SIAM J. APPL. ALGEBRA GEOMETRY
Vol. 4, No. 2, pp. 297–321

Two Hilbert Schemes in Computer Vision*

Max Lieblich† and Lucas Van Meter†

Abstract. We study multiview moduli problems that arise in computer vision. We show that these moduli spaces are always smooth and irreducible, in both the calibrated and uncalibrated cases, for any number of views. We also show that these moduli spaces always admit open immersions into Hilbert schemes for more than two views, extending and refining work of Aholt, Sturmfels, and Thomas. We use these moduli spaces to study and extend the classical twisted pair covering of the essential variety.

Key words. multivew geometry, Hilbert scheme, computer vision, moduli theory, deformation theory

AMS subject classifications. 14D23, 14B10, 14C05, 68T45

DOI. 10.1137/18M1200117

1. Introduction.

In this paper, we discuss a functorial approach to multiview geometry, a subfield of computer vision. The literature on multiview geometry is vast, although this is the first attempt that we know of to use the techniques of modern functorial algebraic geometry to approach the subject. As we hope to demonstrate here and elsewhere, this approach has a great deal of promise. A beautiful introduction to the subject can be found in [6]. Earlier versions of this paper (available as arXiv preprints) also contain a condensed introduction to the subject suitable for algebraic geometers.

1.1. Our results.

The main result of this paper is the following, proven in sections 3 and 4.

Theorem 1.1. *There are smooth irreducible varieties* Cam_n *and* CalCam_n *parametrizing* n*-view camera configurations and* n*-view calibrated camera configurations, respectively.*

1. *The variety* Cam_n *has dimension* $11n - 15$*. For all* $n > 1$*, sending a configuration to its joint image defines a locally closed embedding*

$$\mathrm{Cam}_n \hookrightarrow \mathrm{Hilb}_{(\mathbf{P}^2)^n}\,.$$

If $n > 2$*, then this morphism is an open immersion, so that* Cam_n *is identified with an open subscheme of the smooth locus of* $\mathrm{Hilb}_{(\mathbf{P}^2)^n}$*.*

2. *The variety* CalCam_n *has dimension* $6n - 7$*. For all* $n > 1$*, there is a natural locally closed embedding*

$$\mathrm{CalCam}_n \hookrightarrow \mathrm{Hilb}_{C_1 \times \cdots \times C_n \subset (\mathbf{P}^2)^n}$$

*Received by the editors July 12, 2018; accepted for publication (in revised form) November 7, 2019; published electronically May 20, 2020.
https://doi.org/10.1137/18M1200117
Funding: The authors were partially supported by NSF CAREER grants DMS-1056129 and DMS-1600813.
†Mathematics, University of Washington, Seattle, WA 98195 (lieblich@uw.edu, lucasvm@uw.edu).

(where the latter is a diagram Hilbert scheme; see subsection 3.3). If $n > 2$, then this morphism is an open immersion.

3. *The natural decalibration morphism $\nu_n : \mathrm{CalCam}_n \to \mathrm{Cam}_n$ is finite, proper, and unramified. The morphism ν_2 is an étale cover of its image with general fiber of order 2. For $n > 2$ the morphism ν_n is generically injective but not injective.*

The statements on Hilbert schemes generalize and refine the results of [1]. In particular, our methods show that the formation of the multiview variety gives an open immersion into the Hilbert scheme at all points, identifying the moduli space with an open subscheme of the Hilbert scheme.

1.2. Methodological contributions. There are a few basic principles that set this work apart from other work on multiview geometry.

1. The *functorial method*, common in modern algebraic geometry, gives us insight into the intrinsic geometry of natural moduli problems growing out of the classical constructions. While [1] uses the Geometric Invariant Theory (GIT) quotient to construct the moduli of uncalibrated camera configurations, this method does not obviously generalize to a construction for calibrated cameras. Additionally, by developing the functorial theory of cameras we hope to make the field of multiview geometry accessible to a wider audience in pure mathematics.
2. The *geometric view of calibration* via calibration data gives us insight into the structure of the space of calibrated cameras in a way that seems not to have been considered before. In particular, by restricting camera configurations to morphisms between calibrating conics, we get a fibration structure on the moduli space of calibrated camera configurations that is quite useful for studying the moduli space. In subsection 3.4, there's a third Hilbert scheme—the Hilbert scheme of the product of calibrating conics—that is the base of this fibration. This way of thinking about calibration can also be used to understand the essential variety in new ways. In [12], this is used to reproduce results of both [2] and [3] (which itself used the results of [2]) from first principles, among other things.
3. The use of *diagram Hilbert schemes* allows us to treat the case of calibrated cameras similarly to how uncalibrated cameras are treated in [1]. Instead of closed subschemes, as were used for the calibrated case, we use a type of flag to keep track of the calibration data. This transparently recovers the result that the moduli space is open in a Hilbert scheme.

This paper also opens up many new lines of inquiry and leaves many questions unanswered. We discuss a few of these questions in section 5.

2. The algebraic geometry of pinhole cameras. In this section we review the basic theory of pinhole cameras, with a geometric emphasis. We include a canonical treatment of calibrated cameras with a greater focus on the geometry of the calibrating conics. For the sake of clarity, we focus in subsections 2.1 and 2.2 on the geometry over an algebraically closed field. In subsection 2.3 we study what happens over a general base scheme, as a preparation for the study of moduli and deformation theory in section 3.

2.1. Basic definitions.

Definition 2.1. *A* pinhole camera *is a surjective rational map* $\varphi : \mathbf{P}^3 \dashrightarrow \mathbf{P}^2$ *given by three linearly independent sections of* $\mathscr{O}_{\mathbf{P}^3}(1)$*. The* center *of the camera is the unique point* $p \in \mathbf{P}^3$ *at which* φ *is undefined.*

Definition 2.2. *A* calibrated plane *is a pair* $(\mathbf{P}^2, D)$ *with* D *a smooth conic.*

Definition 2.3. *A* calibration datum *for a pinhole camera* φ *is a pair of planar degree* 2 *curves* $C \subset \mathbf{P}^3$ *and* $D \subset \mathbf{P}^2$ *such that* D *is a smooth conic and the restriction* $\varphi_C : C \dashrightarrow \mathbf{P}^2$ *factors through the inclusion* $D \subset \mathbf{P}^2$*.*

If C *is smooth, the calibration datum will be called* smooth *or* nondegenerate*; otherwise, it will be called* degenerate*. If a calibrated plane* $(\mathbf{P}^2, D)$ *is fixed, a* relative calibration datum *for a pinhole camera* $\mathbf{\Phi}$ *is a curve* $C \subset \mathbf{P}^3$ *such that* (C, D) *is a calibration datum for* $\mathbf{\Phi}$*.*

Remark 2.4. If C is smooth, then it follows from the linearity of the camera projection that $\mathbf{\Phi}$ must map C isomorphically to D, and that the center of $\mathbf{\Phi}$ is not contained in the plane spanned by C. If C is degenerate, it must be a divisor-theoretic sum of two lines on the quadric cone in $\mathbf{P}^3$ generated by D under the projection $\mathbf{\Phi}$ (i.e., a union of two distinct rulings or a double ruling). When the two cone points are distinct (i.e., the configuration is general), a union of two distinct rulings cannot occur as a limit of calibration data.

Remark 2.5. A given camera with calibrated image plane $(\mathbf{P}^2, D)$ has infinitely many relative calibration data: one can take any plane section of the quadric cone in $\mathbf{P}^3$ lying over D. Once we look at configurations of two or more cameras, there will be at most two calibration data (smooth or degenerate). This is described at length in subsection 4.1.2.

Degenerate calibrations give us closures of natural moduli spaces, including the closure of the classical twisted pair moduli space $\mathrm{SO}(3) \times \mathbf{P}^2$ to a finite étale cover of the essential variety described in subsection 4.2. Imagining the system of plane sections of the cone over D, one readily sees that degenerate calibration data arise as limits of smooth calibration data.

Definition 2.6. *A* calibrated camera *is a pair* $(\varphi, (C, D))$ *where* φ *is a pinhole camera and* (C, D) *is a calibration datum for* φ*.*

Remark 2.7. In the classical literature, a camera is called *calibrated* (or sometimes *normalized*) when it takes the absolute conic to the Euclidean conic: more precisely, we can endow $\mathbf{P}^3$ with coordinates x, y, z, w and $\mathbf{P}^2$ with coordinates X, Y, Z, and then we take the curves C and D to be given by the equations $\{w = 0, x^2 + y^2 + z^2 = 0\}$ and $\{X^2 + Y^2 + Z^2 = 0\}$, respectively. Note that any camera as described here with a smooth calibration datum can be transformed to a classically calibrated camera by applying suitable automorphisms to $\mathbf{P}^3$ and $\mathbf{P}^2$. (This is not unique.) The degenerate calibrations cannot.

There are two reasons to use this more flexible approach.

(1) It leads to the "right definition" of the moduli space of calibrated camera configurations (subsection 3.4).
(2) By always forcing the absolute conic to map to the Euclidean conic, one makes it impossible to study modular boundary points where the absolute conic is flattened until it collapses (yielding degenerate calibrations). As we will describe below, these degenerate calibrations give geometrically meaningful compactifications of the space

of calibrated camera configurations.

2.2. Multiview configurations. In this section, we describe some of the geometry attached to a collection of cameras with distinct centers.

2.2.1. Uncalibrated cameras.

Definition 2.8. *A* multiview configuration *is a collection of cameras*

$$\varphi_1, \ldots, \varphi_n : \mathbf{P}^3 \dashrightarrow \mathbf{P}^2.$$

Notation 2.9. We will generally use $\mathbf{\Phi} : \mathbf{P}^3 \dashrightarrow (\mathbf{P}^2)^n$ to denote a multiview configuration, writing $\mathbf{\Phi}_i = \mathrm{pr}_i \circ \mathbf{\Phi}$ for its components when necessary. The *length* of $\mathbf{\Phi}$ is the number of cameras; we will denote it $\mathrm{len}(\mathbf{\Phi})$. Write $\mathrm{Center}(\mathbf{\Phi}) \subset \mathbf{P}^3$ for the tuple of camera centers. Write $\pi : \mathrm{Res}(\mathbf{\Phi}) \to \mathbf{P}^3$ for the blowup of $\mathbf{P}^3$ at the reduced closed subscheme supported at the camera centers; if two cameras have the same center, we only count it once. Given an index i, let E_i denote the exceptional divisor over the ith camera center, with canonical inclusion $\iota_i : E_i \hookrightarrow \mathrm{Res}(\mathbf{\Phi})$. By the previous convention, this means that there can be $i \neq j$ for which $E_i = E_j$.

Definition 2.10. *A multiview configuration $\mathbf{\Phi}$ is* general *if the camera centers are all distinct. It is* noncollinear *if the camera centers do not all lie on a single line, and* collinear *otherwise.*

Definition 2.11. *An* isomorphism *between multiview configurations $\mathbf{\Phi}^1$ and $\mathbf{\Phi}^2$ of common length n is an automorphism $\varepsilon : \mathbf{P}^3 \to \mathbf{P}^3$ fitting into a commutative diagram*

$$\begin{tikzcd}
\mathbf{P}^3 \arrow[dd, "\varepsilon"'] \arrow[dr, dashed, "\mathbf{\Phi}^1"] & \\
& (\mathbf{P}^2)^n \\
\mathbf{P}^3 \arrow[ur, dashed, "\mathbf{\Phi}^2"'] &
\end{tikzcd}$$

Lemma 2.12. *Let Y be a scheme, and let $(\mathscr{L}, s_0, \ldots, s_n)$ be an invertible sheaf with n sections. If Z is the zero scheme of $s_0, \ldots, s_n$, then the rational map induced by this linear series extends uniquely to a morphism $\mathrm{Bl}_Z Y \to \mathbf{P}^n$.*

Proof. By definition the sections $s_0, \ldots, s_n$ define a surjection

$$\mathscr{O}_Y^{n+1} \twoheadrightarrow \mathscr{L} \otimes \mathscr{I}_Z,$$

which extends to a surjective map of $\mathscr{O}_Y$-algebras

$$\mathrm{Sym}^*(\mathscr{L}^\vee)^{\oplus n+1} \twoheadrightarrow \bigoplus \mathscr{I}^n.$$

The induced map on relative **Proj** constructions gives the desired morphism. ■

Proposition 2.13. *Given a multiview configuration* $\mathbf{\Phi}$, *there is a unique commutative diagram*

$$\begin{tikzcd} \mathrm{Res}(\mathbf{\Phi}) \arrow[dr, "\rho"] & \\ & (\mathbf{P}^2)^{\mathrm{len}(\mathbf{\Phi})} \\ \mathbf{P}^3 \arrow[uu, dashed, "\pi^{-1}"] \arrow[ur, dashed, "\mathbf{\Phi}"'] & \end{tikzcd}$$

The diagram has the property that for each i, *the composition*

$$E_i \xrightarrow{\iota_i} \mathrm{Res}(\mathbf{\Phi}) \xrightarrow{\rho} (\mathbf{P}^2)^{\mathrm{len}(\Phi)} \xrightarrow{\mathrm{pr}_i} \mathbf{P}^2$$

is an isomorphism.

Proof. Lemma 2.12 shows the existence and uniqueness of the desired diagram. To check that the composition is an isomorphism on exceptional divisors, one can see that each map is locally isomorphic to the morphism $\mathrm{Bl}_0\,\mathbf{A}^3 \to \mathbf{P}^2$ that resolves the canonical presentation $\mathbf{A}^3 \setminus \{0\} \to \mathbf{P}^2$, and here one can simply check that the induced map from the exceptional divisor to the plane is an isomorphism. We omit the details. ■

2.2.2. Calibrated cameras. When the cameras are adorned with calibration data, we track these data through the diagrams.

Definition 2.14. *Given a multiview configuration* $\mathbf{\Phi} : \mathbf{P}^3 \dashrightarrow (\mathbf{P}^2)^n$, *a* multiview calibration datum *is a pair* $(C, (C_1, \ldots, C_n))$ *such that for each* $i = 1, \ldots, n$ *the pair* (C, C_i) *is a calibration datum for* $\mathbf{\Phi}_i$. *Given a tuple of calibrated planes* $(\mathbf{P}^2, C_i)$ *for* $i = 1, \ldots, n$, *a* relative calibration datum *for* $\mathbf{\Phi}$ *is a curve* $C \subset \mathbf{P}^3$ *such that* $(C, (C_1, \ldots, C_n))$ *is a calibration datum for* $\mathbf{\Phi}$.

Notation 2.15. We will write $\mathbf{C}$ for a calibration datum $(C, (C_i))$, and then $\mathbf{C}_0 = C$ and $\mathbf{C}_i = C_i$ for $i = 1, \ldots, n$.

Notation 2.16. A calibrated multiview configuration $(\mathbf{\Phi}, \mathbf{C})$ will be called *nondegenerate* if the calibration datum is nondegenerate.

Definition 2.17. *An* isomorphism *between multiview configurations with calibration data* $(\mathbf{\Phi}^1, \mathbf{C}^1)$ *and* $(\mathbf{\Phi}^2, \mathbf{C}^2)$ *of common length* n *is an isomorphism* $\varepsilon : \mathbf{\Phi}^1 \to \mathbf{\Phi}^2$ *of multiview configurations as in Definition 2.11 such that* $\varepsilon(\mathbf{C}_0^1) = \mathbf{C}_0^2$ *and such that for* $i = 1, \ldots, n$ *we have* $\mathbf{C}_i^1 = \mathbf{C}_i^2$.

2.2.3. A characterization of isomorphic general configurations. In this section we briefly consider when two multiview configurations $\mathbf{\Phi}^1$ and $\mathbf{\Phi}^2$ are isomorphic (and similarly when they are endowed with calibration data). This will play a role in studying a particular map from the moduli space to Hilbert schemes in later sections of this paper.

Definition 2.18. *Given a multiview configuration* $\mathbf{\Phi}$, *the* associated multiview scheme, *also known as the* joint image [1, 16], *is the scheme-theoretic image of the resolution* $\mathrm{Res}(\mathbf{\Phi})$ *under the canonical extension* ρ *of Proposition 2.13. It is denoted* $\mathrm{Sch}(\mathbf{\Phi})$. *Working over a field (as* *we temporarily are here), the multiview scheme is a variety, and is called the "multiview variety" in* [1].

In the following, an *n-term flag of schemes* will be a sequence of closed immersions

$$X_0 \hookrightarrow X_1 \hookrightarrow X_2 \hookrightarrow \cdots \hookrightarrow X_{n-1}.$$

Definition 2.19. *Given a calibrated multiview configuration* $(\boldsymbol{\Phi}, C)$ *with calibrated image planes* $(\mathbf{P}^2, C_i)$, $i = 1, \ldots, n$, *the* associated multiview flag, *denoted* $\mathrm{Flag}(\boldsymbol{\Phi}, C)$, *is the* 2*-term flag of schemes* $C \subset \mathrm{Sch}(\boldsymbol{\Phi})$ *contained in* $C_1 \times \cdots \times C_n \subset (\mathbf{P}^2)^n$.

As we will gradually see, the following lemma is the key result connecting the abstract moduli problems we study here to Hilbert schemes.

Lemma 2.20. *The canonical map* $\mathscr{O}_{\mathrm{Sch}(\boldsymbol{\Phi})} \to \mathbf{R}\,\rho_* \mathscr{O}_{\mathrm{Res}(\boldsymbol{\Phi})}$ *is a quasi-isomorphism. Equivalently, the canonical map* $\rho^\sharp : \mathscr{O}_{(\mathbf{P}^2)^n} \to \rho_* \mathscr{O}_{\mathrm{Res}(\boldsymbol{\Phi})}$ *is an isomorphism, and all higher direct images* $\mathbf{R}^i\,\rho_* \mathscr{O}_{\mathrm{Res}(\boldsymbol{\Phi})}$ *(with* $i > 0$*) vanish.*

Proof. For the first statement, note that $\rho_* \mathscr{O}_{\mathrm{Res}(\boldsymbol{\Phi})}$ is a finite $\mathscr{O}_{(\mathbf{P}^2)^n}$-algebra by properness. Moreover, since every nonempty fiber of ρ is geometrically integral (it being an intersection of lines, and hence either a point or a line), we see that $\rho^\sharp$ is surjective after base change to any point of $(\mathbf{P}^2)^n$. By Nakayama's lemma, $\rho^\sharp$ is surjective.

Now we show that the higher direct images vanish. By the Theorem on Formal Functions [4, Théorème 4.1.5], the completion of $\mathbf{R}^i\,\rho_* \mathscr{O}$ at a point p is isomorphic to $\lim \mathrm{H}^i(X_m, \mathscr{O}_{X_m})$, where X_m is the mth infinitesimal neighborhood of the fiber of ρ over p. When the fiber is empty or a point, this vanishes. The only interesting case is the unique singular point that is the image of the strict transform of the line through all camera centers in the collinear case. Note that $\mathscr{O}_{X_m}$ is filtered by subquotients that are symmetric powers of the ideal sheaf $\mathscr{I}_{X_0}$ restricted to X_0. Given a line L in $\mathbf{P}^3$, we have that $\mathscr{I}_L|_L \cong \mathscr{O}_L(-1)^{\oplus 2}$. For each point on L that we blow up, the ideal sheaf gets twisted by 1 (functions from $\mathbf{P}^3$ vanish to extra order on the strict transform along the intersection with the exceptional divisor). In fact, if we are blowing up n points, we have that $\mathscr{I}_{X_0}|_{X_0} \cong \mathscr{O}_{X_0}(n-1)^{\oplus 2}$. The ℓth symmetric power will be a sum of copies of $\mathscr{O}_{X_0}(\ell(n-1))$. All such sheaves have vanishing H^i for all $i > 0$.

Write $\mathscr{I}_m$ for the ideal sheaf of X_m in $\mathrm{Res}(\boldsymbol{\Phi})$. Consider the standard exact sequences

$$0 \to \mathscr{I}_{m-1}/\mathscr{I}_m \to \mathscr{O}_{X_m} \to \mathscr{O}_{X_{m-1}} \to 0.$$

The above calculations show inductively that $\mathrm{H}^i(X_n, \mathscr{O}_{X_n}) = 0$ for all $n \geq 0$ and all $i > 0$. This concludes the proof. ∎

Corollary 2.21. *If* Φ *is a noncollinear multiview configuration, then the map* $\rho : \mathrm{Res}(\Phi) \to (\mathbf{P}^2)^n$ *is a closed immersion.*

Proof. By the noncollinearity assumption, the geometric fibers of ρ all have length at most 1. Thus, ρ is proper and quasi-finite, and hence finite. Applying Lemma 2.20 then shows that ρ is a closed immersion. ∎

Lemma 2.22. *Suppose* $\varphi_1, \varphi_2 : \mathbf{P}^3 \dashrightarrow \mathbf{P}^2$ *are cameras and* $\alpha : \mathbf{P}^3 \dashrightarrow \mathbf{P}^3$ *is a birational automorphism such that* $\varphi_2 = \varphi_1 \circ \alpha$. *If* α *and* $\varphi_1 \circ \alpha$ *are both regular on an open subset*

$U \subset \mathbf{P}^3$ whose complement has codimension at least 2, then α extends to a unique regular automorphism $\mathbf{P}^3 \to \mathbf{P}^3$.

Proof. Removing the center of φ_1 if necessary, we may assume that there is an open subscheme $U \subset \mathbf{P}^3$ on which φ_1, φ_2, and α are all regular and $\operatorname{codim}(\mathbf{P}^3, \mathbf{P}^3 \setminus U) \geq 2$. By assumption, $\varphi_i^* \mathscr{O}(1) = \mathscr{O}_U(1)$. Thus, $\alpha^* \mathscr{O}(1) = \mathscr{O}(1)$. Since $\Gamma(U, \mathscr{O}(1)) = \Gamma(\mathbf{P}^3, \mathscr{O}(1))$, we conclude from the universal property of projective space that the morphism $\alpha : U \to \mathbf{P}^3$ extends to a unique endomorphism $\widetilde{\alpha}$ of $\mathbf{P}^3$. Since α is birational, $\widetilde{\alpha}$ is an isomorphism, as desired. ■

Proposition 2.23. *Two multiview configurations $\mathbf{\Phi}^1$ and $\mathbf{\Phi}^2$ of length n are isomorphic if and only if their associated multiview schemes in $(\mathbf{P}^2)^n$ are equal. Two calibrated multiview configurations $(\mathbf{\Phi}^1, C_1)$ and $(\mathbf{\Phi}^2, C_2)$ are isomorphic if and only if their associated multiview flags* $\operatorname{Flag}(\mathbf{\Phi}^1, C_1)$ *and* $\operatorname{Flag}(\mathbf{\Phi}^2, C_2)$ *are equal.*

Proof. Since $\mathbf{\Phi}^i$ is birational onto its image for $i = 1, 2$, we see that if $\operatorname{Sch}(\mathbf{\Phi}^1) = \operatorname{Sch}(\mathbf{\Phi}^2)$, then there is a birational automorphism $\alpha : \mathbf{P}^3 \dashrightarrow \mathbf{P}^3$ such that $\mathbf{\Phi}^2 = \mathbf{\Phi}^1 \circ \alpha$. Moreover, $\operatorname{pr}_1 \circ \mathbf{\Phi}^1$, α, and $\operatorname{pr}_1 \circ \mathbf{\Phi}^2 \circ \alpha$ are all regular on the open subscheme of $\mathbf{P}^3$ that is the complement of the line joining the centers of the two cameras $\operatorname{pr}_1 \circ \mathbf{\Phi}^1$ and $\operatorname{pr}_1 \circ \mathbf{\Phi}^2$ (as this maps isomorphically into the smooth locus of $\operatorname{Sch}(\mathbf{\Phi}^1)$). Applying Lemma 2.22, we see that α is regular, as desired. The calibrated case follows, once we note that the calibrating curves lie in the regular locus of all cameras. ■

2.3. Relativization.

In this section we describe how to generalize the results of subsections 2.1 and 2.2 to families of cameras over an arbitrary base space. This is a necessary step towards defining the moduli of camera configurations.

Definition 2.24. *Given a scheme S, a* relative pinhole camera *over S is a rational map $p : \mathbf{P} \dashrightarrow \mathbf{P}^2_S$ over S uniquely determined by the following information:*

1. *the scheme $\mathbf{P}$ is a Zariski $\mathbf{P}^3_S$-bundle (i.e., has the form $\mathbf{P}(V)$ for a locally free $\mathscr{O}_S$-module of rank 4);*
2. *there is a map $\sigma : \mathscr{O}_{\mathbf{P}}^{\oplus 3} \to \mathscr{O}_{\mathbf{P}}(1)$ whose cokernel is an invertible sheaf supported exactly over a section Z of $\mathbf{P} \to S$, called the* camera center*;*
3. *a representative of p is given by the morphism $\mathbf{P} \setminus Z \to \mathbf{P}^2_S$ determined by the quotient $\sigma_{\mathbf{P} \setminus Z}$ and the universal property of projective space.*

Throughout this section, when the base scheme S is clear, we will often simply write $\mathbf{P}^2$ for $\mathbf{P}^2_S$, etc.

Definition 2.25. *Given a scheme S, a* relative multiview configuration of length n *over S is given by a proper S-scheme $\mathbf{P} \to S$ of finite presentation and a rational map $\mathbf{\Phi} : \mathbf{P} \dashrightarrow (\mathbf{P}^2_S)^n$ over S such that for each i the composition $\operatorname{pr}_i \circ \mathbf{\Phi}$ is a relative pinhole camera as in Definition 2.24.*

Two relative multiview configurations

$$\mathbf{\Phi}^i : \mathbf{P}_i \dashrightarrow \mathbf{P}^2_S, \quad i = 1, 2,$$

are isomorphic *if there is an S-isomorphism $\varepsilon : \mathbf{P}_1 \xrightarrow{\sim} \mathbf{P}_2$ such that $\mathbf{\Phi}^2 = \mathbf{\Phi}^1 \circ \varepsilon$.*

In what follows, we will write $\mathbf{P}^2$ for $\mathbf{P}^2_S$, etc., when the base scheme S is understood.

Notation 2.26. Given a multiview configuration $\mathbf{\Phi} : \mathbf{P} \to (\mathbf{P}^2)^n$ of length n, we will write

1. $S(\Phi)$ for the domain $\mathbf{P}$ of $\mathbf{\Phi}$;
2. $Z_1(\mathbf{\Phi}), \dots, Z_n(\mathbf{\Phi}) \subset \mathbf{\Phi}$ for the camera centers;
3. $Z(\mathbf{\Phi})$ for the scheme-theoretic union $Z_1(\mathbf{\Phi}) \cup \cdots \cup Z_n(\mathbf{\Phi})$;
4. $\mathrm{Res}(\mathbf{\Phi})$ for the blowup of $S(\mathbf{\Phi})$ in Z.

Definition 2.27. *A relative multiview configuration $\mathbf{\Phi}$ over S is* general *if the camera centers $Z_1, \dots, Z_{\mathrm{len}(\mathbf{\Phi})}$ are pairwise disjoint closed subschemes of $\mathbf{P}$.*

Definition 2.28. *A relative multiview configuration $\mathbf{\Phi} : \mathbf{P} \dashrightarrow (\mathbf{P}^2_S)^n$ over S is* collinear *if there is a closed subscheme $\mathbf{L} \subset S(\mathbf{\Phi})$ that is a relative line over S and that contains $Z(\mathbf{\Phi})$. It is* nowhere-collinear *if it is not collinear upon any basechange $S' \to S$.*

Definition 2.29. *Given a relative multiview configuration $\mathbf{\Phi}$ of length n over S, a* calibration datum *for $\mathbf{\Phi}$ is a pair $(C, (C_1, \dots, C_n))$ where*

1. *$C \subset \mathbf{P}$ is a relative degree 2 curve over S;*
2. *$C_i \subset \mathbf{P}^2_S$ is a relative smooth conic over S for $i = 1, \dots, n$;*
3. *for $i = 1, \dots, n$, the induced morphism $(\mathrm{pr}_i \circ \mathbf{\Phi})_C$ factors through C_i.*

If C is smooth, the calibration datum will be called smooth *or* nondegenerate*; otherwise, it will be called* degenerate.

Proposition 2.30. *Given a general relative multiview configuration $\mathbf{\Phi}$ over S, there is a unique commutative diagram*

$$\begin{tikzcd}
\mathrm{Res}(\mathbf{\Phi}) \arrow[dr, "\rho"] & \\
& (\mathbf{P}^2)^{\mathrm{len}(\mathbf{\Phi})} \\
\mathbf{P} \arrow[uu, dashed, "\pi^{-1}"] \arrow[ur, dashed, "\mathbf{\Phi}"'] &
\end{tikzcd}$$

The diagram has the property that for each i, the composition

$$E_i \xrightarrow{\iota_i} \mathrm{Res}(\mathbf{\Phi}) \xrightarrow{\rho} (\mathbf{P}^2)^{\mathrm{len}(\Phi)} \xrightarrow{\mathrm{pr}_i} \mathbf{P}^2$$

is an isomorphism. Moreover, this diagram is compatible with arbitrary base change on S.

Proof. The arrow ρ exists again by Lemma 2.12, and the functoriality follows from the functoriality of Lemma 2.12 and the flatness of everything over S. Finally, the isomorphism condition can be checked on geometric fibers, which reduces it to Proposition 2.13. ∎

Definition 2.31. *Given a general multiview configuration $\mathbf{\Phi}$ of length n, the scheme-theoretic image of the morphism ρ described in Proposition 2.30 is the* multiview scheme *of $\mathbf{\Phi}$. Similarly, given a calibrated multiview configuration $(\mathbf{\Phi}, C, (C_1, \dots, C_n))$, there is an associated flag* $\mathrm{Flag}(\mathbf{\Phi}, C)$ *sitting inside the flag scheme $C_1 \times \cdots \times C_n \subset (\mathbf{P}^2)^n$.*

Notation 2.32. The multiview scheme of $\mathbf{\Phi}$ will be denoted $\mathrm{Sch}(\mathbf{\Phi})$. It is a closed subscheme of $(\mathbf{P}^2)^{\mathrm{len}(\mathbf{\Phi})}$.

In the following, we fix conics in $\mathbf{P}^2$ and only record the curve $C \subset \mathbf{P}^3$ when considering calibrations.

Proposition 2.33. *Two general multiview configurations $\mathbf{\Phi}^1, \mathbf{\Phi}^2$ of length n over S are isomorphic if and only if* $\mathrm{Sch}(\mathbf{\Phi}^1) = \mathrm{Sch}(\mathbf{\Phi}^2)$ *as closed subschemes of $(\mathbf{P}^2_S)^n$. Similarly, two general calibrated multiview configurations $(\mathbf{\Phi}_1, C_1)$ and $(\mathbf{\Phi}_2, C_2)$ are isomorphic if and only if their flags* $\mathrm{Flag}(\mathbf{\Phi}_1, C_1)$ *and* $\mathrm{Flag}(\mathbf{\Phi}_2, C_2)$ *are equal.*

The proof of Proposition 2.33 is a modification of that of Proposition 2.23. We require a modification of Lemma 2.22.

Lemma 2.34. *Suppose A is a ring and $U \subset \mathbf{P}^3_A$ is an open subset such that for every geometric point $A \to \kappa$ the fiber $U_\kappa \subset \mathbf{P}^3_\kappa$ has complement of codimension at least 2. Suppose $\alpha : U \to \mathbf{P}^3_A$ is a morphism such that $\alpha^*\mathscr{O}(1) = \mathscr{O}_U(1)$. Then α extends to a unique automorphism of $\mathbf{P}^3_A$.*

Proof. By the universal property of projective space, it suffices to show that restriction defines an isomorphism

$$\Gamma(\mathbf{P}^3_A, \mathscr{O}(1)) \xrightarrow{\sim} \Gamma(U, \mathscr{O}(1)).$$

To show this, it suffices to show that the adjunction map $\nu(1) : \mathscr{O}_{\mathbf{P}^3}(1) \to \iota_*\mathscr{O}_U(1)$ is an isomorphism of sheaves. By the projection formula, it suffices to show that the adjunction map for the structure sheaf

$$\nu : \mathscr{O}_{\mathbf{P}^3_A} \to \iota_*\mathscr{O}_U$$

is an isomorphism. But this is precisely Proposition 3.5 of [7]. ∎

Proposition 2.35. *If $\mathbf{\Phi}$ is a general multiview configuration over S, then for all base changes $T \to S$ we have that the natural morphism*

$$\mathrm{Sch}(\mathbf{\Phi}) \times_S T \to \mathrm{Sch}(\mathbf{\Phi} \times_S T)$$

is an isomorphism; that is, formation of the associated multiview scheme is compatible with base change. Furthermore, $\mathrm{Sch}(\mathbf{\Phi})$ *is flat over the base.*

Proof. By Lemma 2.20 the structure morphism $\mathscr{O}_{(\mathbf{P}^2)^n} \to \rho_*\mathscr{O}_{\mathrm{Res}(\mathbf{\Phi})}$ is surjective. Consider the triangle in the derived category

$$I \to \mathscr{O}_{(\mathbf{P}^2)^n} \to \mathbf{R}\,\rho_*\mathscr{O}_{\mathrm{Res}(\mathbf{\Phi})} \to I[1].$$

Let $i : (\mathbf{P}^2)^n_q \to (\mathbf{P}^2)^n$ be an embedding of a fiber. Pulling back to the fiber and using cohomology and base change, we have

$$\begin{aligned}\mathbf{L}\,i^*\,\mathbf{R}\,\rho_*\mathscr{O}_{\mathrm{Res}(\mathbf{\Phi})} &\simeq \mathbf{R}\,\rho_*\,\mathbf{L}\,i^*_{\mathrm{Res}(\mathbf{\Phi})}\mathscr{O}_{\mathrm{Res}(\mathbf{\Phi})}\\ &\simeq \mathbf{R}\,\rho_*(\mathscr{O}_{\mathrm{Res}(\mathbf{\Phi})})_q\\ &\simeq (\mathscr{O}_{\mathrm{Res}(\mathbf{\Phi})})_q.\end{aligned}$$

Applying [9, Lemma 3.31] to $\mathbf{R}\,\rho_*\mathscr{O}_{\mathrm{Res}(\mathbf{\Phi})}$, we see that it is quasi-isomorphic to a sheaf flat over the base. But $\mathscr{H}^0(\mathbf{R}\,\rho_*\mathscr{O}_{\mathrm{Res}(\mathbf{\Phi})})$ is $\rho_*\mathscr{O}_{\mathrm{Res}(\mathbf{\Phi})}$. Thus, we conclude that the short exact sequence

$$0 \to \mathscr{I} \to \mathscr{O}_{(\mathbf{P}^2)^n} \to \rho_*\mathscr{O}_{\mathrm{Res}(\mathbf{\Phi})} \to 0$$

consists of S-flat sheaves and is compatible with arbitrary base change. This establishes the result. ■

3. Moduli and deformation theory.

3.1. Moduli of uncalibrated camera configurations. In this section we describe the basic moduli problem attached to uncalibrated camera configurations. In subsection 3.2 we will study the deformation theory of a configuration $\mathbf{\Phi}$, especially as it relates to the deformation theory of the associated scheme $\mathrm{Sch}(\mathbf{\Phi})$. Ultimately this will allow us to embed the moduli space into the Hilbert scheme.

Definition 3.1. *Given a positive integer n, the* functor of camera configurations of length *n, denoted* Cam_n*, has as value over a scheme S the set of isomorphism classes of general relative multiview configurations of length n.*

Since a camera configuration of length at least 2 has trivial automorphism group, it follows from standard descent theory that Cam_n is a sheaf in the fppf topology. In this section we will show that it is a quasi-projective variety.

Notation 3.2. Let $\mathrm{M}^n \subset \mathrm{M}^n_{3\times 4}$ be the locus of n-tuples of full rank 3×4 matrices whose kernels are pairwise distinct. Let T be the torus given by the kernel of the multiplication map $\mathbf{G}^n_m \to \mathbf{G}_m$. There is a natural free action of $T \times \mathrm{GL}_4$ on M^n (where the torus T acts diagonally by scaling). Moreover, since $T \times \mathrm{GL}_4$ is reductive over $\mathbf{Z}$ and $M^n_{3\times 4}$ is affine, we can realize the quotient sheaf $\mathrm{M}^n/T \times \mathrm{GL}_4$ as an open subvariety of the GIT quotient $M^n_{3\times 4} /\!\!/ T \times \mathrm{GL}_4$. In particular, the quotient $\mathrm{M}^n/T \times \mathrm{GL}_4$ is a smooth quasi-projective variety. Because the action is free, we also know the functor of points of $\mathrm{M}^n/T \times \mathrm{GL}_4$: the S-valued points are given by pairs $(L \to S, \varphi : S \to \mathrm{M}^n)$, where $L \to S$ is a $T \times \mathrm{GL}_4$-torsor and φ is a $T \times \mathrm{GL}_4$-equivariant map. In particular, a morphism $\mathrm{M}^n/T \times \mathrm{GL}_4 \to Y$ to a scheme Y is the same thing as a $T \times \mathrm{GL}_4$-invariant morphism $\mathrm{M}^n \to Y$.

Proposition 3.3. *There is a natural isomorphism of functors $c : \mathrm{M}^n/T \times \mathrm{GL}_4 \to \mathrm{Cam}_n$.*

Proof. Sending a 3×4-matrix to its associated camera defines a morphism $\mathrm{M}^n \to \mathrm{Cam}_n$. This is $T \times \mathrm{GL}_4$-equivariant, since, by definition, projective automorphisms of $\mathbf{P}^3$ do not affect the isomorphism class of a camera configuration. To see that c is an isomorphism, it suffices to show that $c(R)$ is a bijection for any strictly Henselian local ring R. In this case, every form of $\mathbf{P}^3$ is trivial, so we see that any camera configuration is given by a tuple of matrices, showing that c is surjective. On the other hand, by definition, two such configurations are isomorphic if and only if they differ by an automorphism of $\mathbf{P}^3$ and individual scalings of the factors, which says precisely that they lie in the same $T \times \mathrm{GL}_4(R)$-orbit in $\mathrm{M}^n(R)$. The result follows. ■

Corollary 3.4. *If $n > 1$, then the space Cam_n is a smooth quasi-projective scheme over* $\mathrm{Spec}\,\mathbf{Z}$.

Proof. This follows immediately from Proposition 3.3 and the remarks in Notation 3.2. ■

3.2. Deformations of multiview configurations. In this section, we study the relationship between the infinitesimal deformation theory of a camera configuration and the deformation theory of its associated multiview scheme. As we will see in subsection 4.3, the deformationtheoretic approach gives strong results on the relationship between Cam_n and $\mathrm{Hilb}_{(\mathbf{P}^2)^n}$, clarifying and improving the groundbreaking results of [1]. In particular, our infinitesimal analysis will apply at all points. These methods are very different from the ideal-theoretic methods of [1]. It would be especially interesting to understand how the cotangent complex argument of subsection 3.2.3 relates to the Gröbner basis calculations in [1].

Definition 3.5. *Fix a ring A containing an ideal I such that $I^2 = 0$, and let $A_0 = A/I$. Suppose $\mathbf{\Phi}^0$ is a relative multiview configuration of length n over A_0. An* infinitesimal deformation of $\mathbf{\Phi}^0$ to A *is a pair $(\mathbf{\Phi}, \varepsilon)$, where $\mathbf{\Phi}$ is a multiview configuration of length n over A and $\varepsilon : \mathbf{\Phi} \otimes_A A_0 \xrightarrow{\sim} \mathbf{\Phi}^0$ is an isomorphism of relative multiview configurations.*

An isomorphism between infinitesimal deformations *$(\mathbf{\Phi}, \varepsilon)$ and $(\mathbf{\Phi}', \varepsilon')$ of $\mathbf{\Phi}^0$ is an isomorphism $\alpha : \mathbf{\Phi} \xrightarrow{\sim} \mathbf{\Phi}'$ of relative multiview configurations such that $\varepsilon' \circ \alpha \otimes_A A_0 = \varepsilon$.*

Notation 3.6. We will write $\mathrm{Def}_{\mathbf{\Phi}^0}$ for the functor of isomorphism classes of infinitesimal deformations of $\mathbf{\Phi}^0$, and $\mathrm{Def}_{\mathrm{Sch}(\mathbf{\Phi}^0) \subset (\mathbf{P}^2)^{\mathrm{len}(\mathbf{\Phi}^0)}}$ for the usual functor of infinitesimal deformations of the point $[\mathrm{Sch}(\mathbf{\Phi}^0)]$ of the Hilbert scheme $\mathrm{Hilb}_{(\mathbf{P}^2)^{\mathrm{len}(\mathbf{\Phi}^0)}}$.

Our goal in this section is to prove the following, which is the key step in our generalization of the results of [1].

Proposition 3.7. *If $\mathbf{\Phi}$ is a general multiview configuration of length $n > 2$, then the morphism*

$$\mathrm{Sch} : \mathrm{Def}_{\mathbf{\Phi}^0} \to \mathrm{Def}_{\mathrm{Sch}(\mathbf{\Phi}^0) \subset (\mathbf{P}^2)^n}$$

is an isomorphism of deformation functors.

Proof. The proof will be developed through this section. In particular, the injectivity of Sch follows from Proposition 2.33, and surjectivity follows from Proposition 3.11. ∎

That is, if $\mathbf{\Phi}$ is a general multiview configuration of length $n > 2$ with associated multiview variety $V \subset (\mathbf{P}^2)^n$, then we have that the infinitesimal deformations of $\mathbf{\Phi}$ are in bijection with the infinitesimal deformations of V as a closed subscheme of $(\mathbf{P}^2)^n$. The proof will work roughly as follows.

1. First, we will recall the well-known description of abstract deformations of V as a scheme. As we will see, V has a property that we will call *essential rigidity*.
2. Using this essential rigidity, we will show that any deformation of V as a closed subscheme of $(\mathbf{P}^2)^n$ arises from a deformation of $\mathbf{\Phi}$. In the collinear case this is nontrivial, because $\mathrm{Res}(\Phi) \to (\mathbf{P}^2)^n$ contracts a line, but a simple argument with the cotangent complex gives the desired result.
3. Using Proposition 2.33, we have that two deformations of $\mathbf{\Phi}$ give rise to the same deformation of V if and only if they are isomorphic, completing the proof.

It is worth noting (as hinted at in this outline) that the proof we give here is almost purely geometric. We do not rely on dimension estimates, ideal-theoretic calculations, etc. The arguments are simple variants of classical Italian geometric arguments, first used to study the geometry of projective surfaces. Proposition 3.7 is ultimately the reason that the space of multiview configurations admits an open immersion into the Hilbert scheme, as we will see in subsection 4.3.

3.2.1. Essential rigidity of blowups of $\mathbf{P}^3$. In this section we fix a commutative ring A_0, a square-zero extension

$$I \subset A \to A_0,$$

and a collection of pairwise everywhere-disjoint sections

$$\sigma_i : \operatorname{Spec} A_0 \to \mathbf{P}^3_{A_0}.$$

We write P_0 for the blowup $\operatorname{Bl}_{Z_0} \mathbf{P}^3_{A_0}$, where Z_0 is the reduced closed subscheme of $\mathbf{P}^3_{A_0}$ supported on the union of the images of the σ_i. For the most part, these results are well known. Unfortunately, the available literature tends not work in sufficient generality (for example, [14] works over a fixed field $\mathbf{k}$).

Proposition 3.8. *Given a deformation P of P_0 over A, there is a unique morphism*

$$\beta : P \to \mathbf{P}^3_A$$

deforming the canonical blowdown map

$$\beta_0 : P_0 \to \mathbf{P}^3_{A_0},$$

up to infinitesimal automorphism of $\mathbf{P}^3_A$. Moreover, β realizes P as the blowup of $\mathbf{P}^3_A$ at a closed subscheme Z that deforms Z_0 (and Z is a union of n sections of $\mathbf{P}^3_A$).

Proof. If one is willing to work entirely over a field (although here we are working over $\mathbf{Z}$), one can extract this from [14, Proposition 3.4.25(ii)]. It is not difficult to prove this in full generality for blowups of projective spaces along collections of sections by showing that the blowdown map admits a canonical deformation, and each deformed exceptional divisor maps to a section under this deformed blowdown. We omit the details for the sake of space. ∎

3.2.2. Lifting deformation for noncollinear configurations. In this section, we explain how any deformation of a noncollinear multiview scheme lifts to a deformation of the associated multiview configuration. Fix a deformation situation

$$I \subset A \to A_0$$

and a noncollinear multiview configuration $\mathbf{\Phi}^0$ of length n over A_0 with scheme $\operatorname{Sch}(\mathbf{\Phi}^0)$.

Proposition 3.9. *If $X \subset (\mathbf{P}^2)^n_A$ is an A-flat deformation of $\operatorname{Sch}(\mathbf{\Phi}^0)$, then there is a deformation $\mathbf{\Phi}$ of $\mathbf{\Phi}^0$ such that $\operatorname{Sch}(\mathbf{\Phi}) = X$ as closed subschemes of $(\mathbf{P}^2)^n$. Moreover, $\mathbf{\Phi}$ is unique up to unique isomorphism of deformations of $\mathbf{\Phi}^0$ over A.*

Proof. Since $\mathbf{\Phi}^0$ is noncollinear, the natural morphism

$$\operatorname{Res}(\mathbf{\Phi}^0) \to \operatorname{Sch}(\mathbf{\Phi}^0) \subset (\mathbf{P}^2)^n$$

is an isomorphism. By Proposition 3.8, any deformation of $\operatorname{Sch}(\mathbf{\Phi}^0)$ is a blowup P of $\mathbf{P}^3_A$ at n disjoint sections over $\operatorname{Spec} A$. The deformation thus results in a rational map

$$\mathbf{\Phi} : \mathbf{P}^3_A \dashrightarrow (\mathbf{P}^2_A)^n$$

extending $\mathbf{\Phi}^0$. We wish to show that $\mathbf{\Phi}$ is a relative multiview configuration in the sense of Definition 2.27. To do this, it suffices to check that composition with each projection is a relative pinhole camera. Write $p : \mathbf{P}^3_A \dashrightarrow \mathbf{P}^2_A$ for one such projection; we will abuse notation and also write p for the corresponding map $P \to \mathbf{P}^2_A$ from the blowup. We will write E for the exceptional divisor associated to p and Z for the section blown up to make E, that is, we assume that p is the ith projection of $\mathbf{\Phi}$ and that E is the preimage of the ith section in $\mathbf{P}^3_A$, which we call Z, uniformly omitting i from the notation. By the pinhole camera assumptions on $\mathbf{\Phi}^0$, $p|E_{A_0}$ maps E isomorphically to $\mathbf{P}^2_{A_0}$. It follows from Nakayama's lemma that $p|_E$ maps E isomorphically to $\mathbf{P}^2_A$.

Write $U \subset \mathbf{P}^3_A$ for the complement of the sections that are blown up to resolve $\mathbf{\Phi}$. By the previous paragraph, we see that $U_{A_0} \subset \mathbf{P}^3_{A_0}$ is precisely the complement of the camera centers of $\mathbf{\Phi}^0$. By the universal property of projective space, the morphism p is given by a surjective morphism

$$\lambda : \mathscr{O}_P^{\oplus 3} \to \mathscr{L}$$

for some $\mathscr{L}$ in $\mathrm{Pic}(P)$. Write $\pi : P \to \mathbf{P}^3_A$ for the blowdown map. We know from the definition of pinhole cameras, the rigidity of invertible sheaves on P, and the canonical way to extend morphisms generically across blowups that $\mathscr{L} \cong \pi^*(\mathscr{O}(1))(-E)$. Moreover, the resulting arrow

$$f : \pi_* \mathscr{O}^{\oplus 3} \to \mathscr{O}_{\mathbf{P}^3_A}(1)$$

has the property that its image is precisely $\mathscr{O}_{\mathbf{P}^3_A}(1) \otimes \mathscr{I}_Z$, where $\mathscr{I}_Z$ is the ideal sheaf of Z. (This follows from the universal property of blowing up.) This shows that the cokernel of f is an invertible sheaf supported on Z, showing that p is a relative pinhole camera, as desired.

It remains to show that any two such realizations $\mathbf{\Phi}_1$ and $\mathbf{\Phi}_2$ are conjugate by an infinitesimal automorphism of $\mathbf{P}^3$. But this follows immediately from Proposition 2.33. ∎

3.2.3. Lifting deformations for collinear configurations. For the sake of computational ease, in this section we consider a deformation situation $I \subset A \to A_0$ in which A is an Artinian local ring with maximal ideal $\mathfrak{m}$ and $\mathfrak{m}I = 0$. Write $k = A/\mathfrak{m}$.

We start with a multiview configuration $\mathbf{\Phi} : \mathbf{P}^3_{A_0} \dashrightarrow (\mathbf{P}^2)^n$ whose special fiber $\mathbf{\Phi}_k$ is collinear. Thus, the morphism

$$\mathrm{Res}(\mathbf{\Phi}_k) \to \mathrm{Sch}(\mathbf{\Phi}_k) \subset (\mathbf{P}^2)^n$$

contracts a line $\ell \subset \mathrm{Res}(\mathbf{\Phi}_k)$. To make things easier to read, write $R = \mathrm{Res}(\mathbf{\Phi}_k)$ and $B = \mathrm{Sch}(\mathbf{\Phi}_k)$. Write $L_{R/B}$ for the cotangent complex of the morphism $R \to B$. In addition, write $E_1, \ldots, E_n \subset R$ for the exceptional divisors. The usual calculations show that $K_R = \pi^* K_{\mathbf{P}^3} + 2E_1 + \cdots + 2E_n$.

Lemma 3.10. *If $n > 2$, then $\mathrm{Ext}^2_R(L_{R/B}, \mathscr{O}_R) = 0$.*

Proof. Consider the standard spectral sequence

$$E_2^{pq} = \mathrm{Ext}^p(\mathscr{H}^{-q}(L_{R/B}, \mathscr{O}_R)) \Rightarrow \mathrm{Ext}^{p+q}(L_{R/B}, \mathscr{O}_R). \tag{3.1}$$

We know that $\mathscr{H}^0(L_{R/B}) = \Omega^1_{R/B}$, and that $\mathscr{H}^{-j}(L_{R/B})$ is supported on ℓ for all $j \geq 0$. By Serre duality, we can compute the terms in the spectral sequence as

$$\mathrm{Ext}^p(\mathscr{H}^{-q}(L_{R/B}), \mathscr{O}_R) = \mathrm{H}^{3-p}(R, \mathscr{H}^{-q}(L_{R/B})(K_R))^\vee.$$

Since the cohomology sheaves of $L_{R/B}$ are all supported on ℓ, all columns of the E^2_{pq} page (3.1) vanish except (possibly) for $p = 2, 3$. It follows that

$$\mathrm{Ext}^2_R(L_{R/B}, \mathscr{O}_R) \cong \mathrm{H}^1(R, \Omega^1_{R/B}(K_R))^\vee.$$

A local calculation shows that $\Omega^1_{R/B}$ is annihilated by the ideal of ℓ, so that $\Omega^1_{R/B} = \Omega^1_{\ell/\operatorname{Spec} k}$, and thus

$$\mathrm{H}^1(R, \Omega^1_{R/B}(K_R))^\vee \cong \mathrm{H}^1(\ell, \mathscr{O}_\ell(K_\ell + K_R))^\vee \cong \mathrm{H}^0(\ell, \mathscr{O}_\ell(-K_R)) = \mathrm{H}^0(\ell, \mathscr{O}(4-2n)) = 0,$$

as desired. ∎

Proposition 3.11. *Suppose $n > 2$. If $X \subset (\mathbf{P}^2)^n_A$ is an A-flat deformation of* $\mathrm{Sch}(\mathbf{\Phi}^0)$*, then there is a deformation $\mathbf{\Phi}$ of $\mathbf{\Phi}^0$ such that* $\mathrm{Sch}(\mathbf{\Phi}) = X$ *as closed subschemes of* $(\mathbf{P}^2)^n$*. Moreover, $\mathbf{\Phi}$ is unique up to the unique isomorphism of deformations of $\mathbf{\Phi}^0$ over A.*

Proof. By Lemma 3.10 and [10, III.2.2.4], the obstruction to deforming the morphism

$$\mathrm{Res}(\Phi^0) \to \mathrm{Sch}(\Phi^0)$$

over A vanishes, resulting in a deformation $R \to X$. Applying the results of subsection 3.2.1, we see that this arises from a deformation Φ, as desired. The uniqueness of Φ up to isomorphism is an immediate consequence of Proposition 2.33. ∎

3.3. Diagram Hilbert schemes. In this section, we briefly explain a basic idea that is hard to find in the literature: diagram Hom-schemes and diagram Hilbert schemes. They are a mild elaboration of the idea of a flag Hilbert scheme. By remembering not only the data of the image but also the calibrating conics, the moduli space of calibrated cameras maps to a diagram Hilbert scheme in the same way that the moduli space of uncalibrated cameras maps to a Hilbert scheme.

3.3.1. Definition and examples. Fix a base scheme S, a category I, and a functor $\underline{X} : I \to \mathfrak{AlgSp}_S$, where $\mathfrak{AlgSp}_S$ denotes the category of algebraic spaces over S.

Definition 3.12. *The* diagram Hilbert functor

$$\mathrm{Hilb}_{\underline{X}} : \mathfrak{Sch}_S^\circ \to \mathfrak{Sets}$$

is the functor whose value on an S-scheme T is the set of isomorphism classes of natural transformations $\underline{Y} \to \underline{X} \times_S T$ of functors $I \to \mathfrak{Sch}_T$ where for each $i \in I$ the associated arrow $\underline{Y}(i) \to \underline{X}(i) \times_S T$ is a T-flat family of proper closed subschemes of $\underline{X}(i)$ of finite presentation over T.

Example 3.13. The usual Hilbert scheme is an example: Just take I to be the singleton category. So is the flag Hilbert scheme of length n: In this case the category I is the category $\underline{n}$ associated to the poset $\{1, \dots, n\}$, and the functor $\underline{X}$ is the constant functor $X \to X$. A natural transformation $\underline{Y} \to \underline{X}$ defines a nested sequence of closed subschemes of X. This is the flag Hilbert scheme (of length 2 flags).

There is also a stricter kind of flag scheme: Suppose $X_1 \subset X_2$ is a closed immersion and one wants to parameterize pairs $Y_i \subset X_i$ such that $Y_1 \subset Y_2$, that is precisely the diagram Hilbert functor associated to the poset-category $\underline{2} = \{1 < 2\}$ with the functor $\underline{2} \to \mathfrak{Sch}_S$ sending i to X_i. This last example is the one that will arise naturally for us in the context of calibrated cameras. (We record more general results here in case someone in the future needs this general idea of a diagram Hilbert scheme.)

Notation 3.14. If the diagram in question is a single morphism $X \to Y$, we will write $\mathrm{Hilb}_{X\to Y}$ for the associated Hilbert functor.

3.3.2. Representability. The main result about diagram Hilbert functors is that they are representable. We prove this in a high degree of generality, in case this is of independent interest.

Proposition 3.15. *Let I be a finite category and $\underline{X} : I \to \mathfrak{AlgSp}_S$ a functor whose components are separated algebraic spaces. Then the diagram Hilbert functor $\mathrm{Hilb}_{\underline{X}}$ is representable by an algebraic space locally of finite presentation over S. If the $\underline{X}(i)$ are locally quasi-projective schemes, then $\mathrm{Hilb}_{\underline{X}}$ is represented by a locally quasi-projective S-scheme.*

Proof. There is a natural functor

$$F : \mathrm{Hilb}_{\underline{X}} \to \prod_{i\in I} \mathrm{Hilb}_{\underline{X}(i)},$$

and we know that the latter is representable by algebraic spaces (resp., schemes) satisfying the desired conditions. It thus suffices to show the same for F, i.e., that F is representable by spaces of the required type.

For each $i \in I$, let

$$Z_i \subset \underline{X}(i) \times \prod \mathrm{Hilb}_{\underline{X}(i)}$$

denote the universal closed subscheme (pulled back over the product). Let A denote the set of arrows in I; for an arrow $a \in A$, let $s(a)$ and $t(a)$ denote the source and target of a. Consider the scheme

$$H := \prod_{a\in A} \mathrm{Hom}_{\mathrm{Hilb}_{\prod \underline{X}(i)}}(Z(s(a)), Z(t(a))),$$

which naturally fibers over $\prod \mathrm{Hilb}_{\underline{X}(i)}$. The standard theory of Hom-schemes shows that $H \to \prod \mathrm{Hilb}_{\underline{X}(i)}$ is representable by spaces of the desired type.

The final observation to make is that composition of two arrows gives equations $b \circ a = c$ in A, and these translate into *closed* conditions on H because all of the subschemes $Z(i)$ are separated. Since the conditions desired are stable under taking closed subspaces, we have proven the result. ■

3.4. Moduli of calibrated camera configurations. Let $\mathscr{C}$ denote the space of smooth conics in $\mathbf{P}^2_{\mathrm{Spec}\,\mathbf{Z}[1/2]}$, and let $C_{\mathrm{univ}} \subset \mathbf{P}^2_{\mathscr{C}}$ denote the universal smooth conic. (The space $\mathscr{C}$ is an open subscheme of the bundle of sections of $\mathscr{O}_{\mathbf{P}^2_{\mathrm{Spec}\,\mathbf{Z}[1/2]}}(2)$.) The tuple of conics $(C_{\mathrm{univ}}, \ldots, C_{\mathrm{univ}})$ inside $(\mathbf{P}^2)^n$ will be called the *universal calibration*.

Definition 3.16. *Given a positive integer n, the* sheaf of calibrated camera configurations of length n, *denoted* CalCam_n, *is the sheaf over the Cartesian power $\mathscr{C}^n$ whose value over a point $t : S \to \mathscr{C}^n$ consists of the set of isomorphism classes of general relative calibrated multiview configurations of length n with calibration datum of the form $(C, t^*(C_{\mathrm{univ}}, \ldots, C_{\mathrm{univ}}))$.*

In down-to-earth terms, we are just describing the space of n-tuples of calibrated cameras with pairwise nonintersecting centers, together with *arbitrary but specified* calibration data. In the existing literature, the word "calibrated" usually means that one has fixed the calibrating conics to be the canonical absolute conic in space (attached to the Euclidean distance form on $\mathbf{P}^3$) and the circle in the plane. Since any two smooth conics are conjugate under a homography, this seems harmless. As we hope to describe in this section, thinking more geometrically and *tracking the conics as data instead of normalizing them* gives us a great deal of insight into the underlying moduli problem. The point of the universal conic in $\mathbf{P}^2$ is that we only want to allow the conic in $\mathbf{P}^3$ to vary; that is, we fix calibration data on the image planes when we define the moduli problem. By working with the universal conic, we allow those fixed planar data to be arbitrary.

Notation 3.17. Since we are fixing the calibration data on the image planes to be the universal conic, we will omit them from the notation for a calibration datum. Thus, we will write $(\mathbf{\Phi}, C)$ for a calibrated configuration. When we need to refer to the image plane calibrating curves, we will use C_i for the curve in the ith plane. It is key to remember that while C_i can vary as the base varies (depending upon how it maps to $\mathscr{C}^n$), this is determined solely by the base and not by the object of CalCam_n over that point of the base.

The main result of this section is the following.

Proposition 3.18. *The sheaf CalCam_n is a smooth scheme of finite type over $\mathscr{C}^n$.*

Let $\tau_n : \mathrm{CalCam}_n \to \mathrm{CalCam}_{n-1} \times_{\mathscr{C}^{n-1}} \mathscr{C}^n$ be the morphism given by forgetting the last camera (and retaining the last calibrating plane conic).

Lemma 3.19. *The morphism τ_n is representable by separated schemes of finite presentation.*

Proof. Let $((\varphi_1, \ldots, \varphi_{n-1}, C), C_n)$ be a T-valued point of $\mathrm{CalCam}_{n-1} \times_{\mathscr{C}^{n-1}} \mathscr{C}^n$. The fiber of τ_n is given by the set of cameras φ_n with the same domain $\mathbf{P} \to T$ as the first $n-1$ cameras, with the following additonal properties.

1. The center of φ_n avoids the centers of φ_i for $i = 1, \ldots, n-1$.
2. The restriction $\varphi_n|C$ factors through the closed subscheme $C_n \subset \mathbf{P}$.

The space of camera centers satisfying the first condition is an open subscheme $\mathbf{P}^\circ \subset \mathbf{P}$, and taking the center gives a natural map

$$\mathrm{CalCam}_n \to \mathbf{P}^\circ \times \mathrm{CalCam}_{n-1} \times_{\mathscr{C}^{n-1}} \mathscr{C}^n.$$

It suffices to show that this map is representable, and thus we may assume that the center is a given section $\sigma : T \to \mathbf{P}$. Blowing up along $\sigma(T)$ to yield $\widetilde{\mathbf{P}}$, with exceptional divisor E, we can then realize the cameras inside the open locus of the Hom-scheme $\mathrm{Hom}(\widetilde{\mathbf{P}}, \mathbf{P}^2)$ parametrizing maps $f : \widetilde{\mathbf{P}} \to \mathbf{P}^2$ for which $f^*\mathscr{O}_{\mathbf{P}^2}(1)$ is isomorphic to $\mathscr{O}(1)(-E)$ on each geometric fiber over T. This locus is of finite type. Finally, the condition that C lands in C_n is closed (and of finite presentation), completing the proof. ■

Proposition 3.20. *The morphism τ_n is smooth.*

Proof. By Lemma 3.19 and [15, Tag 02H6], it suffices to show that τ_n is formally smooth. Let $A \to A_0$ be a square-zero extension of rings, and suppose that

$$(\varphi_1, \ldots, \varphi_n, C) \in \mathrm{CalCam}_n(A_0)$$

is fixed. To show formal smoothness we can work Zariski-locally and thus assume that the domains of $\varphi_1, \ldots, \varphi_n$ are $\mathbf{P}^3_{A_0}$. Now suppose that we fix a deformation

$$((\varphi'_1, \ldots, \varphi'_{n-1}, C_A), C_n) \in \mathrm{CalCam}_{n-1}(A) \times_{\mathscr{C}^{n-1}(A)} \mathscr{C}^n(A).$$

(Because we are working over the universal conic in each image plane, we have to specify the deformation of the conic that we will use in attempting to deform the nth calibrated camera.) To show formal smoothness it suffices to extend φ_n to a morphism φ'_n that maps C_A to C_n.

The choice of deformation of C to C_A induces a lift of $C \to \mathbf{P}^2_{A_0}$ to $C_A \to \mathbf{P}^2_A$. This is because $\mathrm{H}^1(C, \mathscr{O}_C(1)) = 0$, so sections defining a map can always be lifted. We will show that we can extend this to a camera that acts on C_A in the given way.

We are thus reduced to the following: We are given a tuple of three sections $\sigma_0, \sigma_1, \sigma_2 \in \Gamma(\mathbf{P}^3_{A_0}, \mathscr{O}(1))$, a planar curve $C_A \subset \mathbf{P}^3_A$ of degree 2, and lifts of the $\sigma_j|_C$ to $\Gamma(C_A, \mathscr{O}(1))$. We wish to lift these extensions to sections $\widetilde{\sigma}_j \in \Gamma(\mathbf{P}^3_A, \mathscr{O}(1))$. We can do this one section at a time. By Definition 2.3, the curve C_A is contained in a canonically defined family of planes in $\mathbf{P}^3_A$; we will write $C_A \subset \mathbf{P}^2_A \subset \mathbf{P}^3_A$ and similarly for A_0. (If the plane is not trivial, we can further shrink A to make it so; this is immaterial for the calculations and is only a notational device.)

Consider the diagrams

$$\begin{tikzcd}
0 \arrow[r] & \Gamma(\mathbf{P}^3_{A_0}, \mathscr{O}) \otimes_{A_0} I \arrow[r] \arrow[d] & \Gamma(\mathbf{P}^3_A, \mathscr{O}) \arrow[r] \arrow[d] & \Gamma(\mathbf{P}^3_{A_0}, \mathscr{O}) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \Gamma(\mathbf{P}^3_{A_0}, \mathscr{O}(1)) \otimes_{A_0} I \arrow[r] \arrow[d] & \Gamma(\mathbf{P}^3_A, \mathscr{O}(1)) \arrow[r] \arrow[d] & \Gamma(\mathbf{P}^3_{A_0}, \mathscr{O}(1)) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \Gamma(\mathbf{P}^2_{A_0}, \mathscr{O}(1)) \otimes_{A_0} I \arrow[r] & \Gamma(\mathbf{P}^2_A, \mathscr{O}(1)) \arrow[r] & \Gamma(\mathbf{P}^2_{A_0}, \mathscr{O}(1)) \arrow[r] & 0
\end{tikzcd}$$

and

$$\begin{tikzcd}
0 \arrow[r] & \Gamma(\mathbf{P}^2_{A_0}, \mathscr{O}(-1)) \otimes_{A_0} I \arrow[r] \arrow[d] & \Gamma(\mathbf{P}^2_A, \mathscr{O}(-1)) \arrow[r] \arrow[d] & \Gamma(\mathbf{P}^2_{A_0}, \mathscr{O}(-1)) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \Gamma(\mathbf{P}^2_{A_0}, \mathscr{O}(1)) \otimes_{A_0} I \arrow[r] \arrow[d] & \Gamma(\mathbf{P}^2_A, \mathscr{O}(1)) \arrow[r] \arrow[d] & \Gamma(\mathbf{P}^2_{A_0}, \mathscr{O}(1)) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \Gamma(C, \mathscr{O}(1)) \otimes_{A_0} I \arrow[r] & \Gamma(C_A, \mathscr{O}(1)) \arrow[r] & \Gamma(C, \mathscr{O}(1)) \arrow[r] & 0
\end{tikzcd}$$

By the usual calculations of the cohomology of projective space, these two diagrams have exact columns. A simple diagram chase then shows that we can lift sections to $\mathbf{P}^3_A$ given values on $\mathbf{P}^3_{A_0}$ and C_A, completing the proof. ■

Proof of Proposition 3.18. It remains to show smoothness. We use Proposition 3.20 and induction on n. For $n = 1$, we see that CalCam_1 is smooth over $\mathscr{C}$, which is itself open in a projective space, and hence smooth. ■

3.5. Deformation theory of calibrated camera configurations. In this section we prove the following analogue of Proposition 3.7.

Theorem 3.21. *If* $(\mathbf{\Phi}, C)$ *is a nondegenerate calibrated general multiview configuration of length* $n > 2$ *with associated multiview flag*

$$(C \subset V) \hookrightarrow (C_1 \times \cdots \cdot C_n \subset (\mathbf{P}^2)^n),$$

then we have that the infinitesimal deformations of $(\mathbf{\Phi}, C)$ *are in bijection with the infinitesimal deformations of* $C \subset V$ *as a closed subscheme diagram of* $C_1 \times \cdots \times C_n \subset (\mathbf{P}^2)^n$*.*

Proof. The proof leverages the proof of Proposition 3.7. In particular, we can forget the calibrations and apply Proposition 3.7 to see that under the given hypotheses any deformation of $\mathrm{Flag}(\mathbf{\Phi}, C)$ induces a deformation of $\mathrm{Sch}(\mathbf{\Phi})$ that is the image of a deformation $\widetilde{\mathbf{\Phi}}$ of $\mathbf{\Phi}$. The assumption that the deformation of $\mathrm{Sch}(\mathbf{\Phi})$ arises from a deformation of $\mathrm{Flag}(\mathbf{\Phi}, C)$ means that there is also an associated deformation of C. Since $\mathbf{\Phi}$ is an isomorphism onto its image in a neighborhood of C, this deformation of C canonically lifts to give a calibration of $\widetilde{\mathbf{\Phi}}$. ■

4. Comparison morphisms. In subsection 4.1 we compare Cam_n and CalCam_n by the natural decalibration morphism. In subsection 4.2 we focus on the case of two cameras, leading to a 2-1 cover of the essential variety that compactifies the twisted pair covering. Finally, in subsection 4.3 we state how both moduli spaces of cameras map to appropriate Hilbert schemes.

4.1. The decalibration morphism $\nu_n : \mathrm{CalCam}_n \to \mathrm{Cam}_n \times \mathscr{C}^n$. In this section, we study a natural morphism

$$\mathrm{CalCam}_n \to \mathrm{Cam}_n \times \mathscr{C}^n$$

given by forgetting the camera calibration datum.

Definition 4.1. *The* decalibration morphism *is the morphism*

$$\nu_n : \mathrm{CalCam}_n \to \mathrm{Cam}_n \times \mathscr{C}^n$$

given by sending $(\mathbf{\Phi}, C)$ *to* $\mathbf{\Phi}$*.*

4.1.1. Intersections of conic cones. Before we delve into the geometry of ν_n, we need a few preliminaries about intersections of conic cones in $\mathbf{P}^3$.

Proposition 4.2. *Let* X_1 *and* X_2 *be two conic cones in* $\mathbf{P}^3$ *with distinct cone points* P_1 *and* P_2*. Suppose* $C \subset X_1 \cap X_2$ *is a plane curve of degree* 2*, so that* $X_1 \cap X_2 = C \cup D$ *with* D *a curve of degree* 2*. Then* D *must be planar and have support distinct from the support of* C*. More precisely, one of the following must occur.*

1. *C and D are smooth conics meeting at two distinct points.*
2. *C is a smooth conic and D is a doubled planar line.*
3. *C is a doubled planar line and D is a smooth conic.*

In particular, we can never have $C = D$ (i.e., $X_1 \cap X_2$ cannot be a doubled smooth conic).

Proof. This is a standard result, and it can be extracted from the material in [8, Chapter 13, section 11]. We briefly describe a proof in modern language for the reader's convenience. By assumption, C is either a smooth conic or a planar doubled line. It is easy to write down examples where the intersection $X_1 \cap X_2$ is a union of two smooth conics meeting at two points (e.g., in characteristic different from 2 the pair $X^2 + Y^2 + Z^2 = 0$ and $Y^2 + Z^2 + W^2 = 0$ is such an example).

If $X_1 \cap X_2$ contains a doubled planar line, then X_1 and X_2 must be tangent along a ruling. Since $P_1 \neq P_2$, the residual curve must be a smooth conic.

Suppose $X_1 \cap X_2 = C \cup D$ with C a smooth conic and D a singular curve. We wish to show that D is a doubled planar line. Since D has degree 2 in $\mathbf{P}^3$, it must be the case that the reduced structure on D is a line. The only doubled lines contained in a conic cone are planar: They are given by intersecting with the tangent plane along rulings.

It remains to rule out the possibility that $X_1 \cap X_2$ is a doubled conic. Note that a doubled conic is the intersection of X_1 with a doubled plane $2P \in \mathscr{O}_{\mathbf{P}^3}(2)$. We can rule out this case if we can show that the pencil spanned by X_1 and a doubled plane not containing its cone point does not contain any more conic cones. We can represent the cone X_1 and an aribtrary doubled plane missing the cone point by the matrices

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} a^2 & ab & ac & a \\ ab & b^2 & bc & b \\ ac & bc & c^2 & c \\ a & b & c & 1 \end{pmatrix}$$

for $a, b, c \in k$. Searching for a conic cone in the pencil corresponds to finding λ such that the following matrix has rank 3:

$$\begin{pmatrix} a^2+\lambda & ab & ac & a \\ ab & b^2+\lambda & bc & b \\ ac & bc & c^2+\lambda & c \\ a & b & c & 1 \end{pmatrix} \text{ with row reduction } \begin{pmatrix} \lambda & 0 & 0 & 0 \\ 0 & \lambda & 0 & 0 \\ 0 & 0 & \lambda & 0 \\ a & b & c & 1 \end{pmatrix},$$

but the latter matrix can never have rank 3. ■

4.1.2. The geometry of ν_n. Fix a point ξ of $\mathrm{Cam}_n \times \mathscr{C}^n$. That is, fix conics $C_1, \ldots, C_n$ in $\mathbf{P}^2$ and a multiview configuration $\mathbf{\Phi}$. In this section we compute the fiber of ν_n over ξ.

Proposition 4.3. *The scheme-theoretic fiber $\nu_n^{-1}(\xi)$ is a reduced $\kappa(\xi)$-scheme of length at most 2.*

Proof. The fiber $\nu_n^{-1}(\xi)$ is precisely the scheme of smooth conics in the intersection of the cones over the image conics C_i inside the ambient $\mathbf{P}^3$. The result is thus immediate from Proposition 4.2. (In particular, the lack of doubled conic means that the fibers are discrete.) ■

Corollary 4.4. *The morphism ν_n is unramified.*

Proof. The proof is an immediate consequence of Proposition 4.3. ■

Proposition 4.5. *The morphism ν_n is proper.*

Proof. Suppose we have a multiview configuration $\mathbf{\Phi}$ of length 2 over a complete dvr R with fraction field K, degree 2 curves $C_1, \ldots, C_n \subset \mathbf{P}^2_R$, and a degree 2 curve $C_K \subset \mathbf{P}^3_K$ such that $\mathbf{\Phi}_K$ maps C_K isomorphically to the generic fiber of each C_i. By the valuative criterion for properness, it suffices to extend C_K to a degree 2 curve C_R.

Assume we have a multiview configuration $\mathbf{\Phi}$ of length 2 over a complete dvr R with fraction field K, and suppose we have conics $C_1, \ldots, C_n \subset \mathbf{P}^2_{\underline{R}}$ in each image plane. Write $\overline{C}_i \subset \mathbf{P}^3$ for the cone over C_i under $\mathrm{pr}_i \circ \mathbf{\Phi}$ and $I = \overline{C}_1 \cap \cdots \cap \overline{C}_n$. Finally, assume that there is a conic $C_K \subset \mathbf{P}^3_K$ such that $\mathbf{\Phi}_K$ maps C_K isomorphically to the generic fiber of each C_i; that is, $C_K \subset I_K$. Let C_R be the specialization of C_K in the closed fiber C_0. The curve C_R is degree 2, giving us a calibrated configuration over R. ∎

Note that even if C_k is a nondegenerate conic, C_0 need not be. This is why we need to add degenerate conics.

Proposition 4.6. *The morphism ν_2 has smooth image and general fiber of length 2. For any $n > 2$ the morphism ν_n is generically injective.*

Proof. The projective closure of the image of a fiber of CalCam_2 over $\mathscr{C}^2$ under ν_2 is known as the "essential variety," and its singularities are well known (see [3, section 2.1]); none of its singular points lie in the image of ν_2. To study the general fiber, it suffices by the irreducibility of all spaces involved to produce a single example of a camera configuration of length 2 such that the fiber of ν_2 has length 2. To do this, it further suffices to find a single example of two conic cones $C_1, C_2 \subset \mathbf{P}^3$ whose intersection is a pair of smooth conics. One such example is given by the cones $X^2 + Y^2 + Z^2 = 0$ and $Y^2 + Z^2 + W^2 = 0$.

We now show that ν_n is generically injective for $n > 2$. Given a smooth conic C in $\mathbf{P}^3$, the locus in $|\mathscr{O}_{\mathbf{P}^3}(2)|$ consisting of conic cones containing C is 3-dimensional (since such a cone is determined by its vertex). Thus, we can find three noncollinear conic cones that contain any given smooth conic C. On the other hand, given two conic cones C_1, C_2, the set of conic cones that vanish on their entire intersection $C_1 \cap C_2$ is contained in the pencil spanned by C_1 and C_2. We conclude that if $C_1 \cap C_2$ is reducible, then we can choose general cones $C_3, \ldots, C_n$ containing a smooth conic in $C_1 \cap C_2$ such that C_i is not in the pencil spanned by C_1 and C_2 for each $i > 2$. The joint vanishing locus $C_1 \cap C_2 \cap C_3 \cap \cdots \cap C_n$ is a smooth conic. Since this is generic behavior, this shows that ν_n is generically injective for all $n > 2$. ∎

It is a potentially interesting problem to characterize the locus over which ν_n is not injective, and the singular locus of its image (the "variety of calibrated n-focal tensors," which is studied for $n = 3$ in coordinatized form in [11]).

Corollary 4.7. *The morphism ν_n is finite.*

Proof. We have shown that ν_n is quasi-finite and proper and, thus, finite. ∎

Question 4.8. *Is the singular locus of the image of* CalCam_n*, for $n > 2$, equal to the locus over which the fiber of ν_n has length 2?*

4.2. Twisted pairs and moduli. In this section we study the morphism ν_2 in more detail, showing how the Hilbert scheme gives a natural compactification of the classical "twisted pair" construction. To explicitly compare this new treatment with the literature, in this section we will fix the calibrating conics to be $v(x_0^2 + x_1^2 + x_2^2) \subset \mathbf{P}^2$. Also, we will often think of an essential matrix as the corresponding pair of calibrated cameras in normalized coordinates. In these coordinates we can fix notation $P_1 = [I|0]$ and $P_2 = R[I|t]$, where $t = (a, b, c)$.

4.2.1. Twisted pairs. As shown in section 5.2 of [13], the locus $\mathcal{M}$ of essential matrices is smooth (over $\mathbf{C}$) and admits an étale surjection $\mathrm{SO}(3) \times \mathbf{P}^2 \to \mathcal{M}$, coming from composing a camera with a rotation and a translation, up to scaling. In terms of matrices we send (R, t) to the camera pair $P = [I|0], Q = [R|t]$ which has essential matrix $[t]_\times R$. One can check in local coordinates that the map is étale [2, Proposition 3.2].

For any *real* essential matrix $M \in \mathcal{M}(\mathbf{R})$, the fiber of π over M contains two points: One can take a pair of cameras P_1, P_2 and replace it with the pair $P_1, \widetilde{P}_2$ where $\widetilde{P}_2$ results from rotating P_2 by 180 degrees around the axis connecting the centers of P_1 and P_2. In normalized coordinates, the matrix

$$R_t = \begin{pmatrix} 2a^2-1 & 2ab & 2ac & 0 \\ 2ab & 2b^2-1 & 2bc & 0 \\ 2ac & 2bc & 2c^2-1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

is rotation by 180 degrees and $\widetilde{P}_2 = R[I|t]R_t$. (Note that over the reals we can always rescale t so that $a^2 + b^2 + c^2 = 1$.) The pair $(P_1, P_2), (P_1, \widetilde{P}_2)$ is called a *twisted pair*; what we have described is a well-known construction in computer vision [6, Result 9.19]. The key thing to note is that the rotation construction described above *preserves calibrations* for real cameras. For complex cameras, things get more complicated, and for displacements (a, b, c) such that $a^2 + b^2 + c^2 = 0$, the corresponding transformation produces a new camera pair $(P_1, \widetilde{P}_2)$ for which $\widetilde{P}_2$ is no longer calibrated.

4.2.2. Compactification of the twisted pair construction. The morphism

$$\nu_2 : \mathrm{CalCam}_2 \to \mathrm{Cam}_2 \times \mathscr{C}^2$$

gives a double covering of a closed subscheme that generalizes the twisted pair covering of the essential variety. A point of CalCam_2 is the datum (P_1, P_2, C) where P_1 and P_2 are cameras and C is a planar curve of degree 2 contained in the intersection of the cones defined by the preimage of C_{univ} via P_1 and P_2. Proposition 4.2 tells us that this intersection must contain either another nondegenerate conic or a doubled line. In either case denote this other degree 2 curve by $\widetilde{C}$. The general fibers of ν_2 are the triples (P_1, P_2, C) and $(P_1, P_2, \widetilde{C})$.

This double covering agrees with the twisted pair covering on real points. In normalized coordinates $\widetilde{C}$ is defined by the simultaneous vanishing of

$$x^2 + y^2 + z^2 = 0 \quad \text{and} \quad (a^2 + b^2 + c^2)w - 2(ax + by + cz) = 0.$$

When $a^2 + b^2 + c^2 = 1$, as it must over $\mathbf{R}$ (up to scaling), one can check that changing coordinates on $\mathbf{P}^3$ via the automorphism

$$H = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -2a & -2b & -2c & 1 \end{pmatrix}$$

sends the triple $(P_1, \widetilde{P}_2, C)$ to the triple $(P_1, P_2, \widetilde{C})$.

However, over the complex numbers there exist essential matrices such that $a^2+b^2+c^2 = 0$. This is exactly the condition that $\widetilde{C}$ is a doubled line. In this situation the twisted pair construction fails because the camera $\widetilde{P}$ no longer has a trivial calibration. Mathematically speaking, we are really discussing the fact that the twisted pair morphism π, while always étale, is *not* finite. Allowing degenerate calibrations (doubled lines) extends the twisted pair morphism π to ν_2.

Proposition 4.9. *There exists a fixed-point free involution, $\chi : \mathrm{CalCam}_2 \to \mathrm{CalCam}_2$ over Cam_2 given by fixing the cameras and swapping calibrating curves. More precisely, $\nu_2 \circ \chi = \nu_2$.*

Proof. Given a pair of cameras $\mathbf{\Phi} \to \mathbf{P}^2 \times \mathbf{P}^2$ and smooth conics $D_1, D_2 \subset \mathbf{P}^2$, we can pull back to get two cones $X_1, X_2 \subset \mathbf{P}^3$. Let $F = X_1 \cap X_2$. Blowing up the camera centers, the strict transforms of these cones, $\tilde{X}_1, \tilde{X}_2 \subset \mathrm{Bl}_{Z_1,Z_2} \mathbf{P}^3$, are smooth surfaces in $\mathbf{P}^3$. The intersection is a relative effective Cartier divisor and $\tilde{X}_1 \cap \tilde{X}_2 \simeq F$ since the cone centers are distinct.

A point in CalCam_2 is a pair $(\mathbf{\Phi}, C)$ where C is a relative effective Cartier divisor contained in F. By [15, Tag 0B8V] there exists another relative effective Cartier divisor C' such that $C' + C = F$. Checking at a geometric point, Proposition 4.2 shows that C' is a degree 2 curve, and that no geometric point of CalCam_2 is fixed by χ. This argument is functorial and so induces the desired involution. Since χ only changes the calibrating conic we have $\nu_2 \circ \chi = \nu_2$. ∎

Theorem 4.10. *The morphism ν_2 factors as a finite étale morphism followed by a closed immersion.*

Proof. By Corollary 4.7, ν_2 is a finite morphism, and hence closed. This yields a factorization $\mathrm{CalCam}_2 \to Z \to \mathrm{Cam}_2$ with the second arrow a closed immersion and the first scheme-theoretically surjective. Let A be a strictly Henselian local ring and $\operatorname{Spec} A \to Z$ a morphism. The finiteness of ν_2 yields a diagram

$$\begin{tikzcd} \operatorname{Spec} B \arrow[r] \arrow[d, "\psi"] & \mathrm{CalCam}_2 \arrow[d, "\nu_2"] \\ \operatorname{Spec} A \arrow[r] & \mathrm{Cam}_2 \end{tikzcd}$$

By [15, Tag 04GH], B is the product of local Henselian rings. By Proposition 4.6, the general fibers of ψ are length 2, corresponding to the two possible calibrating conics, so $\operatorname{Spec} B \simeq \operatorname{Spec} B_1 \sqcup \operatorname{Spec} B_2$. By Corollary 4.4, ψ is unramified, and thus (by [15, Tag 04GL]) restricts to a closed embedding on each $\operatorname{Spec} B_i$.

$$\begin{tikzcd} \operatorname{Spec} B_i \arrow[r] \arrow[dr, hook] & \operatorname{Spec} B_1 \sqcup \operatorname{Spec} B_2 \arrow[r] \arrow[d, "\psi"] & \mathrm{CalCam}_2 \arrow[d] \\ & \operatorname{Spec} A \arrow[r] & \mathrm{Cam}_2 \end{tikzcd}$$

The involution described in Proposition 4.9 induces an isomorphism $f : \operatorname{Spec} B_1 \to \operatorname{Spec} B_2$. In other words, both components map isomorphically to the image, so ν_2 is étale over Z, as claimed. ∎

4.3. Morphisms to Hilbert schemes. The following describes the main result relating the moduli problems Cam_n and CalCam_n to Hilbert schemes. This gives the generalization of the results of [1, Theorem 6], leveraging the novel methods of this paper to give more information about the uncalibrated case and the appropriate result in the calibrated case.

Proposition 4.11. *The associations*

$$\mathbf{\Phi} \mapsto \mathrm{Sch}(\mathbf{\Phi})$$

and

$$(\mathbf{\Phi}, C) \mapsto \mathrm{Flag}(\mathbf{\Phi}, C)$$

define monomorphisms

$$\mathrm{Sch} : \mathrm{Cam}_n \to \mathrm{Hilb}_{(\mathbf{P}^2)^n / \operatorname{Spec} \mathbf{Z}[1/2]}$$

and

$$\mathrm{Flag} : \mathrm{CalCam}_n \to \mathrm{Hilb}_{C^n_{\mathrm{univ}} \subset (\mathbf{P}^2)^n / \mathscr{C}^n}$$

such that

1. *when $n > 2$, the morphism* Sch *(resp.,* Flag*) itself is an open immersion into* $\mathrm{Hilb}^{\mathrm{sm}}_{(\mathbf{P}^2)^n / \operatorname{Spec} \mathbf{Z}[1/2]}$ *(resp.,* $\mathrm{Hilb}^{\mathrm{sm}}_{C^n_{\mathrm{univ}} \subset (\mathbf{P}^2)^n / \mathscr{C}^n}$*);*
2. *the arrows* Sch *and* Flag *together with the forgetful maps give a commutative diagram*

$$\begin{tikzcd}
\mathrm{CalCam}_n \arrow[r, "\mathrm{Flag}"] \arrow[d, "\nu_n"'] & \mathrm{Hilb}_{C^n_{\mathrm{univ}} \subset (\mathbf{P}^2_{\mathscr{C}^n})^n / \mathscr{C}^n} \arrow[d] \\
\mathrm{Cam}_n \times_{\operatorname{Spec} \mathbf{Z}[1/2]} \mathscr{C}^n \arrow[r, "\mathrm{Sch}"'] & \mathrm{Hilb}_{(\mathbf{P}^2_{\mathscr{C}^n})^n / \mathscr{C}^n}
\end{tikzcd}$$

In particular, every geometric fiber of Sch *over* $\operatorname{Spec} \mathbf{Z}[1/2]$ *is an open immersion of* Cam_n *into the smooth locus of a single irreducible component of the Hilbert scheme, and similarly for geometric fibers of* Flag *and components of the diagram Hilbert scheme.*

Proof. Propositions 2.33 and 2.35 show that Flag is a well-defined monomorphism. Since CalCam_n is smooth over $\mathscr{C}^n$, we have that Flag is an open immersion in a neighborhood of any point where it induces an isomorphism of deformation functors. Theorem 3.21 then applies to give the two desired statements. ∎

5. Questions.

In this section, we briefly discuss questions raised by this work and suggest some directions for future investigation.

Question 5.1. *What concrete computational consequences follow from functorial methods?*

We believe that the techniques described here may be useful for studying the numerical properties of multiview geometry. For example, in [12], we will give an explicit equation for the fiber of CalCam_2 over the pair of standard Euclidean conics, which appears as a double cover of the essential variety extending the twisted pair construction. It is given by the vanishing of a single bilinear form on $\mathbf{P}^3 \times \mathbf{P}^3$. This can be used to rederive the main results of [2], and to rephrase the five-point algorithm in terms of intersections of six bilinear forms in $\mathbf{P}^3 \times \mathbf{P}^3$ instead of the nine Demazure cubics and five linear forms. This is also related to the results of [3], but the derivations are completely different and independent of [2] (which is used in an essential way in [3]).

Question 5.2. *What is the correct boundary for* Cam_n *(resp.,* CalCam_n*)?*

Is there a extension of our moduli theory to handle degenerate configurations, where camera centers collide? Should these models include degenerations of image planes along the lines of Hacking's approach [5]? Is there a good moduli theory for pairs (X, C) consisting of a threefold with an embedded curve? These might be useful for studying degenerations of the ambient space together with its calibrating curve.

Question 5.3. *What is the right general formulation of Carlsson–Weinshall duality?*

Carlsson–Weinshall duality is somewhat mysterious from the point of view taken here. One can think about it in terms of birational isomorphisms of universal correspondences. It would be interesting to get a deeper understanding of this phenomenon.

Acknowledgments. We had interesting and helpful conversations with many people during the course of this work: Sameer Agarwal, Roya Beheshti, Dustin Cartwright, Charles Godfrey, Richard Hartley, Jonathan Hauenstein, Fredrik Kahl, Joe Kileel, Irina Kogan, Luke Oeding, Peter Olver, Brian Osserman, Tomas Pajdla, Jean Ponce, Jessica Sidman, Bernd Sturmfels, Rekha Thomas, Matthew Trager, and Bianca Viray. Rekha Thomas gave especially valuable remarks that helped us significantly improve our exposition. We benefitted greatly from the Berlin Algebraic Vision meeting in October 2015, hosted at TU Berlin, and the AIM meeting on Algebraic Vision in May 2016, held in San Jose, CA. Last, we thank the referees and editors for patiently giving us numerous helpful suggestions and comments.

REFERENCES

[1] C. Aholt, B. Sturmfels, and R. Thomas, *A Hilbert scheme in computer vision*, Canad. J. Math., 65 (2013), pp. 961–988, https://doi.org/10.4153/CJM-2012-023-2.

[2] M. Demazure, *Sur deux problemes de reconstruction*, Tech. Report RR-0882, INRIA, Rocquencourt, France, 1988, https://hal.inria.fr/inria-00075672.

[3] G. Fløystad, J. Kileel, and G. Ottaviani, *The Chow form of the essential variety in computer vision*, J. Symbolic Comput., 86 (2018), pp. 97–119, https://doi.org/10.1016/j.jsc.2017.03.010.

[4] A. Grothendieck, *Éléments de géométrie algébrique.* III. *Étude cohomologique des faisceaux cohérents.* I, Inst. Hautes Études Sci. Publ. Math., 11 (1961), pp. 5–167, http://www.numdam.org/item/PMIHES_1961__11__5_0/.

[5] P. Hacking, *Compact moduli of plane curves*, Duke Math. J., 124 (2004), pp. 213–257, https://doi.org/10.1215/S0012-7094-04-12421-2.

[6] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, Cambridge University Press, Cambridge, 2003.

[7] B. Hassett and S. J. Kovács, *Reflexive pull-backs and base extension*, J. Algebraic Geom., 13 (2004), pp. 233–247, https://doi.org/10.1090/S1056-3911-03-00331-X.

[8] W. V. D. Hodge and D. Pedoe, *Methods of Algebraic Geometry. Vol.* II. *Book* III: *General Theory of Algebraic Varieties in Projective Space, Book* IV: *Quadrics and Grassmann Varieties*, Cambridge Mathematical Library, Cambridge University Press, Cambridge, 1994, reprint of the 1952 original, https://doi.org/10.1017/CBO9780511623899.

[9] D. Huybrechts, *Fourier-Mukai Transforms in Algebraic Geometry*, Oxford University Press, Oxford, 2006.

[10] L. Illusie, *Complexe cotangent et déformations.* I, Lecture Notes in Math. 239, Springer-Verlag, Berlin, New York, 1971.
[11] J. Kileel, *Minimal problems for the calibrated trifocal variety*, SIAM J. Appl. Algebra Geom., 1 (2017), pp. 575–598, https://doi.org/10.1137/16M1104482.
[12] M. Lieblich, L. Van Meter, and B. Viray, *A new approach to the essential variety*, in preparation.
[13] S. Maybank, *Theory of Reconstruction from Image Motion*, Springer Ser. Inform. Sci. 28, Springer-Verlag, Berlin, 1993, https://doi.org/10.1007/978-3-642-77557-4.
[14] E. Sernesi, *Deformations of Algebraic Schemes*, Grundlehren Math. Wiss. 334, Springer-Verlag, Berlin, 2007.
[15] *The Stacks project*, http://stacks.math.columbia.edu, 2017.
[16] M. Trager, M. Hebert, and J. Ponce, *The joint image handbook*, in Proceedings of the 2015 IEEE International Conference on Computer Vision, IEEE, 2015, pp. 909–917.