Chirality in Multiview Geometry

Andrew Pryhuber

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2021

Reading Committee:

Rekha R. Thomas, Chair

Sameer Agarwal

Dmitriy Drusvyatskiy

Program Authorized to Offer Degree:
Mathematics

University of Washington

Abstract

Chirality in Multiview Geometry

Andrew Pryhuber

Chair of the Supervisory Committee:
Professor Rekha R. Thomas
Mathematics

This thesis studies mathematical problems associated with reconstructing a three dimensional scene from images. Using the traditional pinhole camera model and tools from multiview geometry, we pose these problems from an algebraic perspective. We then build on this model to respect the physical constraint that cameras can only see points in front of them. We explore this constraint, known as *chirality*, using techniques from topology, real algebraic geometry, and optimization. Our results provide a complete semialgebraic description of the set of world points in front of an arrangement of cameras. We use this to classify when a set of point pairs is the image of a scene in front of two cameras.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

Thank you to my advisers Rekha Thomas and Sameer Agarwal for their mentoring. Your guidance and encouragement kept me going and made this thesis possible. Thank you also to Rainer Sinn, Lucas Van Meter, and Max Lieblich for many invaluable discussions.

I am grateful to my friends outside of math who supported me and kept me going throughout my PhD. Thanks for getting me out on a bike ride when I most needed it.

To my parents, thank you for believing in me and teaching me that I can accomplish whatever I put my mind to.

Chapter 1

INTRODUCTION

A fundamental problem in computer vision is *3D reconstruction*, which involves understanding a three-dimensional scene from two-dimensional images of the scene. The field of *multiview geometry* studies this problem from a mathematical point of view, drawing from linear algebra and classical projective geometry [19, 23, 46, 49]. Recent work has applied tools of algebraic geometry to give a new perspective on various aspects of 3D reconstruction [1, 4, 5, 27, 37]. A drawback of the algebraic approach is that it ignores the physical constraint that cameras can only image points that lie in front of them. This constraint is referred to as *chirality* [24]. This thesis will explore chirality and how it fits into the mathematical model of imaging. First, we recall the building blocks of this model.

Definition 1.0.1. *A* ***(pinhole) camera*** *is a full-rank linear map* $\mathbb{P}^3_{\mathbb{R}} \dashrightarrow \mathbb{P}^2_{\mathbb{R}}$ *represented by a* 3×4 *matrix* A*, up to scale. The camera is defined on the open set* $\mathbb{P}^3_{\mathbb{R}} \setminus \{\mathbf{c}_A\}$ *where the* ***center*** *of* A*,* $\mathbf{c}_A$*, spans the one-dimensional kernel of* A*.*

Going forward, we will refer to the matrix representation A itself as a camera. We say the domain $\mathbb{P}^3_{\mathbb{R}}$ is the set of *world points* while the codomain $\mathbb{P}^2_{\mathbb{R}}$ is the *image plane* of A. We say a world point $\mathbf{q} = (q_1, q_2, q_3, q_4)$ is *finite* if and only if its fourth coordinate q_4 is non-zero. The dehomogenization of a finite point is denoted as $\widetilde{\mathbf{q}} := (q_1/q_4, q_2/q_4, q_3/q_4)^\top$. We will generally be interested in those cameras whose centers are finite points, which we call *finite cameras.* A finite camera matrix A can be decomposed as $A = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}$ where $G \in \mathrm{GL}_3$ and $\mathbf{t} \in \mathbb{R}^3$. Geometrically, a camera A sends a world point $\mathbf{q} \in \mathbb{P}^3_{\mathbb{R}}$ to the intersection of the affine line spanned by $\mathbf{c}_A$ and $\mathbf{q}$ with the image plane of A. Figure 1.1 depicts this for a finite camera and finite world point.

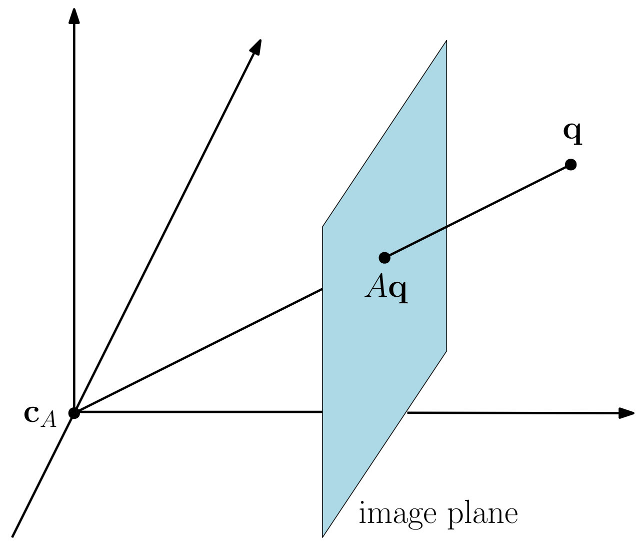

Figure 1.1: Image of world point $\mathbf{q}$ in the camera A.

A fundamental object we study is a *projective reconstruction* of a set of point correspondences $\mathcal{P} = \{(\mathbf{p}_{1k}, \ldots, \mathbf{p}_{mk}) : k = 1, \ldots, n\} \subseteq (\mathbb{P}^2_{\mathbb{R}})^m$. A projective reconstruction of $\mathcal{P}$ is an arrangement of cameras $\mathcal{A} = (A_1, \ldots, A_m)$ and a set of world points $\mathcal{Q} = (\mathbf{q}_1, \ldots, \mathbf{q}_n)$ such that $\mathbf{p}_{ij}$ is the image of $\mathbf{q}_j$ in the camera A_i. Further discussion of projective reconstructions can be found in [23, Chapters 9 & 10].

This thesis consists of three papers. We summarize the results of these papers in the following three sections. In Section 1.1, we give an algebraic description of the image of all world points under a camera arrangement. In Section 1.2, we discuss how to constrain world points to be in front of cameras, giving rise to the notion of chirality. This allows us to define a more restrictive type of projective reconstruction, called a *chiral reconstruction*. We discuss when a projective reconstruction can be transformed into a chiral reconstruction. In Section 1.3, we give a complete answer to when a set of point pairs has a chiral reconstruction.

1.1 *Ideals of the Multiview Variety*

Chapter 2 comes from a paper co-authored with Sameer Agarwal and Rekha Thomas, which appears in *IEEE Transactions on Pattern Analysis and Machine Intelligence* [3].

1.1.1 Background

This work was motivated by the *triangulation* problem, which arises in 3D reconstruction. This problem involves reconstructing a world point from noisy image observations in a known arrangement of cameras. To answer this question, we first consider a noiseless version. That is, given an arrangement of m cameras $\mathcal{A} = (A_1, \ldots, A_m)$, we want to know when a point in $(\mathbb{P}^2_{\mathbb{R}})^m$ is the image of a world point in $\mathbb{P}^3_{\mathbb{R}}$ under $\mathcal{A}$. To that end, consider the image formation map of the arrangement $\mathcal{A} := (A_1, \ldots, A_m)$

$$\varphi_{\mathcal{A}} : \mathbb{P}^3_{\mathbb{R}} \dashrightarrow (\mathbb{P}^2_{\mathbb{R}})^m,$$

which sends a world point $\mathbf{q} \in \mathbb{P}^3_{\mathbb{R}}$ to its images $(\mathbf{p}_1 = A_1\mathbf{q}, \ldots, \mathbf{p}_m = A_m\mathbf{q}) \in (\mathbb{P}^2_{\mathbb{R}})^m$. The image of all world points $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$ is called the *joint image* [50] of $\mathcal{A}$. The Zariski closure of the joint image is called the *multiview variety* of $\mathcal{A}$. With the multiview variety description in hand, the triangulation problem can be posed geometrically as minimizing the Euclidean distance from a noisy observation in $(\mathbb{P}^2_{\mathbb{R}})^m$ to the multiview variety [5].

1.1.2 Results

We give a complete ideal-theoretic description of the multiview variety. The ideal of the multiview variety of $\mathcal{A}$, denoted by $M_{\mathcal{A}}$, is called the *multiview ideal* of $\mathcal{A}$.

Past research has focused on various systems of polynomials that vanish on $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$. In the computer vision literature these are known as *multiview constraints* [18, 23, 27, 35, 37]. The ideals generated by these systems of polynomials are contained in $M_{\mathcal{A}}$. However, there hasn't been much discussion of whether these polynomials generate $M_{\mathcal{A}}$ since the focus of all these papers has been on the multiview variety and not its vanishing ideal. It can be difficult to determine the vanishing ideal of a variety; however, there are various advantages to knowing it. In particular, to be able to do any computations with a variety or to study its structure using algebra, we need a description in terms of polynomials and the vanishing ideal is the optimal algebraic description.

We show that $M_{\mathcal{A}}$ is generated by *k-focal polynomials*. These polynomials are determinantal and occur as maximal minors of partially symbolic $3k \times (4+k)$ matrices

$$\mathcal{A}(p) = \begin{pmatrix} A_1 & p_1 & 0 & \dots & 0 \\ A_2 & 0 & p_2 & \ddots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ A_k & 0 & \dots & 0 & p_k \end{pmatrix} \tag{1.1}$$

where p_i represents the tuple of variables (x_i, y_i, z_i), denoting the coordinates on the ith image plane. Each k-focal polynomial is multilinear in the entries of the cameras A_i and hence is of total degree k. We show that $M_{\mathcal{A}}$ is generated by the set of all 2-focal and 3-focal polynomials among cameras in $\mathcal{A}$.

We discuss three different systems of polynomials proposed to cut out the multiview variety by Heyden-Åström [27], Faugeras *et al.* [18] and Ma *et al.* [35]. These systems consist of homogeneous polynomials of total degree n, four and three, respectively. We show that while the ideals they generate are all contained in $M_{\mathcal{A}}$, none of them actually coincide with $M_{\mathcal{A}}$. We establish precise algebraic relationships between these ideals and $M_{\mathcal{A}}$, using the characterization of $M_{\mathcal{A}}$ in terms of k-focal polynomials above.

The multiview ideal has been studied in detail before in [5]. This work makes an algebraic assumption on the camera matrices $A_1, \dots, A_m$ in $\mathcal{A}$ called *minor-genericity*. For minor-generic camera arrangements $\mathcal{A}$, the set of all 2-focal, 3-focal, and 4-focal polynomials is shown to form a universal Gröbner basis of $M_{\mathcal{A}}$. Our results rely on the less restrictive assumption that the camera centers of $\mathcal{A}$ are pairwise distinct. Under this assumption, we show that the 2-focal and 3-focal polynomials alone form a generating set of $M_{\mathcal{A}}$.

When camera centers in $\mathcal{A}$ are assumed to be non-coplanar, we achieve an even simpler description of the multiview ideal. In this case, we show that only the 2-focal polynomials are needed to generate $M_{\mathcal{A}}$. This provides justification for the tightness of the relaxation of the triangulation problem described in [4]. In that work, the multiview variety is approximated by imposing only epipolar constraints, which are exactly our 2-focal polynomials. Our work shows that, whenever camera centers are non-coplanar, the epipolar constraints alone exactly describe the multiview variety.

1.2 The Chiral Domain of a Camera Arrangement

Chapter 3 comes from a paper co-authored with Sameer Agarwal, Rainer Sinn, and Rekha Thomas, and is currently under review for the *Journal of Mathematical Imaging and Vision* [2].

1.2.1 Background

We introduce the concept of *chirality* for arrangements of finite cameras. First, we recall some basic definitions needed for our results. The *principal plane* of a finite camera $A = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}$ is the hyperplane $L_A := \{\mathbf{q} \in \mathbb{P}^3 \,:\, A_{3,\bullet}\mathbf{q} = 0\}$, where $A_{3,\bullet}$ is the third row of A, i.e. it is the set of points in $\mathbb{P}^3$ that image to infinite points in $\mathbb{P}^2$. Note that the camera center $\mathbf{c}_A$ lies on L_A. We regard L_A as an oriented hyperplane in $\mathbb{R}^4$ with normal vector $\mathbf{n}_A := \det(G)A_{3\bullet}^\top$, which we call the *principal ray* of A. See Figure 1.2 for a dehomogenized picture in $\mathbb{R}^3$.

The *depth* of a finite point $\mathbf{q}$ in a finite camera A is essentially the projection of the dehomogenized vector $\widetilde{\mathbf{q}} - \widetilde{\mathbf{c}_A}$ along the principal ray, see [23]. Formally, we define it as

$$\mathrm{depth}(\mathbf{q}; A) := \left(\frac{1}{|\det(G)| \|G_{3,\bullet}\|}\right) \frac{(\mathbf{n}_A^\top \mathbf{q})}{(\mathbf{n}_\infty^\top \mathbf{q})}. \tag{1.2}$$

We say that a finite point $\mathbf{q}$ is *in front of* the camera A if $\mathrm{depth}(\mathbf{q}; A) > 0$, see [23, 24]. Since only the sign of $\mathrm{depth}(\mathbf{q}; A)$ matters, we refer to this sign as the *chirality* of $\mathbf{q}$ in A, denoted as $\chi(\mathbf{q}; A)$, which is either 1 or -1. This definition does not immediately generalize to infinite points or to an arrangement of multiple cameras. Both of these technicalities are dealt with below by taking closures.

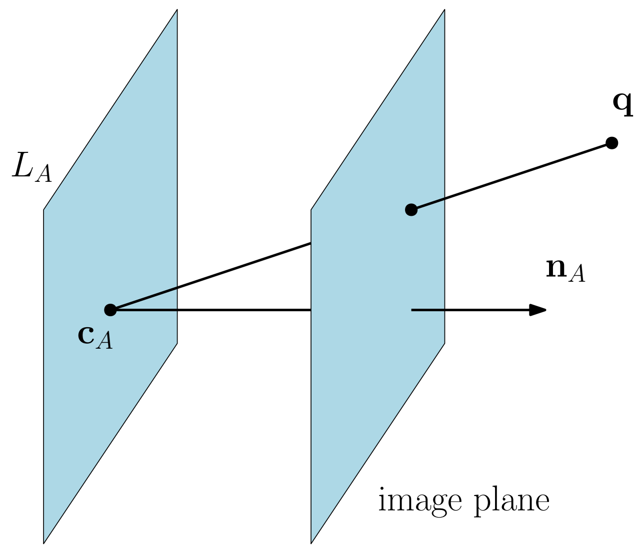

Figure 1.2: Principal plane and principal ray of a finite camera A.

1.2.2 Results

We derive the chiral analogue of the multiview variety of $\mathcal{A}$. While the joint image considers the image of all of $\mathbb{P}^3_{\mathbb{R}}$ under $\mathcal{A}$, here we restrict to only the world points in front of $\mathcal{A}$. To make this precise, we first need a notion of *in front of* an arrangement of multiple cameras. We make the following definition.

Definition 1.2.1. *Let $\mathcal{A}$ be an arrangement of finite cameras. Then the* ***chiral domain*** *of $\mathcal{A}$, denoted as $D_{\mathcal{A}}$ is the (Euclidean) closure of the set*

$$\{\mathbf{q} \in \mathbb{P}^3_{\mathbb{R}} \,|\, \mathbf{q} \text{ finite},\ \ \forall A \in \mathcal{A},\ \operatorname{depth}(\mathbf{q}; A) > 0\}.$$

Moreover, a point $\mathbf{q} \in \mathbb{P}^3$ is said to have chirality 1 with respect to $\mathcal{A}$*, denoted as $\chi(\mathbf{q}; \mathcal{A}) = 1$, if and only if $\mathbf{q} \in D_{\mathcal{A}}$. In this case, we say $\mathbf{q}$ is in front of $\mathcal{A}$.*

This generalizes the classical definition of chirality of finite points in a single camera to apply to all of $\mathbb{P}^3_{\mathbb{R}}$ in an arrangement of cameras. Using the chiral domain, we then define the *chiral joint image* of $\mathcal{A}$ to be the set $\varphi_{\mathcal{A}}(D_{\mathcal{A}})$. We give a complete semialgebraic description of the chiral joint image of $\mathcal{A}$ using the coordinates on the product of image planes $(\mathbb{P}^2_{\mathbb{R}})^m$.

The notion of the chiral domain also allows us to sharpen the notion of projective reconstruction discussed above. We say that a projective reconstruction $(\mathcal{A}, \mathcal{Q})$ is a *chiral reconstruction* if it satisfies the additional condition that $\mathbf{q}_j \in D_{\mathcal{A}}$ for all j, i.e., all world points are in front of the camera arrangement. Observe that the definition of chirality above implicitly depends on the choice of coordinates on $\mathbb{P}^3_{\mathbb{R}}$. Therefore, it is natural to ask whether some choice of coordinates transforms a given projective reconstruction into a chiral reconstruction. In the language of classical projective geometry, this translates to asking whether there exists a *homography* of the world $\mathbb{P}^3_{\mathbb{R}}$, which transforms a projective reconstruction into a chiral one. The effect of a homography on chirality is completely determined by which hyperplane the homography sends to the hyperplane at infinity. Therefore the search for a chiral reconstruction in is equivalent to finding a hyperplane satisfying certain separation conditions with respect to world points and camera centers.

Hartley discusses these conditions for reconstructions with arbitrarily many views in [23]. We show his conditions always reduce to checking whether two polyhedral cones have nonempty intersection (Theorem 3.5.1). This is equivalent to a linear programming feasibility problem and hence can be solved in a computationally efficient way. Our method provides a constructive algorithm to find the homography necessary to achieve a chiral reconstruction.

Hartley also presents a sufficient condition for a two view chiral reconstruction to exist [24]. Restating this condition in terms of the principal rays of the given cameras, we give a simpler proof of Hartley's result using polyhedral geometry Theorem 3.5.3.

1.3 Existence of Two View Chiral Reconstructions

Chapter 4 comes from a paper co-authored with Rainer Sinn and Rekha Thomas, and is currently under review for the *SIAM Journal on Applied Algebra and Geometry* [42].

1.3.1 Background

In Chapter 3, we described the set of images of world points that lie in front of a given camera arrangement. Further, we discussed when a given non-chiral reconstruction be made chiral.

In Chapter 4, we consider the problem of deciding if a set of point pairs $\mathcal{P} = \{(\mathbf{u}_i, \mathbf{v}_i) : i = 1, \ldots, k\}$ has a chiral reconstruction $(A_1, A_2, \mathcal{Q})$. It is well known that the existence of a two view projective reconstruction is intimately related to the existence of *fundamental matrices* [1, 23]. Recall that a fundamental matrix of $\mathcal{P}$ is a real, rank two, 3×3 matrix X which satisfies the epipolar equations

$$\mathbf{v}_i^\top X \mathbf{u}_i = 0, i = 1, \ldots, k.$$

Every fundamental matrix X is rank deficient, hence it satisfies $\det(X) = 0$. Intersecting the zero set of $\det(X)$ with the zero sets of the k epipolar equations imposed by the point pairs $\mathcal{P}$, we get an algebraic variety called the *epipolar variety* of $\mathcal{P}$. A necessary condition for the existence of a projective reconstruction of $\mathcal{P}$ is the existence of a real rank two matrix X in the epipolar variety of $\mathcal{P}$

It is well-known that a fundamental matrix determines a projective reconstruction $(A_1, A_2, \mathcal{Q})$ up to projective equivalence, i.e., up to a choice of $\mathbb{P}^3_\mathbb{R}$ coordinates [23]. Hence, there is an entire orbit of projective reconstructions $(A_1H^{-1}, A_2H^{-1}, H\mathcal{Q})$ for $H \in \mathrm{GL}_4$ associated to a given fundamental matrix. To address the question of when a two view chiral reconstruction exists, we must determine if any of the projective reconstructions in this orbit is a chiral reconstruction.

1.3.2 Results

In Chapter 3, we present necessary and sufficient conditions for a two view projective reconstruction $(A_1, A_2, \mathcal{Q})$ to be projectively equivalent to a chiral reconstruction. In Chapter 4, we show that, in fact, these conditions can be stated in the form of polynomial inequalities on the fundamental matrix associated to A_1 and A_2 (Theorem 4.4.1). We call these the *chiral epipolar inequalities.* These inequalities precisely determine if there is a chiral reconstruction in the orbit over a fundamental matrix. The existence of a two view chiral reconstruction of $\mathcal{P}$ is therefore equivalent to the non-emptiness of the semialgebraic set in the epipolar variety of $\mathcal{P}$. This set, called the *chiral epipolar region* of $\mathcal{P}$, is defined by the chiral epipolar inequalities along with an additional regularity condition.

Our work builds on the results of [1], which classifies the sets of point pairs which have a fundamental matrix. To have a chiral reconstruction, we need more than a fundamental matrix; we need a fundamental matrix in the chiral epipolar region of $\mathcal{P}$. We develop tools to study this region in $\mathbb{P}^8_{\mathbb{R}}$ geometrically by considering its image under the adjoint map $\mathbb{P}^8_{\mathbb{R}} \dashrightarrow \mathbb{P}^8_{\mathbb{R}}$ sending $X \mapsto \mathrm{adj}(X)$. Using this machinery, we provide a complete classification of sets of point pairs for which a chiral reconstruction exists. In particular, we show that:

1. Up to three point pairs always have a chiral reconstruction.

2. Four point pairs have a chiral reconstruction unless they belong to two non-generic combinatorial types, in which case they may or may not.

3. For generic sets of five point pairs, the chiral epipolar region is bounded by line segments in a Schläfli double six on a cubic surface with 27 real lines.

4. Five or more point pairs do not have a chiral reconstruction with positive probability.

Chapter 2

IDEALS OF THE MULTIVIEW VARIETY

2.1 Introduction

A general projective camera is a rank three matrix in $\mathbb{R}^{3\times 4}$. Given a camera arrangement $\mathcal{A} = (A_1, \ldots, A_n)$, the image formation map

$$\varphi_{\mathcal{A}} \,:\, \mathbb{P}^3_{\mathbb{R}} \dashrightarrow (\mathbb{P}^2_{\mathbb{R}})^n$$

sends a homogenized world point $\mathbf{q} \in \mathbb{P}^3_{\mathbb{R}}$ to its images $(\mathbf{p}_1 = A_1\mathbf{q}, \ldots, \mathbf{p}_n = A_n\mathbf{q}) \in (\mathbb{P}^2_{\mathbb{R}})^n$. The ith copy of $\mathbb{P}^2_{\mathbb{R}}$ in the codomain of $\varphi_{\mathcal{A}}$ is the homogenized image plane of camera i. The unique point $\mathbf{c}_i \in \mathbb{P}^3_{\mathbb{R}}$ in the kernel of A_i is the focal point of camera i. The map $\varphi_{\mathcal{A}}$ is defined at all points in $\mathbb{P}^3_{\mathbb{R}}$ except at the foci $\mathbf{c}_1, \ldots, \mathbf{c}_n$. Triggs called $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$ the *joint image* [50] and Heyden-Åström call it the *natural descriptor* [27]. We are interested in studying the complete set of polynomials that vanish on $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$.

Definition 2.1.1. *Given a set $S \subseteq \mathbb{P}^{d-1}_{\mathbb{C}}$, the collection of all polynomials in $\mathbb{C}[x_1, \ldots, x_d]$ that vanish on S is a homogeneous ideal, known as the* ***vanishing ideal*** *of S, and denoted as $\mathbf{I}(S)$. The variety $\mathbf{V}(\mathbf{I}(S))$ is the the smallest complex projective variety that contains S, known as the* ***Zariski closure*** *of S.*

We refer the reader to [14] for the basics on ideals and varieties. In this paper we will be interested in the vanishing ideal of the joint image $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$.

Definition 2.1.2. *The* ***multiview ideal*** *of $\mathcal{A}$, denoted $M_{\mathcal{A}}$, is the vanishing ideal of $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$ in $\mathbb{C}[p_1, \ldots, p_n]$ where $p_i = (x_i, y_i, z_i)$ are the coordinates on the ith copy of $\mathbb{P}^2_{\mathbb{C}}$. The Zariski closure of $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$ in $(\mathbb{P}^2_{\mathbb{C}})^n$ is the complex projective variety $\mathbf{V}(M_{\mathcal{A}})$, which we call the* ***multiview variety*** *of $\mathcal{A}$.*

The terminology *multiview ideal* and *multiview variety* comes from [5]. Following Triggs [50], Trager *et al.* refer to the multiview variety as the *joint image variety.*

Starting with the seminal work of Longuet-Higgins [34], researchers have studied various systems of polynomials that vanish on $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$. In the computer vision literature these equations are known as *multiview constraints* [18, 23, 27, 35, 37]. Obviously, the ideals generated by these systems of polynomials are contained in $M_{\mathcal{A}}$. However, there hasn't been much discussion of whether these polynomials generate $M_{\mathcal{A}}$ since the focus of all these papers has been on the multiview variety and not its vanishing ideal. The aim of this paper is to provide a complete description of the multiview ideal and study its relationship to the above sets of polynomials.

It can be difficult to determine the vanishing ideal of a variety. However, there are various advantages to knowing it. To be able to do any computations with a variety or to study its structure using algebra, we need a description in terms of polynomials and the vanishing ideal is the optimal algebraic description. This manifests itself in a number of ways.

The set of all polynomial functions on X is precisely $\mathbb{C}[x_1, \ldots, x_d]/\mathbf{I}(X)$, known as the coordinate ring of X. In particular, a polynomial g vanishes on X if and only if g belongs to $\mathbf{I}(X)$. Knowledge of a generating set $\{g_1, \ldots, g_k\}$ of $\mathbf{I}(X)$ also informs us about the local structure of X, since a point $x \in X$ is *smooth* if and only if the Jacobian matrix $(\frac{\partial g_i}{\partial x_j})$ has rank equal to the codimension of X. More generally, if $X \subset \mathbb{P}^{d-1}_{\mathbb{C}}$ is a projective variety then $\mathbf{I}(X)$ carries all the geometric information about X allowing algebra (and algebraic algorithms) to infer geometric properties of X. For example, the dimension and degree of X can be read off from the *Hilbert polynomial* of $\mathbf{I}(X)$ which also carries many more sophisticated invariants of X. See [14] for all the above.

In multiview geometry, many estimation problems can be phrased as polynomial optimization problems over varieties [5, 29]. In particular, the triangulation problem under Gaussian noise amounts to projecting a point onto the multiview variety[4].

In general, polynomial optimization on a variety $X \subseteq \mathbb{R}^n$ boils down to certifying the non-negativity of a polynomial f on X by expressing it as a sum-of-squares (sos) modulo an ideal J vanishing on X [7]. This means finding a sos polynomial $s = \sum p_i^2$ such that $f - s$ lies in J. This expressibility is maximized, and the algorithms terminate in the lowest possible degree, when $J = \mathbf{I}(X)$. We illustrate this on a very small example.

Example 2.1.1. *The polynomial $x + 1$ is non-negative on $X = \{0\} \subset \mathbb{R}$. The ideal $\langle x^2 \rangle$ cuts out X but $\mathbf{I}(X) = \langle x \rangle$. Now $(x + 1) - 1 \in \langle x \rangle$ allowing $s = 1$ as the sos certificate. On the other hand, if $x + 1 - s \in \langle x^2 \rangle$ then s has to have degree at least 2; for instance $(x + 1) - (1 + \frac{1}{2}x)^2 \in \langle x^2 \rangle$.*

The above phenomenon can have a major impact on the number of rounds of convex relaxations needed to solve a polynomial optimization problem such as the well-known *Lasserre/sos hierarchies* [30, 41], where each round looks for sos certificates of a fixed degree with degrees increasing monotonically with rounds. In each round the semidefinite program being solved is of size $O(n^d)$, where n is the number of variables and d is degree in that round. As a result, in many cases only the first round maybe computationally feasible and having access to $\mathbf{I}(X)$ can make the difference between the problem being tractable or not.

The rest of the paper is structured as follows. After a brief discussion of the notation used in this paper we begin in Section 2.2 by introducing a family of ideals associated with every camera arrangement $\mathcal{A}$ which we call the k-focal ideals. We describe how these ideals behave under change of coordinates, and dispel the popular myth that, under a change of image coordinates, k-focal polynomials go to k-focal polynomials. In Section 2.3, we prove our first main theorem (Theorem 2.3.1), that the well-known bifocal (epipolar constraints) and trifocal polynomials generate $M_{\mathcal{A}}$ when the camera foci in $\mathcal{A}$ are distinct. Next, in Section 2.4, we consider three different types of determinantal polynomials proposed to cut out the multiview variety by Heyden-Åström [27], Faugeras *et al.* [18] and Ma *et al.* [35]. We show that while the ideals they generate are all contained in $M_{\mathcal{A}}$, none of them actually coincide with $M_{\mathcal{A}}$. We establish their precise algebraic relationship with $M_{\mathcal{A}}$. In Section 2.5, we consider the relationship of the multiview ideal to bifocal polynomials and prove the algebraic analog of the statement that the bifocal polynomials cut out the multiview variety when the camera foci are noncoplanar. In Section 2.6, we study how the various ideals relate to each other when we restrict our attention to finite images, *i.e.* exclude points at infinity. We conclude in Section 2.7 with a summary.

Many results in this paper require explicit computation. We recommend the reader have a copy of Macaulay2 [21] (or equivalent symbolic algebra software) handy. The Macaulay2 codes for our computations can be found at `https://sites.math.washington.edu/~thomas/papers/Multiview_Ideal.zip`

2.1.1 Notation

In the rest of the paper, we will use $\mathbb{P}$ to denote $\mathbb{P}_{\mathbb{C}}$. The ideal generated by the polynomials $f_1, \dots, f_s$ will be denoted as $\langle f_1, \dots, f_s \rangle$.

We will use A for cameras and G for matrices in GL_n. $\mathcal{A}$ and $\mathcal{G}$ will denote arrangements of corresponding matrices. Bold, lower-case roman letters will be used to indicate vectors, and lower-case greek letters will be used for functions. Given a partial symbolic matrix M, $minors(k, M)$ will denote the ideal generated by all $k \times k$ minors of the matrix M. The symbol $[n]$ denotes the set $\{1, \dots, n\}$ and $\binom{[n]}{m}$ denotes the set of all size m subsets of $[n]$.

2.2 The k-focal ideals of a camera arrangement

Let p_i be the tuple of variables (x_i, y_i, z_i) denoting the coordinates associated to the projective plane $\mathbb{P}^2_{\mathbb{R}}$ corresponding to the ith camera image. Write $p = (p_1, \dots, p_n)$, and consider the partially symbolic matrix

$$\mathcal{A}(p) := \begin{bmatrix} A_1 & p_1 & & & \\ A_2 & & p_2 & & \\ \vdots & & & \ddots & \\ A_n & & & & p_n \end{bmatrix}. \tag{2.1}$$

Let $\mathcal{A}(\mathbf{p})$ denote the evaluation of $\mathcal{A}(p)$ at $p = \mathbf{p}$. If $\mathbf{p} := (\mathbf{p}_1, \dots, \mathbf{p}_n) \in \varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$ then there exists some $\mathbf{q} \in \mathbb{P}^3_{\mathbb{R}}$ and scalars $\lambda_i \in \mathbb{R}$ such that $A_i \mathbf{q} = \lambda_i \mathbf{p}_i$ for all $i = 1, \dots, n$.

Therefore, $\mathcal{A}(\mathbf{p})$ has a non-trivial kernel since it contains the point $(\mathbf{q}, -\lambda_1, \ldots, -\lambda_n)$, and hence the maximal minors of $\mathcal{A}(p)$, which are polynomials in $p_1, \ldots, p_n$, vanish on $\mathbf{p}$. Since $\mathbf{p}$ was arbitrary, these maximal minors vanish on all of $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$ and on the multiview variety. Therefore,

$$minors(4+n, \mathcal{A}(p)) \subseteq M_{\mathcal{A}}.$$

In this section, we describe further minors of $\mathcal{A}(p)$ and the ideals they generate, which will play an important role in the description of $M_{\mathcal{A}}$.

Definition 2.2.1. *For a subset $\sigma = \{\sigma_1, \ldots, \sigma_k\} \subseteq [n]$ where $k \geq 2$, consider the partially symbolic matrix*

$$\mathcal{A}_\sigma(p) = \begin{pmatrix} A_{\sigma_1} & p_{\sigma_1} & 0 & \ldots & 0 \\ A_{\sigma_2} & 0 & p_{\sigma_2} & \ddots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ A_{\sigma_k} & 0 & \ldots & 0 & p_{\sigma_k} \end{pmatrix} \tag{2.2}$$

*of size $3k \times (4+k)$. A maximal $(4+k)\times(4+k)$ minor of $\mathcal{A}_\sigma(p)$ is called a k-**focal polynomial** of $\mathcal{A}$. The k-**focal ideal** of $\mathcal{A}$, $H^k_{\mathcal{A}}$, is the ideal sum*

$$H^k_{\mathcal{A}} = \sum_{\sigma \in \binom{[n]}{k}} minors(4+k, \mathcal{A}_\sigma(p)).$$

Trager *et al.* also study the k-focal polynomials and refer to them as *k-linearities* [46, 47]. Note that every k-focal polynomial is multilinear and of total degree k. Such a minor involves choosing $4+k$ rows of $\mathcal{A}_\sigma(p)$, and by a pigeonhole argument, at most four cameras may contribute more than one row to the minor when $k > 4$. Indeed, if more than four cameras contributed at least two rows each, then at least 10 rows are accounted for, which leaves at most $k-6$ rows to take from the remaining $k-5$ cameras. So at least one camera will be left out entirely which means that the submatrix of that $4+k$ minor has a zero column and the minor is zero.

A useful fact for us will be that for two positive integers $l > k \geq 2$, there is a simple way to "bump up" a k-focal polynomial to an l-focal polynomial by multiplying the k-focal polynomial with a monomial.

Lemma 2.2.1. *Suppose f is a k-focal polynomial from cameras $\sigma = \{\sigma_1, \dots, \sigma_k\} \subset [n]$ where $k \geq 2$. For any $l > k$ cameras $\tau = \{\sigma_1, \dots, \sigma_k, \tau_1, \dots, \tau_{l-k}\}$, there is a l-focal polynomial g such that $(\prod_{i=1}^{l-k} w_{\tau_i})f = g$ for any choice of variables $w_{\tau_i} \in \{x_{\tau_i}, y_{\tau_i}, z_{\tau_i}\}$, one for each camera.*

Proof. Add the row and column associated to coordinate w_{τ_i} to $\mathcal{A}_\sigma(p)$ for $\tau_1, \dots, \tau_{l-k}$ as follows

$$\begin{pmatrix} A_{\sigma_1} & p_{\sigma_1} & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ A_{\sigma_k} & 0 & \dots & p_{\sigma_k} & 0 & \dots & 0 \\ (A_{\tau_1})_{w_{\tau_1}} & 0 & \dots & 0 & w_{\tau_1} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & 0 \\ (A_{\tau_{l-k}})_{w_{\tau_{l-k}}} & 0 & \dots & 0 & 0 & \dots & w_{\tau_{l-k}} \end{pmatrix}.$$

Taking the determinant of this matrix yields the l-focal polynomial $g = (\prod_{i=1}^{l-k} w_{\tau_i})f$. □

Combining the above facts we get that any l-focal polynomial for $l > 4$ is of the form $(\prod_{i=1}^{l-k} w_{\tau_i})f$ where f is a $k \leq 4$ focal polynomial. This is a generalization of Proposition 2 in [46] that showed that every n-focal polynomial is a monomial multiple of a k-focal polynomial for $k \leq 4$. As a result, we will primarily focus on the ideals $H^2_{\mathcal{A}}$, $H^3_{\mathcal{A}}$, and $H^4_{\mathcal{A}}$, called the *bifocal*, *trifocal*, and *quadrifocal ideals* of $\mathcal{A}$.

A closer look at $H^2_{\mathcal{A}}$ reveals that it is the ideal generated by the $\binom{n}{2}$ epipolar constraints, since $\mathcal{A}_{\{i,j\}}$ is a 6×6 matrix, whose determinant is the epipolar constraint between images i and j. By Lemma 2.2.1, $H^3_{\mathcal{A}}$ contains the bumped up version of $H^2_{\mathcal{A}}$ and for every triplet of images $\{i, j, k\}$, the 27 trifocals implied by the three trifocal tensors relating them. And finally, $H^4_{\mathcal{A}}$ contains the bumped up versions of $H^2_{\mathcal{A}}$ and $H^3_{\mathcal{A}}$ and the 81 quadrifocals implied by the quadrifocal tensor. The fact that we only need to study $H^2_{\mathcal{A}}$, $H^3_{\mathcal{A}}$, and $H^4_{\mathcal{A}}$ lines up with the well known fact in multiview geometry that when studying n-view constraints, one only needs to study the epipolar matrix, the trifocal tensor and the quadrifocal tensor. See Chapter 17 in the book by Hartley & Zisserman [23] for explicit computations of the generators of $H^2_{\mathcal{A}}, H^3_{\mathcal{A}}$, and $H^4_{\mathcal{A}}$ and their history.

In the remainder of this section, we will investigate how k-focal ideals transform under certain linear transformations on cameras. It is widely known that, from image data, the geometry of a camera arrangement can only be determined up to an arbitrary choice of $\mathbb{P}^3$ coordinates. This is reflected in the following lemma.

Lemma 2.2.2 (Projective Ambiguity). *Suppose $G \in \mathrm{GL}_4$. Then for any k, $H^k_{\mathcal{A}} = H^k_{\mathcal{A}G}$ where $\mathcal{A}G = (A_1G, A_2G, \ldots, A_kG)$.*

Proof. This follows since $(\mathcal{A}G)_\sigma(p) = \mathcal{A}_\sigma(p)\,\mathrm{diag}(G, I_k)$ for any k-element subset $\sigma \subset [n]$ which implies that any k-focal of $\mathcal{A}G$ differs from the same k-focal of $\mathcal{A}$ by a factor of $\det(G) \neq 0$. $\square$

From the proof of Lemma 2.2.2, we see that a $\mathbb{P}^3$ coordinate change that sends $\mathbf{q} \mapsto G\mathbf{q}$ maps k-focals to k-focals, picking up only a scalar factor $\det G \neq 0$. We will now see that change of coordinates on the image planes $\mathbb{P}^2$ affect the k-focals in a more subtle way.

Let $\mathcal{G} = (G_1, \ldots, G_n) \in (GL_3)^n$ be a sequence of invertible matrices and consider the camera arrangement $\mathcal{G}\mathcal{A} := (G_1A_1, \ldots, G_nA_n)$ obtained from a given arrangement $\mathcal{A}$ by left-multiplying A_i with G_i. Note that the focal point of the camera A_i is the same as the focal point of the camera G_iA_i. Since $p_i = (x_i, y_i, z_i)$, we denote the ring $\mathbb{C}[x_1, y_1, z_1, \ldots, x_i, y_i, z_i, \ldots, x_n, y_n, z_n]$ by $\mathbb{C}[p_1, \ldots, p_n]$ and a polynomial in it by $f(p_1, \ldots, p_n)$. The sequence $\mathcal{G}$ induces a camera-wise linear change of coordinates $\chi_{\mathcal{G}}$ on $\mathbb{C}[p_1, \ldots, p_n]$ by sending

$$\chi_{\mathcal{G}} : \begin{pmatrix} x_i \\ y_i \\ z_i \end{pmatrix} \mapsto G_i^{-1} \begin{pmatrix} x_i \\ y_i \\ z_i \end{pmatrix} \tag{2.3}$$

Note that this amounts to a change of coordinates in the image planes $\mathbb{P}^2$ of the cameras in $\mathcal{A}$. Let $G^{-1}p$ denote $\chi_{\mathcal{G}}(p) = (G_1^{-1}p_1, \ldots, G_n^{-1}p_n)$. In what follows we will also need the notation $\mathcal{G}^{-1} := (G_1^{-1}, \ldots, G_n^{-1})$, $\mathcal{G}^{-1}\mathcal{A} := (G_1^{-1}A_1, \ldots, G_n^{-1}A_n)$ and $\chi_{\mathcal{G}^{-1}}(p_i) = G_i p_i$.

To analyze the effect of $\chi_{\mathcal{G}}$ on k-focal ideals, we recall the classical Cauchy-Binet formula, a proof of which can be found in [9].

Lemma 2.2.3 (Cauchy-Binet)**.** *If A and B are rectangular matrices of size $m \times n$ and $n \times m$, respectively, where $m \leq n$, then the determinant of the square matrix AB is:*

$$\det(AB) = \sum_{\sigma \in \binom{[n]}{m}} \det(A_{[:,\sigma]}) \det(B_{[\sigma,:]})$$

where $:$ indicates that all rows/columns are taken.

Lemma 2.2.4. *For the k-focal ideal $H^k_{\mathcal{A}}$, $\chi_{\mathcal{G}}(H^k_{\mathcal{A}}) = H^k_{\mathcal{G}\mathcal{A}}$. Similarly, $\chi_{\mathcal{G}^{-1}}(H^k_{\mathcal{G}\mathcal{A}}) = H^k_{\mathcal{A}}$.*

Proof. We prove the first statement and the other follows similarly. We will show that the k-focal ideal of $\mathcal{A}_{[k]}$ is sent to the k-focal ideal of $(\mathcal{G}\mathcal{A})_{[k]}$. The result then follows for the full k-focal ideal $H^k_{\mathcal{A}}$ by summing the k-focal ideals of all $\mathcal{A}_\sigma$ as σ varies over all k-subsets of $[n]$.

Recall that a k-focal polynomial of $\mathcal{A}_{[k]} := (A_1, \ldots, A_k)$ is a maximal minor of:

$$\mathcal{A}_{[k]}(p) = \begin{bmatrix} A_1 & p_1 & & & \\ A_2 & & p_2 & & \\ \vdots & & & \ddots & \\ A_k & & & & p_k \end{bmatrix}.$$

Applying $\chi_{\mathcal{G}}$ to this maximal minor is the same as taking the same maximal minor of

$$\mathcal{A}_{[k]}(\chi_{\mathcal{G}}(p)) = \begin{bmatrix} A_1 & G_1^{-1}p_1 & & & \\ A_2 & & G_2^{-1}p_2 & & \\ \vdots & & & \ddots & \\ A_k & & & & G_k^{-1}p_k \end{bmatrix}.$$

The corresponding k-focal polynomial of $\mathcal{G}\mathcal{A}$ is the same maximal minor of

$$(\mathcal{G}\mathcal{A})_{[k]}(p) = \operatorname{diag}(G_1, \ldots, G_k)\mathcal{A}_{[k]}(p) \tag{2.4}$$

The ideal $\chi_{\mathcal{G}}(H^k_{\mathcal{A}_{[k]}})$ is generated by the maximal minors of $\mathcal{A}_{[k]}(\chi_{\mathcal{G}}(p))$, namely

$$\left\{\det(\mathcal{A}_{[k]}(G^{-1}p)_{[\sigma,:]}) \, : \, \sigma \in \binom{[3k]}{4+k}\right\},$$

while $H^k_{(\mathcal{G}\mathcal{A})_{[k]}}$ is generated by the maximal minors of $(\mathcal{G}\mathcal{A})_{[k]}(p)$. We need to show that these ideals coincide.

Let G denote the block diagonal matrix with blocks $G_1, \ldots, G_n$. A $(4+k)$-minor of $(\mathcal{G}\mathcal{A})_{[k]}(p)$ is the determinant of a submatrix with $4+k$ rows indexed by some $\tau \in \binom{[3k]}{4+k}$. Such a submatrix has the form $G_\tau \mathcal{A}_{[k]}(G^{-1}p)$ where G_τ is the submatrix of G consisting of the rows of G indexed by τ. By the Cauchy-Binet formula,

$$\det(G_\tau \mathcal{A}_{[k]}(G^{-1}p)) = \sum_{\sigma \in \binom{[3k]}{4+k}} \det((G_\tau)_{[:,\sigma]}) \det(\mathcal{A}_{[k]}(G^{-1}p)_{[\sigma,:]}).$$

This implies that $\det(G_\tau \mathcal{A}_{[k]}(G^{-1}p))$ is in the ideal $\chi_{\mathcal{G}}(H^k_{\mathcal{A}_{[k]}})$, hence, $H^k_{(\mathcal{G}\mathcal{A})_{[k]}} \subseteq \chi_{\mathcal{G}}(H^k_{\mathcal{A}_{[k]}})$.

The reverse containment follows by applying the same argument to $\mathcal{A}_{[k]}(p) = G^{-1}G\mathcal{A}_{[k]}(p)$ and $G\mathcal{A}_{[k]}(p)$ where G^{-1} is the block diagonal matrix with blocks $G_1^{-1}, \ldots, G_k^{-1}$.

Summing over all k camera subsets, the result follows:

$$\chi_{\mathcal{G}}(H^k_{\mathcal{A}}) = \chi_{\mathcal{G}}(\sum_{\sigma \in \binom{[n]}{k}} H^k_{\mathcal{A}_\sigma}) = \sum_{\sigma \in \binom{[n]}{k}} \chi_{\mathcal{G}}(H^k_{\mathcal{A}_\sigma}) = \sum_{\sigma \in \binom{[n]}{k}} H^k_{(\mathcal{G}\mathcal{A})_\sigma} = H^k_{\mathcal{G}\mathcal{A}}.$$

□

This proof shows that, contrary to popular belief, it is not true that k-focal polynomials go to k-focal polynomials under the change of coordinates given by $\chi_{\mathcal{G}}$, but the ideals do as in Lemma 2.2.4.

2.3 The Multiview Ideal

Recall from Definition 2.1.2 that the multiview ideal $M_{\mathcal{A}}$ of the camera arrangement $\mathcal{A}$ is the vanishing ideal of $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$, meaning that it is the set of all polynomials in $\mathbb{C}[p_1, \ldots, p_n]$ that vanish on $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$. Since $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$ is a subset of $(\mathbb{P}^2_{\mathbb{R}})^n$, $M_{\mathcal{A}}$ is, in fact, generated by polynomials with real coefficients[1].

The complex projective variety $\mathbf{V}(M_{\mathcal{A}}) \subset (\mathbb{P}^2)^n$, which is the complex Zariski closure of $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$, is the multiview variety of $\mathcal{A}$. One might wonder if it is better to study the real Zariski closure of $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$ and its vanishing ideal since complex points in the multiview variety do not have any physical meaning, and hence no relevance to multiview geometry. However, observe that if the real Zariski closure was strictly smaller than the set of real points in $\mathbf{V}(M_{\mathcal{A}})$, then there would be a polynomial not in $M_{\mathcal{A}}$ that vanishes on $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$, which would contradict that $M_{\mathcal{A}}$ is the vanishing ideal of $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$. Therefore, $M_{\mathcal{A}}$ is also the vanishing ideal of the real Zariski closure of $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$, and hence a *real radical ideal* [36, §12.5].

Further, since $\varphi_{\mathcal{A}}$ is a polynomial map and $\mathbb{P}^3_{\mathbb{R}}$ is irreducible, $\mathbf{V}(M_{\mathcal{A}})$ is an irreducible three-dimensional variety in $(\mathbb{P}^2)^n$. Hence $M_{\mathcal{A}}$ is a *prime* (homogeneous) ideal, meaning that if $fg \in M_{\mathcal{A}}$ then either f or g is in $M_{\mathcal{A}}$.

It was shown in [5] that the bifocals, trifocals and quadrifocals of $\mathcal{A}$ form a *universal Gröbner basis* of $M_{\mathcal{A}}$ under a certain genericity assumption on the cameras. This means that this collection of polynomials form a *Gröbner basis* for $M_{\mathcal{A}}$ with respect to any term order [14]. We will use this result to establish a generating set for $M_{\mathcal{A}}$ when the camera foci are distinct.

We first note what happens to $M_{\mathcal{A}}$ under the change of coordinates $\chi_{\mathcal{G}}$ defined in the previous section. Recall that $\chi_{\mathcal{G}}$ sends a polynomial $f(p_1, \ldots, p_n) \in \mathbb{C}[p_1, \ldots, p_n]$ to $f(G_1^{-1}p_1, \ldots, G_n^{-1}p_n)$.

Lemma 2.3.1. *The image of the multiview ideal $M_{\mathcal{A}}$ under the map $\chi_{\mathcal{G}}$ is $M_{\mathcal{G}\mathcal{A}}$, the multiview*

[1] Let $h(x) = f(x) + ig(x)$ be a complex polynomial, where $f(x)$ and $g(x)$ are real polynomials. Then if $h(x)$ vanish on a set of real points, then so must $f(x)$ and $g(x)$.

ideal of $\mathcal{GA}$. i.e. , $\chi_{\mathcal{G}}(M_{\mathcal{A}}) = M_{\mathcal{GA}}$. Similarly, $\chi_{\mathcal{G}^{-1}}(M_{\mathcal{GA}}) = M_{\mathcal{A}}$.

Proof. Again, we will prove that $\chi_{\mathcal{G}}(M_{\mathcal{A}}) = M_{\mathcal{GA}}$. The proof that $\chi_{\mathcal{G}^{-1}}(M_{\mathcal{GA}}) = M_{\mathcal{A}}$ is similar.

From the definition we see that a polynomial $f(p_1, \dots, p_n)$ vanishes on the multiview variety $\mathbf{V}(M_{\mathcal{A}})$ if and only if $f(A_1\mathbf{q}, \dots, A_n\mathbf{q}) = 0$ for all $\mathbf{q} \in \mathbb{P}^3 \setminus \{\mathbf{c}_1, \dots, \mathbf{c}_n\}$, equivalently, if and only if

$$f(G_1^{-1}(G_1A_1\mathbf{q}), \dots, G_n^{-1}(G_nA_n\mathbf{q})) = 0$$

for all $\mathbf{q} \in \mathbb{P}^3 \setminus \{\mathbf{c}_1, \dots, \mathbf{c}_n\}$. The multiview variety of $\mathcal{GA}$ is the Zariski closure of the points $(G_1A_1\mathbf{q}, \dots, G_nA_n\mathbf{q})$ as $\mathbf{q}$ varies over $\mathbb{P}^3 \setminus \{\mathbf{c}_1, \dots, \mathbf{c}_n\}$. Therefore, f vanishes on $\mathbf{V}(M_{\mathcal{A}})$ if and only if $\chi_{\mathcal{G}}(f)$ vanishes on $\mathbf{V}(M_{\mathcal{GA}})$. This proves that $\chi_{\mathcal{G}}(M_{\mathcal{A}}) \subseteq M_{\mathcal{GA}}$.

To finish the proof we need to argue that if $g(p_1, \dots, p_n) \in M_{\mathcal{GA}}$ then $g = \chi_{\mathcal{G}}(f)$ for some $f \in M_{\mathcal{A}}$. A polynomial $g \in M_{\mathcal{GA}}$ if and only if $g(G_1A_1\mathbf{q}, \dots, G_nA_n\mathbf{q}) = 0$ for all $\mathbf{q} \in \mathbb{P}^3 \setminus \{\mathbf{c}_1, \dots, \mathbf{c}_n\}$ if and only if $g(G_1\mathbf{p}_1, \dots, G_n\mathbf{p}_n) = 0$ for all $(\mathbf{p}_1, \dots, \mathbf{p}_n) \in \mathbf{V}(M_{\mathcal{A}})$. Define $g(G_1p_1, \dots, G_np_n) =: f \in M_{\mathcal{A}}$. Then $\chi_{\mathcal{G}}(f) = g(p_1, \dots, p_n)$. $\square$

We will use the results obtained so far to give an elementary proof that the bifocals and trifocals generate the multiview ideal $M_{\mathcal{A}}$ for any arrangement $\mathcal{A}$ of cameras with pairwise distinct foci. An important tool will be *translational* cameras.

Definition 2.3.1. *A camera T is said to be translational if its left 3×3 block is the identity matrix, i.e. , $T = [I\ \mathbf{t}]$ for some $\mathbf{t} \in \mathbb{R}^3$.*

Lemma 2.3.2. *If $\mathcal{T}$ is an arrangement of translational cameras, then $H^4_{\mathcal{T}} \subseteq H^3_{\mathcal{T}}$.*

Proof. Using Macaulay2, this statement can be checked for $n = 4$ translational cameras with foci represented symbolically as $(t_{i1}, t_{i2}, t_{i3}, -1)$. For $n \geq 4$, since $H^4_{\mathcal{T}} = \sum_{\sigma \in \binom{[n]}{4}} H^4_{\mathcal{T}_\sigma}$ and $H^3_{\mathcal{T}} = \sum_{\sigma \in \binom{[n]}{3}} H^3_{\mathcal{T}_\sigma}$, the statement follows. $\square$

We now use translational cameras to show that the quadrifocals are not needed in a generating set of $M_{\mathcal{A}}$. This is done by extending the result for translational cameras to finite cameras. Recall that a *finite* camera is a camera whose left 3×3 block is invertible, or equivalently a camera whose focal point is not a point at infinity. Observe that any finite camera can be obtained by multiplying some translational camera on the left by an invertible 3×3 matrix.

Corollary 2.3.1. *If $\mathcal{A}$ is any arrangement of cameras, then $H^4_{\mathcal{A}} \subseteq H^3_{\mathcal{A}}$.*

Proof. If $\mathcal{A}$ is an arrangement of finite cameras, then $A_i = G_i[I\ \mathbf{t}_i]$ for some $G_i \in GL_3$. Therefore $\mathcal{A} = \mathcal{G}\mathcal{T}$ where $\mathcal{T}$ is an arrangement of translational cameras. By Lemma 2.3.2, $H^4_{\mathcal{T}} \subseteq H^3_{\mathcal{T}}$. Hence, Lemma 2.2.4 implies

$$H^4_{\mathcal{A}} = H^4_{\mathcal{G}\mathcal{T}} = \chi_{\mathcal{G}}(H^4_{\mathcal{T}}) \subseteq \chi_{\mathcal{G}}(H^3_{\mathcal{T}}) = H^3_{\mathcal{G}\mathcal{T}} = H^3_{\mathcal{A}}.$$

For any four cameras indexed by $\sigma \in \binom{[n]}{4}$, there exists some $G \in \mathrm{GL}_4$ which takes the foci of $\mathcal{A}_\sigma$ off of the plane at infinity, *i.e.* , so that $\mathcal{A}_\sigma G$ is an arrangement of finite cameras. Inverting this $\mathbb{P}^3$-coordinate change does not change ideal containment by Lemma 2.2.2. The general result follows since $H^4_{\mathcal{A}} = \sum_{\sigma \in \binom{[n]}{4}} H^4_{\mathcal{A}_\sigma} \subseteq \sum_{\sigma \in \binom{[n]}{3}} H^3_{\mathcal{A}_\sigma} = H^3_{\mathcal{A}}$. ∎

To get to our main result, we will need a result from [5] about camera arrangements $\mathcal{A}$ that are generic in the sense that all 4×4 minors of $[A_1^\top\ A_2^\top\ \cdots\ A_n^\top]$ are non-zero. We call such an $\mathcal{A}$ *minor-generic*.

Corollary 2.3.2. *Suppose $\mathcal{A}$ is minor-generic. Then $M_{\mathcal{A}} = H^2_{\mathcal{A}} + H^3_{\mathcal{A}}$.*

Proof. Theorem 2.1 in [5] says that if $\mathcal{A}$ is minor-generic, then the bifocals, trifocals and quadrifocals form a universal Gröbner basis of $M_{\mathcal{A}}$. In particular, this implies that $M_{\mathcal{A}} = H^2_{\mathcal{A}} + H^3_{\mathcal{A}} + H^4_{\mathcal{A}}$. The statement is then immediate from Corollary 2.3.1. ∎

Minor-genericity is a purely algebraic condition on camera arrangements. The following statement, which appears as a brief comment in [5] without proof, gives a geometric reinterpretation of this condition.

Lemma 2.3.3. *If $\mathcal{A}$ is minor-generic, then the foci of the cameras in $\mathcal{A}$ are pairwise distinct. Conversely, if the cameras in $\mathcal{A}$ have pairwise distinct foci, then there exist $G_i \in \mathrm{GL}_3$ such that $\mathcal{G}\mathcal{A}$ is minor-generic.*

Proof. Let $L_i \subset \mathbb{C}^4$ denote the three-dimensional row span of A_i. If A_i and A_j have the same focal point then $L_i = L_j$ and hence any four of the six rows of A_i and A_j are linearly dependent and $\mathcal{A}$ is not minor-generic. This proves the first statement.

Now suppose the foci of cameras in $\mathcal{A}$ are pairwise distinct. This means that the planes L_i are pairwise distinct. For any $G_i \in \mathrm{GL}_3$, the rows of G_iA_i form a basis of L_i. By choosing G_i appropriately, the three rows of A_i can be sent to any choice of three linearly independent vectors in L_i. We need to show that there is a choice of G_i such that no four rows from the matrices G_iA_i are linearly dependent.

Consider the $3n \times 4$ matrix obtained by vertically stacking the cameras in $\mathcal{A}$, as a point in $(\mathbb{C}^4)^{3n}$, with coordinates x^i_{kl} representing the (k, l)-entry of the ith camera. We will identify this point in $(\mathbb{C}^4)^{3n}$ with the corresponding $3n \times 4$ matrix, and stack of n cameras, and call all of them $\mathcal{A}$. Let $\mathcal{A}(x)$ denote the symbolic $3n \times 4$ matrix with entries x^i_{kl}. For $\sigma \in \binom{[3n]}{4}$, let d_σ denote the determinant of the 4×4 submatrix of $\mathcal{A}(x)$ with rows indexed by σ. These cut out $\binom{3n}{4}$ quartic hypersurfaces $\mathbf{V}(d_\sigma)$ in $(\mathbb{C}^4)^{3n}$. Let v_i denote the normal of the hyperplane $L_i \subset \mathbb{C}^4$. Impose linear conditions saying that the rows of $\mathcal{A}(x)$, numbered $3i, 3i + 1, 3i + 2$, dot to zero with v_i. These $3n$ equations determine a subspace L in $(\mathbb{C}^4)^{3n}$ of dimension at least $9n = 12n - 3n$. The given point $\mathcal{A}$ lies in L. We need to show that there is a choice of $\mathcal{G} \in (\mathrm{GL}_3)^n$ such that $\mathcal{G}\mathcal{A}$ (which again lies in L) avoids the determinantal surfaces. This is equivalent to picking a basis for each L_i that stack together to a $\mathcal{B} \in L \setminus \bigcup_\sigma \mathbf{V}(d_\sigma)$.

We first show that L is not contained in any $\mathbf{V}(d_\sigma)$ by exhibiting a point in $L \setminus \mathbf{V}(d_\sigma)$ for each σ. Since at most four cameras can be involved in any d_σ, we may assume without loss of generality that σ involves only rows of the first four cameras. There are four cases to consider depending on how many rows these four cameras contribute to σ — the possibilities being $(3, 1, 0, 0)$, $(2, 2, 0, 0)$, $(2, 1, 1, 0)$, and $(1, 1, 1, 1)$. In each case we will produce a $\mathcal{B} \in L \setminus \mathbf{V}(d_\sigma)$.

A key observation is that A_i and A_j having distinct foci implies $L_i \cap L_j$ is a proper subspace of both L_i and L_j for all i, j. Our starting point in each case below is $\mathcal{A} \in L$ which we modify to the needed $\mathcal{B}$ by replacing the bases of L_i that provide the rows of A_i.

Case 1. (3,1,0,0): Modify $\mathcal{A}$ to $\mathcal{B}$ by choosing a basis for L_2 to be the three rows of B_2 so that no element in this basis lies in $L_1 \cap L_2$. Then $\mathcal{B}$ does not vanish on d_σ.

Case 2. (2,2,0,0): Choose a basis for L_1 such that the two rows v_1, v_2 contributing to σ from the first camera are chosen from $L_1 \setminus L_2$. Then $L_2 \cap \mathrm{Span}\{v_1, v_2\}$ is a proper subspace of L_2 of dimension at most one. Therefore taking two linearly independent vectors v_3, v_4 outside of this subspace as the two rows from L_2 creates a $\mathcal{B}$ that does not vanish on d_σ.

Case 3. (2,1,1,0): Choose a basis for L_1 such that the two contributing rows v_1, v_2 from the first camera lie in $L_1 \setminus (L_2 \cup L_3)$. Choose the row v_3 from L_2 such that $v_3 \in L_2 \setminus (\mathrm{Span}\{v_1, v_2\} \cup L_3)$, which forces $L_3 \cap \mathrm{Span}\{v_1, v_2, v_3\}$ to be a proper subspace of L_3. Taking v_4 outside this subspace, we get a point $\mathcal{B} \in L$ at which d_σ does not vanish.

Case 4. (1,1,1,1): Choose $v_1 \in L_1 \smallsetminus (L_2 \cup L_3 \cup L_4)$, $v_2 \in L_2 \smallsetminus (\mathrm{Span}\{v_1\} \cup L_3 \cup L_4)$, $v_3 \in L_3 \smallsetminus (\mathrm{Span}\{v_1, v_2\} \cup L_4)$, and $v_4 \in L_4 \smallsetminus (\mathrm{Span}\{v_1, v_2, v_3\})$. By construction, we get a point in L at which d_σ does not vanish.

Therefore, $L \cap \mathbf{V}(d_\sigma)$ is a proper subvariety of L for each σ, and a generic choice of $\mathcal{G}$ will put $\mathcal{G}\mathcal{A} \in L \smallsetminus \bigcup_\sigma \mathbf{V}(d_\sigma)$. □

We note that $\mathcal{A}$ having distinct foci does not imply that $\mathcal{A}$ is minor-generic. A simple example would be an arrangement of four translational cameras; the submatrix consisting of the four first rows in each camera has zero determinant. However, having distinct foci allows the camera arrangement to be made minor-generic by the action of a tuple $\mathcal{G}$. We are now ready to prove the main theorem of this section.

Theorem 2.3.1. *Let $\mathcal{A}$ be an arrangement of cameras with distinct foci. Then $M_\mathcal{A} = H^2_\mathcal{A} + H^3_\mathcal{A}$.*

Proof. By Lemma 2.3.3, there exists $\mathcal{G} \in (GL_3)^n$ such that $\mathcal{G}\mathcal{A}$ is minor-generic. Then, by

Corollary 2.3.2, $M_{\mathcal{GA}} = H^2_{\mathcal{GA}} + H^3_{\mathcal{GA}}$. Therefore, by Lemmas 2.3.1 and 2.2.4, we get

$$\begin{aligned} M_{\mathcal{A}} = \chi_{\mathcal{G}^{-1}}(M_{\mathcal{GA}}) &= \chi_{\mathcal{G}^{-1}}(H^2_{\mathcal{GA}}) + \chi_{\mathcal{G}^{-1}}(H^3_{\mathcal{GA}}) \\ &= H^2_{\mathcal{A}} + H^3_{\mathcal{A}} \end{aligned}$$

□

Proposition 5(1) in [46] says that the $H^2_{\mathcal{A}}$ and $H^3_{\mathcal{A}}$ together cut out the multiview variety which implies that $H^2_{\mathcal{A}} + H^3_{\mathcal{A}} \subseteq \mathcal{M}_{\mathcal{A}}$. Theorem 2.3.1 shows that these polynomials also generate the multiview ideal providing the analogous ideal-theoretic statement.

Theorem 2.3.1 improves on Corollary 2.7 in [5] which states that when the foci of the cameras A_i are in linearly general position, then $M_{\mathcal{A}}$ is generated by the bifocals and trifocals. Theorem 2.3.1 requires no sophisticated condition on the cameras beyond the foci being pairwise distinct.

Conca *et al.* [13] and Li [33] also consider the vanishing ideal of the image of a linear map from a projective space to a product of projective spaces. It is shown in [13] that this ideal is Cartwright-Sturmfels, meaning that its initial ideal is radical after a generic change of coordinates. Both of these works allow for projective spaces of arbitrary dimension. Specializing to our situation, Li's results show that $M_{\mathcal{A}} = \sum_{k=2}^{n} H^k_{\mathcal{A}}$ while we prove that $M_{\mathcal{A}} = H^2_{\mathcal{A}} + H^3_{\mathcal{A}}$.

Just like in [46] where the results automatically generalized from projective cameras to Euclidean cameras, Theorem 2.3.1 also generalizes to Euclidean cameras. Recall that a camera A_i is Euclidean if it is of the form $A_i = [R_i \; t_i]$ where $R_i \in \mathrm{SO}_3$.

Corollary 2.3.3. *Let $\mathcal{A}$ be an arrangement of Euclidean cameras with pairwise distinct foci. Then $M_{\mathcal{A}} = H^2_{\mathcal{A}} + H^3_{\mathcal{A}}$.*

We state one more consequence of Theorem 2.3.1 which will be needed in the next section.

Corollary 2.3.4. *Let $\mathcal{A}$ be a camera arrangement with pairwise distinct foci. Then for any $\mathbf{p}_i \in \mathbb{P}^2$, the points $(A_1\mathbf{c}_i, A_2\mathbf{c}_i, \ldots, \mathbf{p}_i, \ldots, A_n\mathbf{c}_i)$ lie in $\mathbf{V}(M_{\mathcal{A}})$ where $\mathbf{c}_i$ is the focal point of A_i.*

Proof. By Theorem 2.3.1, it suffices to show that for any i, the bifocals and trifocals vanish on the points $(A_1\mathbf{c}_i, A_2\mathbf{c}_i, \ldots, \mathbf{p}_i, \ldots, A_n\mathbf{c}_i)$. For any pair of cameras $\{i, j\}$, observe that $(\mathbf{c}_i, 0, -1)$ is a nonzero element of $\ker \mathcal{A}_{\{i,j\}}(\mathbf{p}_i, A_j\mathbf{c}_i)$. For any pair $\{j, k\}$ not containing camera i, $(\mathbf{c}_i, -1, -1)$ is a nonzero element of $\ker \mathcal{A}_{\{j,k\}}(A_j\mathbf{c}_i, A_k\mathbf{c}_i)$. Hence all polynomials of $H^2_{\mathcal{A}}$ vanish on $(A_1\mathbf{c}_i, A_2\mathbf{c}_i, \ldots, \mathbf{p}_i, \ldots, A_n\mathbf{c}_i)$. A similar argument applies to any triples of cameras, from which it follows that all polynomials in $H^3_{\mathcal{A}}$ vanish on $(A_1\mathbf{c}_i, A_2\mathbf{c}_i, \ldots, \mathbf{p}_i, \ldots, A_n\mathbf{c}_i)$. $\square$

The image of focal point i in image j, *i.e.* , $A_j\mathbf{c}_i$, is called the *epipole in image j relative to image i*. Corollary 2.3.4 shows that while the product of an arbitrary point in image i with all epipoles relative to image i does not appear in the image of $\varphi_{\mathcal{A}}$, these points appear in the multiview variety after taking Zariski closure. See also Proposition 1 in [46].

We conclude this section by showing that the hypothesis in Theorem 2.3.1 cannot be relaxed, namely if a pair of foci of cameras in $\mathcal{A}$ coincide, then the multiview ideal is strictly larger than the ideal generated by bifocals and trifocals.

Example 2.3.1. *Consider the four translational camera arrangement $\mathcal{A}$ where $\mathbf{t}_1, \mathbf{t}_2 = (0, 0, 0)$, $\mathbf{t}_3 = (1, 1, 1)$, $\mathbf{t}_4 = (-1, -1, -1)$. Eliminating the variables q and λ_i from the ideal $\langle A_i q - \lambda_i p_i : i = 1, \ldots, n\rangle$, we can directly obtain $M_{\mathcal{A}}$. Computing a primary decomposition of $H^2_{\mathcal{A}} + H^3_{\mathcal{A}}$, we find that*

$$H^2_{\mathcal{A}} + H^3_{\mathcal{A}} = M_{\mathcal{A}} \cap \langle y_4 - z_4, y_3 - z_3, x_4 - z_4, x_3 - z_3\rangle.$$

The extra component $\langle y_4 - z_4, y_3 - z_3, x_4 - z_4, x_3 - z_3\rangle$ cuts out the points $(\mathbf{p}_1, \mathbf{p}_2, A_3\mathbf{c}_1, A_4\mathbf{c}_1)$, and from the primary decomposition we see that the projective variety they form is not contained in $\mathbf{V}(M_{\mathcal{A}})$.

2.4 More Ideals for the Multiview Variety

In the computer vision literature, there are several sets of polynomials that have been shown to vanish on the space of images $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})$, and hence they also vanish on the multiview variety. We now consider three such sets of polynomials and the ideals they generate, and compare them to the multiview ideal $M_{\mathcal{A}}$.

2.4.1 Heyden and Åström [27]

Heyden and Åström were the first to do an algebraic study of the multiview variety, by studying the n-focal ideal $H^n_{\mathcal{A}}$ [27]. The variety of this ideal is indeed the multiview variety.

Lemma 2.4.1. *For any camera arrangement $\mathcal{A}$ with pairwise distinct foci, $\mathbf{V}(M_{\mathcal{A}}) = \mathbf{V}(H^n_{\mathcal{A}})$.*

Proof. Recall from the image formation equations, $A_i\mathbf{q} = \lambda_i\mathbf{p}_i$ for all $i = 1, \ldots, n$, that if $\mathbf{p} = (\mathbf{p}_1, \ldots, \mathbf{p}_n)$ lies in the image of $\varphi_{\mathcal{A}}$ then the matrix $\mathcal{A}(\mathbf{p})$ has a non-trivial kernel. This means that all maximal minors of $\mathcal{A}(p)$ vanish on the image of $\varphi_{\mathcal{A}}$, and therefore also on its Zariski closure, which is the multiview variety. Therefore, $\mathbf{V}(M_{\mathcal{A}}) \subseteq \mathbf{V}(H^n_{\mathcal{A}})$.

To see the reverse inclusion, suppose $\mathbf{p} = (\mathbf{p}_1, \ldots, \mathbf{p}_n) \in \mathbf{V}(H^n_{\mathcal{A}})$ which means that $\mathcal{A}(\mathbf{p})$ is rank deficient and there is a nonzero vector of the form $(\mathbf{q}, -\lambda_1, \ldots, -\lambda_n)$ in the kernel of $\mathcal{A}(\mathbf{p})$. If $\mathbf{q} = 0$, then we will get that $\lambda_i\mathbf{p}_i = 0$ for all i. However, since $\mathbf{p}_i \neq 0$, it must be that $\lambda_i = 0$ for all i and hence the vector in the kernel is the zero vector which is a contradiction. Therefore, there is a nonzero vector $\mathbf{q}$ such that $A_i\mathbf{q} = \lambda_i\mathbf{p}_i$ for some λ_i. If $\mathbf{q}$ is not the focal point of any camera, then $\mathbf{p}$ lies in $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{C}})$. Since $\varphi_{\mathcal{A}}$ is continuous, $\varphi_{\mathcal{A}}(\overline{\mathbb{P}^3_{\mathbb{R}}}) \subseteq \overline{\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{R}})}$. It follows that $\varphi_{\mathcal{A}}(\mathbb{P}^3_{\mathbb{C}}) \subseteq \mathbf{V}(M_{\mathcal{A}})$ because $\overline{\mathbb{P}^3_{\mathbb{R}}} = \mathbb{P}^3_{\mathbb{C}}$ and so $\mathbf{p} \in \mathbf{V}(M_{\mathcal{A}})$. On the other hand, if $\mathbf{q}$ is the focal point $\mathbf{c}_i$ of camera i, then $\mathbf{p}_j = A_j\mathbf{c}_i$ for all $j \neq i$, and by Corollary 2.3.4, $\mathbf{p} \in \mathbf{V}(M_{\mathcal{A}})$. Thus we get that $\mathbf{V}(M_{\mathcal{A}}) \supseteq \mathbf{V}(H^n_{\mathcal{A}})$.

□

Example 2.3.1 shows that the assumption of distinct foci is necessary for Lemma 2.4.1. In this example, $n = 4$ and $H^4_{\mathcal{A}} = H^2_{\mathcal{A}} + H^3_{\mathcal{A}}$ by Corollary 2.3.1. We see that $\mathbf{V}(H^4_{\mathcal{A}})$ has a component other than $\mathbf{V}(M_{\mathcal{A}})$.

2.4.2 Faugeras et al. [19].

The second set of polynomials we will study were constructed by Faugeras & Mourrain while proving that the multiview variety is cut out by epipolar/bifocal and trifocal polynomials, and that the quadrifocal constraints corresponding to the quadrifocal tensor were not needed [18, 19].

Observe that $A_i\mathbf{q} = \lambda_i\mathbf{p}_i$ implies $A_i\mathbf{q} \times \mathbf{p}_i = 0$, for each i, or equivalently, $[p_i]_\times A_i\mathbf{q} = 0$, where

$$[p_i]_\times = \begin{pmatrix} 0 & -z_i & y_i \\ z_i & 0 & -x_i \\ -y_i & x_i & 0 \end{pmatrix} \tag{2.5}$$

represents taking cross product with p_i, *i.e.* , $[p_i]_\times v = p_i \times v$. Stacking all 3×4 matrices $[p_i]_\times A_i$, we get the $3n \times 4$ partially symbolic matrix

$$\mathcal{A}^F(p) := \begin{pmatrix} [p_1]_\times A_1 \\ [p_2]_\times A_2 \\ \vdots \\ [p_n]_\times A_n \end{pmatrix}. \tag{2.6}$$

If there is a world point $\mathbf{q}$ satisfying $A_i\mathbf{q} \times \mathbf{p}_i = 0$, then this matrix is rank deficient and all maximal minors of $\mathcal{A}_F(p)$ vanishes on the multiview variety.

Definition 2.4.1. *The ideal of all maximal 4×4 minors of $\mathcal{A}^F(p)$, denoted by $F_\mathcal{A}$, will be called the* **Faugeras ideal** *of the arrangement $\mathcal{A}$. We denote the subideals of $F_\mathcal{A}$ generated by minors involving only two and three cameras by $F^2_\mathcal{A}$ and $F^3_\mathcal{A}$, respectively.*

We now describe a sequence of matrix transformations that allow us to obtain $\mathcal{A}^F(p)$ from $\mathcal{A}(p)$. Let $P(p) := \mathrm{diag}([p_1]_\times, \ldots, [p_n]_\times)$ be the symbolic block diagonal matrix of size $3n \times 3n$. Multiplying $\mathcal{A}(p)$ on the left by the block diagonal matrix $P(p)$ and dropping the rightmost n columns of the resulting matrix, we obtain $\mathcal{A}^F(p)$:

$$\mathcal{A}^F(p) = P(p)\mathcal{A}(p) \begin{bmatrix} I_4 \\ 0_{n\times 4} \end{bmatrix} = P(p)\mathcal{A} \tag{2.7}$$

where as before, we abuse notation to let $\mathcal{A}$ also represents the $3n \times 4$ matrix $[A_1; \ldots; A_n]$ obtained by stacking the cameras vertically. From the matrix constructions of $H^n_{\mathcal{A}}$ and $F_{\mathcal{A}}$, we observe that their projective vanishing sets in $(\mathbb{P}^2)^n$ coincide.

Lemma 2.4.2. *For any camera arrangement $\mathcal{A}$ with pairwise distinct foci,* $\mathbf{V}(M_{\mathcal{A}}) = \mathbf{V}(F_{\mathcal{A}})$.

Proof. The proof will follow from Lemma 2.4.1 if we can show that $\mathbf{V}(F_{\mathcal{A}}) = \mathbf{V}(H^n_{\mathcal{A}})$. If $\mathbf{p} \in (\mathbb{P}^2)^n$ is such that $\mathcal{A}^F(\mathbf{p})$ drops rank, then there exists a nonzero $\mathbf{q} \in \ker(\mathcal{A}^F(\mathbf{p}))$ so that $A_i\mathbf{q} \times \mathbf{p}_i = 0$ for all i. This means there exist nonzero scale factors λ_i such that $A_i\mathbf{q} = \lambda_i \mathbf{p}_i$. The vector $(\mathbf{q}, -\lambda_1, \ldots, -\lambda_n)$ is a nontrivial element in $\ker(\mathcal{A}(\mathbf{p}))$, so $\mathcal{A}(\mathbf{p})$ is rank deficient. Therefore $\mathbf{V}(F_{\mathcal{A}}) \subseteq \mathbf{V}(H^n_{\mathcal{A}})$

For the other inclusion, if there is a nontrivial $(\mathbf{q}, -\lambda_1, \ldots, -\lambda_n) \in \ker(\mathcal{A}(\mathbf{p}))$ for some $\mathbf{p} \in (\mathbb{P}^2)^n$, then as in the proof of Lemma 2.4.1, $\mathbf{q}$ must be nonzero, and so $\mathbf{q}$ is a nontrivial element of $\ker(\mathcal{A}^F(\mathbf{p}))$. This shows that $\mathbf{V}(F_{\mathcal{A}}) \supseteq \mathbf{V}(H^n_{\mathcal{A}})$, hence $\mathbf{V}(F_{\mathcal{A}}) = \mathbf{V}(H^n_{\mathcal{A}}) = \mathbf{V}(M_{\mathcal{A}})$.

□

2.4.3 Ma et al. [35]

The third and final set of polynomials we will study are the so called *multiview rank constraints* which were proposed by Ma and collaborators [35] as an alternative to the multilinear constraints studied for example in Hartley & Zisserman [23].

Suppose $A_1 = [I \ 0]$ and $A_i = [B_i \ \mathbf{t}_i]$ for $i \geq 2$. Starting with $\mathcal{A}(p)$, a series of matrix operations are described in Chapter 8 in [35] to arrive at a new set of determinantal polynomials, arising as maximal minors of

$$\mathcal{A}^Y(p) := \begin{bmatrix} p_1 \times (Ip_1) & p_1 \times 0 \\ p_2 \times (B_2 p_1) & p_2 \times t_2 \\ \vdots & \vdots \\ \vdots & \vdots \\ p_n \times (B_n p_1) & p_n \times t_n \end{bmatrix}. \tag{2.8}$$

Definition 2.4.2. *The ideal of all maximal* 2×2 *minors of* $\mathcal{A}^Y(p)$*, denoted by* $Y_{\mathcal{A}}$*, will be called the* **Ma ideal** *of the arrangement* $\mathcal{A}$*.*

We observe that $\mathcal{A}^Y(p)$ can be obtained from $\mathcal{A}^F(p)$ by multiplying by a single matrix on the right:

$$\mathcal{A}^Y(p) = \mathcal{A}^F(p) \begin{bmatrix} p_1 & 0 \\ 0 & 1 \end{bmatrix}. \tag{2.9}$$

From this we observe that $Y_{\mathcal{A}}$ has the same projective vanishing set as $F_{\mathcal{A}}$, and hence $H^n_{\mathcal{A}}$ and $M_{\mathcal{A}}$.

Lemma 2.4.3. *For any camera arrangement* $\mathcal{A}$ *with pairwise distinct foci and* $A_1 = [I\ 0]$*,* $\mathbf{V}(M_{\mathcal{A}}) = \mathbf{V}(Y_{\mathcal{A}})$*.*

Proof. If $\mathbf{p} \in (\mathbb{P}^2)^n$ is such that $\mathcal{A}^Y(\mathbf{p})$ drops rank, then there exists a nontrivial $(v_1, v_2) \in \ker(\mathcal{A}^Y(\mathbf{p}))$. Therefore, $\mathbf{q} = (v_1\mathbf{p}_1, v_2) \in \ker(\mathcal{A}^F(\mathbf{p}))$ is nontrivial. Note that it is necessary that we assume $A_1 = [I\ 0]$ so that $[\mathbf{p}_1]_\times A_1(v_1\mathbf{p}_1, v_2) = v_1[\mathbf{p}_1]_\times \mathbf{p}_1 = 0$. This shows that $\mathbf{V}(Y_{\mathcal{A}}) \subseteq \mathbf{V}(F_{\mathcal{A}})$.

For the other inclusion, if $0 \neq \mathbf{q} \in \ker(\mathcal{A}^F(\mathbf{p}))$ for some $\mathbf{p} \in (\mathbb{P}^2)^n$, then since $\mathbf{p}_1 \times [I\ 0]\mathbf{q} = 0$, there exists a scalar v_1 such that $v_1\mathbf{p}_1 = (\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3)$. This means that $(v_1, \mathbf{q}_4) \in \ker(\mathcal{A}^Y(\mathbf{p}))$, which is nontrivial because if $v_1 = 0$, then $(\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3) = 0$, so $\mathbf{q}_4 \neq 0$. This shows $\mathbf{V}(Y_{\mathcal{A}}) \supseteq \mathbf{V}(F_{\mathcal{A}})$, and the desired result follows from Lemma 2.4.2.

□

Observe that $Y_{\mathcal{A}}$ is generated by polynomials of total degree 3. This fact has an interesting consequence. As we mentioned earlier, $Y_{\mathcal{A}}$ has been proposed as an alternate algebraic foundation for multi-view geometry. From Lemma 2.4.3, we know that it cuts out the multiview variety. Since $M_{\mathcal{A}}$ is the vanishing ideal of the multiview variety, we get that $Y_{\mathcal{A}} \subseteq M_{\mathcal{A}}$. However, from Theorem 2.3.1 we know that $M_{\mathcal{A}} = H^2_{\mathcal{A}} + H^3_{\mathcal{A}}$, *i.e.* it is generated by polynomials of degree two and three, which means that in general $Y_{\mathcal{A}} \neq M_{\mathcal{A}}$ and instead $Y_{\mathcal{A}} \subset M_{\mathcal{A}}$ or equivalently $Y_{\mathcal{A}} \subset H^2_{\mathcal{A}} + H^3_{\mathcal{A}}$. This means that the bifocals and trifocals imply the multiview rank constraints, but not the other way around. Similarly, $H^n_{\mathcal{A}}$ and $F_{\mathcal{A}}$, which are generated by polynomials of total degree n and four respectively, are properly contained in $M_{\mathcal{A}}$. We see this in Example 2.4.1 below.

2.4.4 Relationships to the Multiview Ideal

We now compute the three ideals on an example, foreshadowing their structural properties, which we examine next.

Example 2.4.1. *Consider the translational arrangement $\mathcal{A}$ where $\mathbf{t}_1 = (0,0,0)$, $\mathbf{t}_2 = (1,0,0)$, $\mathbf{t}_3 = (0,1,0)$ whose multiview ideal is:*

$$\begin{aligned} M_{\mathcal{A}} = \langle & y_1z_2 - y_2z_1, x_2z_3 - x_3z_2 + y_2z_3 - y_3z_2, \\ & x_1z_3 - x_3z_1, x_1x_3y_2 + x_1y_2y_3 - x_2x_3y_1 - x_3y_1y_2\rangle. \end{aligned}$$

The primary decompositions of $H^n_{\mathcal{A}}$, $F_{\mathcal{A}}$, and $Y_{\mathcal{A}}$ are

$$\begin{aligned} H^n_{\mathcal{A}} &= M_{\mathcal{A}} \cap \langle z_1, y_1, x_1\rangle \cap \langle z_2, y_2, x_2\rangle \cap \langle z_3, y_3, x_3\rangle, \\ Y_{\mathcal{A}} &= M_{\mathcal{A}} \cap \langle z_1, y_1, x_1\rangle \cap \langle y_3, y_2, x_3, x_2, z_3^2, z_2z_3, z_2^2\rangle \\ &\quad \cap \langle z_1, y_2, x_3, x_2, x_1, z_3^2, z_2z_3, z_2^2\rangle \\ &\quad \cap \langle z_1, y_3, y_2, y_1, x_3, z_3^2, z_2z_3, z_2^2\rangle, \end{aligned}$$

$$\begin{aligned}
F_{\mathcal{A}} = M_{\mathcal{A}} &\cap \langle y_2, y_1, x_2, x_1, z_2^2, z_1z_2, z_1^2\rangle \\
&\cap \langle y_3, y_1, x_3, x_1, z_3^2, z_1z_3, z_1^2\rangle \\
&\cap \langle z_2, z_1, y_3, y_2, y_1, x_3, z_3^2\rangle \\
&\cap \langle z_3, z_2, y_1, x_3 + y_3, x_2 + y_2, x_1, z_1^2\rangle \\
&\cap \big\langle y_2, x_2, z_3^2, z_2z_3, z_2^2, y_3z_3, y_3z_2, y_3^2, x_3z_3, \\
&\qquad x_3z_2, x_3y_3, x_3^2, x_1x_3 + x_1y_3 - x_3y_1\big\rangle \\
&\cap \big\langle z_3, y_2, x_2, z_2^2, z_1z_2, z_1^2, x_3z_2, x_3z_1, x_3^2, \\
&\qquad x_1z_2, x_1z_1, x_1y_3 - x_3y_1, x_1x_3, x_1^2\big\rangle \cap C
\end{aligned}$$

where C is a component minimally generated by 133 polynomials of total degree up to eight.

While each of $H^n_{\mathcal{A}}$, $F_{\mathcal{A}}$, and $Y_{\mathcal{A}}$ notably contains $M_{\mathcal{A}}$ as a component, the nature of their other components is worth further investigation. □

To analyze the extra components, we rely on several notions from commutative algebra, which we define next. The first notion is that of a *multigraded* ring. Consider the ring $\mathbb{C}[p_1, \ldots, p_n]$ endowed with the $\mathbb{Z}^n$-grading $\deg(w_i) = \mathbf{e}_i$ where $w_i \in \{x_i, y_i, z_i\}$ and $\mathbf{e}_i$ is the ith standard basis vector in $\mathbb{R}^n$. We say a polynomial in this ring is homogeneous if each of its terms have the same multidegree.

The *irrelevant ideal* in this grading, which we denote by $\mathfrak{m}$, is the intersection of the ideals $\mathfrak{m}_i := \langle x_i, y_i, z_i\rangle$:

$$\mathfrak{m} := \bigcap_{i=1}^{n} \mathfrak{m}_i = \bigcap_{i=1}^{n} \langle x_i, y_i, z_i\rangle. \tag{2.10}$$

Observe that $\mathfrak{m}$ is generated by all multilinear monomials of multidegree $(1, 1, \ldots, 1)$ and total degree n. It is the maximal ideal in the ring $\mathbb{C}[p_1, \ldots, p_n]$ generated by homogeneous elements of strictly positive multidegree.

The *radical* of an ideal I is the ideal $\sqrt{I} := \{f : f^k \in I \text{ for some } k \in \mathbb{N}\}$. If I is a homogeneous ideal then so is its radical, and $I \subseteq \sqrt{I}$. The *colon* of an ideal I with the ideal

J, denoted as $(I : J)$ is the set of all polynomials f such that $fg \in I$ for all $g \in J$, *i.e.* , $I : J = \{f \,:\, fJ \subseteq I\}$.

Recall that the projective varieties of the ideals $H^n_{\mathcal{A}}$, $F_{\mathcal{A}}$, and $Y_{\mathcal{A}}$ all agree and equal the multiview variety $\mathbf{V}(M_{\mathcal{A}})$. We can now state a first relationship among the ideals that follows easily from the projective Nullstellensatz in our multigraded setting, whose statement and proof will appear in Section 2.8.1.

Theorem 2.4.1. *For any $\mathcal{A}$ with pairwise distinct foci,*

(a) $\sqrt{H^n_{\mathcal{A}}} : \mathfrak{m} = M_{\mathcal{A}}$.

(b) $\sqrt{F_{\mathcal{A}}} : \mathfrak{m} = M_{\mathcal{A}}$.

(c) $\sqrt{Y_{\mathcal{A}}} : \mathfrak{m} = M_{\mathcal{A}}$ *when* $A_1 = [I \ 0]$.

Proof. See Section 2.8.1 □

In the language of algebraic geometry what this says is that $\sqrt{H^n_{\mathcal{A}}}, \sqrt{F_{\mathcal{A}}}$ and $\sqrt{Y_{\mathcal{A}}}$ all cut out the multiview variety *scheme-theoretically.* They are not equal as ideals but they agree in high enough multidegree with $M_{\mathcal{A}}$, see [22, pp 50].

We now strengthen Theorem 2.4.1 (a) and (b) to show that the operation of taking the radical is not needed, *i.e.* , $H^n_{\mathcal{A}} : \mathfrak{m} = M_{\mathcal{A}}$ and $F_{\mathcal{A}} : \mathfrak{m} = M_{\mathcal{A}}$. This means that $H^n_{\mathcal{A}}$ and $F_{\mathcal{A}}$ already cut out the multiview variety scheme-theoretically. Experimental evidence suggests that when $A_1 = [I \mid 0]$, such a result is also true for $Y_{\mathcal{A}}$, but an explicit proof is made difficult by the convoluted structure of the 2×2 minors of $\mathcal{A}^Y(p)$.

We first show that the simple structure of the primary decomposition of $H^n_{\mathcal{A}}$ observed in Example 2.4.1 holds in general.

Lemma 2.4.4. *For any camera arrangement $\mathcal{A}$ with pairwise distinct foci, $H^n_{\mathcal{A}} = M_{\mathcal{A}} \cap \mathfrak{m}$. In particular, $H^n_{\mathcal{A}}$ is a radical ideal with prime decomposition $M_{\mathcal{A}} \cap \mathfrak{m}_1 \cap \mathfrak{m}_2 \cap \cdots \cap \mathfrak{m}_n$.*

Proof. Suppose f is a generator of $\in H^n_{\mathcal{A}}$, *i.e.* , a maximal minor of $\mathcal{A}(p)$. Then $f \in \mathfrak{m}$. Also, since f vanishes on $\mathbf{V}(M_{\mathcal{A}})$, $f \in M_{\mathcal{A}}$. Therefore, $H^n_{\mathcal{A}} \subseteq M_{\mathcal{A}} \cap \mathfrak{m}$.

Now suppose $f \in M_{\mathcal{A}} \cap \mathfrak{m}$. Since $M_{\mathcal{A}}$ is generated by bifocals and trifocals $f = \sum \lambda_i r_i b_i + \sum \mu_j s_j t_j$ where b_i's are bifocals, t_j's are trifocals, r_i, s_j are monomials, and λ_i, μ_j are scalars. Further, since $f \in \mathfrak{m}$, every term in f is divisible by some generator $\prod_{i=1}^n w_i$ of $\mathfrak{m}$ where $w_i \in \{x_i, y_i, z_i\}$. Now consider $r_i b_i$. Since b_i involves only two cameras, it must be that r_i contains a variable w_i from each of the other $n-2$ cameras so that each term of $r_i b_i$ lies in $\mathfrak{m}$. This makes $r_i b_i$ a monomial multiple of a n-focal by Lemma 2.2.1. The same argument holds for $s_j t_j$. Thus, $f \in H^n_{\mathcal{A}}$. □

Proposition b3 in [47] proves that when $\mathcal{A}$ is minor-generic, $H^n_{\mathcal{A}}$ is a radical ideal. Lemma 2.4.4 shows that $H^n_{\mathcal{A}}$ is always a radical ideal under the weaker assumption of distinct foci.

Theorem 2.4.2. *For any camera arrangement $\mathcal{A}$ with pairwise distinct foci, $H^n_{\mathcal{A}} : \mathfrak{m} = M_{\mathcal{A}}$.*

Proof. We first note that $M_{\mathcal{A}} : \mathfrak{m} = M_{\mathcal{A}}$. Suppose $f \in M_{\mathcal{A}} : \mathfrak{m}$. Then $fu \in M_{\mathcal{A}}$ for any monomial generator u of $\mathfrak{m}$. Since $M_{\mathcal{A}}$ is prime and does not contain any monomials, $f \in M_{\mathcal{A}}$. Since $H^n_{\mathcal{A}} = M_{\mathcal{A}} \cap \mathfrak{m}$ by Lemma 2.4.4, $H^n_{\mathcal{A}} : \mathfrak{m} = M_{\mathcal{A}} : \mathfrak{m} = M_{\mathcal{A}}$. □

We now consider the Faugeras ideal $F_{\mathcal{A}}$ and prove that $F_{\mathcal{A}} : \mathfrak{m} = M_{\mathcal{A}}$. The nontrivial part is to argue that $M_{\mathcal{A}}$ is contained in $F_{\mathcal{A}} : \mathfrak{m}$. This fact relies on the following technical lemma, similar in flavor to Lemma 2.2.1, which shows that bifocals and trifocals can both be multiplied by any generator of $\mathfrak{m}$ to fall into $F_{\mathcal{A}}$.

Lemma 2.4.5. *(a) For $n = 2$ cameras, and any monomial $p_{1j}p_{2k}$, there exists a 4×4 minor f of $\mathcal{A}^F(p)$ such that $f = (-1)^{j+k} p_{1j}p_{2k} \det \mathcal{A}(p)$.*

(b) Let $n = 3$ and i_1, i_2, i_3 be pairwise distinct. Then for any trifocal $\det \mathcal{A}(p)_{\{p_{i_1 j_1} p_{i_2 j_2}\}}$ and any coordinate $p_{i_3 k}$, there exists a 4×4 minor f of $\mathcal{A}^F(p)$ such that

$$f = (-1)^k p_{i_3 k} \det \mathcal{A}(p)_{\{p_{i_1 j_1} p_{i_2 j_2}\}}.$$

Proof. See Section 2.8.2 both for the notation and the proof. $\square$

Theorem 2.4.3. *For any camera arrangement $\mathcal{A}$ with pairwise distinct foci, $F_{\mathcal{A}} : \mathfrak{m} = M_{\mathcal{A}}$.*

Proof. The containment $F_{\mathcal{A}} : \mathfrak{m} \subseteq M_{\mathcal{A}}$ follows as in Theorem 2.4.2 because $F_{\mathcal{A}} \subseteq M_{\mathcal{A}}$ and hence, $F_{\mathcal{A}} : \mathfrak{m} \subseteq M_{\mathcal{A}} : \mathfrak{m} = M_{\mathcal{A}}$. The other containment will follow by showing $H^2_{\mathcal{A}}, H^3_{\mathcal{A}} \subseteq F_{\mathcal{A}} : \mathfrak{m}$. For general camera arrangements with n cameras, recall that $F^2_{\mathcal{A}}$ (resp. $F^3_{\mathcal{A}}$) is the ideal generated by all 4×4 minors of $\mathcal{A}^F(p)$ that involve only two (resp. three) cameras. By Lemma 2.4.5(a), for any multilinear monomial $(\prod_{m=1}^n w_m)$ and any bifocal b_{ij}, $(\prod w_m)b_{ij} \in (f)$ for some Faugeras minor $f \in F^2_{\mathcal{A}}$, hence $H^2_{\mathcal{A}} \subseteq F_{\mathcal{A}} : \mathfrak{m}$. We address the trifocals in two cases. First consider the case when the two rows eliminated from $\mathcal{A}_{\{i,j,k\}}(p)$ to form a trifocal $t \in H^3_{\{i,j,k\}}$ come from the same camera, say without loss of generality, from camera i. In this case, $t = w_i b_{jk}$ for some w_i, and Lemma 2.4.5(a) again implies $t \in F_{\mathcal{A}} : \mathfrak{m}$. For the case when the two rows from $\mathcal{A}_{\{i,j,k\}}(p)$ to form $t \in H^3_{\{i,j,k\}}$ come from different cameras, Lemma 2.4.5(b) implies that, for any $(\prod w_m)$, $(\prod w_m)t \in (f)$ for some $f \in F^3_{\mathcal{A}}$. We conclude that $H^3_{\mathcal{A}} \subseteq F_{\mathcal{A}} : \mathfrak{m}$, as desired. $\square$

2.5 The Bifocal Ideal

We saw in Theorem 2.3.1 that the bifocals and trifocals together generate the multiview ideal when the camera foci are pairwise distinct. In this section, we investigate how imposing further conditions on the cameras can lead to an even simpler description of the multiview ideal. Heyden and Åström [27] and Trager *et al.* [46] show that when the camera foci are not all on a plane, the bifocals are necessary and sufficient to cut out the multiview variety. There has also been work to further reduce this description by considering the minimal number of bifocals needed ([27], [48]), though we will not address this question here. In this section, we focus on the ideal-theoretic relationship between the bifocal ideal $H^2_{\mathcal{A}}$ and the multiview ideal $M_{\mathcal{A}}$ when the camera foci are noncoplanar.

To motivate our investigation, we start with some examples. We say that a camera arrangement $\mathcal{A}$ is coplanar, noncoplanar or collinear if their foci have the corresponding property.

Example 2.5.1. *Consider the four noncoplanar translational camera arrangement $\mathcal{A}_1$ where* $\mathbf{t}_1 = (0,0,0)$*,* $\mathbf{t}_2 = (1,0,0)$*,* $\mathbf{t}_3 = (0,1,0)$*,* $\mathbf{t}_4 = (0,0,1)$*. Eliminating the variables q and λ_i from the ideal $\langle A_i q - \lambda_i p_i : i = 1,\ldots,n\rangle$, we observe $M_{\mathcal{A}_1}$ occurs as a component in $H^2_{\mathcal{A}_1}$*

$$\begin{aligned} H^2_{\mathcal{A}_1} = M_{\mathcal{A}_1} &\cap \langle x_2, y_2, z_2, x_1, x_3, x_4\rangle \cap \langle x_1, y_1, z_1, x_2 + y_2 + z_2, x_3 + y_3 + z_3, x_4 + y_4 + z_4\rangle \\ &\cap \langle x_3, y_3, z_3, y_1, y_2, y_4\rangle \cap \langle x_4, y_4, z_4, z_1, z_2, z_3\rangle. \end{aligned}$$

Example 2.5.2. *Consider the four coplanar translational camera arrangement $\mathcal{A}_2$ where* $\mathbf{t}_1 = (1,0,0)$*,* $\mathbf{t}_2 = (0,1,0)$*,* $\mathbf{t}_3 = (0,0,1)$*,* $\mathbf{t}_4 = (1/3,1/3,1/3)$*. We observe that $H^2_{\mathcal{A}_2} = M_{\mathcal{A}_2} \cap C$ where*

$$C = \langle x_4 + y_4 + z_4, x_3 + y_3 + z_3, x_2 + y_2 + z_2, x_1 + y_1 + z_1\rangle.$$

In Example 2.5.1, each extra component of $H^2_{\mathcal{A}_1}$ contains an irrelevant ideal $\mathfrak{m}_i$ and hence does not contribute to $\mathbf{V}(H^2_{\mathcal{A}_1})$. Saturating the bifocal ideal $H^2_{\mathcal{A}_1}$ with respect to the full irrelevant ideal $\mathfrak{m}$ removes these components. We will prove that this is always true when camera foci are noncoplanar. We begin by proving a series of three lemmas.

Lemma 2.5.1. *Suppose $\mathcal{A}$ is an arrangement of $n \geq 4$ cameras with pairwise distinct foci. Then $\mathcal{A}$ is noncoplanar $\implies H^n_{\mathcal{A}} \subseteq H^2_{\mathcal{A}}$.*

Proof. $\mathbf{n = 4, 5, 6}$. If $\mathcal{A}$ is noncoplanar, then there is some subset of four cameras that is noncoplanar. Order the cameras in $\mathcal{A}$ so that these are the cameras $A_1, \ldots, A_4$. By a change of coordinates on $\mathbb{P}^3$, we can send the foci of the cameras $A_1, \ldots, A_4$ to the foci of the cameras in $\mathcal{A}_1$ from Example 2.5.1. Then, by Lemma 2.2.4, applying $\mathbb{P}^2$ coordinate changes using some $\mathcal{G} \in (GL_3)^n$, we can assume that $\mathcal{A}$ is an arrangement of translational cameras. These transformations fix the first four cameras, and we think of the cameras A_i for $i \geq 5$ as variable, represented symbolically by their translations, and the implication can confirmed by direct calculation in Macaulay2.

$\mathbf{n = 7}$. In this case, the full computation is too expensive. To make the computation feasible, we split the proof into two cases, depending on whether the arrangement has five collinear cameras or not.

Case I: If a noncoplanar arrangement of seven cameras has at most four collinear cameras, then every four camera subset can be augmented with two additional cameras to get a noncoplanar arrangement of six cameras. Thus every 7-focal of such an arrangement, which looks like $w_i w_j w_k q$ for some quadrifocal q, has the form of a 6-focal from a noncoplanar arrangement, say $w_i w_j q$, multiplied by a coordinate w_k. The $n = 6$ case shows that $w_i w_j q$ is generated by 2-focals, hence $w_i w_j w_k q$ is generated by 2-focals.

Case II: We now consider the case of noncoplanar seven camera arrangements in which five cameras are collinear. In this case, by a proper choice of camera ordering and $\mathbb{P}^3$ coordinate change, we can assume the translations of A_5, A_6, A_7 are of the form $\mathbf{t}_5 = (\lambda_5, 0, 0)^\top, \mathbf{t}_5 = (\lambda_6, 0, 0)^\top, \mathbf{t}_5 = (\lambda_7, 0, 0)^\top$ where the λ_i are symbolic. This makes A_1, A_2, A_5, A_6, A_7 collinear. The choice to take the line that the cameras lie on to be the x axis is arbitrary, but can be made without loss of generality. This arrangement is now described by few enough variables to enable a direct computation showing that $H^7_{\mathcal{A}} \subseteq H^2_{\mathcal{A}}$.

$\mathbf{n \geq 8}$. Now suppose $n \geq 8$ and f is an n-focal of $\mathcal{A}$. Recall that f involves all n cameras but at most four cameras can contribute two rows to the matrix whose determinant is f. At one extreme, these four cameras maybe $A_1, \ldots, A_4$ and at the other extreme they might be four cameras different from the first four, which we call $A_5, \ldots, A_8$. Thus the n-focal $f \in H^n_{\mathcal{A}}$ is a monomial multiple of a 8-focal $g = mq$ of $\{A_1, \ldots, A_4, A_5, \ldots, A_8\}$ where where q is a quadrifocal and m is a monomial.

If the four cameras contributing to q involve $A_1, \ldots, A_4$, then g is a multiple of a 7-focal from noncoplanar cameras. On the other hand, if $q \in H^4_{A_5,\ldots,A_8}$, then q can be generated by the trifocals of $A_5, \ldots, A_8$ by Lemma 2.3.2:

$$g = m \left(\sum_{t_i \in H^3_{A_5,\ldots,A_8}} h_i t_i \right) = \sum_{t_i \in H^3_{A_5,\ldots,A_8}} h_i (m t_i).$$

In particular, this shows that g can be generated from 7-focals, mt_i. These come from noncoplanar seven camera arrangements because $A_1, \dots, A_4$ are noncoplanar. In either case, we know that such 7-focals can be generated by 2-focals, hence $g \in H^2_{\mathcal{A}}$. It follows that $f \in H^2_{\mathcal{A}}$, as desired. $\square$

Lemma 2.5.2. *Suppose $\mathcal{A}$ is an arrangement of $n \geq 4$ cameras with pairwise distinct foci. Then $H^n_{\mathcal{A}} \subseteq H^2_{\mathcal{A}} \implies M_{\mathcal{A}} = H^2_{\mathcal{A}} : \mathfrak{m}$.*

Proof. If $f \in H^2_{\mathcal{A}} : \mathfrak{m}$, then $f(\prod z_i) \in H^2_{\mathcal{A}} \subseteq M_{\mathcal{A}}$, vanishes on $\mathbf{V}(M_{\mathcal{A}})$. Since $M_{\mathcal{A}}$ is prime and does not contain any monomials, $f \in M_{\mathcal{A}}$. Therefore, $H^2_{\mathcal{A}} : \mathfrak{m} \subseteq M_{\mathcal{A}}$. For the other containment, by Theorem 2.3.1, it suffices to show that $H^2_{\mathcal{A}}$ and $H^3_{\mathcal{A}}$ are contained in $H^2_{\mathcal{A}} : \mathfrak{m}$. It is clear that $H^2_{\mathcal{A}} \subseteq H^2_{\mathcal{A}} : \mathfrak{m}$. By Lemma 2.2.1, multiplying any $t \in H^3_{\mathcal{A}}$ by a generator $\prod w_i$ of $\mathfrak{m}$ yields a monomial multiple of an n-focal. By assumption, this n-focal lies in $H^2_{\mathcal{A}}$. Thus, $t \in H^2_{\mathcal{A}} : \mathfrak{m}$ and $M_{\mathcal{A}} \subseteq H^2_{\mathcal{A}} : \mathfrak{m}$. $\square$

Lemma 2.5.3. *Suppose $\mathcal{A}$ is an arrangement of $n \geq 4$ cameras with pairwise distinct foci. Then $M_{\mathcal{A}} = H^2_{\mathcal{A}} : \mathfrak{m} \implies \mathcal{A}$ is noncoplanar.*

Proof. We prove the contrapositive, namely that if $\mathcal{A}$ is coplanar then $M_{\mathcal{A}} \neq H^2_{\mathcal{A}} : \mathfrak{m}$. We will construct a point $\mathbf{p} \in \mathbf{V}(H^2_{\mathcal{A}} : \mathfrak{m}) \smallsetminus \mathbf{V}(M_{\mathcal{A}})$, from which the result will follow.

Let $\mathbf{n} \in \mathbb{P}^3$ be the normal vector of a plane containing the foci of the cameras in $\mathcal{A}$. If the foci are not collinear then $\mathbf{n}$ is unique, otherwise we choose any plane containing the foci and its normal $\mathbf{n}$. Let $l_i \subseteq \mathbb{P}^2$ denote the image of the plane $\mathbf{n}^\perp$ in camera i, and let $\mathbf{e}_{i,j}$ denote the image of the focal point of camera j in image i. Then $\mathbf{e}_{i,j} \in l_i$ since the focal point of camera j lies in $\mathbf{n}^\perp$. Choose $\mathbf{p}_1 \in l_1 \smallsetminus \{\mathbf{e}_{1,2}, \mathbf{e}_{1,3}\}$ and $\mathbf{p}_2 \in l_2 \smallsetminus \{\mathbf{e}_{2,1}, \mathbf{e}_{2,3}\}$. Then there is a unique world point $\mathbf{q}$ on $\mathbf{n}^\perp$ whose images in cameras 1 and 2 are $\mathbf{p}_1$ and $\mathbf{p}_2$. Let $\widetilde{\mathbf{p}}_3 \in l_3$ be the (unique) image of $\mathbf{q}$ in camera 3. Then $\mathbf{p}_1, \mathbf{p}_2, \widetilde{\mathbf{p}}_3$ satisfy trifocal constraints. Choose $\mathbf{p}_3 \in l_3 \smallsetminus \{\widetilde{\mathbf{p}}_3\}$ and some $\mathbf{p}_i \in l_i$ for $i \geq 4$. By construction, $\mathbf{p} \notin \mathbf{V}(M_{\mathcal{A}})$. Since the cameras are coplanar, the epipolar plane given by $\mathbf{q}$ and any two cameras i and j is $\mathbf{n}^\perp$ for any pair i, j. By choosing $\mathbf{p}_i \in l_i$ for all i, we force every bifocal polynomial to vanish on

$\mathbf{p}$. Therefore by construction, $\mathbf{p} \in \mathbf{V}(H^2_{\mathcal{A}}) \smallsetminus \mathbf{V}(M_{\mathcal{A}})$, but since $\mathbf{V}(H^2_{\mathcal{A}}) = \mathbf{V}(H^2_{\mathcal{A}} : \mathfrak{m})$, we conclude that $H^2_{\mathcal{A}} : \mathfrak{m} \neq M_{\mathcal{A}}$. □

Together, Lemmas 2.5.1, 2.5.2, 2.5.3 imply the following theorem.

Theorem 2.5.1. *Suppose $\mathcal{A}$ is an arrangement of $n \geq 4$ cameras with pairwise distinct foci. Then the following are equivalent.*

(a) $\mathcal{A}$ is noncoplanar.

(b) $H^n_{\mathcal{A}} \subseteq H^2_{\mathcal{A}}$.

(c) $M_{\mathcal{A}} = H^2_{\mathcal{A}} : \mathfrak{m}$.

We now make some observations about Theorem 2.5.1.

Theorem 6.1 in [27] observes that $\mathbf{V}(H^2_{\mathcal{A}}) = \mathbf{V}(M_{\mathcal{A}})$ for noncoplanar $\mathcal{A}$ while Proposition 5 (2) in [46] further shows that $\mathbf{V}(H^2_{\mathcal{A}}) = \mathbf{V}(M_{\mathcal{A}})$ is equivalent to the foci of $\mathcal{A}$ being noncoplanar. Our Theorem 2.5.1 proves the analogous ideal statement, namely that noncoplanarity of foci is equivalent to $M_{\mathcal{A}} = H^2_{\mathcal{A}} : \mathfrak{m}$.

Example 2.5.2 shows how Theorem 2.5.1 fails when $\mathcal{A}$ is coplanar. The bifocal ideal $H^2_{\mathcal{A}_2}$ contains the component $\langle x_1 + y_1 + z_1, x_2 + y_2 + z_2, x_3 + y_3 + z_3, x_4 + y_4 + z_4 \rangle$, which cannot be removed by saturating with respect to $\mathfrak{m}$. Its variety cuts out the projections of the plane containing the foci of $\mathcal{A}_2$ in each camera image. This plane in $\mathbb{P}^3$ has normal vector $(1, 1, 1, -1)$. The following example shows that further degeneracy occurs when camera foci are collinear.

Example 2.5.3. *Consider the four collinear translational camera arrangement $\mathcal{A}_3$ where $\mathbf{t}_1 = (0, 0, 0)$, $\mathbf{t}_2 = (1, 0, 0)$, $\mathbf{t}_3 = (2, 0, 0)$, $\mathbf{t}_4 = (3, 0, 0)$. Here, $H^2_{\mathcal{A}_3} \subseteq M_{\mathcal{A}_3}$, but both ideals are prime, so $M_{\mathcal{A}_3}$ cannot occur as a component of $H^2_{\mathcal{A}_3}$. In addition, the dimension of $H^2_{\mathcal{A}_3}$ is one larger than that of $M_{\mathcal{A}_3}$. This is explained by the fact that there is an entire one-dimensional family of planes that contains the camera centers of $\mathcal{A}_3$.*

As seen in the above examples and discussion, the relation between $H^2_{\mathcal{A}}$ and $M_{\mathcal{A}}$ can be complicated when camera centers are coplanar or collinear. Determining the exact relationship between ideals in these degenerate settings would be an interesting problem for the future.

In Theorem 2.5.1 we showed that when cameras are noncoplanar, the n-focal ideal becomes a subset of the 2-focal ideal. We now give an example to show that this containment need not hold for $H^k_{\mathcal{A}}$ where $n > k > 2$. The construction relies on having three of five cameras being collinear.

Example 2.5.4. *Consider the five translational camera arrangement $\mathcal{B}$ with $\mathbf{t}_1 = (0,0,0), \mathbf{t}_2 = (0,0,1), \mathbf{t}_3 = (0,0,2), \mathbf{t}_4 = (0,1,0), \mathbf{t}_5 = (0,0,1)$. Theorem 2.5.1 shows that $H^5_{\mathcal{B}} \subseteq H^2_{\mathcal{B}}$ since $\mathcal{B}$ is noncoplanar. However the following trifocal from B_1, B_2, B_3,*

$$t = -x_1y_2y_3 + 2\,x_2y_1y_3 - x_3y_1y_2$$

is not in $H^2_{\mathcal{B}}$. Similarly, the quadrifocal,

$$q = x_4t = -x_1x_4y_2y_3 + 2\,x_2x_4y_1y_3 - x_3x_4y_1y_2,$$

from cameras B_1, B_2, B_3, B_4 is not in $H^2_{\mathcal{B}}$.

2.6 Finite Images

The results of the previous sections have important practical consequences when we restrict attention to the set of all finite images, that is to all $(\mathbf{p_1}, \ldots, \mathbf{p_n}) \in \mathbf{V}(M_{\mathcal{A}})$ with $z_i \neq 0$ for all i. The vanishing ideal of this affine patch is obtained by dehomogenizing $M_{\mathcal{A}}$ with respect to the variables z_i from each image plane. We call this the *affine multiview ideal* of $\mathcal{A}$ and denote it $\pi(M_{\mathcal{A}})$, where $\pi : \mathbb{C}[x_i, y_i, z_i] \to \mathbb{C}[x_i, y_i]$ is the map setting each z_i to 1. From Theorem 2.3.1, we see that $\pi(M_{\mathcal{A}})$ is generated by dehomogenized bifocals and dehomogenized trifocals when the foci of $\mathcal{A}$ are pairwise distinct.

Corollary 2.6.1. *If $\mathcal{A}$ is a camera arrangement with pairwise distinct foci, then $\pi(M_{\mathcal{A}}) = \pi(H^2_{\mathcal{A}}) + \pi(H^3_{\mathcal{A}})$.*

Using the following fact about dehomogenizing colon ideals, the results of Section 2.4 yield a nice relation among $\pi(H^n_{\mathcal{A}}), \pi(F_{\mathcal{A}}), \pi(Y_{\mathcal{A}})$, and the affine multiview ideal, $\pi(M_{\mathcal{A}})$.

Lemma 2.6.1. *For ideals $I, J \subset \mathbb{C}[x_i, y_i, z_i]$, $\pi(I : J) = \pi(I) : \pi(J)$.*

Proof. If $f \in \pi(I : J)$, then $f = \pi(g)$ for some g which satisfies $gh \in I$ for all $h \in J$. Therefore $f\pi(h) = \pi(g)\pi(h) = \pi(gh) \in \pi(I)$ for any $h \in J$, proving $f \in \pi(I) : \pi(J)$. If $f \in \pi(I) : \pi(J)$, then for any $h \in J$, $f\pi(h) \in \pi(I)$, *i.e.* , there exists $g \in I$ such that $f\pi(h) = \pi(g)$. Denote the homogenization of f with respect to $z_1, \dots, z_n$ by $\widetilde{f}$. We claim that $\widetilde{f} \in I : J$. Indeed for any $h \in J$, $\pi(\widetilde{f}h) = \pi(\widetilde{f})\pi(h) = f\pi(h) = \pi(g)$ for some $g \in I$. Homogenizing both sides, we get $\widetilde{f}h = g \in I$, and we conclude that $\pi(I) : \pi(J) \subseteq \pi(I : J)$ □

Corollary 2.6.2. *If $\mathcal{A}$ is a camera arrangement with pairwise distinct foci, then $\pi(M_{\mathcal{A}}) = \pi(H^n_{\mathcal{A}}) = \pi(F_{\mathcal{A}}) = \pi(\sqrt{Y_{\mathcal{A}}})$.*

Proof. Lemma 2.6.1 implies that $\pi(I : \mathfrak{m}) = \pi(I) : (1) = \pi(I)$ for any ideal I. Dehomogenizing Theorems 2.4.2, 2.4.3, and 2.4.1, each equality follows. □

Observe that the last equality in Corollary 2.6.2 requires $A_1 = [I\ 0]$. Geometrically, Corollary 2.6.2 shows that while the homogenous ideals $H^n_{\mathcal{A}}, F_{\mathcal{A}}, Y_{\mathcal{A}}$, and $M_{\mathcal{A}}$ do not coincide, they are the same away from the origin in each image plane. In particular, this is the case on the affine patch $\{\mathbf{p} \in \mathbb{P}^{2n} : z_1 = \cdots = z_n = 1\}$ corresponding to finite image data.

Using Theorem 2.5.1 we see that, when $\mathcal{A}$ is noncoplanar, the dehomogenized bifocals alone suffice to generate the affine multiview ideal $\pi(M_{\mathcal{A}})$.

Corollary 2.6.3. *Suppose $\mathcal{A}$ is a noncoplanar camera arrangement with pairwise distinct foci. Then*

$$\pi(M_{\mathcal{A}}) = \pi(H^2_{\mathcal{A}}).$$

Proof. Dehomogenizing the result of Theorem 2.5.1, we get $\pi(M_{\mathcal{A}}) = \pi(H^2_{\mathcal{A}} : \mathfrak{m}) = \pi(H^2_{\mathcal{A}}) : \pi(\mathfrak{m}) = \pi(H^2_{\mathcal{A}})$. □

Corollary 2.6.3 shows that $\pi(M_{\mathcal{A}})$ is generated by quadratics whenever $\mathcal{A}$ satisfies the noncoplanarity assumption. This observation was used in [4] to create a semidefinite programming relaxation of the triangulation problem which is can be seen as minimizing Euclidean distance from an observed noisy data point to the affine multiview variety. It was shown that when the noise is small, the semidefinite relaxation solves triangulation. Of course, Corollary 2.6.2 needs the foci of the cameras to be noncoplanar and indeed, the experiments in [4] show that the quality of the semidefinite programming solution deteriorates as the foci become coplanar and then collinear.

Geometrically, we can understand how the quality of the relaxation deteriorates because the bifocal ideal cuts out more than the multiview variety for coplanar arrangements. In the coplanar case, the bifocal ideal cuts out the image of the plane that contains the camera centers. These points are not the images of true 3D points. It is therefore possible that the nearest point problem yields a spurious solution on this extra component. Similarly, in the collinear case, the bifocal ideal cuts out a strictly larger variety than just the multiview variety. In this case, the dimension of the vanishing set of the bifocal ideal is one larger than the multiview variety.

2.7 Summary

The multiview variety is a foundational geometric object in multiview geometry and understanding its vanishing ideal $M_{\mathcal{A}}$ precisely is important for any algebraic algorithm that solves problems on this variety. There have been many partial results about the algebraic structure of the multiview variety. The aim of our paper is to put them all into a unified algebraic setting and give a complete description of $M_{\mathcal{A}}$.

Our main result is that when the foci of the cameras are pairwise distinct, $M_{\mathcal{A}}$ is generated by the bifocal and trifocal polynomials of $\mathcal{A}$ (Theorem 2.3.1). The proof requires an understanding of the behavior of coordinate changes on k-focal ideals (Lemma 2.2.4), and translational cameras (Lemma 2.3.2). The main result holds for Euclidean cameras as well (Corollary 2.3.3). We also give an example to illustrate that the assumption of distinct foci cannot be relaxed for this result to hold (Example 2.3.1).

Next we study three sets of polynomials that have been proposed to cut out the multiview variety, by Heyden-Åström, Faugeras and Ma et. al. respectively. We show that the ideals generated by these polynomials are all properly contained in $M_{\mathcal{A}}$. We establish the exact algebraic relationships between the above ideals and $M_{\mathcal{A}}$ (Theorems 2.4.1, 2.4.2 and 2.4.3).

We then prove that if the camera foci are assumed to be noncoplanar, then in fact $M_{\mathcal{A}}$ is the saturation of the bifocal ideal by the irrelevant ideal (Theorem 2.5.1). In this situation the n-focal ideal is a subset of the bifocal ideal.

Finally we prove that the dehomogenization of the ideals by Heyden-Åström, Faugeras and Ma et. al. all agree with the dehomogenization of $M_{\mathcal{A}}$ (Corollary 2.6.2). Similarly, under noncoplanarity of foci, the bifocal ideal also has the same dehomogenization (Corollary 2.6.3). This means that all of these ideals cut out the space of finite images.

2.8 Technical Proofs

2.8.1 Multigraded Projective Nullstellensatz

In this section, we state and prove the projective Nullstellensatz in our multigraded setting, which we use to prove Theorem 2.4.1 in Section 2.4. Let $I \subseteq \mathbb{C}[p_1, \ldots, p_n]$ be homogeneous with respect to the $\mathbb{Z}^n$-grading $\deg w_i = \mathbf{e}_i$. To be clear about projective versus affine varieties, we define $\mathbf{V}_{\mathbb{P}}(I) := \mathbf{V}(I) = \{\mathbf{p} \in (\mathbb{P}^2)^n : f(\mathbf{p}) = 0 \text{ for all } f \in I\}$, and for a set $S \subseteq (\mathbb{P}^2)^n$, we define

$$\mathbf{I}_{\mathbb{P}}(S) = \{f \in \mathfrak{m} : f(\mathbf{p}) = 0 \text{ for all } \mathbf{p} \in S\}.$$

We say that $\mathbf{V}_{\mathbb{P}}(I)$ is the projective vanishing set of I in $(\mathbb{P}^2)^n$ and $\mathbf{I}_{\mathbb{P}}(S)$ is the largest homogeneous ideal vanishing on S contained in $\mathfrak{m}$. While we force $\mathbf{I}_{\mathbb{P}}(S) \subseteq \mathfrak{m}$, it also makes sense to consider the largest homogeneous ideal vanishing on S without intersecting with $\mathfrak{m}$. As before we denote this ideal by $\mathbf{I}(S)$, and notice that $\mathbf{I}_{\mathbb{P}}(S) = \mathbf{I}(S) \cap \mathfrak{m}$. In the usual grading on $\mathbb{C}[p_1, \ldots, p_n]$, a vanishing ideal $\mathbf{I}(S)$ is homogeneous in the usual sense which means that it is contained in the usual irrelevant ideal $\langle x_1, y_1, z_1, \ldots, x_n, y_n, z_n \rangle$. Under the multi-grading, $\mathbf{I}_{\mathbb{P}}(S)$ is required to be in the corresponding irrelevant ideal $\mathfrak{m}$. We will use the following variant of the Nullstellensatz.

Lemma 2.8.1. *For any homogeneous ideal $I \subseteq \mathbb{C}[p_1, \dots, p_n]$ such that $I \subseteq \mathfrak{m}$, $\mathbf{I}_{\mathbb{P}}(\mathbf{V}_{\mathbb{P}}(I)) = \sqrt{I}$.*

Proof. Define the affine operations

$$\mathbf{V}_{\mathbb{A}}(I) = \{\mathbf{p} \in (\mathbb{A}^3)^n : f(\mathbf{p}) = 0 \text{ for all } f \in I\}$$
$$\mathbf{I}_{\mathbb{A}}(S) = \{f \in \mathbb{C}[p_1, \dots, p_n] : f(\mathbf{p}) = 0 \text{ for all } \mathbf{p} \in S\}$$

where we treat S as a subset of $(\mathbb{A}^3)^n$. We will use the affine version of the Nullstellensatz on the cone over $V := \mathbf{V}_{\mathbb{P}}(I)$, *i.e.* , the set $C_V = \mathbf{V}_{\mathbb{A}}(I) \subseteq (\mathbb{A}^3)^n$. We claim that

$$\mathbf{I}_{\mathbb{A}}(C_V) = \mathbf{I}_{\mathbb{P}}(V). \tag{2.11}$$

First suppose $f \in \mathbf{I}_{\mathbb{A}}(C_V)$. Given $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_n) \in V$, all homogeneous coordinates of $\mathbf{p}$, represented by scalings $(\lambda_1\mathbf{p}_1, \dots, \lambda_n\mathbf{p}_n)$, lie in C_V, so f vanishes for all homogeneous coordinates of $\mathbf{p}$. This means that the homogeneous components $f_{i_1,\dots,i_n}$ of f, consisting of all terms with multidegree $(i_1, \dots, i_n)$, vanish at $\mathbf{p}$, so $f \in \mathbf{I}(V)$, hence $\mathbf{I}_{\mathbb{A}}(C_V) \subseteq \mathbf{I}(V)$. By the Nullstellensatz in $(\mathbb{A}^3)^n$, $\mathbf{I}_{\mathbb{A}}(C_V) = \mathbf{I}_{\mathbb{A}}(\mathbf{V}_{\mathbb{A}}(I)) = \sqrt{I}$, and by the assumption that $I \subseteq \mathfrak{m}$, $\sqrt{I} \subseteq \sqrt{\mathfrak{m}} = \mathfrak{m}$. This shows that $\mathbf{I}_{\mathbb{A}}(C_V) \subseteq \mathbf{I}(V) \cap \mathfrak{m} = \mathbf{I}_{\mathbb{P}}(V)$.

Conversely, suppose $f \in \mathbf{I}_{\mathbb{P}}(V)$. Since any point $\mathbf{p}$ of C_V such that $\mathbf{p}_i \neq 0$ for all i gives homogeneous coordinates for a point in V, it follows that f vanishes on $C_V \setminus \bigcup_{i=1}^n \mathbb{A}^3 \times \cdots \times \{0\}_i \times \cdots \times \mathbb{A}^3$. We need to show that f vanishes on each of the sets $\mathbb{A}^3 \times \cdots \times \{0\}_i \times \cdots \times \mathbb{A}^3$. Since $f \subseteq \mathfrak{m}$, it has strictly positive multidegree, and every monomial in f contains at least one coordinate from each copy of $\mathbb{A}^3$. Setting all 3 coordinates to zero in any $\mathbb{A}^3$ forces f to be zero, so we conclude that $f \in \mathbf{I}_{\mathbb{A}}(C_V)$. Finally, from (2.11), we conclude

$$\sqrt{I} = \mathbf{I}_{\mathbb{A}}(\mathbf{V}_{\mathbb{A}}(I)) = \mathbf{I}_{\mathbb{A}}(C_V) = \mathbf{I}_{\mathbb{P}}(V) = \mathbf{I}_{\mathbb{P}}(\mathbf{V}_{\mathbb{P}}(I)).$$

$\square$

Corollary 2.8.1. *For any homogeneous ideal* $I \subseteq \mathbb{C}[p_1, \ldots, p_n]$, $\mathbf{I}_{\mathbb{P}}(\mathbf{V}_{\mathbb{P}}(I)) = \sqrt{I} \cap \mathfrak{m}$.

Proof. Observe that

$$\mathbf{V}_{\mathbb{P}}(I \cap \mathfrak{m}) = \mathbf{V}_{\mathbb{P}}(I) \cup \mathbf{V}_{\mathbb{P}}(\mathfrak{m}) = \mathbf{V}_{\mathbb{P}}(I)$$

and

$$\sqrt{I \cap \mathfrak{m}} = \sqrt{I} \cap \sqrt{\mathfrak{m}} = \sqrt{I} \cap \mathfrak{m} \subseteq \mathfrak{m}$$

Therefore by Lemma 2.8.1, $\mathbf{I}_{\mathbb{P}}(\mathbf{V}_{\mathbb{P}}(I)) = \sqrt{I} \cap \mathfrak{m}$. □

Corollary 2.8.2. *For any* $\mathcal{A}$ *with pairwise distinct foci,*

$$M_{\mathcal{A}} \cap \mathfrak{m} = \sqrt{H^n_{\mathcal{A}}} \cap \mathfrak{m} = \sqrt{F_{\mathcal{A}}} \cap \mathfrak{m} = \sqrt{Y_{\mathcal{A}}} \cap \mathfrak{m}.$$

Proof. We have already shown in Section 2.4 that $\mathbf{V}_{\mathbb{P}}(H^n_{\mathcal{A}}) = \mathbf{V}_{\mathbb{P}}(F_{\mathcal{A}}) = \mathbf{V}_{\mathbb{P}}(Y_{\mathcal{A}}) = \mathbf{V}_{\mathbb{P}}(M_{\mathcal{A}})$. Since $M_{\mathcal{A}}$ is radical, the result follows by Corollary 2.8.1. □

We can now prove Theorem 2.4.1, restated here, from the main body of the paper.

Theorem 2.4.1. *For any* $\mathcal{A}$ *with pairwise distinct foci,*

(a) $\sqrt{H^n_{\mathcal{A}}} : \mathfrak{m} = M_{\mathcal{A}}$

(b) $\sqrt{F_{\mathcal{A}}} : \mathfrak{m} = M_{\mathcal{A}}$

(c) $\sqrt{Y_{\mathcal{A}}} : \mathfrak{m} = M_{\mathcal{A}}$ *when* $A_1 = [I \mid 0]$

Proof. Taking colon ideal with $\mathfrak{m}$, the desired result follows from Corollary 2.8.2 and the fact that $M_{\mathcal{A}} : \mathfrak{m} = M_{\mathcal{A}}$, which was proven in Theorem 2.4.2. □

2.8.2 Proof of Lemma 2.4.5

In this section, we elaborate on the technical details used to prove Theorem 2.4.3. Recall that the nontrivial statement there was that bifocals and trifocals can be multiplied by any generator of $\mathfrak{m}$ to fall into $F_{\mathcal{A}}$. This requires understanding the 4×4 minors of $\mathcal{A}^F(p)$ for which we once again invoke the Cauchy-Binet formula and the observation that $\mathcal{A}^F(p) = P(p)\mathcal{A}$ from (2.7).

First we characterize certain 4×4 minors of $P(p)$. Let p_{ij} denote the jth coordinate of p_i, *i.e.* , $p_{i1} = x_i$, $p_{i2} = y_i$, and $p_{i3} = z_i$. Having the subscript (resp. superscript) p_{ij} on $P(p)$ indicates eliminating from $P(p)$ the unique row (resp. column) of $[p_i]_\times$ that does not contain p_{ij}. On the other hand, having the subscript p_{ij} on the matrices $\mathcal{A}$ and $\mathcal{A}(p)$ will stand for eliminating the unique row of the matrix containing p_{ij}.

We will only need to consider the 4×4 minors of $P(p)$ when $n = 2$ and $n = 3$. Let $R_i, C_i \subseteq \{p_{i1}, p_{i2}, p_{i3}\}$ denote collections of coordinates, and write $R = \bigcup_{i=1}^n R_i$, $C = \bigcup_{i=1}^n C_i$. When $n = 2$, a 4×4 minor of $P(p)$ is $\det(P(p)_R^C)$ for some R, C of size $|R| = |C| = 2$, and when $n = 3$, $|R| = |C| = 5$. Observe that if $|R_i| \neq |C_i|$ for any i, then the submatrix $P(p)_R^C$ has at least two linearly dependent rows or columns, yielding a zero minor. When $|R_i| = |C_i|$ for all i, $P(p)_R^C$ is block diagonal, so $\det(P(p)_R^C) = \prod_{i=1}^n \det(([p_i]_\times)_{R_i}^{C_i})$.

Lemma 2.8.2. *Let $n = 2$. The nonzero 4×4 minors of $P(p)$ are determined by collections of coordinates R, C with $|R_1| = |C_1| = |R_2| = |C_2| = 1$. For $R = \{p_{1j}, p_{2k}\}$ and $C = \{p_{1l}, p_{2m}\}$, the 4×4 minor $\det(P(p)_R^C)$ is the monomial*

$$\det(P(p)_R^C) = (-1)^{j+k+l+m} p_{1j} p_{2k} p_{1l} p_{2m}.$$

Proof. As noted above, if $|R_i| \neq |C_i|$ for either i, then $\det(P(p)_R^C) = 0$, whereas if $|R_i| = |C_i| = 2$ for either i, then $P(p)_R^C$ has a rank 2 block on its diagonal, hence $\det(P(p)_R^C) = 0$, proving the first statement. For $R = \{p_{1j}, p_{2k}\}$ and $C = \{p_{1l}, p_{2m}\}$, the 4×4 minor $\det P(p)_R^C$ is

$$\begin{aligned}
\det P(p)_R^C &= \det(([p_1]_\times)_{R_1}^{C_1} \det(([p_2]_\times)_{R_2}^{C_2}) \\
&= ((-1)^{j+l} p_{1j} p_{1l})((-1)^{k+m} p_{2k} p_{2m}) \\
&= (-1)^{j+k+l+m} p_{1j} p_{2k} p_{1l} p_{2m}.
\end{aligned}$$

□

Lemma 2.8.3. *Let $n = 3$. Suppose $|R_3| = |C_3| = 1$, and $|R_1| = |C_1| = |R_2| = |C_2| = 2$. For $R_3 = \{p_{3j}\}, C_3 = \{p_{3k}\}$, the 4×4 minor $\det(P(p)_R^C)$ is the monomial*

$$\begin{cases} (-1)^{j+k+l+m} p_{3j} p_{3k} p_{1l} p_{2m} & \text{if } R_1 \neq C_1, R_2 \neq C_2 \\ 0 & \text{otherwise.} \end{cases}$$

where p_{1l} is the coordinate common to R_1 and C_1 and p_{2m} is the coordinate common to R_2 and C_2.

Proof. When $R_i = C_i$ as sets for $i = 1$ or $i = 2$, then $([p_i]_\times)_{R_i}^{C_i} = 0$, hence $\det P(p)_R^C = \prod_{i=1}^n \det(([p_i]_\times)_{R_i}^{C_i}) = 0$. On the other hand, when $R_1 \neq C_1$, $\det(([p_1]_\times)_{R_1}^{C_1}) = (-1)^l p_{1l}$ where $p_{1l} = R_1 \cap C_1$. Similarly $\det(([p_2]_\times)_{R_2}^{C_2}) = (-1)^m p_{2m}$ where $p_{2m} = R_2 \cap C_2$ when $R_2 \neq C_2$. $\square$

We now show that bifocals and trifocals can both be multiplied by any generator of $\mathfrak{m}$ to fall into $F_{\mathcal{A}}$.

Lemma 2.8.4. *(a) For $n = 2$ cameras, and any monomial $p_{1j} p_{2k}$, there exists a 4×4 minor f of $\mathcal{A}^F(p)$ such that $f = (-1)^{j+k} p_{1j} p_{2k} \det(\mathcal{A}(p))$.*

(b) Let $n = 3$ and i_1, i_2, i_3 be pairwise distinct. Then for any trifocal $\det(\mathcal{A}(p)_{\{p_{i_1 j_1} p_{i_2 j_2}\}})$ and any coordinate $p_{i_3 k}$, there exists a 4×4 minor f of $\mathcal{A}^F(p)$ such that

$$f = (-1)^k p_{i_3 k} \det(\mathcal{A}(p)_{\{p_{i_1 j_1} p_{i_2 j_2}\}}).$$

Proof. (a) Fix some $p_{1j} p_{2k}$. Since $n = 2$, $P(p)\mathcal{A}$ is a 6×4 matrix and we need to delete two rows to get a 4×4 minor. Using Lemma 2.8.2 and Cauchy-Binet, the result follows from the computation below:

$$
\begin{aligned}
f &= \det\left(P(p)_{\{p_{1j},p_{2k}\}}\mathcal{A}\right) \\
&= \sum_{|C|=2} \det\left(P(p)^{C}_{\{p_{1j},p_{2k}\}}\right)\det\left(\mathcal{A}_C\right) \\
&= \sum_{|C_1|=|C_2|=1} \det\left(P(p)^{C}_{\{p_{1j},p_{2k}\}}\right)\det\left(\mathcal{A}_C\right) \\
&= \sum_{1\le l,m\le 3} \det\left(([p_1]_\times)^{\{p_{1l}\}}_{\{p_{1j}\}}\right)\det\left(([p_2]_\times)^{\{p_{2m}\}}_{\{p_{2k}\}}\right)\times\det\left(\mathcal{A}_{\{p_{1l},p_{2m}\}}\right) \\
&= \sum_{1\le l,m\le 3} (-1)^{j+k+l+m} p_{1j}p_{2k}p_{1l}p_{2m}\det\left(\mathcal{A}_{\{p_{1l},p_{2m}\}}\right) \\
&= (-1)^{j+k}p_{1j}p_{2k}\sum_{1\le l,m\le 3}(-1)^{l+m}p_{1l}p_{2m}\det\left(\mathcal{A}_{\{p_{1l},p_{2m}\}}\right) \\
&= (-1)^{j+k}p_{1j}p_{2k}\det\left(\mathcal{A}(p)\right).
\end{aligned}
$$

where the last equality follows from expanding the determinant of $\mathcal{A}(p)$ along the last two columns.

(b) Without loss of generality, let $i_1 = 1$, $i_2 = 2$, $i_3 = 3$ and let p_{3k} be arbitrary. For simplicity, suppose $j_1 = j_2 = 1$. Therefore, we consider the trifocal $\det(\mathcal{A}(p)_{\{p_{11},p_{21}\}})$. Using Lemma 2.8.3 and Cauchy-Binet, we expand $f = \det(P(p)_R\mathcal{A})$ where $R_1 = \{p_{12}, p_{13}\}$, $R_2 = \{p_{22}, p_{23}\}$, $R_3 = \{p_{3k}\}$ as follows:

$$
\begin{aligned}
f &= \det\left(P(p)_R\mathcal{A}\right) \\
&= \sum_{C\,:\,|C_1|=|C_2|=2,|C_3|=1} \det\left(P(p)^{C}_{\{R_1,R_2,R_3\}}\right)\det(\mathcal{A}_C) \\
&= \sum_{|C_3|=1}\det\left(([p_3]_\times)^{C_3}_{R_3}\right)\times\sum_{|C_1|=|C_2|=2}\left(\det\left(([p_1]_\times)^{C_1}_{R_1}\right)\times\det\left(([p_2]_\times)^{C_2}_{R_2}\right)\det\left(\mathcal{A}_C\right)\right) \\
&= \sum_{i=1}^{3}(-1)^{i+k}p_{3k}p_{3i}\times\sum_{\substack{|C_1|=|C_2|=2\\ C_1\neq R_1\\ C_2\neq R_2}}\left(\det\left(([p_1]_\times)^{C_1}_{R_1}\right)\times\det\left(([p_2]_\times)^{C_2}_{R_2}\right)\det\left(\mathcal{A}_{\{C_1,C_2,p_{3i}\}}\right)\right)
\end{aligned}
$$

$$= (-1)^k p_{3k} \sum_{i=1}^{3} (-1)^i p_{3i} \times$$

$$\sum_{2 \le l,m \le 3} \left(\det\left(([p_1]_\times)^{\{p_{11},p_{1l}\}}_{\{p_{12},p_{13}\}} \right) \times \det\left(([p_2]_\times)^{\{p_{21},p_{2m}\}}_{\{p_{22},p_{23}\}} \right) \times \det\left(\mathcal{A}_{\{p_{11},p_{1l},p_{21},p_{2m},p_{3i}\}} \right) \right)$$

$$= (-1)^k p_{3k} \sum_{i=1}^{3} (-1)^i p_{3i} \times \sum_{2 \le l,m \le 3} (-1)^{l+m} p_{1l} p_{2m} \det\left(\mathcal{A}_{\{p_{11},p_{1l},p_{21},p_{2m},p_{3i}\}} \right)$$

$$= (-1)^k p_{3k} \det(\mathcal{A}(p)_{\{p_{11},p_{21}\}})$$

Observe that the final equality follows from expanding the determinant of $\mathcal{A}(p)_{\{p_{11},p_{21}\}}$ on the p_3 column.

For general j_1, j_2, performing the same computation with $R_1 = \{p_{11}, p_{12}, p_{13}\} \setminus \{p_{1j_1}\}$, $R_2 = \{p_{21}, p_{22}, p_{23}\} \setminus \{p_{1j_2}\}$ and $R_3 = \{p_{3k}\}$ yields $\det(P(p)_R \mathcal{A}) = (-1)^k p_{3k} \det\left(\mathcal{A}(p)_{\{p_{1j_1}, p_{2j_2}\}}\right)$.

$\square$

Chapter 3

THE CHIRAL DOMAIN OF A CAMERA ARRANGEMENT

3.1 Introduction

In computer vision, chirality refers to the constraint that for a scene point to be visible in a camera, it must lie *in front* of it [24]. There is now a mature theory of multiview geometry that ignores this constraint [23] modeling image formation only via algebraic constraints coming from a (projective) camera being a (rational) linear map from $\mathbb{P}^3$ to $\mathbb{P}^2$. Chirality imposes additional semialgebraic conditions. The study of chirality was initiated by Hartley in his seminal paper [24], much of which forms [23, Chapter 21]. Further studies of chirality can be found in papers such as [31, 51, 52, 53, 54].

In this paper we develop a general theory of *multiview chirality* for an arrangement of projective cameras $\mathcal{A} = \{A_1, \ldots, A_m\}$ with distinct centers. Our central contribution is the notion of the *chiral domain* of $\mathcal{A}$ which is the subset of $\mathbb{P}^3$ (the world) that is visible in the cameras of $\mathcal{A}$. This is a multiview generalization of the classical definition of chirality, and covers all of $\mathbb{P}^3$ including infinite points and points on the principal planes of the cameras. The previous definition only covered finite world points. The extension is easy for one camera but subtle for multiple cameras as we explain. We show that the chiral domain admits a simple semialgebraic description by quadratic inequalities determined by the principal planes of the cameras and the plane at infinity. This description is the workhorse of the paper.

Recall that the *joint image* of $\mathcal{A}$ due to Triggs [49], [50], is the set of all "images" of $\mathbb{P}^3$ in the cameras of $\mathcal{A}$, ignoring chirality. Starting with the seminal work of Longuet-Higgins [34], there is now a complete algebraic and set theoretic characterization of the joint image [3, 18, 23, 27, 35, 37, 46]. The closure of the joint image is an algebraic variety in $(\mathbb{P}^2)^m$ called called the *joint image variety* in [46] and the *multiview variety* in [3], [5].

We define and describe the *chiral joint image* of $\mathcal{A}$, which is the analog of the joint image under the requirement of chirality. It is the image of the chiral domain of $\mathcal{A}$ in the cameras of $\mathcal{A}$, and is thus the true image of the world in $\mathcal{A}$. The chiral joint image is a subset of the joint image. One of our main results is a semialgebraic description (using polynomial equalities and inequalities) of the chiral joint image. This semialgebraic description is a refinement of the multiview constraints (epipolar & trifocal), when the world points are constrained to lie in front of the cameras.

Hartley studies the question of when a projective reconstruction of a set of point correspondences in $(\mathbb{P}^2)^m$ can be turned into a chiral reconstruction by applying a $\mathbb{P}^3$-homography. His answer is in terms of the feasibility of a system of linear inequalities called *chiral inequalities* [23, Chapter 21] which boils down to solving two linear programs. The chiral domain recovers this result using polyhedral geometry. We are also able to interpret quasi-affine transformations in this language. Analogous results for Euclidean cameras are also stated.

In [24], Hartley also characterizes the existence of a chiral reconstruction for two-views in terms of a sign condition on the given projective reconstruction. Werner et. al. also study the two-view case, considering both minimal and nonminimal configurations [52, 53]. Nistér & Schafflitzky consider the minimial problem in the Euclidean case [40]. Our polyhedral approach provides a simple proof of this two-view result and explains why such a result does not extend to more than two cameras.

Hartley develops chirality using classical projective geometry [23, 24]. Other authors have used *oriented projective geometry* [45] to model chirality (see for example, [31], [54]). This approach requires the choice of an orientation of the cameras involved based on a world point that is known to be in front of the cameras. The initial choice of orientation percolates down a chain of subsequent choices. In particular, a projective camera A is considered different from $-A$ in this theory.

Following Hartley, we use classical projective geometry in this paper. By staying in the original projective setup of projective cameras and projective reconstructions we, like Hartley, are able to avoid all of the choices needed in oriented projective geometry, and still obtain a perfectly valid theory of chirality. Working in the projective framework is especially handy when describing the chiral joint image in this paper which is naturally a subset of the joint image, a quasi-projective algebraic variety. The trick is to derive meaningful inequalities in projective space, by which we mean inequalities that are valid on both $\mathbf{x}$ and $-\mathbf{x}$ in Euclidean space.

This paper is organized as follows. In Section 3.2, we provide some background and set the notation. Section 3.3 introduces the chiral domain of a camera arrangement. This is then used to define and describe the chiral joint image of the camera arrangement in Section 3.4. In Section 3.5 we establish the connections to Hartley's results about when a projective reconstruction can be made chiral using a homography. We also show how our results connect to quasi-affine transformations and Hartley's two-view results on chirality. The case of Euclidean reconstructions is treated in Section 3.5.2. Section 3.6 summarizes our contributions. Many of the technical proofs in the three main sections can be found in Section 3.7.

3.2 Background and Notation

The sets of nonnegative integers, nonnegative real numbers, and positive real numbers are denoted by $\mathbb{N}, \mathbb{R}_+$, and $\mathbb{R}_{++}$, respectively. $\mathbb{P}^n$ denotes n-dimensional projective space over the reals, which is $\mathbb{R}^{n+1} \setminus \{0\}$ modulo the equivalence relation $\sim$ where $\mathbf{x} \sim \mathbf{y}$ if $\mathbf{x}$ is a scalar multiple of $\mathbf{y}$. If $\mathbf{x} \sim \mathbf{y}$, then we say that $\mathbf{x}$ and $\mathbf{y}$ are equal in $\mathbb{P}^n$, or $\mathbf{x}$ is *identified with* $\mathbf{y}$. We use $=$ to denote coordinate wise equality in $\mathbb{R}^n$.

In multiview geometry, we focus on $\mathbb{P}^3, \mathbb{P}^2$ and $\mathbb{R}^3, \mathbb{R}^2$, where $\mathbb{P}^n$ is a compactificaction of $\mathbb{R}^n$ with respect to the embedding $\mathbb{R}^n \to \mathbb{P}^n$, $\mathbf{x} \mapsto \widehat{\mathbf{x}} = (\mathbf{x}, 1)$. So points whose last coordinate is nonzero are said to be *finite*, whereas points whose last coordinate is 0 form the *hyperplane at infinity*. We write the plane at infinity as $L_\infty := \{\mathbf{q} \in \mathbb{P}^3 \,:\, \mathbf{n}_\infty^\top \mathbf{q} = 0\}$, where we fix the normal $\mathbf{n}_\infty = (0,0,0,1)^\top$.

We denote points in $\mathbb{P}^3$ and $\mathbb{R}^3$ by $\mathbf{q}$ allowing the context to decide where $\mathbf{q}$ lies. Similarly we denote points in $\mathbb{P}^2$ and $\mathbb{R}^2$ by $\mathbf{p}$. The dehomogenization of a finite point $\mathbf{q} \in \mathbb{P}^3$ is denoted

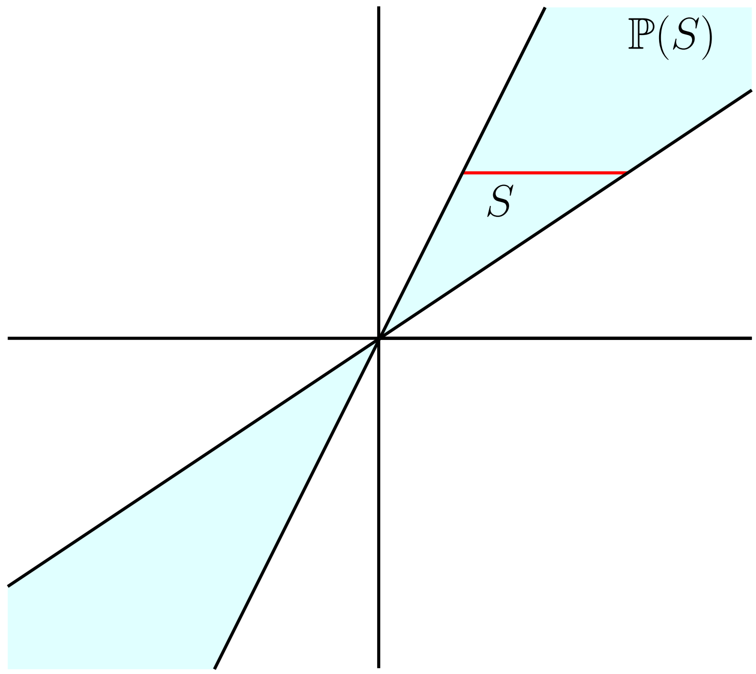

Figure 3.1: Projectivization of a set $S \subset \mathbb{R}^2$.

$\widetilde{\mathbf{q}} := (q_1/q_4, q_2/q_4, q_3/q_4)^\top$.

The *projectivization* of a set $S \subseteq \mathbb{R}^{n+1}$ is the union of all lines through the origin in $\mathbb{R}^{n+1}$ that intersect S. For an example see Figure 3.1.

A *projective* camera is a matrix $A = \begin{bmatrix} G & \mathbf{t} \end{bmatrix} \in \mathbb{R}^{3\times 4}$ of rank 3. The camera A is *finite* if $\det(G) \neq 0$. The *center* of the camera A is the unique point $\mathbf{c} \in \mathbb{P}^3$ such that $A\mathbf{c} = 0$. The camera A is finite if and only if its center $\mathbf{c} \in \mathbb{P}^3$ is finite. All cameras considered in this paper are finite. For consistency, we will choose the $\mathbb{R}^4$ representative $\mathbf{c}_A = \begin{bmatrix} -G^{-1}\mathbf{t} \\ 1 \end{bmatrix}$ for the center of A .

The *principal plane* of a finite camera $A = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}$ is the hyperplane $L_A := \{\mathbf{q} \in \mathbb{P}^3 : A_{3,\bullet}\mathbf{q} = 0\}$, where $A_{3,\bullet}$ is the third row of A, i.e. L_A is the set of points in $\mathbb{P}^3$ that image to infinite points in $\mathbb{P}^2$ under camera A. Note that the camera center $\mathbf{c}$ lies on L_A. We regard L_A as an oriented hyperplane in $\mathbb{R}^4$ with normal vector $\mathbf{n}_A := \det(G)A_{3\bullet}^\top$, which we call the *principal ray* of A. The $\det(G)$ factor makes sure that if we pass from A to λA for some nonzero scalar $\lambda \in \mathbb{R}$, the normal vector of the principal plane does not change sign.

The *world* $\mathbb{R}^3$, which is to be imaged by A, is modeled as the affine patch in $\mathbb{P}^3$ with $q_4 = 1$. This allows the identification of a finite point $\mathbf{q} \in \mathbb{P}^3$ with the world point $\widetilde{\mathbf{q}} \in \mathbb{R}^3$, and a world point $\mathbf{q} \in \mathbb{R}^3$ with the finite point $\widehat{\mathbf{q}} \in \mathbb{P}^3$. The image of $\mathbf{q} \in \mathbb{P}^3$, in the camera A is $A\mathbf{q} \in \mathbb{P}^2$. The rational map $A\colon \mathbb{P}^3 \dashrightarrow \mathbb{P}^2$, $\mathbf{q} \mapsto A\mathbf{q}$, is defined for all $\mathbf{q} \in \mathbb{P}^3$ except the center $\mathbf{c}$ of A[1].

Let $\mathcal{A} = \{A_1, \ldots, A_m\}$ denote an arrangement of m cameras with distinct centers. We use the shorthand $\mathbf{n}_i$ for the principal ray of camera A_i. Given a pair of cameras A_i, A_j with centers $\mathbf{c}_i$ and $\mathbf{c}_j$, let $\mathbf{e}_{ij}$ denote the image of $\mathbf{c}_j$ in A_i. The points $\mathbf{e}_{ij}$ are called *epipoles*. The line through $\mathbf{c}_i, \mathbf{c}_j$ is called the *baseline* of the pair of cameras $\{A_i, A_j\}$. All points on the base line (except for the centers themselves) will image in the two cameras at their respective epipoles $(\mathbf{e}_{ij}, \mathbf{e}_{ji})$.

We now recall some basics of convex geometry. Further details can be found in [8]. The (polyhedral) *cone* K_U spanned by a set of vectors $U = \{\mathbf{u}_1, \ldots, \mathbf{u}_\ell\} \subseteq \mathbb{R}^n$ is the set of all nonnegative linear combinations of the vectors in U, so

$$\begin{aligned} K_U &:= \operatorname{cone}(\mathbf{u}_1, \ldots, \mathbf{u}_\ell) \\ &= \{x_1\mathbf{u}_1 + \cdots + x_\ell \mathbf{u}_\ell \mid x_1, \ldots, x_\ell \geq 0\}. \end{aligned}$$

The *dimension* of a cone $K_U \subseteq \mathbb{R}^n$ is the dimension of the smallest vector space containing it. Let $\operatorname{int} K_U$ denote the relative interior of K_U, namely the interior of K_U in the linear span of U. A cone is *pointed* if it does not contain a line.

The *dual cone* K_U^* to K_U is the set of all vectors that make nonnegative inner product with every vector of K_U, so

$$K_U^* := \{\mathbf{y} \mid \forall\, \mathbf{x} \in K_U \colon \mathbf{y}^\top \mathbf{x} \geq 0\}. \tag{3.1}$$

[1]The broken arrow ($\dashrightarrow$) and the phrase "rational map" mean here that the domain of the map A is not actually $\mathbb{P}^3$ but rather $\mathbb{P}^3 \setminus \{\mathbf{c}\}$.

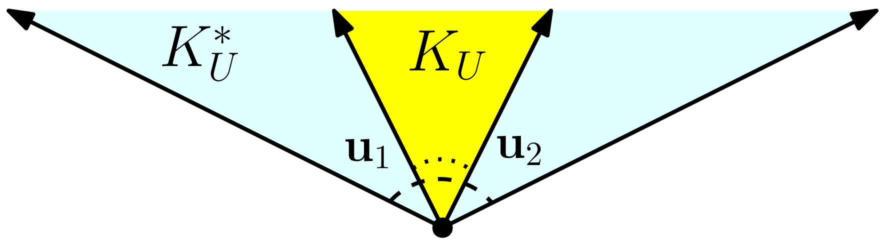

Figure 3.2: A cone K_U and its dual cone K_U^*.

Polyhedral cones are bidual in the sense that $K_U = (K_U^*)^*$. The cone K_U is pointed if and only if K_U^* is full-dimensional, and because of biduality, K_U is full-dimensional if and only if K_U^* is pointed. If K_U^* is full-dimensional, then $\operatorname{int} K_U^* = \{\mathbf{y} \mid \mathbf{y}^\top \mathbf{u}_i > 0,\ \forall\, i = 1, \ldots, \ell\}$. Otherwise, $K_U = L + K_W$ where L is a subspace and $K_W = \operatorname{cone}(\mathbf{w}_1, \ldots, \mathbf{w}_p)$ is a pointed cone. Writing $L^\perp$ as the kernel of a matrix A, $K_U^* = L^\perp \cap K_W^* = \{\mathbf{y} : A\mathbf{y} = 0, \mathbf{y}^\top \mathbf{w}_j > 0\ \forall\, j = 1, \ldots, p\}$. See Figure 3.3.

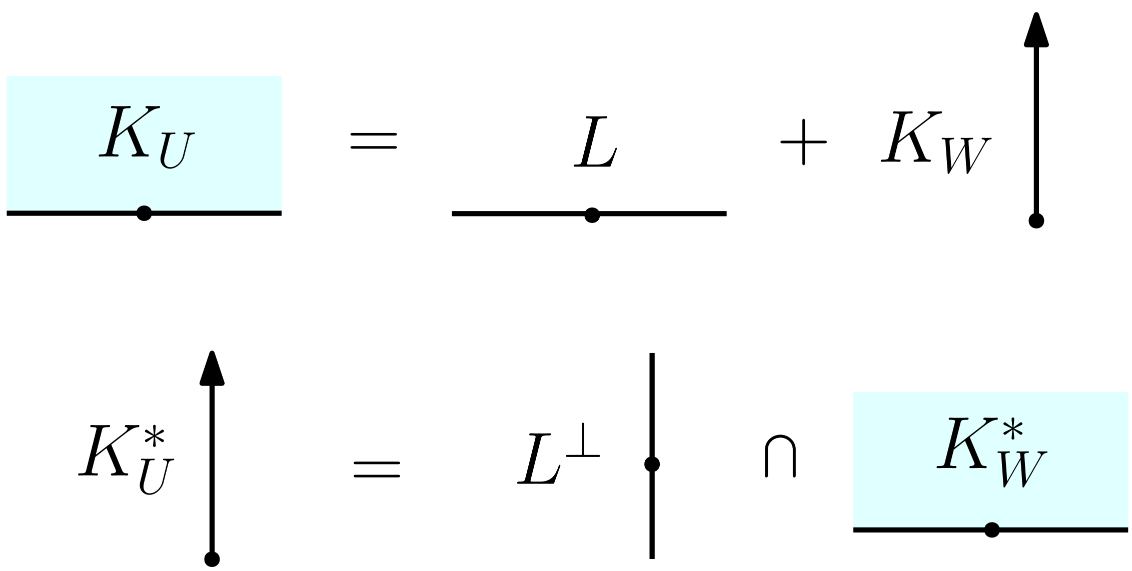

Figure 3.3: The non-pointed cone K_U on the top left is the sum of the subspace L and the pointed cone K_W. Its dual cone K_U^* in the second row is the intersection of $L^\perp$ and K_W^*.

Remark 3.2.1. *Membership in* $\operatorname{int} K_U^*$ *can be determined via linear programming. Indeed, the optimal value of the linear program*

$$\max\{\varepsilon \,:\, A\mathbf{y} = 0,\ \varepsilon \le 1,\ \mathbf{y}^\top \mathbf{w}_j \ge \varepsilon,\ j = 1, \dots, p\}$$

is positive if and only if $\operatorname{int} K_U^* \neq \emptyset$.

3.3 The Chiral Domain of an Arrangement of Cameras

We begin by recalling the definition of the depth of a finite point $\mathbf{q}$ in a finite camera A. It is essentially the projection of $\widetilde{\mathbf{q}} - \widetilde{\mathbf{c}}$ along the principal ray, see [23]. Formally, it is defined as

$$\operatorname{depth}(\mathbf{q}; A) := \left(\frac{1}{|\det(G)| \|G_{3,\bullet}\|}\right) \frac{(\mathbf{n}_A^\top \mathbf{q})}{(\mathbf{n}_\infty^\top \mathbf{q})}. \tag{3.2}$$

Notice that the sign of $\operatorname{depth}(\mathbf{q}; A)$ is unaffected by scaling $\mathbf{q}$. In fact, this definition of depth as a rational function of degree 0 in $\mathbf{q}$ (meaning that the degree of numerator and denominator are equal) underlines the inherently projective nature of the notion of depth. Furthermore, scaling A also does not affect the sign of $\operatorname{depth}(\mathbf{q}; A)$ because the orientation of $\mathbf{n}_A = \det(G) A_{3,\bullet}$ is independent of scaling.

We say that a finite point $\mathbf{q}$ not on the principal plane is *in front of* the camera if $\operatorname{depth}(\mathbf{q}; A) > 0$ [24]. Since only the sign of $\operatorname{depth}(\mathbf{q}; A)$ matters, we define the *chirality* of $\mathbf{q}$ in A to be $\chi(\mathbf{q}; A) = \operatorname{sgn}(\operatorname{depth}(\mathbf{q}; A))$. This is either 1 or -1. This definition of chirality excludes finite points with zero depth and points at infinity. Also, it treats chirality as a per camera concept.

In this section we will extend the above notion of chirality to all points in $\mathbb{P}^3$ with respect to one or more finite cameras. This will then lead to the central concept of this paper, the *chiral domain of an arrangement of cameras*, which we use to develop a unified theory of multiview chirality.

The exclusion of points on the principal plane and the plane at infinity makes the algebraic treatment of chirality complicated because it forces us to work with strict inequalities to avoid boundary points where depth is not defined. Our generalization below remedies this situation and in particular leads to an algebraic description of the *chiral joint image* of a camera arrangement as a subset of the classical *joint image* [3, 46].

We now discuss how we might decide the chirality of points on the principal plane and the plane at infinity. We first discuss the case of one camera: Traveling along the line through a point $\mathbf{a} \in \mathbb{R}^3$ in direction $\mathbf{q} \in \mathbb{R}^3$ corresponding to a point $\widehat{\mathbf{q}} \in \mathbb{P}^3$, we see from the definition that chirality changes only when we cross the principal plane or the plane at infinity (or it is always 0). So any point on either plane is arbitrarily close to finite points of chirality 1 in the camera A. More precisely, we should argue in $\mathbb{P}^3$ directly: Suppose $\mathbf{q} = (q_1, q_2, q_3, 0)^\top \in L_\infty$ and $\mathbf{a} = (a_1, a_2, a_3, 1)^\top$ has positive depth in A. We may assume that $\mathbf{n}_A^\top \mathbf{q} \geq 0$ since otherwise we can work with $-\mathbf{q}$. Then the points $\mathbf{q}_t := \frac{1}{t}\mathbf{a} + \mathbf{q}$ have positive depth for $t > 0$, and since $\mathbf{q} = \lim_{t \to \infty} \mathbf{q}_t$, it is arbitrarily close to points with chirality 1 in A. Now suppose $\mathbf{q} \in L_A \backslash L_\infty$. We may assume that $q_4 > 0$ since otherwise we can take $-\mathbf{q}$. As before, pick a finite point $\mathbf{a} = (a_1, a_2, a_3, 1)^\top$ with positive depth in A. Then the sequence of points $\mathbf{q}_t := \frac{1}{t}\mathbf{a} + \mathbf{q}$ have chirality 1 for all $t > 0$ and again, $\mathbf{q}$ is arbitrarily close to finite points of positive depth in A. Thus it seems natural to consider all points on $L_\infty \cup L_A$ to have chirality 1 in A. We could just as well argue that all points in $L_A \cup L_\infty$ should have chirality -1 with respect to the camera A by a similar argument to the above. However, our physical intuition is that L_∞ is visible in A, so points on L_∞ should have chirality 1 in A.

For arrangements of cameras, even just two cameras, the situation is more complicated. Consider Figure 3.4, where two cameras A_1, A_2 are placed on a train track looking in opposite directions so that $\mathbf{n}_{A_1} = -\mathbf{n}_{A_2}$. For instance, we could choose the camera facing right to be $A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}$ and the camera facing left to be

$$A_2 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 1 \\ 0 & 0 & -1 & 0 \end{bmatrix}.$$

Then $\mathbf{c}_1 = (0, 0, 0, 1)^\top$ and $\mathbf{c}_2 = (0, 1, 0, 1)^\top$ which differ in the vertical direction and the principal rays drawn in $\mathbb{R}^3$ are $(0,0,1)$ and $(0,0,-1)$. For any finite $\mathbf{q} \notin L_{A_1} = L_{A_2}$, $\chi(\mathbf{q}; A_1) = -\chi(\mathbf{q}; A_2)$, and hence there is no sequence of finite points that each have positive depth in the two cameras that can approach points on L_∞ or $L_{A_1} = L_{A_2}$. This is in contrast to the single camera case where all points can be approached by finite points with positive depth in the camera.

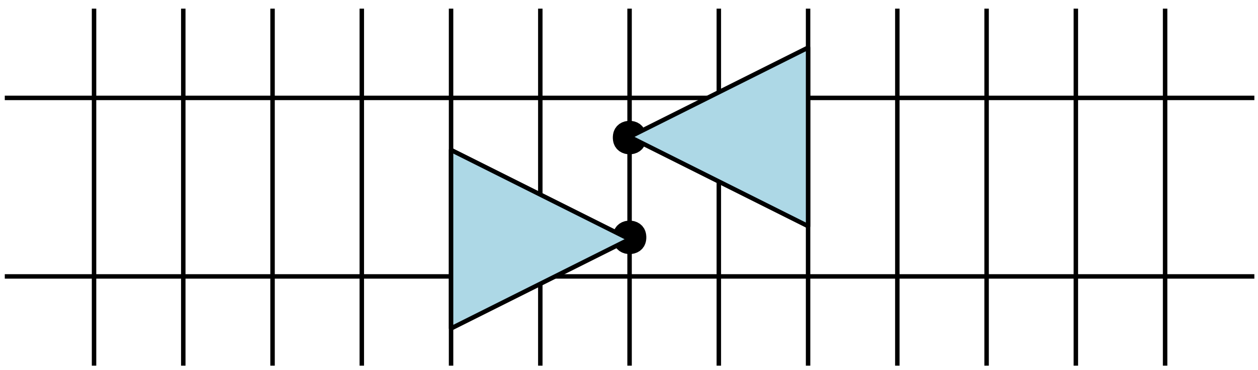

Figure 3.4: Cameras facing in opposite directions on train tracks.

This brings us to the following definition, that declares a point to be in front of a collection of cameras if and only if the point can be approached by a sequence of finite points that have chirality 1 in each camera.

Definition 3.3.1 (Chiral Domain of $\mathcal{A}$ and Chirality)**.** *Let $\mathcal{A}$ be an arrangement of finite cameras. Then the* chiral domain *of $\mathcal{A}$, denoted as $D_{\mathcal{A}}$, is the closure of the set*

$$\{\mathbf{q} \in \mathbb{P}^3 \mid \mathbf{q} \textit{ finite}, \ \forall A \in \mathcal{A}, \ \operatorname{depth}(\mathbf{q}; A) > 0\}.$$

Moreover, a point $\mathbf{q} \in \mathbb{P}^3$ is said to have chirality 1 with respect to $\mathcal{A}$*, denoted as $\chi(\mathbf{q}; \mathcal{A}) = 1$, if and only if $\mathbf{q} \in D_{\mathcal{A}}$.*

The limit in the above definition is defined using the natural topology in $\mathbb{P}^3$ induced by the quotient map $\pi : \mathbb{R}^4 \smallsetminus \{0\} \to \mathbb{P}^3$ in which $\pi(\mathbf{v}) = \pi(\mathbf{w})$ if and only if $\mathbf{v} \sim \mathbf{w}$. In this quotient topology, a set $U \subseteq \mathbb{P}^3$ is open if and only if its preimage $\pi^{-1}(U)$ is open in $\mathbb{R}^4 \smallsetminus \{0\}$ in the Euclidean topology. Thus $\mathbf{q} \in \mathbb{P}^3$ is a limit point of a sequence $\{\mathbf{q}_i\} \subseteq \mathbb{P}^3$ if and only if all open sets $\pi^{-1}(U)$ containing the line $\pi^{-1}(\mathbf{q})$ contains the line $\pi^{-1}(\mathbf{q}_i)$ for some i. The closure of a set $S \subseteq \mathbb{P}^3$ is the set of all limit points of sequences in S.

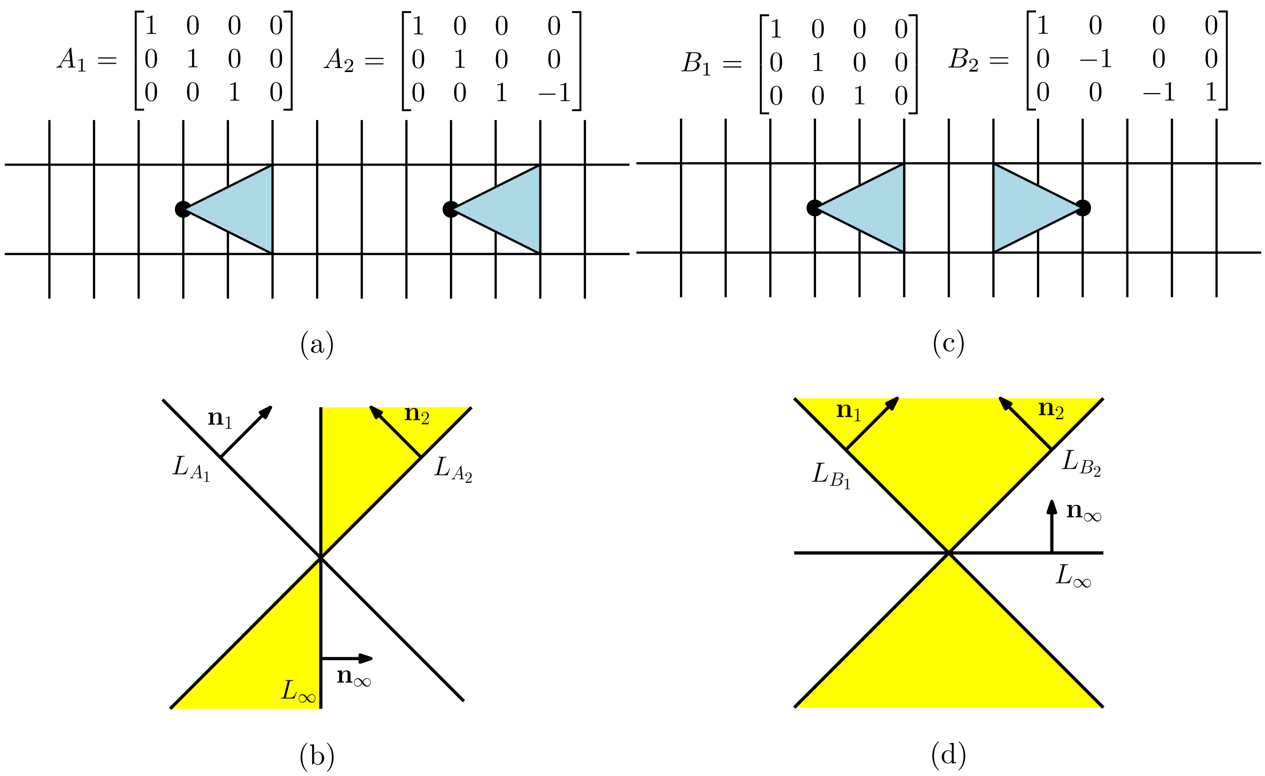

Figure 3.5: Two cameras on parallel tracks facing the same direction in Figure 3.5a and opposing directions in Figure 3.5c. The shaded area in Figure 3.5b and Figure 3.5d represents the chiral domain of the arrangements $\mathcal{A}$ and $\mathcal{B}$ around their intersection points with the hyperplane at infinity (that is after projective transformation of the corresponding Figure 3.5a and Figure 3.5c). Every infinite point is in the chiral domain of $\mathcal{A}$. The only infinite points in the chiral domain of $\mathcal{B}$ are those in the intersection $L_\infty \cap L_{B_1} \cap L_{B_2}$.

We illustrate Definition 3.3.1 using the camera arrangements in Figure 3.5. In Figure 3.5a $\mathbf{n}_1 = (0,0,1,0)$ and $\mathbf{n}_2 = (0,0,1,-1)$ and hence the principal rays point in the same direction $(0,0,1)$ in $\mathbb{R}^3$. The principal planes are parallel but distinct in $\mathbb{R}^3$, hence L_{A_1} and L_{A_2} intersect L_∞ in a two-dimensional subspace of $\mathbb{R}^4$ (line in $\mathbb{P}^3$). Figure 3.5b shows this intersection after projecting from a common point of L_{A_1}, L_{A_2} and L_∞ in order to get a 2-dimensional picture, where the projections of these three planes intersect in a point. The vanishing point of the train tracks has chirality 1 with respect to both cameras as there is a sequence of finite points with chirality 1 in both cameras, that converges to it. The chiral domain of the pair of cameras is shaded yellow and all of L_∞ is visible in both cameras.

The cameras B_1, B_2 in Figure 3.5c have principal rays $\mathbf{n}_{B_1} = (0,0,1,0)$ and $\mathbf{n}_{B_2} = (0,0,-1,1)$. In $\mathbb{R}^3$, they are pointed in opposite directions. The principal planes are again parallel but distinct in $\mathbb{R}^3$, hence L_{B_1} and L_{B_2} intersect L_∞ in a projective line. The only part of L_∞ in the chiral domain (shaded in yellow) is the intersection $L_\infty \cap L_{B_1} \cap L_{B_2}$ as seen in Figure 3.5d. These examples show that when there are multiple cameras, the chirality of points on L_{A_i} and L_∞ need more care.

Observe that $D_\mathcal{A}$ is nonempty if and only if it has nonempty interior. Indeed, if $D_\mathcal{A}$ is nonempty, there is a finite point $\mathbf{q} \in \mathbb{P}^3$ that has positive depth in all cameras of $\mathcal{A}$. Since depth depends continuously on the finite point, there is a neighborhood $U \subseteq \mathbb{P}^3$ of finite points with positive depth in all cameras.

We now give an algebraic description of $D_\mathcal{A}$ in terms of quadratic inequalities. This explicit description will play a central role in the rest of the paper.

Theorem 3.3.1. *Let $\mathcal{A} = \{A_1, \ldots, A_m\}$ be an arrangement of finite cameras. If the chiral domain, $D_\mathcal{A}$, is nonempty, then it has the semi-algebraic description*

$$D_\mathcal{A} = \left\{ \mathbf{q} \in \mathbb{P}^3 \mid \forall i,j, \begin{array}{c} (\mathbf{n}_\infty^\top \mathbf{q})(\mathbf{n}_i^\top \mathbf{q}) \geq 0, \\ (\mathbf{n}_i^\top \mathbf{q})(\mathbf{n}_j^\top \mathbf{q}) \geq 0 \end{array} \right\}, \tag{3.3}$$

where $\mathbf{n}_i$ is the principal ray of A_i.

Note that the inequalities $(\mathbf{n}_\infty^\top \mathbf{q})(\mathbf{n}_i^\top \mathbf{q}) \geq 0$ (and also $(\mathbf{n}_i^\top \mathbf{q})(\mathbf{n}_j^\top \mathbf{q}) \geq 0$) are well-defined in $\mathbb{P}^3$ because in each case, the polynomial f on the left hand side has even degree (namely 2) in $\mathbf{q}$, that is $f(\lambda\mathbf{q}) = \lambda^2 f(\mathbf{q})$ for every real scalar λ. Therefore, the sign of $f(\mathbf{q})$ is constant along the line spanned by $\mathbf{q}$ in $\mathbb{R}^4$.

of Theorem 3.3.1. If $D_{\mathcal{A}}$ is nonempty, then it has a nonempty interior. Let S be the interior of $D_{\mathcal{A}}$, i.e., the set of finite points in $\mathbb{P}^3$ that have positive depth in all cameras in $\mathcal{A}$. Such points correspond to lines through the origin in $\mathbb{R}^4$, consisting of points $\mathbf{q}$ that have ($\mathbf{n}_i^\top\mathbf{q} > 0$ and $\mathbf{n}_\infty^\top\mathbf{q} > 0$) or ($\mathbf{n}_i^\top\mathbf{q} < 0$ and $\mathbf{n}_\infty^\top\mathbf{q} < 0$). Let $Q \subseteq \mathbb{R}^4$ be the polyhedral cone defined by the inequalities $\mathbf{n}_i^\top\mathbf{q} \geq 0$ and $\mathbf{n}_\infty^\top\mathbf{q} \geq 0$. Then we see that $S = \mathbb{P}(\operatorname{int} Q \cup -\operatorname{int} Q)$ and hence the projectivization of $\operatorname{int} Q$. This implies that $D_{\mathcal{A}}$ is the projectivization of Q which can be defined by the quadratic inequalities $(\mathbf{n}_i^\top\mathbf{q})(\mathbf{n}_j^\top\mathbf{q}) \geq 0$, where $\{i, j\}$ ranges over all 2-element subsets of $\{1, 2, \ldots, m, \infty\}$. □

Remark 3.3.1. *Note from the proof of Theorem 3.3.1 that $D_{\mathcal{A}}$ is the projectivization of the polyhedral cone $Q \subseteq \mathbb{R}^4$ defined by the linear inequalities $q_4 \geq 0$ and $\mathbf{n}_i^\top\mathbf{q} \geq 0$ ($i = 1, 2, \ldots, m$). In other words, the lines in $\mathbb{R}^4$ corresponding to the points in $D_{\mathcal{A}}$ are exactly the lines through the origin in $Q \cup -Q$. Therefore, $D_{\mathcal{A}}$ is inherently a polyhedral set even though Theorem 3.3.1 describes $D_{\mathcal{A}}$ using quadratic inequalities.*

Remark 3.3.2. *The inequality description of $D_{\mathcal{A}}$ in* (3.3) *is only valid when $D_{\mathcal{A}}$ is nonempty. Indeed, the set on the right hand side can be nonempty even if $D_{\mathcal{A}}$ is empty. For example, for the cameras in Figure 3.4, $D_{\mathcal{A}} = \emptyset$ as there is no finite point with positive depth in both cameras. However, all points that lie on the line that is the intersection of L_∞ with the (common) principal plane of the cameras satisfies the (non-strict) inequalities on the right hand side of* (3.3) *even though there is no point in $\mathbb{P}^3$ where the inequalities are satisfied strictly. In general, the nonstrict inequalities admit all points that are on the principal planes of some cameras in $\mathcal{A}$ and have nonnegative depth in the others.*

Remark 3.3.3. *Specializing Theorem 3.3.1 to one camera, we get $D_{\{A\}} = \{\mathbf{q} \in \mathbb{P}^3 \colon (\mathbf{n}_A^\top\mathbf{q})(\mathbf{n}_\infty^\top\mathbf{q}) \geq 0\}$, which implies that $\chi(\mathbf{q}; \{A\}) = 1$ if $\mathbf{q} \in L_\infty \cup L_A$ for one camera A. This specialization matches our expectation for one camera that we explained earlier.*

We now give a criterion for the non-emptiness of $D_{\mathcal{A}}$ in terms of linear programming.

Theorem 3.3.2. *Let $\mathcal{A} = \{A_1, A_2, \ldots, A_m\}$ be an arrangement of finite cameras. Then $D_{\mathcal{A}} \neq \emptyset$ if and only if the row space of the $4 \times (m+1)$ matrix N with columns $\mathbf{n}_1, \ldots, \mathbf{n}_m, \mathbf{n}_\infty$ intersects the positive orthant $\mathbb{R}^{m+1}_{++}$.*

In particular, for all arrangements of $m \leq 3$ cameras such that $\mathbf{n}_\infty$ and the principal rays $\mathbf{n}_1, \ldots, \mathbf{n}_m$ are linearly independent, $D_{\mathcal{A}} \neq \varnothing$.

Proof. The set $D_{\mathcal{A}} \neq \varnothing$ if and only if there is a finite point with positive depth in all cameras. Equivalently, if and only if there is a $\mathbf{q} \in \mathbb{R}^4$ such that $q_4 \neq 0$ where $\mathbf{q}^\top N = [\mathbf{q}^\top \mathbf{n}_1, \ldots, \mathbf{q}^\top \mathbf{n}_m, q_4]$ lies in the positive or negative orthant. Thus $D_{\mathcal{A}} \neq \emptyset$ if and only if the row space of N has an intersection with $\mathbb{R}^{m+1}_{++}$.

If $m \leq 3$ and the columns of N are linearly independent then N has row rank $m+1$, and the rows of N span $\mathbb{R}^{m+1}$. So the rowspace of N intersects $\mathbb{R}^{m+1}_{++}$. $\square$

The following example shows that our result is tight in the sense that when $m > 3$, the chiral domain may be empty.

Example 3.3.1. *Consider an arrangement $\mathcal{A}$ with principal rays $\mathbf{n}_1 = (-1, 0, 0, 0)^\top, \mathbf{n}_2 = (1, -1, 0, 0)^\top, \mathbf{n}_3 = (0, 1, -1, 0)^\top$, and $\mathbf{n}_4 = (0, 0, 1, -1)^\top$. The matrix*

$$N = \begin{bmatrix} -1 & 1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & -1 & 1 \end{bmatrix}$$

has full rank. However, the row space of N has empty intersection with $\mathbb{R}^5_{++}$, so $D_{\mathcal{A}}$ is empty.

Theorem 3.3.2 provides an efficient method for checking if $D_{\mathcal{A}}$ is nonempty by checking the feasibility of a linear program (see Remark 3.2.1), whose size scales linearly with the number of cameras.

3.4 The Chiral Joint Image

Recall that world points are imaged in an arrangement of finite cameras $\mathcal{A} = \{A_1, \ldots, A_m\}$ via the rational map[2]

$$\varphi_{\mathcal{A}} : \begin{cases} \mathbb{P}^3 \dashrightarrow (\mathbb{P}^2)^m \\ \mathbf{q} \mapsto (A_1\mathbf{q}, A_2\mathbf{q}, \ldots, A_m\mathbf{q}) \end{cases} \tag{3.4}$$

Triggs calls $\varphi_{\mathcal{A}}(\mathbb{P}^3) =: \mathcal{J}_{\mathcal{A}}$ the joint image [49, 50] and Heyden-Åström call it the *natural descriptor* [27]. In this section we will describe the chiral analog of the joint image, i.e. the set of images of points that lie in front of an arrangement of cameras.

Definition 3.4.1 (Chiral Joint Image)**.** *The* chiral joint image *of a camera arrangement $\mathcal{A}$ is $\mathcal{X}_{\mathcal{A}} := \varphi_{\mathcal{A}}(D_{\mathcal{A}})$, the image of the chiral domain of $\mathcal{A}$ under $\varphi_{\mathcal{A}}$.*

Our goal will be to get an algebraic description of the chiral joint image $\mathcal{X}_{\mathcal{A}}$ and its Euclidean closure. Since $\mathcal{X}_{\mathcal{A}}$ lies in the joint image $\mathcal{J}_{\mathcal{A}}$, we begin by looking at $\mathcal{J}_{\mathcal{A}}$ and its closures.

A *variety* in $\mathbb{P}^n$ is the set of solutions to a finite set of homogeneous polynomial equations. The varieties in $\mathbb{P}^n$ are the closed sets of the Zariski topology on $\mathbb{P}^n$. The Zariski closure of a set $S \subset \mathbb{P}^n$, denoted by $\overline{S}^{Zar}$ is the smallest variety containing S.

Trager et al. [46] refer to $\overline{\mathcal{J}}_{\mathcal{A}}^{Zar}$ in $(\mathbb{P}^2)^m$ as the *joint image variety of* $\mathcal{A}$. Recall that the *epipolar* and *trifocal* constraints cut out the joint image variety [3]. These are known as the *multiview constraints* on the image points. Trager et al. also characterize the points added to $\mathcal{J}_{\mathcal{A}}$ by the closure operation, that is the difference between $\mathcal{J}_{\mathcal{A}}$ and $\overline{\mathcal{J}}_{\mathcal{A}}^{Zar}$. They do this using the following sets.

Definition 3.4.2. *Given an arrangement of cameras $\mathcal{A} = \{A_1, \ldots, A_m\}$, let $E_j = \mathbf{e}_{1j} \times \ldots \mathbb{P}^2_j \ldots \times \mathbf{e}_{mj}$, where $\mathbb{P}^2_j$ represents a copy of $\mathbb{P}^2$ in the jth slot. Set $E_{\mathcal{A}} = \bigcup_{j=1}^m E_j$, where $A_i\mathbf{c}_j = \mathbf{e}_{ij}$ are the epipoles of $\mathcal{A}$.*

[2]Again, the broken arrow ($\dashrightarrow$) and the words "rational map" refer to the fact that the domain of the map $\varphi_{\mathcal{A}}$ is not $\mathbb{P}^3$ but rather $\mathbb{P}^3 \setminus \{\mathbf{c}_1, \ldots, \mathbf{c}_m\}$.

Theorem 3.4.1 ([46, Proposition 1])**.** *Given an arrangement of cameras* $\mathcal{A} = \{A_1, \ldots, A_m\}$*, with distinct camera centers,*

$$\overline{\mathcal{J}}_{\mathcal{A}}^{Zar} = \mathcal{J}_{\mathcal{A}} \cup E_{\mathcal{A}}.$$

While the Zariski topology is natural for algebraic sets, it is too coarse for semialgebraic sets. For example, consider the set $C = \{(x, y) \in \mathbb{R}^2 | x^2 + y^2 = 1, x > 0, y > 0\}$, i.e., the unit circle restricted to the positive quadrant. When talking about its closure, one would want to talk about the set $\overline{C} = \{(x, y) \in \mathbb{R}^2 \,|\, x^2 + y^2 = 1, x \geq 0, y \geq 0\}$. However, the Zariski closure of C, $\overline{C}^{Zar} = \{(x, y) \in \mathbb{R}^2 | x^2 + y^2 = 1\}$ is the entire unit circle. Luckily for us, as the following theorem shows, the Euclidean closure $\overline{\mathcal{J}}_{\mathcal{A}}$ and the Zariski closure $\overline{\mathcal{J}}_{\mathcal{A}}^{Zar}$ of the joint image are the same [3].

Theorem 3.4.2. $\overline{\mathcal{J}}_{\mathcal{A}} = \overline{\mathcal{J}}_{\mathcal{A}}^{Zar} = \mathcal{J}_{\mathcal{A}} \cup E_{\mathcal{A}}$.

Proof. Recall that $\mathbf{p} = (\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m) \in E_j$ is of the form $(\mathbf{e}_{1j}, \ldots \mathbf{p}_j, \ldots, \mathbf{e}_{mj})$ for some $\mathbf{p}_j \in \mathbb{P}^2$. So all coordinates of $\mathbf{p}$ except $\mathbf{p}_j$ are the images of $\mathbf{c}_j = \begin{bmatrix} -G_j^{-1}\mathbf{t}_j \\ 1 \end{bmatrix}$. Consider now the curve

$$\mathbf{v}(s) = \begin{pmatrix} sG_j^{-1}\mathbf{p}_j - G_j^{-1}\mathbf{t}_j \\ 1 \end{pmatrix} = \begin{pmatrix} sG_j^{-1}\mathbf{p}_j \\ 0 \end{pmatrix} + \mathbf{c}_j$$

as s varies over $\mathbb{R}$. Then $\lim_{s \to 0} \varphi_{\mathcal{A}}(\mathbf{v}(s)) = \mathbf{p}$, since for $i \neq j$, $A_i\mathbf{v}(s) = sG_iG_j^{-1}\mathbf{p}_j + \mathbf{e}_{ij}$ and $A_j\mathbf{v}(s) = s\mathbf{p}_j \sim \mathbf{p}_j$. So $\mathbf{p} \in \overline{\mathcal{J}}_{\mathcal{A}}$, and hence $E_j \subseteq \overline{\mathcal{J}}_{\mathcal{A}}$. Therefore, by Theorem 3.4.1, $\overline{\mathcal{J}}_{\mathcal{A}}^{Zar} \subseteq \overline{\mathcal{J}}_{\mathcal{A}}$. This means that

$$\mathcal{J}_{\mathcal{A}} \subseteq \overline{\mathcal{J}}_{\mathcal{A}}^{Zar} \subseteq \overline{\mathcal{J}}_{\mathcal{A}} \tag{3.5}$$

and taking Euclidean closure throughout and noting that Zariski closed sets are also closed in the Euclidean topology, we get the first equality. The second equality is Theorem 3.4.1. □

[3] Recall that the topology we use on $\mathbb{P}^n$ is induced by the Euclidean topology on $\mathbb{R}^{n+1} \setminus \{0\}$. This induces a topology on the product of real projective spaces $(\mathbb{P}^2)^m$. Explicitly, a set $U_1 \times U_2 \times \ldots \times U_m \subseteq (\mathbb{P}^2)^m$ is open if and only if the sets $U_i \subseteq \mathbb{P}^2$ are all open sets.

Remark 3.4.1. *In general, over the complex numbers, the Euclidean closure and the Zariski closure of a constructible set are equal [38, Theorem I.10.1]. Theorem 3.4.2, however, is about real numbers, and the real Euclidean closure of a set is not always equal to the real part of the complex Euclidean closure. This is because, the Euclidean closure of the complex points recovers the real points and can sometimes produce isolated real points (real points that can be separated from the real points in the constructible set). For example, consider the map* $(a : b) \mapsto (a^3 + ab^2 : a^2b + b^3 : a^3)$ *from* $\mathbb{P}^1$ *to* $\mathbb{P}^2$*. The image is a rational curve in* $\mathbb{P}^2$ *defined by the equation* $y^2z = x^2(x - z)$ *and has an isolated real singularity at* $(0 : 0 : 1)$*, which is not in the closure of the image of* $\mathbb{P}^1$*. Our proof of Theorem 3.4.2 rests on the multi-linearity of the image formation map.*

We will now focus on giving a semialgebraic description of the chiral joint image $\mathcal{X}_\mathcal{A}$ and its Euclidean closure. The essential inequalities enforcing chirality in the space of images are given in the following definition.

Definition 3.4.3. *Given an arrangement of finite cameras* $A_i = \begin{bmatrix} G_i & \mathbf{t}_i \end{bmatrix}$*, define* $C_\mathcal{A}$ *to be the set*

$$\left\{ \mathbf{p} : \begin{array}{l} \det(G_i) p_{i3} (\mathbf{a}_i \times \mathbf{a}_j)^\top (\mathbf{b}_{ij} \times \mathbf{a}_j) \geq 0, \\ \det(G_i)\det(G_j) p_{i3} p_{j3} (\mathbf{b}_{ij} \times \mathbf{a}_i)^\top (\mathbf{b}_{ij} \times \mathbf{a}_j) \geq 0 \end{array} \right\}$$

where $\mathbf{p} = (\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m) \in (\mathbb{P}^2)^m$*,* $\mathbf{b}_{ij} = G_i^{-1}\mathbf{t}_i - G_j^{-1}\mathbf{t}_j$ *is a direction of the baseline connecting the centers of cameras* A_i *and* A_j*, and* $\mathbf{a}_i = G_i^{-1}\mathbf{p}_i$*.*

The inequalities describing $C_\mathcal{A}$ come from those describing the chiral domain $D_\mathcal{A}$. See proof of Lemma 3.4.1 in Section 3.7.

Note that each inequality in $C_\mathcal{A}$ involves only two cameras in the arrangement. They are well-defined on $(\mathbb{P}^2)^m$ because every inequality has even degree in the coordinates on the $\mathbb{P}^2$-factors. In fact, the inequalities are all biquadratic, i.e. of degree $(2, 2)$. Moreover, the sign does not depend on the choice of the order of the cameras in the arrangement because this choice is implicit in $\mathbf{b}_{ij}$ and explicit in the terms $(\mathbf{a}_i \times \mathbf{a}_j)$ in the inequalities. So a relabeling of the cameras will not change the signs involved.

The following lemma (whose proof can be found in Section 3.7) gives the set theoretic relationship between the joint image, the chiral joint image and the set $C_{\mathcal{A}}$.

Lemma 3.4.1. *Let $\mathcal{A} = \{A_1, \ldots, A_m\}$ be an arrangement of finite cameras such that $D_{\mathcal{A}}$ is nonempty. If the centers of $\mathcal{A}$ are not collinear, then*

$$\mathcal{X}_{\mathcal{A}} = \mathcal{J}_{\mathcal{A}} \cap C_{\mathcal{A}}. \tag{3.6}$$

If the centers are collinear, then set $\mathbf{e} := (\mathbf{e}_1, \ldots, \mathbf{e}_m)$ to be the image of the common baseline under $\varphi_{\mathcal{A}}$. Then

$$\mathcal{X}_{\mathcal{A}} \setminus \{\mathbf{e}\} = (\mathcal{J}_{\mathcal{A}} \cap C_{\mathcal{A}}) \setminus \{\mathbf{e}\} \tag{3.7}$$

In both cases, $\overline{\mathcal{X}}_{\mathcal{A}} \subseteq C_{\mathcal{A}}$.

Lemma 3.4.1, while interesting, is not useful in practice, since it does not give us a way of algebraically representing $\mathcal{X}_{\mathcal{A}}$. This is because it is stated using $\mathcal{J}_{\mathcal{A}}$. A more useful description involves $\overline{\mathcal{J}}_{\mathcal{A}}$ which is algebraic. Recall that going from $\mathcal{J}_{\mathcal{A}}$ to $\overline{\mathcal{J}}_{\mathcal{A}}$ brings in the set $E_{\mathcal{A}}$. For the chiral joint image the relevant part of $E_{\mathcal{A}}$ is

$$E_{\mathcal{A}}^{+} := \bigcup_{j | \mathbf{c}_j \in D_{\mathcal{A}}} E_j$$

The set $E_{\mathcal{A}}^{+}$ can be divided into two parts as follows:

- $E_{\mathcal{A}}^{++}$, the union of the sets E_j such that $\mathbf{c}_j$ has positive depth in every camera A_i with $i \neq j$, and

- $E_{\mathcal{A}}^{0}$, the union of all E_j such that $\mathbf{c}_j \in D_{\mathcal{A}}$ and the depth of $\mathbf{c}_j$ is zero in some camera A_i with $i \neq j$.

Of these, the second set $E_{\mathcal{A}}^{0}$ causes the most technical issues as it may intersect (but not be contained in) $\overline{\mathcal{X}}_{\mathcal{A}}$, whereas it is always contained in $C_{\mathcal{A}}$. Armed with these definitions we state the main theorem of this section.

Theorem 3.4.3. *Let $\mathcal{A}$ be an arrangement of finite cameras with distinct centers. Further, assume that the chiral domain $D_{\mathcal{A}}$ is nonempty. Let $E^+_{\mathcal{A}}$ be the union of all E_j such that $\mathbf{c}_j$ lies in $D_{\mathcal{A}}$. If the centers are not collinear then*

$$\overline{\mathcal{J}}_{\mathcal{A}} \cap C_{\mathcal{A}} = \mathcal{X}_{\mathcal{A}} \cup E^+_{\mathcal{A}} = \overline{\mathcal{X}}_{\mathcal{A}} \cup E^0_{\mathcal{A}}. \tag{3.8}$$

Otherwise,

$$\overline{\mathcal{J}}_{\mathcal{A}} \cap C_{\mathcal{A}} = \mathcal{X}_{\mathcal{A}} \cup E^+_{\mathcal{A}} \cup \{\mathbf{e}\} = \overline{\mathcal{X}}_{\mathcal{A}} \cup E^0_{\mathcal{A}} \cup \{\mathbf{e}\}. \tag{3.9}$$

where $\mathbf{e} = (\mathbf{e}_1, \ldots, \mathbf{e}_m)$ is the image of the common baseline under $\varphi_{\mathcal{A}}$.

By Theorem 3.4.3, the epipolar and trifocal constraints together with the inequalities defining $C_{\mathcal{A}}$ are the *chiral multiview constraints*. Specializations of Theorem 3.4.3 to Euclidean cameras are straightforward.

The proof of Theorem 3.4.3 relies on the following additional lemmas, the proofs of which can be found in Section 3.7.

Lemma 3.4.2. *Let $\mathcal{A} = \{A_1, \ldots, A_m\}$ be an arrangement of finite cameras. If the centers $\mathbf{c}_i$ are not collinear then $E_{\mathcal{A}} \cap C_{\mathcal{A}} = E^+_{\mathcal{A}}$. Otherwise, $E_{\mathcal{A}} \cap C_{\mathcal{A}} = E^+_{\mathcal{A}} \cup \{\mathbf{e}\}$.*

Lemma 3.4.3. *Let $\mathcal{A} = \{A_1, \ldots, A_m\}$ be an arrangement of finite cameras. Then $E^{++}_{\mathcal{A}} \subseteq \overline{\mathcal{X}}_{\mathcal{A}}$.*

of Theorem 3.4.3. We will first prove Equation (3.8) which is the case of noncollinear centers. By Theorem 3.4.2, $\overline{\mathcal{J}_{\mathcal{A}}} = \mathcal{J}_{\mathcal{A}} \cup E_{\mathcal{A}}$. Therefore,

$$\overline{\mathcal{J}}_{\mathcal{A}} \cap C_{\mathcal{A}} = (\mathcal{J}_{\mathcal{A}} \cap C_{\mathcal{A}}) \cup (E_{\mathcal{A}} \cap C_{\mathcal{A}}) = \mathcal{X}_{\mathcal{A}} \cup E^+_{\mathcal{A}}$$

where the last equality follows from Lemmas 3.4.1 and 3.4.2. This proves the first equality in Equation (3.8).

By Lemma 3.4.3, $\overline{\mathcal{J}}_{\mathcal{A}} \cap C_{\mathcal{A}} \subseteq \overline{\mathcal{X}}_{\mathcal{A}} \cup E^0_{\mathcal{A}}$ since

$$\overline{\mathcal{J}}_{\mathcal{A}} \cap C_{\mathcal{A}} = \mathcal{X}_{\mathcal{A}} \cup E^+_{\mathcal{A}} = \mathcal{X}_{\mathcal{A}} \cup E^{++}_{\mathcal{A}} \cup E^0_{\mathcal{A}} \subseteq \overline{\mathcal{X}}_{\mathcal{A}} \cup E^0_{\mathcal{A}}.$$

By Lemma 3.4.1, we get $\overline{\mathcal{X}}_{\mathcal{A}} \subseteq C_{\mathcal{A}}$. Therefore, since $\mathcal{X}_{\mathcal{A}} \subseteq \mathcal{J}_{\mathcal{A}}$, it follows that $\overline{\mathcal{X}}_{\mathcal{A}} \subseteq \overline{\mathcal{J}}_{\mathcal{A}} \cap C_{\mathcal{A}}$. Lemma 3.4.2 shows that $E^+_{\mathcal{A}}$ (so in particular $E^0_{\mathcal{A}}$) is contained in $C_{\mathcal{A}}$, but since all of $E_{\mathcal{A}}$ is inside $\overline{\mathcal{J}}_{\mathcal{A}}$, we also have $E^0_{\mathcal{A}} \subseteq \overline{\mathcal{J}}_{\mathcal{A}} \cap C_{\mathcal{A}}$. Therefore, $\overline{\mathcal{X}}_{\mathcal{A}} \cup E^0_{\mathcal{A}} \subseteq \overline{\mathcal{J}}_{\mathcal{A}} \cap C_{\mathcal{A}}$. Putting both these containments together we get the second equality in Equation (3.8).

The collinear case follows the same proof mechanics as above, utilizing the parts of Lemmas 3.4.1, 3.4.2, and 3.4.3 that deal with collinear cameras. □

We now illustrate the chiral joint image with the help of an example.

Example 3.4.1. *Consider the pair of cameras* $A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}$, $A_2 = \begin{bmatrix} I & \mathbf{t} \end{bmatrix}$ *where* $\mathbf{t} = (1,1,1)^\top$. *We depict the set* $C_{\mathcal{A}}$ *and its intersection with* $\overline{\mathcal{J}}_{\mathcal{A}}$ *in Figure 3.6 and Figure 3.7. Fixing* $\mathbf{p}_1$, *we plot the points on the corresponding epipolar line in the second image satisfying the inequalities defining* $C_{\mathcal{A}}$. *Restricting to* $p_{23} = 1$, *we observe that one inequality is quadratic in the* $\mathbf{p}_2$ *factor while the other two inequalities are linear.*

Observe how strong of a constraint chirality is by comparing the length of the epipolar line (dashed black) in each image to the length of the chiral joint image (solid black).

The set of inequalities defining $C_{\mathcal{A}}$ offer a powerful new tool to constrain triangulation and multiview stereo matching algorithms.

3.5 Chiral Reconstructions

We now revisit results of Hartley from [24] and [23, Chapter 21] on conditions under which a projective reconstruction of a collection of image correspondences

$$\mathcal{P} := \{(\mathbf{p}_{1k}, \ldots, \mathbf{p}_{mk}) \in (\mathbb{R}^2)^m,\ k = 1, \ldots, n\}$$

can be made chiral by a homography of $\mathbb{P}^3$. Hartley proves that deciding when a reconstruction can be made chiral reduces to solving a pair of linear programs. We give a conic interpretation of this in $\mathbb{R}^4$ and relate to quasi-affine transformations. Furthermore, we show that transforming to a chiral reconstruction naturally gives rise to the algebraic notion of signing a reconstruction. We describe how projective reconstructions that can be signed interpolate between projective and chiral reconstructions. In the special case of two cameras, Hartley shows that reconstructions that can be signed coincide with those that can be made chiral. We construct an example to show that when $m \geq 3$, these sets do not coincide in the sense that signed reconstructions may fail to become chiral.

Our chirality framework in projective space works seamlessly with projective reconstructions and directly leads to Hartley's results. Our approach, based on the chiral domain, leads to algebraically concise arguments based on geometric intuition. In particular, we give a simple polyhedral proof of Hartley's two view result.

3.5.1 Projective Reconstructions

A *projective reconstruction* of $\mathcal{P}$ is a pair $(\mathcal{A}, \mathcal{Q})$ consisting of an arrangement of m finite cameras $\mathcal{A} := \{A_1, \ldots, A_m\}$ and a set of n points $\mathcal{Q} := \{\mathbf{q}_1, \ldots, \mathbf{q}_n\} \subseteq \mathbb{R}^4 \setminus \{\mathbf{0}\}$ such that $A_i \mathbf{q}_k = w_{ik} \widehat{\mathbf{p}}_{ik}$ for some scalars w_{ik}.

Recall that $\mathbf{n}_i^\top \mathbf{q}_k = \det(G_i) w_{ik}$ for all i, k. Since all $\widehat{\mathbf{p}}_{ik}$ have last coordinate 1, $w_{ik} \neq 0$, and since all cameras are finite $\det(G_i) \neq 0$ for all i. Therefore, $\mathbf{n}_i^\top \mathbf{q}_k \neq 0$ for any i, k, or equivalently, no point $\mathbf{q}_k \in \mathcal{Q}$ lies on the principal plane of any camera A_i. Define $\sigma_{ik} = \operatorname{sgn}(\mathbf{n}_i^\top \mathbf{q}_k) \in \{-1, 1\}$ for all i, k.

Definition 3.5.1. *A* chiral reconstruction *of $\mathcal{P}$ is a projective reconstruction $(\mathcal{A}, \mathcal{Q})$ of $\mathcal{P}$ such that $\mathbf{q}_k \in D_{\mathcal{A}}$ for all k.*

We call a reconstruction *finite* if the points in $\mathcal{Q}$ are finite. Recall that cameras are already required to be finite. In [24] and [53] a *finite* projective reconstruction of $\mathcal{P}$ is called a *weak realization*, and a *finite* chiral reconstruction a *strong realization*. If there is a projective reconstruction of $\mathcal{P}$, there is always a finite one [32].

Consider now the question of when a given projective reconstruction $(\mathcal{A}, \mathcal{Q})$ of $\mathcal{P}$ can be transformed to a *projectively equivalent* reconstruction $(\mathcal{A}H^{-1}, H\mathcal{Q})$ that is chiral, by a homography $H \in \mathrm{GL}_4$ of $\mathbb{P}^3$ where

$$\mathcal{A}H^{-1} := \{A_1H^{-1}, \ldots, A_mH^{-1}\},$$
$$H\mathcal{Q} := \{H\mathbf{q}_1, \ldots, H\mathbf{q}_n\}.$$

The following lemma (parts of which appear in [24], and the proof of which given in Section 3.7) describes the effects of a homography H on a reconstruction.

Lemma 3.5.1. *Let $A = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}$ be a finite camera with center $\mathbf{c}_A$. Let $H \in \mathrm{GL}_4$ with last row $\mathbf{h}^\top$ and $\delta := \det(H^{-1})$. Then*

1. *Under the homography $\mathbf{q} \mapsto H\mathbf{q}$, the plane $\mathbf{h}^\top\mathbf{q} = \mathbf{0}$ maps to the plane at infinity.*

2. *The camera AH^{-1} is finite if and only if $\mathbf{h}^\top\mathbf{c}_A \neq 0$. Its center then is $\mathbf{c}_{AH^{-1}} = \frac{1}{\mathbf{h}^\top\mathbf{c}_A}H\mathbf{c}_A$.*

3. *The principal ray of AH^{-1} is*
$$\mathbf{n}_{AH^{-1}} = \delta(\mathbf{h}^\top\mathbf{c}_A)H^{-\top}\mathbf{n}_A.$$

4. *For all $\mathbf{q} \in \mathbb{R}^4$, we have*
$$\mathbf{n}_{AH^{-1}}^\top(H\mathbf{q}) = \delta(\mathbf{h}^\top\mathbf{c}_A)(\mathbf{n}_A^\top\mathbf{q}).$$

For a $H \in \mathrm{GL}_4$ with last row $\mathbf{h}^\top$, the reconstruction $(\mathcal{A}H^{-1}, H\mathcal{Q})$ of $\mathcal{P}$ is chiral if and only if for each k, $H\mathbf{q}_k$ lies in the chiral domain $D_{\mathcal{A}H^{-1}}$ of the camera arrangement $\mathcal{A}H^{-1}$. Therefore, from Theorem 3.3.1, Lemma 3.5.1, and the requirement that cameras in the chiral reconstruction need to be finite, i.e., $\mathbf{h}^\top\mathbf{c}_i \neq 0$ for all i, $(\mathcal{A}H^{-1}, H\mathcal{Q})$ is chiral if and only if one of the following sets $\mathcal{S}_1$ or $\mathcal{S}_2$ is nonempty:

$$\mathcal{S}_1 = \left\{ \mathbf{h} \mid \forall i, j, k, \quad \begin{array}{c} (\mathbf{h}^\top\mathbf{q}_k)(\mathbf{h}^\top\mathbf{c}_i)\sigma_{ik} \geq 0, \\ (\mathbf{h}^\top\mathbf{c}_i)(\mathbf{h}^\top\mathbf{c}_j)\sigma_{ik}\sigma_{jk} > 0 \end{array} \right\} \tag{3.10}$$

$$\mathcal{S}_2 = \left\{ \mathbf{h} \mid \forall i, j, k, \quad \begin{array}{c} (\mathbf{h}^\top\mathbf{q}_k)(\mathbf{h}^\top\mathbf{c}_i)\sigma_{ik} \leq 0, \\ (\mathbf{h}^\top\mathbf{c}_i)(\mathbf{h}^\top\mathbf{c}_j)\sigma_{ik}\sigma_{jk} > 0 \end{array} \right\} \tag{3.11}$$

This inspires the algebraic notion of *signing* a reconstruction.

Definition 3.5.2. *A* signed reconstruction $(\mathcal{A}, \mathcal{Q}^s)$ *of* $\mathcal{P}$ *is a projective reconstruction of* $\mathcal{P}$ *in which for each camera* i*, there exist constants* $\sigma_i^s \in \{-1, 1\}$ *such that* $\operatorname{sgn}\left(\mathbf{n}_i^\top \mathbf{q}_k^s\right) = \sigma_i^s$ *for all* k*. We say that a projective reconstruction* $(\mathcal{A}, \mathcal{Q})$ *can be signed if there exist* $\mathbf{q}_k^s \in \mathbb{R}^4$ *such that* $\mathbf{q}_k^s \sim \mathbf{q}_k$ *in* $\mathbb{P}^3$ *and* $(\mathcal{A}, \mathcal{Q}^s)$ *is a signed reconstruction.*

Note that signing a projective reconstruction $(\mathcal{A}, \mathcal{Q})$ only amounts to changing the sign of some world points. It does not affect the cameras or chirality of the world points in these cameras. Geometrically, signing a reconstruction puts all chosen representatives of the world points in the same half space in $\mathbb{R}^4$, of the principal plane of each camera.

The following result shows that being able to sign a projective reconstruction is necessary to transform to a chiral reconstruction.

Lemma 3.5.2. *Suppose that a projective reconstruction* $(\mathcal{A}, \mathcal{Q})$ *of* $\mathcal{P}$ *is projectively equivalent to a chiral reconstruction of* $\mathcal{P}$*. Then for each pair* i, j*, the product* $\sigma_{ik}\sigma_{jk}$ *is constant for all* k*, and* $(\mathcal{A}, \mathcal{Q})$ *can be signed.*

Proof. We saw that the projective reconstruction $(\mathcal{A}, \mathcal{Q})$ can be made chiral only if either $\mathcal{S}_1$ or $\mathcal{S}_2$ is non-empty which happens only if for each pair i, j, the product $\sigma_{ik}\sigma_{jk}$ is constant for every k. In this case, we show that $(\mathcal{A}, \mathcal{Q})$ can be signed. For each k, define $\mathbf{q}_k^s := \mathbf{q}_k$ if $\sigma_{1k} = 1$ or $\mathbf{q}_k^s := -\mathbf{q}_k$ if $\sigma_{1k} = -1$. By construction, $\sigma_{1k}^s := \operatorname{sgn}(\mathbf{n}_1^\top \mathbf{q}_k^s) = 1$ for all k. After this change, we still have $(\sigma_{1k}^s \sigma_{ik}^s)$ is constant for all k. Then it follows that for each i, σ_{ik}^s is constant for all k, and $(\mathcal{A}, \mathcal{Q}^s)$ is a signed reconstruction of $\mathcal{P}$. $\square$

The main result of this section is the following.

Theorem 3.5.1. *Given a signed reconstruction* $(\mathcal{A}, \mathcal{Q})$ *of* $\mathcal{P}$*, there is a* $H \in \mathrm{GL}_4$ *such that* $(\mathcal{A}H^{-1}, H\mathcal{Q})$ *is a chiral reconstruction if and only if*

$$K_{\mathcal{Q}}^* \cap (\operatorname{int} K_{\sigma C}^* \cup \operatorname{int} K_{-\sigma C}^*) \neq \{\mathbf{0}\} \tag{3.12}$$

where $K_{\sigma C} = \operatorname{cone}\{\sigma_1 \mathbf{c}_1, \ldots, \sigma_m \mathbf{c}_m\}$*, and* $\operatorname{int} K_{\sigma C}^*$ *is the interior of its dual cone, and* $K_{\mathcal{Q}} = \operatorname{cone}\{\mathbf{q}_1, \ldots, \mathbf{q}_n\}$*, and* $K_{\mathcal{Q}}^*$ *is its dual cone.*

Proof. (Theorem 3.5.1) Since $(\mathcal{A}, \mathcal{Q})$ is a signed reconstruction, we may substitute the constants σ_i for σ_{ik}. Then $\mathcal{S}_1$ is the union of the cones $(K^*_{\mathcal{Q}} \cap \operatorname{int} K^*_{\sigma C})$ and $(K^*_{-\mathcal{Q}} \cap \operatorname{int} K^*_{-\sigma C})$. Similarly, $\mathcal{S}_2$ is the union of $(K^*_{-\mathcal{Q}} \cap \operatorname{int} K^*_{\sigma C})$ and $(K^*_{\mathcal{Q}} \cap \operatorname{int} K^*_{-\sigma C})$. Since $K^*_{\mathcal{Q}} \cap \operatorname{int} K^*_{\sigma C} \neq \{\mathbf{0}\}$ if and only if $K^*_{-\mathcal{Q}} \cap \operatorname{int} K^*_{-\sigma C} \neq \{\mathbf{0}\}$, and $K^*_{-\mathcal{Q}} \cap \operatorname{int} K^*_{\sigma C} \neq \{\mathbf{0}\}$ if and only if $K^*_{\mathcal{Q}} \cap \operatorname{int} K^*_{-\sigma C} \neq \{\mathbf{0}\}$, finding a chiral reconstruction reduces to checking whether $K^*_{\mathcal{Q}}$ intersects one of the cones $\operatorname{int} K^*_{\sigma C}$ or $\operatorname{int} K^*_{-\sigma C}$. $\square$

The inequalities presented by the cone conditions in Theorem 3.5.1 are essentially Hartley's chiral inequalities [23, Equation 21.5]. In the rest of this section, we make several remarks about Theorem 3.5.1 and show how it relates to Hartley's results.

Quasi-affine transformations

We now interpret Equation (3.10) and Equation (3.11) geometrically. Lemma 3.5.1 shows that the effect of a homography on chirality is determined by the hyperplane it sends to infinity and its position relative to the camera centers and world points. The last row of a 4×4 matrix H representing a homography is the normal vector $\mathbf{h}$ of an oriented hyperplane in $\mathbb{R}^4$. The second conditions of Equation (3.10), $(\mathbf{h}^\top \mathbf{c}_i)(\mathbf{h}^\top \mathbf{c}_j)\sigma_{ik}\sigma_{jk} > 0$, express that the camera centers should all be in the same (open) half-space given by $\mathbf{h}$. The first conditions of Equation (3.10), $(\mathbf{h}^\top \mathbf{q}_k)(\mathbf{h}^\top \mathbf{c}_i)\sigma_{ik} \geq 0$, say that the (possibly resigned by σ_{ik}) world points and camera centers also lie in the same half spaces of this oriented hyperplane (or, in case of equality, combined with the second conditions, the world point lies on the hyperplane). Hence, these conditions encode a linear separation condition on the given points in $\mathbb{R}^4$, which can be checked via linear programming. The geometric interpretation of the inequalities in Equation (3.11) is analogous: In this case, the (possibly resigned) world points lie in the opposite closed half-space defined by $\mathbf{h}$ to the the camera centers.

Hartley presents this geometry in terms of *quasi-affine transformations* in both [24] and [23, Chapter 21]. In [23, Definition 21.3], a homography H is said to be quasi-affine with respect to a set $\mathcal{X} \subseteq \mathbb{R}^4$, with elements having last coordinate 1, if no point in the convex hull of $\mathcal{X}$ is sent to infinity by H. We observe that this is equivalent to saying that $\mathbf{h}$, the last row of H, lies in $\operatorname{int} K^*_{\mathcal{X}}$ or $\operatorname{int} K^*_{-\mathcal{X}}$. To accommodate infinite points, we make a more general definition of a quasi-affine transformation.

Definition 3.5.3. *A linear map $H \in \mathrm{GL}_4$ is* quasi-affine *with respect to $\mathcal{X} \subseteq \mathbb{R}^4$ if the last row $\mathbf{h}$ of H is in $K^*_{\mathcal{X}} \cup K^*_{-\mathcal{X}}$. Further, H is* strictly quasi-affine *with respect to $\mathcal{X}$ if $\mathbf{h} \in \operatorname{int} K^*_{\mathcal{X}} \cup \operatorname{int} K^*_{-\mathcal{X}}$.*

Geometrically, H is quasi-affine with respect to $\mathcal{X}$ if $H\mathcal{X}$ lies in one of the closed halfspaces defined by the hyperplane $\mathbf{h}^\perp = \{\mathbf{x} \in \mathbb{R}^4 : \mathbf{h}^\top \mathbf{x} = 0\}$, which is the plane sent to infinity by the homography H. If $H\mathcal{X}$ lies in a open halfspace of $\mathbf{h}^\perp$ (as in Hartley's setup) then H is strictly quasi-affine with respect to $\mathcal{X}$.

Recall that in a signed reconstruction $(\mathcal{A}, \mathcal{Q})$ we have fixed the sign of the last coordinates of all $\mathbf{q}_k \in \mathcal{Q} \subseteq \mathbb{R}^4$ and of all $\sigma_i \mathbf{c}_i$, and all points in $\mathcal{Q}$ and σC are considered to be in $\mathbb{R}^4$. We first show that Theorem 3.5.1 can be interpreted in terms of quasi-affine transformations.

Theorem 3.5.2. *Suppose $(\mathcal{A}, \mathcal{Q})$ is a signed reconstruction of $\mathcal{P}$. Then there exists a chiral reconstruction of $\mathcal{P}$ if and only if there is a homography H that is quasi-affine with respect to $\mathcal{Q}$ and strictly quasi-affine with respect to σC.*

Proof. The intersection $K^*_Q \cap \left(\operatorname{int}(K^*_{\sigma C}) \cup \operatorname{int}(K^*_{-\sigma C})\right)$ is nonempty if and only if

$$\left(K^*_Q \cup K^*_{-Q}\right) \cap \left(\operatorname{int}(K^*_{\sigma C}) \cup \operatorname{int}(K^*_{-\sigma C})\right)$$

is nonempty since

$$\mathbf{x} \in K^*_Q \cap \left(\operatorname{int}(K^*_{\sigma C}) \cup \operatorname{int}(K^*_{-\sigma C})\right) \iff -\mathbf{x} \in K^*_{-Q} \cap \left(\operatorname{int}(K^*_{\sigma C}) \cup \operatorname{int}(K^*_{-\sigma C})\right).$$

The statement now follows from Theorem 3.5.1. ☐

Hartley's chiral inequalities are strict while our cone conditions in Theorem 3.5.1 allow vectors $\mathbf{h}$ in the boundary of $K^*_{\mathcal{Q}}$. Theorem 3.3.1 leads to non-strict inequalities. The reason to pass to $\operatorname{int} K_{\sigma C}$ and $\operatorname{int} K_{-\sigma C}$ was because of the need for finite cameras in a chiral reconstruction. In fact, we could have restricted to $\operatorname{int} K^*_{\mathcal{Q}}$ in Theorem 3.5.1 which would exactly give Hartley's chiral inequalities. This is because if the $\mathbf{h}$ produced in Theorem 3.5.1 lies in the boundary of $K^*_{\mathcal{Q}}$, we may replace it by one in $\operatorname{int} K^*_{\mathcal{Q}}$ by continuity. Geometrically this means that if $\mathcal{P}$ has a chiral reconstruction, then it has one in which all world points are finite and do not lie on any principal planes.

Hartley's work was done with the aim of upgrading a two view projective reconstruction to a metric reconstruction. In follow up work, Nistér addresses this question for multiple views [39]. He does this by transforming the projective reconstruction into one which is quasi-affine with respect to the camera centers. As can be seen from Theorem 3.5.2 above, quasi-affineness with respect to the camera centers is a necessary condition for chirality. He does not enforce quasi-affineness with respect to the scene points, because they are often noisy and their chirality may change as part of the metric upgrade. Nistér shows that enforcing the quasi-affineness on camera centers makes the iterative algorithm used to perform the subsequent metric upgrade easier and more reliable.

Two-view chirality

In this section, we show that a two view reconstruction can be made chiral if and only if it can be signed.

Suppose $(\{A_1, A_2\}, \mathcal{Q})$ is a two-view reconstruction of $\mathcal{P}$ such that A_1 and A_2 have distinct centers, $A_1\mathbf{q}_k = w_{1k}\widehat{\mathbf{p}}_{1k}$, and $A_2\mathbf{q}_k = w_{2k}\widehat{\mathbf{p}}_{2k}$. Theorem 17 in [24] (also [53, Theorem 1]) gives a necessary and sufficient condition for when a two-view projective reconstruction can be transformed by a homography to a chiral reconstruction. We have restated this result in our language in Theorem 3.5.3 below. For the translation, recall that $\mathbf{n}_1^\top\mathbf{q}_k = \det(G_1)w_{1k}$, $\mathbf{n}_2^\top\mathbf{q}_k = \det(G_2)w_{2k}$ and $\sigma_{ik} = \operatorname{sgn}(\mathbf{n}_i^\top\mathbf{q}_k)$. Therefore, the products $w_{1k}w_{2k}$ have the same sign for all k if and only if $(\mathbf{n}_1^\top\mathbf{q}_k)(\mathbf{n}_2^\top\mathbf{q}_k)$ have the same sign for all k, i.e., $\sigma_{1k}\sigma_{2k}$ is constant for all k.

The "only if" direction of Theorem 3.5.3 appears in both [24, Theorem 17] and [23, Theorem 21.7 (i)], and the proof is straightforward. We also give the argument in Lemma 3.5.2.

The "if" direction appears in [24] with a rather complicated proof, and not in [23]. We provide a short polyhedral proof of the "if" direction using Theorem 3.5.1.

Theorem 3.5.3. *[24, Theorem 17] A projective reconstruction $(\{A_1, A_2\}, \mathcal{Q})$ of $\mathcal{P}$ can be transformed by a homography H to a chiral reconstruction if and only if $(\mathbf{n}_1^\top \mathbf{q}_k)(\mathbf{n}_2^\top \mathbf{q}_k)$ have the same sign for all k.*

Proof. Suppose $(\mathbf{n}_1^\top \mathbf{q}_k)(\mathbf{n}_2{}^\top \mathbf{q}_k)$ have the same sign for all k. Then by Lemma 3.5.2, $\mathbf{n}_1^\top \mathbf{q}_k = \sigma_{1k} = \sigma_1$ and $\mathbf{n}_2^\top \mathbf{q}_k = \sigma_{2k} = \sigma_2$ for all k.

We first note that $\sigma_1 \mathbf{n}_1$ is a nonzero element of either $K^*_{\sigma C}$ or $K^*_{-\sigma C}$. We have that $(\sigma_1 \mathbf{n}_1)^\top(\sigma_1 \mathbf{c}_1) = 0$. If $\mathrm{sgn}(\sigma_1 \mathbf{n}_1)^\top(\sigma_2 \mathbf{c}_2) = 1$ or $\mathrm{sgn}(\sigma_1 \mathbf{n}_1)^\top(\sigma_2 \mathbf{c}_2) = 0$, then $\sigma_1 \mathbf{n}_1 \in K^*_{\sigma C}$. Otherwise if $\mathrm{sgn}(\sigma_1 \mathbf{n}_1)^\top(\sigma_2 \mathbf{c}_2) = -1$, then $\sigma_1 \mathbf{n}_1 \in (K^*_{-\sigma C})$.

Also, since the centers $\mathbf{c}_1$ and $\mathbf{c}_2$ are distinct, $\sigma_1 \mathbf{c}_1$ is not a scalar multiple of $\sigma_2 \mathbf{c}_2$, hence $K_{\sigma C}$ is a pointed cone. This implies that $K^*_{\sigma C}$ is full-dimensional and hence has an interior. The same is true for $K^*_{-\sigma C}$.

Without loss of generality suppose $\sigma_1 \mathbf{n}_1 \in K^*_{\sigma C}$. Since $\mathrm{sgn}(\mathbf{n}_1^\top \mathbf{q}_k) = \sigma_1$, we have that $\sigma_1 \mathbf{n}_1^\top \mathbf{q}_k > 0$ for all k, and so $\sigma_1 \mathbf{n}_1 \in \operatorname{int} K^*_{\mathcal{Q}}$. Let U be a neighborhood of $\sigma_1 \mathbf{n}_1$ contained in $\operatorname{int} K^*_{\mathcal{Q}}$. Since $\sigma_1 \mathbf{n}_1$ is also in $K^*_{\sigma C}$, there is some $\mathbf{h} \in U$ that lies in the $\operatorname{int} K^*_{\mathcal{Q}} \cap \operatorname{int} K^*_{\sigma C}$. This $\mathbf{h}$ is in $\mathcal{S}_1$, so by Theorem 3.5.1, $\mathcal{P}$ has a chiral reconstruction. $\square$

Theorem 3.5.3 shows that a chiral reconstruction exists if and only if $(\mathbf{n}_1^\top \mathbf{q}_k)(\mathbf{n}_2^\top \mathbf{q}_k)$ has the same sign for all k. Lemma 3.5.2 shows that this is equivalent to being able to sign the reconstruction $(\{A_1, A_2\}, \mathcal{Q})$. Hence a two view reconstruction can be made chiral if and only if it can be signed. Our notion of signing readily generalizes to multiple views. However, the following example shows that the "if" direction of Theorem 3.5.3 does not generalize to multiple views. In other words, it may not be possible to transform a signed reconstruction with three or more cameras into a chiral one.

Example 3.5.1. *Consider the reconstruction*

$$(\mathcal{A} = \{A_1, A_2, A_3\}, \mathcal{Q} = (\mathbf{q}_1, \mathbf{q}_2))$$

where

$$A_1 = \begin{bmatrix} 0 & 0 & -1 & -1 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix}, A_2 = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 0 & -1 & 1 \\ 0 & 1 & 0 & 0 \end{bmatrix}, A_3 = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & 0 \end{bmatrix},$$

and $\mathbf{q}_1 = (1, 1, 2, -6)^\top$ *and* $\mathbf{q}_2 = (1, 1, 2, 6)^\top$.

The reconstruction is signed and $\sigma_1 = \sigma_2 = \sigma_3 = 1$ *because* $\sigma_{ik} = \operatorname{sgn}(\mathbf{n}_i^\top \mathbf{q}_k) = 1$ *for all* i, k. *However,* $(\mathcal{A}, \mathcal{Q})$ *is not chiral because* $\mathbf{q}_1$ *is not in* $D_{\mathcal{A}}$. *Indeed, check that* $(\mathbf{n}_\infty^\top \mathbf{q}_1)(\mathbf{n}_i^\top \mathbf{q}_1) < 0$ *for all* $i = 1, 2, 3$.

We argue that $(\mathcal{A}, \mathcal{Q})$ *is not projectively equivalent to a chiral reconstruction using the conditions of Theorem 3.5.1. Consider the matrices* $M^+ = \begin{bmatrix} \mathbf{c}_1 & \mathbf{c}_2 & \mathbf{c}_3 & \mathbf{q}_1 & \mathbf{q}_2 \end{bmatrix}$ *and* $M^- = \begin{bmatrix} -\mathbf{c}_1 & -\mathbf{c}_2 & -\mathbf{c}_3 & \mathbf{q}_1 & \mathbf{q}_2 \end{bmatrix}$, *or explicitly*

$$M^+ = \begin{bmatrix} 0 & 1 & -1 & 1 & 1 \\ -1 & 0 & 1 & 1 & 1 \\ -1 & 1 & 0 & 2 & 2 \\ 1 & 1 & 1 & -6 & 6 \end{bmatrix}, \quad M^- = \begin{bmatrix} 0 & -1 & 1 & 1 & 1 \\ 1 & 0 & -1 & 1 & 1 \\ 1 & -1 & 0 & 2 & 2 \\ -1 & -1 & -1 & -6 & 6 \end{bmatrix}.$$

Both M^+ *and* M^- *have a strictly positive kernel element. In particular,* $M^+\mathbf{v}^+ = 0$ *and* $M^-\mathbf{v}^- = 0$ *where* $\mathbf{v}^+ = (11, 1, 6, 4, 1)^\top > 0$ *and* $\mathbf{v}^- = (1, 11, 6, 1, 4)^\top > 0$. *Existence of* $\mathbf{v}^+$ *and* $\mathbf{v}^-$ *shows that the linear systems*

$$\{\mathbf{h} : \mathbf{h} \neq 0, \mathbf{h}^\top M^+ \geq 0\} \text{ and } \{\mathbf{h} : \mathbf{h} \neq 0, \mathbf{h}^\top M^- \geq 0\}$$

are infeasible. Indeed, suppose there is a $\mathbf{h} \neq 0$ *such that* $\mathbf{h}^\top M^+ = \mathbf{y} \geq 0$. *As* M^+ *has full rank,* $\mathbf{y} \neq 0$. *Since* $\mathbf{v}^+$ *is in the null space of* M^+, *it is orthogonal to the row space of* M^+. *It follows that* $\mathbf{y}^\top \mathbf{v}^+ = 0$, *but this is a contradiction because* $\mathbf{v}^+$ *is strictly positive and* $\mathbf{y} \neq 0$. *An analogous argument applies to the linear system involving* M^-.

Translated into cone language, this means $(K_{\mathcal{Q}}^* \cap \operatorname{int} K_C^*) \cup (K_{\mathcal{Q}}^* \cap \operatorname{int} K_{-C}^*) = \varnothing$. *By Theorem 3.5.1,* $(\mathcal{A}, \mathcal{Q})$ *is not projectively equivalent to a chiral reconstruction.*

In general, if we have $m \geq 3$ cameras, then the hyperplanes with normals $\sigma_i \mathbf{c}_i$ partition $\mathbb{R}^4$ into 2^m (possibly empty) regions, each indexed by an element of $\{+,-\}^m$. It can be that K_Q^* lies entirely in a region of mixed signs forcing $K_Q^* \cap (\operatorname{int} K_{\sigma C}^* \cup \operatorname{int} K_{-\sigma C}^*) = \emptyset$.

For two cameras, this does not happen as we saw in the proof of Theorem 3.5.3; $\sigma_1 \mathbf{n}_1$ is a non-zero element in $(K_Q^* \cap K_{\sigma C}^*) \cup (K_Q^* \cap K_{-\sigma C}^*)$. This relied crucially on the fact that $\sigma_1 \mathbf{n}_1$ is on the hyperplane with normal $\sigma_1 \mathbf{c}_1$ which is divided into two halfspaces by the hyperplane with normal $\sigma_2 \mathbf{c}_2$. Regardless of which half space $\sigma_1 \mathbf{n}_1$ lies in, it belongs to either $K_{\sigma C}^*$ or $K_{-\sigma C}^*$, i.e., it is automatically in either the $++$ or $--$ regions of hyperplanes with normal $\sigma_1 \mathbf{c}_1$ and $\sigma_2 \mathbf{c}_2$. This argument works for any number of cameras if $K_{\sigma N} := \operatorname{cone}(\sigma_1 \mathbf{n}_1, \ldots, \sigma_m \mathbf{n}_m)$ intersects $K_{\sigma C}^*$ or $K_{-\sigma C}^*$ because $K_{\sigma N} \subseteq K_Q^*$ for any signed reconstruction. For $m > 2$ cameras, it can be that $K_{\sigma N}$, and even all of K_Q^*, lies in a region of mixed signs of the hyperplane arrangement with oriented normals $\sigma_i \mathbf{c}_i$, as in Example 3.5.1.

3.5.2 Euclidean Reconstructions

In the previous section, we asked when a projective reconstruction can be transformed to a chiral reconstruction. We now ask the same question for a Euclidean reconstruction of $\mathcal{P}$, by which we mean a projective reconstruction $(\mathcal{A}, \mathcal{Q})$ in which each camera has the form $\begin{bmatrix} R & \mathbf{t} \end{bmatrix}$ where $R \in SO(3)$.

Unlike for projective reconstructions, it is not true that if a Euclidean reconstruction exists, there is always one that is finite. However, this is not a problem since our definition of chiral reconstruction allows world points to be infinite thus generalizing the old notion of a strong realization.

Proposition 10 in [24] shows that we can assume $A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}$ by applying an appropriate similarity, without affecting chirality. Under this assumption, the following two theorems (whose proofs appear in Section 3.7) answer the above question for $m = 2$ and $m > 2$ views respectively.

Theorem 3.5.4. *Let* $(\{A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}, A_2 = \begin{bmatrix} R & \mathbf{t} \end{bmatrix}\}, \mathcal{Q})$ *be a signed Euclidean reconstruction* *of $\mathcal{P}$ with distinct centers. There exists a chiral Euclidean reconstruction of $\mathcal{P}$ if and only if $\mathbf{n}_\infty \in K^*_{\mathcal{Q}} \cup K^*_{-\mathcal{Q}}$ or $\mathbf{r} := \begin{bmatrix} -\frac{2}{\|\mathbf{t}\|^2} R^\top \mathbf{t} \\ 1 \end{bmatrix} \in K^*_{\mathcal{Q}} \cup K^*_{-\mathcal{Q}}$. Equivalently, if exactly one of the following holds for all $\mathbf{q}_i$:*

$$q_{i4} \geq 0 \;\; \text{or} \;\; q_{i4} \leq 0 \;\; \text{or} \;\; \mathbf{r}^\top \mathbf{q}_i \geq 0 \;\; \text{or} \;\; \mathbf{r}^\top \mathbf{q}_i \leq 0.$$

Theorem 3.5.5. *Let $(\mathcal{A}, \mathcal{Q})$ be a signed Euclidean reconstruction of $\mathcal{P}$ with $m > 2$ cameras, distinct centers, and $A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}$. There exists a chiral Euclidean reconstruction of $\mathcal{P}$ if and only if $\sigma_i = \sigma_j$ for all $1 \leq i < j \leq m$ and either $q_{i4} \geq 0$ for all i or $q_{i4} \leq 0$ for all i.*

These theorems are specializations of Theorem 3.5.1. Their proofs are based on the observation that restricting the cameras to be Euclidean restricts the class of homographies in Theorem 3.5.1 to four ($m = 2$) and two ($m > 2$) discrete choices respectively. The four choices for $m = 2$ correspond to the well known twisted pair transformations and the two choices for $m > 2$ correspond to reflection.

3.6 Summary

We introduce the chiral domain of an arrangement of cameras — a multiview generalization of the definition of chirality that covers all of $\mathbb{P}^3$ — and give a semialgebraic description of this set.

We define the chiral joint image of a camera arrangement to be the image of the chiral domain in the cameras; it is the true image of the world in the cameras. The chiral joint image lives naturally in the joint image of the camera arrangement, a classical quasi-projective variety in multiview geometry. We provide a complete semialgebraic description of the chiral joint image.

The equations and inequalities describing the chiral joint image are the chiral analogs of the familiar multiview constraints. They lay the foundations for the development of a theory of chiral reconstruction. Our algebraic descriptions of the chiral domain and the chiral joint image can be used to enforce chirality when solving reconstruction or triangulation problems. Similarly, the chiral joint image can be used to constrain the region used for stereo matching.

The chiral domain also provides a polyhedral derivation of Hartley's chiral inequalities whose feasibility characterizes when a projective reconstruction can be made chiral by a homography. A similar result is given for Euclidean cameras. Our polyhedral formulation also extends to quasi-affine transformations which are central to Hartley's work on chirality. Our approach provides a simple proof of the hard direction of Hartley's theorem that says that a two-view reconstruction can be made chiral by a homography if and only if the reconstruction satisfies a sign condition. We provide an example to show that such a sign condition does not suffice when there are more than two cameras.

3.7 Technical Proofs

This section contains the proofs of statements not proved in the main body of the paper due to reasons of narrative clarity. The numbering of theorems and lemmas matches those in the main paper and they are presented here in the order in which they appear in the main paper. In some cases, these proofs rely on additional lemmas (Lemmas 3.7.2, 3.7.3, 3.7.4, and 3.7.5) which are only present in this section. As a result, some lemmas appear out of order because they are presented in the order they are needed.

Proofs from Section 3.4

Recall from Definition 3.4.3 that $\mathbf{a}_i = G_i^{-1}\mathbf{p}_i$ and $\mathbf{b}_{ij} = G_i^{-1}\mathbf{t}_i - G_j^{-1}\mathbf{t}_j$. Throughout this section, we denote the baseline of finite cameras A_i and A_j by l_{ij}, i.e.,

$$l_{ij} := \{\mathbf{q} \in \mathbb{P}^3 : \exists \lambda_1, \lambda_2 \in \mathbb{R} \text{ s.t. } \mathbf{q} = \lambda_1 \mathbf{c}_1 + \lambda_2 \mathbf{c}_2\} \tag{3.13}$$

Lemma 3.7.1. *Let $\mathcal{A} = \{A_1, A_2\}$ be a pair of finite cameras $A_i = \begin{bmatrix} G_i & \mathbf{t}_i \end{bmatrix}$ with distinct centers. Fix $\mathbf{q} \in \mathbb{P}^3 \setminus \{\mathbf{c}_1, \mathbf{c}_2\}$ and write $(\mathbf{p}_1, \mathbf{p}_2) \sim \varphi_{\mathcal{A}}(\mathbf{q})$ with $A_i\mathbf{q} = \lambda_i \mathbf{p}_i$. The vectors $\mathbf{a}_1, \mathbf{a}_2$ and $\mathbf{b}_{12}$ are collinear in $\mathbb{R}^3$ if and only if $\mathbf{q} \in l_{12}$. In this case, $\mathbf{p}_1 = \mathbf{e}_{12}$ and $\mathbf{p}_2 = \mathbf{e}_{21}$.*

Proof. We argue geometrically. The vector $(\mathbf{b}_{12}^\top, 0)^\top \in \mathbb{P}^3$ is the intersection point of the baseline l_{12} with the hyperplane L_∞. Moreover, a point $\mathbf{p} \in \mathbb{P}^2$ has a 1-dimensional family of preimages under a finite camera A which is the span of the vector $((G^{-1}\mathbf{p})^\top, 0)^\top = (\mathbf{a}^\top, 0)^\top \in \mathbb{P}^3$ and the (finite) center of camera A, where the camera center cannot be imaged in A, of course. Indeed, this follows from the fact that $A(\mathbf{a}^\top, 0)^\top = \mathbf{p}$ and the fact that the kernel of A is spanned by its center $\mathbf{c} \in \mathbb{P}^3$. Of course, we could also take the span of any other two points on this line. Below we will see the span of the center $\mathbf{c}$ and a point $\mathbf{q} \in \mathbb{P}^3$ with $A\mathbf{q} \sim \mathbf{p}$. We now apply these geometric facts to the epipoles and the baseline to show the two implications.

So first suppose that $\mathbf{a}_1$, $\mathbf{a}_2$, and $\mathbf{b}_{12}$ are collinear. Geometrically, this means that the lines of preimages $\{\mathbf{q} \in \mathbb{P}^3 \colon A_i\mathbf{q} \sim \mathbf{p}_i\}$ have the same intersection points with the hyperplane at infinity and that point is also the intersection point of l_{12} and L_∞. Since the baseline contains both centers $\mathbf{c}_1$ and $\mathbf{c}_2$ and the intersection points at infinity coincide, all three lines are equal to the baseline.

Conversely, if $\mathbf{q}$ is on the baseline but not a camera center, then the line spanned by $\mathbf{q}$ and any camera center $\mathbf{c}_i$ is the baseline l_{ij} and the line of preimages $\{\mathbf{q} \in \mathbb{P}^3 \colon A_i\mathbf{q} \sim \mathbf{p}_i\}$. □

Lemma 3.7.2. *Let $\mathcal{A} = \{A_1, A_2\}$ be a pair of finite cameras with distinct centers. Fix $\mathbf{q} \in \mathbb{P}^3 \setminus \{\mathbf{c}_1, \mathbf{c}_2\}$ and write $(\mathbf{p}_1, \mathbf{p}_2) \sim \varphi_{\mathcal{A}}(\mathbf{q})$ with $A_i\mathbf{q} = \lambda_i\mathbf{p}_i$. If $\mathbf{q} \notin l_{ij}$, then the following conditions hold:*

$$\mathbf{b}_{12}^\top(\mathbf{a}_1 \times \mathbf{a}_2) = 0, \tag{3.14}$$

$$\operatorname{sgn}(\lambda_1 q_4) = \operatorname{sgn}\left((\mathbf{a}_1 \times \mathbf{a}_2)^\top(\mathbf{b}_{12} \times \mathbf{a}_2)\right), \tag{3.15}$$

$$\operatorname{sgn}(\lambda_2 q_4) = \operatorname{sgn}\left((\mathbf{a}_1 \times \mathbf{a}_2)^\top(\mathbf{b}_{12} \times \mathbf{a}_1)\right), \tag{3.16}$$

$$\operatorname{sgn}(\lambda_1 \lambda_2) = \operatorname{sgn}\left((\mathbf{b}_{12} \times \mathbf{a}_1)^\top(\mathbf{b}_{12} \times \mathbf{a}_2)\right). \tag{3.17}$$

Proof. Since $\mathbf{q} \neq \mathbf{c}_1, \mathbf{c}_2$, we know $\lambda_1 \neq 0, \lambda_2 \neq 0$. Write $\mathbf{q} = (\mathbf{r}, q_4)$ and $\mathbf{b}$ for $\mathbf{b}_{12}$. Eliminating $\mathbf{r}$ by taking the difference of the equations $\lambda_1\mathbf{p}_1 = G_1\mathbf{r} + q_4\mathbf{t}_1$ and $\lambda_2\mathbf{p}_2 = G_2\mathbf{r} + q_4\mathbf{t}_2$, we get

$\lambda_1 G_1^{-1}\mathbf{p}_1 - \lambda_2 G_2^{-1}\mathbf{p}_2 = q_4\left(G_1^{-1}\mathbf{t}_1 - G_2^{-1}\mathbf{t}_2\right)$, equivalently,

$$\lambda_1 \mathbf{a}_1 - \lambda_2 \mathbf{a}_2 = q_4 \mathbf{b}. \tag{3.18}$$

Taking cross products with $\mathbf{a}_1, \mathbf{a}_2$, and $\mathbf{b}$ on both sides of (3.18), we get

$$-\lambda_2(\mathbf{a}_2 \times \mathbf{a}_1) = q_4(\mathbf{b} \times \mathbf{a}_1), \tag{3.19}$$

$$\lambda_1(\mathbf{a}_1 \times \mathbf{a}_2) = q_4(\mathbf{b} \times \mathbf{a}_2), \tag{3.20}$$

$$\lambda_1(\mathbf{b} \times \mathbf{a}_1) = \lambda_2(\mathbf{b} \times \mathbf{a}_2). \tag{3.21}$$

We now consider two cases:

Case a: Suppose $q_4 \neq 0$. Equation (3.18) implies that $\mathbf{b}$, $\mathbf{a}_1$, and $\mathbf{a}_2$ are coplanar in $\mathbb{R}^3$, so that (3.14) is satisfied. Further, it is straightforward to see from the cross product equations that either $\mathbf{a}_1 \times \mathbf{a}_2$, $\mathbf{b} \times \mathbf{a}_1$ and $\mathbf{b} \times \mathbf{a}_2$ are all equal to zero or none of them are, i.e. either $\mathbf{a}_1, \mathbf{a}_2$, and $\mathbf{b}$ are all pairwise collinear or not. From Lemma 3.7.1, our assumption that $\mathbf{q} \notin l_{ij}$ implies that $\mathbf{a}_1, \mathbf{a}_2$, and $\mathbf{b}$ are not all collinear. Hence, the equations (3.15),(3.16),(3.17) follow from multiplying each equality above by the transpose of the left hand side.

Case b: Suppose $q_4 = 0$. This implies that $\mathbf{a}_1$ and $\mathbf{a}_2$ are collinear in $\mathbb{R}^3$ and $\mathbf{a}_1 \times \mathbf{a}_2 = 0$. This proves (3.14), (3.15) and (3.16). Since by assumption $\mathbf{b}$ is not collinear with $\mathbf{a}_1$ and $\mathbf{a}_2$ (3.17) follows by multiplying the third equality above by its right hand side. $\square$

Lemma 3.4.1. *Let $\mathcal{A} = \{A_1, \ldots, A_m\}$ be an arrangement of finite cameras such that $D_{\mathcal{A}}$ is nonempty. If the centers of $\mathcal{A}$ are not collinear, then*

$$\mathcal{X}_{\mathcal{A}} = \mathcal{J}_{\mathcal{A}} \cap C_{\mathcal{A}}. \tag{3.22}$$

If the centers are collinear, then set $\mathbf{e} := (\mathbf{e}_1, \ldots, \mathbf{e}_m)$ to be the image of the common baseline under $\varphi_{\mathcal{A}}$. Then

$$\mathcal{X}_{\mathcal{A}} \setminus \{\mathbf{e}\} = (\mathcal{J}_{\mathcal{A}} \cap C_{\mathcal{A}}) \setminus \{\mathbf{e}\} \tag{3.23}$$

In both cases, $\overline{\mathcal{X}}_{\mathcal{A}} \subseteq C_{\mathcal{A}}$.

Proof. Suppose $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_m) = \varphi_{\mathcal{A}}(\mathbf{q})$ for a $\mathbf{q} = (\mathbf{r}, q_4)$ in $\mathbb{P}^3$. As before, we write $\mathbf{b}_{ij} = G_i^{-1}\mathbf{t}_i - G_j^{-1}\mathbf{t}_j$ and $\mathbf{a}_i = G_i^{-1}\mathbf{p}_i$. It suffices to show that $\mathbf{p} \in \mathcal{J}_{\mathcal{A}} \cap C_{\mathcal{A}}$ if and only if $\mathbf{p} \in \mathcal{X}_{\mathcal{A}}$. Recall that the principal ray of camera A_i is given by $\det(G_i)(A_i)_{3\bullet}$ and $p_{i3} = (A_i)_{3,\bullet}\mathbf{q}$. By Theorem 3.3.1, we know $\mathbf{p} \in \mathcal{X}_{\mathcal{A}}$ if and only if

$$\det(G_i)\lambda_i p_{i3} q_4 \geq 0 \tag{3.24}$$

$$\det(G_i)\det(G_j)\lambda_i p_{i3}\lambda_j p_{j3} \geq 0 \tag{3.25}$$

for all i, j. We will show that $\mathbf{p}$ satisfies these inequalities if and only if $\mathbf{p} \in \mathcal{J}_{\mathcal{A}} \cap C_{\mathcal{A}}$, i.e., if $\mathbf{p}$ satisfies the inequalities

$$\det(G_i)p_{i3}(\mathbf{a}_i \times \mathbf{a}_j)^\top(\mathbf{b}_{ij} \times \mathbf{a}_i) \geq 0 \tag{3.26}$$

$$\det(G_i)\det(G_j)p_{i3}p_{j3}(\mathbf{b}_{ij} \times \mathbf{a}_i)^\top(\mathbf{b}_{ij} \times \mathbf{a}_j) \geq 0 \tag{3.27}$$

for all i, j. We make some observations that follow for all $\mathbf{q}$ when the cameras in $\mathcal{A}$ are noncollinear.

1. For each camera A_i, there is some camera A_j such that $\mathbf{q} \notin l_{ij}$. Indeed, fixing i, the pencil of lines l_{ij} are not all identical, hence $\mathbf{q}$ cannot lie on all of them. Lemma 3.7.2 therefore implies that

$$\operatorname{sgn}(\lambda_i q_4) = \operatorname{sgn}(\mathbf{a}_i \times \mathbf{a}_j)^\top(\mathbf{b}_{ij} \times \mathbf{a}_i). \tag{3.28}$$

2. For each pair of cameras $\{A_i, A_j\}$ such that $\mathbf{q} \in l_{ij}$, there is a camera A_k such that $\mathbf{q} \notin l_{ik}$ and $\mathbf{q} \notin l_{jk}$. This follows from noncollinearity because for every line l_{ij}, there must be some center $\mathbf{c}_k$ not on this line. Lemma 3.7.2 therefore implies that

$$\begin{aligned} \operatorname{sgn}(\lambda_i\lambda_j) &= \operatorname{sgn}(\lambda_i\lambda_j\lambda_k^2) && (3.29)\\ &= \operatorname{sgn}(\lambda_i\lambda_k)\operatorname{sgn}(\lambda_j\lambda_k) && (3.30)\\ &= \operatorname{sgn}(\mathbf{b}_{ik} \times \mathbf{a}_i)^\top(\mathbf{b}_{ik} \times \mathbf{a}_k)(\mathbf{b}_{jk} \times \mathbf{a}_j)^\top(\mathbf{b}_{jk} \times \mathbf{a}_k). && (3.31) \end{aligned}$$

Suppose $\mathbf{p}$ satisfies the inequalities (3.24) and (3.25). Then for every pair i, j either

$$\det(G_i)p_{i3}(\mathbf{a}_i \times \mathbf{a}_j)^\top(\mathbf{b}_{ij} \times \mathbf{a}_i) = 0 \text{ if } \mathbf{q} \in l_{ij} \text{ or} \tag{3.32}$$

$$\operatorname{sgn}\left(\det(G_i)p_{i3}(\mathbf{a}_i \times \mathbf{a}_j)^\top(\mathbf{b}_{ij} \times \mathbf{a}_i)\right) = \operatorname{sgn}(\det(G_i)\lambda_i p_{i3} q_4) \geq 0 \tag{3.33}$$

Similar reasoning shows that (3.27) holds for all i, j. Conversely, suppose $\mathbf{p}$ satisfies the inequalities (3.26) and (3.27). From the observations above, we see that $\operatorname{sgn}(\lambda_i q_4)$ and $\operatorname{sgn}(\lambda_i \lambda_j)$ can be inferred from the inequalities (3.26) and (3.27). Hence (3.24) and (3.25) hold, meaning $\mathbf{p} \in \mathcal{X}_\mathcal{A}$. We conclude that

$$\mathcal{J}_\mathcal{A} \cap C_\mathcal{A} = \mathcal{X}_\mathcal{A}$$

In particular, this means $\mathcal{X}_\mathcal{A} \subseteq C_\mathcal{A}$, so $\overline{\mathcal{X}_\mathcal{A}} \subseteq C_\mathcal{A}$ because $C_\mathcal{A}$ is closed. The above argument holds for all $\mathbf{p}$ such that its preimage under $\varphi_\mathcal{A}$ is a unique $\mathbf{q}$. For collinear cameras this is true for all $\mathbf{p} \neq \mathbf{e}$, hence the only point that must be removed from (3.22) is $\mathbf{e}$. □

Lemma 3.7.3. *Let $\mathcal{A} = \{A_1, A_2\}$ be a pair of finite cameras. If $\mathbf{c}_j$ has nonnegative depth in the other camera A_i, then E_j is contained in $C_\mathcal{A}$. Otherwise $(\mathbf{e}_{12}, \mathbf{e}_{21})$ is the only point in E_j that lies in $C_\mathcal{A}$.*

Proof. Without loss of generality, we can assume $j = 2$. Let $A_i = \begin{bmatrix} G_i & \mathbf{t}_i \end{bmatrix}$ for $i = 1, 2$. We write $\widetilde{\mathbf{c}}_i = -G_i^{-1}\mathbf{t}_i$ and $\mathbf{b} = \widetilde{\mathbf{c}}_2 - \widetilde{\mathbf{c}}_1$. Let $\mathbf{s}_1 = G_1\widetilde{\mathbf{c}}_2 + \mathbf{t}_1 = G_1\left(-G_2^{-1}\mathbf{t}_2 + G_1^{-1}\mathbf{t}_1\right)$ $= G_1\mathbf{b}$. Then, the image of $\mathbf{c}_2$ in A_1 is $\mathbf{e}_{12} = \lambda_1\mathbf{s}_1$. Similarly, let $\mathbf{s}_2 = G_2\widetilde{\mathbf{c}}_1 + \mathbf{t}_2 =$ $G_2(-G_1^{-1}\mathbf{t}_1 + G_2^{-1}\mathbf{t}_2) = -G_2\mathbf{b}$. Then, the image of $\mathbf{c}_1$ in A_2 is $\mathbf{e}_{21} = \lambda_2\mathbf{s}_2$.

Now if $\mathbf{p}_1 = \mathbf{e}_{12} = \lambda_1 G_1\mathbf{b}$, then $\mathbf{a}_1 = \lambda_1\mathbf{b}$ and $\mathbf{b} \times \mathbf{a}_1 = 0$. Which means that the only inequality defining $C_\mathcal{A}$ not identically equal to zero is

$$\det(G_1)p_{13}(\mathbf{a}_1 \times \mathbf{a}_2)^\top(\mathbf{b} \times \mathbf{a}_2) \geq 0. \tag{3.34}$$

Plugging $\mathbf{a}_1 = \lambda_1\mathbf{b}$ and $\mathbf{p}_1 = \mathbf{e}_{12} = \lambda_1\mathbf{s}_1$ in the above we get

$$\det(G_1)\lambda_1 s_{13}(\lambda_1\mathbf{b} \times \mathbf{a}_2)^\top(\mathbf{b} \times \mathbf{a}_2) \geq 0 \tag{3.35}$$

$$\det(G_1)\lambda_1^2 s_{13}\|\mathbf{b} \times \mathbf{a}_2\|^2 \geq 0 \tag{3.36}$$

$$\det(G_1)s_{13}\|\mathbf{b} \times \mathbf{a}_2\|^2 \geq 0 \tag{3.37}$$

This can be satisfied in two ways, namely $\mathbf{b} \times \mathbf{a}_2 = 0$ or $\det(G_1)s_{13} \geq 0$.

Case 1: Suppose $\mathbf{b} \times \mathbf{a}_2 = 0$. Then since

$$\mathbf{b} \times \mathbf{a}_2 = 0 \iff \mathbf{a}_2 \sim \mathbf{b} \iff \mathbf{p}_2 \sim G_2\mathbf{b} \sim \mathbf{e}_{21}, \tag{3.38}$$

the condition $\mathbf{b} \times \mathbf{a}_2 = 0$ is the same as $\mathbf{p}_2 \sim \mathbf{e}_{21}$.

Case 2: Now suppose $\det(G_1)s_{13} \geq 0$. Observe that the depth of $\mathbf{c}_2$ in A_1 is

$$\operatorname{depth}(\mathbf{c}_2; A_1) = \frac{1}{|\det(G_1)| \|G^1_{3,\bullet}\|} \det(G_1)s_{13}. \tag{3.39}$$

Therefore, $\det(G_1)s_{13} \geq 0$ if and only if $\mathbf{c}_2$ has nonnegative depth in A_1. In this case, the inequality (3.34) imposes no constraints on $\mathbf{p}_2$ as claimed. □

Lemma 3.4.2. *Let $\mathcal{A} = \{A_1, \ldots, A_m\}$ be an arrangement of finite cameras. If the centers $\mathbf{c}_i$ are not collinear then $E_{\mathcal{A}} \cap C_{\mathcal{A}} = E^+_{\mathcal{A}}$. Otherwise, $E_{\mathcal{A}} \cap C_{\mathcal{A}} = E^+_{\mathcal{A}} \cup \{\mathbf{e}\}$.*

Proof. We first show that $E^+_{\mathcal{A}}$ is in $C_{\mathcal{A}}$. Since the inequalities defining $C_{\mathcal{A}}$ only depend on pairs of cameras, we can restrict to the case of every pair $\{A_k, A_\ell\}$. If none of the indices are equal to j the cameras A_k and A_ℓ see the center of camera A_j and the $C_{\{A_k, A_\ell\}}$ inequalities are satisfied if $\mathbf{c}_j \in D_{\mathcal{A}}$. If one of the indices is equal to j, we use the previous Lemma 3.7.3. We conclude that if $\mathbf{c}_j \in D_{\mathcal{A}}$, then

$$E_j \subseteq \bigcap_{k,\ell} C_{\{A_k, A_\ell\}} = C_{\mathcal{A}}.$$

Conversely, we argue that $\mathbf{p} \in E_{\mathcal{A}} \setminus E^+_{\mathcal{A}}$ cannot lie in $C_{\mathcal{A}}$. If $\mathbf{p} \in E_{\mathcal{A}} \setminus E^+_{\mathcal{A}}$, then $\mathbf{p} = (\mathbf{e}_{1j}, \ldots, \mathbf{p}_j, \ldots, \mathbf{e}_{mj})$ where $\mathbf{c}_j$ has negative depth in some camera $A_k \in \mathcal{A}$. From the definition of depth, this means

$$\mathbf{n}_k^\top \mathbf{c}_j = \det(G_k)\lambda_3 p_{k3} c_{j4} < 0.$$

If the camera centers are not collinear, we can choose a camera A_ℓ with $\ell \neq j$ such that $\mathbf{c}_j$, $\mathbf{c}_k$, and $\mathbf{c}_\ell$ do not lie on a line. By Lemma 3.7.2,

$$\operatorname{sgn}\left(\det(G_k)p_{k3}(\mathbf{a}_k \times \mathbf{a}_\ell)^\top(\mathbf{b}_{k\ell} \times \mathbf{a}_\ell)\right) = \operatorname{sgn}(\det(G_k)p_{k3}\lambda_k c_{j4}) < 0, \tag{3.40}$$

violating one of the inequalities of $C_{\mathcal{A}}$. Hence, $\mathbf{p} \notin C_{\mathcal{A}}$.

If the camera centers are collinear, then the point of epipoles $(\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_m)$ is the image of the line connecting the centers. This point trivially lies in $C_{\mathcal{A}}$ because all defining inequalities evaluate to 0 on this point. Again, let A_k be a camera such that $\mathbf{c}_j$ has negative depth in A_k. Lemma 3.7.3 shows that $(\mathbf{e}_i, \mathbf{e}_k)$ is the only point in E_i that lies in $C_{\{A_i, A_k\}}$. $\square$

Lemma 3.4.3. *Let $\mathcal{A} = \{A_1, A_2, \ldots, A_m\}$ be an arrangement of finite cameras with distinct centers. Let $E_{\mathcal{A}}^{++}$ be the union of the sets E_j such that $\mathbf{c}_j$ has positive depth in every camera $A_i \in \mathcal{A} \setminus \{A_j\}$, then $E_{\mathcal{A}}^{++} \subseteq \overline{\mathcal{X}}_{\mathcal{A}}$.*

Proof. We can approach $\mathbf{p} = (\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m) \in E_j$ by the sequence of points $\varphi_{\mathcal{A}}(\mathbf{v}(s))$ as s goes to 0, where

$$\mathbf{v}(s) = \begin{pmatrix} sG_j^{-1}\mathbf{p}_j \\ 0 \end{pmatrix} + \mathbf{c}_j.$$

Indeed, $A_i\mathbf{v}(s) = sG_iG_j^{-1}\mathbf{p}_j + \mathbf{e}_{ij}$ which approaches $\mathbf{e}_{ij}$ as $s \to 0$, and $A_j\mathbf{v}(s) = s\mathbf{p}_j \sim \mathbf{p}_j$ for all $s \neq 0$. Since depth is continuous and $\mathbf{c}_j$ has positive depth in A_i, $i \neq j$, the point $\mathbf{v}(s)$ has positive depth in A_i for sufficiently small $s \in \mathbb{R}$. The depth of $\mathbf{v}(s)$ with respect to camera A_j changes sign at $s = 0$ (or it is identically 0, that is $\mathbf{v}(s)$ lies on the principal plane of camera A_j for all s). Therefore, $\mathbf{v}(s)$ is in $D_{\mathcal{A}}$ for sufficiently small positive or negative s. So $\mathbf{p}$ lies in the closure of $\mathcal{X}_{\mathcal{A}}$. $\square$

Proofs from Section 3.5.2

Applying techniques from Section 3.5.1, we show when a Euclidean reconstruction can be made chiral using a homography. As we argue in Section 3.5.2, we may assume that our starting and target reconstructions have $A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}$. This choice of the first cameras restricts the homographies we need to consider to H such that $H^{-1} = \begin{bmatrix} I & 0 \\ \mathbf{v}^\top & \delta \end{bmatrix}$ for some $\mathbf{v} \in \mathbb{R}^3$ and nonzero $\delta \in \mathbb{R}$. Note that $\delta = \det H^{-1}$.

We now introduce the notion of a quasi-Euclidean camera.

Definition 3.7.1. *A camera $A = \begin{bmatrix} U & \mathbf{t} \end{bmatrix}$ is* quasi-Euclidean *if $UU^\top = I$.*

While we are interested in transforming a Euclidean reconstruction into a chiral Euclidean reconstruction, a homography may only be able to yield a reconstruction where the transformed cameras are quasi-Euclidean. However, since scaling a camera does not change chirality, a chiral quasi-Euclidean reconstruction can be turned into a chiral Euclidean reconstruction by multiplying A_i by $\operatorname{sgn}(\det(U_i))$. As a result, we only need to search for a homography H that sends our starting Euclidean reconstruction to one where every camera is quasi-Euclidean, which bring us to the following lemma.

Lemma 3.7.4. *Given a Euclidean camera $A = \begin{bmatrix} R & \mathbf{t} \end{bmatrix}$ such that $\mathbf{t} \neq \mathbf{0}$ and a homography H such that $H^{-1} = \begin{bmatrix} I & 0 \\ \mathbf{v}^\top & \delta \end{bmatrix}$ for some vector $\mathbf{v} \in \mathbb{R}^3$ and $\delta \neq 0$, the camera AH^{-1} is quasi-Euclidean if and only if $\mathbf{v} = \mathbf{0}$ or $\mathbf{v} = -\frac{2}{\|\mathbf{t}\|^2} R^\top \mathbf{t}$.*

Proof. The requirement that AH^{-1} be quasi-Euclidean translates to

$$\begin{aligned} I &= (R + \mathbf{t}\mathbf{v}^\top)^\top (R + \mathbf{t}\mathbf{v}^\top) \\ &= R^\top R + \mathbf{v}\mathbf{t}^\top R + R^\top \mathbf{t}\mathbf{v}^\top + \mathbf{v}\mathbf{t}^\top \mathbf{t}\mathbf{v}^\top \\ &= I + \mathbf{v}\mathbf{t}^\top R + R^\top \mathbf{t}\mathbf{v}^\top + \|\mathbf{t}\|^2 \mathbf{v}\mathbf{v}^\top \end{aligned}$$

For the fixed vector $\widetilde{\mathbf{c}} := -R^\top \mathbf{t} \neq \mathbf{0}$, this system is equivalent to finding $\mathbf{v}$ such that $M := -\mathbf{v}\widetilde{\mathbf{c}}^\top - \widetilde{\mathbf{c}}\mathbf{v}^\top + (\mathbf{v}\widetilde{\mathbf{c}}^\top)(\widetilde{\mathbf{c}}\mathbf{v}^\top) = 0$. Certainly $\mathbf{v} = \mathbf{0}$ is one solution. Otherwise, applying M to $\mathbf{v}$, we get that

$$\mathbf{0} = M\mathbf{v} = -(\widetilde{\mathbf{c}}^\top \mathbf{v})\mathbf{v} - (\mathbf{v}^\top \mathbf{v})\widetilde{\mathbf{c}} + (\widetilde{\mathbf{c}}^\top \widetilde{\mathbf{c}})(\mathbf{v}^\top \mathbf{v})\mathbf{v} \tag{3.41}$$

$$= ((\widetilde{\mathbf{c}}^\top \widetilde{\mathbf{c}})(\mathbf{v}^\top \mathbf{v}) - (\widetilde{\mathbf{c}}^\top \mathbf{v}))\mathbf{v} - (\mathbf{v}^\top \mathbf{v})\widetilde{\mathbf{c}}. \tag{3.42}$$

If $(\widetilde{\mathbf{c}}^\top\widetilde{\mathbf{c}})(\mathbf{v}^\top\mathbf{v}) - (\widetilde{\mathbf{c}}^\top\mathbf{v}) = \mathbf{0}$ for some $\mathbf{v} \neq 0$, then $M\mathbf{v} = (\mathbf{v}^\top\mathbf{v})\widetilde{\mathbf{c}} \neq \mathbf{0}$. Therefore, Equation (3.42) implies that $\mathbf{v} = \lambda\widetilde{\mathbf{c}}$ for some $\lambda \neq 0$. Solving for λ, we get $\lambda = \frac{2}{\widetilde{\mathbf{c}}^\top\widetilde{\mathbf{c}}}$ Which gives us the only additional solution $\mathbf{v} = \frac{2}{\|\widetilde{\mathbf{c}}\|^2}\widetilde{\mathbf{c}} = -\frac{2}{\|\mathbf{t}\|^2}R^\top\mathbf{t}$. □

Without loss of generality, we may assume the homographies in Lemma 3.7.4 have $|\delta| = 1$, leaving us with the following four possibilities for two view Euclidean reconstructions:

$$H_1^{-1} := \begin{bmatrix} I & \mathbf{0} \\ \mathbf{0}^\top & 1 \end{bmatrix}, \; H_2^{-1} := \begin{bmatrix} I & \mathbf{0} \\ \mathbf{0}^\top & -1 \end{bmatrix}, \tag{3.43}$$

$$H_3^{-1} := \begin{bmatrix} I & \mathbf{0} \\ \mathbf{v}^\top & 1 \end{bmatrix}, \; H_4^{-1} := \begin{bmatrix} I & \mathbf{0} \\ \mathbf{v}^\top & -1 \end{bmatrix} \tag{3.44}$$

where $\mathbf{v} = -\frac{2}{\|\mathbf{t}\|^2}R^\top\mathbf{t}$. These have the following inverses.

$$H_1 = \begin{bmatrix} I & \mathbf{0} \\ \mathbf{0}^\top & 1 \end{bmatrix}, \quad H_2 = \begin{bmatrix} I & \mathbf{0} \\ \mathbf{0}^\top & -1 \end{bmatrix}, \tag{3.45}$$

$$H_3 = \begin{bmatrix} I & \mathbf{0} \\ -\mathbf{v}^\top & 1 \end{bmatrix}, \quad H_4 = \begin{bmatrix} I & \mathbf{0} \\ \mathbf{v}^\top & -1 \end{bmatrix} \tag{3.46}$$

A Euclidean reconstruction $(\{A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}, A_2 = \begin{bmatrix} R & \mathbf{t} \end{bmatrix}\}, \mathcal{Q})$ can be made chiral if and only if one of $(\mathcal{A}H_i^{-1}, H_i\mathcal{Q})$ is chiral. Just as in the projective case, we assume we start with a signed reconstruction. Let $\mathbf{h}_i$ be the last row of H_i. From Theorem 3.5.1, we know we need only check if one of $\mathbf{h}_i$ lies in the cone intersection $K_{\mathcal{Q}}^* \cap (\operatorname{int} K_{\sigma C}^* \cup \operatorname{int} K_{-\sigma C}^*)$. As the following lemma shows, the special structure of $\mathbf{h}_i$ causes the cone conditions to simplify.

Lemma 3.7.5. *Let* $(\{A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}, A_2 = \begin{bmatrix} R & \mathbf{t} \end{bmatrix}\}, \mathcal{Q})$ *be a signed Euclidean reconstruction of* $\mathcal{P}$ *such that* $\mathbf{t} \neq \mathbf{0}$.

1. *If* $\sigma_1 = \sigma_2$, *then* $\mathbf{h}_1, \mathbf{h}_2 \in \operatorname{int} K_{\sigma C}^* \cup \operatorname{int} K_{-\sigma C}^*$ *and* $\mathbf{h}_3, \mathbf{h}_4 \notin \operatorname{int} K_{\sigma C}^* \cup \operatorname{int} K_{-\sigma C}^*$.

2. *If* $\sigma_1 \neq \sigma_2$ *then* $\mathbf{h}_3, \mathbf{h}_4 \in \operatorname{int} K_{\sigma C}^* \cup \operatorname{int} K_{-\sigma C}^*$ *and* $\mathbf{h}_1, \mathbf{h}_2 \notin \operatorname{int} K_{\sigma C}^* \cup \operatorname{int} K_{-\sigma C}^*$.

Proof. We first compute $\mathbf{h}_i^\top \sigma_j \mathbf{c}_j$ for all i, j:

$$\mathbf{h}_1^\top \sigma_1 \mathbf{c}_1 = \sigma_1,\ \mathbf{h}_1^\top \sigma_2 \mathbf{c}_2 = \sigma_2 \tag{3.47}$$

$$\mathbf{h}_2^\top \sigma_1 \mathbf{c}_1 = -\sigma_1,\ \mathbf{h}_2^\top \sigma_2 \mathbf{c}_2 = -\sigma_2 \tag{3.48}$$

$$\mathbf{h}_3^\top \sigma_1 \mathbf{c}_1 = \sigma_1,\ \mathbf{h}_3^\top \sigma_2 \mathbf{c}_2 = (-\mathbf{v}^\top(-R^\top \mathbf{t}) + 1)\sigma_2 = -\sigma_2 \tag{3.49}$$

$$\mathbf{h}_4^\top \sigma_1 \mathbf{c}_1 = -\sigma_1,\ \mathbf{h}_4^\top \sigma_2 \mathbf{c}_2 = (\mathbf{v}^\top(-R^\top \mathbf{t}) - 1)\sigma_2 = \sigma_2 \tag{3.50}$$

The vectors $\mathbf{h}_1$ and $\mathbf{h}_2$ make the same sign inner product with $\sigma_1 \mathbf{c}_1$ and $\sigma_2 \mathbf{c}_2$ if and only if $\sigma_1 = \sigma_2$. Similarly the vectors $\mathbf{h}_3$ and $\mathbf{h}_4$ make the same sign inner product with $\sigma_1 \mathbf{c}_1$ and $\sigma_2 \mathbf{c}_2$ if and only if $\sigma_1 = -\sigma_2$. □

Theorem 3.5.4. *Let* $(\{A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}, A_2 = \begin{bmatrix} R & \mathbf{t} \end{bmatrix}\}, \mathcal{Q})$ *be a signed Euclidean reconstruction of* $\mathcal{P}$ *with distinct centers. There exists a chiral Euclidean reconstruction of* $\mathcal{P}$ *if and only if* $\mathbf{n}_\infty \in K^*_{\mathcal{Q}} \cup K^*_{-\mathcal{Q}}$ *or* $\mathbf{r} := \begin{bmatrix} -\frac{2}{\|\mathbf{t}\|^2} R^\top \mathbf{t} \\ 1 \end{bmatrix} \in K^*_{\mathcal{Q}} \cup K^*_{-\mathcal{Q}}$. *Equivalently, if exactly one of the following holds for all* $\mathbf{q}_i$*:*

$$q_{i4} \geq 0 \ \ \textit{or} \ \ q_{i4} \leq 0 \ \ \textit{or} \ \ \mathbf{r}^\top \mathbf{q}_i \geq 0 \ \ \textit{or} \ \ \mathbf{r}^\top \mathbf{q}_i \leq 0.$$

Proof. By Theorem 3.5.1, a chiral Euclidean reconstruction exists if and only if one of the $\mathbf{h}_i$ lies in the cone intersection $K^*_{\mathcal{Q}} \cap (\operatorname{int} K^*_{\sigma C} \cup \operatorname{int} K^*_{-\sigma C})$. By Lemma 3.7.5, if $\sigma_1 = \sigma_2$, it is necessary and sufficient that either $\mathbf{h}_1 = \mathbf{n}_\infty \in K^*_{\mathcal{Q}}$ or $\mathbf{h}_2 = -\mathbf{n}_\infty \in K^*_{\mathcal{Q}}$. On the other hand, if $\sigma_1 \neq \sigma_2$, it is necessary and sufficient that either $\mathbf{h}_3 = \begin{bmatrix} -\frac{2}{\|\mathbf{t}\|^2} R^\top \mathbf{t} \\ 1 \end{bmatrix} \in K^*_{\mathcal{Q}}$ or $\mathbf{h}_4 = -\begin{bmatrix} -\frac{2}{\|\mathbf{t}\|^2} R^\top \mathbf{t} \\ 1 \end{bmatrix} \in K^*_{\mathcal{Q}}$, proving the statement. □

Theorem 3.5.5. *Let* $(\mathcal{A}, \mathcal{Q})$ *be a signed Euclidean reconstruction of* $\mathcal{P}$ *with* $m > 2$ *cameras, distinct centers, and* $A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}$. *There exists a chiral Euclidean reconstruction of* $\mathcal{P}$ *if and only if* $\sigma_i = \sigma_j$ *for all* $1 \leq i < j \leq m$ *and either* $q_{i4} \geq 0$ *for all* i *or* $q_{i4} \leq 0$ *for all* i.

Proof. Since the cameras have distinct centers, the vectors $-\frac{2}{\|\mathbf{t}_i\|^2} R_i^\top \mathbf{t}_i$ will not coincide, so by Lemma 3.7.4, the only homographies we can consider are H_1 and H_2. As in Lemma 3.7.5,

$\mathbf{h}_1 = \mathbf{n}_\infty, \mathbf{h}_2 = -\mathbf{n}_\infty \in \operatorname{int} K^*_{\sigma C} \cup \operatorname{int} K^*_{-\sigma C}$ if and only if $\sigma_i = \sigma_j$ for all i, j. When this is the case, a chiral reconstruction exists if and only if $\mathbf{n}_\infty \in K^*_{\mathcal{Q}}$ or $-\mathbf{n}_\infty \in K^*_{\mathcal{Q}}$, proving the statement. $\square$

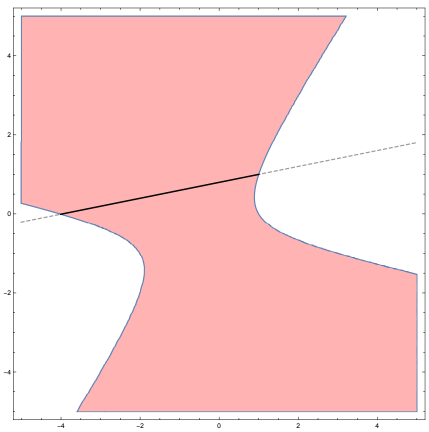

(a) The red region satisfies $(\mathbf{a}_1 \times \mathbf{a}_2)^\top(\mathbf{b}_{12} \times \mathbf{a}_2) \geq 0$.

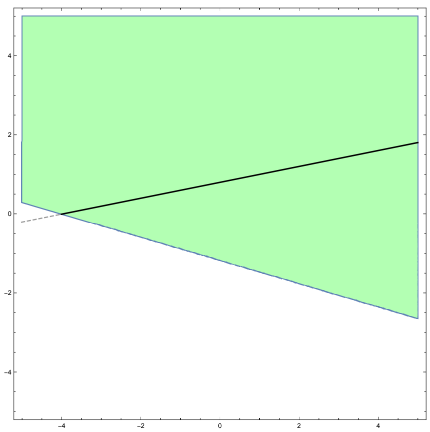

(b) The green region satisfies $(\mathbf{a}_1 \times \mathbf{a}_2)^\top(\mathbf{b}_{12} \times \mathbf{a}_1) \geq 0$.

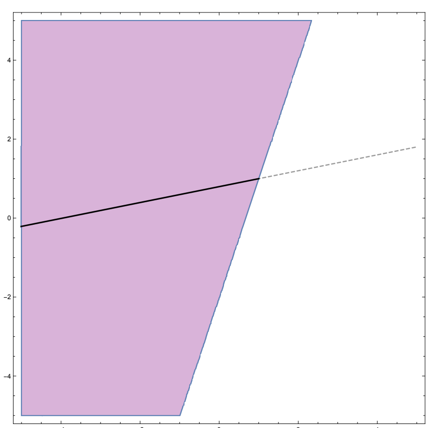

(c) The purple region satisfies $(\mathbf{b}_{12} \times \mathbf{a}_1)^\top(\mathbf{b}_{12} \times \mathbf{a}_2) \geq 0$.

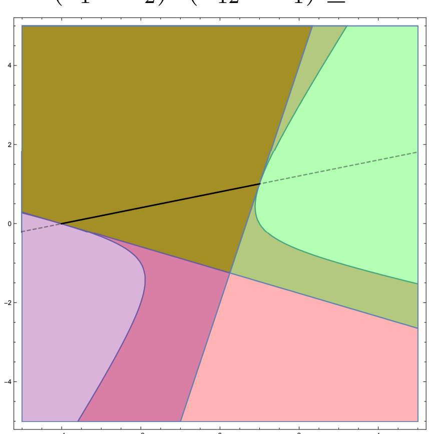

(d) The intersection of all regions satisfies all $C_\mathcal{A}$ inequalities.

Figure 3.6: An illustration of how the various inequalities of $C_\mathcal{A}$ cut out the chiral joint image $\mathcal{X}_\mathcal{A}$ in Example 3.4.1. In each figure, we fix $\mathbf{p}_1 = (-4, 0, 1)^\top$ and consider corresponding points in the second image. The dashed black line is in the epipolar line $\overline{\mathcal{J}}_\mathcal{A}$. The solid black line segments in figures (a), (b) and (c) is the intersection of the epipolar line with the (colored) region defined by one of the inequalities. The three inequalities are combined in figure (d) and their joint intersection with the epipolar line is shown in solid black. This is the closure of the chiral joint image $\overline{\mathcal{X}}_\mathcal{A}$ with $\mathbf{p}_1 = (-4, 0, 1)^\top$.

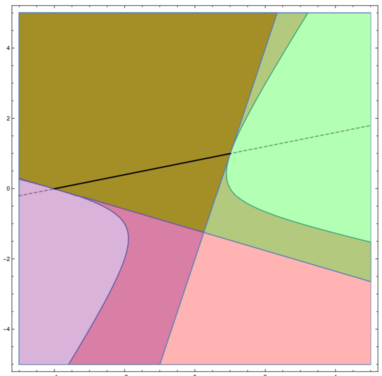

(a) $\mathbf{p}_1 = (-4, 0, 1)^\top$.

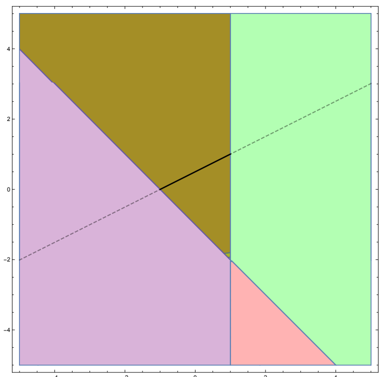

(b) $\mathbf{p}_1 = (-1, 0, 1)^\top$.

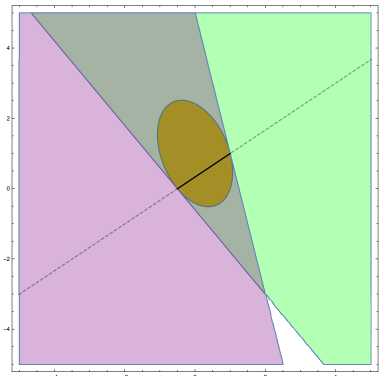

(c) $\mathbf{p}_1 = (-.5, 0, 1)^\top$.

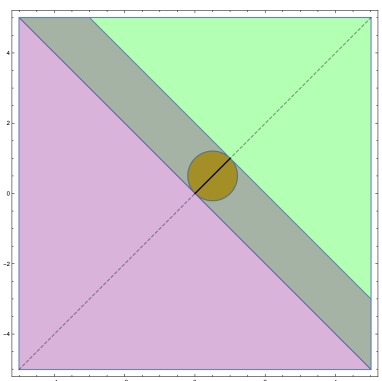

(d) $\mathbf{p}_1 = (0, 0, 1)^\top$.

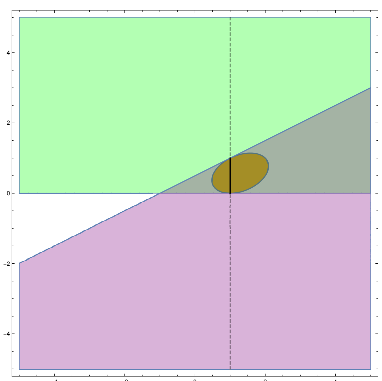

(e) $\mathbf{p}_1 = (1, 0, 1)^\top$.

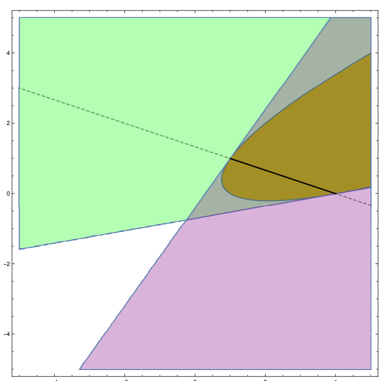

(f) $\mathbf{p}_1 = (4, 0, 1)^\top$.

Figure 3.7: The chiral joint image for Example 3.4.1 as $\mathbf{p}_1$ varies along a line. The dashed black line is the epipolar line in $(\overline{\mathcal{J}}_\mathcal{A})$. The solid black line segment is in the closure of the chiral joint image $(\overline{\mathcal{X}}_\mathcal{A})$. See Figure 3.6 for complete description of the regions.

Chapter 4

EXISTENCE OF TWO VIEW CHIRAL RECONSTRUCTIONS

4.1 Introduction

A fundamental question in computer vision is whether a set of point pairs $\mathcal{P} = \{(\mathbf{u}_i, \mathbf{v}_i) : i = 1, \ldots, k\}$ is the image of a set of world points $\mathbf{q}_i$ that are visible in two cameras. If we ignore the constraint (as is commonly done) that the points $\mathbf{q}_i$ need to lie in front of the cameras, we get a *projective reconstruction* [23]. In reality though, cameras can only image points in front of them. A reconstruction that obeys this additional constraint is known as a *chiral reconstruction* [2, 24]. The aim of this paper is to develop a theory of existence of two view chiral reconstructions.

Under the assumption that the points in each image are distinct, we prove the following:

1. A set of at most three point pairs always has a chiral reconstruction.

2. A set of four point pairs has a chiral reconstruction unless the configurations are of two specific non-generic types, in which case a chiral reconstruction may not exist.

3. Five or more point pairs can fail to have a chiral reconstruction with positive probability (in particular even if they are in general position).

We rely on complex and real algebraic geometry to prove these results. For five sufficiently generic point pairs, the problem translates to finding points in semialgebraic regions on a real cubic surface whose boundaries are segments of real lines in a Schläfli double six of real lines.

Formal definitions of projective and chiral reconstructions can be found in Section 4.3. In Section 4.4, we introduce the inequalities imposed by chirality and develop tools to certify them. Along the way we prove that a set of at most three point pairs always has a chiral reconstruction. In Section 4.5, we prove that $\mathcal{P}$ has a chiral reconstruction when the point configurations in each view have sufficiently similar geometry, and in particular, when they are in general position. The bad cases fall into two non-generic combinatorial types. In particular, the probability of choosing four point pairs that fail to have a chiral reconstruction is zero.

In Section 4.6 we show that when $k > 4$, point pairs that are in general position may not have a chiral reconstruction. Specific examples of this type when $k = 5$ were known to Werner [52] and there are close connections between our work and that of Werner's [51, 52, 53, 54]. We make two new contributions. In Section 4.6 we show that one can decide the existence of a chiral reconstruction when $k = 5$ and $\mathcal{P}$ is generic, by checking 20 discrete points. We use this test to show that the set of five point pairs that do not admit a chiral reconstruction is Zariski-dense. In other words, five point pairs do not have a chiral reconstruction with positive probability. A set of six or more point pairs can have a chiral reconstruction only if any subset of five point pairs among them have one. Hence, for any value of $k > 5$, there will be point configurations without a chiral reconstruction. Our second contribution in Section 4.7 is to show that the case of $k = 5$ is intimately related to the theory of cubic surfaces from classical algebraic geometry. Indeed, $\mathcal{P}$ creates a Schläfli double six of 12 lines on a cubic surface, all of whose 27 lines are real. These lines determine the boundary of the semialgebraic regions corresponding to chiral reconstructions.

4.2 Background and Notation

We now introduce some background and notation that will be needed in the paper. Let $\mathbb{P}^n$ and $\mathbb{P}^n_{\mathbb{R}}$ denote n-dimensional projective space over complex and real numbers respectively. We write $\mathbf{a} \sim \mathbf{b}$ if $\mathbf{a}$ and $\mathbf{b}$ are the same points in projective space, and reserve $\mathbf{a} = \mathbf{b}$ to mean coodinate-wise equality. A *projective camera* is a matrix in $\mathbb{P}^{3\times 4}_{\mathbb{R}}$ of rank three. Usually, an affine representative of a projective camera is fixed and we block-partition such a matrix as $A = \begin{bmatrix} G & \mathbf{t} \end{bmatrix} \in \mathbb{R}^{3\times 4}$ where $G \in \mathbb{R}^{3\times 3}$ and $\mathbf{t} \in \mathbb{R}^3$. The *center* of A is the unique point

$\mathbf{c}_A \in \mathbb{P}^3_{\mathbb{R}}$ such that $A\mathbf{c}_A = 0$. The camera A is a rational map from the "world" $\mathbb{P}^3$ to the "camera plane" $\mathbb{P}^2$ sending a "world point" $\mathbf{q}$ to its "image" $A\mathbf{q}$. It is defined everywhere except at $\mathbf{c}_A$. Consider the hyperplane at infinity in $\mathbb{P}^3$, $L_\infty = \{\mathbf{q} \in \mathbb{P}^3 \,:\, \mathbf{n}_\infty^\top \mathbf{q} = 0\}$, as an oriented hyperplane in $\mathbb{R}^4$ with fixed normal $\mathbf{n}_\infty = (0,0,0,1)^\top$. The camera A is said to be *finite* if $\mathbf{c}_A$ is a finite point, i.e., $\mathbf{c}_A \notin L_\infty$. A special representative of a camera center can be obtained by Cramer's rule where the ith coordinate of $\mathbf{c}_A$ is the determinant of the submatrix of A obtained by dropping the ith column. In particular, for a finite camera $A = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}$, the Cramer's rule center is $\mathbf{c}_A = \det(G)(-G^{-1}\mathbf{t}, 1)^\top$, and hence A is finite if and only if $\det(G) \neq 0$. Throughout this paper we use the Cramer's rule representation of $\mathbf{c}_A$, and all cameras will be finite.

The *principal plane* of a finite camera $A = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}$ is the hyperplane $L_A := \{\mathbf{q} \in \mathbb{P}^3 \,:\, A_{3,\bullet}\mathbf{q} = 0\}$, where $A_{3,\bullet}$ is the third row of A, i.e. it is the set of points in $\mathbb{P}^3$ that image to infinite points in $\mathbb{P}^2$. Note that the camera center $\mathbf{c}_A$ lies on L_A. We regard L_A as an oriented hyperplane in $\mathbb{R}^4$ with normal vector $\mathbf{n}_A := \det(G)A_{3\bullet}^\top$, which we call the *principal ray* of A. The $\det(G)$ factor ensures that the normal vector of the principal plane does not flip sign under a scaling of A. The depth of a finite point $\mathbf{q}$ in a finite camera A is defined as (see [23])

$$\operatorname{depth}(\mathbf{q}; A) := \left(\frac{1}{|\det(G)| \|G_{3,\bullet}\|} \right) \frac{(\mathbf{n}_A^\top \mathbf{q})}{(\mathbf{n}_\infty^\top \mathbf{q})}. \tag{4.1}$$

Note that the sign of $\operatorname{depth}(\mathbf{q}; A)$ is unaffected by scaling $\mathbf{q}$ and A. The depth of a finite point $\mathbf{q} \in \mathbb{P}^3$ in a finite camera A defined in Equation (4.1) is zero if and only if $\mathbf{n}_A^\top \mathbf{q} = 0$ which happens if and only if $\mathbf{q}$ lies on the principal plane L_A. Otherwise, $\mathbf{n}_A^\top \mathbf{q} \neq 0$ and $\operatorname{sgn}(\operatorname{depth}(\mathbf{q}; A)) = \operatorname{sgn}((\mathbf{n}_A^\top \mathbf{q})(\mathbf{n}_\infty^\top \mathbf{q}))$ is either positive or negative. It is then natural to say that a finite point $\mathbf{q}$ is *in front of* the camera A if $\operatorname{depth}(\mathbf{q}; A) > 0$, see [24]. Since only the sign of $\operatorname{depth}(\mathbf{q}; A)$ matters, we refer to this sign as the *chirality* of $\mathbf{q}$ in A, denoted as $\chi(\mathbf{q}; A)$, which is either 1 or -1.

The above notion of chirality was introduced by Hartley in the seminal paper [24], where he was concerned with a pair of cameras, see also [23, Chapter 21]. In [2], the definition of chirality was extended to to all points in $\mathbb{P}^3$, finite and infinite, and defined for an arrangement of cameras. Here is the two camera version we need.

Definition 4.2.1. *Let (A_1, A_2) be a pair of finite projective cameras. Then the* chiral domain *of (A_1, A_2), is the closure in $\mathbb{P}^3$ of the set*

$$\{\mathbf{q} \in \mathbb{P}^3 \mid \mathbf{q} \text{ finite}, \ \text{depth}(\mathbf{q}, A_1) > 0, \ \text{depth}(\mathbf{q}, A_2) > 0\}.$$

A point $\mathbf{q} \in \mathbb{P}^3$ is said to have chirality 1 with respect to (A_1, A_2), *denoted as $\chi(\mathbf{q}; (A_1, A_2)) = 1$, if and only if $\mathbf{q}$ lies in the chiral domain of (A_1, A_2).*

In this paper we will be concerned with a pair of finite *non-coincident* cameras (A_1, A_2) by which we mean that their centers are distinct. We will see that one can always take $A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}$, and then the conditions of finite and non-coincident imply that $A_2 = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}$ where $G \in \text{GL}_3$ and $\mathbf{t} \neq \mathbf{0}$. The pair (A_1, A_2) gives rise to the unique real, rank two *fundamental matrix* $X = [\mathbf{t}]_\times G$ where

$$[\mathbf{t}]_\times = \begin{bmatrix} 0 & -t_3 & t_2 \\ t_3 & 0 & -t_1 \\ -t_2 & t_1 & 0 \end{bmatrix}.$$

The skew-symmetric matrix $[\mathbf{t}]_\times$ represents the cross product with $\mathbf{t}$ as a linear map; that means that for $\mathbf{t}, \mathbf{r} \in \mathbb{R}^3$ we have $\mathbf{t} \times \mathbf{r} = [\mathbf{t}]_\times \mathbf{r} = [\mathbf{r}]_\times^\top \mathbf{t}$. Also, $\text{rank}([\mathbf{t}]_\times) = 2$ if and only if $\mathbf{t} \neq \mathbf{0}$.

The *epipoles* of the camera pair (A_1, A_2) are $\mathbf{e}_1, \mathbf{e}_2 \in \mathbb{P}^2$ such that $\mathbf{e}_1$ is the image of the center $\mathbf{c}_2$ in A_1, and $\mathbf{e}_2$ is the image of the center $\mathbf{c}_1$ in A_2. The line joining the centers $\mathbf{c}_1$ and $\mathbf{c}_2$ is called the *baseline* of the camera pair (A_1, A_2). Note that all world points on the baseline image to the epipole in each camera.

4.3 Projective and Chiral Reconstructions

Throughout this paper our input is a collection of point pairs $\mathcal{P} = \{(\mathbf{u}_i, \mathbf{v}_i) : i = 1, \ldots, k\} \subset \mathbb{P}^2_\mathbb{R} \times \mathbb{P}^2_\mathbb{R}$. Each $\mathbf{u}_i$ (and $\mathbf{v}_i$) is the homogenization of a point in $\mathbb{R}^2$ by adding a last coordinate one. Hence $(\mathbf{u}_i, \mathbf{v}_i)$ is a pair of finite points in $\mathbb{P}^2_{\mathbb{R}} \times \mathbb{P}^2_{\mathbb{R}}$ with a fixed representation. We will also assume that all $\mathbf{u}_i$ (and all $\mathbf{v}_i$) are distinct.

In this section we formally define projective and chiral reconstructions of $\mathcal{P}$, and characterize their existence. We then set up the geometric framework within which these reconstructions will be studied in this paper.

4.3.1 Projective reconstructions

Definition 4.3.1. *A* projective reconstruction *of $\mathcal{P}$ consists of a pair of projective cameras A_1, $A_2 \in \mathbb{P}^{3\times 4}$, world points $\mathcal{Q} = \{\mathbf{q}_1, \ldots, \mathbf{q}_k\} \subset \mathbb{P}^3$ and non-zero scalars w_{1i}, w_{2i} such that $A_1\mathbf{q}_i = w_{1i}\mathbf{u}_i$ and $A_2\mathbf{q}_i = w_{2i}\mathbf{v}_i$ for $i = 1, \ldots, k$. If the cameras and world points are all finite, then $(A_1, A_2, \mathcal{Q})$ is called a* finite projective reconstruction *of $\mathcal{P}$.*

The basics of projective reconstructions can be found in [23, Chapters 9 & 10]. Theorem 3.1 in [32] proves that if $\mathcal{P}$ has a projective reconstruction then it also has a finite projective reconstruction with $A_1 = \begin{bmatrix} I & 0 \end{bmatrix}$.

We now recall the necessary and sufficient conditions for the existence of a finite projective reconstruction of $\mathcal{P}$. For a point $\mathbf{e} \in \mathbb{P}^2$, let $\mathbf{e}(\mathbf{u}_1, \ldots, \mathbf{u}_k)$ denote the pencil of lines joining $\mathbf{e}$ to each $\mathbf{u}_i$. The following geometric characterization is well-known [23, 37, 52, 53].

Theorem 4.3.1. *[23, Section 9.4], [37, Section 2.4] The set of point pairs $\mathcal{P}$ has a projective reconstruction $(A_1, A_2, \mathcal{Q})$ if and only if there exist points $\mathbf{e}_1, \mathbf{e}_2 \in \mathbb{P}^2$ and a homography which sends the pencil of lines $\mathbf{e}_1(\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_k)$ to the pencil of lines $\mathbf{e}_2(\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k)$.*

The points $\mathbf{e}_1, \mathbf{e}_2$ in the above theorem are the epipoles of the camera pair (A_1, A_2) in the reconstruction. We now give a second characterization of the existence of a projective reconstruction.

Definition 4.3.2. *1. A* fundamental matrix *of $\mathcal{P}$ is a rank two matrix $X \in \mathbb{P}^{3\times 3}_{\mathbb{R}}$ such that*

$$\mathbf{v}_i^\top X \mathbf{u}_i = 0, \quad \text{for } i = 1, \ldots, k. \tag{4.2}$$

The linear equations (4.2) *in X are called the* epipolar equations *of $\mathcal{P}$.*

2. *Given a rank two matrix $X \in \mathbb{P}^{3\times 3}$ and a pair $(\mathbf{u}, \mathbf{v}) \in \mathbb{P}^2 \times \mathbb{P}^2$, such that $\mathbf{v}^\top X \mathbf{u} = 0$, we say that X is* $(\mathbf{u}, \mathbf{v})$-regular *if $\mathbf{v}^\top X = 0$ if and only if $X\mathbf{u} = 0$. i.e., $\mathbf{u}$ and $\mathbf{v}$ simultaneously generate the right and left kernels of X, or neither generate a kernel.*

3. *A fundamental matrix of $\mathcal{P}$ is* $\mathcal{P}$-regular *if it is $(\mathbf{u}_i, \mathbf{v}_i)$-regular for each point pair in $\mathcal{P}$.*

It is commonly believed that $\mathcal{P}$ has a projective reconstruction if and only if it has a fundamental matrix. However, a bit more care is needed as in the following theorem.

Theorem 4.3.2. *[32, Theorem 4.6] There exists a finite projective reconstruction of $\mathcal{P}$ with two non-coincident cameras, and $A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}$, if and only if there exists a $\mathcal{P}$-regular fundamental matrix.*

Remark 4.3.1. 1. *Let $A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}$ and $A_2 = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}$ be the two finite non-coincident cameras in the projective reconstruction. Then recall that $G \in \mathrm{GL}_3$, $\mathbf{t} \neq \mathbf{0}$, and the cameras correspond to the unique fundamental matrix $X = [\mathbf{t}]_\times G$ up to scale. The epipoles of the cameras are $\mathbf{e}_1 \sim G^{-1}\mathbf{t}$ and $\mathbf{e}_2 \sim \mathbf{t}$ which generate the right and left kernels of X. Further, G defines a $\mathbb{P}^2$ homography that sends $\mathbf{e}_1$ to $\mathbf{e}_2$, and the line joining $\mathbf{e}_1$ to $\mathbf{u}_i$ to the line joining $\mathbf{e}_2$ to $\mathbf{v}_i$. Hence G encodes the epipolar line homography of Theorem 4.3.1.*

2. *Conversely, any (rank two) fundamental matrix X of $\mathcal{P}$ can be factored as $X = [\mathbf{t}]_\times G$ for some $\mathbf{t} \in \mathbb{R}^3$ and $G \in \mathrm{GL}_3$ and yield a pair of cameras $(A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}, A_2 = \begin{bmatrix} G & \mathbf{t} \end{bmatrix})$ whose epipoles generate the left and right kernels of X. See [23, Section 9.5] for more details on the correspondence between fundamental matrices and camera pairs. These cameras can then be used to reconstruct a set of world points Q if X is $\mathcal{P}$-regular. The resulting projective reconstruction is said to be* associated to *X.*

3. *A fundamental matrix is $\mathcal{P}$-regular if and only if for each i, either $\mathbf{u}_i$ and $\mathbf{v}_i$ are both epipoles of the cameras or neither are. Indeed, this is a necessary condition for a projective reconstruction since if $\mathbf{u}_i \sim \mathbf{e}_1$, then its reconstruction $\mathbf{q}_i \in \mathbb{P}^3$ lies on the baseline of the cameras and hence images to $\mathbf{e}_2$ in camera A_2 requiring $\mathbf{v}_i \sim \mathbf{e}_2$. This subtlety is often overlooked, and it is common to equate the existence of a projective reconstruction of $\mathcal{P}$ to the existence of a fundamental matrix of $\mathcal{P}$.*

4. *Lastly, we remark that [32, Theorem 4.6] is stated using a different notion of regularity. However, a fundamental matrix X is $\mathcal{P}$-regular in our sense if and only if it is $([\mathbf{t}]_\times X, \mathbf{t})$-regular in the sense of [32] and hence the above theorem is exactly [32, Theorem 4.6].*

We now discuss the geometry encoded in Theorem 4.3.2 which will set the foundation for the work in this paper. Even though fundamental matrices are real, we will work over $\mathbb{C}$ to allow for methods from complex algebraic geometry, and will specialize to $\mathbb{R}$ as needed. A matrix $X \in \mathbb{P}^{3\times 3}$ can be identified with a point in $\mathbb{P}^8$ by concatenating its rows. Under this identification we let $\mathcal{R}_2 \subset \mathbb{P}^8$ be the determinantal hypersurface of matrices in $\mathbb{P}^{3\times 3}$ of rank at most two, and $\mathcal{R}_1$ be its subvariety of rank one matrices. As projective subvarieties of $\mathbb{P}^8$, $\dim \mathcal{R}_2 = 7$ and degree $\mathcal{R}_2 = 3$ while, $\dim \mathcal{R}_1 = 4$ and degree $\mathcal{R}_1 = 6$. For a point pair $(\mathbf{u}, \mathbf{v}) \in \mathbb{P}^2 \times \mathbb{P}^2$, let $L_{(\mathbf{u},\mathbf{v})}$ denote the hyperplane in $\mathbb{P}^8$:

$$L_{(\mathbf{u},\mathbf{v})} = \{X \in \mathbb{P}^{3\times 3} \,:\, \mathbf{v}^\top X \mathbf{u} = \langle X, \mathbf{v}\mathbf{u}^\top \rangle := \mathrm{Tr}(X^\top \mathbf{v}\mathbf{u}^\top) = 0\}$$

where $\langle \cdot, \cdot \rangle$ denotes the Frobenius inner product on matrices. Let $\mathcal{L}_\mathcal{P} = \bigcap_{i=1}^k \{L_{(\mathbf{u}_i,\mathbf{v}_i)} : (\mathbf{u}_i, \mathbf{v}_i) \in \mathcal{P}\}$. Generically, $\mathcal{L}_\mathcal{P}$ is a linear space in $\mathbb{P}^8$ of codimension k.

Definition 4.3.3. *The variety $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$ in $\mathbb{P}^8$ is the* epipolar variety *of $\mathcal{P}$.*

Under sufficient genericity of $\mathcal{P}$, $\dim(\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}) = 7 - k$ and degree$(\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}) = 3$. Hence the epipolar variety is empty when $k \geq 8$, consists of three points when $k = 7$, and infinitely many points when $k < 7$.

For $(\mathbf{u}_i, \mathbf{v}_i) \in \mathcal{P}$, consider the following the five-dimensional linear spaces of $\mathbb{P}^8$ that are in fact in $\mathcal{R}_2$:

$$\widetilde{W}_{\mathbf{u}_i} := \{X \in \mathbb{P}^{3\times 3} \,:\, X\mathbf{u}_i = 0\} \quad \text{and} \quad \widetilde{W}^{\mathbf{v}_i} := \{X \in \mathbb{P}^{3\times 3} \,:\, {\mathbf{v}_i}^\top X = 0\}.$$

Their intersections with the epipolar variety are the linear spaces: $W_{\mathbf{u}_i} := \mathcal{L}_\mathcal{P} \cap \widetilde{W}_{\mathbf{u}_i}$ and $W^{\mathbf{v}_i} := \mathcal{L}_\mathcal{P} \cap \widetilde{W}^{\mathbf{v}_i}$, each of which generically has dimension $6 - k$ since $\widetilde{W}_{\mathbf{u}_i}, \widetilde{W}^{\mathbf{v}_i} \subset L_{(\mathbf{u}_i, \mathbf{v}_i)}$.

Definition 4.3.4. *The linear space $W_{\mathbf{u}_i}$ (resp. $W^{\mathbf{v}_j}$) will be called the $\mathbf{u}_i$ (resp. $\mathbf{v}_j$)* wall *and the intersection $W_{\mathbf{u}_i} \cap W^{\mathbf{v}_j}$ will be called the $(\mathbf{u}_i, \mathbf{v}_j)$* corner.

When $i \neq j$, a $(\mathbf{u}_i, \mathbf{v}_j)$ corner has dimension $5 - k$ generically since there are at most 5 independent equations among $\mathbf{v}_j^\top X = 0 = X\mathbf{u}_i$. Thus a wall has codimension one and a corner has codimension two in the epipolar variety, generically. For non-generic data $\mathcal{P}$, all the dimensions computed above may be larger.

The second condition in Theorem 4.3.2 can now be rephrased as the existence of a real rank two matrix X in the epipolar variety of $\mathcal{P}$ such that for each i, X is either in the $(\mathbf{u}_i, \mathbf{v}_i)$ corner or in the complement of $W_{\mathbf{u}_i} \cup W^{\mathbf{v}_i}$. Note that a rank two X can lie in at most one $(\mathbf{u}_i, \mathbf{v}_i)$ corner.

Going forward, we will work both in $\mathbb{P}^8$, the space of fundamental matrices, and in $\mathbb{P}^2 \times \mathbb{P}^2$, the space of epipoles. These spaces are related by the *adjoint map*,

$$\operatorname{adj} : \mathbb{P}^8 \dashrightarrow \mathbb{P}^8, \quad \operatorname{adj}(X) = \operatorname{cof}(X)^\top$$

where $\operatorname{cof}(X)$ is the cofactor matrix of X. If $X \in \mathcal{R}_2$ then $X \cdot \operatorname{adj}(X) = \operatorname{adj}(X) \cdot X = 0$ and thus, if $\operatorname{rank}(X) = 2$, then all non-zero rows (columns) of $\operatorname{adj}(X)$ are multiples of each other and generate the left (right) kernel of X. Since generators of the right and left kernels of a fundamental matrix represent epipoles, the adjoint map provides a convenient connection between epipole space and fundamental matrix space.

4.3.2 Chiral reconstructions

A physical constraint on a true reconstruction $(A_1, A_2, \mathcal{Q})$ is that the reconstructed world points in $\mathcal{Q}$ must lie in front of the cameras A_1 and A_2. Recall from the Introduction that this means we require $\mathcal{Q}$ to lie in the chiral domain of (A_1, A_2) or equivalently, $\chi(\mathbf{q}_i; (A_1, A_2)) = 1$ for all $\mathbf{q}_i \in \mathcal{Q}$. A full development of multiview chirality can be found in [2]. For this paper, we use the following result for two views from [2, Theorem 2], as a definition.

Definition 4.3.5. *A* chiral reconstruction *of $\mathcal{P}$ is a projective reconstruction $(A_1, A_2, \mathcal{Q})$ of $\mathcal{P}$ with finite non-coindent cameras such that for all i,*

$$(\mathbf{n}_\infty^\top \mathbf{q}_i)(\mathbf{n}_1^\top \mathbf{q}_i) \geq 0, \;\; (\mathbf{n}_\infty^\top \mathbf{q}_i)(\mathbf{n}_2^\top \mathbf{q}_i) \geq 0, \;\; \textit{and} \;\; (\mathbf{n}_1^\top \mathbf{q}_i)(\mathbf{n}_2^\top \mathbf{q}_i) \geq 0$$

where $\mathbf{n}_\infty = (0,0,0,1)^\top$ and $\mathbf{n}_i$ is the principal ray of A_i.

Recall that two projective reconstructions $(A_1, A_2, \mathcal{Q})$ and $(A_1', A_2', \mathcal{Q}')$ are *projectively equivalent* if they are related by a homography of $\mathbb{P}^3$, i.e., there is a $H \in \mathrm{GL}_4$ such that $A_i' = A_i H^{-1}$ and $\mathcal{Q}' = H\mathcal{Q} := \{H\mathbf{q}_i,\, i = 1, \ldots, k\}$. A projective reconstruction which is not chiral can sometimes be transformed into a chiral reconstruction by a homography [2, 24, 53]. We recall the conditions under which this is possible.

Theorem 4.3.3. *Consider a finite projective reconstruction of $\mathcal{P}$ with non-coincident cameras $A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}$, $A_2 = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}$, and world points $\mathcal{Q} = \{\mathbf{q}_1, \ldots, \mathbf{q}_k\} \subset \mathbb{P}^3$. Then the following are equivalent.*

1. *There exists a projectively equivalent chiral reconstruction $(A_1 H^{-1}, A_2 H^{-1}, H\mathcal{Q})$ of $\mathcal{P}$.*

2. *$(\mathbf{n}_1^\top \mathbf{q}_i)(\mathbf{n}_2^\top \mathbf{q}_i)$ has the same sign for all i.*

3. *$w_{1i} w_{2i}$ has the same sign for all i.*

Furthermore, if no $\mathbf{q}_i$ lies on the baseline of (A_1, A_2), then (1), (2), (3) are equivalent to

4. $(\mathbf{t} \times \mathbf{v}_i)^\top(\mathbf{t} \times G\mathbf{u}_i)$ *has the same sign for all* i*.*

The equivalence of (1) and (2) is Theorem 8 in [2]. The equivalence of (1) and (3) is Theorem 17 in [24]. The epipoles of the given cameras are $\mathbf{e}_1 \sim G^{-1}\mathbf{t}$ and $\mathbf{e}_2 \sim \mathbf{t}$, and hence if no world point lies on the baseline, $G^{-1}\mathbf{t} \not\sim \mathbf{u}_i$ (equivalently, $\mathbf{t} \not\sim G\mathbf{u}_i$) and $\mathbf{t} \not\sim \mathbf{v}_i$ for any i. Also, since A_1, A_2 are non-coincident, $\mathbf{t} \neq \mathbf{0}$, and from Theorem 4.3.1 $\mathbf{t}, \mathbf{v}_i, G\mathbf{u}_i$ are collinear. Therefore, $(\mathbf{t} \times \mathbf{v}_i) \sim (\mathbf{t} \times G\mathbf{u}_i)$ and so $(\mathbf{t} \times \mathbf{v}_i)^\top(\mathbf{t} \times G\mathbf{u}_i) \neq 0$ for all i. The equivalence of (3) and (4) can then be derived from the same arguments as in Lemma 3.7.2.

If a world point lies on the baseline, then its images $(\mathbf{u}_i, \mathbf{v}_i)$ in A_1, A_2 are the epipoles of the cameras, and the expression in (4) becomes zero. However, since any point on the baseline can serve as the world point $\mathbf{q}_i$ imaging to $(\mathbf{u}_i, \mathbf{v}_i)$, we can control the sign of $(\mathbf{n}_1^\top \mathbf{q}_i)(\mathbf{n}_2^\top \mathbf{q}_i)$ as shown next.

Lemma 4.3.1. *For a pair of non-coincident cameras* (A_1, A_2) *whose baseline is not contained in either principal plane* L_{A_1} *and* L_{A_2}*, there exist* $\mathbf{q}_+$ *and* $\mathbf{q}_-$ *on the baseline such that* $(\mathbf{n}_1^\top \mathbf{q}_+)(\mathbf{n}_2^\top \mathbf{q}_+) > 0$ *and* $(\mathbf{n}_1^\top \mathbf{q}_-)(\mathbf{n}_2^\top \mathbf{q}_-) < 0$*.*

Proof. Since $\mathbf{c}_1 \in L_{A_1}$, $\mathbf{n}_1^\top \mathbf{c}_1 = 0$. On the other hand, since the baseline is not contained in L_{A_2} and $\mathbf{c}_1 \neq \mathbf{c}_2$, $\mathbf{c}_1 \notin L_{A_2}$, and so $(\mathbf{n}_2^\top \mathbf{c}_1) \neq 0$. By continuity, there exist perturbations $\mathbf{q}_+$ and $\mathbf{q}_-$ of $\mathbf{c}_1$ on the baseline such that $(\mathbf{n}_1^\top \mathbf{q}_+)(\mathbf{n}_2^\top \mathbf{q}_+) > 0$ and $(\mathbf{n}_1^\top \mathbf{q}_-)(\mathbf{n}_2^\top \mathbf{q}_-) < 0$. ☐

Remark 4.3.2. *For reconstructions where both epipoles are finite, the hypothesis of Lemma 4.3.1 is satisfied. Indeed, if for instance the baseline was contained in the principal plane* L_{A_1} *then* $\mathbf{c}_2 \in L_{A_1}$ *and so* $\mathbf{e}_1 \sim A_1\mathbf{c}_2$ *would be an infinite point.*

We now have a necessary and sufficient condition for the existence of a chiral reconstruction.

Lemma 4.3.2. *There exists a chiral reconstruction of* $\mathcal{P}$ *if and only if there exist* $\mathbf{t} \in \mathbb{R}^3 \setminus \{\mathbf{0}\}$ *and* $G \in GL_3$ *such that* $[\mathbf{t}]_\times G$ *is a* $\mathcal{P}$*-regular fundamental matrix and*

$$\left((\mathbf{t} \times \mathbf{v}_i)^\top(\mathbf{t} \times G\mathbf{u}_i)\right)\left((\mathbf{t} \times \mathbf{v}_j)^\top(\mathbf{t} \times G\mathbf{u}_j)\right) \geq 0 \;\forall\; i, j. \tag{4.3}$$

Proof. Suppose $(A_1, A_2, \mathcal{Q})$ is a chiral reconstruction of $\mathcal{P}$ with non-coincident finite cameras where $A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}$. Then $A_2 = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}$ for some $\mathbf{t} \in \mathbb{R}^3 \setminus \{\mathbf{0}\}$ and $G \in GL_3$, and by Theorem 4.3.2, $[\mathbf{t}]_\times G$ is a $\mathcal{P}$-regular fundamental matrix associated to A_1 and A_2. For all i such that $\mathbf{q}_i$ is not on the baseline, $(\mathbf{t} \times \mathbf{v}_i)^\top (\mathbf{t} \times G\mathbf{u}_i)$ has the same sign by Theorem 4.3.3. If some world point $\mathbf{q}_i$ is on the baseline, then its image $(\mathbf{u}_i, \mathbf{v}_i)$ is the pair of epipoles $(G^{-1}\mathbf{t}, \mathbf{t})$, and hence $(\mathbf{t} \times \mathbf{v})^\top (\mathbf{t} \times G\mathbf{u}) = 0$. Therefore, the inequalities in (4.3) hold.

Conversely, suppose there exist $\mathbf{t} \in \mathbb{R}^3 \setminus \{\mathbf{0}\}$ and $G \in GL_3$ such that $[\mathbf{t}]_\times G$ is a $\mathcal{P}$-regular fundamental matrix and the inequalities (4.3) hold. By Theorem 4.3.2, there exist world points $\mathcal{Q}$ such that $(A_1 = \begin{bmatrix} I & \mathbf{0} \end{bmatrix}, A_2 = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}, \mathcal{Q})$ is a projective reconstruction of $\mathcal{P}$. Let $\widehat{\mathcal{Q}} \subseteq \mathcal{Q}$ be the set of world points not on the baseline of A_1 and A_2. Since the inequalities (4.3) hold, the quadruple products $(\mathbf{t} \times \mathbf{v}_i)^\top (\mathbf{t} \times G\mathbf{u}_i)$ have the same sign for all $\mathbf{q}_i \in \widehat{\mathcal{Q}}$. By Theorem 4.3.3, there exists a constant $\sigma \in \{-1, 1\}$ such that $\sigma = \operatorname{sgn}(\mathbf{n}_1^\top \mathbf{q}_i)(\mathbf{n}_2^\top \mathbf{q}_i)$ for all $\mathbf{q}_i \in \widehat{\mathcal{Q}}$.

If some point $\mathbf{q}_j \in \mathcal{Q}$ lies on the baseline, then $\mathbf{q}_j$ images to the pair of epipoles $G^{-1}\mathbf{t}$ and $\mathbf{t}$ in the two cameras and hence $(\mathbf{t} \times \mathbf{v}_j)^\top (\mathbf{t} \times G\mathbf{u}_j) = 0$. By Lemma 4.3.1, we may replace $\mathbf{q}_j$ by some world point $\mathbf{q}'_j$ on the baseline such that $\operatorname{sgn}(\mathbf{n}_1^\top \mathbf{q}'_j)(\mathbf{n}_2^\top \mathbf{q}'_j) = \sigma$. Let $\mathcal{Q}'$ be the modification of $\mathcal{Q}$ obtained by replacing all world points on the baseline as above, but keeping all other world points intact. By construction, $\operatorname{sgn}(\mathbf{n}_1^\top \mathbf{q}'_i)(\mathbf{n}_2^\top \mathbf{q}'_i) = \sigma$ for all $\mathbf{q}'_i \in \mathcal{Q}'$. The transformed reconstruction $(A_1, A_2, \mathcal{Q}')$ is projectively equivalent to a chiral reconstruction by Theorem 4.3.3. $\square$

Lemma 4.3.2 implies that for a chiral reconstruction to exist, there must be $\mathcal{P}$-regular fundamental matrices that satisfy the inequalities (4.3). In the next section, we examine these inequalities to understand the regions of the epipolar variety in which fundamental matrices that lead to chiral reconstructions live.

4.4 Chiral Tools

In this section we develop tools to prove the existence of chiral reconstructions. In Section 4.4.1, we describe the semialgebraic *chiral epipolar region* of fundamental matrices associated to chiral reconstructions of $\mathcal{P}$. In Section 4.4.2, we show how inequalities defining the chiral epipolar region can be checked in epipole space. Even if not stated explicitly, we are working over $\mathbb{P}_{\mathbb{R}}$ when dealing with inequalities. We combine these tools in Section 4.4.3 to prove that a set of three point pairs has a chiral reconstruction. In Section 4.4.4, we show how the walls and corners of the epipolar variety can be used to decide if a set of more than three point pairs has a chiral reconstruction.

4.4.1 The chiral epipolar region

By Lemma 4.3.2, a fundamental matrix must satisfy (4.3) to yield a chiral reconstruction. In this section, we describe the strict subset of the epipolar variety satisfying these constraints.

Definition 4.4.1. *Let $X \in \mathbb{P}^8_{\mathbb{R}}$. For each point pair $(\mathbf{u}_i, \mathbf{v}_i)$, the i-th* chiral polynomial *will be denoted*

$$g_i(X) := \mathbf{v}_i^\top [-\mathbf{t}]_\times X \mathbf{u}_i$$

where $\mathbf{t}^\top X = \mathbf{0}$. The set of all $g_i g_j(X) \geq 0$ are called the chiral epipolar inequalities of $\mathcal{P}$.

The non-zero rows of $\mathrm{adj}(X)$ are possible choices for $\mathbf{t}$, and hence $g_i(X)$ is a cubic polynomial in the entries of X. The rows of $\mathrm{adj}(X)$ define a map from $\mathcal{R}_2$ to $\mathbb{P}^2$. We can evaluate the inequality $g_i g_j(X) \geq 0$ locally on the affine charts given by the conditions that a certain row of $\mathrm{adj}(X)$ is not $\mathbf{0}$. We now show that the chiral polynomial $g_i(X)$ records the quadruple product $(\mathbf{t} \times \mathbf{v}_i)^\top(\mathbf{t} \times G\mathbf{u}_i)$ from Lemma 4.3.2.

Lemma 4.4.1. *If $X = [\mathbf{t}]_\times G$, then $g_i(X) = (\mathbf{t} \times \mathbf{v}_i)^\top(\mathbf{t} \times G\mathbf{u}_i)$ for each i.*

Proof.

$$g_i(X) = \mathbf{v}_i^\top[-\mathbf{t}]_\times X\mathbf{u}_i = ([-\mathbf{t}]_\times^\top \mathbf{v}_i)^\top [\mathbf{t}]_\times G\mathbf{u}_i = ([\mathbf{t}]_\times \mathbf{v}_i)^\top [\mathbf{t}]_\times G\mathbf{u}_i = (\mathbf{t} \times \mathbf{v}_i)^\top(\mathbf{t} \times G\mathbf{u}_i).$$

□

The next theorem, which is analogous to Theorem 3 in [53], now follows from Lemma 4.3.2 and Lemma 4.4.1.

Theorem 4.4.1. *There exists a chiral reconstruction of $\mathcal{P}$ if and only if there exists a $\mathcal{P}$-regular fundamental matrix X such that $g_i g_j(X) \geq 0$ for all $1 \leq i < j \leq k$.*

Definition 4.4.2. *The* chiral epipolar region of $\mathcal{P}$ *is the set of $\mathcal{P}$-regular fundamental matrices X such that $g_i g_j(X) \geq 0$ for all $1 \leq i < j \leq k$.*

The chiral epipolar region of $\mathcal{P}$ is contained in the semialgebraic subset of the real part of the epipolar variety $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ that is cut out by the chiral epipolar inequalities. It is not necessarily equal to this set because the chiral epipolar region additionally requires the fundamental matrices to be $\mathcal{P}$-regular. However, since $\mathcal{P}$-regularity only fails on a proper algebraic subset, if the chiral epipolar region has non-empty interior, the boundary of the interior is determined by the points where the chiral epipolar inequalities change sign, which we study next.

Lemma 4.4.2. *Let X be a fundamental matrix of $\mathcal{P}$. Then $g_i(X) = 0$ if and only if $X \in W_{\mathbf{u}_i}$ or $X \in W^{\mathbf{v}_i}$.*

Proof. Clearly, $g_i(X) = \mathbf{v}_i^\top [-\mathbf{t}]_\times X \mathbf{u}_i = 0$ if $X\mathbf{u}_i = 0$. If $\mathbf{v}_i^\top X = 0$, then $\mathbf{v}_i$ and $\mathbf{t}$ are collinear and therefore $\mathbf{v}_i[-\mathbf{t}]_\times = \mathbf{0}$, which implies $g_i(X) = 0$. For the other implication, we know that $\mathbf{v}_i^\top X \mathbf{u}_i = 0$ and $\mathbf{v}_i^\top [-\mathbf{t}]_\times X \mathbf{u}_i = 0$, where the three vectors $\mathbf{v}_i$, $\mathbf{u}_i$, and $\mathbf{t}$ are real and non-zero. We assume that $X\mathbf{u}_i \neq \mathbf{0}$ and show that $\mathbf{v}_i^\top X = \mathbf{0}$. We know $X\mathbf{u}_i$ is orthogonal to $\mathbf{v}_i$ and $-\mathbf{t} \times \mathbf{v}_i$. Therefore, it must be collinear with $\mathbf{v}_i \times (\mathbf{v}_i \times \mathbf{t})$, which is the same as $(\mathbf{v}_i^\top \mathbf{t})\mathbf{v}_i - (\mathbf{v}_i^\top \mathbf{v}_i)\mathbf{t}$. We also know that $\mathbf{t}^\top X = 0$, which implies that $\mathbf{t}$ is also orthogonal to $X\mathbf{u}_i$, hence also to $\mathbf{v}_i \times (\mathbf{v}_i \times \mathbf{t})$. The dot product $\mathbf{t}^\top \left((\mathbf{v}_i^\top \mathbf{t})\mathbf{v}_i - (\mathbf{v}_i^\top \mathbf{v}_i)\mathbf{t}\right) = 0$, i.e. $(\mathbf{v}_i^\top \mathbf{t})^2 = (\mathbf{v}_i^\top \mathbf{v}_i)(\mathbf{t}^\top \mathbf{t})$. The Cauchy-Schwarz inequality implies that $\mathbf{t}$ and $\mathbf{v}_i$ are collinear, which implies the claim $\mathbf{v}_i^\top X = \mathbf{0}$. □

The goal of the paper is to understand when the chiral epipolar region of $\mathcal{P}$ is non-empty, or equivalently, when $\mathcal{P}$ has a chiral reconstruction. When $k = 7$, generically $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ consists of three points and it is easy to check if the real points lie in the chiral epipolar region of $\mathcal{P}$. Therefore, our focus will be on values of $k < 7$.

4.4.2 Translating to epipole space

In this section, we show how we can check the validity of chiral epipolar inequalities in $\mathbb{P}^2_{\mathbb{R}} \times \mathbb{P}^2_{\mathbb{R}}$, the space of epipoles. Consider the $(1,1)$-homogeneous quadratic polynomial

$$D_{ij}(\mathbf{u}, \mathbf{v}) := \det \begin{bmatrix} \mathbf{u}_i & \mathbf{u}_j & \mathbf{u} \end{bmatrix} \det \begin{bmatrix} \mathbf{v}_i & \mathbf{v}_j & \mathbf{v} \end{bmatrix} \tag{4.4}$$

where $(\mathbf{u}, \mathbf{v}) \in \mathbb{P}^2_{\mathbb{R}} \times \mathbb{P}^2_{\mathbb{R}}$. Note that $D_{ij}(\mathbf{u}, \mathbf{v}) = 0$ if and only if either factor is zero which is if and only if $\mathbf{u}_i, \mathbf{u}_j, \mathbf{u}$ are collinear or $\mathbf{v}_i, \mathbf{v}_j, \mathbf{v}$ are collinear. Werner uses the quantities $D_{ij}(\mathbf{u}, \mathbf{v})$ to impose an orientation constraint on the epipolar line homography described in Theorem 4.3.1 [52]. We will show that $D_{ij}(\mathbf{u}, \mathbf{v})$ is closely related to the products of chiral polynomials $g_i g_j(X)$ where $\mathbf{u}$ and $\mathbf{v}$ generate the right and left kernels of a fundamental matrix X. We rely on the following well known identity [12, 51].

Lemma 4.4.3. *Suppose* $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3 \in \mathbb{R}^4$. *Let* $A = \begin{bmatrix} G & \mathbf{t} \end{bmatrix}$ *be a finite camera with Cramer's rule center* $\mathbf{c}_A = \det(G)(-G^{-1}\mathbf{t}, 1)$. *Then* $\det \begin{bmatrix} A\mathbf{q}_1 & A\mathbf{q}_2 & A\mathbf{q}_3 \end{bmatrix} = \det \begin{bmatrix} \mathbf{q}_1 & \mathbf{q}_2 & \mathbf{q}_3 & \mathbf{c}_A \end{bmatrix}$.

Lemma 4.4.4. *Consider a projective reconstruction* $(A_1, A_2, \mathcal{Q})$ *of* $\mathcal{P}$, *i.e.,* $A_1\mathbf{q}_i = w_{1i}\mathbf{u}_i$ *and* $A_2\mathbf{q}_i = w_{2i}\mathbf{v}_i$ *where* $w_{ij} \neq 0$. *Suppose* $D_{ij}(-A_1\mathbf{c}_2, A_2\mathbf{c}_1) \neq 0$ *where* $\mathbf{c}_i$ *is the Cramer's rule center of* A_i. *Then*

$$\operatorname{sgn} D_{ij}(-A_1\mathbf{c}_2, A_2\mathbf{c}_1) = \operatorname{sgn}(w_{1i}w_{2i})(w_{1j}w_{2j}).$$

Proof. Expand $D_{ij}(-A_1\mathbf{c}_2, A_2\mathbf{c}_1)$ as follows.

$$\begin{aligned}
D_{ij}(-A_1\mathbf{c}_2, A_2\mathbf{c}_1) &= -\det\begin{bmatrix}\mathbf{u}_i & \mathbf{u}_j & A_1\mathbf{c}_2\end{bmatrix}\det\begin{bmatrix}\mathbf{v}_i & \mathbf{v}_j & A_2\mathbf{c}_1\end{bmatrix} && (4.5)\\
&= -\det\begin{bmatrix}\frac{1}{w_{1i}}A_1\mathbf{q}_i & \frac{1}{w_{1j}}A_1\mathbf{q}_j & A_1\mathbf{c}_2\end{bmatrix}\det\begin{bmatrix}\frac{1}{w_{2i}}A_2\mathbf{q}_i & \frac{1}{w_{2j}}A_2\mathbf{q}_j & A_2\mathbf{c}_1\end{bmatrix} && (4.6)\\
&= -\det\begin{bmatrix}\frac{1}{w_{1i}}\mathbf{q}_i & \frac{1}{w_{1j}}\mathbf{q}_j & \mathbf{c}_2 & \mathbf{c}_1\end{bmatrix}\det\begin{bmatrix}\frac{1}{w_{2i}}\mathbf{q}_i & \frac{1}{w_{2j}}\mathbf{q}_j & \mathbf{c}_1 & \mathbf{c}_2\end{bmatrix} && (4.7)\\
&= \frac{1}{w_{1i}w_{1j}w_{2i}w_{2j}}\det\begin{bmatrix}\mathbf{q}_i & \mathbf{q}_j & \mathbf{c}_1 & \mathbf{c}_2\end{bmatrix}\det\begin{bmatrix}\mathbf{q}_i & \mathbf{q}_j & \mathbf{c}_1 & \mathbf{c}_2\end{bmatrix} && (4.8)\\
&= \frac{1}{w_{1i}w_{1j}w_{2i}w_{2j}}(\det\begin{bmatrix}\mathbf{q}_i & \mathbf{q}_j & \mathbf{c}_1 & \mathbf{c}_2\end{bmatrix})^2. && (4.9)
\end{aligned}$$

Equation (4.7) follows from Equation (4.6) by applying Lemma 4.4.3 to both determinants in the product. Since $D_{ij}(-A_1\mathbf{c}_2, A_2\mathbf{c}_1) \neq 0$ by assumption, we conclude that

$$\operatorname{sgn} D_{ij}(-A_1\mathbf{c}_2, A_2\mathbf{c}_1) = \operatorname{sgn}(w_{1i}w_{2i})(w_{1j}w_{2j}).$$

□

Lemma 4.4.5. *Let X be a fundamental matrix of $\mathcal{P}$. Suppose $D_{ij}(\operatorname{adj}(X)\mathbf{t}, \mathbf{t}) \neq 0$ where $\mathbf{t}^\top X = \mathbf{0}$. Then $\operatorname{sgn} D_{ij}(\operatorname{adj}(X)\mathbf{t}, \mathbf{t}) = \operatorname{sgn} g_i g_j(X)$.*

Proof. Write $X = [\mathbf{t}]_\times G$ for $\mathbf{t} \in \mathbb{R}^3 \setminus \{\mathbf{0}\}$ and some $G \in \mathrm{GL}_3$. We know that $\mathbf{t}$ and $\operatorname{adj}(X)\mathbf{t}$ generate the one-dimensional left and right kernels of X, respectively. Since $D_{ij}(\operatorname{adj}(X)\mathbf{t}, \mathbf{t}) \neq 0$, $\operatorname{adj}(X)\mathbf{t}$ is not collinear with $\mathbf{u}_i$ and $\mathbf{u}_j$ and $\mathbf{t}$ is not collinear with $\mathbf{v}_i$ and $\mathbf{v}_j$. In particular, this means that neither $\mathbf{u}_i$ nor $\mathbf{u}_j$ is the right kernel of X and neither $\mathbf{v}_i$ nor $\mathbf{v}_j$ is the left kernel of X. It follows that X is $(\mathbf{u}_i, \mathbf{v}_i)$ regular and $(\mathbf{u}_j, \mathbf{v}_j)$ regular and neither is the epipole pair. By Theorem 4.3.2, there exists a finite projective reconstruction $(A_1 = \begin{bmatrix} I & \mathbf{0}\end{bmatrix}, A_2 = \begin{bmatrix} G & \mathbf{t}\end{bmatrix}, \{\mathbf{q}_i, \mathbf{q}_j\})$ of $\{(\mathbf{u}_i, \mathbf{v}_i), (\mathbf{u}_j, \mathbf{v}_j)\}$ such that the world points $\mathbf{q}_i, \mathbf{q}_j$ are not on the baseline.

Lemma 4.4.4 implies that $\operatorname{sgn} D_{ij}(-A_1\mathbf{c}_2, A_2\mathbf{c}_1) = \operatorname{sgn}(w_{1i}w_{1j})(w_{2i}w_{2j})$. By Theorem 4.3.3, $(w_{1i}w_{1j})(w_{2i}w_{2j}) > 0$ if and only if $\left[(\mathbf{t}\times\mathbf{v}_i)^\top(\mathbf{t}\times G\mathbf{u}_i)\right]\left[(\mathbf{t}\times\mathbf{v}_j)^\top(\mathbf{t}\times G\mathbf{u}_j)\right] > 0$. Combining this fact with Lemma 4.4.1, it follows that $\operatorname{sgn} D_{ij}(-A_1\mathbf{c}_2, A_2\mathbf{c}_1) = \operatorname{sgn} g_i g_j(X)$. Finally note that $A_2\mathbf{c}_1 = \mathbf{t}$ and $-A_1\mathbf{c}_2 = \det(G)(G^{-1}\mathbf{t})$ which is a positive multiple of $\operatorname{adj}(X)\mathbf{t}$. Indeed,

$$\operatorname{adj}(X)\mathbf{t} = \operatorname{adj}([\mathbf{t}]_\times G)\mathbf{t} = \operatorname{adj}(G)\operatorname{adj}([\mathbf{t}]_\times))\mathbf{t} = \det(G)G^{-1}(\mathbf{t}\mathbf{t}^\top)\mathbf{t} = \|\mathbf{t}\|^2\det(G)(G^{-1}\mathbf{t}).$$

Substituting $\operatorname{adj}(X)\mathbf{t}$ for $-A_1\mathbf{c}_2$ and $\mathbf{t}$ for $A_2\mathbf{c}_1$, the result follows. □

Note that $\det\begin{bmatrix}\mathbf{u}_i & \mathbf{u}_j & A_1\mathbf{c}_2\end{bmatrix}$ can be zero without $\mathbf{u}_i$ or $\mathbf{u}_j$ being the epipole $A_1\mathbf{c}_2$. Indeed, by Lemma 4.4.3 this happens whenever $\mathbf{q}_i, \mathbf{q}_j, \mathbf{c}_1, \mathbf{c}_2$ are coplanar. On the other hand, Lemma 4.4.2 implies that g_i vanishes at X if and only if $\mathbf{u}_i$ or $\mathbf{v}_i$ is an epipole of X. Therefore, $D_{ij}(\operatorname{adj}(X)\mathbf{t}, \mathbf{t})$ may vanish even when $g_ig_j(X) \neq 0$.

Lemma 4.4.5 shows that knowing the specific generators of the kernels of X, i.e., $\mathbf{t}$ and $\operatorname{adj}(X)\mathbf{t}$, respectively, is enough to compute the sign of the chiral epipolar inequalities. When $D_{ij}(\operatorname{adj}(X)\mathbf{t}, \mathbf{t})$ does not vanish, we can use it to infer the validity of chiral epipolar inequalities via Lemma 4.4.5, and hence argue for the existence of a chiral reconstruction of $\mathcal{P}$. We now identify a situation where we can use *any* generators of the kernels of X in D_{ij}.

Definition 4.4.3. *Suppose X is a fundamental matrix of $\mathcal{P}$. Define $I(X)$ to be the set of indices i such that $g_i(X) \neq 0$, i.e., the index set of inactive chiral polynomials at X. Let $\mathcal{P}_{I(X)}$ be the subset of point pairs in $\mathcal{P}$ indexed by $I(X)$.*

Theorem 4.4.2. *Let X be a fundamental matrix of $\mathcal{P}$ where $\mathbf{e}_1$ and $\mathbf{e}_2$ generate the right and left kernels of X, respectively. Suppose $|I(X)| \geq 3$, and $D_{ij}(\mathbf{e}_1, \mathbf{e}_2) \neq 0$ for all $i, j \in I(X)$. Then there exists a chiral reconstruction of $\mathcal{P}_{I(X)}$ associated to X if and only if $D_{ij}(\mathbf{e}_1, \mathbf{e}_2)$ has the same sign for all $i, j \in I(X)$.*

Proof. Suppose there exists a chiral reconstruction of $\mathcal{P}_{I(X)}$ associated to X. Then by Theorem 4.4.1, $g_ig_j(X) \geq 0$ for all $i, j \in I(X)$. In fact, $g_ig_j(X) > 0$ for all $i, j \in I(X)$ since if $g_ig_j(X) = 0$ for some i, j while $D_{ij}(\operatorname{adj}(X)\mathbf{t}, \mathbf{t}) \neq 0$, we would contradict Lemma 4.4.5. Note that if $D_{ij}(\mathbf{e}_1, \mathbf{e}_2) \neq 0$ for some kernel representatives $\mathbf{e}_1, \mathbf{e}_2$, it remains non-zero for any other pair of kernel representatives. By Theorem 4.4.1, $D_{ij}(\operatorname{adj}(X)\mathbf{t}, \mathbf{t})$ has the same sign for all i, j, and since $\text{adj}(X)\mathbf{t}$ and $\mathbf{t}$ are representatives of the right and left kernels of X, the result follows.

Conversely, suppose $D_{ij}(\mathbf{e}_1, \mathbf{e}_2)$ has the same non-zero sign for all $i, j \in I(X)$ where $\mathbf{e}_1$ and $\mathbf{e}_2$ generate the right and left kernels of X, respectively. Then $\mathbf{e}_1 = \lambda\, \text{adj}(X)\mathbf{e}_2$ for some non-zero λ. By Theorem 4.4.1, we know $\text{sgn}\, D_{ij}(\mathbf{e}_1, \mathbf{e}_2) = \lambda\, \text{sgn}\, D_{ij}(\text{adj}(X)\mathbf{e}_2, \mathbf{e}_2) = \lambda\, \text{sgn}\, g_i g_j(X)$ for all i, j. This shows that $g_i g_j(X)$ has the same sign for all $i, j \in I(X)$. Since $|I(X)| \geq 3$, this common sign cannot be negative and hence $g_i g_j(X) > 0$ for all $i, j \in I(X)$. These strict inequalities also imply that X is $\mathcal{P}_{I(X)}$-regular. Then by Theorem 4.4.1 there is a chiral reconstruction of $\mathcal{P}_{I(X)}$ associated to X. $\square$

We remark that $D_{ij}(\mathbf{e}_1, \mathbf{e}_2)$ does not have a well-defined sign on $\mathbb{P}^2_{\mathbb{R}} \times \mathbb{P}^2_{\mathbb{R}}$ because it is linear in $\mathbf{e}_1$ and $\mathbf{e}_2$. To get an inequality description of chirality in epipole space, we can take pairwise products $D_{ij}(\mathbf{e}_1, \mathbf{e}_2)D_{ik}(\mathbf{e}_1, \mathbf{e}_2)$ which are quadratic in each $\mathbb{P}^2_{\mathbb{R}}$ factor. If $D_{ij}(\mathbf{e}_1, \mathbf{e}_2)D_{ik}(\mathbf{e}_1, \mathbf{e}_2) > 0$, then $g_j g_k(X) > 0$ for any fundamental matrix X with epipoles $\mathbf{e}_1$ and $\mathbf{e}_2$. However, since $D_{ij}(\mathbf{e}_1, \mathbf{e}_2)$ may vanish even when $g_i g_j(X)$ does not, we observe that $D_{ij}(\mathbf{e}_1, \mathbf{e}_2)D_{ik}(\mathbf{e}_1, \mathbf{e}_2) \geq 0$ for all triples i, j, k is not equivalent to $g_i g_j(X) \geq 0, g_i g_k(X) \geq 0$ and $g_i g_k(X) \geq 0$. Due to this subtlety, we primarily study chirality using $g_i g_j \geq 0$ in $\mathbb{P}^8$ as opposed to $D_{ij} D_{ik} \geq 0$ in $\mathbb{P}^2 \times \mathbb{P}^2$.

4.4.3 Three point pairs always have a chiral reconstruction

In this section, we apply the tools developed so far to show that there is always a chiral reconstruction when $|\mathcal{P}| = 3$, and hence also when $|\mathcal{P}| \leq 3$ since $\mathcal{P}$ can have a chiral reconstruction only if all its subsets have one. We begin with two technical lemmas.

Lemma 4.4.6. *Suppose $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$ are three non-collinear points in $\mathbb{R}^3$. Then for each of the eight elements in $\sigma \in \{+, -\}^3$, there is an $\mathbf{e} \in \mathbb{R}^3$ such that $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, \mathbf{e}$ are in general position (no three in a line) and*

$$\sigma = (\text{sgn}(\det[\mathbf{a}_1\, \mathbf{a}_2\, \mathbf{e}]), \text{sgn}(\det[\mathbf{a}_1\, \mathbf{a}_3\, \mathbf{e}]), \text{sgn}(\det[\mathbf{a}_2\, \mathbf{a}_3\, \mathbf{e}])).$$

Further, for each $\sigma \in \{+,-\}^3$, the corresponding choices of $\mathbf{e}$ *come from an open polyhedral cone in* $\mathbb{R}^3$*.*

Proof. The expression $\det[\mathbf{a}_i\,\mathbf{a}_j\,\mathbf{e}] = l_{ij}(\mathbf{e})$ is the linear form whose kernel is the span of $\mathbf{a}_i$ and $\mathbf{a}_j$. Since $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$ are non-collinear, the hyperplanes cut out by $l_{12}(\mathbf{e}), l_{13}(\mathbf{e}), l_{23}(\mathbf{e})$ divide $\mathbb{R}^3$ into eight regions, each of which is a polyhedral cone. The interiors of these cones correspond to the eight sign patterns σ. □

Lemma 4.4.7. *Suppose* $\mathbf{v}_l, \mathbf{v}_r, \mathbf{t}$ *are finite points in* $\mathbb{P}^2$ *normalized to have last coordinate 1 on a line* L*. If* $\mathbf{t} \in L \setminus [\mathbf{v}_l, \mathbf{v}_r]$*, then for all* $\mathbf{w}_1, \mathbf{w}_2 \in [\mathbf{v}_l, \mathbf{v}_r]$*,* $(\mathbf{t} \times \mathbf{w}_1)^\top(\mathbf{t} \times \mathbf{w}_2) > 0$*.*

Proof. If $\mathbf{t} \notin [\mathbf{v}_l, \mathbf{v}_r]$, then either $\mathbf{v}_l \in [\mathbf{t}, \mathbf{v}_r]$ or $\mathbf{v}_r \in [\mathbf{v}_l, \mathbf{t}]$. Suppose $\mathbf{v}_l \in [\mathbf{t}, \mathbf{v}_r]$. Since $\mathbf{w}_1, \mathbf{w}_2 \in [\mathbf{v}_l, \mathbf{v}_r]$ and $[\mathbf{v}_l, \mathbf{v}_r] \subseteq [\mathbf{t}, \mathbf{v}_r]$, we know $\mathbf{w}_1, \mathbf{w}_2 \in [\mathbf{t}, \mathbf{v}_r]$. Write $\mathbf{w}_1 = \lambda_1\mathbf{t} + \lambda_2\mathbf{v}_r$ and $\mathbf{w}_2 = \mu_1\mathbf{t} + \mu_2\mathbf{v}_r$ where $\lambda_i, \mu_j \geq 0$. Since $\mathbf{w}_i \neq \mathbf{t}$, $\lambda_2, \mu_2 > 0$. The result follows from direct computation using that $\mathbf{t} \neq \mathbf{v}_r$:

$$(\mathbf{t} \times \mathbf{w}_1)^\top(\mathbf{t} \times \mathbf{w}_2) = (\mathbf{t} \times (\lambda_1\mathbf{t} + \lambda_2\mathbf{v}_r))^\top(\mathbf{t} \times (\mu_1\mathbf{t} + \mu_2\mathbf{v}_r)) \tag{4.10}$$

$$= \lambda_2\mu_2(\mathbf{t} \times \mathbf{v}_r)^\top(\mathbf{t} \times \mathbf{v}_r) > 0. \tag{4.11}$$

Similar reasoning applies if $\mathbf{v}_r \in [\mathbf{v}_l, \mathbf{t}]$. □

Theorem 4.4.3. *If* $|\mathcal{P}| = 3$ *then* $\mathcal{P}$ *has a chiral reconstruction.*

Proof. We break the proof into two parts:

1. Suppose $\mathcal{U} = \{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3\}$ or $\mathcal{V} = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ is in general position, say $\mathcal{U}$ is non-collinear. Choose $\mathbf{e}_2$ not on the line spanned by $\mathbf{v}_i$ and $\mathbf{v}_j$ for any i, j, so that $\det\begin{bmatrix}\mathbf{v}_i & \mathbf{v}_j & \mathbf{e}_2\end{bmatrix} \neq 0$ for all i, j. By Lemma 4.4.6, there exists an $\mathbf{e}_1$ such that $D_{ij}(\mathbf{e}_1, \mathbf{e}_2)$ has the same non-zero sign for all i, j. Since $k = 3$ and $\mathbf{e}_1$ and $\mathbf{e}_2$ are chosen from open regions, $W_{\mathbf{e}_1} \cap W^{\mathbf{e}_2}$ contains at least one rank two matrix X. By construction, this X is a $\mathcal{P}$-regular fundamental matrix with epipoles $\mathbf{e}_1, \mathbf{e}_2$ and $I(X) = \{1, 2, 3\}$. By Theorem 4.4.2, X yields a chiral reconstruction of $\mathcal{P}$.

2. Suppose both $\mathcal{U}$ and $\mathcal{V}$ are collinear and the lines they span are $L_{\mathcal{U}}$ and $L^{\mathcal{V}}$. Let $\mathbf{u}_l, \mathbf{u}_r$ be the furthest left and right points on the $L_{\mathcal{U}}$ line, so that the third point lies strictly between $\mathbf{u}_l, \mathbf{u}_r$. Similarly let $\mathbf{v}_l, \mathbf{v}_r$ be the furthest left and right points on the $L^{\mathcal{V}}$ line. Let $\mathbf{t} \in L^{\mathcal{V}} \setminus [\mathbf{v}_l, \mathbf{v}_r]$ and choose $G \in \mathrm{GL}_3$ such that $G\mathbf{u}_l = \mathbf{v}_l$ and $G\mathbf{u}_r = \mathbf{v}_r$. Define $X = [\mathbf{t}]_\times G$. Since $\mathbf{t}, \mathbf{v}_i, G\mathbf{u}_i$ are collinear for all i, the ith epipolar equation is satisfied. Since the chosen epipoles for X do not coincide with any data points, X is a $\mathcal{P}$-regular fundamental matrix. By construction $G\mathbf{u}_i \in [\mathbf{v}_l, \mathbf{v}_r]$ for each i. Combining Lemma 4.4.1 and Lemma 4.4.7, it follows that $g_i(X) > 0$ for each i, and there is a chiral reconstruction of $\mathcal{P}$ associated to X by Theorem 4.4.1.

□

4.4.4 Walls and Corners

To understand the existence of chiral reconstructions when $|\mathcal{P}| \geq 4$, we need one more tool that we now develop. Recall that the chiral epipolar region of $\mathcal{P}$ is the set of $\mathcal{P}$-regular fundamental matrices that live in the semialgebraic subset of the real epipolar variety cut out by the chiral epipolar inequalities. Lemma 4.4.2 implies that the chiral epipolar region is bounded by the $W_{\mathbf{u}_i}$, $W^{\mathbf{v}_j}$ walls. The fundamental matrices on walls are generally $\mathcal{P}$-irregular and do not correspond to a reconstruction. However, we show that $\mathcal{P}$-irregular fundamental matrices that are smooth points of the epipolar variety and yield partial chiral reconstructions, can be perturbed to $\mathcal{P}$-regular fundamental matrices that yield chiral reconstructions of $\mathcal{P}$.

Lemma 4.4.8. *Suppose $\mathcal{R}_2 \cap \mathcal{L}_P$ is irreducible. If X is a smooth fundamental matrix that is $(\mathbf{u}_i, \mathbf{v}_i)$-irregular, then there is a tangent direction $\mathbf{d} \in T_X(\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}})$ such that the directional derivative $D_{\mathbf{d}} g_i(X) \neq 0$.*

Proof. Suppose X is a smooth fundamental matrix of $\mathcal{P}$. Smoothness implies that the tangent space at X to the epipolar variety has the same dimension as the variety. If X is $(\mathbf{u}_i, \mathbf{v}_i)$-irregular for some i, then X is in exactly one of $W_{\mathbf{u}_i}$ or $W^{\mathbf{v}_i}$. Since $\mathcal{R}_2 \cap \mathcal{L}_P$ is irreducible, each wall must be an embedded component of strictly smaller dimension. This means that the wall's tangent space is strictly contained in the tangent space of the epipolar variety at X. Therefore, we can choose a direction $\mathbf{d}$ tangent to the epipolar variety at X which is not tangent to the wall which contains X. Lemma 4.4.2 implies that g_i vanishes on the real part of the epipolar variety only on the walls. By construction $D_{\mathbf{d}} g_i(X) \neq 0$. □

The following lemma is needed for Theorem 4.4.4 below, but its proof might be best understood after Section 4.7.

Lemma 4.4.9. *Suppose $|\mathcal{P}| \leq 5$. If a wall $W_{\mathbf{u}_i}$ (or $W^{\mathbf{v}_i}$) contains a matrix of rank two, then a general point Y on the wall is a smooth point of $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$.*

Proof. Recall that a wall $W_{\mathbf{u}_i}$ has codimension at most one in $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ because $\mathbf{v}_i \mathbf{u}_i^\top$ is in $\mathcal{L}_{\mathcal{P}}^\perp$. If the wall is an irreducible component of $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$, almost all points on it are smooth points of $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$. Otherwise, the wall has codimension one in the epipolar variety. If this wall contains a smooth point, then so will its intersection with generic data planes $L_{(\mathbf{u}_j, \mathbf{v}_j)}$. Therefore, cutting with sufficiently many of these, we can assume that $\mathcal{L}_{\mathcal{P}}$ has dimension three, $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ is an irreducible cubic surface in $\mathbb{P}^3$, and $W_{\mathbf{u}_i}$ is a line on it. Suppose for contradiction that $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ is singular at every point in $W_{\mathbf{u}_i}$.

The cubic surfaces which are singular along a line have been classified, see e.g. [10, in particular Case E]. We show that $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ cannot be any of these types, essentially because it contains too many intersecting lines. Indeed, $W_{\mathbf{u}_j}$ intersects $W^{\mathbf{v}_l}$ as long as $j \neq l$ because the equations $X\mathbf{u}_j = \mathbf{0}$ and $\mathbf{v}_l^\top X = \mathbf{0}$ impose at most three additional conditions on the three dimensional $\mathcal{L}_{\mathcal{P}}$. Additionally, the assumption that the wall $W_{\mathbf{u}_i}$ contains a matrix of rank two implies that this wall does not coincide with $W_{\mathbf{u}_j}$ for $j \neq i$.

The first examples of cubic surfaces singular along a line are the cones over a singular plane cubic curve. Our epipolar variety cannot be such a surface because it contains intersecting lines, which these cones do not. There are only two other types of cubic surfaces (up to change of coordinates) that are singular along a line.

The first type is singular along a line, contains a one-dimensional family of lines, and one more line. A representative is given by the equation $w^2y + x^2z$, which contains the lines $\mathcal{V}(w,x)$, $\mathcal{V}(y,z)$ and a family of lines that form a twisted cubic in the Plücker quadric $\mathrm{Gr}(2,4)$ of lines in $\mathbb{P}^3$. The lines in the family are pairwise skew. This type of singular cubic surface only contains two lines that intersect lines in the family, which is inconsistent with the intersection pattern of lines on $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$. The second type is represented by the equation $w^2y+wxz+x^3$. It is singular along the line $\mathcal{V}(x,w)$ and it contains a one-dimensional family of lines. This family of lines is a twisted cubic curve in $\mathrm{Gr}(2,4)$ and they are mutually skew, which also does not fit the intersection pattern on our epipolar variety.

This discussion of cases shows that the epipolar variety cannot be singular along an entire line if it is an irreducible cubic surface in $\mathbb{P}^3$. □

Theorem 4.4.4. *Suppose $|\mathcal{P}| \leq 5$ and $\mathcal{R}_2 \cap \mathcal{L}_P$ is irreducible. Then there exists a chiral reconstruction of $\mathcal{P}$ if and only if $\mathcal{P}_{I(X)}$ has a chiral reconstruction associated to some smooth $X \in \mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$.*

Proof. Recall $I(X)$ and $\mathcal{P}_{I(X)}$ from Definition 4.4.3. Suppose there exists a chiral reconstruction of $\mathcal{P}$. By Theorem 4.4.1, there exists a fundamental matrix X, which is $(\mathbf{u}_i, \mathbf{v}_i)$-regular for all i and $g_ig_j(X) \geq 0$ for all $1 \leq i < j \leq k$. If $g_ig_j(X) > 0$ for all i,j, then the semialgebraic subset of the epipolar variety described by the chiral epipolar inequalities has non-empty interior. Since every open semialgebraic set in an algebraic variety intersects its smooth locus, there is a smooth $X' \in \mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ associated to a chiral reconstruction. Otherwise, $\mathcal{P}$-regularity implies that at most one $(\mathbf{u}_i, \mathbf{v}_i)$ can be the epipole pair of X, i.e., at most one g_i can vanish at X. Without loss of generality, suppose $g_1(X) = 0$ so that $X \in W_{\mathbf{u}_1} \cap W^{\mathbf{v}_1}$. By Lemma 4.4.9, a generic point in $W_{\mathbf{u}_1}$ is a smooth point of $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$. We can move away from the corner $W_{\mathbf{u}_i} \cap W^{\mathbf{v}_i}$ along the wall to a smooth point of $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ without changing the signs of any of the chiral epipolar inequalities, but keeping $g_1 = 0$. This gives a chiral reconstruction of $\mathcal{P}_{I(X)}$ associated to a smooth fundamental matrix, finishing the proof of the first implication.

Conversely, suppose there is a chiral reconstruction of $\mathcal{P}_{I(X)}$ associated to a smooth $X \in \mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$. If $\mathcal{P} = \mathcal{P}_{I(X)}$, there is nothing to show. Otherwise, we have $g_i(X) = 0$ for some i, which means that X lies on some walls by Lemma 4.4.2. Smoothness of X implies that X must have rank two, so its left and right kernels are one-dimensional. Since the $\mathbf{u}_i$'s and $\mathbf{v}_j$'s are distinct, this means X may lie on at most one wall $W_{\mathbf{u}_i}$ and at most one wall $W^{\mathbf{v}_j}$. In other words, $I(X)$ must be one of $[k] \setminus \{i\}$ or $[k] \setminus \{i, j\}$.

Suppose $I(X) = [k] \setminus \{i\}$. We argue locally: Choose an affine chart around X such that $g_j(X) > 0$ for all $j \in [k] \setminus \{i\}$, which is possible because we know that $g_j g_\ell(X) > 0$ for all pairs $j, \ell \in [k] \setminus \{i\} = I(X)$. By Lemma 4.4.8, there is a tangent direction $\mathbf{d} \in T_X(\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}})$ such that $D_{\mathbf{d}} g_i(X) \neq 0$. Then we have $D_{\mathbf{d}}(g_i g_j)(X) = (D_{\mathbf{d}} g_i)(X) g_j(X)$ because $g_i(X) = 0$. Since $g_j(X) > 0$ in our affine chart, the sign of the directional derivative of $g_i g_j$ is determined by $(D_{\mathbf{d}} g_i)(X)$, which is independent of j. By flipping the sign of $\mathbf{d}$, if necessary, we can assume that $(D_{\mathbf{d}} g_i)(X) > 0$. By Taylor approximation, there is a nearby smooth point $X' \in \mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ such that $g_j g_\ell(X) > 0$ for all $j, \ell \in [k]$. By Theorem 4.4.1, there is a chiral reconstruction of $\mathcal{P}$ corresponding to X'.

Suppose $I(X) = [k] \setminus \{i, j\}$. Choose an affine chart around X such that $g_l(X) > 0$ for all $l \in [k] \setminus \{i, j\}$, which is possible because we know that $g_l g_m(X) > 0$ for all pairs $l, m \in [k] \setminus \{i, j\} = I(X)$. Choose linearly independent tangent vectors $\mathbf{d}_i, \mathbf{d}_j \in T_X(\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}})$ such that the directional derivatives $D_{\mathbf{d}_i} g_i(X) \neq 0$ and $D_{\mathbf{d}_j} g_j(X) \neq 0$. Choose a linear combination $\mathbf{d} \in \mathrm{Span}\{\mathbf{d}_i, \mathbf{d}_j\}$ such that $(D_{\mathbf{d}} g_i)(X) > 0$ and $(D_{\mathbf{d}} g_j)(X) > 0$. Since $g_l(X) > 0$ in our affine chart, the sign of the directional derivatives $(D_{\mathbf{d}} g_i g_l)(X)$ and $(D_{\mathbf{d}} g_j g_l)(X)$ at X are determined by $(D_{\mathbf{d}} g_i)(X)$ and $(D_{\mathbf{d}} g_j)(X)$, respectively. By Taylor approximation, there is a nearby smooth point $X' \in \mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ such that $g_l g_m(X) > 0$ for all $l, m \in [k]$. By Theorem 4.4.1, there is a chiral reconstruction of $\mathcal{P}$ associated to X'. $\square$

To apply Theorem 4.4.4, it is useful to understand the smooth locus of the epipolar variety which we describe next. For a more general discussion of tangent spaces to rank varieties we refer to [22, Example 14.16].

Lemma 4.4.10. *Suppose $X \in \mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$ is a rank two matrtix such that $\mathbf{u}$ spans the right kernel of X and $\mathbf{v}$ spans the left kernel of X. Then X is a smooth point of $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$ if and only if $\mathbf{v}\mathbf{u}^\top$ does not lie in the span of $\{\mathbf{v}_i\mathbf{u}_i^\top : i = 1, \dots, k\} \subset \mathbb{P}^8$.*

Proof. The gradient of the determinant of X is the cofactor matrix of X, which is $\mathrm{adj}(X)^\top$. Since $\mathbf{u}\mathbf{v}^\top$ and $\mathrm{adj}(X)$ are collinear, the tangent hyperplane to $\mathcal{R}_2$ at X is

$$T_X\mathcal{R}_2 = \{M \in \mathbb{P}^8 \colon \langle M, \mathbf{v}\mathbf{u}^\top\rangle = 0\}.$$

Since it is a hyperplane, it intersects $\mathcal{L}_\mathcal{P}$ transversely if and only if $\mathcal{L}_\mathcal{P} \not\subset T_X\mathcal{R}_2$, or equivalently, if and only if $\mathbf{v}\mathbf{u}^\top$ does not lie in the span of $\{\mathbf{v}_i\mathbf{u}_i^\top : i = 1, \dots, k\} \subset \mathbb{P}^8$. $\square$

4.5 Four Point Pairs

In the previous section we saw that any set of three point pairs will always have a chiral reconstruction. In this section we completely analyze the case of four point pairs. In Theorem 4.5.4 we prove that if the point configurations in the two views have the same *rank*, then they admit a chiral reconstruction. Otherwise, a chiral reconstruction can fail to exist (Theorem 4.5.5). Our main tool will be Theorem 4.4.4 which requires understanding when the epipolar variety reduces. We do this in Section 4.5.1 which leads to our chirality results in Section 4.5.2.

4.5.1 Irreducibility of the epipolar variety

Suppose the epipolar variety $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$ is reducible. Then $\mathcal{L}_\mathcal{P}$ is not contained in $\mathcal{R}_2$ since otherwise, $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P} = \mathcal{L}_\mathcal{P}$ which is irreducible. Therefore, $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$ is a proper subvariety of $\mathcal{L}_\mathcal{P}$. Further, any irreducible component of $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$ has codimension one in $\mathcal{L}_\mathcal{P}$ because $\det(X) = 0$ is the only additional condition to those defining $\mathcal{L}_\mathcal{P}$. Conversely, if C is a proper subvariety of $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$ of codimension one in $\mathcal{L}_\mathcal{P}$, then C is an irreducible component of $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$. Indeed, it lies in some irreducible component of $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$ of dimension larger than $\dim(C)$, which means C has codimension at least two in $\mathcal{L}_\mathcal{P}$. Lastly, since $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$ has degree three, it must have a linear component if it is reducible and this linear component is a subspace of $\mathcal{R}_2$. Thus to understand whether $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ is reducible, we need to understand the linear subspaces of $\mathcal{R}_2$ that have codimension one in $\mathcal{L}_{\mathcal{P}}$.

A subspace $C \subset \mathbb{C}^{3\times 3}$ is said to have rank at most k if k is the maximum rank of matrices in C. The following definition is from [17].

Definition 4.5.1. *Let $C \subset \mathbb{C}^{3\times 3}$ be a subspace of rank at most k. Then C is called a* compression space *if there exists a subspace $V \subseteq \mathbb{C}^3$ of codimension k_1 and a subspace $W \subseteq \mathbb{C}^3$ of dimension k_2 such that $k_1+k_2 = k$, and every $X \in C$ maps V into W (*compresses V into W*). We refer to C as a (k_1, k_2)-compression space.*

Note that every subspace of a (k_1, k_2)-compression space is again a (k_1, k_2)-compression space. The following theorem from [17] attributed to Atkinson [6] tells us what subspaces of $\mathcal{R}_2$ look like.

Theorem 4.5.1. *[17, Theorem 1.1] A vector space of matrices in $\mathbb{C}^{3\times 3}$ of rank at most two is either a compression space or a subspace of the linear space of 3×3 skew-symmetric matrices.*

Theorem 4.5.2. *Suppose $|\mathcal{P}| = 4$. Then the epipolar variety $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ is reducible if and only if $\mathcal{L}_{\mathcal{P}}$ contains a codimension one compression space of rank at most two.*

Proof. By the above discussion, if $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ is reducible, then it has a linear component C that is a subspace of $\mathcal{R}_2$ and has codimension one in $\mathcal{L}_{\mathcal{P}}$. Since $\dim(\mathcal{L}_{\mathcal{P}}) \geq 4$ when $|\mathcal{P}| = 4$, and the space of 3×3 skew-symmetric matrices has dimension two, by Theorem 4.5.1, C must be a compression space of rank at most two. Conversely, if C is a codimension one compression space in $\mathcal{L}_{\mathcal{P}}$ of rank at most two, then $C \subseteq \mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$, and since C has codimension one in $\mathcal{L}_{\mathcal{P}}$, C must be an irreducible component of $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$. $\square$

Example 4.5.1. *Consider the three point pairs given by the columns of the following matrices:*

$$U = \begin{pmatrix} 1 & 2 & 3 \\ 0 & 0 & 0 \\ 1 & 1 & 1 \end{pmatrix}, \quad V = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 2 & 3 \\ 1 & 1 & 1 \end{pmatrix}.$$

Using Macaulay2 one can compute that $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}} = Q \cup C$ *where* Q *has degree two and* C *has degree one. The ideal of* C *is* $\langle x_{21}, x_{23}, x_{31}, x_{33} \rangle$*, so every* $X \in C$ *has the form*

$$\begin{pmatrix} x_{11} & x_{12} & x_{13} \\ 0 & x_{22} & 0 \\ 0 & x_{32} & 0 \end{pmatrix}.$$

Since $\dim(C) = 4$ *and* $\dim(\mathcal{L}_{\mathcal{P}}) = 5$*,* C *has codimension one in* $\mathcal{L}_{\mathcal{P}}$*. The linear space* C *is a* $(1,1)$*-compression space of rank at most two; each* $X \in C$ *compresses the line spanned by the columns of* U *into the orthogonal complement of the line spanned by the columns of* V*.*

Suppose we now enlarge the set of point pairs to

$$U = \begin{pmatrix} 1 & 2 & 3 & 5 \\ 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix}, \quad V = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 1 & 2 & 3 & 7 \\ 1 & 1 & 1 & 1 \end{pmatrix}$$

Now $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ *has a linear component* C' *consisting of matrices of the form*

$$\begin{pmatrix} x_{11} & x_{12} & x_{13} \\ 0 & x_{22} & 0 \\ 0 & x_{32} & 0 \end{pmatrix} \quad \text{such that} \quad 5x_{11} + x_{12} + x_{13} + 7x_{22} + x_{32} = 0.$$

This subspace C' *is clearly a subspace of the compression space* C *and has codimension one in the new* $\mathcal{L}_{\mathcal{P}}$*. The linear condition arises from the epipolar equation* $\langle X, \mathbf{v}_4 \mathbf{u}_4^\top \rangle = 0$ *imposed by the new point pair.*

The main result of this subsection is the following which will allow us to use Theorem 4.4.4 to establish chirality results for four point pairs in the next subsection. A set of points in $\mathbb{R}^2$ is said to be in *general position* if no three of them are collinear.

Theorem 4.5.3. *Suppose $|\mathcal{P}| = 4$. If for all triples of distinct indices i, j, k, either $\mathbf{u}_i, \mathbf{u}_j, \mathbf{u}_k$ or $\mathbf{v}_i, \mathbf{v}_j, \mathbf{v}_k$ are in general position, then $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ is irreducible.*

Theorem 4.5.3 will follow from Theorem 4.5.2 if under the assumptions of the theorem, there are no compression spaces of rank at most two and codimension one in $\mathcal{L}_{\mathcal{P}}$. We will now prove that this is indeed the case by understanding the orthogonal complements of compression spaces of rank at most two.

Two vector spaces C_1 and C_2 in $\mathbb{C}^{3\times 3}$ are *equivalent* if they have the same linear transformations up to a change of bases in the domain $\mathbb{C}^3$ and codomain $\mathbb{C}^3$, i.e., if there exists $A, B \in \mathrm{GL}_3$ such that $C_2 = \{BXA : X \in C_1\}$. The following classification is straightforward.

Lemma 4.5.1. *Let C be a (k_1, k_2)-compression space of rank at most two in $\mathbb{C}^{3\times 3}$.*

1. *If $(k_1, k_2) = (2, 0)$ then C is equivalent to the vector space of matrices with a zero last column.*

2. *If $(k_1, k_2) = (0, 2)$ then C is equivalent to the vector space of matrices with a zero last row.*

3. *If $(k_1, k_2) = (1, 1)$ then C is equivalent to the vector space of matrices of the form*
$$\begin{pmatrix} 0 & 0 & * \\ 0 & 0 & * \\ * & * & * \end{pmatrix}.$$

Call a (k_1, k_2)-compression space *maximal* if the equations cutting it out are precisely the conditions that impose the zeros guaranteed in their standard forms. Non-maximal compression spaces satisfy further conditions on their potentially non-zero coordinates. The following lemmas help us understand maximal compression spaces.

Lemma 4.5.2. *Suppose $\mathbf{u}, \mathbf{v} \in \mathbb{C}^3 \setminus \{0\}$. Then $\mathbf{v}\mathbf{u}^\top \in S := \mathrm{Span}\{\mathbf{b}_1\mathbf{a}_1^\top, \mathbf{b}_1\mathbf{a}_2^\top, \ldots, \mathbf{b}_n\mathbf{a}_{m-1}^\top, \mathbf{b}_n\mathbf{a}_m^\top\}$ if and only if $\mathbf{u} \in \mathrm{Span}\{\mathbf{a}_i\}_{i=1}^m$ and $\mathbf{v} \in \mathrm{Span}\{\mathbf{b}_i\}_{i=1}^n$.*

Proof. Suppose $\mathbf{v}\mathbf{u}^\top \in S$. Then multiplying $\mathbf{v}\mathbf{u}^\top$ on the right by $\mathbf{u}$, we get that $\mathbf{v} \in \operatorname{Span}\{\mathbf{b}_1, \ldots, \mathbf{b}_n\}$. Similarly, multiplying $\mathbf{v}\mathbf{u}^\top$ on the left by $\mathbf{v}^\top$, we get that $\mathbf{u} \in \operatorname{Span}\{\mathbf{a}_1, \ldots, \mathbf{a}_m\}$. Conversely, if $\mathbf{u} = \lambda_1\mathbf{a}_1 + \ldots, \lambda_m\mathbf{a}_m$ and $\mathbf{v} = \mu_1\mathbf{v}_1 + \ldots, \mu_n\mathbf{v}_n$, then it is immediate that $\mathbf{v}\mathbf{u}^\top \in S$. Note that for this last statement to hold, it was important that all possible outerproducts of the form $\mathbf{b}_i\mathbf{a}_j^\top$ are present in S. □

Lemma 4.5.3. *Let $C \subseteq \mathbb{C}^{3\times 3}$ be a maximal (k_1, k_2)-compression space of rank at most two and let $C^\perp \subseteq \mathbb{C}^{3\times 3}$ be its orthogonal complement.*

1. *If $(k_1, k_2) = (2, 0)$, then $C^\perp = \operatorname{Span}\{\mathbf{b}_1\mathbf{a}^\top, \mathbf{b}_2\mathbf{a}^\top, \mathbf{b}_3\mathbf{a}^\top\}$ where $\mathbf{a} \in \mathbb{C}^3 \setminus \{0\}$ and $\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3$ are linearly independent vectors in $\mathbb{C}^3$. Similarly if $(k_1, k_2) = (0, 2)$, then $C^\perp = \operatorname{Span}\{\mathbf{b}\mathbf{a}_1^\top, \mathbf{b}\mathbf{a}_2^\top, \mathbf{b}\mathbf{a}_3^\top\}$ where $\mathbf{b} \in \mathbb{C}^3 \setminus \{0\}$ and $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$ are linearly independent vectors in $\mathbb{C}^3$. In both cases, as a projective space, $\dim(C^\perp) = 2$ and all matrices in $C^\perp$ have rank one.*

2. *If $(k_1, k_2) = (1, 1)$, then $C^\perp = \operatorname{Span}\{\mathbf{b}_1\mathbf{a}_1^\top,\ \mathbf{b}_1\mathbf{a}_2^\top,\ \mathbf{b}_2\mathbf{a}_1^\top,\ \mathbf{b}_2\mathbf{a}_2^\top\}$ where $\mathbf{b}_1, \mathbf{b}_2$ are linearly independent and $\mathbf{a}_1, \mathbf{a}_2$ are linearly independent vectors in $\mathbb{C}^3$. As a projective space, $\dim(C^\perp) = 3$ and the rank one matrices in it cut out a variety isomorphic to $\mathbb{P}^1 \times \mathbb{P}^1$.*

Proof. 1. Let $\overline{C}$ be the maximal $(2, 0)$-compression space in standard coordinates, consisting of all matrices with a zero last column. Then $(\overline{C})^\perp$ is is spanned by $\mathbf{e}_1\mathbf{e}_3^\top, \mathbf{e}_2\mathbf{e}_3^\top, \mathbf{e}_3\mathbf{e}_3^\top$. By Lemma 4.5.1, a general $(2, 0)$-compression space $C = \{X \in \mathbb{C}^{3\times 3} : BXA \in \overline{C}\}$ for a pair of fixed invertible matrices A, B. Therefore, $0 = \langle BXA, \mathbf{e}_1\mathbf{e}_3^\top\rangle = \langle BXA, \mathbf{e}_2\mathbf{e}_3^\top\rangle = \langle BXA, \mathbf{e}_3\mathbf{e}_3^\top\rangle$ for all $X \in C$. Taking the first dot product:

$$0 = \langle BXA, \mathbf{e}_1\mathbf{e}_3^\top\rangle = \operatorname{Tr}(A^\top X^\top B^\top \mathbf{e}_1\mathbf{e}_3^\top) = \operatorname{Tr}(X^\top (B^\top\mathbf{e}_1)(A\mathbf{e}_3)^\top) = \langle X, (B^\top\mathbf{e}_1)(A\mathbf{e}_3)^\top\rangle.$$

Setting $A\mathbf{e}_i = \mathbf{a}_i$ and $B^\top\mathbf{e}_j = \mathbf{b}_j$, we get that $C^\perp = \operatorname{Span}\{\mathbf{b}_1\mathbf{a}_3^\top, \mathbf{b}_2\mathbf{a}_3^\top, \mathbf{b}_3\mathbf{a}_3^\top\}$. Since B is invertible, $\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3$ are independent, and no row of A is all zero. Hence, $C^\perp$ is spanned by 3 independent rank one matrices and as a projective space, $\dim(C^\perp) = 2$. Further, any linear combination of these three rank one matrices looks like $\lambda\mathbf{b}_1\mathbf{a}_3^\top +$

$\mu\mathbf{b}_2\mathbf{a}_3^\top + \nu\mathbf{b}_3\mathbf{a}_3^\top = (\lambda\mathbf{b}_1 + \mu\mathbf{b}_2 + \nu\mathbf{b}_3)\mathbf{a}_3^\top$ which is a rank one matrix. Therefore, all matrices in $C^\perp$ have rank one.

2. By the same reasoning as in the previous case, if C is a $(1,1)$-compression space, then $C^\perp = \mathrm{Span}\{\mathbf{b}_1\mathbf{a}_1^\top,\ \mathbf{b}_1\mathbf{a}_2^\top, \mathbf{b}_2\mathbf{a}_1^\top, \mathbf{b}_2\mathbf{a}_2^\top\}$. Since C is maximal, $\dim(C^\perp) = 3$ as a projective space, and hence $C^\perp \cong \mathbb{P}^3$. Since the linear combinations

$$\lambda\gamma\mathbf{b}_1\mathbf{a}_1^\top + \lambda\nu\mathbf{b}_1\mathbf{a}_2^\top + \mu\gamma\mathbf{b}_2\mathbf{a}_1^\top + \mu\nu\mathbf{b}_2\mathbf{a}_2^\top = (\lambda\mathbf{b}_1 + \mu\mathbf{b}_2)(\gamma\mathbf{a}_1 + \nu\mathbf{a}_2)^\top$$

are rank one matrices, there is a $\mathbb{P}^1 \times \mathbb{P}^1$ worth of rank 1 matrices in $C^\perp$ which forms a subvariety of dimension two. By Lemma 4.5.2, there are no more rank one matrices in $C^\perp$.

□

From now on we will denote a maximal compression space of rank at most two by C and a subspace of it by C'. If a non-maximal compression space C' appears as an irreducible component of $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$, then $(C')^\perp$ will be spanned by the generators of $C^\perp$ along with some *data matrices* $\mathbf{v}_i\mathbf{u}_i^\perp$ as in Example 4.5.1. In particular a basis of $(C')^\perp$ is the union of the basis for $C^\perp$ from Lemma 4.5.3 and some data matrices.

Lemma 4.5.4. *Suppose $|\mathcal{P}| = 4$ and for all distinct indices i, j, k, either $\mathbf{u}_i, \mathbf{u}_j, \mathbf{u}_k$ or $\mathbf{v}_i, \mathbf{v}_j, \mathbf{v}_k$ are in general position. Let C' be a (k_1, k_2)-compression space contained in a maximal (k_1, k_2) compression space C.*

1. *If C' is of type $(2, 0)$ or $(0, 2)$, then $C^\perp$ contains at most one data matrix.*

2. *If C' is of type $(1, 1)$, then $C^\perp$ contains at most two data matrices.*

Proof. 1. If C' is a (2,0) compression space, then by Lemma 4.5.3,

$$C^\perp = \mathrm{Span}\{\mathbf{b}_1\mathbf{a}^\top, \mathbf{b}_2\mathbf{a}^\top, \mathbf{b}_3\mathbf{a}^\top\}$$

for some $\mathbf{a} \neq 0$ and $\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3$ which are linearly independent. By Lemma 4.5.2, $\mathbf{v}_i\mathbf{u}_i^\top \in C^\perp$ if and only if $\mathbf{u}_i \in \mathrm{Span}\{\mathbf{a}\}$. Since point pairs are distinct and have last coordinate equal to one, at most one $\mathbf{u}_i \in \mathrm{Span}\{\mathbf{a}\}$, hence at most one $\mathbf{v}_i\mathbf{u}_i^\top \in C^\perp$. The case $(0,2)$ is argued similarly.

2. If C' is a (1,1) compression space, then by Lemma 4.5.3

$$C^\perp = \mathrm{Span}\{\mathbf{b}_1\mathbf{a}_1^\top, \mathbf{b}_1\mathbf{a}_2^\top, \mathbf{b}_2\mathbf{a}_1^\top, \mathbf{b}_2\mathbf{a}_2^\top\}$$

for some linearly independent $\mathbf{a}_1, \mathbf{a}_2$ and $\mathbf{b}_1, \mathbf{b}_2$. By Lemma 4.5.2, $\mathbf{v}_i\mathbf{u}_i^\top \in C^\perp$ if and only if $\mathbf{u}_i \in \mathrm{Span}\{\mathbf{a}_1, \mathbf{a}_2\}$ and $\mathbf{v}_i \in \mathrm{Span}\{\mathbf{b}_1, \mathbf{b}_2\}$. If there are three data matrices in $C^\perp$, say corresponding to indices i, j, k, then $\mathbf{u}_i, \mathbf{u}_j, \mathbf{u}_k \in \mathrm{Span}\{\mathbf{a}_1, \mathbf{a}_2\}$ and $\mathbf{v}_i, \mathbf{v}_j, \mathbf{v}_k \in \mathrm{Span}\{\mathbf{b}_1, \mathbf{b}_2\}$. However this contradicts our general position assumption, so we conclude that $C^\perp$ may contain at most two data matrices.

□

We need one final lemma to show that the general position assumption in Theorem 4.5.3 prevents the existence of codimension one compression spaces of rank at most two in $L_{\mathcal{P}}$ when $|\mathcal{P}| = 4$. The following simple fact about our input point pairs will be useful.

Lemma 4.5.5. *If for $i \neq j$, $\mathbf{w}\mathbf{u}_i^\top = \mathbf{v}_j\mathbf{u}_j^\top$ then $\mathbf{w} = \mathbf{v}_j$ and $\mathbf{u}_i = \mathbf{u}_j$. Similarly, if for $i \neq j$, $\mathbf{v}_i\mathbf{u}_i^\top = \mathbf{v}_j\mathbf{w}^\top$ then $\mathbf{w} = \mathbf{u}_i$ and $\mathbf{v}_i = \mathbf{v}_j$.*

Proof. Since the $\mathbf{u}_i$'s have last coordinate 1, if $\mathbf{w}\mathbf{u}_i^\top = \mathbf{v}_j\mathbf{u}_j^\top$, the last columns of the two matrices are $\mathbf{w}$ and $\mathbf{v}_j$. Therefore, $\mathbf{w} = \mathbf{v}_j$ and hence $\mathbf{u}_i = \mathbf{u}_j$. The second statement is proved similarly. □

Lemma 4.5.6. *Suppose $|\mathcal{P}| = 4$ and for all distinct indices i, j, k, either $\mathbf{u}_i, \mathbf{u}_j, \mathbf{u}_k$ or $\mathbf{v}_i, \mathbf{v}_j, \mathbf{v}_k$ are in general position. Let C' be a (k_1, k_2)-compression space of rank at most two and codimension one in $\mathcal{L}_{\mathcal{P}}$, and let C be a maximal (k_1, k_2)-compression space containing C'. Then,*

1. *if C' is of type $(2,0)$ or $(0,2)$, then $C^\perp$ contains at least two data matrices, and*

2. *if C' is of type $(1,1)$, then $C^\perp$ contains at least three data matrices.*

Proof. By the general position assumption, at least three of the data matrices are linearly independent, and hence $\dim(\mathcal{L}_{\mathcal{P}}) = 4$ or 5.

1. Suppose C' is a maximal $(2,0)$-compression space of rank at most two and codimension one in $\mathcal{L}_{\mathcal{P}}$. Then by Lemma 4.5.3, $(C')^{\perp} = \operatorname{Span}\{\mathbf{b}_1\mathbf{a}^{\top}, \mathbf{b}_2\mathbf{a}^{\top}, \mathbf{b}_3\mathbf{a}^{\top}\}$ and $\dim(C') = 5$. This implies that $\dim(\mathcal{L}_{\mathcal{P}}) = 6$ which is a contradiction.

   So suppose C' is a non-maximal $(2,0)$-compression space of rank at most two and codimension one in $\mathcal{L}_{\mathcal{P}}$, and C is a maximal $(2,0)$-compression space containing C'. Define

$$M_C = \begin{pmatrix} \mathbf{b}_1\mathbf{a}^{\top} \\ \mathbf{b}_2\mathbf{a}^{\top} \\ \mathbf{b}_3\mathbf{a}^{\top} \end{pmatrix}, \quad M_D = \begin{pmatrix} \mathbf{v}_1\mathbf{u}_1^{\top} \\ \mathbf{v}_2\mathbf{u}_2^{\top} \\ \mathbf{v}_3\mathbf{u}_3^{\top} \\ \mathbf{v}_4\mathbf{u}_4^{\top} \end{pmatrix}, \quad \text{and} \quad M = \begin{pmatrix} M_C \\ M_D \end{pmatrix}.$$

   (a) $\dim(\mathcal{L}_{\mathcal{P}}) = 5$: Then $\dim(C') = 4$ and $(C')^{\perp}$ has a basis of cardinality four. We may assume without loss of generality that $(C')^{\perp} = \operatorname{Span}\{\mathbf{b}_1\mathbf{a}^{\top}, \mathbf{b}_2\mathbf{a}^{\top}, \mathbf{b}_3\mathbf{a}^{\top}, \mathbf{v}_1\mathbf{u}_1^{\top}\}$. By Lemma 4.5.2, $\mathbf{u}_1$ is not a multiple of $\mathbf{a}$. Since C' has codimension one in $\mathcal{L}_{\mathcal{P}}$, $\operatorname{rank}(M) = 4$ and so $\mathbf{v}_2\mathbf{u}_2^{\top}, \mathbf{v}_3\mathbf{u}_3^{\top}, \mathbf{v}_4\mathbf{u}_4^{\top} \in (C')^{\perp}$. We need to argue that for $i \geq 1$, $\mathbf{v}_i\mathbf{u}_i^{\top} \in C^{\perp} = \operatorname{Span}\{\mathbf{b}_1\mathbf{a}^{\top}, \mathbf{b}_2\mathbf{a}^{\top}, \mathbf{b}_3\mathbf{a}^{\top}\}$. Suppose for $i > 1$, $\mathbf{v}_i\mathbf{u}_i^{\top}$ equals the linear combination

$$\alpha_1\mathbf{b}_1\mathbf{a}^{\top} + \alpha_2\mathbf{b}_2\mathbf{a}^{\top} + \alpha_3\mathbf{b}_3\mathbf{a}^{\top} + \beta\mathbf{v}_1\mathbf{u}_1^{\top} = \begin{pmatrix} \sum \alpha_i\mathbf{b}_i & \beta\mathbf{v}_1 \end{pmatrix} \begin{pmatrix} \mathbf{a}^{\top} \\ \mathbf{u}_1^{\top} \end{pmatrix}. \tag{4.12}$$

   If $\beta = 0$ then we are done, so suppose $\beta \neq 0$. If all the α_i's are zero, then $\mathbf{v}_i\mathbf{u}_i^{\top}$ is a multiple of $\mathbf{v}_1\mathbf{u}_1^{\top}$ which is impossible by Lemma 4.5.5. If at least one $\alpha_i \neq 0$, then (4.12) has rank two unless $\mathbf{v}_1 \sim \sum \alpha_i\mathbf{b}_i$. In this case, (4.12) looks like $\mathbf{v}_1(\gamma\mathbf{a}^{\top} + \delta\mathbf{u}_1^{\top})$. Again, by Lemma 4.5.5, this cannot equal $\mathbf{v}_i\mathbf{u}_i^{\top}$ for any $i = 2, 3, 4$. Therefore, it must be that for $i \geq 1$, $\mathbf{v}_i\mathbf{u}_i^{\top} \in C^{\perp}$.

   (b) $\dim(\mathcal{L}_{\mathcal{P}}) = 4$: In this case, $\dim(C') = 3$, $\operatorname{rank}(M_D) = 4$ and $\operatorname{rank}(M) = 5$. Without loss of generality, $\mathbf{v}_3\mathbf{u}_3^{\top}, \mathbf{v}_4\mathbf{u}_4^{\top} \in (C')^{\perp} = \operatorname{Span}\{\mathbf{b}_1\mathbf{a}^{\top}, \mathbf{b}_2\mathbf{a}^{\top}, \mathbf{b}_3\mathbf{a}^{\top}, \mathbf{v}_1\mathbf{u}_1^{\top}, \mathbf{v}_2\mathbf{u}_2^{\top}\}$.

We need to argue that $\mathbf{v}_3\mathbf{u}_3^\top, \mathbf{v}_4\mathbf{u}_4^\top \in \mathrm{Span}\{\mathbf{b}_1\mathbf{a}^\top, \mathbf{b}_2\mathbf{a}^\top, \mathbf{b}_3\mathbf{a}^\top\}$. Consider a linear combination of the form

$$\alpha_1\mathbf{b}_1\mathbf{a}^\top + \alpha_2\mathbf{b}_2\mathbf{a}^\top + \alpha_3\mathbf{b}_3\mathbf{a}^\top + \beta_1\mathbf{v}_1\mathbf{u}_1^\top + \beta_2\mathbf{v}_2\mathbf{u}_2^\top = \begin{pmatrix} \sum \alpha_i\mathbf{b}_i & \beta_1\mathbf{v}_1 & \beta_2\mathbf{v}_2 \end{pmatrix} \begin{pmatrix} \mathbf{a}^\top \\ \mathbf{u}_1 \\ \mathbf{u}_2 \end{pmatrix}. \tag{4.13}$$

If all the α_i's are 0 then $\mathbf{v}_3\mathbf{u}_3^\top, \mathbf{v}_4\mathbf{u}_4^\top$ lies in the span of the other two data matrices which contradicts that $\mathrm{rank}(M_D) = 4$. If all β's are 0 then we are done. So assume that at least one α_i and one β_i are non-zero. If only one β_i is non-zero, the claim follows from the same argument as in the previous case. So suppose $\beta_1, \beta_2 \neq 0$. By Lemma 4.5.2, neither $\mathbf{u}_1$ nor $\mathbf{u}_2$ are multiples of $\mathbf{a}$ which means that the second matrix in the product has rank at least two. The first matrix also has rank at least two since $\mathbf{v}_1$ and $\mathbf{v}_2$ are linearly independent and hence (4.13) has rank at least two and cannot equal $\mathbf{v}_i\mathbf{u}_i^\top$ for $i = 3, 4$.

The $(0, 2)$ case is argued similarly.

2. Suppose $\mathcal{R}_2 \cap \mathcal{L}_\mathcal{P}$ has an irreducible component C' which is a $(1, 1)$-compression space.

   (a) $\dim(\mathcal{L}_\mathcal{P}) = 5$: In this case, $\dim(C') = 4$ and $C' = C$ is a maximal compression space. Therefore, $(C')^\perp = C^\perp = \mathrm{Span}\{\mathbf{b}_1\mathbf{a}_1^\top, \mathbf{b}_1\mathbf{a}_2^\top, \mathbf{b}_2\mathbf{a}_1^\top, \mathbf{b}_2\mathbf{a}_2^\top\}$ and $\mathrm{rank}(M) = 4$. The matrix M_C now has rows $\mathbf{b}_1\mathbf{a}_1^\top, \mathbf{b}_1\mathbf{a}_2^\top, \mathbf{b}_2\mathbf{a}_1^\top, \mathbf{b}_2\mathbf{a}_2^\top$. Since it already has rank four, it must be that all the data matrices are in $C^\perp$. They are in fact in the $\mathbb{P}^1 \times \mathbb{P}^1$ subvariety of $C^\perp$ containing all the rank matrices.

   (b) $\dim(\mathcal{L}_\mathcal{P}) = 4$: In this case $\dim(C') = 3$ which means that without loss of generality, $(C')^\perp = \mathrm{Span}\{\mathbf{b}_1\mathbf{a}_1^\top, \mathbf{b}_1\mathbf{a}_2^\top, \mathbf{b}_2\mathbf{a}_1^\top, \mathbf{b}_2\mathbf{a}_2^\top, \mathbf{v}_1\mathbf{u}_1^\top\}$. Also, $\mathrm{rank}(M) = 5$ and hence $\mathbf{v}_2\mathbf{u}_2^\top, \mathbf{v}_3\mathbf{u}_3^\top, \mathbf{v}_4\mathbf{u}_4^\top \in (C')^\perp$. We need to argue that these matrices in fact lie in

$C^{\perp} = \text{Span}\{\mathbf{b}_1\mathbf{a}_1^{\top}, \mathbf{b}_1\mathbf{a}_2^{\top}, \mathbf{b}_2\mathbf{a}_1^{\top}, \mathbf{b}_2\mathbf{a}_2^{\top}\}$. Consider a linear combination of the form

$$\alpha_{11}\mathbf{b}_1\mathbf{a}_1^{\top} + \alpha_{12}\mathbf{b}_1\mathbf{a}_2^{\top} + \alpha_{21}\mathbf{b}_2\mathbf{a}_1^{\top} + \alpha_{22}\mathbf{b}_2\mathbf{a}_2^{\top} + \beta\mathbf{v}_1\mathbf{u}_1^{\top} \tag{4.14}$$

$$= \begin{pmatrix} \mathbf{b}_1 & \mathbf{b}_2 & \mathbf{v}_1 \end{pmatrix} \begin{pmatrix} \alpha_{11}\mathbf{a}_1^{\top} + \alpha_{12}\mathbf{a}_2^{\top} \\ \alpha_{21}\mathbf{a}_1^{\top} + \alpha_{22}\mathbf{a}_2^{\top} \\ \beta\mathbf{u}_1^{\top} \end{pmatrix} \tag{4.15}$$

with $\beta \neq 0$. As before, some α_{ij} must also be non-zero. By Lemma 4.5.2, $\mathbf{v}_1 \notin \text{Span}\{\mathbf{b}_1, \mathbf{b}_2\}$ or $\mathbf{u}_1 \notin \text{Span}\{\mathbf{a}_1, \mathbf{a}_2\}$. Suppose $\mathbf{v}_1 \notin \text{Span}\{\mathbf{b}_1, \mathbf{b}_2\}$. Then the first matrix in (4.14) has rank two and hence (4.14) has rank two unless all rows of the second matrix are multiples of $\mathbf{u}_1$. Then as in the previous cases, (4.14) cannot equal $\mathbf{v}_i\mathbf{u}_i^{\top}$ for any $i = 2, 3, 4$. If $\mathbf{u}_1 \notin \text{Span}\{\mathbf{a}_1, \mathbf{a}_2\}$ then the second matrix in (4.14) has rank two and hence the linear combination has rank two unless all columns of the first matrix are multiples of $\mathbf{v}_1$. Again by Lemma 4.5.5, (4.14) cannot be $\mathbf{v}_i\mathbf{u}_i^{\top}$ for any $i = 2, 3, 4$.

□

Proof of Theorem 4.5.3. If $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ is reducible, then by Theorem 4.5.2 there is a codimension one compression space of rank at most two in $L_{\mathcal{P}}$. However, if for all distinct indices i, j, k, either $\mathbf{u}_i, \mathbf{u}_j, \mathbf{u}_k$ or $\mathbf{v}_i, \mathbf{v}_j, \mathbf{v}_k$ are in general position, then Lemma 4.5.4 and Lemma 4.5.6 together prohibit the existence of a codimension one compression space of rank at most two in $\mathcal{L}_{\mathcal{P}}$, proving the theorem. □

4.5.2 Chiral Reconstructions for Four Point Pairs

We now have the tools to determine when a set of four distinct point pairs have a chiral reconstruction. Up to permuting the two images, every set of four point pairs belong to one of the combinatorial types shown in Figure 4.1.

We can further group these types by the *rank* of the points in $\mathcal{U} = \{\mathbf{u}_1, \ldots, \mathbf{u}_4\}$ and $\mathcal{V} = \{\mathbf{v}_1, \ldots, \mathbf{v}_4\}$.

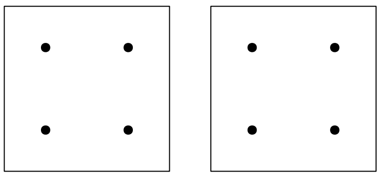

(a) $\mathcal{U}$ and $\mathcal{V}$ in general position.

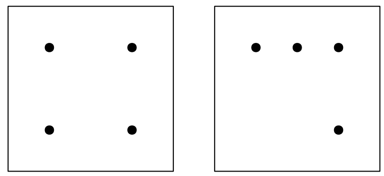

(b) $\mathcal{U}$ general position, three $\mathbf{v}_i$ collinear.

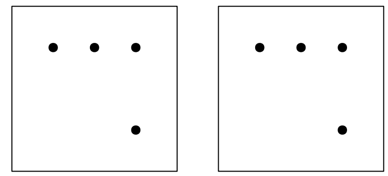

(c) Three $\mathbf{u}_i$ collinear, three $\mathbf{v}_i$ collinear.

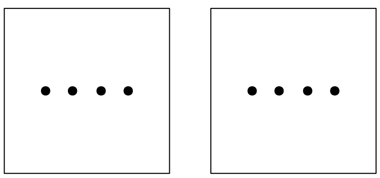

(d) Four $\mathbf{u}_i$ collinear, four $\mathbf{v}_i$ collinear.

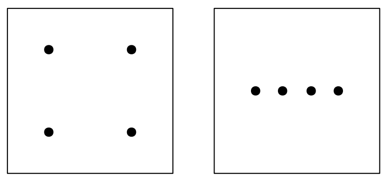

(e) $\mathcal{U}$ general position, four $\mathbf{v}_i$ collinear.

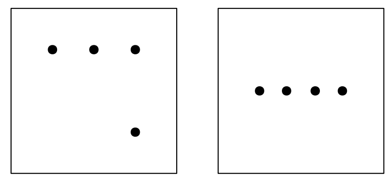

(f) Three $\mathbf{u}_i$ collinear, four $\mathbf{v}_i$ collinear.

Figure 4.1

Definition 4.5.2. *We say that* $\operatorname{rank}(\mathcal{U}) = r$ *if and only if the* 3×4 *matrix* U *with columns* $\mathbf{u}_1, \ldots, \mathbf{u}_4$ *has rank* r. *Equivalently,* $\operatorname{rank}(\mathcal{V}) = r$ *if and only if the* 3×4 *matrix* V *with columns* $\mathbf{v}_1, \ldots, \mathbf{v}_4$ *has rank* r.

Rank captures the geometry of $\mathcal{U}$ and $\mathcal{V}$. Indeed, if $rank(\mathcal{U}) = 1$, then all $\mathbf{u}_i$ are coincident as points in $\mathbb{P}^2$. If $rank(\mathcal{U}) = 2$, then all $\mathbf{u}_i$ are collinear in $\mathbb{P}^2$, and if $rank(\mathcal{U}) = 3$, then some three $\mathbf{u}_i$ are non-collinear in $\mathbb{P}^2$. The assumption that points in $\mathcal{U}$ and $\mathcal{V}$ are distinct implies that $\operatorname{rank}(\mathcal{U}), \operatorname{rank}(\mathcal{V}) \geq 2$. The spirit of Theorem 4.5.4, the main result of this section, is that when $\mathcal{U}$ and $\mathcal{V}$ have similar geometry, then $\mathcal{P}$ has a chiral reconstruction. In particular, a chiral reconstruction exists for the combinatorial types in Figure 4.1a, Figure 4.1b, and Figure 4.1c where $\operatorname{rank}(\mathcal{U}) = \operatorname{rank}(\mathcal{V}) = 3$ and for the type in Figure 4.1d where $\operatorname{rank}(\mathcal{U}) = \operatorname{rank}(\mathcal{V}) = 2$. We will present examples of configurations of the types in Figure 4.1e and Figure 4.1f for which a chiral reconstruction can fail to exist.

Theorem 4.5.4. *If* $|\mathcal{P}| = 4$ *and* $\operatorname{rank}(\mathcal{U}) = \operatorname{rank}(\mathcal{V})$, *then* $\mathcal{P}$ *has a chiral reconstruction.*

Proof. We break the proof into three parts:

1. Suppose $\text{rank}(\mathcal{U}) = \text{rank}(\mathcal{V}) = 2$. Then the points in both $\mathcal{U}$ and $\mathcal{V}$ are collinear as in Figure 4.1d. Assume without loss of generality that the $\mathbf{v}_i$ points appear in the order $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4$ along the line L they span. The ordering of $\mathbf{v}_i$ induces an ordering of the $\mathbf{u}_i$. Let $l, r \in \{1, 2, 3, 4\}$ be such that for all $i \in \{1, 2, 3, 4\}$, $\mathbf{u}_i \in [\mathbf{u}_l, \mathbf{u}_r]$, and define $G \in GL_3$ such that $G\mathbf{u}_l = \mathbf{v}_1$ and $G\mathbf{u}_r = \mathbf{v}_4$. This forces $\mathbf{v}_i, G\mathbf{u}_i \in [\mathbf{v}_1, \mathbf{v}_4]$ for all i. For a $\mathbf{t} \in L \smallsetminus [\mathbf{v}_1, \mathbf{v}_4]$, $(\mathbf{t} \times \mathbf{v}_i)^\top (\mathbf{t} \times G\mathbf{u}_i) > 0$ by Lemma 4.4.7. Then, Lemma 4.4.1 implies that $g_i([\mathbf{t}]_\times G) > 0$ for all i, proving the result.

2. Suppose $\text{rank}(\mathcal{U}) = \text{rank}(\mathcal{V}) = 3$ and for all i, j, k, either $\mathbf{u}_i, \mathbf{u}_j, \mathbf{u}_k$ or $\mathbf{v}_i, \mathbf{v}_j, \mathbf{v}_k$ are in general position. Then by Theorem 4.5.3, $\mathcal{R}_2 \cap L_{\mathcal{P}}$ is irreducible, and hence by Theorem 4.4.4, it suffices to show that there is a smooth rank two matrix $X \in \mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ for which $\mathcal{P}_{I(X)}$ has a chiral reconstruction.

   We first claim that there is an ordering of point pairs such that $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$ are in general position and $\mathbf{v}_4$ does not lie on any of the lines L^{12}, L^{13}, and L^{23} where L^{ij} is the line spanned by $\mathbf{v}_i$ and $\mathbf{v}_j$. Indeed, since $\text{rank}(\mathcal{V}) = 3$, there is an ordering of the point pairs so that $\mathbf{v}_4$ does not lie on any of the lines L^{12}, L^{13}, or L^{23}. If $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$ are in general position for this ordering, we are done. Otherwise, $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$ collinear and so $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ must be in general position by assumption, and hence $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4$ are in general position. Also, since $\text{rank}(\mathcal{U}) = 3$, $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_4$ must be non-collinear. Since $\mathbf{v}_3$ cannot be on the lines L^{12}, L^{14}, L^{24}, swapping point pairs with indices 3 and 4 gives the desired result.

   Reorder point pairs so that $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$ are in general position and $\mathbf{v}_4$ does not lie on any of L^{12}, L^{13}, and L^{23}. Then choosing $\mathbf{e}_2 = \mathbf{v}_4$ forces $\det[\mathbf{v}_1\, \mathbf{v}_2\, \mathbf{e}_2]$, $\det[\mathbf{v}_1\, \mathbf{v}_3\, \mathbf{e}_2]$, and $\det[\mathbf{v}_2\, \mathbf{v}_3\, \mathbf{e}_2]$ to all be non-zero. Removing $\mathbf{u}_4$ from $\mathcal{U}$ leaves three points in general position, so we may pick $\mathbf{e}_1$ by Lemma 4.4.6 so that $\det[\mathbf{u}_i\, \mathbf{u}_j\, \mathbf{e}_1]\det[\mathbf{v}_i\, \mathbf{v}_j\, \mathbf{e}_2] > 0$ for all $1 \leq i < j \leq 3$. Since the positivity of these determinant products is an open condition, there is a open polyhedral cone from which such an $\mathbf{e}_1$ can be chosen. Consider the linear system

$$X\mathbf{e}_1 = 0,\ \mathbf{v}_4^\top X = 0,\ \mathbf{v}_i^\top X\mathbf{u}_i = 0,\ i = 1,2,3,4$$

which consists of at most eight linearly independent equations. Therefore, this system has a solution $X \in \mathbb{P}^8$. From Lemma 4.4.10, we know the tangent space to $\mathcal{R}_2$ at X has normal $\mathbf{v}_4\mathbf{e}_1^\top$ and X is a smooth point of the epipolar variety if $\mathbf{v}_4\mathbf{e}_1^\top$ does not lie in Span$\{\mathbf{v}_i\mathbf{u}_i^\top,\ i = 1,2,3,4\}$. If $\mathbf{v}_4\mathbf{e}_1^\top \in$ Span$\{\mathbf{v}_i\mathbf{u}_i^\top,\ i = 1,2,3,4\}$, there is a minimal subset of the five matrices, including $\mathbf{v}_4\mathbf{e}_1^\top$ that are dependent. Placing the vectorizations of these matrices in the rows of a matrix, we get that all its maximal minors are zero. Each of these minors, set to zero, is a linear equation in $\mathbf{e}_1$ and together they cut out a subspace in $\mathbb{P}^2$. Since we were able to choose $\mathbf{e}_1$ from an open polyhedral cone, there is a choice of $\mathbf{e}_1$ that avoids this subspace and the rank one variety. Hence, X can be chosen (by choosing $\mathbf{e}_1$ appropriately) to be a smooth rank two fundamental matrix and we are done by Theorem 4.4.4.

3. Suppose rank$(\mathcal{U})$ = rank$(\mathcal{V}) = 3$ and there exist i, j, k for which both $\mathbf{u}_i, \mathbf{u}_j, \mathbf{u}_k$ and $\mathbf{v}_i, \mathbf{v}_j, \mathbf{v}_k$ are collinear. Without loss of generality we may assume that $\{i, j, k\} = \{1, 2, 3\}$, and that $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ appear in this order on the line L they span. Let $l, r \in \{1, 2, 3\}$ be such that $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3 \in [\mathbf{u}_l, \mathbf{u}_r]$, and define G by $G(\mathbf{u}_l) = \mathbf{v}_1$, $G(\mathbf{u}_r) = \mathbf{v}_3$, and $G(\mathbf{u}_4) = \mathbf{v}_4$. Then for a $\mathbf{t} \in L \smallsetminus [\mathbf{v}_1, \mathbf{v}_3]$, set $X = [\mathbf{t}]_\times G$. Since, $\mathbf{v}_i, \mathbf{t}$, and $G\mathbf{u}_i$ are collinear for each i, X satisfies all the epipolar equations $\mathbf{v}_i^\top X\mathbf{u}_i = 0$. Also, since $\mathbf{v}_i, G\mathbf{u}_i \in [\mathbf{v}_1, \mathbf{v}_3]$ for $i = 1, 2, 3$, $g_1(X), g_2(X), g_3(X) > 0$ by Lemma 4.4.1 and Lemma 4.4.7. Finally, $g_4(X) = (\mathbf{t} \times \mathbf{v}_4)^\top(\mathbf{t} \times G\mathbf{u}_4) = (\mathbf{t} \times \mathbf{v}_4)^\top(\mathbf{t} \times \mathbf{v}_4) > 0$, and so X strictly satisfies all chiral epipolar inequalities and $\mathcal{P}$ has a chiral reconstruction by Theorem 4.4.1.

$\square$

We conclude by showing that for four point pairs, Theorem 4.5.4 is best possible in the following sense.

Theorem 4.5.5. *For the two combinatorial types where* $\operatorname{rank}(\mathcal{U}) \neq \operatorname{rank}(\mathcal{V})$*, there are instances of* $\mathcal{P}$ *without a chiral reconstruction.*

We give examples to prove Theorem 4.5.5. Recall that the epipolar line homography of Theorem 4.3.1 cannot send coincident lines to distinct lines or vice versa. Therefore if X is a $\mathcal{P}$-regular fundamental matrix with right and left kernels generated by $\mathbf{e}_1$ and $\mathbf{e}_2$, respectively, then $\mathbf{e}_1, \mathbf{u}_i, \mathbf{u}_j$ are collinear if and only if $\mathbf{e}_2, \mathbf{v}_i, \mathbf{v}_j$ are collinear.

Example 4.5.2. *Consider the arrangement in Figure 4.2 where* $\operatorname{rank}(\mathcal{U}) = 3 \neq \operatorname{rank}(\mathcal{V}) = 2$.

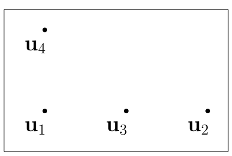

$$U = \begin{pmatrix} 1 & 3 & 2 & 1 \\ 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \qquad V = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{pmatrix}$$

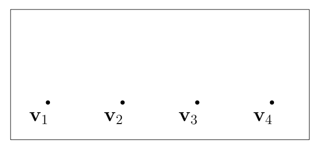

Figure 4.2

Let L *be the line in* $\mathbb{P}^2$ *spanned by* $(0,0,1)^\top$ *and* $(1,0,1)^\top$. *Suppose there is a fundamental matrix* X *for* $\mathcal{P}$ *with* $\mathbf{e}_2^\top X = 0 = X\mathbf{e}_1$, *and suppose* $\mathbf{e}_2 = (e_{21}, e_{22}, e_{23})^\top$ *and* $\mathbf{e}_1 = (e_{11}, e_{12}, e_{13})^\top$. *Then* $\mathbf{e}_2 \notin L$ *(equivalently,* $e_{22} \neq 0$*) because then the pencil* $\mathbf{e}_2(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4)$ *consists of a repeated line, while the pencil* $\mathbf{e}_1(\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \mathbf{u}_4)$ *contains at least two distinct lines for any choice of* $\mathbf{e}_1$. *If* $\mathbf{e}_2 \in \mathbb{P}^2 \setminus L$, *then* $\mathbf{e}_2(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4)$ *consists of four distinct lines, and hence* $\mathbf{e}_1(\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \mathbf{u}_4)$ *must also have four distinct lines. In particular,* $\mathbf{e}_1 \notin L$, *or equivalently,* $e_{12} \neq 0$. *These restrictions imply that* $D_{12}(\mathbf{e}_1, \mathbf{e}_2) = 2e_{12}e_{22}$, $\quad D_{13}(\mathbf{e}_1, \mathbf{e}_2) = 2e_{12}e_{22}$, $\quad D_{23}(\mathbf{e}_1, \mathbf{e}_2) = -e_{12}e_{22}$, *are all non-zero. It is then clear that they cannot all have the same sign. Therefore, by Theorem 4.4.2, there is no chiral reconstruction of* $\mathcal{P}$.

The epipolar variety in Example 4.5.2 is reducible; $W^{\mathbf{v}_4}$ is a linear component containing the wall $W_{\mathbf{u}_2}$. Let $\mathbf{t} = \mathbf{v}_4$ and define $G \in \mathrm{GL}_3$ such that $G\mathbf{u}_1 = \mathbf{v}_1$ and $G\mathbf{u}_2 = \mathbf{v}_4$. Then $X = [\mathbf{t}]_\times G$ is a fundamental matrix of $\mathcal{P}$. Indeed, $\mathbf{v}_4^\top X = 0$ and $X\mathbf{u}_2 = [\mathbf{v}_4]_\times G\mathbf{u}_2 = [\mathbf{v}_4]_\times \mathbf{v}_4 = 0$

and therefore X satisfies the second and fourth epipolar equations. It is straightforward to check that $\mathbf{v}_1, \mathbf{t}, G\mathbf{u}_1$ are collinear and $\mathbf{v}_3, \mathbf{t}, G\mathbf{u}_3$ are collinear, so X also satisfies the first and third epipolar equations. By construction, $g_2(X) = g_4(X) = 0$, and by Lemma 4.4.1 and Lemma 4.4.7, $g_1 g_3(X) > 0$. Thus, X is a fundamental matrix of $\mathcal{P}$ that satisfies all the chiral epipolar inequalities $g_i g_j(X) \geq 0$ for all $1 \leq i < j \leq 4$. However, X lies in the corner $W_{\mathbf{u}_2} \cap W^{\mathbf{v}_4}$ and is not $\mathcal{P}$-regular and all neighboring matrices to X on $W^{\mathbf{v}_4}$ are also $\mathcal{P}$-irregular. Therefore, we cannot use Theorem 4.4.4 to perturb X to a regular fundamental matrix.

Example 4.5.2 illustrates why the irreducibility of the epipolar variety is needed in Theorem 4.4.4. The next example shows that a chiral reconstruction may not exist even when the epipolar variety is irreducible.

Example 4.5.3. *Consider the arrangement in Figure 4.3 where* $\operatorname{rank}(\mathcal{U}) = 3 \neq \operatorname{rank}(\mathcal{V}) = 2$.

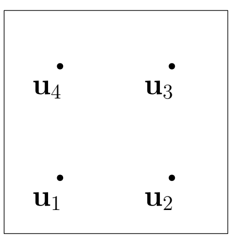

$$U = \begin{pmatrix} 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \qquad V = \begin{pmatrix} 1 & 3 & 2 & 4 \\ 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{pmatrix}$$

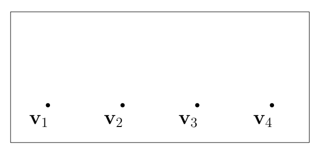

Figure 4.3

By Theorem 4.5.3, this epipolar variety is irreducible. Suppose $\mathcal{P}$ *has a fundamental matrix* X *with left epipole* $\mathbf{e}_2 = (e_{21}, e_{22}, e_{23})^\top$ *and right epipole* $\mathbf{e}_1 = (e_{11}, e_{12}, e_{13})^\top$*. By Theorem 4.3.1,* $\mathbf{e}_2$ *cannot be on the line spanned by the* $\mathbf{v}$ *points since there is no* $\mathbf{e}_1$ *for which the pencil* $\mathbf{e}_1(\mathbf{u}_1, \ldots, \mathbf{u}_4)$ *consists of coincident lines. For any* $\mathbf{e}_2 = (e_{21}, e_{22}, e_{23})^\top$ *not on the line spanned by the* $\mathbf{v}$ *points, the vector*

$$(\det[\mathbf{v}_i\, \mathbf{v}_j\, \mathbf{e}_2] \,:\, 1 \leq i < j \leq 4) = (2e_{22}, e_{22}, 3e_{22}, -e_{22}, e_{22}, 2e_{22})$$

which has sign pattern $+++-++$ *(up to sign) since* $e_{22} \neq 0$*, and for any* $\mathbf{e}_1 = (e_{11}, e_{12}, e_{13})^\top$*,*

$(\det[\mathbf{u}_i\,\mathbf{u}_j\,\mathbf{e}_1] \,:\, 1 \leq i < j \leq 4) = (e_{12}, -e_{11}+e_{12}, -e_{11}, -e_{11}+e_{13}, -e_{11}-e_{12}+e_{13}, -e_{12}+e_{13})$.

We can guarantee a chiral reconstruction if

$$e_{12} > 0, -e_{11} + e_{12} > 0, -e_{11} > 0, -e_{11} + e_{13} < 0, -e_{11} - e_{12} + e_{13} > 0, -e_{12} + e_{13} > 0.$$

However, this implies that $e_{12} > e_{11} > e_{13} > e_{12}$ *which is a contradiction. Similarly, we cannot achieve the sign pattern* $---+--$ *either. What might still be possible is to choose* $\mathbf{e}_1$ *so that some of the determinants* $\det[\mathbf{u}_i\,\mathbf{u}_j\,\mathbf{e}_1]$ *are zero, or equivalently,* $\mathbf{e}_1$ *lies on the line joining two of the* $\mathbf{u}$ *points. This will not work since then two coincident rays will have to map to two non-coincident rays under the homography in Theorem 4.3.1. Therefore,* $\mathcal{P}$ *has no chiral reconstruction.*

Unlike in Example 4.5.2, there is no fundamental matrix for Example 4.5.3 that satisfies all chiral epipolar inequalities. Indeed, Theorem 4.4.4 implies that there cannot be a smooth X such that $\mathcal{P}_{I(X)}$ has a chiral reconstruction associated to X. This rules out the possibility that there is an irregular X with respect to some index that satisfies the chiral epipolar inequalities.

4.6 Five Point Pairs

We have seen so far that three point pairs unconditionally have a chiral reconstruction as do four point pairs unless their geometry is special and dissimilar. In this section we will see that five point pairs can fail to have a chiral reconstruction even if both sets of points are in general position. Specific examples of such point pairs were known to Werner [52]. We will show that, in fact, five point pairs do not have a chiral reconstruction with positive probability. We will use Theorem 4.4.4 to devise a simple test for chirality in this section, and connect to the classical theory of cubic surfaces in the next section.

Throughout this section, let $\mathcal{P} = \{(\mathbf{u}_i, \mathbf{v}_i) : i = 1, \ldots, 5\}$ be a set of five generic point pairs in the sense that $\mathcal{L}_\mathcal{P}$ is a generic codimension five plane in $\mathbb{P}^8$. In this case, $\mathcal{L}_\mathcal{P}$ misses the four-dimensional variety $\mathcal{R}_1$, and by Bertini's Theorem, the epipolar variety $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ is a smooth, irreducible cubic surface. Each wall $W_{\mathbf{u}_i}$ or $W^{\mathbf{v}_j}$ is a line on $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$. We first see how these lines intersect.

Lemma 4.6.1. *Let $\mathcal{P}$ be a set of five generic point pairs. Then*

1. *$W_{\mathbf{u}_i}$ and $W_{\mathbf{u}_j}$, and similarly $W^{\mathbf{v}_i}$ and $W^{\mathbf{v}_j}$, do not intersect for all $i \neq j$.*

2. *$W_{\mathbf{u}_i}$ and $W^{\mathbf{v}_i}$ do not intersect for all i.*

3. *The corner $W_{\mathbf{u}_i} \cap W^{\mathbf{v}_j}$ consists of a unique real rank two matrix for $i \neq j$.*

Proof.

1. Recall that for $i \neq j$, $\mathbf{u}_i \not\sim \mathbf{u}_j$, and so there is no rank two X such that $X\mathbf{u}_i = 0 = X\mathbf{u}_j$.

2. Since $\widetilde{W}_{\mathbf{u}_i}, \widetilde{W}^{\mathbf{v}_i} \subset L_{(\mathbf{u}_i,\mathbf{v}_i)}$, $W_{\mathbf{u}_i} \cap W^{\mathbf{v}_i}$ is the intersection of the three-dimensional linear space $\widetilde{W}_{\mathbf{u}_i} \cap \widetilde{W}^{\mathbf{v}_i}$ with the generic codimension four linear space $\bigcap_{l \neq i} L_{(\mathbf{u}_l,\mathbf{v}_l)}$, and hence empty.

3. When $i \neq j$, $W_{\mathbf{u}_i} \cap W^{\mathbf{v}_j}$ is the intersection of the three-dimensional linear space $\widetilde{W}_{\mathbf{u}_i} \cap \widetilde{W}^{\mathbf{v}_j}$ with the codimension three linear space $\bigcap_{l \neq i,j} L_{(\mathbf{u}_l,\mathbf{v}_l)}$. This intersection is zero-dimensional and since the defining equations are linear with real coefficients, it consists of a single real matrix X. Generically, the intersection will miss $\mathcal{R}_1$, so we conclude that X is a real rank two matrix.

□

Remark 4.6.1. *Concretely, when $i \neq j$, the rank two matrix $W_{\mathbf{u}_i} \cap W^{\mathbf{v}_j}$ is $X = [\mathbf{v}_j]_\times H$ where H is the unique homography sending $\mathbf{u}_i$ to $\mathbf{v}_j$ and $\mathbf{u}_l$ to $\mathbf{v}_l$ for all $l \neq i, j$.*

Definition 4.6.1. *Let C be the set of all wall intersections on the epipolar variety:*

$$C = \bigcup_{1 \leq i,j \leq k} W_{\mathbf{u}_i} \cap W^{\mathbf{v}_j}.$$

We call the points in C the corners *of $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$.*

Lemma 4.6.1 shows that C consists of 20 distinct fundamental matrices of $\mathcal{P}$. However, they do not correspond to full projective reconstructions of $\mathcal{P}$ because necessarily the $(\mathbf{u}_i, \mathbf{v}_j)$ corner is neither $(\mathbf{u}_i, \mathbf{v}_i)$ regular nor $(\mathbf{u}_j, \mathbf{v}_j)$ regular. Despite this fact, we argue that checking the signs of the chiral epipolar inequalities at these corners suffice to determine whether $\mathcal{P}$ has a chiral reconstruction.

Theorem 4.6.1. *Let $\mathcal{P}$ be a generic set of five point pairs. Then $\mathcal{P}$ has a chiral reconstruction if and only if $\mathcal{P}_{I(X)}$ has a chiral reconstruction for one of the 20 corners $X \in C$.*

Proof. Suppose $\mathcal{P}$ has a chiral reconstruction. Then Theorem 4.4.1 implies that there is a $\mathcal{P}$-regular fundamental matrix X such that $g_i g_j(X) \geq 0$ for all i, j. By Lemma 4.6.1, the epipolar variety does not contain the $(\mathbf{u}_i, \mathbf{v}_i)$ corner for any i. Therefore regularity with respect to all image pairs implies $g_i(X) \neq 0$ for all i, meaning $g_i g_j(X) > 0$ for all i, j. Let U be the connected component of the semialgebraic subset of the epipolar variety described by $g_i g_j \geq 0$ that contains X. The chiral epipolar inequalities only change sign along walls. Either U is the entire epipolar variety or its boundary is contained in the walls. Every wall contains four corners and every interval on a wall that is contained in U is bounded by corners (or unbounded, in which case it contains all four corners on the wall). So U contains corners and every corner $X \in U$ lies in C and corresponds to a chiral reconstruction of the subset $\mathcal{P}_{I(X)}$ of point pairs.

Conversely, suppose $\mathcal{P}_{I(X)}$ has a chiral reconstruction associated to some $X \in C$. The epipolar variety $\mathcal{R}_2 \cap L_{\mathcal{P}}$ is smooth, irreducible, and consists only of rank two matrices. Hence the points in C are smooth fundamental matrices and Theorem 4.4.4 implies there is a chiral reconstruction of $\mathcal{P}$. □

Using the results in Section 4.4.2, it is easy to evaluate the chiral epipolar inequalities at the corners of the epipolar variety. In particular, Theorem 4.4.2 allows us to infer the sign of all non-zero chiral epipolar inequalities at $X \in C$ directly using the input point pairs. The next corollary makes this precise.

Corollary 4.6.1. *Let $\mathcal{P}$ be a set of five generic point pairs. Then $\mathcal{P}$ has a chiral reconstruction if and only if there is some $(\mathbf{u}_i, \mathbf{v}_j)$ corner such that $D_{lm}(\mathbf{u}_i, \mathbf{v}_j), D_{ln}(\mathbf{u}_i, \mathbf{v}_j)$, and $D_{mn}(\mathbf{u}_i, \mathbf{v}_j)$ have the same sign for $l, m, n \in [5] \setminus \{i, j\}$.*

Proof. Since $\mathcal{P}$ is generic, no three $\mathbf{u}_i$ are collinear and no three $\mathbf{v}_i$ are collinear. Hence, the expressions $D_{lm}(\mathbf{u}_i, \mathbf{v}_j)$, $D_{ln}(\mathbf{u}_i, \mathbf{v}_j)$, and $D_{mn}(\mathbf{u}_i, \mathbf{v}_j)$ are all non-zero for pairwise distinct $l, m, n \in [k] \setminus \{i, j\}$. Suppose there is some $(\mathbf{u}_i, \mathbf{v}_j)$ corner $X \in C$ such that $D_{lm}(\mathbf{u}_i, \mathbf{v}_j), D_{ln}(\mathbf{u}_i, \mathbf{v}_j)$, and $D_{mn}(\mathbf{u}_i, \mathbf{v}_j)$ have the same sign. By Theorem 4.4.2, $\mathcal{P}_{I(X)}$ has a chiral reconstruction associated to X. Theorem 4.6.1 then implies that a chiral reconstruction of $\mathcal{P}$ exists. $\square$

Corollary 4.6.1 gives an algorithm to check whether five generic point pairs have a chiral reconstruction. We evaluate the sign of three polynomials on the input point pairs at each of the 20 corners. The signs are the same at a corner if and only if it bounds part of the chiral epipolar region of $\mathcal{P}$. We illustrate the procedure on an example.

Example 4.6.1. *Let*

$$U = \begin{pmatrix} 0 & 0 & 4 & 2 & 2 \\ 0 & 4 & 0 & 1 & 3 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix} \quad \text{and} \quad V = \begin{pmatrix} 2 & 2 & 4 & 0 & 1 \\ 1 & 3 & 0 & 4 & 1 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix}, \tag{4.16}$$

and $\mathcal{P} = \{(\mathbf{u}_i, \mathbf{v}_i)\}$ where $\mathbf{u}_i$ is the ith column of U and $\mathbf{v}_i$ is the ith column of V. By Corollary 4.6.1 and the following table, $\mathcal{P}$ has no chiral reconstruction.

i	j	$D_{lm}(\mathbf{u}_i,\mathbf{v}_j)$	$D_{ln}(\mathbf{u}_i,\mathbf{v}_j)$	$D_{mn}(\mathbf{u}_i,\mathbf{v}_j)$
1	2	−16	−84	20
1	3	−32	−56	32
1	4	64	40	−96
1	5	112	−40	32
2	1	−16	−4	12
2	3	−32	8	32
2	4	64	−24	−32
2	5	−16	24	−32
3	1	16	−8	−12
3	2	16	24	−20

i	j	$D_{lm}(\mathbf{u}_i,\mathbf{v}_j)$	$D_{ln}(\mathbf{u}_i,\mathbf{v}_j)$	$D_{mn}(\mathbf{u}_i,\mathbf{v}_j)$
3	4	−64	36	20
3	5	−32	−12	20
4	1	16	−8	−4
4	2	16	8	−28
4	3	32	−4	−28
4	5	−16	−4	28
5	1	−16	16	−16
5	2	48	−16	16
5	3	32	−16	16
5	4	−32	48	−16

Now consider the following modification of the above example obtained by perturbing $\mathbf{v}_5$*:*

$$U = \begin{pmatrix} 0 & 0 & 4 & 2 & 2 \\ 0 & 4 & 0 & 1 & 3 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix} \quad \text{and} \quad V = \begin{pmatrix} 2 & 2 & 4 & 0 & 4 \\ 1 & 3 & 0 & 4 & 4 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix}. \tag{4.17}$$

We compute the products D_{lm}, D_{ln}, D_{mn} *at all 20 corners* $(\mathbf{u}_i, \mathbf{v}_j)$ *and find that*

$$D_{14}(\mathbf{u}_2,\mathbf{v}_3) = -32, \quad D_{15}(\mathbf{u}_2,\mathbf{v}_3) = -64, \quad D_{45}(\mathbf{u}_2,\mathbf{v}_3) = -64,$$

and hence the $(\mathbf{u}_2, \mathbf{v}_3)$ *corner bounds the chiral epipolar region, so these point pairs have a chiral reconstruction. We point out that, the same-signness property is not symmetric in the sense that if it holds for* $(\mathbf{u}_i, \mathbf{v}_j)$*, it need not hold for* $(\mathbf{u}_j, \mathbf{v}_i)$*. Indeed,*

$$D_{14}(\mathbf{u}_3,\mathbf{v}_2) = 16, \quad D_{15}(\mathbf{u}_3,\mathbf{v}_2) = -48, \quad D_{45}(\mathbf{u}_3,\mathbf{v}_2) = 16.$$

The point pairs in (4.16) are a slight modification of [52, Figure 1] for which the epipolar variety is singular. Our modification makes the variety non-singular, something we will discuss in the next section. Furthermore, since sufficiently small perturbations of point pairs do not change the signs of D_{lm}, D_{ln}, and D_{mn} at any $(\mathbf{u}_i, \mathbf{v}_j)$ corner, our methods show that having no chiral reconstruction is an open condition. Hence, there are Zariski dense subsets of five point pairs without a chiral reconstruction. Finally, embedding the point pairs in (4.16) into instances of six or more point pairs one can construct larger configurations with no chiral reconstruction.

4.7 Connections to Classical Algebraic Geometry

In this section, we discuss the connection between five point pairs and the classical theory of real cubic surfaces in projective 3-space. This point of view is yet another way to study the space of epipoles, which goes back to Klein and Segre [44]. General references for cubic surfaces are still mostly classical books like [26],[28], or [44]. More modern accounts of classical facts about cubic surfaces can be found in [16] or [43]. Some of the history going back to Cayley and Salmon is discussed in [15].

As in Section 4.6, we consider sets of five point pairs $\mathcal{P}$ that are generic though some of our discussion below might generalize to more singular situations. Throughout this section, we use the following running example.

Example 4.7.1. *Consider the point correspondences* $\mathcal{P} = \{(\mathbf{u}_i, \mathbf{v}_i),\ i = 1, \ldots, 5\}$ *where* $\mathbf{u}_i$ *is the ith column of* U *and* $\mathbf{v}_i$ *is the ith column of* V *shown below:*

$$U = \begin{pmatrix} 0 & 0 & 1 & 1 & 2 \\ 1 & 0 & 1 & 2 & -1 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix}, \qquad V = \begin{pmatrix} 3 & 5 & -1 & -3 & 1 \\ 0 & 0 & -2 & -2 & 4 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix}$$

4.7.1 From two images to a cubic surface

Given five generic point pairs $\mathcal{P} = \{(\mathbf{u}_i, \mathbf{v}_i) \in \mathbb{P}^2_{\mathbb{R}} \times \mathbb{P}^2_{\mathbb{R}} \mid i = 1, 2, 3, 4, 5\}$, the epipolar variety $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}} \subset \mathcal{L}_{\mathcal{P}} \simeq \mathbb{P}^3_{\mathbb{R}}$ is a smooth cubic surface and its defining equation comes with a determinantal representation. Indeed, $\mathcal{R}_2 \subset \mathbb{P}^8_{\mathbb{R}}$ is defined by $\det(X) = 0$; to get the equation of the intersection with the linear space $\mathcal{L}_{\mathcal{P}} \subset \mathbb{P}^8_{\mathbb{R}}$, we compute a basis (M_0, M_1, M_2, M_3) of

$\mathcal{L}_{\mathcal{P}}$ and plug a general linear combination $z_0M_0 + z_1M_1 + z_2M_2 + z_3M_3$ of this basis into the equation $\det(X)$ to obtain

$$\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}} = \{(z_0, z_1, z_2, z_3) \in \mathbb{P}^3_{\mathbb{R}} \mid \det(M(z_0, z_1, z_2, z_3) = 0\},$$

where the entries of M are linear forms in z_0, z_1, z_2, z_3.

Example 4.7.2. *The variety $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$ is cut out by $\det M(\mathbf{z}) = 0$ where*

$$M(\mathbf{z}) = \begin{pmatrix} -z_0 - z_1 + z_2 + 4z_3 & -z_2 - z_3 & -z_2 + z_3 \\ z_0 - z_1 & z_0 + z_1 + 3z_2 & -z_0 + 2z_1 + z_2 - z_3 \\ z_0 + 3z_1 + z_2 + 2z_3 & z_2 + 5z_3 & 5z_2 - 5z_3 \end{pmatrix}.$$

The coefficient matrices M_0, M_1, M_2, M_3 in the linear matrix polynomial $M(\mathbf{z}) = z_0M_0 + z_1M_1 + z_2M_2 + z_3M_3$ form a basis of $\mathcal{L}_{\mathcal{P}}$.

4.7.2 The 27 lines on a cubic surface

The cubic surface $S = \mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}} \subset \mathbb{P}^3_{\mathbb{R}}$ contains special lines coming from the input point pairs, namely the walls $W_{\mathbf{u}_i}$ and $W^{\mathbf{v}_j}$. From Lemma 4.6.1, we know the lines $W_{\mathbf{u}_i}$ do not intersect each other and the lines $W^{\mathbf{v}_j}$ do not intersect each other. Additionally the lines $W_{\mathbf{u}_i}$ intersect the lines $W^{\mathbf{v}_j}$ in corners, but they do not all intersect pairwise. We have 10 lines in a very special configuration: It turns out that this is only possible if these 10 lines are contained in a real *Schläfli double six* on S.

Definition 4.7.1. *A Schläfli double six on a smooth cubic surface $S \subset \mathbb{P}^3_{\mathbb{R}}$ is a pair of six-tuples of lines*

$$\{(\ell_1, \ell_2, \ldots, \ell_6), (\ell'_1, \ell'_2, \ldots, \ell'_6)\}$$

on S such that the six lines in each tuple are pairwise disjoint and $\ell_i \cap \ell'_j$ is non-empty if and only if $i \neq j$.

Every smooth cubic surface contains 27 complex lines in total and they can be organized into 36 different Schläfli double six configurations. The combinatorics behind this has been studied extensively.

The fact that our cubic surfaces, that appear as epipolar varieties, always contain a Schläfli double six consisting of real lines, means that they are all of the same real topological type. Schläfli, Klein, and Zeuthen classified the real topologies of cubic surfaces. There are five possible types. The real classification of cubic surfaces is summarized in [11, Theorem 5.4] and discussed in detail in [20, pp. 40–55]. Epipolar varieties are always of type F_1 containing 27 real lines. Indeed, cubic surfaces of type F_3, F_4, or F_5 do not contain enough real lines to have a real Schläfli double six (namely only seven, three, and three). Type F_2 surfaces contain 15 real lines, which would be enough for a Schläfli double six. These surfaces are the blow-up of $\mathbb{P}^2_{\mathbb{R}}$ in four real points and one conjugate pair of complex points. The 15 lines are the four real exceptional divisors over the four real points, the strict transforms of the six real lines joining any pair of the four real points, the strict transform of the one real line joining the conjugate pair of complex points and the strict transforms of the four conics passing through all sets of five points that is complementary to one of the real points. It is straightforward to check that there are no six mutually skew lines among these 15 real lines.

Every smooth cubic surface that arises as the epipolar variety of five point pairs is the blow-up of the real projective plane $\mathbb{P}^2_{\mathbb{R}}$ in *six* real points in general linear position. (The other types are obtained by blowing up $\mathbb{P}^2_{\mathbb{R}}$ in six complex points that are invariant under complex conjugation with different numbers of fixed points, namely 0, 2, and 4. The last remaining type is not isomorphic to a blow-up of $\mathbb{P}^2_{\mathbb{R}}$ over the reals.) This is curious because we start with five point pairs that specify 10 lines on the epipolar variety, but these 10 lines determine two other lines via the unique Schläfli double six containing the 10. This prescribes one more point in every image, as we will see below. In [52], Werner shows one way to construct this sixth point in each image using the five points pairs. We discuss two other ways to explain the sixth point (see Example 4.7.3 and Example 4.7.4).

Example 4.7.3. *The four-dimensional linear space* $\mathrm{Span}\{\mathbf{v}_1\mathbf{u}_1^\top, \mathbf{v}_2\mathbf{u}_2^\top, \ldots, \mathbf{v}_5\mathbf{u}_5^\top\}$ *intersects the four-dimensional, degree six variety* $\mathcal{R}_1$ *in six points. The sixth point is* $\mathbf{v}_0\mathbf{u}_0^\top$ *where* $\mathbf{v}_0 = (-3, -12, 5)$ *and* $\mathbf{u}_0 = (18, 11, 17)^\top$.

A discussion of the Schläfli double sixes is summarized in [11, Section 2]. For a development via the blow-up construction, see [25, Chapter V, Section 4]. With this approach, it is straightforward to check that a pair of five tuples of lines $(\ell_1, \ldots, \ell_5), (\ell'_1, \ldots, \ell'_5))$ on a smooth cubic surface with $\ell_i \cap \ell_j = \emptyset$, $\ell'_i \cap \ell'_j = \emptyset$ (for all $i \neq j$), and $\ell_i \cap \ell'_j = \emptyset$ if and only if $i = j$, uniquely determines a Schläfli double six.

The Schläfli double six specified by the lines $W_{\mathbf{u}_i}$ and $W^{\mathbf{v}_j}$ for $0 \leq i, j \leq 5$ accounts for twelve of the 27 lines on the epipolar variety. For $i \neq j$, $W_{\mathbf{u}_i}$ and $W^{\mathbf{v}_j}$ span a tritangent plane, π_i^j. The tritangent plane intersects the epipolar variety in $W_{\mathbf{u}_i}$, $W^{\mathbf{v}_j}$ and one additional line

$$W_i^j := \{X \in \mathcal{R}_2 \cap L_{\mathcal{P}} : X\mathbf{u} = 0, \mathbf{v}^\top X = 0, \mathbf{u} \in \operatorname{Span}\{\mathbf{u}_i, \mathbf{u}_j\}, \mathbf{v} \in \operatorname{Span}\{\mathbf{v}_i, \mathbf{v}_j\}\}.$$

The plane π_i^j coincides with π_j^i, hence the line W_i^j coincides with W_j^i for all i, j. The distinct lines W_i^j for $0 \leq i < j \leq 5$ account for the remaining 15 real lines of the epipolar variety.

4.7.3 The determinantal representations of a cubic surface

From the classical point of view, we can start with six real points in the real projective plane $\mathbb{P}^2_{\mathbb{R}}$ and obtain a cubic surface of the correct topological type. Each of these surfaces is the epipolar variety for five generic points pairs but every cubic surface is compatible with many second images. To determine the second image, we go via a determinantal representation; different choices of a determinantal representation result in different second images (even modulo projective transformations).

Given six real points $\mathbf{u}_0, \ldots, \mathbf{u}_5$ in $\mathbb{P}^2_{\mathbb{R}}$, we get a cubic surface $S \subset \mathbb{P}^3_{\mathbb{R}}$ that is the blow up of $\mathbb{P}^2_{\mathbb{R}}$ in these six points by considering the rational map $\mathbb{P}^2_{\mathbb{R}} \dashrightarrow \mathbb{P}^3_{\mathbb{R}}$ given by the four-dimensional space of cubics that vanish at the six points. Fixing a basis c_0, c_1, c_2, c_3 of this space, the map is given by $\mathbf{x} \mapsto (c_0(\mathbf{x}), c_1(\mathbf{x}), c_2(\mathbf{x}), c_3(\mathbf{x}))$ and is defined on $\mathbb{P}^2_{\mathbb{R}} \setminus \{\mathbf{u}_0, \ldots, \mathbf{u}_5\}$. The closure of the image of this map is the cubic surface S in $\mathbb{P}^3_{\mathbb{R}}$ (which is determined up to change of coordinates on $\mathbb{P}^3_{\mathbb{R}}$ by the six points in $\mathbb{P}^2_{\mathbb{R}}$).

We cannot determine the second image from this surface alone. However, this information is specified by a determinantal representation of the surface. Determinantal representations are closely related to the *Hilbert-Burch matrix* as is explained in [11, Section 3] (see also [20]). The vanishing ideal I of six points in $\mathbb{P}^2_{\mathbb{R}}$ in general linear position is generated by the four-dimensional linear space of cubics vanishing on it. Its minimal free resolution looks like

$$0 \to F \to G \to I \to 0$$

where F is a free, graded $\mathbb{R}[x_0, x_1, x_2]$-module of rank three and the map from F to I is given by a 4×3 Hilbert-Burch matrix $L(x_0, x_1, x_2)^\top$ with entries linear in x_0, x_1, x_2. Every such matrix gives a determinantal representation of S via a simple computation: Determine the 3×3 matrix M with entries in $\mathbb{R}[x_0, x_1, x_2]_1$ such that

$$L(x_0, x_1, x_2) \begin{pmatrix} z_0 \\ z_1 \\ z_2 \\ z_3 \end{pmatrix} = M(z_0, z_1, z_2, z_3) \begin{pmatrix} x_0 \\ x_1 \\ x_2 \end{pmatrix}$$

and then $S = \{(z_0, z_1, z_2, z_3) \in \mathbb{P}^3_{\mathbb{R}} \mid \det(M(\mathbf{z})) = 0\}$. Since $M(\mathbf{z})^\top$ is also a determinantal representation of S, we get another Hilbert-Burch matrix $L'(x_0, x_1, x_2)^\top$ with linear entries such that

$$L'(x_0, x_1, x_2) \begin{pmatrix} z_0 \\ z_1 \\ z_2 \\ z_3 \end{pmatrix} = [M(z_0, z_1, z_2, z_3)]^\top \begin{pmatrix} x_0 \\ x_1 \\ x_2 \end{pmatrix}.$$

Therefore a determinantal representation determines six real points $\mathbf{u}_i$ in the first image cut out by the 3×3 minors of L and six real points $\mathbf{v}_i$ in the second image cut out by the 3×3 minors of L'.

Example 4.7.4. *The Hilbert Burch matrix of $M(\mathbf{z})$ is*

$$L = \begin{pmatrix} -x_0 & -x_0 & x_0 - x_1 - x_2 & 4\,x_0 - x_1 + x_2 \\ x_0 + x_1 - x_2 & -x_0 + x_1 + 2\,x_2 & 3\,x_1 + x_2 & -x_2 \\ x_0 & 3\,x_0 & x_0 + x_1 + 5\,x_2 & 2\,x_0 + 5\,x_1 - 5\,x_2 \end{pmatrix}$$

while the Hilbert-Burch matrix of $M(\mathbf{z})^\top$ is

$$L' = \begin{pmatrix} -x_0 + x_1 + x_2 & -x_0 - x_1 + 3\,x_2 & x_0 + x_2 & 4\,x_0 + 2\,x_2 \\ x_1 & x_1 & -x_0 + 3\,x_1 + x_2 & -x_0 + 5\,x_2 \\ -x_1 & 2\,x_1 & -x_0 + x_1 + 5\,x_2 & x_0 - x_1 - 5\,x_2 \end{pmatrix}.$$

We can compute the six real points in each image from the Hilbert-Burch matrices, because the 3×3 minors vanish exactly at these points. The zeros of these ideals of minors are $\mathbf{u}_0, \ldots, \mathbf{u}_5$ where $\mathbf{u}_0 = (18, 11, 17)^\top$ and $\mathbf{v}_0, \ldots, \mathbf{v}_5$ where $\mathbf{v}_0 = (-3, -12, 5)^\top$, respectively. This gives yet another way to compute $\mathbf{u}_0$ and $\mathbf{v}_0$.

4.7.4 From a determinantal representation to epipoles

Given a determinantal representation $M(\mathbf{z})$ of a cubic surface $S \subset \mathbb{P}^3_{\mathbb{R}}$, consider the map $S \to \mathbb{P}^2_{\mathbb{R}} \times \mathbb{P}^2_{\mathbb{R}}$, $\mathbf{z} \mapsto \operatorname{adj}(M(\mathbf{z}))$. Since the matrix $M(\mathbf{z})$ has rank two for every point $\mathbf{z} \in \mathbb{P}^3_{\mathbb{R}}$, the image lies indeed inside $\mathbb{P}^2_{\mathbb{R}} \times \mathbb{P}^2_{\mathbb{R}}$ embedded via the standard Segre embedding in $\mathbb{P}^8_{\mathbb{R}}$. With its given determinantal representation, the image of the cubic surface S under the adjoint map is the space of epipoles in two images consistent with $\mathcal{P}$.

From our study of the chiral epipolar region in Section 4.6, we are interested in the lines $W_{\mathbf{u}_i}$ and $W^{\mathbf{v}_j}$ on the cubic surface corresponding to $\mathcal{P}$. We compute the images of these lines under the adjoint map, connecting our work with Werner's results in epipole space [52]. The image of $W_{\mathbf{u}_i}$ under the adjoint map is $\{\mathbf{u}_i\} \times C^i$, where $C^i \subset \mathbb{P}^2_{\mathbb{R}}$ is a curve of degree two, namely the conic passing through $\mathbf{v}_j$ for $i \neq j$. Indeed, for every matrix $X \in W_{\mathbf{u}_i}$, the vector $\mathbf{u}_i$ is in the right kernel, by definition, whereas the left kernel varies. This left kernel varies along a conic in $\mathbb{P}^2_{\mathbb{R}}$ because every entry of the adjoint matrix of a 3×3 matrix is a quadric. So intersecting the image of $W_{\mathbf{u}_i}$ under the adjoint map with a line in $\mathbb{P}^2_{\mathbb{R}}$ translates, by pulling back via the adjoint map, to the intersection of a quadric hypersurface in $\mathbb{P}^3_{\mathbb{R}}$ with the line $W_{\mathbf{u}_i}$, which therefore consists of two intersection points. Symmetrically, the image of $W^{\mathbf{v}_j}$ is $C_j \times \{\mathbf{v}_j\}$, where $C_j \subset \mathbb{P}^2_{\mathbb{R}}$ is the conic passing through $\mathbf{u}_j$ for $i \neq j$. This behaviour is special and occurs only for the walls in the cubic surface. The image of the other lines W_i^j under the adjoint is $\operatorname{Span}\{\mathbf{u}_i, \mathbf{u}_j\} \times \operatorname{Span}\{\mathbf{v}_i, \mathbf{v}_j\}$.

The conics C_i in the first image and C^j in the second image are described in Section 4 of [52]. These conics divide each image into cells of possible epipoles with piecewise conic boundaries. Each cell is uniquely characterized by the subset of conics that participate to form its boundary or equivalently by the subset of points that belong to its boundary. Werner's test for the existence of a chiral reconstruction involves looking for "allowed segments" of the conics C_i and C^j. In our work, we translate this question to studying the intersections of lines, i.e., the $(\mathbf{u}_i, \mathbf{v}_j)$ corners in the cubic surface in the preimage of the adjoint. In the following example, we show how our chirality test on corners relates to Werner's allowed segments in the space of epipoles.

Example (4.6.1 continued). *Consider the five point pairs from* (4.17)*:*

$$U = \begin{pmatrix} 0 & 0 & 4 & 2 & 2 \\ 0 & 4 & 0 & 1 & 3 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix} \quad \text{and} \quad V = \begin{pmatrix} 2 & 2 & 4 & 0 & 4 \\ 1 & 3 & 0 & 4 & 4 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix}.$$

We find that the same sign condition of Corollary 4.6.1 holds at the following corners: $(\mathbf{u}_2, \mathbf{v}_3)$*,* $(\mathbf{u}_2, \mathbf{v}_4)$*,* $(\mathbf{u}_3, \mathbf{v}_1)$*,* $(\mathbf{u}_3, \mathbf{v}_1)$*,* $(\mathbf{u}_4, \mathbf{v}_1)$*, and* $(\mathbf{u}_4, \mathbf{v}_3)$*. These corners are in the boundary of the chiral epipolar region of* $\mathcal{P}$ *which lives in the cubic surface* $\mathcal{R}_2 \cap \mathcal{L}_{\mathcal{P}}$*. Blowing down with the adjoint map, we get the non-empty region of epipoles in* $\mathbb{P}^2_{\mathbb{R}} \times \mathbb{P}^2_{\mathbb{R}}$ *that correspond to a chiral reconstruction. The boundary of the region in the first image is defined by* C_3 *and* C_4 *at* $\mathbf{u}_2$*,* C_1 *and* C_4 *at* $\mathbf{v}_3$*, and* C_1 *and* C_3 *at* $\mathbf{u}_4$*. The boundary of the region in the second image is defined by* C^3 *and* C^4 *at* $\mathbf{v}_1$*,* C^2 *and* C^4 *at* $\mathbf{v}_3$*, and* C^2 *and* C^3 *at* $\mathbf{v}_4$*. See Figure 4.4.*

We remark that our test for chirality requires more than just the cubic surface S determined by six real points $\mathbf{u}_i$ in one image $\mathbb{P}^2_{\mathbb{R}}$. For different determinantal representations of S, the six points in the other image may differ. In order to check the chiral epipolar inequalities, it is necessary to fix the second image, for example, by fixing a determinantal representation of S.

It is also necessary to specify five out of the six point pairs to test for chirality. As the next example shows, while some sets of five point pairs may have a chiral reconstruction,

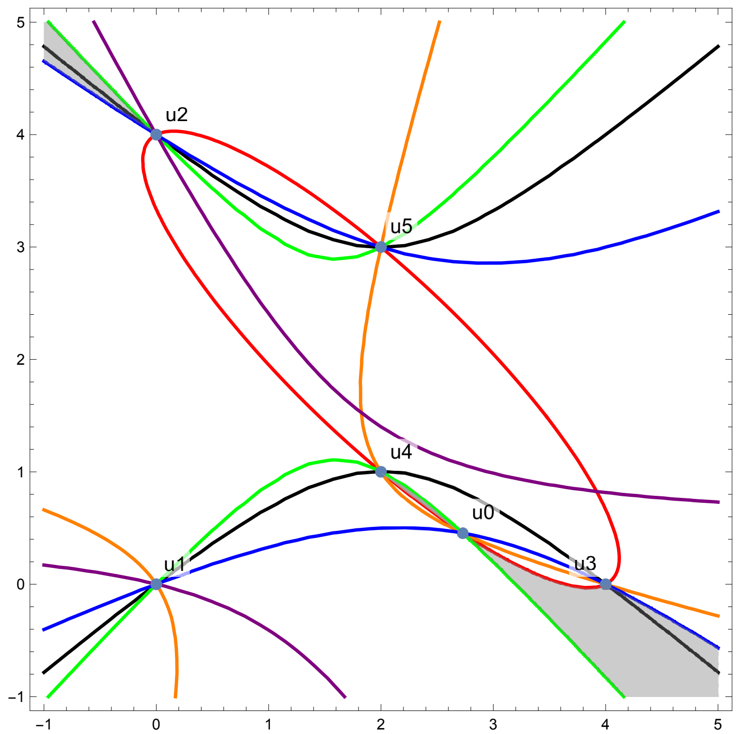

(a) Shaded area represents region of allowed epipoles $\mathbf{e}_1$.

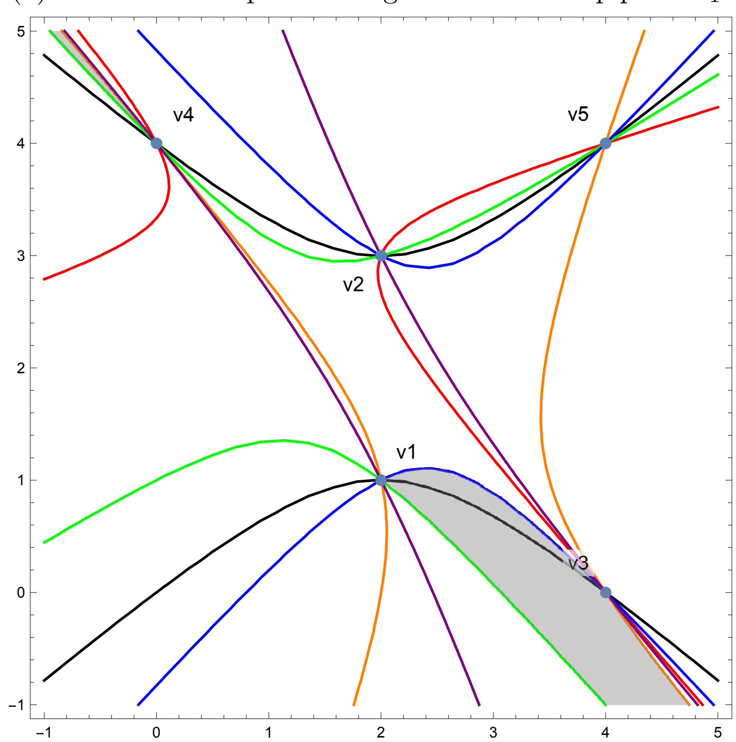

(b) Shaded area represents region of allowed epipoles $\mathbf{e}_2$.

Figure 4.4

another set associated with the same determinantal representation may not.

Example (**4.6.1 continued**). *Consider the five point pairs from* (4.16)*:*

$$U = \begin{pmatrix} 0 & 0 & 4 & 2 & 2 \\ 0 & 4 & 0 & 1 & 3 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix} \qquad V = \begin{pmatrix} 2 & 2 & 4 & 0 & 1 \\ 1 & 3 & 0 & 4 & 1 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix}.$$

We compute $\mathbf{u}_0 = (504, 300, 281)^\top$ *and* $\mathbf{v}_0 = (68, 300, 97)^\top$. *By computing the relevant products* D_{lm}, D_{ln}, D_{mn} *at the 20 corners, we find that* $\{(\mathbf{u}_i, \mathbf{v}_i)\}_{i=1}^5$ *does not have a chiral reconstruction. However, suppose we consider a different subset of five point pairs*

$$\widehat{\mathcal{P}}_5 = \{(\mathbf{u}_i, \mathbf{v}_i) : i \in \{0, 1, 2, 3, 4\}\}$$

as the original five point pairs. Some corners satisfy the same sign condition of Corollary 4.6.1, hence $\widehat{\mathcal{P}}_5$ *has a chiral reconstruction. In fact, the subset of five point pairs* $\widehat{\mathcal{P}}_i$ *which omits index* i *has a chiral reconstruction when* $i = 1, 2, 3, 5$ *and has no chiral reconstruction when* $i = 0, 4$. *The cubic surface* $\mathcal{R}_2 \cap L_{\widehat{\mathcal{P}}_i}$ *is the same for all* i *and the determinantal representations are all equivalent, but whether a chiral reconstruction exists or not depends on which five point pairs we choose.*

BIBLIOGRAPHY

[1] Sameer Agarwal, Hon-Leung Lee, Bernd Sturmfels, and Rekha R. Thomas. On the existence of epipolar matrices, 2015.

[2] Sameer Agarwal, Andrew Pryhuber, Rainer Sinn, and Rekha R. Thomas. The chiral domain of a camera arrangement, 2021.

[3] Sameer Agarwal, Andrew Pryhuber, and Rekha R Thomas. Ideals of the multiview variety. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(4):1279–1292, 2021.

[4] Chris Aholt, Sameer Agarwal, and Rekha Thomas. A QCQP approach to triangulation. In *European Conference on Computer Vision*, pages 654–667. Springer, 2012.

[5] Chris Aholt, Bernd Sturmfels, and Rekha Thomas. A Hilbert scheme in computer vision. *Canadian Journal of Mathematics*, 65(5):961–988, 2013.

[6] MD Atkinson. Primitive spaces of matrices of bounded rank. II. *Journal of the Australian Mathematical Society*, 34(3):306–315, 1983.

[7] Grigoriy Blekherman, Pablo A. Parrilo, and Rekha R. Thomas. *Semidefinite Optimization and Convex Algebraic Geometry*. SIAM, 2012.

[8] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge university press, 2004.

[9] Joel G Broida and Stanley Gill Williamson. *A Comprehensive Introduction to Linear Algebra*. Addison-Wesley Redwood City, CA, 1989.

[10] J. W. Bruce and C. T. C. Wall. On the classification of cubic surfaces. *J. London Math. Soc. (2)*, 19(2):245–256, 1979.

[11] Anita Buckley and Tomaž Košir. Determinantal representations of smooth cubic surfaces. *Geometriae Dedicata*, 125(1):115–140, 2007.

[12] Ondřej Chum, Tomáš Werner, and Tomáš Pajdla. Joint orientation of epipoles. In *Proceedings of the British Machine Vision Conference*, 2003.

[13] Aldo Conca, Emanuela De Negri, and Elisa Gorla. Cartwright-sturmfels ideals associated to graphs and linear spaces. *arXiv preprint arXiv:1705.00575*, 2017.

[14] David A. Cox, John Little, and Donal O'Shea. *Ideals, Varieties, and Algorithms.* Springer, 4th edition, 2015.

[15] Igor V. Dolgachev. Luigi Cremona and cubic surfaces. In *Luigi Cremona (1830–1903) (Italian)*, volume 36 of *Incontr. Studio*, pages 55–70. Istituto Lombardo di Scienze e Lettere, Milan, 2005.

[16] Igor V. Dolgachev. *Classical Algebraic Geometry.* Cambridge University Press, Cambridge, 2012. A modern view.

[17] David Eisenbud and Joe Harris. Vector spaces of matrices of low rank. *Advances in Mathematics*, 70(2):135–155, 1988.

[18] Olivier Faugeras, Quang-Tuan Luong, and T. Papadopoulou. *The Geometry of Multiple Images: The Laws that Govern the Formation of Images of a Scene and Some of Their Applications.* MIT Press, Cambridge, MA, USA, 2001.

[19] Olivier Faugeras and Bernard Mourrain. On the geometry and algebra of the point and line correspondences between n images. In *Computer Vision, 1995. Proceedings., Fifth International Conference on*, pages 951–956. IEEE, 1995.

[20] Anthony V. Geramita. Lectures on the nonsingular cubic surface in $\mathbf{P}^3$. In *The Curves Seminar at Queen's, Vol. VI (Kingston, ON, 1989)*, volume 83 of *Queen's Papers in Pure and Appl. Math.*, pages Exp. No. A, 106. Queen's Univ., Kingston, ON, 1989. With appendixes by the author and E. Kani and by Robert McCann.

[21] Daniel R Grayson and Michael E Stillman. Macaulay 2, a software system for research in algebraic geometry, 2002.

[22] J. Harris. *Algebraic Geometry: A First Course.* Graduate Texts in Mathematics. Springer-Verlag New York, 1995.

[23] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision.* Cambridge University Press, second edition, 2004.

[24] Richard I. Hartley. Chirality. *Int. J. Comput. Vision*, 26(1):41–61, January 1998.

[25] Robin Hartshorne. *Algebraic Geometry; Graduate Texts in Mathematics.* Springer, 1977.

[26] Archibald Henderson. *The Twenty-Seven Lines Upon the Cubic Surface.* Reprinting of Cambridge Tracts in Mathematics and Mathematical Physics, No. 13. Hafner Publishing Co., New York, 1960.

[27] A. Heyden and K. Aström. Algebraic Properties of Multilinear Constraints. *Mathematical Methods in the Applied Sciences*, 20:1135–1162, September 1997.

[28] D. Hilbert and S. Cohn-Vossen. *Geometry and the Imagination.* Chelsea Publishing Company, New York, N. Y., 1952. Translated by P. Neményi.

[29] Fredrik Kahl and Didier Henrion. Globally optimal estimates for geometric reconstruction problems. *International Journal of Computer Vision*, 74(3), 2007.

[30] Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2000/01.

[31] Stéphane Laveau and Olivier Faugeras. Oriented projective geometry for computer vision. In *Proceedings of the European Conference on Computer Vision*. Springer, 1996.

[32] Hon-Leung Lee. On the existence of a projective reconstruction. *CoRR*, abs/1608.05518, 2016.

[33] Binglin Li. Images of rational maps of projective spaces. *International Mathematics Research Notices*, 2018(13):4190–4228, 2018.

[34] H. Christopher Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293(5828):133, 1981.

[35] Yi Ma, Stefano Soatto, Jana Kosecka, and S Shankar Sastry. *An Invitation to 3-d Vision: From Images to Geometric Models*. Springer, 2012.

[36] Murray Marshall. *Positive polynomials and sums of squares*, volume 146 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2008.

[37] Stephen Maybank. *Theory of Reconstruction from Image Motion*. Springer-Verlag, 1993.

[38] David Mumford. *The Red Book of Varieties and Schemes*, volume 1358. 1996.

[39] David Nistér. Untwisting a projective reconstruction. *International Journal of Computer Vision*, 60(2):165–183, 2004.

[40] David Nistér and Frederik Schaffalitzky. Four points in two or three calibrated views: Theory and practice. *International Journal of Computer Vision*, 67(2):211–231, 2006.

[41] Pablo A Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical programming*, 96(2):293–320, 2003.

[42] Andrew Pryhuber, Rainer Sinn, and Rekha R Thomas. On the existence of two view chiral reconstructions. *arXiv preprint arXiv:2011.07197*, 2020.

[43] Miles Reid. *Undergraduate Algebraic Geometry*, volume 12 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1988.

[44] B. Segre. *The Non-singular Cubic Surfaces*. Oxford University Press, Oxford, 1942.

[45] J. Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computations*. Academic Press, 1991.

[46] Matthew Trager, Martial Hebert, and Jean Ponce. The joint image handbook. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.

[47] Matthew Trager, Martial Hebert, and Jean Ponce. The joint image handbook supplementary material. In *Proceedings of the IEEE International Conference on Computer Vision*, December 2015.

[48] Matthew Trager, Brian Osserman, and Jean Ponce. On the solvability of viewing graphs. In *European Conference on Computer Vision*, 2018.

[49] B Triggs. The geometry of projective reconstruction I: Matching constraints and the joint image. submitted to. *Int. J. Computer Vision*, 1995.

[50] B. Triggs. Matching constraints and the joint image. In *IEEE International Conference on Computer Vision*, pages 338–343, June 1995.

[51] Tomáš Werner. Combinatorial constraints on multiple projections of a set of points. In *ICCV*, pages 1011–1016, 2003.

[52] Tomáš Werner. Constraint on five points in two images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2003.

[53] Tomáš Werner and Tomáš Pajdla. Cheirality in epipolar geometry. In *Proceedings of the IEEE International Conference on Computer Vision*, 2001.

[54] Tomáš Werner and Tomáš Pajdla. Oriented matching constraints. In *Proceedings of the British Machine Vision Conference*, 2001.