

# Multiview Reconstruction of Space Curves

Fredrik Kahl

Centre for Mathematical Sciences  
Lund University, Sweden  
fredrik@maths.lth.se

Jonas August

Robotics Institute  
Carnegie Mellon University, USA  
jonas@cs.cmu.edu

## Abstract

*Is the real problem in resolving correspondence using current stereo algorithms the lack of the “right” matching criterion? In studying the related task of reconstructing three-dimensional space curves from their projections in multiple views, we suggest that the problem is more basic: matching and reconstruction are coupled, and so reconstruction algorithms should exploit this rather than assuming that matching can be successfully performed before reconstruction. To realize this coupling, a generative model of curves is introduced which has two key components: (i) a prior distribution of general space curves and (ii) an image formation model which describes how 3D curves are projected onto the image plane. A novel aspect of the image formation model is that it uses an exact description of the gradient field of a piecewise constant image. Based on this forward model, a fully automatic algorithm for solving the inverse problem is developed for an arbitrary number of views. The resulting algorithm is robust to partial occlusion, deficiencies in image curve extraction and it does not rely on photometric information. The relative motion of the cameras is assumed to be given. Several experiments are carried out on various realistic scenarios. In particular, we focus on scenes where traditional correlation-based methods would fail.*

## 1. Introduction

The key difficulty of binocular stereo vision is the resolution of the correspondence problem: structures in the left image must be matched to those of the right to compute a disparity, and thus depth, for the related point in the three-dimensional scene. Numerous criteria such as image correlation and local orientation differences have been proposed to guide the matching process, but their success has been limited. For more than two images—multiview reconstruction—where wide-baseline geometries predominate, the weaknesses of matching criteria are even more apparent (Fig. 1). In this paper we focus on space curves to develop an inverse problem approach to multiview reconstruction. Instead of first seeking a correspondence of image structures and then computing 3D structure (a space curve), we do the opposite: we seek that space curve that is consistent with the image structures observed. Correspondence falls out as an easy consequence by reprojection of the sought-after space curve.



Figure 1: At first, the image curves of the steel-iron face mask may seem easy to automatically match and reconstruct. However, the mask has little texture and the background along a contour differs between images, interfering with correlation-based matching. In addition, many image curves, e.g., from the nose, do not exist in 3D and are due only to the projection, i.e., apparent contours [2, 6].

Not only do space curves provide a nontrivial problem domain for this study, they are also a rich source of information. While point features reveal little about surface topology, 3D curves provide such geometrical cues. The inference of 3D shape and scene geometry is perhaps the main application area [10, 18, 1]; others include object recognition [21], motion estimation [5] and determining whether a given image curve is the projection of a space curve or a surface contour [6, 20]. Curves with known correspondences in multiple views have been studied in e.g. [12].

This work draws upon continuing progress on the application of projective geometry to computer vision. For example, there now exist systems capable of computing both the motion of the camera and the structure of the scene from images alone, e.g. [11]. The images can be taken by uncalibrated cameras undergoing an unknown motion of an unknown scene. See [9] and the references therein. Developing these ideas beyond point sets and simple analytic curves motivates us to explore and analyze the multiview geometry of space curves.

### 1.1. Related work

As mentioned in §1, researchers in stereo vision and multiview reconstruction have employed many criteria for

matching structures in different images to obtain correspondence. The criteria include cross-correlation, sum-of-squared-differences, or sum-of-absolute-differences of image patches; an ordering constraint [17] which forces corresponding structures to have the same ordering along epipolar lines; a penalty on the orientation difference between corresponding tangents in different images; or a bound on the allowable disparity. While each of these criteria may occasionally be true, they are all generally false, as counter-examples demonstrate. For example, a line that is tilted in depth may appear vertical in one image but almost horizontal in another, violating the orientation difference constraint.

In [18], Robert and Faugeras present an edge-based trinocular stereo algorithm using geometric matching principles. They show that given the image curvature of corresponding curve points in two views, it is possible to predict the curvature in the third one and use that as a matching criterion. In [19], Schmid and Zisserman present an algorithm which can be viewed as an extension to [18]. In addition to geometric constraints, photometric information is used to disambiguate matches. Both these approaches apply a lot of heuristics and the algorithms do not generalize naturally to more than three views. As curve segments are first extracted from the images, they are dependent on that image edges are linked together in exactly the same way in different images. Hence, a space curve which is partly occluded in one view will necessarily be fragmented.

To escape these limitations, Alibhai and Zucker [1] have exploited the local differential structure of the space curve and its projection in the images, and have emphasized the importance of a helix aligned with the viewing direction, the so-called *z-helix*, as a local model of the unknown space curve. Although the *z-helix* is view-dependent, for small baseline stereo pairs it can be an effective structure for articulating the local consistency image curves via the projection of a quasi-single *z-helix* in each image plane. For general multiview imaging geometries, where the baseline can be arbitrarily large, the *z-helices* associated with different views differ greatly, making the very notion of consistency undefined, since there is no single space curve from which the image curves were projected. Our approach resolves this non-uniqueness issue by postulating a single smooth curve that induces the observed projections.

## 1.2. Overview of Our Approach

In order to avoid the multitude of matching criteria and *z-helices*, we propose a deceptively simple inverse problem approach: solve for that smooth space curve whose projections account for the observed image curves. Elaborating this idea requires definitions for smoothness, for the geometry of projection to a planar curve, and for what an “image of” such a planar curve actually is. In §2, we provide a brief review of the relevant 3D geometry that is used to describe the projec-



Figure 2: A tree with a lot of (self)-occluding branches. In addition, the background around the branches is different from view to view.

tion of the space curve to the plane. In §3.1, we articulate a contour smoothness prior using a 3D generalization of elastica [15] and relate it to a generative model of random space curves. In §3.2 we introduce an elaborate image formation model (likelihood function) that provides a definition of the “perfect” gradient field of an image with edges described by the projected space curve. Our goal is to compute a maximum a-posteriori (MAP) estimate of the unknown space curve.

Our approach allows for an arbitrary number of images, and can easily be altered to allow for variations in image corruption and imaging geometry. The resulting projected contours are defined to subpixel precision, but still allow for singular points such as cusps that occur when the smooth 3D curve contains the view direction. Further, it does not rely on the identical extraction of edge segments in different images; for example, portions of contour may be missing in some images, perhaps due to occlusion, see Fig. 2. We also note that because a complete probabilistic model is specified, fewer heuristics were needed during implementation than in previous attempts. Perhaps the most important aspect of this work is that the matching and reconstruction of curves are not separated—they are coupled and natural consequences of the model. Our experimental results (§5) show that the inverse problem approach is very attractive for multiview reconstruction.

## 2. Background

In this section we provide mathematical background for our camera model and for space curves.

### 2.1. Projective Geometry of Camera

We use the standard pinhole camera model. An object point  $\mathbf{X}$  in  $\mathbb{R}^3$  is projected to the image point  $\mathbf{x}$  in  $\mathbb{R}^2$  according to:

$$\lambda \begin{bmatrix} \mathbf{x} \\ 1 \end{bmatrix} = \underbrace{\begin{bmatrix} \gamma f & sf & x_0 \\ 0 & f & y_0 \\ 0 & 0 & 1 \end{bmatrix}}_P [R \mid -RT] \begin{bmatrix} \mathbf{X} \\ 1 \end{bmatrix}, \quad (1)$$

where  $f$  is the focal length,  $\gamma$  is the aspect ratio,  $s$  is the skew,  $(x_0, y_0)$  is the principal point, and rotation matrix  $R$  and translation vector  $\mathbf{T}$  determine the rigid motion between the camera and object coordinate systems. All the camera-dependent parameters are grouped together into the camera matrix  $P$ . The scalar  $\lambda$  accounts for depth.

Given a sequence of images, there are systems that automatically can calculate the corresponding camera matrices  $P_i$  for  $i = 1, \dots, N$  where  $N$  is the number of images. Such systems depend on several components, including feature extraction and tracking (in general, using point features), the finding of initial solutions and their refinement [9]. In addition to the estimated camera matrices, a sparse set of 3D point features is also obtained. The refinement of the solution (usually based on the maximum likelihood principle) is often referred to as **bundle adjustment**. Typically, the re-projection errors of the 3D points are around a pixel. These errors are due to both noise in the measurements and imperfections of the camera model.

## 2.2. Space Curves

A space curve can be defined as a mapping:

$$\mathbf{C} : \Omega = [0, l] \longrightarrow \mathbb{R}^3 \quad \text{such that} \quad s \mapsto \mathbf{C}(s). \quad (2)$$

We assume that the curves are parameterized by arclength, so the scalar  $l$  is the length of the curve.

## 3. Inverse Problem Formulation

Suppose we are given several images and their corresponding pinhole camera models. For each camera, a curve in 3D space  $\mathbf{C}(s)$  is projected to a curve  $\mathbf{c}(s)$  in the image plane. We are interested in solving the inverse problem, i.e. given the image curves  $\mathbf{c}(s)$  estimate the space curve  $\mathbf{C}(s)$ .

Inspired by the success of active contour models [13, 7], we formulate this as an energy minimizing problem. Let  $Ad$  be a space of admissible deformations of 3D curves and  $E$  an energy functional of the form:

$$\begin{aligned} E : Ad &\longrightarrow \mathbb{R} \\ \mathbf{C} &\mapsto E(\mathbf{C}) = E_{prior}(\mathbf{C}) + E_{image}(\mathbf{C}). \end{aligned} \quad (3)$$

$E_{prior}$  is the potential associated to our prior knowledge of space curves and  $E_{image}$  is associated to the image formation model. There are several aspects of this particular problem which makes it more complicated than traditional deformable models:

- We do not know the corresponding image curves a priori, and hence no initialization is given. The correspondences should be obtained automatically as a part of the estimation process.
- A space curve may be (partly or completely) occluded in some of the projections.

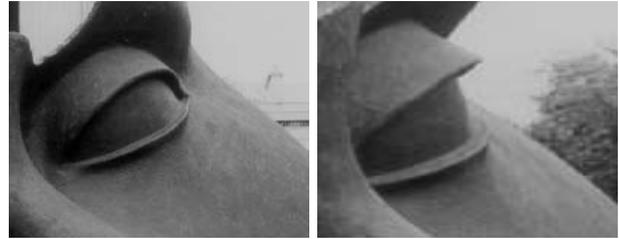


Figure 3: Even though a space curve for the upper eye lid is smooth, its projections may (left) or may not (right) be. Observe how the upper eye lid contour may appear smooth (left) or sharp (right, where projected lid contour has high curvature, approximately a cusp), depending on viewpoint.

- The imperfections of the estimated camera matrices are often quite strong. This is both due to noise and the approximative model of a true camera. See §5.1.
- An apparent image curve is often not the projection of a space curve. For example, occluding contours and specularities cause such “false” image curves.

### 3.1. Prior Model of Contour Smoothness

A common assumption for the inference of contours in the image plane is smoothness, which is often formulated through a penalty on large curvatures [15, 22]. Here the prior is instead on the space curve, for which this we select the three-dimensional form of smoothness captured by

$$E_{prior}(\mathbf{C}) = \frac{1}{\sigma_{prior}^2} \int_{\Omega} \kappa_{\mathbf{C}}(s)^2 ds,$$

where  $\kappa_{\mathbf{C}}(s)$  is the curvature of  $\mathbf{C}(s)$ . While smooth space curves frequently project to smooth planar curves, occasionally they do not, such as at a cusp or corner in the projection (Fig. 3). In a later paper we shall develop an embellishment to this model where contour torsion is also penalized, in order to ensure the resulting curves are approximately locally planar.

A probabilistic interpretation of this prior may be obtained by constructing the probability density  $p_{prior}(\mathbf{C}) = Z^{-1} \exp[-E_{prior}(\mathbf{C})]$ , where  $Z$  is a normalizing constant. This describes a Markov process model of the unknown contour where curvature is white-noise. Intuitively, this model describes the motion of a particle which is moving forward at constant speed, but whose direction is slightly perturbed at each time instant. Although minimizers of  $E_{prior}$  are spline-like curves, the connection to probability emphasizes the model’s flexibility: typical realizations from  $p_{prior}$  will have low  $E_{prior}$  and an approximate bound on curvature, but are otherwise unconstrained.

## 3.2 Generative Edge Model

Now we introduce a generative model for image gradients as a means of deriving a form for  $E_{image}$  in terms of the unknown curve. Let  $\mathbf{c}(s)$  be the image projection of the curve  $\mathbf{C}(s)$  according to pinhole model. Imagine that the image curve  $\mathbf{c}$  describes the border between two regions having constant brightness in the image. We thus suppose a model of image formation where the projected image  $I_{step}(\mathbf{x})$  is piecewise constant (a two-tone “step” image), where  $\mathbf{c}(s)$  is the border. Now, the gradient of this image will have zero value except along the contour, where it will be proportional to  $b \mathbf{n}(s)$ , where  $b$  is the brightness difference between object and background and  $\mathbf{n}(s)$  is the normal to the actual projected curve. In other words, the ideal gradient image is an extended impulse along the border (since there the image is discontinuous), where the impulse has a local vector weighting of  $b\mathbf{n}(s)$ :

$$\nabla I_{step}(\mathbf{x}) = \int_{\Omega} \delta(\mathbf{x} - \mathbf{c}(s)) b \mathbf{n}(s) ds. \quad (4)$$

The above definition is closely related to the curve indicator random field [4], except here the indicator is vector-weighted by  $b\mathbf{n}$ . The actually observed gradient image  $\nabla I$  is imperfect: sensor noise and quantization give rise to error. Here we assume the noise in the gradient is additive. Thus our observation model of the gradient image is

$$\nabla I(\mathbf{x}) = \nabla I_{step}(\mathbf{x}) + noise. \quad (5)$$

Assuming that the noise is 0-mean, independent and identically Gaussian distributed, the negative log-likelihood for  $\nabla I_{step}$  is therefore  $\frac{1}{2\sigma_{image}^2} \|\nabla I_{step} - \nabla I\|^2 + constant$ , where  $\sigma_{image}^2$  is the noise variance,  $\|f\|^2 = \langle f, f \rangle$  is a norm, and the inner product of vector fields  $f, g : \mathbb{R}^2 \rightarrow \mathbb{R}^2$  is  $\langle f, g \rangle = \int f(\mathbf{x}) \cdot g(\mathbf{x}) d\mathbf{x}$ , with  $a_1 \cdot a_2$  the dot product between vectors  $a_1, a_2 \in \mathbb{R}^2$ . Ignoring the constants and expanding the square we get  $\|\nabla I_{step}\|^2 - 2\langle \nabla I_{step}, \nabla I \rangle + \|\nabla I\|^2$ . The last term can be ignored for the purposes of finding the minimal energy for  $\mathbf{c}$ , but the other terms depend on  $\mathbf{c}$ . The first term is a penalty on the self-crossings of the curve [3], but we shall ignore this and focus on the middle term. Calculating by changing order of integration and using the sifting property of Dirac  $\delta$ -functions, we get

$$\begin{aligned} \langle \nabla I_{step}, \nabla I \rangle &= \int \left[ \int \delta(\mathbf{x} - \mathbf{c}(s)) b \mathbf{n}(s) ds \right] \cdot \nabla I(\mathbf{x}) d\mathbf{x} \\ &= \int b \left[ \int \delta(\mathbf{x} - \mathbf{c}(s)) \nabla I(\mathbf{x}) d\mathbf{x} \right] \cdot \mathbf{n}(s) ds \\ &= \int b \nabla I(\mathbf{c}(s)) \cdot \mathbf{n}(s) ds \end{aligned} \quad (6)$$

In summary, our model of image energy is

$$E_{image}(\mathbf{C}) = -\frac{b}{\sigma_{image}^2} \int_{\Omega} \nabla I(\mathbf{c}(s)) \cdot \mathbf{n}(s) ds, \quad (7)$$

which encourages the image gradient to align with the normal  $\mathbf{n}$  to the actual curve  $\mathbf{c}$ . Although this model may be generalized beyond step images, it naturally exploits both the magnitude and direction of image gradients.

## 3.3. Curves vs. Non-Curves

There are some things which are not modeled in the described approach, e.g. occlusion. And not all image curves can be explained by projections of space curves, e.g. apparent contours [2, 6]. So, even if we have a minimizing curve for the potential  $E$  (with appropriate boundary conditions), we cannot be sure that it truly corresponds to a space curve.

To determine whether a minimizing curve should be regarded as a space curve or not, we introduce some heuristics. We will require that a minimizer of  $E$  should have (i) image support (to be defined) in at least  $N_{support}$  images and (ii) a total projected curve length (in supported images) which is  $\geq L_{total}$  pixels. By *image support*, we mean that the projection of the minimizing curve has (i) on the average at least  $b_{min}$  units of brightness difference (i.e. the mean image gradient in the normal direction is  $\geq b_{min}$ ) and (ii) an image curve length of at least  $L_{image}$  pixels.

In addition, there is a global constraint that should be satisfied, namely that two space curves cannot project to the same image curve.

## 4. Estimation of Space Curves

In this section we present an algorithm that estimates space curves according to our model. Hence, the algorithm should produce minimizers of the energy potential in (3) having enough support in the images.

### 4.1. Parameterization

One way to minimize  $E$  would be to derive the Euler-Lagrange equations and then try to solve them numerically. Instead of doing that we choose to explicitly model the curves by cubic B-splines and hence turning the infinite dimensional problem into a finite dimensional one.

A cubic B-spline [8] is specified by  $N$  control points and comprises  $N - 3$  cubic polynomial curve segments. Using this parameterization, the curve is guaranteed to be continuous and it is easy to compute both first and second order derivatives. Arc-length parameterization is achieved by reparameterization using linear interpolation. As the problem is formulated independently of the actual parameterization, other choices are also possible.

### 4.2. Generation of Initial Hypotheses

In order to start the minimization process, an initial hypothesis of a curve segment is needed. In principle, any technique is fine, for instance [18]. We apply the following technique which is adopted to our model and multiple views.

First, image edge points are extracted using the Canny edge detector. Edgels are then linked into chains, jumping up to one pixel gap. The extracted image curves are never perfect in practice: there will be missing segments, erroneous additional segments, etc. Therefore, to improve robustness to these imperfections we begin with one such segment in one single image and seek confirming evidence in the other images. Empirically we have found that a segment of 20 edge points is a good starting point. Given the depth for a segment, one can calculate the image support for that instance in the other images. If there is support in  $N_{support}$  images according to the criteria in § 3.3 for some positive depth, a hypothesis is generated. At this early stage of the computations,  $b_{min}$  as well as  $L_{image}$  are halved compared to the values used in the final optimization. For each hypothesis, the endpoints of the curve will be used as boundary conditions. An estimate of the expected brightness difference  $b$  is also stored. In the experiments, we have limited the amount of hypotheses generated due to computational and practical reasons. Only edge segments detected in the first image are used to initiate the above procedure.

### 4.3. Algorithm

The hypotheses are just small segments, while we are interested in reconstructing complete minimizing curves (obeying the image support condition). So, we need to extend the boundary conditions of our initial segments.

The following scheme produces minimizing energy curves with the properties described in §3.

1. Minimize the potential  $E$  in (3) over all control points (using a gradient method).
2. Check for image support. If not supported, then stop.
3. Extend the curve (both in terms of boundaries and number of control points). Goto step 1.

The scheme has one weakness which is that each curve is optimized independently and hence the global constraint is not enforced. However, a final postprocessing step remedies this by making sure that each image curve is the projection of at most one space curve. If there are multiple responses, the one with strongest image support is kept. In practice, this reduces the number of mismatches.

## 5. Experimental Validation

**Setup.** There are a couple of parameters that need to be set depending on our prior belief of space curves, expected variance of brightness differences, etc. The estimates for these are based on the first experiment below, by analyzing and learning the characteristics of the image curves:  $\sigma_{prior} = 0.3$  and  $\sigma_{image} = 1$ . All image derivatives are computed using the derivative of a Gaussian kernel with  $\sigma = 2$ .

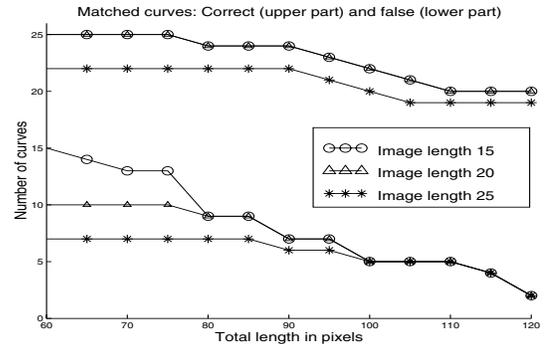


Figure 4: Number of correct and false matches for face statue as a function of total and single image contour length.

When a faulty space curve is reconstructed, what usually happens is that the space curve does not project to corresponding image curves. We say that a curve has been *mismatched* when this happens. The assessment of correctly matched and mismatched curves has been done by hand.

The remaining parameters concern the image support (§3.3). We require that a curve has support in at least  $N_{support} = 3$  images and the minimum image brightness difference required is  $b_{min} = 8$ . In Fig. 4, the number of correct matches and the number of mismatches are graphed depending on two parameters. The first one is the minimum image curve length  $L_{image}$  and the other one is the total image curve length  $L_{total}$ . If we allow very short curves, there will be a lot mismatches. On the other hand, if we are too strict, correct edges will be discarded. A good trade-off can be found with  $L_{image} = 20$  pixels and  $L_{total} = 100$  pixels from the graph. All these parameters are kept fixed for the remaining experiments (if not otherwise stated).

**The face statue.** In total, six images ( $546 \times 884$  pixels) were taken of a black iron statue of a face. Two images are shown in Fig. 5 together with the reconstructed curves (in total 22 correct and 5 mismatches from 331 hypotheses). The lighting variations are quite strong between different images and there is very little texture on the statue. Both apparent image contours and occlusions are frequent.

Notice in the left image how the contour curve on the right side of the face is not affected by the occlusion of the nose. The reason that there are no curves detected on the upper right and the lower left side of the face is that in the majority of the images the image curves correspond to occluding contours. The erroneous image curves can be found in the background and they all have support in just three images.

Recall the projection of the smooth curve forming a corner point in Fig. 3. The final result of that particular space curve projected in the same images are given in Fig. 6.

**Flower on a chair.** This sequence of five images of sizes  $1000 \times 1280$  pixels (Fig. 7) presents a challenge mainly for two objects which we will focus on: (i) the arm rests and (ii) the flower. The cushion of the chair has “texture edges”, but

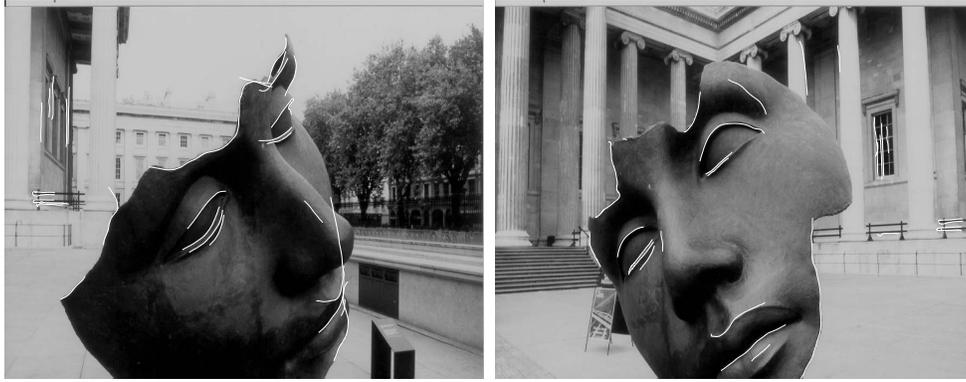


Figure 5: Two images of the steel-iron face statue with reprojected space curves.

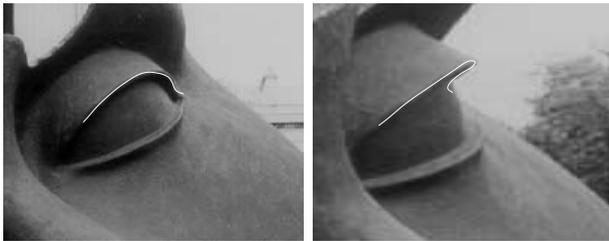


Figure 6: The projected space curve of the left eye lid of the iron mask.

they are not supported by our model. The arm rests (both left and right ones) are textureless and they are partly (self)-occluded. Yet, most (strong) edges are reconstructed well. The (green) leaf on the right hand side, however, lacks image curves. The probable reason is that the background changes from dark to light, and since the leaf's colour can be considered to be something in between, it is not covered by our model.

Out of the 90 reconstructed space curves (from 252 hypotheses), all of them were assessed to be correctly matched. The baselines between the five cameras are rather short, and therefore the contour of the flower pot is considered to be a space curve rather than a surface contour.

**Tree.** Now recall the image of a tree in a snowy courtyard (Fig. 2). In Fig. 8, two out of the total of five images ( $1000 \times 1280$  pixels) are plotted together with the projections of the estimated space curves. Here, we require that each space curve is supported in at least  $N_{support} = 4$  images. There are 58 space curves reconstructed (out of 87 initial hypotheses). The assessment of the correctness of the reprojected curves is not an easy task for a human. Even after a careful inspection, we have not found a single mismatch.

A three-dimensional model of the space curves and the camera positions are shown in Fig. 9. One part which is missing in the reconstruction is the (lower) tree trunk. The most probable reason is that the images of the trunk can be better characterized as apparent contours than images of a true 3D

curve. The smaller branches, however, agree with our curve model.

**Corridor.** A final experiment has been conducted on the well-known corridor sequence, see Fig. 10 for results with  $N_{support} = 4$  images. In this case, the use of cross-correlation is sufficient for matching [19]. The sequence has in total 11 images (sizes  $512 \times 512$  pixels) in a forward moving direction. Even though there are no mismatches among the 58 curves, the results are not impressive for two reasons: (i) There are dominant curves missing which are present in the majority of the images, and (ii) the visual quality of the curves is not good. Often they fail to stay closely to the edges. In light of the previous experiments, these results come as a surprise. After having analyzed the outcome in more detail, we have determined that the cameras were badly calibrated. In certain parts of the images, the projections of a space curve need to be displaced a few pixels in order to match to the actual image curves. A possible remedy to this is presented in the end of the next subsection.

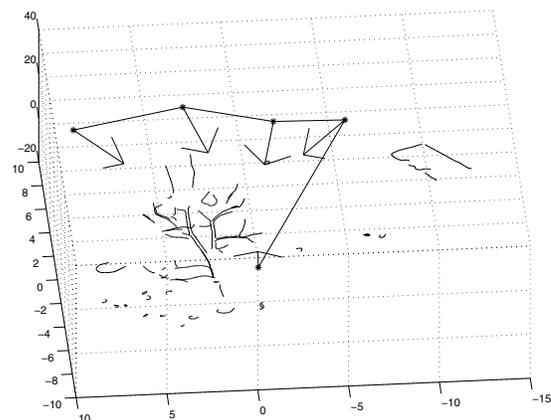


Figure 9: 3D model of reconstructed tree curves. The camera trajectory is also plotted. The arrows indicate the cameras' viewing directions.



Figure 7: Reprojection results for two out of five images of “flower on a chair”.

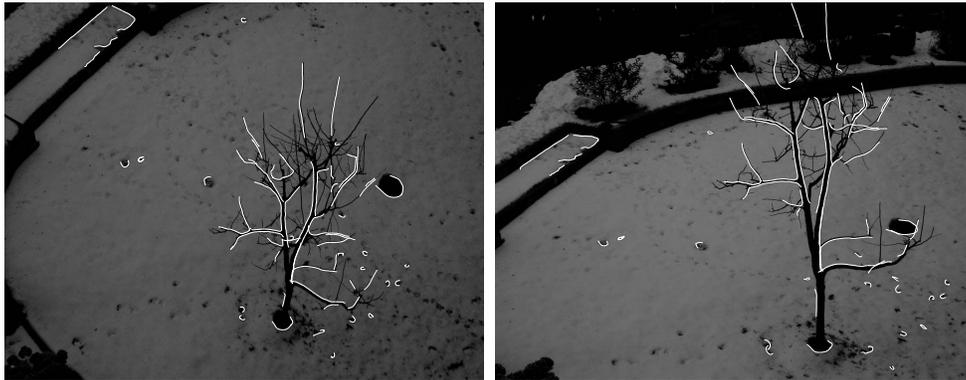


Figure 8: Curve reprojections and images of a tree on a snowy courtyard.

## 5.1. Limitations and Extensions

We note that our approach has a number of limitations that bear consideration. First, our edge model only linearly depends on the image gradient, and thus differs from commonly used image energy terms that depend only on gradient magnitude. Including such nonlinearities maybe fruitful to improve performance, particularly for non-Gaussian noise. Secondly, our current method for distinguishing curves from non-curves and dealing with occlusion amounts to the application of a number of rules (curve lengths and visibility). A more complete probabilistic approach would have an explicit distribution over the number of curves, their length, etc. Thirdly, the (current) computational burden of this method is significant. Fortunately, due to the Markovianity of the smoothness assumption, dynamic programming may be applicable.

Finally, each camera matrix  $P$  is an estimate according to the pinhole model, which in turn, is only an approximation of the true underlying camera. As a result the image curve  $\mathbf{c}$  will be slightly displaced. Specifically, we must include the effect of an unknown geometric distortion vector field  $\mathbf{w}(x, y)$  [9]. Since we are only concerned with the actual projected contour, we need only evaluate the ge-

ometric distortion along it, resulting in the perturbed contour  $\mathbf{c}(s) + \mathbf{w}(\mathbf{c}(s))$ . Now,  $\mathbf{w}(\mathbf{c}(s))$  can be decomposed into a normal and tangential components. Ignoring end effects, the tangential component can be ignored, as it amounts to a reparameterization of  $\mathbf{c}(s)$  [14, 16], and the normal component  $d(s)$  is called the *offset*. The actual image curve  $\mathbf{c}'(s)$  therefore has the form  $\mathbf{c}'(s) = \mathbf{c}(s) + d(s)\mathbf{n}(s)$ , where  $\mathbf{n}(s)$  is the normal of  $\mathbf{c}(s)$ . The offset  $d(s)$  is itself a one-dimensional curve  $d : \Omega \rightarrow \mathbb{R}$  which adjusts the projected curve in the normal direction. Since the geometric distortion  $\mathbf{w}$  is typically small and slowly varying, so too will be  $d(s)$ , and therefore one may include the offset energy  $E_{\text{offset}}(d) = \int_{\Omega} d(s)^2 + d'(s)^2 ds$  to penalize the magnitude and the derivative of the offset. The relative importance of this strong prior knowledge of  $d(s)$  is thus related to the estimated accuracy of  $P$ .

## 6. Conclusion

We have presented a completely automatic approach for matching and reconstructing space curves in multiple views. Our method solves an inverse problem for a generative model, and hence the resulting reconstructions are conse-



Figure 10: Reprojected space curves in the corridor sequence.

quences of the model. In a number of challenging experiments, the performance of the algorithm has been demonstrated.

## References

- [1] S. Alibhai and S.W. Zucker. Contour-based correspondence for stereo. In *European Conf. Computer Vision*, volume I, pages 314–330, Dublin, Ireland, 2000.
- [2] K. Åström and F. Kahl. Motion estimation in image sequences using the deformation of apparent contours. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 21(2):114–127, 1999.
- [3] Jonas August. *The Curve Indicator Random Field*. PhD thesis, Yale University, 2001.
- [4] Jonas August and Steven W. Zucker. Sketches with curvature: The curve indicator random field and markov processes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 25(4):387–400, 2003.
- [5] R. Berthilsson and K. Åström. Reconstruction of 3d-curves from 2d-images using affine shape methods for curves. In *Conf. Computer Vision and Pattern Recognition*, pages 476–481, San Juan, Puerto Rico, 1997.
- [6] R. Cipolla and A. Blake. Surface shape from the deformation of apparent contours. *Int. Journal Computer Vision*, 9(2):83–112, 1992.
- [7] L. D. Cohen. On active contour models and balloons. *Computer Vision, Graphics and Image Processing*, 53(2):211–218, 1991.
- [8] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. *Computer Graphics, Principles and Practice*. Addison-Wesley Publishing Company, 1990.
- [9] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.
- [10] F. Kahl and A. Heyden. Affine structure and motion from points, lines and conics. *Int. Journal Computer Vision*, 33(3):163–180, 1999.
- [11] F. Kahl and A. Heyden. Euclidean reconstruction and auto-calibration from continuous motion. In *Int. Conf. Computer Vision*, volume II, pages 572–577, Vancouver, Canada, 2001.
- [12] J. Y. Kaminski, M. Fryers, A. Shashua, and M. Teicher. Multiple view geometry of non-planar algebraic curves. In *Int. Conf. Computer Vision*, volume II, pages 181–186, Vancouver, Canada, 2001.
- [13] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. Journal Computer Vision*, 1(4):321–331, 1987.
- [14] Benjamin Kimia. *Toward a computational theory of shape*. PhD thesis, McGill University, 1991.
- [15] David Mumford. *Algebraic Geometry and Its Applications*, chapter *Elastica and Computer Vision*, pages 491–506. Springer-Verlag, 1994.
- [16] S. Osher and J. Sethian. Fronts propagating with curvature dependent speed: algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988.
- [17] S.B. Pollard, J.E.W Mayhew, and J.P. Frisby. 449–470. *Perception*, 14:45–50, 1985.
- [18] L. Robert and O.D. Faugeras. Curve-based stereo: figural continuity and curvature. In *Conf. Computer Vision and Pattern Recognition*, pages 57–62, Maui, Hawaii, USA, 1991.
- [19] C. Schmid and A. Zisserman. The geometry and matching of lines and curves over multiple views. *Int. Journal Computer Vision*, 40(3):199–233, 2000.
- [20] R. Vaillant and O. D. Faugeras. Using extremal boundaries for 3D object modelling. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(2):157–173, 1992.
- [21] B. Vijayakumar, D. Kriegman, and J. Ponce. Invariant-based recognition of complex 3D curved objects from image contours. In *Int. Conf. Computer Vision*, pages 508–514, Cambridge Ma, USA, 1995. IEEE Computer Society Press.
- [22] L.R. Williams and D.W. Jacobs. Stochastic completion fields: A neural model of illusory contour shape and salience. *Neural Computation*, 9(4):837–858, 1997.