPVT)*, Padova, Italy, June 2002, pp. 672–675.
- [8] S. Cheng, T.K. Dey, Improved constructions of Delaunay based contour surfaces, in: *Proc. ACM Symp. Solid Modeling and Applications*, 1999, pp. 322–323.
- [9] R. Cipolla, A. Blake, The dynamic analysis of apparent contours, in: *Proc. 3rd International Conference on Computer Vision*, Osaka, Japan, December 1990, pp. 616–623.
- [10] B. Curless, M. Levoy, A volumetric method for building complex models from range images, in: *Proc. Computer Graphics (SIGGRAPH)*, 1996.
- [11] J. Davis, S. Marschner, M. Garr, M. Levoy, Filling holes in complex surfaces using volumetric diffusion, in: *Proc. First International Symposium on 3D Data Processing, Visualization, Transmission*, June 2002.
- [12] F. Dell’Acqua, R. Fisher, Reconstruction of planar surfaces behind occlusions in range images, *IEEE Trans. Pattern Anal. Mach. Intell.* 24 (4) (2002) 569–575.
- [13] T.K. Dey, S. Goswami, Tight cocone: a water-tight surface reconstructor, in: *Proc. of the 8th ACM Symp. on Solid Modeling and Applications*, ACM Press, New York, 2003, pp. 127–134.
- [14] A.R. Dick, P.H.S. Torr, R. Cipolla, A Bayesian estimation of building shape using mcmc, in: *Proc. Eur. Conf. on Computer Vision (ECCV)*, vol. 2, 2002, pp. 852–866.
- [15] P.H.S. Dick, A.R. Torr, S.J. Ruffle, R. Cipolla, Combining single view recognition and multiple view stereo for architectural scenes, in: *Internat. Conf. on Computer Vision (ICCV)*, vol. 1, 2001, pp. 268–274.
- [16] A.A. Efros, T.K. Leung, Texture synthesis by non-parametric sampling, in: *Proc. IEEE Internat. Conf. on Computer Vision*, Corfu, Greece, September 1999, pp. 1033–1038.

- [17] Y.H. Fang, H.L. Chou, Z. Chen, 3D shape recovery of complex objects from multiple silhouette images, *Pattern Recogn. Lett.* 24 (9–10) (2003) 1279–1293.
- [18] O. Faugeras, *Three-Dimensional Computer Vision: A Geometric Viewpoint*, MIT Press, 1993.
- [19] R.B. Fisher, Solving architectural modelling problems using knowledge, in: *Proc. 4th Internat. Conf. on 3-D Digital Imaging and Modeling*, Banff, Canada, 2003, pp. 343–351, invited talk.
- [20] R.B. Fisher, *From surfaces to objects: recognizing objects using surface information and object models*, PhD thesis, Department of Artificial Intelligence, University of Edinburgh, UK, 1986.
- [21] P.J. Flynn, 3-D object recognition with symmetrical models: symmetry extraction and encoding, *Pattern Anal. Mach. Intell.* 16 (8) (1994) 814–818.
- [22] D.A. Forsyth, J. Ponce, *Computer Vision: A Modern Approach*, Prentice-Hall, New Jersey, 2002.
- [23] D. Geiger, H. Pao, N. Rubin, Salient and multiple illusory surfaces, in: *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1998, pp. 118–124.
- [24] G. Guy, G. Medioni, Inferring global perceptual contours from local features, in: *IEEE Image Understanding Workshop*, 1993, pp. 881–892.
- [25] G. Guy, G. Medioni, Inference of surfaces, 3D curves, and junctions from sparse, noisy, 3d data, *Pattern Anal. Mach. Intell.* 19 (11) (1997) 1265–1277.
- [26] R.I. Hartley, A. Zisserman, *Multiple View Geometry in Computer Vision*, Cambridge University Press, Cambridge, 2000.
- [27] A. Hilton, A. Stoddart, J. Illingworth, T. Windeatt, Marching triangles: range image fusion for complex object modelling, in: *Proc. Internat. Conf. on Image Processing 96*, 1996, p. 172.
- [28] T. Horikoshi, S. Suzuki, 3D parts decomposition from sparse range data information criterion, in: *Proc. Computer Vision and Pattern Recognition*, 1993, pp. 168–173.
- [29] D.A. Huffman, Impossible objects as nonsense sentences, *Machine Intelligence* 6 (1971) 295–323.
- [30] R.C. Jain, R. Kasturi, B.G. Schunck, *Machine Vision*, McGraw-Hill, New York, 1995.
- [31] M.W. Jones, PhD Thesis, chapter on ‘Contour Methods.’ Department of Computer Science, University of Wales Swansea, 1994. Singleton Park, Swansea, SA2 8PP, UK.
- [32] M.W. Jones, M. Chen, A new approach to the construction of surfaces from contour data, *Comput. Graph. Forum* 13 (3) (1994) 75–84.
- [33] T. Ju, Robust repair of polygonal models, *ACM Trans. Graph.* 23 (3) (2004) 888–895.
- [34] G. Kanizsa, Subjective Contours, *Scientific American*, 1976.
- [35] G. Kanizsa, *Organization in Vision: Essays on Gestalt Perception*, Praeger, 1979.
- [36] G. Kanizsa, Amodal completion: seeing or thinking? in: J. Beck (Ed.), *Organisation and Representation in Perception*, Erlbaum, New Jersey, 1982, pp. 167–190.
- [37] P.J. Kellman, T.F. Shipley, A theory of visual interpolation in object perception, *Cognitive Psychol.* 23 (2) (1991) 141–221.
- [38] E. Keppel, Approximating complex surfaces by triangulation of contour lines, *IBM J. Res. Dev.* 19 (1975) 2–11.
- [39] R. Klein, A. Schilling, Fast distance field interpolation for reconstruction of surfaces from contours, in: *Eurographics, Short Papers and Demonstrations*, pages University of Bonn, Technical Report WSI-99-15, Milano, Italy, 1999.
- [40] R. Klein, A. Schilling, W. Straßer, Reconstruction and simplification of surfaces from contours, *Graph. Models* 62 (6) (2000) 429–443.
- [41] K. Koffka, *Principles of Gestalt Psychology*, Harcourt, Brace and Co., 1935.
- [42] A. Laurentini, How many 2D silhouettes does it take to reconstruct a 3D object, *Comput. Vis. Image Understand.* 67 (1) (1997) 81–87.
- [43] P. Liepa, Filling holes in meshes, in: *SGP '03: Proc. of the Eurographics/ACM SIGGRAPH Symposium on Geometry processing*, Eurographics Association, 2003, pp. 200–205.
- [44] X. Liu, D. Wang, A boundary-pair representation for perception modeling, in: *Proc. Internat. Joint Conf. on Neural Networks*, 1999, 1999.
- [45] X. Liu, D. Wang, Perceptual organization based on temporal dynamics, in: *Proc. Internat. Joint Conf. on Neural Networks*, 1999, 1999.
- [46] J.L. Locher, *E.C. Escher: Life and Work*, Abrams, 1992.

- [47] W. Lorensen, H. Cline, Marching cubes: a high resolution 3D surface construction algorithm, *Comput. Graph.* 21 (1987) 163–169.
- [48] M. Lounsbery, C. Loop, S. Mann, D. Meyers, J. Painter, T. DeRose, K. Sloan, Testbed for the comparison of parametric surface methods, in: L.A. Ferrari, R.J.P. de Figueiredo, (Eds.), *Curves and Surfaces in Computer Vision and Graphics (Proc. SPIE)*, vol. 1251, 1990, pp. 94–105.
- [49] A. Manassis, A. Hilton, P. Palmer, P. McLauchlan, X. Shen, Reconstruction of scene models from 3D sparse structure, *Comput. Vis. Pattern Recogn.* (2000) 14–15.
- [50] D. Marr, *Vision: a computational investigation into the human representation and processing of visual information*, W.H. Freeman, 1982.
- [51] G. Miller, Reconstruction of unobserved sides of building columns, Technical report, School of Informatics, University of Edinburgh, May 2003. Undergraduate Honours Dissertation, Supervisor: R.B. Fisher.
- [52] D.P. Mukherjee, A. Zisserman, J.M. Brady, Shape from symmetry-detecting and exploiting symmetry in affine images, *Philos. Trans. R. Soc. London A* 351 (1995) 77–106.
- [53] K. Nakayama, S. Shimojo, Z.J. He, Visual surface representation: a critical link between lower-level and higher-level vision, in: *Visual Cognition: An Invitation to Cognitive Science*, vol. 2, MIT Press, Cambridge, 1995, pp. 1–70.
- [54] N. Rubin, The role of junctions in surface completion and contour matching, *Perception* 30 (2001) 339–366.
- [55] R. Sagawa, K. Ikeuchi, Taking consensus of signed distance field for complementing unobservable surface, in: *Proc. 3D Imaging and Modelling (3DIM)*, October 2003, pp. 410–417.
- [56] J. Sato, R. Cipolla, Affine reconstruction of curved surfaces from uncalibrated views of apparent contours, *Pattern Anal. Mach. Intell.* 21 (11) (1999) 1188–1198.
- [57] A.G. Schilling, R. Klein, Fast generation of multiresolution surfaces from contours, in: *Proc. 9th Eurographics Workshop on Visualization in Scientific Computing*, April 1998, pp. 88–99, Conference held in Blaubeuren, Germany.
- [58] A. Sharf, M. Alexa, D. Cohen-Or, Context-based surface completion, *ACM Trans. Graph.* 23 (3) (2004) 878–887.
- [59] F. Stulp, F. Dell’Acqua, R.B. Fisher, Reconstruction of surfaces behind occlusions in range images, in: *Proc. 3rd Internat. Conf. on 3-D Digital Imaging and Modeling (3DIM01)*, 2001, pp. 232–239.
- [60] C.M. Sun, J. Sherrah, 3D symmetry detection using the extended gaussian image, *Pattern Anal. Mach. Intell.* 19 (2) (1997) 164–168.
- [61] R. Szeliski, R. Weiss, Robust shape recovery from occluding contours using a linear smoother, *Int. J. Comput. Vis.* 28 (1) (1998) 27–44.
- [62] H. Takeichi, H. Nakazawa, I. Murakami, S. Shimojo, The theory of the curvature-constraint line for amodal completion, *Perception* 24 (1995) 267–277.
- [63] C.K. Tang, G. Medioni, Inference of integrated surface, curve, and junction descriptions from sparse 3D data, *Pattern Anal. Mach. Intell.* 20 (11) (1998) 1206–1223.
- [64] S. Tari, J. Shah, Local symmetries of shapes in arbitrary dimension, in: *Proc. Internat. Conf. on Computer Vision*, 1998, pp. 1123–1128.
- [65] L.S. Tekumalla, E. Cohen, A hole-filling algorithm for triangular meshes, Technical Report UUCS-04-019, School of Computing, University of Utah, December 2004.
- [66] C.J. Tessa, R.J. van Lier, Global visual completion of quasi-regular shapes, *Perception* 31 (2002) 969–984.
- [67] T. Thorhallsson, Detecting bilateral symmetry of 3D point sets from affine views, in: *Proc. British Machine Vision Conference*, page Shape, 1996.
- [68] L.A. Torres-Mendez, G. Dudel, Range synthesis for 3D environment modeling, in: *Proc. IEEE Workshop on Applications of Computer Vision*, 2002.
- [69] P. Tse, M. Albert, Amodal completion in the absence of image tangent discontinuities, *Perception* 27 (1998) 455–464.
- [70] P.U. Tse, Illusory volumes from conformation, *Perception* 27 (8) (1998) 977–992.
- [71] P.U. Tse, Complete mergability and amodal completion, *Acta Psychol.* 102 (1999) 165–201.

- [72] P.U. Tse, Volume completion, *Cognitive Psychol.* 39 (1999) 37–68.
- [73] P.U. Tse, A contour propagation approach to surface filling-in and volume formation, *Psychol. Rev.* 109 (1) (2002) 91–115.
- [74] R. Vaillant, O.D. Faugeras, Using extremal boundaries for 3-D object modeling, *IEEE Pattern Anal. Mach. Intell.* 14 (2) (1992) 157–173.
- [75] R.J. Van Lier, Investigating global effects in visual occlusion: from a partly occluded square to the back of a tree trunk, *Acta Psychol.* 102 (1999) 203–220.
- [76] R.J. Van Lier, P.A. Vann Der Helm, E.L.J. Leeuwenberg, Integrating global and local aspects of visual occlusion, *Perception* 23 (8) (1994) 883–903.
- [77] R.J. Van Lier, P.A. Vann Der Helm, E.L.J. Leeuwenberg, Competing global and local completions in visual occlusion, *Perception* 21 (3) (1995) 571–583.
- [78] R.J. Van Lier, J. Wagemans, From images to objects: global and local completions of self occluded parts, *J. Exp. Psychol.: Hum. Percept. Perform.* 25 (6) (1999) 1721–1741.
- [79] P.A.C. Varley, R.R. Martin, Constructing boundary representation solid models from a two-dimensional sketch-frontal geometry and sketch categorisation, in: *Proc. First UK–Korea Workshop on Geometric Modeling and Computer Graphics*, Kyung Moon Publishers, 2000, pp. 113–128.
- [80] P.A.C. Varley, R.R. Martin, Constructing boundary representation solid models from a two-dimensional sketch—topology of hidden parts, in: *First Proc. UK–Korea Workshop on Geometric Modeling and Computer Graphics*, Kyung Moon Publishers, 2000, pp. 129–144.
- [81] T. Vetter, T. Poggio, H. Bulthoff, 3D object recognition: symmetry and virtual views, in: *MIT AI Memo*, 1992.
- [82] L.R. Williams, D.W. Jacobs, Stochastic completion fields: a neural model of illusory contour shape and salience, *Neural Comput.* 9 (4) (1997) 837–858.
- [83] D. Wouterlood, F. Boselie, A good continuation model of some occlusion phenomena, *Psychol. Res.* 54 (1992) 267–277.
- [84] H. Zabrodsky, S. Peleg, D. Avnir, Continuous symmetry measures, *Nature* 360 (17) (1992) 626.
- [85] H. Zabrodsky, S. Peleg, D. Avnir, Completion of occluded shapes using symmetry, *Proc. Comput. Vis. Pattern Recogn.* (1993) 678–679.
- [86] H. Zabrodsky, S. Peleg, D. Avnir, Symmetry as a continuous feature, *Pattern Anal. Mach. Intell.* 17 (12) (1995) 1154–1166.
- [87] Q. Zhang, M. Idesawa, K. Mogi, 3D volumetric object perception in binocular vision *Proc. 9th Internat. Conf. on Neural Information Processing*, vol. 4, IEEE Press, 2002, pp. 1644–1648.
- [88] Q. Zhang, M. Idesawa, Y. Sakaguchi, Pantomime effect in the perception of volumetric transparent illusory object with binocular viewing, *Jpn. J. Appl. Phys.* (1998) L329–L332.
- [89] J.Y. Zheng, Acquiring 3-D models from a sequence of contours, *Pattern Anal. Mach. Intell.* 16 (2) (1994) 163–178.