Dissecting the Image of the Absolute Conic*

Imran Junejo and Hassan Foroosh
University of Central Florida, Orlando, FL 32816, U.S.A.

Abstract

In this paper, we revisit the role of the image of the absolute conic (IAC) in recovering the camera geometry. We derive new constraints on IAC that advance our understanding of its underlying building blocks. The new constraints are shown to be intrinsic to IAC, rather than exploiting the scene geometry or the prior knowledge on the camera. We provide geometric interpretations for these new intrinsic constraints, and show their relations to the invariant properties of the IAC. This in turn provides a better insight into the role that IAC plays in determining the camera internal geometry. Since the new constraints are invariant properties of the IAC, they can be used to reparameterize its elements. We show that such reparameterization would allow to recover a more general camera geometry from a single view, compared to existing methods. We apply the new constraints to single view calibration using vanishing points, investigate the error resilience, and compare our results to Liebowitz-Zisserman (1999).

1 Introduction

The absolute conic, $\mathbf{\Omega}_\infty$, and its perspective projection $\boldsymbol{\omega}$, known as the Image of the Absolute Conic (IAC), are among the most important concepts in defining the camera geometry. The importance of $\mathbf{\Omega}_\infty$ arises from the fact that it lies on the plane at infinity, $\mathbf{\Pi}_\infty$, and hence is invariant under Euclidean transformations. This implies that the relative position of $\mathbf{\Omega}_\infty$ with respect to a moving camera is fixed. As a result its image, IAC, remains fixed if the camera internal parameters do not vary. Therefore IAC can be used as a calibration object, i.e. for recovering the intrinsic camera parameters. Knowing the IAC, the camera pose **?**, and the Euclidean geometry of the scene **?** can be recovered directly from image measurements up to a similarity.

*The support of the National Science Foundation(NSF) through the Grant # IIS-0644280 is gratefully acknowledged.

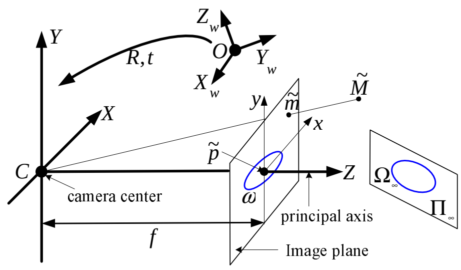

Figure 1. The geometry of a pinhole camera. The absolute conic Ω is a conic on the plane at infinity that is projected into the image plane as the conic ω, which depends only on the intrinsic parameters of the camera.

In this paper, we revisit the role of IAC in determining the camera geometry, and propose new constraints that are intrinsic to it, reflecting its invariant features. We investigate the application of these new constraints in camera calibration. We show that a more general camera model than the one proposed by **?** and formalized in **?** can be recovered from a single view, given an input of three orthogonal vanishing points.

In the next section, we first recall some preliminary notions on the relation between the camera geometry and the IAC. We then dissect the IAC into its constituent components, and provide their geometric meaning and importance. This is followed by an extensive set of experimentations and evaluation of the performance of calibration under noise, and experimental results on real data and comparison with **?**.

2 The Role of IAC

The geometry of imaging the absolute conic in a pinhole camera is shown in Figure **??**. The general pinhole camera projects a 3D point $\tilde{\mathbf{M}}$ to an image point $\tilde{\mathbf{m}}$ via

$$\tilde{\mathbf{m}} \sim \mathbf{KR}\left[\mathbf{I}|-\mathbf{C}\right]\tilde{\mathbf{M}}, \quad \mathbf{K} = \begin{bmatrix} f & s & u_o \\ 0 & \lambda f & v_o \\ 0 & 0 & 1 \end{bmatrix}, \qquad (1)$$

where $\sim$ implies equality up to an unknown non-zero scale factor, $\mathbf{R}$ is the rotation matrix from the world coordinate frame to the camera coordinate frame, $\mathbf{C}$ is the camera projection center, and $\mathbf{K}$ is the camera intrinsic matrix containing the focal length f, the aspect ratio λ, the skew s, and the principal point $\tilde{\mathbf{p}} \sim [u_o \; v_o \; 1]^T$.

The role of IAC in defining the camera geometry is better understood by examining the action of a finite camera on points that lie on the plane at infinity $\boldsymbol{\pi}_\infty$. A point on $\boldsymbol{\pi}_\infty$ can be written as $\tilde{\mathbf{M}}_\infty \sim [\mathbf{d}^T \; 0]^T$, where the 3-vector $\mathbf{d}$ defines the direction of the ray obtained by connecting the image of $\tilde{\mathbf{M}}_\infty$ and the camera projection center. Substituting $\tilde{\mathbf{M}}_\infty$ in (**??**), one can readily verify that $\tilde{\mathbf{m}}_\infty \sim \mathbf{KR}\tilde{\mathbf{M}}_\infty$. It therefore follows that the absolute conic, which is the conic $\boldsymbol{\Omega}_\infty = \mathbf{I}$ on $\boldsymbol{\pi}_\infty$ maps to the image conic

$$\boldsymbol{\omega} \sim (\mathbf{KR})^{-T}\mathbf{I}(\mathbf{KR})^{-1} \sim \mathbf{K}^{-T}\mathbf{K}^{-1} \quad (2)$$

known as the image of the absolute conic (IAC). Conversely, two image points $\tilde{\mathbf{m}}_1$ and $\tilde{\mathbf{m}}_2$ back-project to two rays with directions $\mathbf{d}_1 = \mathbf{K}^{-1}\tilde{\mathbf{m}}_1$ and $\mathbf{d}_2 = \mathbf{K}^{-1}\tilde{\mathbf{m}}_2$ in the camera coordinate system, where the angle between the two rays is given by the familiar cosine formula

$$\cos\theta = \frac{\mathbf{d}_1^T\mathbf{d}_2}{\sqrt{\mathbf{d}_1^T\mathbf{d}_1}\sqrt{\mathbf{d}_2^T\mathbf{d}_2}} = \frac{\tilde{\mathbf{m}}_1^T\boldsymbol{\omega}\tilde{\mathbf{m}}_2}{\sqrt{\tilde{\mathbf{m}}_1^T\boldsymbol{\omega}\tilde{\mathbf{m}}_1}\sqrt{\tilde{\mathbf{m}}_2^T\boldsymbol{\omega}\tilde{\mathbf{m}}_2}} \quad (3)$$

This shows that known angles between vanishing points can be used to impose constraints on the IAC to obtain the camera intrinsic matrix. For instance, given the images of three infinite points $\tilde{\mathbf{v}}_i, i = 1, ..., 3$ along known directions, and assuming zero skew and unit aspect ratio, one can recover the remaining unknown camera intrinsic parameters. In particular if $\tilde{\mathbf{v}}_i$ are the vanishing points along three orthogonal directions then one can write three linear equations of the form $\tilde{\mathbf{v}}_i^T\boldsymbol{\omega}\tilde{\mathbf{v}}_j = 0, i \neq j$ to calibrate the camera **???????**. This is essentially the core idea behind calibration using the vanishing points, which was formalized by **?**. These works showed that only a simplified camera model with three unknown intrinsic parameters can be recovered from the vanishing points of three orthogonal directions, unless additional information is available (e.g. more images or the circular points **??**).

Generally speaking, in recovering the camera geometry from a single view three sources of information have been commonly used in the past to impose constraints on the image of the absolute conic $\boldsymbol{\omega}$:

- metric information about a plane with a known world-to-image homography;
- vanishing points and lines corresponding to known (usually orthogonal) directions and planes;
- *a priori* constraints, such as unit aspect ratio, or zero skew.

These constraints are summarized by Hartley and Zisserman (Table 8.1, page 224 in **?**).

In this paper, we re-examine the problem of recovering the camera geometry from a single view, when three vanishing points along world orthogonal directions are known.

3 Dissecting IAC

In the existing literature on camera calibration the role of IAC is primarily investigated in terms of its relationship with other geometric entities in the image plane, i.e. the vanishing points and the vanishing line. The relation between IAC and the internal parameters is often limited to equation (**??**). In this section and the following one, we present some new constraints and their geometric meaning that are more intrinsic to the IAC, i.e. relate to the internal geometry of camera.

Theorem 3.1 *(Invariance)*
Let $\boldsymbol{\omega}$ be the image of the absolute conic. The principal point $\tilde{\mathbf{p}}$ satisfies

$$\boldsymbol{\omega}\tilde{\mathbf{p}} \sim \mathbf{l}_\infty \quad (4)$$

where $\mathbf{l}_\infty \sim [0 \; 0 \; 1]^T$ is the line at infinity.

The proof is straightforward and follows by performing the Cholesky factorization of the Dual Image of the Absolute Conic (DIAC), ω^*, and direct substitution of $\tilde{\mathbf{p}}$.

In the next section, we also provide an alternative proof, which reveals the geometric meaning of the constraint in (**??**). Note that only two constraints can be derived from (**??**) in the image plane, due to scale ambiguity.

Proposition 3.1 *(Scale)*
Let $\boldsymbol{\omega}$, denote the image of the absolute conic. We have

$$|\omega_{33}| \;\; \tilde{\mathbf{p}}^T\boldsymbol{\omega}\tilde{\mathbf{p}} - \mathtt{det}(\boldsymbol{\omega}) = 0 \quad (5)$$

where $|\omega_{33}|$ denotes the minor of IAC corresponding to its last component, and $\mathtt{det}(\cdot)$ is the determinant.

Proposition 3.2 *(Ortho-Invariance)*
Let $\tilde{\mathbf{v}}_i, i = 1, ..., 3$ denote three vanishing points along mutually orthogonal directions. The image of the absolute conic relates these vanishing points via

$$\sum_i \frac{1}{\tilde{\mathbf{v}}_i^T\boldsymbol{\omega}\tilde{\mathbf{v}}_i} - \frac{1}{\tilde{\mathbf{p}}^T\boldsymbol{\omega}\tilde{\mathbf{p}}} = 0 \quad (6)$$

Proofs for all the above results follow by using the Cholesky decomposition of DIAC, $\boldsymbol{\omega}^*$, and direct substitution and algebraic simplification. Note that the result in (**??**) depends on the orthogonality conditions, and hence is dependent on the familiar linear orthogonality constraints $\tilde{\mathbf{v}}_i^T\boldsymbol{\omega}\tilde{\mathbf{v}}_j = 0, i \neq j$. However, (**??**) and (**??**) reflect some intrinsic properties of the IAC and do not depend on the scene geometry or the prior knowledge on the camera intrinsics. This is the key idea presented in this paper.

3.1 Geometric Interpretation

The result in (**??**) is better understood if we provide its geometric interpretation. Clearly, (**??**) is independent of the image points, i.e. for any two points $\tilde{\mathbf{m}}_1$ and $\tilde{\mathbf{m}}_2$, we have $\tilde{\mathbf{p}}^T \boldsymbol{\omega} \tilde{\mathbf{m}}_1 = \tilde{\mathbf{p}}^T \boldsymbol{\omega} \tilde{\mathbf{m}}_2$. Therefore, it reflects some intrinsic property of the IAC. This intrinsic property is better understood if we rewrite (**??**) as the following two independent constraints

$$\tilde{\mathbf{p}}^T \boldsymbol{\omega}_1 = 0 \tag{7}$$

$$\tilde{\mathbf{p}}^T \boldsymbol{\omega}_2 = 0 \tag{8}$$

where $\boldsymbol{\omega}_i$ are the rows of the IAC (or equivalently its columns due to symmetry).

This shows that

$$\tilde{\mathbf{p}} \sim \boldsymbol{\omega}_1 \times \boldsymbol{\omega}_2 \tag{9}$$

which is true for a general camera model, i.e. no particular assumptions made about the aspect ratio, or the skew.

A geometric interpretation (see Figure **??**) of this result is that the two rows $\boldsymbol{\omega}_1$ and $\boldsymbol{\omega}_2$ of the IAC correspond to two lines in the image plane that always intersect at the principal point regardless of the other intrinsic parameters. We may consider three cases: i.e. varying the skew s, the aspect ratio λ, or the focal length f. Although it is highly unlikely for a CCD camera to change the skew or the aspect ratio, it is useful to evaluate these effects on calibrating a general pinhole camera or a simplified one.

Varying the skew s: We may assume that we deal with two identical cameras that differ only in skew: one zero skew and the other non-zero. Let us denote the two corresponding IAC's by

$$\boldsymbol{\omega} \sim \begin{bmatrix} \boldsymbol{\omega}_1 \\ \boldsymbol{\omega}_2 \\ \boldsymbol{\omega}_3 \end{bmatrix} \quad \text{and} \quad \boldsymbol{\omega}_s \sim \begin{bmatrix} \boldsymbol{\omega}'_1 \\ \boldsymbol{\omega}'_2 \\ \boldsymbol{\omega}'_3 \end{bmatrix} \tag{10}$$

where $\boldsymbol{\omega}_i$ and $\boldsymbol{\omega}'_i$, $i = 1, ..., 3$ are the rows of the corresponding IAC's.

For the IAC with zero skew, i.e. $\boldsymbol{\omega}$, the two lines $\boldsymbol{\omega}_1$ and $\boldsymbol{\omega}_2$ are parallel to the image x and y axes respectively, and intersect at the principal point. For the general IAC with non-zero skew, $\boldsymbol{\omega}'$, the corresponding two lines $\boldsymbol{\omega}'_1$ and $\boldsymbol{\omega}'_2$ are not perpendicular anymore. However, they still intersect at the same image point, i.e. the principal point.

To demonstrate this formally, note that

$$\begin{aligned} \boldsymbol{\omega}' &\sim \mathbf{K}'^{-T}\mathbf{K}'^{-1} \\ &\sim \mathbf{K}'^{-T}\mathbf{K}^{T}\mathbf{K}^{-T}\mathbf{K}^{-1}\mathbf{K}\mathbf{K}'^{-1} \\ &\sim \mathbf{H}_s^{-T}\boldsymbol{\omega}\mathbf{H}_s^{-1} \end{aligned} \tag{11}$$

Therefore the transformation that maps the IAC with zero skew to the general IAC is given by the homography

$$\mathbf{H}_s \sim \mathbf{K}'\mathbf{K}^{-1} \tag{12}$$

It can be shown that this homography is of the form

$$\mathbf{H}_s \sim \begin{bmatrix} 1 & -\frac{\omega_{12}}{\omega_{11}} & \frac{\omega_{12}}{\omega_{11}} v_o \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{13}$$

If we now perform the eigen-decomposition of $\mathbf{H}_s$, we will find that this homography has only two distinct eigenvectors both of which correspond to unit eigenvalues. The two eigenvectors are $[0\ v_o\ 1]^T$ and $[1\ 0\ 0]^T$. Geometrically, this is equivalent to saying that under the transformation $\mathbf{H}_s$ (i.e. if the skew of a camera changes from zero to a non-zero value), the point $[0\ v_o\ 1]^T$ and the vanishing point along the x-axis remain invariant. In other words, these are geometrically fixed points under $\mathbf{H}_s$. Since any linear combination of eigenvectors with a common eigenvalue is an eigenvector with the same eigenvalue, we deduce that the principal point $\tilde{\mathbf{p}}$, which lies on the line joining the two fixed points is also invariant under this transformation. This shows that equation (**??**) conveys an invariant property of the IAC, i.e. upon changing the skew the principal point should still lie on the intersection of the image lines defined by the first two rows of the IAC.

Another illuminating feature of $\mathbf{H}_s$ is that if we do the eigendecomposition of the transposed homography $\mathbf{H}_s^T$, we will find that there are also only two distinct eigenvectors, i.e. $[0\ 1\ -v_o]^T$ and $[0\ 0\ 1]^T$. Geometrically, this implies that the line $[0\ 1\ -v_o]^T$ and the line at infinity are invariant under changes in the skew. Since the principal point lies on the first line, it again confirms that the principal point is a fixed point under variations in the skew.

Varying aspect ratio λ: Interestingly enough, the same process as above can be used to establish that upon changing the aspect ratio λ, the principal point is also an invariant fixed point on the intersection of the two image lines defined by the first two rows of the IAC. Again, if the two IAC's are denoted by $\boldsymbol{\omega}$ and $\boldsymbol{\omega}'$, then their relationship is defined by a homography of the form

$$\mathbf{H}_\lambda \sim \mathbf{K}'\mathbf{K}^{-1} \tag{14}$$

$$\sim \begin{bmatrix} 1 & 0 & 0 \\ 0 & \lambda & v_o(1-\lambda) \\ 0 & 0 & 1 \end{bmatrix} \tag{15}$$

where $\lambda = \sqrt{\frac{\omega_{11}^2}{\omega_{11}\omega_{22} - \omega_{12}^2}}$.

Similar eigen-analysis reveals that $\mathbf{H}_\lambda$ shares the same two eigenvectors $[0\ v_o\ 1]$ and $[1\ 0\ 0]$ corresponding to its repeated unit eigenvalue, and a third eigenvector that corresponds to the point at infinity along the y-axis, i.e. $[0\ 1\ 0]$ with the eigenvalue equal to λ. This shows that the same two points are again geometric fixed points. However this time the infinite point along the y-axis is also fixed. Again

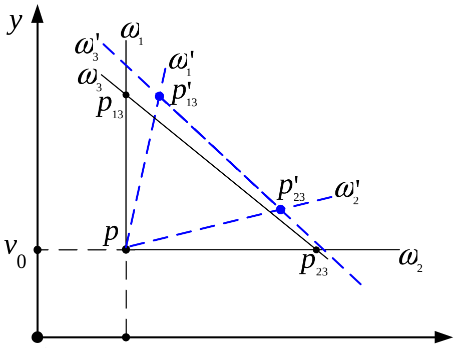

Figure 2. The geometry associated with the IAC: ω_1, ω_2, and ω_3 represent the lines associated with the IAC when the skew is zero, and ω_1', ω_2', and ω_3' illustrate the case when the skew is not zero. In both cases the principal point is located at the intersection of the first two lines, providing two linear constraints on the IAC. The ratio of line segments along the two lines (two rows) are preserved the skew changes.

using the fact that the linear combinations of eigenvectors corresponding to unit eigenvalues is also an eigenvector, we conclude that the principal point, which lies on the line joining the first two eigenvectors, is also geometrically a fixed point under variations of λ.

Varying the focal length f: Finally, if we let the focal length of a camera vary then the homography that relates the two IAC's is given by

$$\mathbf{H}_f \sim \mathbf{K}'\mathbf{K}^{-1} \quad (16)$$

$$\sim \begin{bmatrix} r & (r-1)\frac{\omega_{12}}{\omega_{11}} & (1-r)\left(u_o + v_o\frac{\omega_{12}}{\omega_{11}}\right) \\ 0 & r & (1-r)v_o \\ 0 & 0 & 1 \end{bmatrix} \quad (17)$$

where r is the ratio of two focal lengths.

The eigen decomposition of this homography indicates that the principal point is the eigenvector corresponding to the unit eigenvalue, and hence is a fixed point under $\mathbf{H}_f$. The last two eigenvectors, are repeated and correspond again to the point at infinity along the x-axis $[1\ 0\ 0]$, with the eigenvalue equal to r.

Remark If two cameras differ only by the intrinsic parameters s, λ, or f, then the corresponding IAC's, $\boldsymbol{\omega}$ and $\boldsymbol{\omega}'$, satisfy

$$\boldsymbol{\omega}_1 \times \boldsymbol{\omega}_2 \sim \boldsymbol{\omega}_1' \times \boldsymbol{\omega}_2' \quad (18)$$

Figure **??** illustrates this underlying geometry of IAC for the case of varying skew. As can be seen in Figure **??** the third row of IAC also corresponds to a line in the image plane which intersects the first two lines at two distinct points other than the principal point. These intersection points together with other points along the two lines $\boldsymbol{\omega}_1$ and $\boldsymbol{\omega}_2$ can be used to confirm that the ratio of line segments remain invariant, since all the homographies described above are affine. Unfortunately, the third row of IAC or the resulting invariant ratios do not provide new independent constraints.

Before we close this section, we also formalize the familiar constraint that the principal point is known *a priori* to be close to the centre of the image $\tilde{\mathbf{c}}$, as the following "soft constraint"

$$\hat{\mathbf{p}} = \arg\min(\tilde{\mathbf{p}} - \tilde{\mathbf{c}})^T(\tilde{\mathbf{p}} - \tilde{\mathbf{c}}) \quad (19)$$

This latter constraint is a very practical prior in self-calibration.

4 Single-View Calibration

The results of the previous section are both good news and bad news. The bad news is that we can not find more than two intrinsic constraints on the IAC from its internal geometry. The good news is that the two constraints that we find can be used to reprameterize the IAC. This is rather very useful, since it allows us to recover a more general camera model than the existing single-view calibration techniques such as **?**: e.g. recover f, s and (u_o, v_o) with three vanishing points, or recover f and (u_o, v_o) with two vanishing points.

For instance, let us assume that the camera skew is zero. The IAC is then of the form

$$\boldsymbol{\omega} \sim \begin{bmatrix} 1 & 0 & \omega_{13} \\ 0 & \omega_{22} & \omega_{23} \\ \omega_{13} & \omega_{23} & \omega_{33} \end{bmatrix} \quad (20)$$

Given three orthogonal vanishing points, we can formulate the single-view calibration problem as the solution to the following set of five equations:

$$\tilde{\mathbf{v}}_i^T \boldsymbol{\omega} \tilde{\mathbf{v}}_j = 0, \quad i \neq j, \quad i,j = 1,...,3 \quad (21)$$

$$\tilde{\mathbf{p}}^T \boldsymbol{\omega}_1 = 0, \quad (22)$$

$$\tilde{\mathbf{p}}^T \boldsymbol{\omega}_2 = 0 \quad (23)$$

These equations are linear in terms of the components of $\boldsymbol{\omega}$, and hence any four of them can be used to reparameterize $\boldsymbol{\omega}$ in (**??**) exclusively in terms of the principal point $\tilde{\mathbf{p}}$. Suppose we use the first four equations for reparameterization then the resulting $\boldsymbol{\omega}$, which depends only on $\tilde{\mathbf{p}}$ should minimize

$$\hat{\mathbf{p}} = \arg\min \tilde{\mathbf{p}}^T \mathbf{W} \tilde{\mathbf{p}} \quad \text{where} \quad \mathbf{W} = \boldsymbol{\omega}_2 \boldsymbol{\omega}_2^{\mathbf{T}} \quad (24)$$

We initialize $\hat{\mathbf{p}}$ at the centre of the image, and minimize using a standard optimization method (e.g. Levenberg-Marquardt) in a window around the centre of the image.

Once the principal point is obtained, all components of the IAC can be recovered (since they are expressed in terms of $\tilde{\mathbf{p}}$), and hence the camera intrinsic matrix $\mathbf{K}$ can be computed by the Cholesky decomposition of the DIAC, $\boldsymbol{\omega}^*$. Note that the method recovers a more general camera model of four unknown parameters, e.g. f, λ and (u_o, v_o).

The three columns of the rotation matrix are then given by $\mathbf{r}_i = \pm \frac{\mathbf{K}^{-1}\tilde{\mathbf{v}}_i}{\|\mathbf{K}^{-1}\tilde{\mathbf{v}}_i\|}$ - the sign ambiguity can be removed using the cheirality constraint **?**. The translation of the camera can also be recovered up to an unknown global scale, taking an image point as the projection of the world origin.

5 Results and Noise Resilience

In this section, we show an extensive set of experimental results on both synthetic and real data using the method described above. We have performed detailed experimentation on the effect of noise in the estimation error over substantial number of independent rials (i.e. 1000). The simulated camera has a focal length of $f = 2000$, the aspect ratio $\lambda = \frac{1015}{2000}$, zero skew, and the principal point at $(510, 385)$, for image size of 1024×768.

Performance Versus Noise Level: In this experimentation, we compared estimated camera intrinsic and extrinsic parameters against the ground truth, while adding a zero-mean Gaussian noise varying from 0.1 pixels to 1.5 pixels. The results show the average performance over 1000 independent trials. Figure **??** summarizes the results for intrinsic parameters. For noise level of 1.5 pixels, which is larger than the typical noise in practical calibration **?**, the relative error for the focal length f is 0.7%. The maximum relative error for the aspect ratio is less than 0.01%, while that of the principal point is less than 0.2%. Excellent performance is also achieved for all extrinsic parameters as shown in Figure **??**, i.e. less than 0.4% error for both t_x and t_y relative to f, and absolute errors of less than a tenth of a degree for all rotation angles.

Performance against ?: We performed the comparison using the same setup as above. Figure **??** summarizes our results. Interestingly, it appears that the reparameterization in terms of the principal point improves the performance of estimating the position of the principal point, and increase its noise resilience.

Performance on Real Data: For real data, in order to evaluate our results, we used an approach similar to **?** using the uncertainty associated with the estimated intrinsic parameters characterized by their standard deviation over many images. Figure **??** shows two examples from the set of real images that were used in this experimentation. Results are summarized in table **??**. The uncertainty is reasonable, but could be improved of course if we use more accurate approaches **????** to finding the vanishing points, rather than using an unreliable manual point clicking.

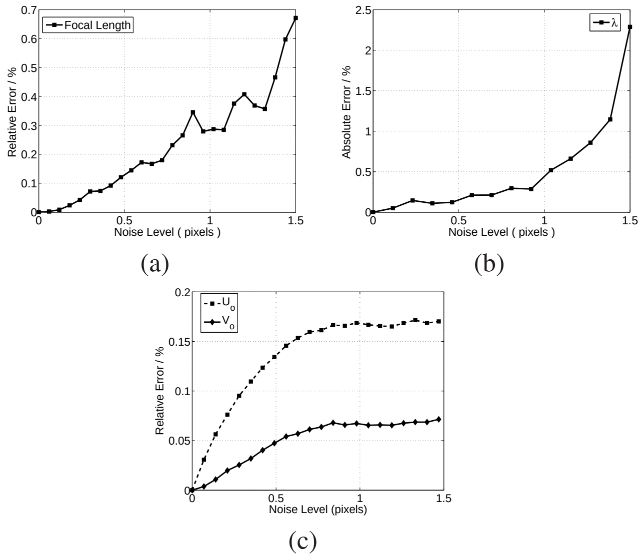

Figure 3. Performance vs noise (in pixels) averaged over 1000 independent trials: (a) relative error for the focal length f, (b) the relative error for the aspect ratio λ, and (c) the relative error in the coordinates of the principal point.

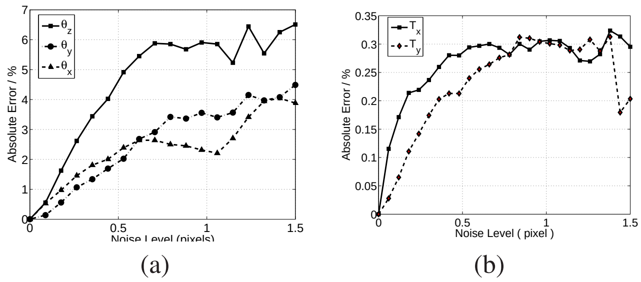

Figure 4. Performance vs noise (in pixels) averaged over 1000 independent trials: (a) absolute error for the rotation angles, (b) absolute error for the translations along x and y axes.

Parameter	Mean	Std.
f	460.52	5.74
λ	1.51	0.24
u_o	318.33	5
v_o	242.77	4.41

Table 1. Uncertainty in experimental results with real data. A set of ten images were used.

6 Conclusion

This paper presents new constraints that are intrinsic to the image of the absolute conic. The constraints reflect the invariant properties of the IAC, and characterize its geometric structure. In particular we show that the rows of the IAC correspond to very specific image lines whose intersections

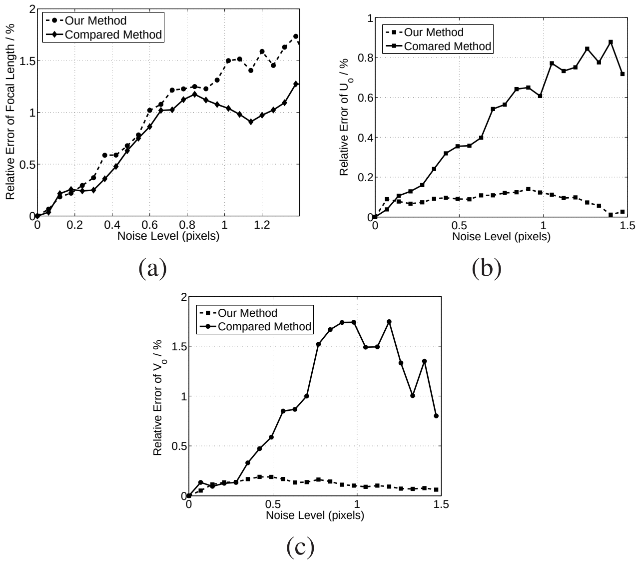

Figure 5. Performance vs ? averaged over 1000 independent trials: (a) relative error for the focal length f, (b) & (c) the relative error in the coordinates of the principal point.

Figure 6. Two of many images used in evaluation with real data

bear the invariant properties of the IAC. An immediate application of this geometric characterization of the IAC is that it can extend our ability to estimate more complete set of camera parameters from a single view. We therefore propose the following table as an addendum to table given by Hartley and Zisserman (Table 8.1, page 224 in **?**). Unfortunately, however, as described in the text, not all the constraints can be used independently. As a result, we believe that it is unlikely that one can recover all the five intrinsic parameters of the camera from a single view of three orthogonal vanishing points, unless some additional information is available.

Condition	Constraint	Linear
Invariance	$\boldsymbol{\omega}\tilde{\mathbf{p}} \sim \mathbf{l}_\infty$	yes
Scale	$\omega_{33}^{*}\ \tilde{\mathbf{p}}^T\boldsymbol{\omega}\tilde{\mathbf{p}} - 1 = 0$	no
Ortho-invariance	$\sum_i \frac{1}{\tilde{\mathbf{v}}_i^T\boldsymbol{\omega}\tilde{\mathbf{v}}_i} - \frac{1}{\tilde{\mathbf{p}}^T\boldsymbol{\omega}\tilde{\mathbf{p}}} = 0$	no
"Soft"	$\tilde{\mathbf{p}} \sim \arg\min(\tilde{\mathbf{p}} - \tilde{\mathbf{c}})^T(\tilde{\mathbf{p}} - \tilde{\mathbf{c}})$	no

Table 2. Intrinsic constraints of IAC. The first two are related to the invariant properties of the principal point, the third constraint cross-correlates this property and the orthogonality constraint (ortho-invariance), and the last one is a "soft constraint" on the position of the principal point in the image plane.