

# From Multiview Image Curves to 3D Drawings

Anil Usumezbas<sup>1</sup>, Ricardo Fabbri<sup>2</sup>(✉), and Benjamin B. Kimia<sup>3</sup>

<sup>1</sup> SRI International, Providence, USA

[anil.usumezbas@sri.com](mailto:anil.usumezbas@sri.com)

<sup>2</sup> Polytechnic Institute, State University of Rio de Janeiro, Rio de Janeiro, Brazil

[rfabbri@iprj.uerj.br](mailto:rfabbri@iprj.uerj.br)

<sup>3</sup> School of Engineering, Brown University, Providence, USA

[benjamin\\_kimia@brown.edu](mailto:benjamin_kimia@brown.edu)

**Abstract.** Reconstructing 3D scenes from multiple views has made impressive strides in recent years, chiefly by correlating isolated feature points, intensity patterns, or curvilinear structures. In the general setting – without controlled acquisition, abundant texture, curves and surfaces following specific models or limiting scene complexity – most methods produce unorganized point clouds, meshes, or voxel representations, with some exceptions producing unorganized clouds of 3D curve fragments. Ideally, many applications require structured representations of curves, surfaces and their spatial relationships. This paper presents a step in this direction by formulating an approach that combines 2D image curves into a collection of 3D curves, with topological connectivity between them represented as a 3D graph. This results in a **3D drawing**, which is complementary to surface representations in the same sense as a 3D scaffold complements a tent taut over it. We evaluate our results against truth on synthetic and real datasets.


**Keywords:** Multiview stereo · 3D reconstruction · 3D curve networks · Junctions

## 1 Introduction

The automated 3D reconstruction of *general* scenes from multiple views obtained using conventional cameras, under uncontrolled acquisition, is a paramount goal of computer vision, ambitious even by modern standards. While a fully complete working system addressing all the underlying challenges is beyond current technology, significant progress has been made in the past few years using approaches that fall into three broad classes, depending on whether one focuses on correlating isolated points, surface patches, or curvilinear structures across views, as described below.

A vast majority of multiview reconstruction methods rely on correlating isolated interest points across views to produce an unorganized 3D cloud of points.

**Electronic supplementary material** The online version of this chapter (doi:[10.1007/978-3-319-46493-0\\_5](https://doi.org/10.1007/978-3-319-46493-0_5)) contains supplementary material, which is available to authorized users.



**Fig. 1.** Our approach transforms calibrated views of a scene into a “3D drawing” – a graph of 3D curves meeting at junctions. Each curve is shown in a different color. (Please zoom in to examine closely. The 3D model is available as supplementary data.) (Color figure online)

The **interest-point-based approach** has been highly successful in reconstructing large-scale scenes with *texture-rich images*, in systems such as in Phototourism and recent large-scale 3D reconstruction work [6, 15, 34, 47]. Despite their manifest usefulness, these methods generally cannot represent smooth, textureless regions (due to the sparsity of interest points in image regions with homogeneous appearance), or regions that change appearance drastically across views. This limits their applicability, especially in man-made environments [28] and objects such as cars [27], non-Lambertian surfaces such as that of the sea, appearance variation due to changing weather [2], and wide baseline [46].

Another approach matches intensity patterns across views using multiview stereo, producing denser point clouds or mesh reconstructions. **Dense multi-view stereo** produces detailed 3D reconstructions of objects imaged under controlled conditions by a large number of precisely calibrated cameras [3, 25, 40–43, 48]. For general, complex scenes with various kinds of objects and surface properties, this approach has shown most promise towards obtaining an accurate and dense 3D model of a given scene. Homogeneous areas, such as walls of a corridor, repeated texture, and areas with view-dependent intensities create challenges for these methods.

A smaller number of techniques correlate and reconstruct image **curvilinear structure** across views, resulting in 3D curvilinear structure. Pipelines based on straight lines (see [11, 20, 31] for recent reviews), algebraic and general curve features [8–10, 21, 23, 29, 36] have been proposed, but some lack generality, *e.g.*, requiring specific curve models [26]. The 3D Curve Sketch system [7, 8, 10] operates on multiple views by pairing curves from two arbitrary “hypothesis views” at a time via epipolar-geometric consistency. A curve pair reconstructs to a 3D curve fragment hypothesis, whose reprojection onto several other “confirmation views” gathers support from subpixel 2D edges. The curve pair hypotheses with enough support result in an unorganized set of 3D curve fragments, the “3D



**Fig. 2. 3D drawings for urban planning and industrial design.** A process from professional practice for communicating solution concepts with a blend of computer and handcrafted renderings [44, 50]. New designs are often based off real object references, mockups or massing models for selecting viewpoints and rough shapes. These can be modeled *manually* in, *e.g.*, Google Sketchup (top-left), in some cases from reference imagery. The desired 2D views are rendered and *manually* traced into a reference curve sketch (center-left, bottom-left) easily modifiable to the designer’s vision. The stylized drawings to be presented to a client are often produced by *manually* tracing and painting over the reference sketch (right). Our system can be used to generate reference *3D curve drawings* from video footage of the real site for urban planning, saving manual interaction, providing initial information such as rough dimensions, and aiding the selection of pose, editing and tracing. The condensed 3D curve drawings make room for the artist to overlay his concept and harness imagery as a clean reference, clear from details to be redesigned.

Curve Sketch”. While the resulting 3D curve segments are visually appealing, they are fragmented, redundant, and lack explicit inter-curve organization.

The plethora of multiview representations, as documented above, arise because 3D structures are geometrically and semantically rich [12, 32]. A building, for example, has walls, windows, doorways, roof, chimneys, etc. The structure can be represented by sample points (*i.e.*, unorganized cloud of points) or a surface mesh where connectivity among points is captured. This representation, especially when rendered with surface albedo or texture, is visually appealing. However, the representation also leaves out a great deal of semantic information: which points or mesh areas represent a window or a wall? Which two walls are adjacent? The representation of such components, or parts, requires an explicit representation of part boundaries such as ridges, as well as where these boundaries come together, such as junctions.

The same point can equally arise if objects in the scene were solely defined by their curve structures. A representation of a building by its ridges may usually give an appealing impression of its structure, but it fails to identify the walls, *i.e.*, which collection of 3D curves bound a wall and what its geometry is. Both surfaces

and curves are important and needed across the board, *e.g.*, in applications such as robotics [4], urban planning and industrial design [44, 50], Fig. 2.

In general, image curve fragments are attractive because they have good localization, they have greater invariance than interest points to changes in illumination, are stable over a greater range of baselines, and are typically denser than interest points. Furthermore, the reflectance or ridge curves provide boundary condition for surface reconstruction, while occluding contour variations across views lead to surfaces [37, 39, 45]. Recent studies strongly support the notion that image curves contain much of the image information [17, 19, 33, 38]. Moreover, curves are structurally rich as reflected by their differential geometry, a fact which is exploited both in recent computer systems [1, 8, 10, 33] and perception studies [13, 33].

This paper develops the technology to process a series of (intrinsic and extrinsically) calibrated multiview images to generate a *3D curve drawing* as a graph of 3D curve segments meeting at junctions. The ultimate goal of this approach is to integrate the 3D curve drawing with the traditional recovery of surfaces so that 3D curves bound the 3D curve segments, towards a more semantic representation of 3D structures. The 3D curve drawing can also be of independent value in applications such as fast recognition of general 3D scenery [23], efficient transmission of general 3D scenes, scene understanding and modeling by reasoning at junctions [22], consistent non-photorealistic rendering from video [5], modeling of branching structures, among others [18, 24, 30].

The paper is organized as follows. In Sect. 2 we review the 3D curve sketch, identify three shortcomings and suggest solutions to each, resulting in the *Enhanced Curve Sketch*. Since the original 3D curve sketch was built around a few views at a time, it did not address fundamental issues surrounding integration of information from numerous views. Section 3 presents as our main contribution the multiview integration of information both at edge- and curve-level, which naturally leads to junctions. Section 4 validates the approach using real and synthetic datasets.

## 2 Enhanced 3D Curve Sketch

Image curve fragments formed from grouped edges are central to our framework. Each image  $V^v$  at view  $v = 1, \dots, N$  contains a number of curves  $\gamma_i^v$ ,  $i = 1, \dots, M^v$ . Reconstructed 3D curve fragments are referred as  $\Gamma_k$ ,  $k = 1, \dots, K$ , whose reprojection onto view  $v$  is  $\gamma^{k,v}$ . Indices may be omitted where clear from context.

The initial stage of our framework is built as an extension of the hypothesize-and-verify 3D Curve Sketch approach [10]. We use the same hypothesis generation mechanism with a novel verification step performing a finer-level analysis of image evidence and significantly reducing the fragmentation and redundancy in the 3D models.

Two image curves  $\gamma_{l_1}^{v_1}$  and  $\gamma_{l_2}^{v_2}$  are paired from two distinct views  $v_1$  and  $v_2$  at a time, the *hypothesis views*, provided they have sufficient epipolar overlap [10]. The verification of these  $K$  curve pair hypotheses, represented as  $\omega_k$ ,



$k = 1, \dots, K$  with the corresponding 3D reconstruction denoted as  $\mathbf{F}_k$ , gauges the extent of edge support for the reprojection  $\gamma^{k,v}$  of  $\mathbf{F}_k$  onto another set of *confirmation views*,  $v = v_{i_3}, \dots, v_{i_n}$ . An image edge in view  $v$  supports  $\gamma^{k,v}$  if it is sufficiently close in distance *and* orientation. The total support a hypothesis  $\omega$  receives from view  $v$  is

$$S_{\omega_k}^v \doteq \int_0^{L^{k,v}} \phi(\gamma^{k,v}(s)) ds, \quad (1)$$

where  $L^{k,v}$  is the length of  $\gamma^{k,v}$ , and  $\phi(\gamma(s))$  is the extent of edge support at  $\gamma(s)$ . A view is considered a *supporting view* for  $\omega_k$  if  $S_{\omega_k}^v > \tau_v$ . Evidence from confirmation views is aggregated in the form


$$S_{\omega_k} \doteq \sum_{v=i_3}^{i_n} [S_{\omega_k}^v > \tau_v] S_{\omega_k}^v. \quad (2)$$

The set of hypotheses  $\omega_k$  whose support  $S_{\omega_k}$  exceeds a threshold are kept and the resulting  $\mathbf{F}_k$  form the unorganized 3D curves.

Despite these advances, three major shortcomings remain: (i) some 3D curve fragments are correct for certain portions of the underlying curve and erroneous in other parts, due to multiview grouping inconsistencies; (ii) gaps in the 3D model, typically due to unreliable reconstructions near epipolar tangencies, where epipolar lines are nearly tangent to the curves; and (iii) multiple, redundant 3D structures. We now document each issue and describe our solutions.


**Problem 1. Erroneous grouping:** inconsistent multiview grouping of edges can lead to reconstructed curves which are veridical only along some portion, which are nevertheless wholly admitted, Fig. 3(a). Also, fully-incorrect hypotheses can accrue support coincidentally, as with repeated patterns or linear structures, Fig. 3(b). Both issues can be addressed by allowing for selective local reconstructions: only those portions of the curve receiving adequate edge support from sufficient views are reconstructed. This ensures that inconsistent 2D groupings do not produce spurious 3D reconstructions. The shift from cumulative global to multi-view local support results in greater selectivity and deals with coincidental alignment of edges with the reconstruction hypotheses.

**Problem 2. Gaps:** The geometric inaccuracy of curve segment reconstructions nearly parallel to epipolar lines led [10] to break off curves at epipolar tangencies, creating 2D gaps leading to gaps in 3D. We observe, however, that while reconstructions near epipolar tangency are *geometrically unreliable*, they are *topologically correct* in that they connect the reliable portions correctly but with highly inaccurate geometry. What is needed is to flag curve segments near epipolar tangency reconstructions as geometrically unreliable. We do this by the integration of support in Eq. 1, giving significantly lower weight to these unreliable portions instead of fully discarding them, which greatly reduces the presence of gaps in the resulting reconstruction.




**Fig. 3.** (a) Due to a lack of consistency in grouping of edges at the image level, a correct 3D curve reconstruction, shown here in blue, can be erroneously grouped with an erroneous reconstruction, shown here in red, leading to partially correct reconstructions. When such a 3D curve is projected in its entirety to a number of image views, we only expect the correct portion to gather sustained image evidence, which argues for a hypothesis verification method that can distinguish between supported segments and outlier segments; (b) An incorrect hypothesis can at times coincidentally gather an extremely high degree of support from a limited set of views. The red 3D line shown here might be an erroneous hypothesis, but because parallel linear structures are common in man-made environments, such an incorrect hypothesis often gathers coincidental strong support from a particular view or two. Our hypothesis verification approach is able to handle such cases by requiring explicit support from a minimum number of viewpoints simultaneously. (Color figure online)

**Problem 3. Redundancy:** A 2D curve can pair up with dozens of curves from other views, all pointing to the same reconstruction, leading to redundant pairwise reconstructions as partially overlapping 3D curve segments, each localized slightly differently. Our solution is to detect and reconcile redundant reconstructions. Since redundancy changes as one traverses a 3D curve, we reconcile redundancy at the local level: each 3D edge is in one-to-one correspondence with a 2D edge of its primary hypothesis view (*i.e.*, the first view from which it was reconstructed), hence 3D edges can be grouped in a one-to-one manner,



**Fig. 4.** (a) Redundant 3D curve reconstructions (orange, green and blue) can arise from a single 2D image curve in the primary hypothesis view. If the redundant curves are put in one-to-one correspondence and averaged, the resulting curve is shown in (b) in purple. Our robust averaging approach, on the other hand, is able to get rid of that bump by eliminating outlier segments, producing the purple curve shown in (c). (Color figure online)



**Fig. 5.** A visual comparison of: (left) the curve sketch results [10], with (right) the results of our enhanced curve sketch algorithm presented in Sect. 2. Notice the significant reduction in both outliers and duplicate reconstructions, without sacrificing coverage.

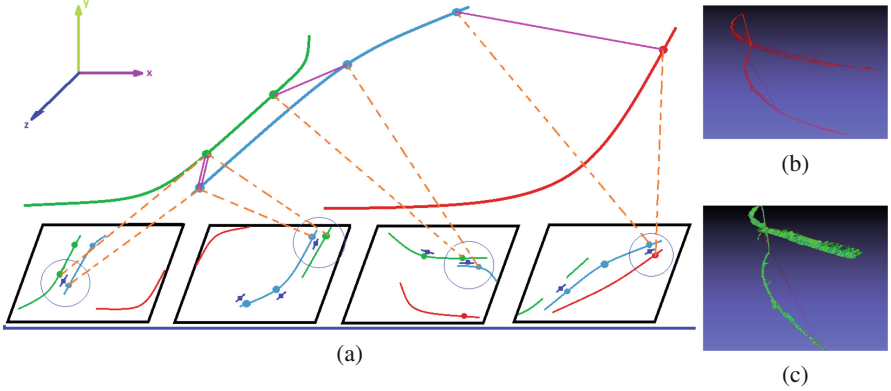
all corresponding to a common 3D source. These are robustly averaged by data-driven outlier removal, where a Gaussian distribution is fit on all pairwise distances between corresponding samples, discarding samples farther than  $2\sigma$  from the average, Fig. 4. Robust averaging improves localization accuracy, removes redundancy, and elongates shorter curve subsegments into longer 3D curves.

### 3 From 3D Curve Sketch to 3D Drawing

Despite the visible improvements of the Enhanced 3D Curve Sketch of Sect. 2, Fig. 5, curves are broken in many places, and there remains redundant overlap. The sketch representation as unorganized clouds of 3D curves are not able to capture the fine-level geometry or spatial organization of 3D curves, *e.g.* by using junction points to characterize proximity and neighborhood relations. The underlying cause of these issues is lack of integration across multiple views. The robust averaging approach of Sect. 2 is one step, anchored on one primary hypothesis view, but integrates evidence within that view only; a scene curve can be visible from multiple hypothesis view pairs, and some redundancy remains.

This lack of multiview integration is responsible for three problems observed in the enhanced curve sketch, Fig. 10: (i) localization inaccuracies, Fig. 10b, due to use of partial information; (ii) reconstruction redundancy, which leads to multiple curves with partial overlap, all arising from the same 3D structure, but remaining distinct, see Fig. 10c; (iii) excessive breaking because each curve segment arises from one curve in one initial view independently.

**Multiview Local Consistency Network:** The key idea underlying integration of reconstructions across views is the detection of a common image structure supporting two reconstruction hypotheses. Two 3D local curve segments depict the same single underlying 3D object feature if they are supported by the same 2D image edge structures. Since the identification of common image structure



**Fig. 6.** The four bottlenecks of Fig. 10 are resolved by integration of information/cues from all views. (a) The shared edge supporting edges, which are marked with circles, create the purple links between the corresponding samples of the 3D curves. These purple bonds will then be used to pull the redundant segments together and reorganize the 3D model into a clean 3D graph. Observe how the determination of common image support can identify portions of the green and blue curves as identical while differentiating the red one as distinct. A real example for a bundle of related curves is shown in (b) and the links among their edges in (c). (Color figure online)

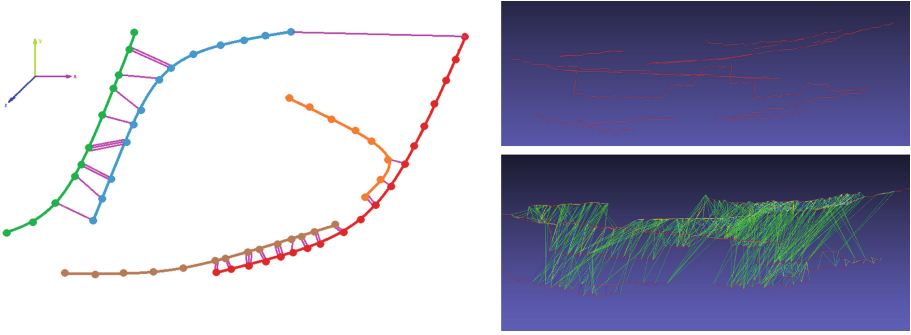
can vary along the curve, it must necessarily be a local process, operating at the level of a 3D local edge and not a 3D curve. Two 3D edge elements (edges) depict the same 3D structure if they receive support from the same 2D edges in a sufficient number of views, so 3D-2D links between a 2D edge to the 3D edge it supports must be kept. Typically, they share supporting image edges in many views; and the number of shared supporting edges is the measure of strength for a 3D-3D link between them.

Formally, we define the Multiview Local geometric consistency Network (MLN) as pointwise alignments  $\phi_{ij}$  between two 3D curves  $\Gamma_i$  and  $\Gamma_j$ : let  $\Gamma_i(s_i)$  and  $\Gamma_j(s_j)$  be two points in two 3D curves, and define

$$S_{ij} \doteq \{v : \gamma^{i,v}(s_i) \text{ and } \gamma^{j,v}(s_j) \text{ share local support}\}. \quad (3)$$

Then the a kernel function  $\phi$  defines a consistency link between these two points, weighted by the extent of multiview image support  $\phi_{ij}(s_i, s_j) \doteq |S_{ij}|$ . When the curves are sampled,  $\phi$  becomes an adjacency matrix of a graph representing links between individual curve samples. The implementation goes through each image edge which votes for a 3D curve point that has received support from it (see the supplementary material for details) (Fig. 6).

**Multiview Curve-Level Consistency Network:** The identification of 3D edges sharing 2D edges leads to high recall operating point with many false links due to accidental alignment of edge support. False positives can be reduced without affecting high recall by employing a notion of curve context for each 3D

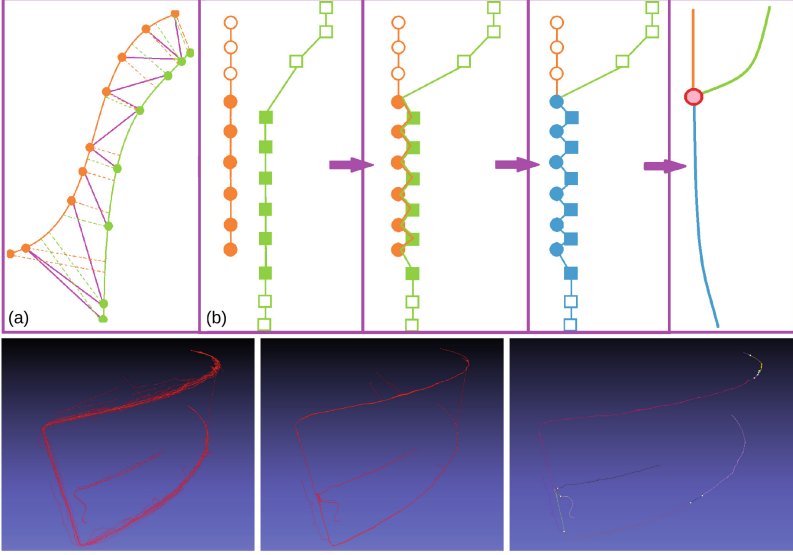


**Fig. 7.** The correspondence between 3D edge samples is skewed along a curve, which is a direct indication that these links cannot be used as-is when averaging and fusing redundant curve reconstructions. Instead, each point is assumed to be in correspondence with the point closest to it on another overlapping curve, during the iterative averaging step. Observe that corrections can be partial along related curves.

edgel: a link between two 3D edgels based on a supporting 2D edgel is more effective if the respective neighbors of the underlying 3D edge on the underlying 3D curve are also linked.

The curve context idea requires establishing new pairwise links between 3D curves using MLN, when there are a sufficient number of links with  $\phi_{ij} > \tau_\epsilon$  between their constituent 3D edges (in our implementation,  $\tau_\epsilon = 3$  and we require 5 such edges or more). The linking of 3D curves is represented by the Multiview Curve-level Consistency network (MCCN), a graph whose nodes are the 3D curves  $\mathbf{I}_j$  and the edges represent the presence of high-weight 3D edge links between these 3D curves. The MCCN graph allows for a clustering of 3D curves by finding connected components; and once a link is established between two curves, there is a high likelihood of their edges corresponding in a regularized fashion, thus fewer common supporting 2D edges are required to establish a link between all their constituent 3D edges. This fact is used to perform gap filling, since even no edge support is acceptable to fill in small gaps and create a continuous and regularized correspondence if both neighbors of the gap are connected (see pseudocode in Supplementary Materials for details). The two stages in tandem, *i.e.*, high recall linking of 3D edges and use of curve context to reduce false positives leads to high recall and high precision, *i.e.*, all the 3D edges which need to be related are related and very few outlier connections remain.

**Integrating Information Across Related Edges:** The identification of a bundle of curves as arising from the same 3D source implies that we can improve the geometric accuracy of this bundle by allowing them to converge to a common solution. While this might appear straightforward, 3D edges are not consistently distributed along related curves, yielding a skew in the correspondence of related samples, Fig. 7, sometimes not a one-to-one correspondence, Fig. 8a. This argues for averaging 3D curves and not 3D edge samples, which in turn requires finding



**Fig. 8.** (a) A schematic of sample correspondence along two related 3D curves, showing skewed correspondences that may not be one-to-one. (b) A sketch of how two curves are integrated. Bottom row: a real case.

a more regularized alignment between the 3D curves, without gaps; we find each curve sample’s closest point on the other curve.

When post averaging a sample with its closest points on related curves, the order of resulting averaged samples is not clear. The order should be inferred from the underlying curves, but this information can be conflicting, unless the distance between two curves is substantially smaller than the sampling distance along the curves. This requires first updating each curve’s geometry separately and iteratively, without merging curves until after convergence, Fig. 8d. This also improves the correspondence of samples at each iteration, as the closest points are continuously updated.

At each stage, the iterative averaging process simply replaces each 3D edge sample with the average of all closest points on curves related to it, Fig. 8b–d. This can be formulated as evolving all 3D curves by averaging along the MCCN using closest points. Formally, each  $\Gamma_i$  is evolved according to

$$\frac{\partial \Gamma_i}{\partial t}(s) = \alpha \underset{(\Gamma_i, \Gamma_j) \in \text{MCCN}}{\text{avg}}_{(i,j) \in L} \{ \Gamma_j(r) : \Gamma_j(r) = \text{cp}_j(\Gamma_i(s)) \}, \quad (4)$$

where  $\text{cp}_i(\mathbf{p})$  is the closest point in  $\Gamma_i$  to  $\mathbf{p}$  and  $L$  is the link set defined as follows: Let the set  $S_{ij}$  of so-called strong local links between curves  $\Gamma_i$  and  $\Gamma_j$  be

$$S_{ij} \doteq \{ (s, t) : \phi_{ij}(s, t) \geq \tau_\epsilon, \phi_{ij} \in \text{MLN}(\Gamma_1, \dots, \Gamma_K) \}. \quad (5)$$



Then the set  $L$  of the MCCN is defined as


$$L \doteq \{(i, j) : |S_{ij}| \geq \tau_{sl}\}. \quad (6)$$

In practice, the averaging is robust and  $\alpha$  is chosen such that in one step we move to the average.

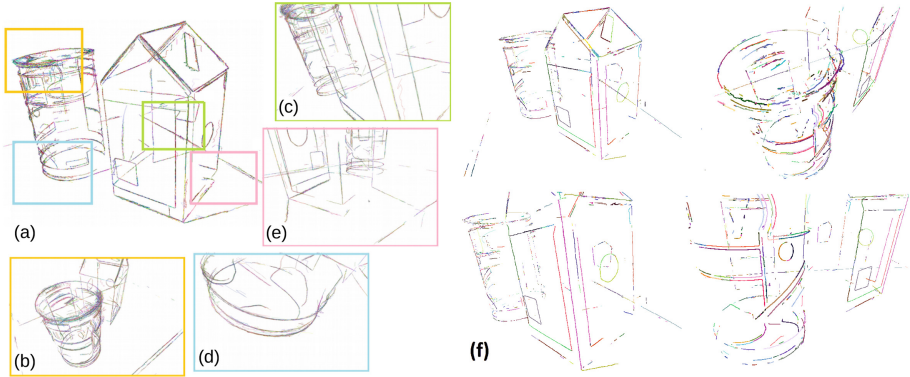
**3D Curve Drawing Graph:** Once all related curves have converged, they can be merged into single curves, separated by junctions where 3 or more curves meet. The order along the resulting curve is also dictated by closest points: The immediate neighbors of any averaged 3D edge are the two closest 3D edges to it among all converged 3D edges in a given MCCN cluster.

This where junctions naturally arise: as two distinct curves may merge along one portion they may diverge at one point, leaving two remaining, non-related subsegments behind, Fig. 8e. This is a *junction node* relating three or more curve segments, and its detection is done using the *merging primitives*, whose complete set are shown in Fig. 9. The intuition is this: a complex merging problem along the full length of two 3D curves actually consists of smaller, simpler and independent merging operations between different segments of each curve. A full merging problem between two complete curves can be expressed as a permutation of any number of simpler merging primitives. These primitives were worked out systematically to serve as the basic building blocks capable of constructing *all possible configurations* of our merging problem.

After iterative averaging, all resulting curves in any given cluster are processed in a pairwise fashion using these primitives: initialize the 3D graph with the longest curve in the cluster, and merge every curve in the cluster one by one into this graph. At each step, any number of these merging primitives arise and are handled appropriately. This process outputs the Multiview Curve Drawing Graph (MDG), which consists of multiple disconnected 3D graphs, one



**Fig. 9.** The complete set of merging primitives, which were systematically worked out to cover all possible merging topologies between a pair of curves whose overlap regions are calculated beforehand. We claim that any configuration of overlap between two curves can be broken down into a series of these primitives along the length of one of the curves. The 5th primitive is representative of a bridge situation, where the connection at either end of the yellow curve can be any one of the first four cases shown, and 6th primitive is representative of a situation where only one end of the yellow curve connects to multiple existing curves, but not necessarily just two. (Color figure online)




**Fig. 10.** (a) The four main issues with the enhanced curve sketch: (b) localization errors along the camera principal axis, which cause loss in accuracy if not corrected, (c) redundant reconstructions due to a lack of integration across different views, (d) the reconstruction of a single long curve as multiple, disconnected (but perhaps overlapping) short curve segments, and (e) the lack of connectivity among distinct 3D curves which naturally form junctions. (f) shows the 3D drawing reconstructed from this enhanced curve sketch, as described in Sect. 3. Observe how each of the four bottlenecks have been resolved. Additional results are evaluated visually and quantitatively, and are reported in Sect. 4 as well as Supplementary Materials.

for each 3D curve cluster in the MCCN. The nodes of each graph are the junctions (with curve endpoints) and the links are curve fragment geometries. This structure is the final 3D curve drawing.

## 4 Experiments and Evaluation

We have devised a number of large real and synthetic multiview datasets, available at [multiview-3d-drawing.sourceforge.net](http://multiview-3d-drawing.sourceforge.net).

**The Barcelona Pavilion Dataset:** a realistic synthetic dataset we created for validating the present approach with control over illumination, geometry and cameras. It consists of: (i) 3D models composing a large, mostly man-made, scene professionally composed by eMirage studios using the 3D modeling software Blender; (ii) ground-truth cameras fly-by’s around chairs with varied reflectance models and cluttered background; (iii) ground-truth videos realistically rendered with high quality ray tracing under 3 extreme illumination conditions (morning, afternoon, and night); (iv) ground-truth 3D *curve* geometry obtained by manually tracing over the meshes. This is the first synthetic 3D ground truth for evaluating multiview reconstruction algorithms that is realistically complex – most existing ground truth is obtained using either laser or structured light methods, both of which suffer from reconstruction inaccuracies and calibration errors. Starting from an existing 3D model ensures that our ground truth is not




**Fig. 11.** Our publicly-available synthetic (left and top-right) and real (bottom-right) 3D ground truths modeled and rendered using Blender for the present work.

polluted by any such errors, since both 3D model and the calibration parameters are obtained from the 3D modeling software, Fig. 11. The result is the first publicly available, high-precision 3D curve ground truth dataset to be used in the evaluation of curve-based multiview stereo algorithms. For the experiments reported in the main manuscript we use 25 views out of 100 from this dataset, evenly distributed around the primary objects of interest, namely the two chairs, see Fig. 11.

**The Vase Dataset:** constructed for this research from the DTU Point Feature Dataset with calibration and 3D ground truth from structured light [16, 35]. The images were taken using an automated robot arm from pre-calibrated positions and our test sequence was constructed using views from different illumination conditions to simulate varying illumination. To the best of our knowledge, these are the most exhaustive public multiview ground truth datasets. To generate ground-truth for curves, we have constructed a GUI based on Blender to manually remove all points of the ground-truth 3D point-cloud that correspond to homogeneous scene structures as observed when projected on all views, Fig. 11(bottom). What remains is a dense 3D point cloud ground truth where the points are restricted to be near abrupt intensity changes on the object, *i.e.* edges and curves. Our results on this real dataset showcase our algorithm’s robustness under varying illumination.

**The Amsterdam House Dataset:** 50 calibrated multiview images, also developed for this research, comprising a wide variety of object properties, including but not limited to smooth surfaces, shiny surfaces, specific close-curve geometries, text, texture, clutter and cast shadows, Fig. 1. The camera reprojection error obtained by Bundler [34] is on average subpixel. There is no ground truth 3D geometry for this dataset; the intent here is: to qualitatively test on a scene that is challenging to approaches that rely on, *e.g.*, point features; and to be able to closely inspect expected geometries and junction arising from simple, known shapes of scene objects.

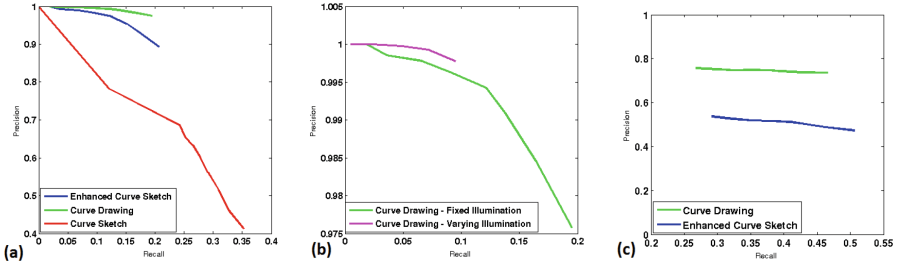


**Fig. 12.** The 3D drawing results on the Barcelona Pavilion, DTU Vase and Capitol Datasets. See Supplementary Materials for more extensive results and comparisons

**The Capitol High Building:** 256HD frames from a high 270° helicopter fly-by of the Rhode Island State Capitol [10]. Camera parameters are from the Matlab Calibration toolbox and tracking 30 points.

**Qualitative Evaluation:** The enhancements of Sect. 2 lead to significant improvements to the 3D curve sketch of [10] in increasing recall while maintaining precision. See Fig. 5 for a qualitative comparison. When the clean clouds of curves are organized into a set of connected 3D graphs, the results are more accurate, more visually pleasing and not redundant, Figs. 10(f) and 12. Each of the issues in Fig. 10(a–e) have been resolved and spatial organization of 3D curves have been captured as junctions, represented by small white spheres.

**Quantitative Evaluation:** Accuracy and coverage of 3D curve reconstructions is evaluated against ground truth. We compare 3 different results to quantify our improvements: (i) Original Curve Sketch [10] run exhaustively on all views, (ii) Enhanced Curve Sketch, Sect. 2, and (iii) Curve Drawing, Sect. 3. Edge maps are obtained using Third-Order Color Edge Detector [49], and are linked using



**Fig. 13.** Precision-recall curves for quantitative evaluation of 3D curve drawing algorithm: (a) Curve sketch, enhanced curve sketch and curve drawing results are compared on Barcelona Pavilion dataset with afternoon rendering, showing significant improvements in reconstruction quality; (b) A comparison of 3D curve drawing results on fixed and varying illumination version of Barcelona Pavilion dataset proves that 3D drawing quality does not get adversely affected by varying illumination; (c) 3D drawing improves reconstruction quality by a large margin in Vase dataset, which consists of images of a real object under slight illumination variation.

Symbolic Linker [14] to extract curve fragments for each view. Edge support thresholds are varied during reconstruction for each method, to obtain precision-recall curves. Here, **precision** is the percentage of accurately reconstructed curve samples: a ground truth curve sample is a true positive if its closer than a proximity threshold to the reconstructed 3D model. A reconstructed 3D sample is deemed a false positive if its not closer than  $\tau_{prox}$  to any ground truth curve. This method ensures that redundant reconstructions aren't rewarded multiple times. All remaining curve samples in the reconstruction are false positives. **Recall** is the fraction of ground truth curve samples covered by the reconstruction. A ground truth sample is marked as a false negative if its farther than  $\tau_{prox}$  to the test reconstruction. The precision-recall curves shown in Fig. 13 *quantitatively* measure the improvements of our algorithm and showcase its robustness under varying illumination.

## 5 Conclusion

We have presented a method to extract a 3D drawing as a graph of 3D curve fragments to represent a scene from a large number of multiview imagery. The 3D drawing is able to pick up contours of objects with homogeneous surfaces where feature and intensity based correlation methods fail. The 3D drawing can act as a scaffold to complement and assist existing feature and intensity based methods. Since image curves are generally invariant to image transformations such as illumination changes, the 3D drawing is stable under such changes. The approach does not require controlled acquisition, does not restrict the number of objects or object properties.

**Acknowledgements.** This research received support from FAPERJ/Brazil E25/2014/204167, UERJ/Brazil Prociencia 2014-2017, and NSF awards 1116140 and 1319914.

## References

1. Abuhashim, T., Sukkarieh, S.: Incorporating geometric information into gaussian process terrain models from monocular images. In: IEEE/RSJ IROS 2012, pp. 4162–4168 (2012)
2. Baatz, G., Saurer, O., Köser, K., Pollefeys, M.: Large scale visual geo-localization of images in mountainous terrain. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012. LNCS, vol. 7573, pp. 517–530. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33709-3\\_37](https://doi.org/10.1007/978-3-642-33709-3_37)
3. Calakli, F., Ulusoy, A.O., Restrepo, M.I., Taubin, G., Mundy, J.L.: High resolution surface reconstruction from multi-view aerial imagery. In: 3DIMPVT 2012, pp. 25–32 (2012)
4. Carlson, F.B., Vuong, N.D., Johansson, R.: Polynomial reconstruction of 3D sampled curves using auxiliary surface data. In: IEEE ICRA 2014, pp. 4411–4416 (2014)

5. Chen, T.-Y., Klette, R.: Animated non-photorealistic rendering in multiple styles. In: Huang, F., Sugimoto, A. (eds.) PSIVT 2013. LNCS, vol. 8334, pp. 12–23. Springer, Heidelberg (2014). doi:[10.1007/978-3-642-53926-8\\_2](https://doi.org/10.1007/978-3-642-53926-8_2)
6. Diskin, Y., Asari, V.: Dense point-cloud representation of a scene using monocular vision. *J. Electron. Imaging* **24**(2), 023003 (2015)
7. Fabbri, R.: Multiview differential geometry in application to computer vision. Ph.D. dissertation. Division of Engineering, Brown University (2010)
8. Fabbri, R., Kimia, B.B., Giblin, P.J.: Camera pose estimation using first-order curve differential geometry. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012. LNCS, vol. 7575, pp. 231–244. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33765-9\\_17](https://doi.org/10.1007/978-3-642-33765-9_17)
9. Fabbri, R., Kimia, B.B.: High-order differential geometry of curves for multiview reconstruction and matching. In: Rangarajan, A., Vemuri, B., Yuille, A.L. (eds.) EMCCVPR 2005. LNCS, vol. 3757, pp. 645–660. Springer, Heidelberg (2005). doi:[10.1007/11585978\\_42](https://doi.org/10.1007/11585978_42)
10. Fabbri, R., Kimia, B.B.: 3D curve sketch: flexible curve-based stereo reconstruction and calibration. In: CVPR (2010)
11. Fathi, H., Dai, F., Lourakis, M.: Automated as-built 3D reconstruction of civil infrastructure using computer vision. *Adv. Eng. Inform.* **29**, 149–161 (2015)
12. Feng, A., Shapiro, A., Ruizhe, W., Bolas, M., Medioni, G., Suma, E.: Rapid avatar capture and simulation using commodity depth sensors. In: SIGGRAPH 2014 (2014)
13. Fleming, R.W., Holtmann-Rice, D., Bülthoff, H.H.: Estimation of 3D shape from image orientations. *PNAS* **108**(51), 20438–20443 (2011)
14. Guo, Y., Kumar, N., Narayanan, M., Kimia, B.: A multi-stage approach to curve extraction. In: CVPR 2014 (2014)
15. Heinly, J., Schönberger, J.L., Dunn, E., Frahm, J.-M.: Reconstructing the world in six days. In: CVPR 2015 (2015)
16. Jensen, R., Dahl, A., Vogiatzis, G., Tola, E., Aanaes, H.: Large scale multi-view stereopsis evaluation. In: CVPR 2014 (2014)
17. Koenderink, J., van Doorn, A., Wagemans, J.: SFS? not likely.... *i-Perception* **4**(5), 299 (2013)
18. Kowdle, A., Batra, D., Chen, W.-C., Chen, T.: iModel: interactive co-segmentation for object of interest 3D modeling. In: Kutulakos, K.N. (ed.) ECCV 2010. LNCS, vol. 6554, pp. 211–224. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-35740-4\\_17](https://doi.org/10.1007/978-3-642-35740-4_17)
19. Kunsberg, B., Zucker, S.W.: Why shading matters along contours. In: Citti, G., Sarti, A. (eds.) Neuromathematics of Vision, pp. 107–129. Springer, Heidelberg (2014)
20. Lebeda, K., Hadfield, S., Bowden, R.: 2D or not 2D: bridging the gap between tracking and structure from motion. In: Cremers, D., Reid, I., Saito, H., Yang, M.-H. (eds.) ACCV 2014. LNCS, vol. 9006, pp. 642–658. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-16817-3\\_42](https://doi.org/10.1007/978-3-319-16817-3_42)
21. Litvinov, V., Yu, S., Lhuillier, M.: 2-manifold reconstruction from sparse visual features. In: IEEE IC3D 2012, pp. 1–8 (2012)
22. Mattingly, W.A., Chariker, J.H., Paris, R., Chang, D., Pani, J.R.: 3D modeling of branching structures for anatomical instruction. *J. Vis. Lang. Comput.* **29**, 54–62 (2015)
23. Pötsch, K., Pinz, A.: 3D geometric shape modeling by ‘3D contour cloud’ reconstruction from stereo videos. In: Computer Vision Winter Workshop, p. 99 (2011)



24. Rao, D., Chung, S.-J., Hutchinson, S.: CurveSLAM: an approach for vision-based navigation without point features. In: IEEE/RSJ IROS, pp. 4198–4204 (2012)
25. Restrepo, M., Ulusoy, A., Mundy, J.: Evaluation of feature-based 3-D registration of probabilistic volumetric scenes. *J. Photogrammetry Remote Sens.* **98**, 1–18 (2014)
26. Rico Espino, J.G., Gonzalez-Barbosa, J.-J., Gómez Loenzo, R.A., Córdova Esparza, D.M., Gonzalez-Barbosa, R.: Vision system for 3D reconstruction with telecentric lens. In: Carrasco-Ochoa, J.A., Martínez-Trinidad, J.F., Olvera López, J.A., Boyer, K.L. (eds.) *MCPR 2012. LNCS*, vol. 7329, pp. 127–136. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-31149-9\\_13](https://doi.org/10.1007/978-3-642-31149-9_13)
27. Shinozuka, Y., Saito, H.: Sharing 3D object with multiple clients via networks using vision-based 3D object tracking. In: *VRIC 2014*, pp. 34:1–34:4 (2014)
28. Simoes, F., Almeida, M., Pinheiro, M., dos Anjos, R.: Challenges in 3D reconstruction from images for difficult large-scale objects. In: *SVAR 2014*, pp. 74–83 (2014)
29. Teney, D., Piater, J.: Sampling-based multiview reconstruction without correspondences for 3D edges. In: *3DIMPVT 2012*, pp. 160–167 (2012)
30. Wang, R., Choi, J., Medioni, G.: 3D modeling from wide baseline range scans using contour coherence. In: *CVPR 2014* (2014)
31. Zhang, L.: Line primitives and their applications in geometric computer vision. Ph.D. thesis (2013)
32. Zia, M., Stark, M., Schindler, K.: Towards scene understanding with detailed 3D object representations. *IJCV* **112**(2), 188–203 (2015)
33. Zucker, S.: Stereo, shading, and surfaces: curvature constraints couple neural computations. *Proc. IEEE* **102**(5), 812–829 (2014)
34. Agarwal, S., Snavely, N., Simon, I., Seitz, S.M., Szeliski, R.: Building Rome in a day. In: *ICCV 2009* (2009)
35. Aanæs, H., Dahl, A.L., Pedersen, K.: Interesting interest points. *IJCV* **97**, 18–35 (2012)
36. Berthilsson, R., Åström, K., Heyden, A.: Reconstruction of general curves, using factorization and bundle adjustment. *IJCV* **41**(3), 171–182 (2001)
37. Cipolla, R., Giblin, P.: *Visual Motion of Curves and Surfaces*. Cambridge University Press, Cambridge (1999)
38. Cole, F., Sanik, K., DeCarlo, D., Finkelstein, A., Funkhouser, T., Rusinkiewicz, S., Singh, M.: How well do line drawings depict shape? In: *SIGGRAPH 2009* (2009)
39. Taubin, G., Crispell, D., Lanman, D., Sibley, P., Zhao, Y.: Shape from depth discontinuities. In: Nielsen, F. (ed.) *ETVC 2008. LNCS*, vol. 5416, pp. 216–237. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-00826-9\\_9](https://doi.org/10.1007/978-3-642-00826-9_9)
40. Furukawa, Y., Ponce, J.: Accurate, dense, and robust multi-view stereopsis. In: *CVPR 2007* (2007)
41. Goesele, M., Snavely, N., Curless, B., Hoppe, H., Seitz, S.: Multi-view stereo for community photo collections. In: *ICCV 2007* (2007)
42. Habbecke, M., Kobbelt, L.: A surface-growing approach to multi-view stereo reconstruction. In: *CVPR 2007* (2007)
43. Esteban, C.H., Schmitt, F.: Silhouette and stereo fusion for 3D object modeling. *CVIU* **96**(3), 367–392 (2004)
44. Leggitt, J.: *Drawing Shortcuts: Developing Quick Drawing Skills Using Today's Technology*. Wiley, New York (2015)
45. Liu, S., Kang, K., Tarel, J.-P., Cooper, D.B.: Freeform object reconstruction from silhouettes, occluding edges and texture edges. *PAMI* **30**, 131–146 (2007)
46. Moreels, P., Perona, P.: Evaluation of features detectors and descriptors based on 3D objects. *IJCV* **73**(3), 263–284 (2007)

47. Pollefeys, M., Gool, L.V., Vergauwen, M., Verbiest, F., Cornelis, K., Tops, J., Koch, R.: Visual modeling with a hand-held camera. *IJCV* **59**(3), 207–232 (2004)
48. Seitz, S., Curless, B., Diebel, J., Scharstein, D., Szeliski, R.: A comparison and evaluation of multi-view stereo reconstruction algorithms. In: *CVPR 2006* (2006)
49. Tamrakar, A., Kimia, B.B.: No grouping left behind: from edges to curve fragments. In: *ICCV 2007* (2007)
50. Yee, R.: *Architectural Drawing: A Visual Compendium of Types and Methods*. Wiley, New York (2012)