

Correspondenceless Structure from Motion

Ameesh Makadia and Kostas Daniilidis*

University of Pennsylvania

Philadelphia, PA 19104

{makadia, kostas}@cis.upenn.edu

Abstract

We present a novel approach for the estimation of 3D-motion directly from two images using the Radon transform. The feasibility of any camera motion is computed by integrating over all feature pairs that satisfy the epipolar constraint. This integration is equivalent to taking the inner product of a similarity function on feature pairs with a Dirac function embedding the epipolar constraint. The maxima in this five dimensional motion space will correspond to compatible rigid motions. The main novelty is in the realization that the Radon transform is a filtering operator: If we assume that the similarity and Dirac functions are defined on spheres and the epipolar constraint is a group action of rotations on spheres, then the Radon transform is a correlation integral. We propose a new algorithm to compute this integral from the spherical harmonics of the similarity and Dirac functions. Generating the similarity function now becomes a preprocessing step which reduces the complexity of the Radon computation by a factor equal to the number of feature pairs processed. The strength of the algorithm is in avoiding a commitment to correspondences, thus being robust to erroneous feature detection, outliers, and multiple motions.

*The authors are grateful for support through the following grants: NSF-IIS-0083209, NSF-IIS-0121293, NSF-EIA-0324977, NSF-CNS-0423891, NSF-IIS-0431070, and ARO/MURI DAAD19-02-1-0383.

1 Introduction

Estimation of 3D-motion from two calibrated views has been exhaustively studied in the case where optical flow or feature correspondences are given and the scene is rigid. Algorithms working over multiple frames yield high-quality motion trajectories and reconstructions when feature matches are cleaned through outlier rejection and motions independent of the camera are excluded. These outlier rejection and segmentation steps are subject to the fundamental coupling of data association and estimation: if we knew the motion estimate, data association would be trivial; if we knew the data association, motion estimation would be easier. Resistance to outliers and independent motions pose severe practical limitations to the wide application of structure from motion as a navigation tool, visual GPS, or a camera tracker.

In this paper, we propose a novel approach for structure from motion applicable in the presence of large motions and many irrelevant features resulting from reduced overlap of the fields of view. Our approach is based on the naive principle that an exhaustive search over all possible correspondence configurations for all motion hypotheses would yield all 3D-motions compatible with these two views. Such a search is intractable when we use a large field of view in an arbitrary, possibly unstructured environment with thousands of features.

The contribution of this paper is in the re-formulation of this Hough-reminiscent approach as a filtering problem: Assuming a similarity function between any two features in the first and second view, we convolve this function with a kernel that checks the compatibility of a correspondence pair with the epipolar constraint for a given motion hypothesis. The resulting integral is a Radon transform known from computer tomography where a material density is integrated over a ray path. In our case, this path is the subset of the cross product of all features that satisfies the epipolar constraint.

The question is: Can we efficiently compute this integral avoiding the combinatorially infeasible summation over all correspondences compatible with the epipolar constraint? The answer is yes, because this is a convolution integral and we can compute it through multiplication in the Fourier domain. The final motion space is obtained through a five dimensional inverse rotational Fourier transform on the motion parameters. An exhaustive search finds the maxima corresponding to rigid motions. The number of spherical Fourier coefficients preserved determines the resolution of the motion space. Obviously, the approach can work on arbitrarily large motions.

We present a complete end-to-end system, from images to motion parameters

where the only tuning parameter is the coupled resolution of the image and the motion space. We extract hundreds of SIFT features [15] for which we define their similarity function proportional to the Euclidean norm of the attribute vectors and we compute the spherical harmonics of the similarity function as the input to the correlation integral. In the experiments, we use as input hemispherical omnidirectional images. A projective plane can always be mapped to the sphere and the field of view has to be large for any structure from motion algorithm to succeed [21, 4]. The results on real sequences are compared to a robust estimation of the Essential Matrix using RANSAC. Before continuing with the related work we summarize the main contributions of this paper:

- We propose a new integral transform that maps a similarity function between two calibrated images to the strength of a motion hypothesis without assuming any correspondences.
- We show that this Radon/Hough transform can be written as a convolution/correlation integral which can be computed from the spherical harmonic coefficients of the image similarity function much faster than computing directly the Hough transform.

The inspiring idea of this work has first been drafted in [11] where a Hough transform is computed on the essential manifold. A short version of the current paper has appeared in [16]. In this paper, we will present a complete theoretical and experimental treatment of our approach. In the next subsection we will discuss related approaches. Then we will motivate the Radon transform by explaining how the well-known Hough line detection can be written as a Radon integral [5]. In section 2 we elaborate on the spherical and rotational Fourier transforms. We extend this to incorporate the epipolar geometry and we show how to compute the Radon transform in the frequency domain. We describe the algorithm in a form that can be easily replicated and we finish with experiments.

1.1 Related Work

Structure from motion without correspondences has a history since the 80's. Most of the approaches, called *direct* motion computation, assumed a temporally dense sequence so that computation of spatio-temporal derivatives is feasible. When assuming the projection of a plane [20, 24], the eight optical flow parameters can be estimated directly from the brightness change constraint equation. When no assumption about structure is made, several computation schemes have been

proposed [12]. The main constraint used is depth-positiveness and usually a variational problem is solved where depth is the unknown function over the image. Direct approaches based on normal optical flow or even just its direction have been thoroughly studied by Fermuller et al. [8] who also established formal conditions for ambiguity and instability of solutions. Jin et al. [13] have applied a direct method for simultaneous matching of regions and 3D-motion estimation over time by exploiting photometric constraints.

Among the approaches which do not use spatio-temporal derivatives and thus can afford any amount of motion, the closest to ours is the ones by Dellaert et al. [6], Antone and Teller [1], and Roy and Cox [22]. In [6], all possible assignments of 3D-points to image features are considered and the correct correspondence is established through an iterative expectation-maximization scheme where the E-step computes assignment weights and the M-step structure and motion parameters. In [1], images are already de-rotated using vanishing point correspondences and the translation is initialized via a Hough transform over all possible feature correspondences. Antone and Teller are the only ones who use the epipolar constraint and address the complexity of such a Hough transform. They propose ways to prune the search space through feature similarity as well as limits in the parameter space. In [22], an exhaustive search in the 5D parameter space is performed where for each motion hypothesis a cost function between points in the first image and segments of the corresponding epipolar line in the second image is computed. Our approach is also related to the learning of the epipolar geometry [26] though ours is not data-driven but requires a calibrated camera. Our approach is superior to [6] and [1] because it is not based on an iterative process which can possibly run through all assignments. While we use an exhaustive search in parameter space, the computation of the associated “likelihood” is accomplished without iteration but directly from the spherical harmonic coefficients. Our approach is superior to Roy and Cox only in the efficient computation of each motion hypothesis. We have not described here work on motion segmentation given correspondences. The reader is referred to the application of normalized cuts [23] and the generalized PCA [25] among tens of other papers on the subject.

2 Radon transform

The first steps of state-of-the-art motion estimation algorithms invariably involve generating and matching features between image pairs. The assumption is that a sufficient number of these hypothesized pairs will reflect true correspondences.




Figure 1: Concept: Instead of searching for corresponding points between images, we consider *all* feature pairs. The motion which is satisfied by the largest subset of feature pairs (weighted by a similarity measure) is considered to be the true camera motion. In the example above a weighting could be generated from the similarity between local blob structure.

Any subsequent processing, such as a RANSAC motion estimation, will then terminate quickly and correctly. The problem arises when this requirement cannot be satisfied. When dealing with image pairs with small overlap, or a particularly noisy scene for feature detection, the true correspondences within a group of matched features may be very small. Our desire to process images with small overlap and to resist outliers leads us to revisit classical robust accumulation algorithms like the Hough transform. In lieu of filtering sets of image features in search of the best matches, we will treat all possible feature pairs between two images. The only discriminating measure we will consider is a similarity between features. Our signal is not an image of greyscale intensities, but rather a function which maps feature pairs to their similarities. We will accomplish our robust accumulation via a filtering which, for any camera motion, collects and *counts* all the feature pairs which satisfy a geometrical motion constraint. The counting will be weighted by the feature similarities (see figure 1). The filtering result provides the score for a particular motion, and in this way we can evaluate all the possible camera motions. Before presenting the concrete specification of our formulation, we introduce necessary notation and definitions which we will use throughout this section.

Consider a camera moving rigidly in space. Assuming the intrinsic calibration parameters of the camera are known (meaning we can associate with each image pixel a ray in space), we can assume that the camera model is spherical perspective. This is useful since many single-viewpoint camera systems ranging from traditional CCD cameras to fish-eye lenses and even omnidirectional cameras can be treated with a spherical projection model. In this setting, points $P \in \mathbb{R}^3$ in the world project to points on the unit sphere: $p \in \mathbb{S}^2$, where $p = P/\|P\|$. We will identify rigid camera motions with elements of the Euclidean motion group $SE(3)$, with one notable irregularity. Since camera translations can only be recovered up to scale, we fix the scale of the translational motion component to have unit length. Although the set of all possible camera movements can be identified with $SE(3)$, we can represent any full observable camera motion with a pair $(R, T) \in \{R \in SO(3), T \in \mathbb{R}^3, \|T\| = 1\}$. We will parameterize $SO(3)$ with ZYZ Euler Angles such that $R(\alpha, \beta, \gamma) = R_z(\gamma)R_y(\beta)R_z(\alpha)$. The projection geometry in stereo pairs has been extensively studied, and it is well known that if points p and q represent projections of the same scene point in cameras separated by a motion (R, T) , they must obey the coplanarity (epipolar) constraint:

$$(Rp \times q)^T T = 0 \tag{1}$$

We are now prepared to concretely develop our accumulation. As we mentioned earlier, we will not be treating an image of intensities for our robust accumulation, but rather a function on feature pairs. We declare $g(p, q)$ to measure the similarity between points pairs in two images. Assuming an image has n pixels, the number of possible point pairs considered would be n^2 , of which clearly no more than n pairs can represent true correspondences. With such a miniscule percentage of inlying point pairs, it is essential that we construct a sufficiently discriminating weighting function $g(p, q)$. In our setting it is clear a simple image-based neighborhood similarity will not suffice. Instead of using intensity information directly, we have chosen to use the popular SIFT features [15], which histogram neighborhood gradient orientations at peaks and valleys of difference-of-Gaussian. These histograms typically make up a 128-dimensional vector (which we will denote with \tilde{p}), which affords us many options in selecting a similarity function. For example, our weighting could depend inversely on the Euclidean distance between two feature vectors:

$$g(p, q) = e^{-\|\tilde{p}-\tilde{q}\|} \tag{2}$$

Alternatively, we could choose a step function:

$$g(p, q) = \begin{cases} 1 & \text{if } \|\tilde{p} - \tilde{q}\| \leq \text{Threshold} \\ 0 & \text{otherwise} \end{cases}$$

To perform our robust accumulation, we need a way to filter and collect all the feature pairs (p, q) from the similarity function g which satisfy the epipolar geometry given by a particular motion. To this end, we introduce the *Epipolar Delta Filter (EDF)*. The EDF has the effect of counting all the feature pairs (p, q) which satisfy the motion constraint (weighted by their feature similarities $g(p, q)$), through an inner product with g . As the EDF captures the geometry of the epipolar constraint, it must encode the possible locations of an image point p after a camera motion. We choose the most straightforward definition constructed from the epipolar constraint:

$$\Delta_{(R,T)}(p, q) = \delta((Rp \times q)^T T) \quad (3)$$

We can now write our robust accumulation as a filtering of a similarity function g with the EDF:

$$G(R, T) = \int_{p \in \mathbb{S}^2} \int_{q \in \mathbb{S}^2} g(p, q) \Delta_{(R,T)}(p, q) dp dq \quad (4)$$

Effectively, $G(R, T)$ is a global likelihood function as the relative likelihoods of all possible motions are computed. The correct camera motion is expected to coincide with the global peak in this grid. Notice that the usual problems encountered with degenerate motions are avoided since in such cases there will not be a single global maximum but a region of maxima. To generate our likelihoods, we must compute the integral (4) as many times as the number of samples we are considering in our discrete motion space.

If N is the number of samples in each dimension of the motion space, and M the number of features identified in each image, then the complexity of this direct approach would be on the order of $O(N^5 M^2)$. This is an unacceptable load for almost any practical application. In the following sections we will demonstrate an efficient algorithm to generate the values of $G(R, T)$.

3 Motion estimation as correlation

In choosing to develop our global likelihood grid as a spherical filtering process, it is naturally revealed that the similarity function g is independent of the motion

parameters and the EDF is independent of any feature information. For now, we will focus our attention on the EDF $\Delta_{(R,T)}$. As the direction of camera translation is the unit vector $T \in \mathbb{S}^2$, we can represent T with a rotation $R_t \in SO(3)$: $T = R_t e_3$. This allows us to parameterize the space of camera motions with a rotation pair $(R, R_t) \in SO(3) \times SO(3)$. The Epipolar Delta Filter can now be redefined as

$$\begin{aligned}\Delta_{(R,R_t)}(p, q) &= \delta((Rp \times q)^T R_t e_3) \\ &= \delta((R_t^{-1}Rp \times R_t^{-1}q)^T e_3)\end{aligned}\quad (5)$$

If we write $R_c = R^{-1}R_t$ for the composite rotation embedding the rotational and translational terms, we see that the EDF simplifies to

$$\Delta_{(R_c,R_t)}(p, q) = \delta((R_c^{-1}p \times R_t^{-1}q)^T e_3) \quad (6)$$

Defining the rotation of a spherical function with the operator Λ_R ($\Lambda_R f(p) = f(R^{-1}p)$), the EDF (6) can be seen as just a spherical rotation of the canonical EDF given by $(R_c, R_t) = (I, I)$:

$$\begin{aligned}\Delta_{(R_c,R_t)}(p, q) &= \delta((R_c^{-1}p \times R_t^{-1}q)^T e_3) \\ &= \Delta_{(I,I)}(R_c^{-1}p, R_t^{-1}q) \\ &= \Lambda_{(R_c,R_t)}\Delta_{(I,I)}(p, q)\end{aligned}\quad (7)$$

To simplify notation, we will write $\Delta(p, q)$ in place of $\Delta_{(I,I)}(p, q)$. Notice that the canonical EDF $\Delta(p, q)$ captures a translation along the Z axis and a rotation of either 0° or 180° about the Z axis. With the evolution of the EDF into equation (7), we can revisit our original formulation of the global likelihood grid (3):

$$G(R_c, R_t) = \int_p \int_q g(p, q) \Lambda_{(R_c,R_t)} \Delta(p, q) dp dq \quad (8)$$

This shows us that our likelihoods can be computed as a correlation between spherical functions. In the next section we will explore the theory of generalized Fourier analysis to help alleviate some of the computational burden in evaluating our likelihood function.

4 Harmonic Analysis

The spherical correlation we are considering recalls the classical signal correlations on the real line or plane. Applications of such methods include standard

techniques in pattern matching. In such problems the search is for a planar shift (translational and/or rotational) which aligns a template pattern with a query image, where the location of highest correlation marks the correct alignment. These methods exploit the fact that correlations on the plane can be expressed as convolutions, and the well-known convolution theorem allows temporal convolutions to be replaced with pointwise multiplication in the spectral domain. Unfortunately, this property does not extend simply to the sphere, as convolutions and correlations on the sphere have different interpretations. Since it is not immediately clear what the relationship between the two formulations are, we will give a brief explanation.

A general definition of convolution can be given as

$$(f \star h)(x) = \int_{g \in G} f(g)h(g^{-1}x)dg$$

Here $f(x)$ and $h(x)$ are defined on some group G , and $g, x \in G$. If we take the real plane \mathbb{R}^2 to be a group with the action of translations, the convolution can be specifically written as

$$(f \star h)(x_1, x_2) = \int_{g_1} \int_{g_2} f(g_1, g_2)h(x_1 - g_1, x_2 - g_2)dg_1dg_2$$

This equation is the traditional form of planar convolution. Unfortunately, although the sphere is a manifold, it is not a group. We must find an alternative definition for the convolution of functions on the sphere. It is well known that the sphere is a homogeneous space of the group of 3D rotations $SO(3)$, with the isotropy subgroup of one dimensional rotations $SO(2)$ which keeps the north pole fixed [9]. A general definition of convolutions on homogeneous spaces can be given as

$$(f \star h)(x) = \int_{g \in G} f(g\eta)h(g^{-1}x)dg$$

Here $f(x)$ and $h(x)$ are defined on some homogeneous space of a group G , and η is given as the fixed point in G of the isotropy subgroup. The convolution of two functions on the sphere is given as

$$(f \star h)(x) = \int_{g \in SO(3)} f(ge_3)h(g^{-1}x)dg \quad x \in \mathbb{S}^2$$

Here e_3 is the standard Euclidean basis vector associated with the Z axis. Looking closely at this definition reveals that spherical convolution betrays the traditional concept of “measuring overlap” which is implied by planar convolution.

Here, points in one sphere ($f(x)$) are integrated through entire circles on the second sphere ($h(x)$). The resulting function $(f \star h)(x)$ is also defined on the sphere, hence spherical convolution reflects the properties of a filtering operator. To achieve the effect of a template matching operation, we must proceed to the general definition of correlation on homogeneous spaces:

$$c(g) = \int_x f(x)h(g^{-1}x)dx$$

As before f, h are defined on a homogeneous space of a group G , and $g \in G$ (alternatively, if we were interested in correlation on groups, we could just specify f, h to be functions on G). Identifying \mathbb{S}^2 as the homogeneous space of $SO(3)$ leads us to this definition of spherical correlation:

$$c(g) = \int_{x \in \mathbb{S}^2} f(x)h(g^{-1}x)dx$$

Here points on the sphere are given as unit vectors, and elements of the rotation group are given with the usual 3×3 rotation matrices. Notice that the resulting function $c(g)$ is defined not on the sphere but the group of rotations. This gives us the desired effect of measuring overlap. We rewrite this definition of spherical correlation using the notation developed earlier:

$$G(R) = \int f(\eta)\Lambda_R h(\eta)d\eta, \quad f, h \in \mathcal{L}^2(\mathbb{S}^2) \quad G(R) \in \mathcal{L}^2(SO(3)) \quad (9)$$

Here $\mathcal{L}^2(\mathbb{S}^2)$ denotes square-integrability, meaning the set of functions f such that $\int |f(\eta)|^2 d\eta$ is finite. If we wish to generalize the convolution theorem to correlation on the sphere, we must be able to answer three questions: (1) *How can we compute the Fourier transform of $f, h \in \mathcal{L}^2(\mathbb{S}^2)$ and $G \in \mathcal{L}^2(SO(3))$?* (2) *How does the spectrum of h change under a rotation $\Lambda_R h$?* (3) *How can we compute the Fourier transform of $G(R)$ efficiently using the answers to questions 1 and 2?* To answer these questions we will present a minimal introduction to spherical and rotational signal processing.

4.1 Fourier Transforms on \mathbb{S}^2 and $SO(3)$

This treatment of spherical harmonics is based on [7, 2]. In traditional Fourier analysis, periodic functions on the line (or equivalently functions on the circle \mathbb{S}^1), are expanded in a basis spanned by the eigenfunctions of the Laplacian. Similarly,

the eigenfunctions of the spherical Laplacian provide a basis for $f(\eta) \in \mathcal{L}^2(\mathbb{S}^2)$. These eigenfunctions are the well known spherical harmonics ($Y_m^l : \mathbb{S}^2 \mapsto \mathbb{C}$), which form an eigenspace of harmonic homogeneous polynomials of dimension $2l + 1$. Consequently, the $2l + 1$ spherical harmonics for each $l \geq 0$ form an orthonormal basis for any $f(\eta) \in \mathcal{L}^2(\mathbb{S}^2)$. The $(2l + 1)$ spherical harmonics of degree l are given as

$$Y_m^l(\theta, \phi) = (-1)^m \sqrt{\frac{(2l+1)(l-m)!}{4\pi(l+m)!}} P_m^l(\cos \theta) e^{im\phi}, \quad m = -l, \dots, l$$

where P_m^l are the associated Legendre functions and the normalization factor is chosen to satisfy the orthogonality relation

$$\int_{\eta \in \mathbb{S}^2} Y_m^l(\eta) Y_{m'}^{l'}(\eta) d\eta = \delta_{mm'} \delta_{ll'}, \quad (10)$$

where δ_{ab} is the Kronecker delta function. Any function $f(\eta) \in \mathcal{L}^2(\mathbb{S}^2)$ can be expanded in a basis of spherical harmonics:

$$f(\eta) = \sum_{l \in \mathbb{N}} \sum_{m=-l}^l \hat{f}_m^l Y_m^l(\eta) \quad (11)$$

$$\text{where } \hat{f}_m^l = \int_{\eta \in \mathbb{S}^2} f(\eta) \overline{Y_m^l(\eta)} d\eta \quad (12)$$

The \hat{f}_m^l are the coefficients of the Spherical Fourier Transform (SFT). Henceforth, we will use \hat{f}^l and Y^l to annotate vectors in \mathbb{C}^{2l+1} containing all coefficients of degree l .

Using a similar approach as seen above, we can develop a Fourier transform on the rotation group $SO(3)$ [3]. When considering functions $f \in \mathcal{L}^2(SO(3))$, the Fourier transform can be described as a change of basis from the group elements to the basis of irreducible matrix representations. The spherical harmonic functions Y_m^l form a complete, orthonormal set providing a basis for the representations of $SO(3)$. Furthermore, Schur's First Lemma from fundamental representation theory shows that they also supply a basis for the irreducible representations of $SO(3)$:

$$\Lambda_R Y^l(\eta) = U^l(R) Y^l(\eta). \quad (13)$$

The matrix elements of U^l are given by

$$U_{mn}^l(R(\alpha, \beta, \gamma)) = e^{-im\gamma} P_{mn}^l(\cos(\beta)) e^{-in\alpha} \quad m, n = -l, \dots, l. \quad (14)$$

The P_{mn}^l are generalized associated Legendre polynomials which can be calculated efficiently using recurrence relations. Such an Euler angle parameterization of the irreducible representations of $SO(3)$ leads to a useful expansion of functions $f \in \mathcal{L}^2(SO(3))$:

$$f(R) = \sum_{l \in \mathbb{N}} \sum_{m=-l}^l \sum_{p=-l}^l \hat{f}_{mp}^l U_{mp}^l(R) \quad (15)$$

$$\text{where } \hat{f}_{mp}^l = \int_{R \in SO(3)} f(R) \overline{U_{mp}^l(R)} dR \quad (16)$$

The \hat{f}_{mp}^l , with $m, p = -l, \dots, l$ are the $(2l+1) \times (2l+1)$ coefficients of degree l of the $SO(3)$ Fourier transform (SOFT, 16).

Now that we have answered our first question, we can try to understand how the spectrum of a function changes under a rotation. Intuitively, we would expect a rotation to manifest itself as a modulation of the Fourier coefficients as is the case in traditional Fourier analysis. This is, in fact, the observed effect. As spherical functions are rotated by elements of the rotation group $SO(3)$, the Fourier coefficients are “modulated” by the irreducible representations of $SO(3)$:

$$f(\eta) \mapsto \Lambda_R f(\eta) \iff \hat{f}^l \mapsto U^l(R)^T \hat{f}^l \quad (17)$$

The U^l matrix representations of $SO(3)$ are the spectral analogue to 3D rotations.

4.2 Rotation Estimation as Correlation

We are now prepared to address the final question regarding a generalized theorem for spherical correlation. Examining (9) more closely, we have developed the necessary tools to treat both $f(\eta)$ and $\Lambda_R h(\eta)$ with their respective Spherical Fourier expansions. Recently, [14, 17] have explored the computation of such a correlation in the spectral domain. Expanding the integral $\int f(\eta) \Lambda_R h(\eta) d\eta$ we have

$$G(R) = \sum_l \sum_{m=-l}^l \sum_n \sum_{p=-n}^n \sum_{k=-n}^n \hat{f}_m^l \overline{\hat{h}_p^n U_{pk}^n(R)} \int_{\eta \in \mathbb{S}^2} \overline{Y_k^n(\eta)} Y_m^l(\eta) d\eta.$$

Given the orthogonality of the spherical harmonic functions (10), the only nonzero terms in the summation appear when $n = l$ and $k = m$, thus

$$G(R) = \sum_l \sum_{m=-l}^l \sum_{p=-l}^l \hat{f}_m^l \overline{\hat{h}_p^l U_{pm}^l(R)}. \quad (18)$$

At this point, a direct application of the SOFT (16) for $G(R)$ produces

$$\hat{G}_{qr}^n = \sum_l \sum_{m=-l}^l \sum_{p=-l}^l \hat{f}_m^l \overline{\hat{h}_p^l} \int_{R \in SO(3)} \overline{U_{pm}^l(R)} U_{qr}^n(R) dR$$

The orthogonality of the matrices $U^l(R)$ ($\int U_{mp}^l(R) \overline{U_{qr}^n(R)} dR = \delta_{ln} \delta_{mq} \delta_{pr}$) yields nonzero terms in the summation only when $l = n$, $m = q$, and $p = r$, resulting in this simpler expression:

$$\hat{G}_{mp}^l = \hat{f}_m^l \overline{\hat{h}_p^l} \quad (19)$$

As we had initially desired, the result of the convolution theorem can indeed be generalized to correlation on the sphere: the $SO(3)$ Fourier coefficients of the correlation of two spherical functions can be obtained directly from the multiplication of the individual SFT coefficients. In vector form, the $(2l + 1) \times (2l + 1)$ matrix of SOFT coefficients \hat{G}^l is equivalent to the outer product of the coefficient vectors \hat{f}^l and \hat{h}^l . Given \hat{G}^l , the inverse SOFT retrieves the desired function $G(R)$ with $(2L + 1)$ samples in each of the three Euler angles, leaving us with accuracy of $\pm \left(\frac{180}{2L+1}\right)^\circ$ in α and γ , and $\pm \left(\frac{90}{2L+1}\right)^\circ$ in β . Here L is the bandwidth of the spherical images, meaning the largest degree for which we retain SFT coefficients.

Recalling our original problem of filtering a feature similarity function with the Epipolar Delta Filter (equation 8), we realize that we are actually correlating two functions on $\mathbb{S}^2 \times \mathbb{S}^2$. As one would expect, the theory we have just introduced extends easily. The Fourier transform for any function $f \in \mathcal{L}^2(\mathbb{S}^2 \times \mathbb{S}^2)$ is given as

$$f(p, q) = \sum_{l_1} \sum_{l_2} \sum_{m_1=-l_1}^{l_1} \sum_{m_2=-l_2}^{l_2} \hat{f}_{m_1 m_2}^{l_1 l_2} Y_{m_1}^{l_1}(p) Y_{m_2}^{l_2}(q) \quad (20)$$

$$\hat{f}_{m_1 m_2}^{l_1 l_2} = \int_p \int_q f(p, q) \overline{Y_{m_1}^{l_1}(p) Y_{m_2}^{l_2}(q)} dp dq \quad (21)$$

The spectrum of $G(R_c, R_t)$ from (8) can be obtained from the Fourier transforms of g, Δ :

$$\hat{G}_{m_1 m_2 k_1 k_2}^{l_1 l_2} = \overline{\hat{f}_{m_1 k_1}^{l_1 l_2}} \hat{\Delta}_{m_2 k_2}^{l_1 l_2} \quad (22)$$

As this last equation (22) shows, the Fourier space of our likelihood grid is six dimensional. However, we know that the space of observable motions is only

five dimensional. This discrepancy arises because we identify the rotation R_t with elements of $SO(3)$ even though the translation direction is independent of the first Euler angle of rotation:

$$R_z(\alpha_1)e_3 = R_z(\alpha_2)e_3 \quad \forall \alpha_1, \alpha_2$$

This issue is resolved easily in the following subsection.

4.3 The Canonical EDF and its Fourier Transform

The canonical Epipolar Delta Filter Δ embeds the epipolar geometry of the motions consistent with a rotation $R = I$ and translation $T = e_3$. As defined, it is only nonzero for point pairs $(p, q) \in \mathbb{S}^2 \times \mathbb{S}^2$ such that $(p \times q)^T e_3 = 0$. For any point p , the points q which satisfy this constraint must all lie on the same great circle. In particular, if we write image points with spherical coordinates θ and ϕ , then the points $p(\theta_1, \phi_1)$ and $q(\theta_2, \phi_2)$ can only satisfy the constraint $(p \times q)^T e_3 = 0$ iff $\phi_2 = \phi_1, \phi_1 + \pi$ or p or $q = \pm e_3$. Armed with this information, we can take a closer look at the Fourier transform of the EDF.

Proposition 1. *The Fourier transform of the EDF ($\hat{\Delta}_{m_1 m_2}^{l_1 l_2}$) is zero if and only if l_1 odd, l_2 odd, $|m_1|$ odd, $|m_2|$ odd, or $m_1 + m_2 \neq 0$.*

Proof. Let us begin by writing out the Fourier transform knowing that $\phi_2 = \phi_1, \phi_1 + \pi$:

$$\hat{\Delta}_{m_1 m_2}^{l_1 l_2} \propto \left[\int P_{m_1}^{l_1}(\cos \theta_1) \sin \theta_1 d\theta_1 \int P_{m_2}^{l_2}(\cos \theta_2) \sin \theta_2 d\theta_2 \int e^{i(m_1+m_2)\phi_1} d\phi_1 \right] (1+e^{im_1\pi}) \quad (23)$$

Immediately we see that if $|m_1|$ is odd, then $e^{im_1\pi} = -1$ and the $\hat{\Delta} = 0$. Equivalently, if we had taken the expansion making a variable substitution for ϕ_1 instead of ϕ_2 , we would have a trailing multiplicative term of $(1 + e^{im_2\pi})$, giving $\hat{\Delta} = 0$ if $|m_2|$ odd. Furthermore, the integral $\int P_{m_1}^{l_1}(\cos \theta_1) \sin \theta_1 d\theta_1 = 0$ if $m_1 = 0$ or $(l_1 + m_1)$ is odd. This means that $\hat{\Delta} = 0$ when l_1 is odd. The same argument shows $\hat{\Delta} = 0$ when l_2 is odd. The remaining integral $\int e^{-i(m_1+m_2)\phi_1} d\phi_1$ is only nonzero when $m_1 + m_2 = 0$, which means $\hat{\Delta} = 0$ whenever $m_1 + m_2 \neq 0$.

Now it remains to show the proposition holds in the other direction. If $\Delta = 0$, then we know at least one of the following must be true:

1. $\int P_{m_1}^{l_1}(\cos \theta_1) \sin \theta_1 d\theta_1 = 0$
2. $\int P_{m_2}^{l_2}(\cos \theta_2) \sin \theta_2 d\theta_2 = 0$

INPUT

1. A pair of spherical images I_1, I_2 .

OFFLINE

1. Compute the Fourier transform $\hat{\Delta}$ of Δ from (21).

ONLINE

1. Detect SIFT feature sets p, q from images I_1, I_2 .
2. From the cross product of the feature sets generate the similarity function g .
3. Compute the Fourier transform \hat{g} of g from (21).
4. Generate the 5D coefficient space $\hat{G}_{m_1 m_2 k_1 - m_2}^{l_1 l_2}$ from \hat{g} and $\hat{\Delta}$ as described in (24).
5. Using inverse Fourier transforms (15) obtain $G(R_c, R_t)$. Note: only a partial 2D inverse transform is needed for $R_t = R(0, \beta, \gamma)$.
6. Locate (R_c, R_t) at the maxima of G .
7. Extract the correct camera motion: relative orientation between cameras is $R = R_t R_c^{-1}$, and the direction of translation is $T = R_t e_3$.

Figure 2: The full motion estimation algorithm.

3. $e^{-im_1\pi} = -1$

4. $\int e^{-i(m_1+m_2)\phi_1} d\phi_1 = 0$

The first option can only be satisfied if $l_1 + m_1$ odd or $m_1 = 0$. The second option requires $l_2 + m_2$ odd or $m_2 = 0$. The third condition requires $|m_1|$ odd (as before, we can also derive the same requirement for $|m_2|$ odd). The final option holds only if $m_1 + m_2 \neq 0$, and this completes our proof. \square

We only have to consider $\hat{\Delta}_{m_1, -m_1}^{l_1 l_2}$ for $l_1, l_2, |m_1|$, even. We can now reduce the Fourier transform of the likelihood grid in (22):

$$\hat{G}_{m_1 m_2 k_1 - m_2}^{l_1 l_2} = \overline{\hat{f}_{m_1 k_1}^{l_1 l_2}} \hat{\Delta}_{m_2, -m_2}^{l_1 l_2} \quad (24)$$

Now that we have made our final simplification, we present an outline of the full algorithm in figure (2).

5 Discretization and Sampling

There are some issues we must address before we can finalize the transition from the continuous environment (integration of functions $f \in \mathcal{L}^2(\mathbb{S}^2)$) to the discrete setting (images and features). The most obvious concern relates to the Spherical Fourier Transform of a discrete spherical image. In addition to the existence of a sampling theorem, we need to be assured that the cost or complexity of the transform does not outweigh the benefits of replacing the correlation with a multiplication in the spectral domain. In other words, we need a Fast Discrete SFT. Driscoll and Healy [7] have developed a fast algorithm to compute the SFT from a uniformly sampled polar grid with complexity $O(L^2 \log^2 L)$, where L is the bandwidth of the signal being transformed. A similar separation-of-variables approach exists for a fast $SO(3)$ Fourier transform in $O(L^3 \log^2 L)$ [14].

The signal bandwidth L specifies the degree such that $\hat{f}^l = 0, \forall l \geq L$. Unfortunately, the signals we are dealing with (impulse responses for the similarity function g , and great circles for the EDF), do not have a frequency limit. Seeing as how we cannot compute and store an infinite number of Fourier coefficients, we can only hope to approximate these functions. Figure (3) shows the approximation of the EDF for different selections of the bandwidth L . The final consideration we will make deals with our motion estimation problem in particular. As the motion space is five-dimensional, the storage requirements for the coefficients \hat{G} and the function samples of G are quite large. For a selected bandwidth L , there will be $32L^5$ samples. This reality forces us to select lower values for L , such as 32. Although we are capturing high resolution images and locating image features with sub-pixel accuracy, the effective resolution of our spherical images is just $2L \times 2L$. The experiment detailed in figure (4) shows just how our algorithm reacts when the feature locations are affected by Gaussian noise when using such low-resolution spherical images.

6 Experiments

In this section we will present the results of the motion estimation algorithm on real image sequences. as well as a simulated result for detecting multiple motions in the scene. Before presenting the results we will address some practical considerations regarding spherical image acquisition and discrete Fourier transforms.

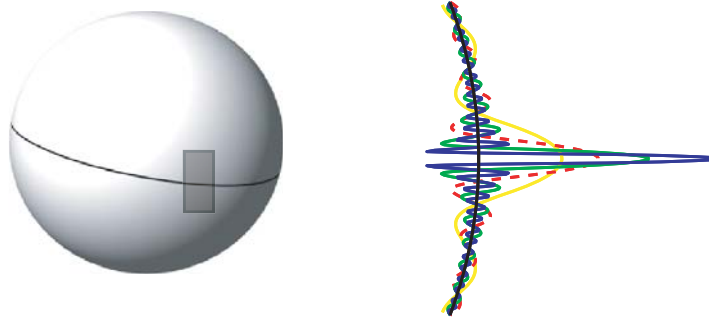


Figure 3: On the left is an spherical image of a great circle. Ideally, the function values are unity for any point on the circle, and zero otherwise. On the right is a slice of the sphere corresponding to the highlighted patch on the left. This shows the reconstructed function values of the delta function at the equator. The four values of the bandwidth L tested were 32, 64, 128, and 256. As L increases, the closer the approximation to an impulse, but because of the discontinuity there is also a greater overshoot (Gibbs phenomenon).

6.1 Spherical image acquisition

One of the benefits of choosing to model our camera with a spherical perspective projection is that it enables us to unite a number of single-viewpoint camera systems. Our experiments were performed with a catadioptric camera system along with a traditional digital camera.

The projection model of a central catadioptric system is equivalent to a spherical projection followed by a projection onto the plane [10]. If calibrated, such a sensor enables us to interpolate spherical perspective images. Our system consisted of a Canon Powershot G2 digital camera fastened to a parabolic mirror attachment from RemoteRealityTM[19]. Being that the mirror's field-of-view is 212° , the camera captures slightly more than a hemisphere of information. Figure (5) shows a sample catadioptric image obtained from a parabolic mirror and its corresponding projection onto the sphere.

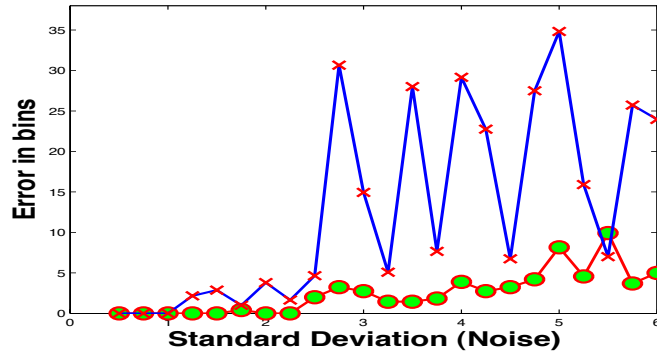


Figure 4: Results of a simulation testing the robust accumulation in the presence of Gaussian noise. The locations of spherical point correspondences are perturbed with Gaussian noise in each spherical coordinate. The standard deviation of the noise distribution is given here in pixels, and the error is computed by measuring the distance of the estimated solution from the correct solution in the 5D motion space. The plot marked with 'x' represents the simulation performed with bandwidth $L = 24$. In this case, a standard deviation of one pixel corresponds to 7.5° and 3.8° in the spherical coordinates ϕ and θ . The plot marked with circles is for $L = 32$, where a standard deviation of one pixel corresponds to 5.6° and 2.8° in the spherical coordinates. For this higher bandwidth, the results are still accurate in presence of significant noise.

6.2 Results

We proceed to show experimental results of our algorithm tested on a sequence of real omnidirectional images. For our tests, we assumed a function bandwidth of $L = 32$, which left us with a spatial resolution of $2L = 64$ samples in each of the five dimensions of our motion space. For comparison, we employed RANSAC to estimate the Essential matrix. Although it seems natural to use RANSAC in the presence of outliers, there are two crucial issues which would prevent a naive implementation from being operative. First is the volume of outliers. As the outliers in the set of possible feature pairs between two typical images is over 99%, the likelihood of selecting a minimal set of true correspondences is negligible. To this end, we discarded all but the best matching pairs during the random sampling stage. We retained only approximately 0.025% of the possible feature pairs (this translates to 250 feature pairs from a set of 10^6 possible pairs). The second issue is in determining the termination threshold of the RANSAC algorithm. In order to perform a proper evaluation of our algorithm, we implemented a best-case



Figure 5: Top Left: a parabolic catadioptric image. Bottom: the corresponding image on a uniformly sampled spherical grid. As the parabolic mirror images only a little more than half the sphere, you can see the lower portion of the spherical image contains no information. Top Right: the spherical image as it would appear on the surface of the sphere.

RANSAC which does not have a termination threshold but rather iterates 50,000 times. The essential matrix which satisfies the most feature pairs (weighted with $g(p, q)$) is selected as the motion. This ensures that a manual selection of the termination threshold may not be set too low to allow termination for an inferior motion. We have evaluated our Radon estimation alongside this modified RANSAC in order to provide an alternative method which is comparable to ours. In some of the following experimental results, the RANSAC performs very well and this is only because we have tuned these parameters quite finely.

We begin with a pure translational sequence of images. By fixing and sliding our camera along a rigid beam, we were able to generate two sequences of translational motion along the X and Z axes of the camera frame. Fixing the magnitude of motion between each frame, we were able to plot the estimated camera trajectory in figure (6). In general, there are four possible rotation and translation pairs which will satisfy a particular epipolar constraint. These solutions correspond to

the true solution, a baseline reversal, a camera rotation of 180° about the baseline (commonly referred to as the “twisted pair” configuration), or a twisted pair with baseline reversal. If the true motion is given by (R, T) , the other three motions which satisfy the same epipolar constraint are given by $(R, -T)$, $(e^{\hat{T}\pi}R, T)$, and $(e^{\hat{T}\pi}R, -T)$ (note that $e^{\hat{T}\pi}$ gives a rotation of 180° about the T axis). In order to identify the expected locations of the four peaks in our likelihood space for the correct motion, we must remember that we identify elements of this five dimensional motion space with the pair (R_c, R_t) where $R = R_t R_c^{-1}, T = R_t e_3$. If we define $R'_t = R_z(\gamma)R_y(\beta)$ so that $-T = R'_t e_3$, then we can expect the four peaks to be located at $(R^{-1}R_t, R_t)$, $(R^{-1}R'_t, R'_t)$, $((e^{\hat{T}\pi}R)^{-1}R_t, R_t)$, and $((e^{\hat{T}\pi}R)^{-1}R'_t, R'_t)$. In the figures, for example, when we show a 2D translational slice with a peak at R'_t , this slice can be generated from the bins corresponding to the rotation $R^{-1}R'_t$. In figure (6), the slice shown depicts a peak at $R_t(0, \frac{\pi}{2}, \pi)e_3 = -e_1$.

A similar experiment was performed with the camera moving along the Z axis. The motion was recovered from pairs of consecutive images, with the estimated camera path shown in figure (7). Our Radon estimation has a smaller deviation from the observed ground truth Z axis than the RANSAC estimation.

In order to test both rotations and translations while recording ground-truth observations, we positioned the camera at the outside edge of a turntable. This allowed us to capture images from the camera moving around in a circle. There was a 45° rotation between each of the images in this sequence, and the estimated camera positions are shown in figure (8). Although the Radon’s trajectory estimate deviates slightly from the plane, the positions as seen from the overhead view coincide with the recorded ground truth more accurately than the RANSAC estimation. After 6 pairwise tests, there was little error accumulation in estimating the trajectory.

We now discuss results of an experiment from a sequence of images from an outdoor environment. Figure (9) shows a representative selection of images from this sequence. Figure (10) shows some results from the motion estimation. Epipolar lines are drawn to allow visual confirmation of the method’s accuracy. Figure (11) displays the obtained camera trajectory using the visualization tools provided in the Epipolar Geometry Toolbox [18]. This same trajectory is projected onto the X-Z plane in figure (12) to show the deviation from the ground plane which is known to be (approximately) the correct plane of motion. This planar motion also restricts the axis of rotation to align with the Z axis, and figure (13) shows just how closely the measured rotations reflect this property.

7 Conclusion

We have presented a novel approach for the computation of 3D-motion from two views without correspondences. It is based on the generation of a global likelihood function on the space of all observable camera motions. Given today's computing power, it is not the search through this likelihood function but rather the combinatorial explosion of all possible correspondences that is intractable. Instead of traversing all possible correspondence assignments, our method computes for each motion hypothesis a correlation function which considers only feature pairs satisfying the epipolar constraint. Such a formulation can be expressed as a correlation integral if the integration path can be written as a group action over the domain of integration. In this case, the integral can be computed as an inner-product in the Fourier domain. The bandwidth limitation affects directly the resolution of the parameter space and it is indeed our future work to establish a "space localization" using wavelets. Such a localization in the parameter space would also allow a constrained search when prior distributions of motion are established causally through time. In that case, we could also achieve near real-time performance which right now is impossible in all correspondenceless approaches.

References

- [1] M. Antone and S. Teller. Scalable, extrinsic calibration of omni-directional image networks. *International Journal of Computer Vision*, 49:143–174, 2002.
- [2] G. Arfken and H. Weber. *Mathematical Methods for Physicists*. Academic Press, 1966.
- [3] G. Chirikjian and A. Kyatkin. *Engineering Applications of Noncommutative Harmonic Analysis: With Emphasis on Rotation and Motion Groups*. CRC Press, 2000.
- [4] K. Daniilidis and M. Spetsakis. Understanding noise sensitivity in structure from motion. In Y. Aloimonos, editor, *Visual Navigation*, pages 61–88. Lawrence Erlbaum Associates, Hillsdale, NJ, 1996.
- [5] S. Deans. Hough transform from the radon transform. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 3:185–188, 1981.
- [6] F. Dellaert, S. Seitz, C. Thorpe, and S. Thrun. Structure from motion without correspondence. In *CVPR*, Hilton Head Island, SC, June 13-15, 2000.
- [7] J. Driscoll and D. Healy. Computing Fourier transforms and convolutions on the 2-sphere. *Advances in Applied Mathematics*, 15:202–250, 1994.
- [8] C. Fermuller and J. Aloimonos. Direct perception of three-dimensional motion from patterns of visual motion. *Science*, 270:1973–1976, 1995.

- [9] J. Gallier. Notes on group actions, manifolds, lie groups, and lie algebras. <http://www.cis.upenn.edu/~cis610/lie1.pdf>, 2005.
- [10] C. Geyer and K. Daniilidis. Catadioptric projective geometry. *International Journal of Computer Vision*, 43:223–243, 2001.
- [11] C. Geyer, S. Sastry, and R. Bajcsy. Euclid meets Fourier. In *Workshop on Omnidirectional Vision*, Prague, 2004.
- [12] B. Horn and E. Weldon. Direct methods for recovering motion. *International Journal of Computer Vision*, 2:51–76, 1988.
- [13] H. Jin, P. Favaro, and S. Soatto. A semi-direct approach to structure from motion. *The Visual Computer*, 19:1–18, 2003.
- [14] P. J. Kostelec and D. N. Rockmore. FFTs on the rotation group. In *Working Paper Series, Santa Fe Institute*, 2003.
- [15] D. Lowe. Sift (scale invariant feature transform): Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60:91–110, 2004.
- [16] A. Makadia, C. Geyer, S. Sastry, and K. Daniilidis. Radon-based structure from motion without correspondences. In *IEEE Conf. Computer Vision and Pattern Recognition*, San Diego, June 20-25, 2005.
- [17] A. Makadia, L. Sorgi, and K. Daniilidis. Rotation estimation from spherical images. In *Proc. Int. Conf. on Pattern Recognition*, Cambridge, UK, 2004.
- [18] G. L. Mariottini and D. Prattichizzo. The epipolar geometry toolbox: multiple view geometry and visual servoing for Matlab. *IEEE Robotics & Automation Magazine*, 2005.
- [19] S. Nayar. Catadioptric omnidirectional camera. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 482–488, Puerto Rico, June 17-19, 1997.
- [20] S. Negahdaripour and B. Horn. Direct passive navigation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 9:168–176, 1987.
- [21] J. Oliensis. A critique of structure from motion algorithms. *Computer Vision and Image Understanding*, 80:172–214, 2000.
- [22] S. Roy and I. Cox. Motion without structure. In *Proc. Int. Conf. on Pattern Recognition*, Vienna, Austria, 1996.
- [23] J. Shi and J. Malik. Motion segmentation and tracking using normalized cuts. In *Proc. Int. Conf. on Computer Vision*, 1998.
- [24] R. Szeliski and S. B. Kang. Direct methods for visual scene reconstruction. In *IEEE Workshop on Representations of Visual Scenes*, pages 26–33, 1995.
- [25] R. Vidal and Y. Ma. A unified algebraic approach to 2-d and 3-d motion segmentation. In *Proc. Eighth European Conference on Computer Vision*, pages 1–15, 2004.
- [26] Y. Wexler, A. Fitzgibbon, and A. Zisserman. Learning epipolar geometry from image sequences. In *IEEE Conf. Computer Vision and Pattern Recognition*, Wisconsin, June 16-22, 2003.

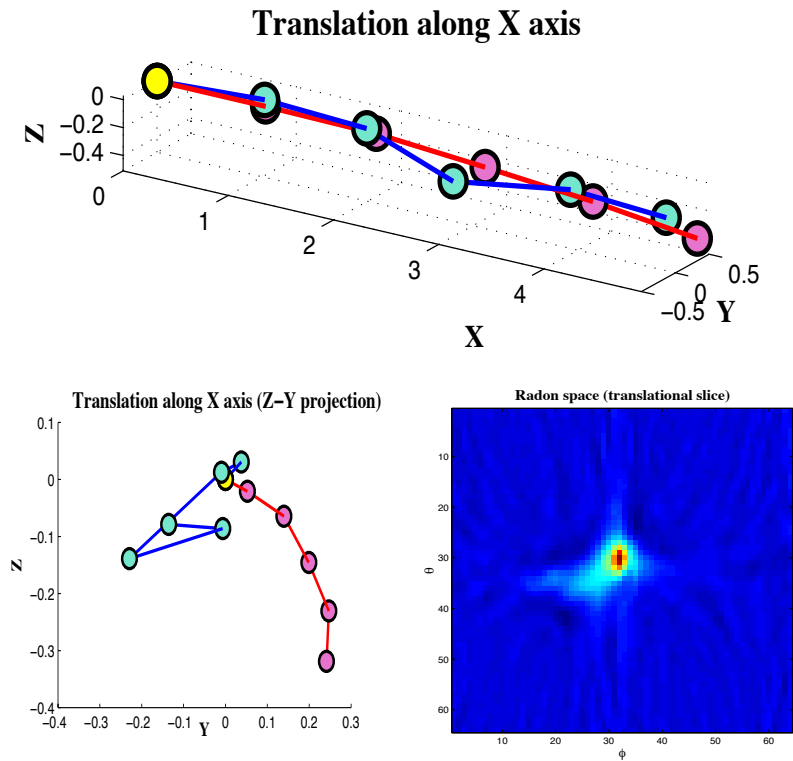


Figure 6: Top: the estimated trajectory of the camera. In blue (light) is the Radon estimation, in red (dark) is the RANSAC computation, and the yellow circle marks the starting position. Bottom Left: A projection of the trajectory onto the Z-Y plane showing the deviation of the estimated positions from the X axis. Bottom Right: the R_t slice of the grid G where the maxima was found.

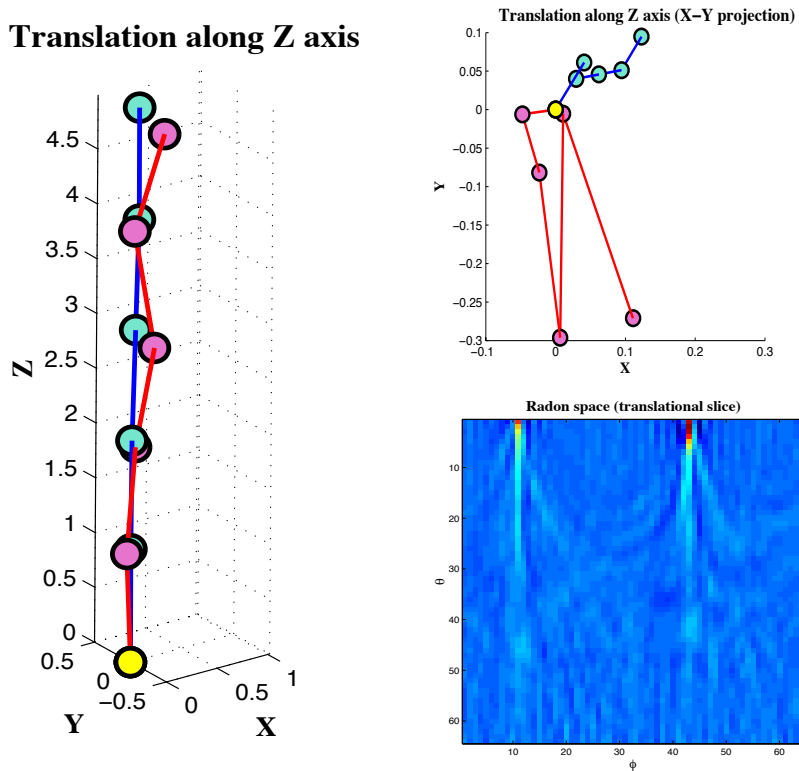


Figure 7: Left: the estimated trajectory of the camera. In blue (light) is the Radon estimation, in red (dark) is the RANSAC computation, and the yellow circle marks the starting position. Top Right: A projection of the trajectory onto the X-Y plane showing the deviation of the estimated positions from the Z axis. Bottom Right: the R_t slice of the grid G where the maxima was found (notice the peak is locate at $\theta \approx 0$, which corresponds to the correct translation along Z).

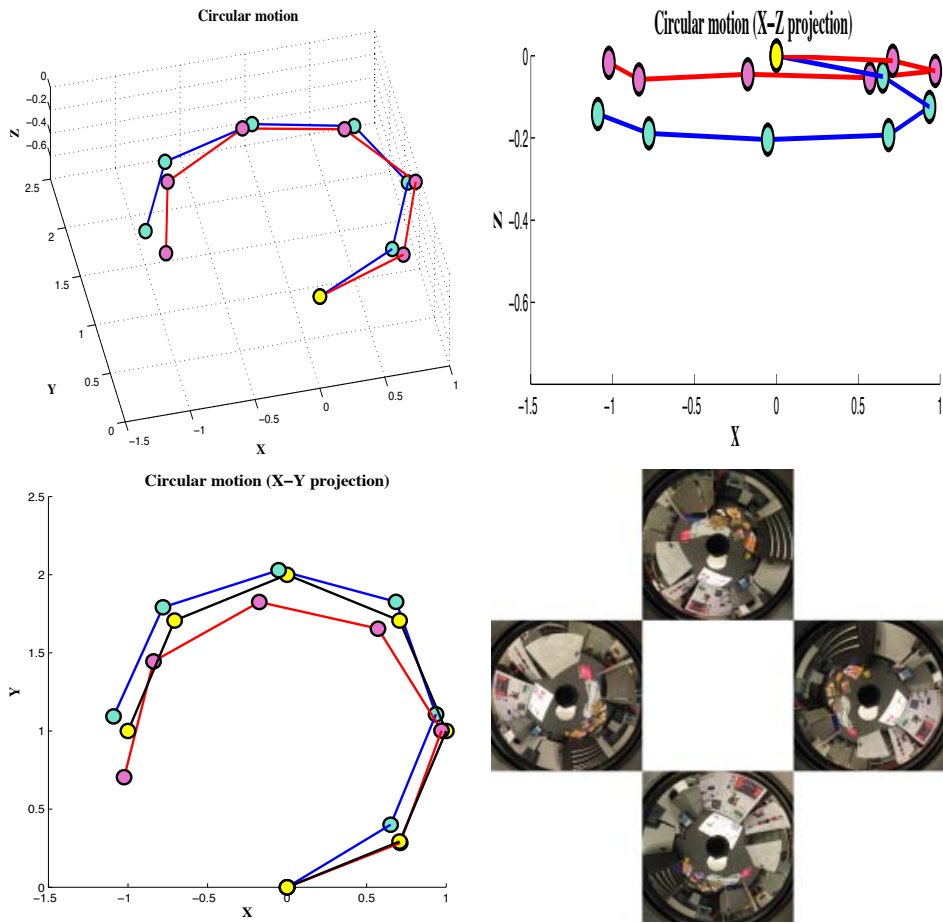


Figure 8: A camera moving along a circular path. Top Left: In blue (light) is the Radon estimation, in red (dark) is the RANSAC. Top Right: A projection onto the X-Z plane showing the deviation from the plane of the turntable. Bottom Left: An overhead view. The yellow circles are the observed ground truth positions of the camera. Bottom Right: four images from the sequence. Even though the dominant motion is rotation, the translation is still effectively detected by the Radon.




Figure 9: Top Row: two representative images from a sequence of outdoor images. The motion between image positions is over five meters. At each position, the equatorial plane of the spherical image is roughly aligned to be parallel with the ground plane to provide a rough, partial ground truth of the motion. The image sequence also contains some dynamic scene content as there were people moving throughout the scene as the images were taken. The bottom two images are the spherical projections of the original omni images. Only the visible band on the sphere is shown here.




Figure 10: On the top row are a pair of images from a sequence for which the motion was estimated. The bottom two rows show the images after they have been rotationally aligned. Epipolar circles have been overlaid onto the images. Since the images have been rotationally aligned, points which lie along these circles in one image will lie along the same circle in the second image. The intersection of these circles mark the focus of expansion and contraction, which define the direction of translation between this image pair. The rotation between image pairs was estimated at approximately 45° , and as the focus of expansion shows the translation was roughly in the equatorial plane.

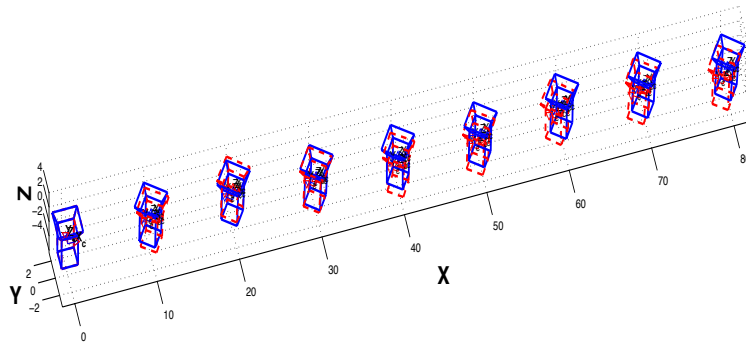


Figure 11: This figure shows the camera trajectory estimated from a sequence of images (see figure (9) for sample images). The camera frames drawn with solid (blue) lines depict the trajectory estimated using the Radon transform, while the dashed (red) lines show the RANSAC trajectory. In this sequence the motion is known to be approximately planar in the equatorial plane. Since the magnitude of camera motion cannot be recovered from pairs of images alone, we have fixed the distance between camera positions to be 10 units for visual purposes.




Figure 12: This figure shows the camera trajectory estimated from a sequence of images (see figure (9) for sample images), projected onto the ZX plane. The camera frames drawn with solid (blue) lines depict the trajectory estimated using the Radon transform, while the dashed (red) lines show the RANSAC trajectory. The motion is known to be planar (on the equatorial plane) and the Radon estimate reflects this more accurately. Since the magnitude of camera motion cannot be recovered from images alone, we have set the distance between camera positions to be 10 units for visual purposes.

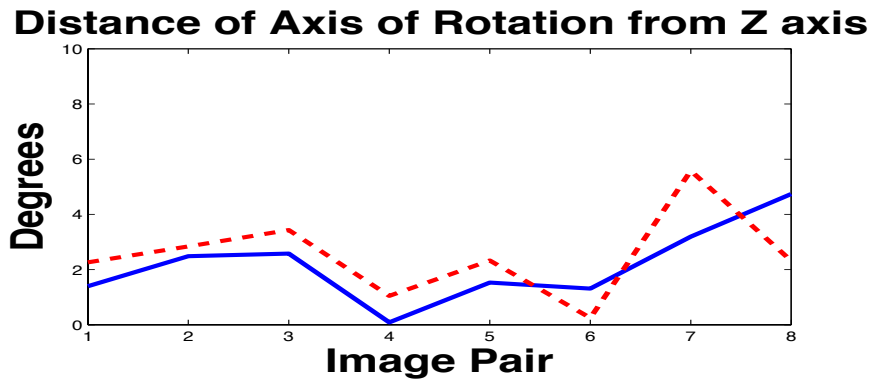


Figure 13: As the camera motion in this sequence (see figure (9) for sample images) is known to be roughly planar, we know that the axis of rotation must align with the Z axis. This plot shows for all eight image pairs in the sequence the distance in degrees of the estimated axis of rotation from the Z axis. The solid line (blue) is the estimate from our Radon integral, and the dashed line (red) is the RANSAC estimate.