

Toward Geometric Visual Servoing

Noah John Cowan¹ and Dong Eui Chang²

¹ University of California, Berkeley, CA 94720, USA

² University of California, Santa Barbara, CA 93106, USA

Abstract. This paper presents a global diffeomorphism from a suitably defined “visible set” of rigid body configurations — a subset of $SE(3)$ — to an image-space. The mapping between the visible set and the image-space is given by the projection of a set of features of a specially designed visual target. The target is a sphere marked with a feature point and a vector tangent to the sphere at the feature point. We show how the construction of a diffeomorphism to image-space should pave the way for developing global, dynamic visual servoing systems using Navigation Functions.

Keywords. Visual servoing, geometric control, occlusions, field-of-view, navigation functions.

1 Introduction

There is a large and growing body of algorithms for “visual servoing” (VS) — motion control using visual feedback. Traditionally, VS algorithms generate motor reference velocities to register a camera’s current view of a scene with a previously stored view (for a tutorial, see [7]).

We seek to move VS toward a systematic theory by characterizing the geometry of “visible” configurations of a visual target relative to a camera. In particular, for a specific target geometry we present a *diffeomorphism* — a smooth and smoothly invertible transformation — from an appropriately defined visible set of configurations to an image space. We believe this transformation will enable the construction of purely image-based, global dynamic VS algorithms.

1.1 Background

A significant challenge involves representing rigid motions in terms of visually measured quantities. Ideally, such a representation should enable effective encoding of

- **Configuration and State**, e.g. position and velocity or position and momentum for Lagrangian or Hamiltonian systems.
- **Tasks and goals**, e.g. trajectories in the state space or points in the configuration space.

- **Obstacles**, e.g. the edge of the field-of-view (FOV) for VS systems.
- **Uncertainty**, e.g. sensor and actuator noise or parametric error.

There are several candidate representations of image-based rigid motion to consider from the literature. The classical approach to “2D VS” employs the projection, treated as a vector in \mathbb{R}^n , of an arbitrary set of feature points [7]. The redundancy of using extra feature points seems to confer robustness to measurement noise in any one of the feature measurements. However, the movement of features is constrained by the underlying rigid motion, rendering image-based control and motion planning in image space challenging for large deviations from a goal. Notwithstanding those challenges, Corke and Hutchinson [2] created a 2D kinematic algorithm for 6DOF VS that seems (empirically) to have a very large basin of attraction while keeping features in the FOV. Their algorithm employs a clever choice of image features which helped motivate the choice of features used in this paper.

A more recent approach uses partial pose reconstruction: given a sufficient number of feature points, the relative pose, up to a scale in translation, between two views may be determined without exploiting a geometrical model of the points. Using this technique, researchers developed six DOF VS algorithms robust to calibration uncertainty [12,15]. It is worth noting that the methods used require sufficient point correspondences between views to fully reconstruct a geometric model of the visual target [10]. Application of this method to contexts besides full six DOF VS remains a challenge.

Alternatively, one may recover the complete pose of a camera with respect to a target by exploiting a model of the target [11]. Vision-based controllers using full pose reconstruction are often referred to as “3D VS” algorithms. Model based pose reconstruction requires fewer feature points than the model-free approach described above, and has the added advantage of fully recovering feature depth, effectively reducing the camera to a “virtual Cartesian sensor.” Representing visibility obstacles, such as the FOV or self-occlusions is less parsimonious, but can be done [5]. Formal results demonstrating parametric robustness of VS systems using this method remain elusive.

Generalized image-based coordinates have proven extremely effective in a few narrow contexts [3,5,16]. Generalized coordinates describe kinematic motion with one variable per mechanical DOF. Lagrange’s equations, for example, are usually written using such coordinates. Hence, this approach enables the expression of *dynamical* equations of motion in terms of measured quantities on the image plane. Obstacles such as the FOV and self-occlusions often appear as the boundary of a compact manifold in image-space and hence their avoidance may be cast as an instance of dynamical obstacle avoidance [5]. Although quite robust in practice, obtaining formal guarantees of robustness to noise or parametric uncertainty for this framework remains an open problem.

1.2 Contribution

To date, global image-based representations of configuration have been applied only to three DOF systems. This paper builds on previous results in a key way: we present an image-based, geometric representation of six DOF rigid motion. Our development of a global representation of “visible” rigid motions viewed through the projection of a set of features should help pave the way for new global, dynamic VS systems.

Organization. In Section 2, we employ a specific target geometry — a sphere with a few markings — to create a global image-based representation of motion for six DOF VS. Included in our development is a simple, purely image-based representation of the so-called image Jacobian (made possible since, as we show, the image and task spaces are diffeomorphic). In Section 4, we suggest a method for using our diffeomorphism for kinematic or dynamic control, although there is much open work to be done in this endeavor. Finally, we give some concluding remarks in Section 5.

2 Six DOF Diffeomorphism to Image-space

We assume a visual target may be designed to our specifications, so we may explore new image-based representations of rigid motion. In cases in which we have the freedom to design visual targets — for example when designing docking stations for space craft, helicopter landing beacons, or visual targets for a factory setting — this approach may lead to novel target designs that ease the control problem. More generally, it is hoped that the insight drawn from taking this approach may enable us to reinterpret target geometries over which we have less design freedom.

Consider the problem of moving a rigid target object in six DOF relative to a perspective camera. The rigid target considered is as follows:

1. **A spherical body.** Consider a spherical body of radius ρ . As the body moves away from the camera, its projection gets smaller. Roughly speaking, the position and size of the body’s image encodes the position of the center of the body relative to the camera.
2. **A single point on the body.** Adding a visible point to the body breaks the visual symmetry, allowing us to resolving two rotational DOF’s from the location of the feature point on the image.
3. **A unit vector tangent to the body.** The final degree of freedom is resolved by considering the orientation on the image of a projected vector attached to our feature point on the body.

Zhang and Ostrowski [16,17] developed the idea of projecting a spherical body to an image plane for VS of a blimp relative to a large ball. Using a “flat” image plane, the resulting image is an ellipse, which they approximate as a circle by assuming that a slice of the spherical body parallel to the image

Table 1. List of symbols.

Symbol	Description
$o, p, b, \dots \in \mathbb{E}^3$	Euclidean points (Roman)
$\mathbf{v}, \mathbf{e}, \dots \in \mathbb{R}^3$	vectors (boldfaced)
$\mathbf{e}_1, \mathbf{e}_2, \dots \in \mathbb{R}^3$	standard basis
$\pi: \{\mathbb{E}^3 - o_c\} \rightarrow \mathbb{S}^2$	image projection model – spherical panoramic camera
$\mathcal{F} = \{o, \mathbf{i}, \mathbf{j}, \mathbf{k}\}$	rigid coordinate frame, $o \in \mathbb{E}^3$ and $\mathbf{i}, \mathbf{j}, \mathbf{k} \in \mathbb{R}^3$
$\mathcal{F}_c, \mathcal{F}_b$	camera frame and body frame
p^b, \mathbf{v}^b	point, p , and vector, \mathbf{v} , with respect to \mathcal{F}_b
$H \in \text{SE}(3)$	rigid transformation of \mathcal{F}_b , relative to \mathcal{F}_c
$R \in \text{SO}(3)$	rotation effected by H , columns $R = [\mathbf{r}_1 \ \mathbf{r}_2 \ \mathbf{r}_3]$
$\mathbf{d} \in \mathbb{R}^3$	translation effected by H
$p^c = Hp^b, \mathbf{v}^c = R\mathbf{v}^b$	point, p , and vector, \mathbf{v} , with respect to \mathcal{F}_c
$\nu: \text{SE}(3) \rightarrow \mathbb{R}$	measure of feature visibility, (5)
$\mathcal{V} \subset \text{SE}(3)$	set of “visible” configurations, $H \in \mathcal{V} \iff \nu(H) > 0$
$\lambda \in (0, 1)$	radius on image sphere of body, (3)
$\mathbf{s} \in \mathbb{S}^2$	unit vector pointing toward body centroid, (3)
$Q \in \text{SO}(3)$	image-based rotation, columns $Q = [\mathbf{q}_1 \ \mathbf{q}_2 \ \mathbf{q}_3]$, (7)
$(Q, \lambda, \mathbf{s}) = c(H)$	camera map, (8)
\mathcal{I}	image feature space, $\mathcal{I} \subset \text{SO}(3) \times (0, 1) \times \mathbb{S}^2$, (9)

plane is projected. The present paper builds on that work, employing a more ‘exact’ diffeomorphism to the image-space, as well as incorporating additional markings on the body whose projection encodes rotational information.

2.1 Notation and Definitions

At the risk of burdening the reader with formalism, we present the following definitions to enable a precise geometric description of the domain and range of a camera viewing rigid motions.

An affine point $p \in \mathbb{A}^3$ has homogeneous coordinates $p = [p_1 \ p_2 \ p_3 \ 1]^T$ with respect to some rigid frame. Note that $T\mathbb{A}^3 = \mathbb{A}^3 \times \mathbb{R}^3$, and that \mathbb{R}^3 acts on points to translate them in the usual way, so that if $\mathbf{v} = [v_1 \ v_2 \ v_3]^T \in \mathbb{R}^3$ and $p \in \mathbb{A}^3$, then $p + \mathbf{v} = [p_1 + v_1 \ p_2 + v_2 \ p_3 + v_3 \ 1]^T \in \mathbb{A}^3$. Two points cannot be “added” together, but if $p, b \in \mathbb{A}^3$ then $\mathbf{v} = p - b \in \mathbb{R}^3$ is the vector such that $p = b + \mathbf{v}$. Adding the usual metric structure to affine space \mathbb{A}^3 yields Euclidean space \mathbb{E}^3 where the distance between two points is given by the two norm of their difference, $\|p - b\|$ (a measure independent of the choice of rigid frame).

A rigid frame, \mathcal{F} , is defined by its origin, $o \in \mathbb{E}^3$, and three mutually orthogonal unit vectors, $\mathbf{i}, \mathbf{j}, \mathbf{k} \in \mathbb{R}^3$, that create a right-handed frame. Consider a full perspective (“pinhole”) camera with frame \mathcal{F}_c such that o_c is located at the pinhole (or optical center), with \mathbf{k}_c aligned with the optical axis. The

pinhole camera projects points in the open half space “in front” of the camera to an image-plane pair, given by via the map, $\pi^+ \{ \mathbb{E}^3 : (p - o_c) \cdot \mathbf{k}_c > 0 \} \rightarrow \mathbb{R}^2$, expressed in camera frame coordinates

$$\pi^+(p) = \frac{f}{p_3} \begin{bmatrix} p_1 \\ p_2 \end{bmatrix} \quad p_3 > 0, \quad (1)$$

where f is the camera focal length. The camera observes features of a rigid body, affixed with rigid frame \mathcal{F}_b . Let

$$H = \begin{bmatrix} R & \mathbf{d} \\ 0^T & 1 \end{bmatrix} \in \text{SE}(3), \quad \text{where} \quad R = [\mathbf{r}_1 \ \mathbf{r}_2 \ \mathbf{r}_3] \in \text{SO}(3), \ \mathbf{d} \in \mathbb{R}^3,$$

denote the rigid transformation of \mathcal{F}_b relative to \mathcal{F}_c . A point expressed with respect to the body-frame as p^b , appears as $p^c = Hp^b$ with respect to the camera frame. Similarly, if \mathbf{v}^b is a vector in the body frame, then $\mathbf{v}^c = R\mathbf{v}^b$ is the same vector with respect to the camera frame.

Hamel *et. al.* [6] remap the image plane to a sphere to recover some symmetry that is “broken” by a flat image plane. This approach has also been used in the structure from motion (SFM) literature [1]. Let $\mathbf{p} = (p - o_c)$ and note that the unit vector, $\mathbf{p}/\|\mathbf{p}\|$ may be recovered from the image-plane pair in (1) since

$$\frac{\mathbf{p}}{\|\mathbf{p}\|} = \begin{bmatrix} \pi^+(p) \\ f \end{bmatrix} / \left\| \begin{bmatrix} \pi^+(p) \\ f \end{bmatrix} \right\|, \quad p_3 > 0$$

with respect to the camera frame. Of course, this assumes that we know the parameter f (or, more generally, all so-called “intrinsic” camera parameters, omitted to simplify the presentation). Motivated by this observation, we consider for convenience a “panoramic” spherical camera

$$\pi : (\mathbb{E}^3 - \{o_c\}) \rightarrow \text{S}^2 \quad (2a)$$

$$: p \mapsto \frac{\mathbf{p}}{\|\mathbf{p}\|} \quad \text{where} \quad \mathbf{p} = (p - o_c). \quad (2b)$$

For the purposes of this paper, $\text{S}^2 = \{ \mathbf{v} \in \mathbb{R}^3 : \mathbf{v} \cdot \mathbf{v} = 1 \} \subset \mathbb{R}^3$. For the camera map, S^2 corresponds to the unit tangent space of \mathbb{E}^3 at o_c , namely “the set of unit vectors originating from the camera origin.” To keep features within a finite FOV, one may introduce an appropriate image-space “obstacle” into the controller design (see Section 4).

2.2 Image-based Translation

Attach the body frame at the center of the sphere, so that the location of the body relative to the camera origin is given by $o_b - o_c = \mathbf{d}$. If $\|\mathbf{d}\| > \varrho$ — *i.e.* the body remains bounded away from the camera origin — then the surface

of the body double-covers a topological disc on S^2 via the map π . The edge of the disc, a planar slice of the image-sphere, is a perfect circle of radius

$$\lambda = \frac{\varrho}{\|\mathbf{d}\|}, \quad \varrho < \|\mathbf{d}\| < \infty.$$

(The circle radius, λ , appears dimensionless because the image-sphere was normalized to unit radius). The center of the circle on the image-sphere is in the direction of

$$\mathbf{s} = \frac{\mathbf{d}}{\|\mathbf{d}\|}$$

and is readily measurable from the projection of the body.

Let $\mathcal{B} := \{\mathbf{d} \in \mathbb{R}^3 : \|\mathbf{d}\| > \varrho\}$ denote the translations of the body origin that keep it a body radius away from the camera. We now have a diffeomorphism — a smooth and smoothly invertible function — from locations of the body to image measurements, $c_1 : \mathcal{B} \rightarrow (0, 1) \times S^2$, given by

$$\boxed{c_1 : \mathbf{d} \mapsto (\lambda, \mathbf{s})}. \quad (3)$$

The inverse of c_1 is given simply by

$$c_1^{-1}(\lambda, \mathbf{s}) = \frac{\varrho}{\lambda} \mathbf{s}. \quad (4)$$

2.3 Image-based Rotation

To break the rotational symmetry of our spherical rigid body, attach a visible feature point, b , to its surface, and a unit vector \mathbf{a} tangent to the body at that point. For convenience, align the body frame so that origin coincides with the center of the body, and the unit vector $(b - o_b)/\varrho$ lies along the negative \mathbf{k}_b axis. Hence, in the body frame $b^b = [0, 0, -\varrho, 1]^T$.

As we will show, the projection of b to the image-sphere, $\mathbf{q}_1 = \pi(b)$, encodes two rotational degrees of freedom. We encode the final degree-of-freedom by projecting a unit vector or “arrow”, \mathbf{a} , tangent to the body at the point b . In practice, the vector \mathbf{a} may be approximated by two distinguishable points on the surface of the sphere. Again for convenience we assume the vectors body-fixed representation is simply $\mathbf{a}^b = \mathbf{e}_2$. Let $\mathbf{b} = b - o_c$ denote the vector from the camera origin to the body point b . Recalling that the rotation matrix R has columns $(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3)$, then with respect to the camera frame, we have

$$b^c = \begin{bmatrix} b^c \\ 1 \end{bmatrix} = Hb^b, \quad \text{where } \mathbf{b}^c = \mathbf{d} - \varrho\mathbf{r}_3, \quad \text{and } \mathbf{a}^c = R\mathbf{a}^b = \mathbf{r}_2.$$




Fig. 1. Projection of a spherical body with a feature point on it to the image-sphere. The image-plane measurement is given by $y = (Q, \lambda, \mathbf{s}) = c(H)$.

Note that $(b - o_b) \cdot \mathbf{a} = -\varrho \mathbf{e}_3 \cdot \mathbf{e}_2 = 0$.

Some configurations cause the body to occlude the feature point, b . This occurs when $(b - o_c) \cdot (o_b - b)$ becomes negative. Hence, we define a “visibility” function [5], ν , and associated “visible set” of rigid transformations, \mathcal{V} by

$$\nu(H) := (\mathbf{d} - \varrho \mathbf{r}_3) \cdot \mathbf{r}_3 \quad \text{and} \quad \mathcal{V} := \{H \in \text{SE}(3) : \nu(H) > 0\}. \quad (5)$$

Note that $\nu(H) > 0 \implies \|\mathbf{d}\| > \varrho$, i.e. $\mathbf{d} \in \mathcal{B} = \{\mathbf{d} \in \mathbb{R}^3 : \|\mathbf{d}\| > \varrho\}$.

The projection of $(b, \mathbf{a}) \in T\mathbb{E}^3$ to the image sphere is modeled by

$$T\pi : (b, \mathbf{a}) \mapsto (\pi(b), T_b\pi \cdot \mathbf{a}) \in TS^2.$$

We are not concerned with the length of the projection of \mathbf{a} , only the direction. Hence, consider the unit tangent map $T^1\pi$ represented in the camera frame by

$$\begin{aligned} T^1\pi : (b, \mathbf{a}) &\mapsto (\mathbf{q}_1, \mathbf{q}_2) \quad \text{where} \\ \mathbf{q}_1 &= \frac{\mathbf{b}^c}{\|\mathbf{b}^c\|} = \frac{\mathbf{d} - \varrho \mathbf{r}_3}{\|\mathbf{d} - \varrho \mathbf{r}_3\|} \quad \text{and} \quad \mathbf{q}_2 = \frac{\Gamma_{\mathbf{q}_1} \mathbf{a}^c}{\|\Gamma_{\mathbf{q}_1} \mathbf{a}^c\|} = \frac{\Gamma_{\mathbf{q}_1} \mathbf{r}_2}{\|\Gamma_{\mathbf{q}_1} \mathbf{r}_2\|} \end{aligned} \quad (6)$$

where $\Gamma_{\mathbf{q}_1} := (I - \mathbf{q}_1 \mathbf{q}_1^T)$.

Geometrically, \mathbf{q}_2 is a unit vector tangent to the image-sphere at the point \mathbf{q}_1 . The unit vectors \mathbf{q}_1 and \mathbf{q}_2 are mutually orthogonal. Consider the plane containing the camera origin o_c , the point b , and the vector \mathbf{a} . The unit vector

$$\mathbf{q}_3 = \mathbf{q}_1 \times \mathbf{q}_2$$

is normal to that plane. Thus, we define a function $c_2: \mathcal{V} \rightarrow \text{SO}(3)$

$$\boxed{c_2: H \mapsto [\mathbf{q}_1 \ \mathbf{q}_2 \ \mathbf{q}_3] = Q}, \quad (7)$$

identifying $T^1\text{S}^2$ with $\text{SO}(3)$.

2.4 Diffeomorphism to Image-space

Claim. The function $c: \mathcal{V} \rightarrow \mathcal{I}$, defined by

$$c(H) := (c_2(H), c_1(\mathbf{d})), \text{ where} \quad (8)$$

$$\mathcal{I} = \left\{ (Q, \lambda, \mathbf{s}) \in \text{SO}(3) \times (0, 1) \times \text{S}^2 : \mathbf{q}_1 \cdot \mathbf{s} > \sqrt{1 - \lambda^2} \right\} \quad (9)$$

$$\text{and } Q = [\mathbf{q}_1 \ \mathbf{q}_2 \ \mathbf{q}_3]$$

is a diffeomorphism, i.e. $\mathcal{V} \simeq \mathcal{I}$. \square

The proof is given in Appendix A

3 Image Jacobian

To be of practical application to VS we present a representation of the tangent map $Tc: T\mathcal{V} \rightarrow T\mathcal{I}$, its inverse Tc^{-1} , and the cotangent map $T^*c: T^*\mathcal{I} \rightarrow T^*\mathcal{V}$, with the following commutative diagram in mind:

$$\begin{array}{ccc} T\mathcal{V} & \xrightleftharpoons[Tc^{-1}]{Tc} & T\mathcal{I} \\ \downarrow & & \downarrow \\ \mathcal{V} & \xrightarrow{c} & \mathcal{I} \\ \uparrow & & \uparrow \\ T^*\mathcal{V} & \xrightleftharpoons[T^*c]{T^*c^{-1}} & T^*\mathcal{I} \end{array}$$

We make the following identification of the tangent space $T\text{SE}(3)$ of the Lie group $\text{SE}(3)$:¹

$$T\text{SE}(3) \simeq \text{SE}(3) \times \mathfrak{se}(3) \simeq \text{SE}(3) \times (\mathbb{R}^3 \otimes \mathbb{R}^3), \quad (10)$$

¹ The Lie algebra $\mathbb{R}^3 \otimes \mathbb{R}^3$ is $\mathbb{R}^3 \times \mathbb{R}^3$ with the Lie bracket structure found in [13].

where $\mathfrak{se}(3)$ is the Lie algebra of $\text{SE}(3)$. The identification occurs via “right translation,” i.e.

$$(H, \dot{H}) \mapsto (H, \dot{H}H^{-1}) \mapsto (H, (\boldsymbol{\omega}, \mathbf{v})) \quad (11)$$

where

$$H = \begin{bmatrix} R & \mathbf{d} \\ 0^T & 1 \end{bmatrix}, \quad \dot{H} = \begin{bmatrix} \dot{R} & \dot{\mathbf{d}} \\ 0^T & 0 \end{bmatrix} \quad \text{and} \quad \begin{cases} \boldsymbol{\omega} = (\dot{R}R^{-1})^\vee \\ \mathbf{v} = -\dot{R}R^{-1}\mathbf{d} + \dot{\mathbf{d}} \end{cases}$$

and the isomorphism $\mathbb{R}^3 \simeq \mathfrak{so}(3)$ is defined by

$$\hat{\cdot}: \begin{bmatrix} \boldsymbol{\omega}_1 \\ \boldsymbol{\omega}_2 \\ \boldsymbol{\omega}_3 \end{bmatrix} \mapsto \begin{bmatrix} 0 & -\boldsymbol{\omega}_3 & \boldsymbol{\omega}_2 \\ \boldsymbol{\omega}_3 & 0 & -\boldsymbol{\omega}_1 \\ -\boldsymbol{\omega}_2 & \boldsymbol{\omega}_1 & 0 \end{bmatrix}, \quad \vee: \begin{bmatrix} 0 & -\boldsymbol{\omega}_3 & \boldsymbol{\omega}_2 \\ \boldsymbol{\omega}_3 & 0 & -\boldsymbol{\omega}_1 \\ -\boldsymbol{\omega}_2 & \boldsymbol{\omega}_1 & 0 \end{bmatrix} \mapsto \begin{bmatrix} \boldsymbol{\omega}_1 \\ \boldsymbol{\omega}_2 \\ \boldsymbol{\omega}_3 \end{bmatrix},$$

where $\mathfrak{so}(3)$ is the Lie algebra of $\text{SO}(3)$. More detail can be found in, for example [13].

Similarly, for each $y = (Q, \lambda, \mathbf{s}) = \mathcal{I} \subset \text{SO}(3) \times (0, 1) \times \text{S}^2$, we have the following identification

$$T_y\mathcal{I} = T_Q\text{SO}(3) \times T_\lambda(0, 1) \times T_{\mathbf{s}}\text{S}^2 \simeq \mathbb{R}^3 \times \mathbb{R} \times T_{\mathbf{s}}\text{S}^2 \quad (12)$$

where we identify $T_Q\text{SO}(3)$ with $\mathfrak{so}(3) \simeq \mathbb{R}^3$, again via right translation

$$(Q, \dot{Q}) \mapsto (Q, \boldsymbol{\xi}) \quad \text{where} \quad \boldsymbol{\xi} = (\dot{Q}Q^{-1})^\vee. \quad (13)$$

Hence, to compute T_Hc we find the mapping relating the tangent space identifications made above in (10) and (12), namely

$$(H, (\boldsymbol{\omega}, \mathbf{v})) \mapsto (y, (\boldsymbol{\xi}, \lambda, \dot{\mathbf{s}}))$$

$$\text{where } y = (Q, \lambda, \mathbf{s}) = c(H), \quad \begin{bmatrix} \boldsymbol{\xi} \\ \lambda \\ \dot{\mathbf{s}} \end{bmatrix} = C(y) \begin{bmatrix} \boldsymbol{\omega} \\ \mathbf{v} \end{bmatrix}$$

and

$$C(y) := T_Hc|_{H=c^{-1}(y)} = \begin{bmatrix} I_{3 \times 3} & \frac{1}{\beta}(\delta \mathbf{q}_1 \mathbf{q}_3^T - \mathbf{q}_2 \mathbf{q}_3^T + \mathbf{q}_3 \mathbf{q}_2^T) \\ 0_{1 \times 3} & -\frac{\lambda^2}{\varrho} \mathbf{s}^T \\ -\hat{\mathbf{s}} & \frac{\lambda}{\varrho}(I_{3 \times 3} - \mathbf{s} \mathbf{s}^T) \end{bmatrix}, \quad (14)$$

where

$$\delta = \frac{\mathbf{s} \cdot \mathbf{q}_2}{\sqrt{\lambda^2 - \sin^2 \phi}}, \quad \beta = \frac{\varrho}{\lambda} \left(\cos \phi - \sqrt{\lambda^2 - \sin^2 \phi} \right),$$

$$\cos \phi = \mathbf{s} \cdot \mathbf{q}_1 \quad \text{and} \quad \sin \phi = \sqrt{1 - (\mathbf{s} \cdot \mathbf{q}_1)^2}.$$

The construction of C is straight forward. The details are given in [4].

To compute Tc^{-1} , and T^*c is now straight forward. Using the above representations, we have

$$T_y c^{-1} = (C(y)^T C(y))^{-1} C(y)^T \quad \text{and} \quad T_y^* c = C(y)^T. \quad (15)$$

Note that the expression for $T_y c^{-1}$ is *not* a pseudo-inverse. The possible confusion arises since the six dimensional tangent space $T_y \mathcal{I}$ is locally embedded in \mathbb{R}^7 . It should be noted that in many image-based visual servoing strategies employ the pseudo-inverse of the image Jacobian since the image feature points are treated as though moving freely in \mathbb{R}^n .

4 Controller

For the present work, we consider the case of so-called “eye-in-hand” VS, wherein the camera moves relative to the body which serves as an inertial reference frame. Let $(\boldsymbol{\Omega}, \mathbf{V})$ denote the angular and linear velocities, respectively, of the camera relative to the fixed, inertial body frame. Let $G = H^{-1}$ denote the transformation of the camera frame, \mathcal{F}_c , relative to the inertial body frame, \mathcal{F}_b . Note that

$$\begin{bmatrix} \widehat{\boldsymbol{\Omega}} & \mathbf{V} \\ 0^T & 0 \end{bmatrix} = G^{-1} \dot{G} = -\dot{H} H^{-1} = -\begin{bmatrix} \widehat{\boldsymbol{\omega}} & \mathbf{v} \\ 0^T & 0 \end{bmatrix}, \quad (16)$$

effectively mapping the identification of TSE(3) given by the right translation of \dot{H} in (10) and (11) to the left translation of $\dot{G} = \frac{d}{dt}(H^{-1})$. Note that this relationship clears up, once and for all, the kinematic distinction between “eye-in-hand” servoing and the so-called “fixed-camera” configuration, wherein the camera is fixed and the body is moving.

For simplicity, we posit a fully actuated purely kinematic plant model

$$\dot{G} = G \begin{bmatrix} \widehat{\boldsymbol{\Omega}} & \mathbf{V} \\ 0^T & 0 \end{bmatrix} \quad (17)$$

where we treat $(\boldsymbol{\Omega}, \mathbf{V}) \in \mathbb{R}^3 \otimes \mathbb{R}^3$ as control inputs. (We generalize this to a dynamical free rigid body in [4]). One possible control strategy involves planning a path $y_d(t) \in \mathcal{I}$ that moves from the initial configuration to the goal state and following the path via

$$\begin{bmatrix} \boldsymbol{\Omega} \\ \mathbf{V} \end{bmatrix} = -T_y c^{-1} \begin{bmatrix} \boldsymbol{\xi}_d \\ \dot{\lambda}_d \\ \dot{\mathbf{s}}_d \end{bmatrix} \quad (18)$$

where $[\boldsymbol{\xi}_d, \dot{\lambda}_d, \dot{\mathbf{s}}_d]^T$ is the desired velocity \dot{y}_d , expressed using the tangent space identification in (12). The minus sign in the above expression arises due to the identification made above in (16).

4.1 Visual Servoing via Navigation Functions

The diffeomorphism c , the visible set \mathcal{V} , and its relatively simple image \mathcal{I} , provide tremendous leverage into the VS problem. Given a desired configuration $G^* = (H^*)^{-1}$, measured through its image $y^* = (Q^*, \lambda^*, \mathbf{s}^*) = c(H^*)$, there are many possible image-based control strategies we can employ to achieve our objective of driving $G \rightarrow G^*$.

An open-loop strategy, such as the one above in (18), may be undesirable. However, the generation of \dot{y}_d can also be conceived as a feedback law, for example by using the method of Navigation Functions (NF's) [8,9,14]. A substantial benefit of using NF's is that they allow us to "lift" our kinematic controller to second order settings with little additional effort, while maintaining similar convergence guarantees (as we do for this problem in [4]). Moreover, these methods have already proven practicable for dynamic VS [5].

Let $\mathcal{D} \subset \mathcal{I}$ be compact "safe" domain. If we carefully design an artificial potential function $\varphi: \mathcal{D} \rightarrow [0, 1]$, then by letting

$$\dot{y}_d = -\nabla\varphi \quad (19)$$

the control law given by (18) drives G so that y converges to y^* , except for a set of measure zero. The following definition, adapted from [8], gives a set of conditions that guarantee essentially global convergence of the above controller (18), with \dot{y}_d given in (19).

Definition 1. Let \mathcal{D} be a smooth compact connected manifold with boundary, and $y^* \in \overset{\circ}{\mathcal{D}}$ be a point in its interior. A Morse function, $\varphi \in C^2[\mathcal{D}, [0, 1]]$ is called an *Navigation Function* if

1. φ takes its unique minimum at $\varphi(y^*) = 0$;
2. φ achieves its maximum of unity uniformly on the boundary, i.e. $\partial\mathcal{D} = \varphi^{-1}(1)$.

For any function satisfying the above definition, the controller given by (18) will ensure convergence $y \xrightarrow{t \rightarrow \infty} y^*$ from all initial conditions in \mathcal{D} . For more information, see [8].

4.2 Computing a Safe Domain and Navigation Function

The next step is to compute a compact domain $\mathcal{D} \subset \mathcal{I}$ that is "safe" with respect to the FOV of our camera system in the sense that if $G^{-1} = H \in c^{-1}(\mathcal{D})$ then all the necessary features are visible. To illustrate, we treat the FOV as a cone originating at the camera origin, with center along \mathbf{e}_3 , as shown in Figure 2. This cone reduces to a constraint on \mathbf{s} and λ , namely

$$f(y) := \lambda \mathbf{s} \cdot \mathbf{e}_3 - \sqrt{(1 - \lambda^2)(1 - (\mathbf{s} \cdot \mathbf{e}_3)^2)} \geq \cos \theta$$

where θ is the angle from \mathbf{e}_3 to the edge of the FOV cone. Additionally, we constrain $\lambda \in [\lambda_{\min}, \lambda_{\max}] \subset (0, 1)$ where the parameters λ_{\min} and λ_{\max}




Fig. 2. A simple model of the FOV cone.

effectively keep the camera from moving too far from or too close to the camera body, respectively. Finally, we keep \mathbf{q}_1 from being too close to the edge of the projected circle, namely $\mathbf{q}_1 \cdot \mathbf{s} + \epsilon \geq \sqrt{1 - \lambda^2}$. Putting these constraints together yields the compact manifold

$$\mathcal{D} = \{y = (Q, \lambda, \mathbf{s}) \in \text{SO}(3) \times [\lambda_{\min}, \lambda_{\max}] \times \text{S}^2 : f(y) \geq \cos \theta, \lambda_{\min} \leq \lambda \leq \lambda_{\max}\} \subset \mathcal{I}, \quad (20)$$

where $\theta \in (0, \pi/2)$, $0 < \lambda_{\min} < \lambda_{\max} < 1$.

Clearly $\mathcal{D} \subset \mathcal{I}$. Given this domain, one must construct an NF on \mathcal{D} . The construction of φ represents work in progress, however, we conjecture that given the relatively simple geometry of \mathcal{D} , that constructing a suitable NF should be straight forward. In fact, we believe (but have not yet formally shown) that $\mathcal{D} \simeq [0, 1]^5 \times \text{S}^1$ which is the same topology for which an NF has already been constructed for VS by the first author and colleagues [5].

5 Conclusion

In this paper, we presented a global diffeomorphism from a large subset of configurations in $\text{SE}(3)$ — those that are “visible” — to an appropriately defined image space. Such constructions provide tremendous leverage because

they shed light on the *geometry* of occlusion free servoing as well as provide a clear pathway to construct *global dynamical* visual servoing systems by using, for example, Navigation Functions.

A global, sensor-based representation of the configuration space leaves many open doors. For example, the control of underactuated and kinematically nonholonomic systems becomes possible in sensor space. Now that we now know it is possible to globally represent rigid motion using image coordinates, the next step is to construct a more general class of diffeomorphisms to the image plane that does not require designing special visual targets. We believe that with proper insight, the projection of a collection of rigidly connected feature points may be interpreted geometrically, again enabling a global representation of visible configurations. For example, perhaps depth can be described in terms of “moments”, as suggested by Hamel and Mahoney [6], and orientation can be described in terms of the projection of two or three feature points.

Acknowledgments

The first author was supported in part by DARPA/ONR under grants N00014-98-1-0747 and N66001-00-C8026, and NSF under grant ECS-9873474.

A Proof That $c: \mathcal{V} \rightarrow \mathcal{I}$ is a Diffeomorphism

The proof proceeds in four parts. First, we show that c is smooth on \mathcal{V} . Next we show that $c(\mathcal{V}) \subset \mathcal{I}$. Third, we show that c is bijective by explicitly computing its inverse, c^{-1} , on \mathcal{I} . Finally, we show that c^{-1} is smooth on \mathcal{I} .

The function c is smooth. The function c is composed of smooth functions away from the set where the arguments of $\|\cdot\|^{-1}$ become zero. But those arguments are nonzero on \mathcal{V} . In particular:

1. Equation (3) depends on $\|\mathbf{d}\|^{-1}$. However, $H \in \mathcal{V}$ implies $\|\mathbf{d}\| > \varrho$.
2. Equation (6) depends on $\|\mathbf{b}\|^{-1} = \|\mathbf{d} - \varrho\mathbf{r}_3\|^{-1}$. Visibility implies $\|\mathbf{d}\| > \varrho$, which in turn implies $\|\mathbf{d} - \varrho\mathbf{r}_3\| \geq \|\mathbf{d}\| - \|\varrho\mathbf{r}_3\| = \|\mathbf{d}\| - \varrho > 0$.
3. Equation (6) depends on $\|\Gamma_{\mathbf{q}_1}\mathbf{r}_2\|^{-1}$. This blows up iff $\mathbf{q}_1 = \pm\mathbf{r}_2$, i.e.

$$\mathbf{q}_1 = \frac{\mathbf{d} - \varrho\mathbf{r}_3}{\|\mathbf{d} - \varrho\mathbf{r}_3\|} = \pm\mathbf{r}_2$$

and hence, from (5)

$$\begin{aligned} \nu(H) &= (\mathbf{d} - \varrho\mathbf{r}_3) \cdot \mathbf{r}_3 = \pm \|\mathbf{d} - \varrho\mathbf{r}_3\| \mathbf{r}_2 \cdot \mathbf{r}_3 = 0, \\ &\implies H \notin \mathcal{V}. \end{aligned}$$

This contradiction implies that $\|\Gamma_{\mathbf{q}_1}R\mathbf{a}^b\| > 0$ for $H \in \mathcal{V}$.

Hence, c is smooth on \mathcal{V} .

The image of \mathcal{V} indeed is contained in \mathcal{I} . To see that $c(\mathcal{V}) \subset \mathcal{I}$, let $(Q, \lambda, \mathbf{s}) = c(H)$. By construction of c , we have that $(Q, \lambda, \mathbf{s}) \in \text{SO}(3) \times (0, 1) \times \text{S}^2$. To show that $\mathbf{q}_1 \cdot \mathbf{s} > \sqrt{1 - \lambda^2}$, note from (6) that

$$\mathbf{q}_1 \cdot \mathbf{s} = \frac{(\mathbf{d} - \varrho \mathbf{r}_3) \cdot \mathbf{s}}{\|\mathbf{d} - \varrho \mathbf{r}_3\|} = \frac{\|\mathbf{d}\| - \varrho \mathbf{s} \cdot \mathbf{r}_3}{\sqrt{\varrho^2 - 2\mathbf{d} \cdot \mathbf{r}_3 + \|\mathbf{d}\|^2}} = \frac{1 - \lambda\alpha}{\sqrt{1 - 2\lambda\alpha + \lambda^2}}$$

where $\alpha = \mathbf{s} \cdot \mathbf{r}_3$. From (5), $\alpha > \lambda$ to ensure $\nu(H) > 0$. Since the right hand side reaches its minimum of $\sqrt{1 - \lambda^2}$ for $\alpha = \lambda$, we have that $\mathbf{q}_1 \cdot \mathbf{s} > \sqrt{1 - \lambda^2}$.

The function c is bijective. Consider any $H \in \mathcal{V}$, with rotation $R = [\mathbf{r}_1 \ \mathbf{r}_2 \ \mathbf{r}_3]$ and translation \mathbf{d} as usual. Let $c(H) = (Q, \lambda, \mathbf{s}) \in \mathcal{I}$, where $Q = [\mathbf{q}_1 \ \mathbf{q}_2 \ \mathbf{q}_3]$. Given (λ, \mathbf{s}) , recovering the translation from (4) is trivial, namely $\mathbf{d} = c_1^{-1}(\lambda, \mathbf{s}) \in \mathcal{B}$. Recovering the rotation from (Q, λ, \mathbf{s}) requires a bit more care, but the result yields the unique inverse on \mathcal{I} (see [4] for details):

$$\begin{aligned} c^{-1}: (Q, \lambda, \mathbf{s}) &\mapsto H = \begin{bmatrix} \mathbf{r}_1 & \mathbf{r}_2 & \mathbf{r}_3 & \mathbf{d} \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \text{where } \mathbf{d} = \frac{\varrho}{\lambda} \mathbf{s}, \\ \mathbf{r}_3 &= \frac{1}{\lambda} (\mathbf{s} - \beta \mathbf{q}_1), \quad \beta = \|\mathbf{b}\| = \frac{\varrho}{\lambda} \left(\cos \phi - \sqrt{\lambda^2 - \sin^2 \phi} \right) \\ \mathbf{r}_2 &= \left(\mathbf{q}_2 - \frac{\mathbf{r}_3 \cdot \mathbf{q}_2}{\mathbf{r}_3 \cdot \mathbf{q}_1} \mathbf{q}_1 \right) \Big/ \left\| \mathbf{q}_2 - \frac{\mathbf{r}_3 \cdot \mathbf{q}_2}{\mathbf{r}_3 \cdot \mathbf{q}_1} \mathbf{q}_1 \right\|, \quad \text{and } \mathbf{r}_1 = \mathbf{r}_2 \times \mathbf{r}_3. \end{aligned} \quad (21)$$

The function c^{-1} is smooth. Finally, we need only show that c^{-1} is smooth. But, c^{-1} is composed of smooth functions. There are two caveats:

1. Equations involving $1/\lambda$. This is fine since $0 < \lambda < 1$.
2. Equation for \mathbf{r}_2 . First, note that $\mathbf{r}_3 \cdot \mathbf{q}_1 = \mathbf{r}_3 \cdot \mathbf{b} / \|\mathbf{b}\| = \nu(H) / \|\mathbf{b}\| > 0$. Also, since \mathbf{q}_2 and \mathbf{q}_1 are linearly independent, the denominator can never be zero, so this equation is smooth on \mathcal{I} .

Hence c^{-1} is smooth, and c is a diffeomorphism $c: \mathcal{V} \approx \mathcal{I}$. \square

References

1. Alessandro Chiuso, Roger Brockett, and Stefano Soatto. Optimal structure from motion: local ambiguities and global estimates. *International Journal of Computer Vision*, pages 195–228, September 2000.
2. Peter I. Corke and Seth A. Hutchinson. A new partitioned approach to image-based visual servo control. *IEEE Transactions on Robotics and Automation*, 17(4):507–515, 2001.
3. Noah Cowan. Binocular visual servoing with a limited field of view. In *Mathematical Theory of Networks and Systems*, Notre Dame, Indiana, August 2002.

4. Noah J. Cowan and Dong Eui Chang. Toward geometric visual servoing. Technical report, University of California, Berkeley, CA 94720, September 2002.
5. Noah J. Cowan, Joel D. Weingarten, and Daniel E. Koditschek. Visual servoing via navigation functions. *Transactions on Robotics and Automation*, 18(4), August 2002.
6. Tarek Hamel and Robert Mahony. Visual servong of an under-actuated dynamic rigid-body system: An image-based approach. *IEEE Transactions on Robotics and Automation*, 18(2):187–198, April 2002.
7. S. Hutchinson, G. D. Hager, and P. I. Corke. A tutorial on visual servo control. *IEEE Transactions on Robotics and Automation*, 12(5):651–670, October 1996.
8. Daniel E. Koditschek. The application of total energy as a Lyapunov function for mechanical control systems. In *Dynamics and control of multibody systems (Brunswick, ME, 1988)*, pages 131–157. Amer. Math. Soc., Providence, RI, 1989.
9. Daniel E. Koditschek and Elon Rimon. Robot navigation functions on manifolds with boundary. *Advances in Applied Mathematics*, 11:412–442, 1990.
10. H. C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293:133–135, 1981.
11. Chien-Ping Lue, Gregory D. Hager, and Eric Mjolsness. Fast and globally convergent pose estimation from video images. *Transactions on Pattern Analysis and Machine Intelligence*, 22(6):610–622, June 2000.
12. Ezio Malis, Francois Chaumette, and Sylvie Boudet. Theoretical improvements in the stability analysis of a new class of model-free visual servoing methods. *IEEE Transactions on Robotics and Automation*, 18(2):176–186, April 2002.
13. Jerrold E. Marsden and Tudor S. Ratiu. *Introduction to Mechanics and Symmetry*, chapter 14.7. Springer, 2 edition, 1999.
14. Elon Rimon and Daniel E. Koditschek. Exact robot navigation using artificial potential fields. *IEEE Transactions on Robotics and Automation*, 8(5):501–518, Oct 1992.
15. Camillo J. Taylor and James P. Ostrowski. Robust vision-based pose control. In *International Conference on Robotics and Automation*, volume 3, pages 2734–2740, San Francisco, CA, 2000. IEEE.
16. Hong Zhang and James P. Ostrowski. Visual servoing with dynamics: Control of an unmanned blimp. In *International Conf. on Robotics and Automation*, 1999.
17. Hong Zhang and Jim Ostrowski. Visual motion planning for mobile robots. *IEEE Transactions on Robotics and Automation*, 18(2):199–208, April 2002.