ECE 661
Purdue Univ
Avi Kak
Lecture 7

World 3D: Quadrics, Transformation Groups, and the Absolute Conic

Reference: "Multiple View Geometry in Computer Vision" by Hartley & Zisserman

- You need to understand the quadrics in 3D in order to grasp the Absolute Conic — an algebraic structure that plays an important role in modern camera calibration algorithms.

- Using homogeneous coordinates, a quadric is a surface in 3D that is defined by $\boxed{x^T Q x = 0}$ where Q is a 4x4 symmetric matrix of parameters (in very much the same way that C was a 3x3 symmetric matrix of parameters for a conic in 2D).

- As you saw in Lecture 2, when the rank of a conic, which is ordinarily 3, becomes 2, we end up with a degenerate conic. Along similar lines when the rank of Q becomes less than 4, we end up with a degenerate quadric.

- As you'd expect, the intersection of a plane Π with a quadric Q is a conic. As an example of the expressive power of homogeneous coordinates, it is interesting to see how easy it is to figure out what this conic is for a given Q and a given Π. But first you must express the points on Π in a parametric form as follows: Let A, B, and C be any three non-collinear points on Π. Any point X on Π can be expressed as a linear combination of the 4-vectors A, B, and C: $\boxed{X = uA + vB + wC}$. We can therefore write
$$\boxed{X = \underbrace{\begin{bmatrix} A & B & C \end{bmatrix}}_{4 \times 3} \begin{bmatrix} u \\ v \\ w \end{bmatrix} = M\vec{p}}$$
(M is 4x3)
Substituting for X in $\boxed{X^T Q X = 0}$, we get $\boxed{\vec{p}^T M^T Q M \vec{p} = 0}$. If we think of $\vec{p}$ as defining the homogeneous coordinates of a point on Π (obviously, with respect to the coordinate frame corresponding to the points A, B, and C), we see that the conic at the intersection of Π and Q is given by $\boxed{C = M^T Q M}$.
You may complain that the conic C is not directly in terms of the four parameters in $\Pi = \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix}$. But note that it is not too hard to specify the points A, B, C for a given Π. You just need to satisfy $\boxed{\Pi^T A = \Pi^T B = \Pi^T C = 0}$. One possible solution: $A = \begin{pmatrix} -b/a \\ 1 \\ 0 \\ 0 \end{pmatrix}$, $B = \begin{pmatrix} -c/a \\ 0 \\ 1 \\ 0 \end{pmatrix}$, $C = \begin{pmatrix} -d/a \\ 0 \\ 0 \\ 1 \end{pmatrix}$

- Let's now see how a quadric is transformed when the physical $\mathbb{R}^3$ space in which it is defined is transformed with a 4x4 homography H. So we are given $\boxed{x' = Hx}$. Substituting in $\boxed{x^T Q x = 0}$ gives us $\boxed{x'^T H^{-T} Q H^{-1} x' = 0}$ which we can express as $\boxed{x'^T Q' x' = 0}$, implying $\boxed{Q' = H^{-T} Q H^{-1}}$.

- The tangent plane to a quadric Q at its surface point x is given by $\boxed{\Pi = Qx}$. The proof of this is similar to the proof of $\boxed{l = Cx}$ being a tangent line to conic C at the perimeter point x.

- The dual of a quadric is the set of tangent planes that satisfy $\boxed{\pi^T Q^* \pi = 0}$ where Q^* is a 4x4 symmetric matrix that can be derived straightforwardly from the point quadric Q and its tangent planes. From $\boxed{\pi = QX}$ for the tangent plane at the surface point X for a point quadric Q, we get $\boxed{X = Q^{-1}\pi}$ for the point of contact at the surface. Substituting this point in $\boxed{X^T Q X = 0}$ gives us $\boxed{\pi^T Q^{-T} Q Q^{-1} \pi = 0}$, implying $Q^* = Q^{-T} = Q^{-1}$ since Q is symmetric. So, the dual quadric Q^* is the inverse of the point quadric Q.

- Earlier it was shown that when the physical $\mathbb{R}^3$ space is transformed by a 4x4 homography H, the point quadric Q is transformed into $H^{-T} Q H^{-1}$. Let's now see how the dual quadric transforms: The planes of a dual quadric must satisfy $\boxed{\pi^T Q^* \pi = 0}$. We do know that when $X' = HX$, we have $\boxed{\pi' = H^{-T}\pi}$, implying $\boxed{\pi = H^T \pi'}$. Substituting this expression for π in the identity $\boxed{\pi^T Q^* \pi = 0}$, we get $\boxed{\pi'^T H Q^* H^T \pi' = 0}$, implying that a dual quadric transforms as $\boxed{Q^{*\prime} = H Q^* H^T}$.

- We will now introduce the reader to the concepts of polar planes and polar lines, the former for the case of quadrics and latter for the case of conics. As you know, when X is a point on the surface of a quadric, $\boxed{\pi = QX}$ is the tangent plane at that point. However, when X is <u>outside</u> the quadric, the entity QX still has a useful geometric interpretation: It is the polar plane of the quadric with respect to point X.

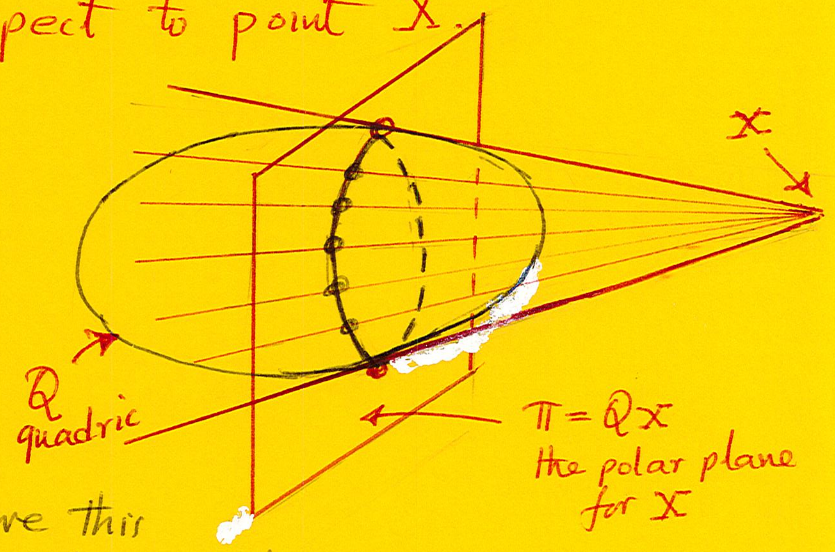

The tangent planes at the intersection of π and Q have one point in common — the point at X. Note that this property is also true for conics: If X is on the perimeter of the conic, CX is the tangent line to C at X. However, if X is <u>outside</u> the conic, CX is a polar line with respect to X as shown below. To prove this is elementary: Let the tangent line dropped from the outside point X to the conic meet the conic at point y. We know that the tangent line at y must be given by the parameters Cy. Since X is on this tangent line, it must be the case that $\boxed{(Cy)^T X = 0}$, implying that $\boxed{y^T C X = 0}$ since C is symmetric. Taking the transpose of both sides, we get $\boxed{(CX)^T y = 0}$ which says that y must be on a line whose parameters are given by $\boxed{l = CX}$, which is our polar line. Basically what we have proved is that every tangent from X to C — there can be only two of them — will touch the conic at a point that is on the polar line. The proof for polar planes is similar.

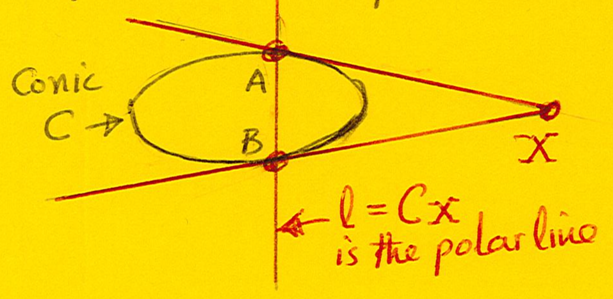

The Plane at Infinity

- The very important role played by l_∞ in 2D projectivity is played by the plane at infinity, Π_∞, for the case of 3D projectivity.
- In physical $\mathbb{R}^3$, all points at infinity lie on a single plane — the plane Π_∞. This assertion can be proved mathematically in a manner similar to our proof in Lecture 2 that all points at infinity in physical $\mathbb{R}^2$ lie on a single straight line (l_∞). You can also establish this result intuitively as follows: Imagine yourself setting out from the origin of the world 3-space in straight lines in different directions and traveling up to a distance R in each direction. The points thus reached would constitute a spherical surface of curvature $\frac{1}{R}$. As $R \to \infty$, the curvature would become zero and surface reached would become flat. The amazing thing is that this "flat" plane "folds on itself".
- Let $X = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}$ be the homogeneous coordinate representation of a point in physical $\mathbb{R}^3$. We represent a point at infinity by setting $x_4 = 0$. Such points will be called Ideal Points. Obviously, Π_∞ is the set of all ideal points in physical $\mathbb{R}^3$.
- In a manner analogus to the representation of l_∞ in 2D projectivity, we represent Π_∞ by $\begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$.

Transformations of World 3-Space

- In homogeneous coordinates, a general transformation of the physical $\mathbb{R}^3$ is achieved with a 4x4 homography: $\boxed{X' = HX}$.
- For the same reasons as for 2D projectivity, a 4x4 homography is best expressed as → $\boxed{X' = \begin{bmatrix} A & \vec{t} \\ \vec{v}^T & v \end{bmatrix} X}$ H

  Such transformations for all non-singular H form the Projective Group.

  Since only the ratios of the elements of H matter, H can be specified fully with 15 parameters. So we say that in the Projective Group H has 15 degrees of freedom.
- Using the same line of reasoning as in Lecture 3, we can show that the Projective Group maps straight lines into straight lines.
- The Affine Subgroup of the general Projective Group is specified by restricting the last row of H to $(0\ 0\ 0\ 1)$. Now we can write → $\boxed{X' = \begin{bmatrix} A & \vec{t} \\ \vec{0}^T & 1 \end{bmatrix} X}$ and give physical meaning to the 3X3 matrix A and the 3-vector $\vec{t}$. Whereas $\vec{t}$ is simple translation, A represents unequal scaling of space along three orthogonal directions plus rotations.
- The only requirement on A is that it be non-singular. Since it takes 9 parameters to specify A and 3 to specify $\vec{t}$, 3D affine has 12 degrees of freedom.

- Following the same line of reasoning as in Lecture 3, we can show that affine transformations map ideal points to ideal points. So, as for world 2D, parallel lines stay parallel. In addition, now parallel planes stay parallel. And, arguing as we did in Lecture 4, we can also prove that a world 3D transformation is affine if and only if π_∞ is mapped to π_∞.

As mentioned on the previous page, 3D affine has 12 DoF. We will argue below that 3D similarity has 7 DoF. That means purely affine distortion in 3D has 5 DoF. With regard to the figure shown, 2 of these are associated with the unequal scaling of ABCD face along two orthogonal directions. Additional 2 DoF are associated with the same for face CDFG. Can you account for the remaining one? (Purely affine distortion in 2D has only 2 DoF)

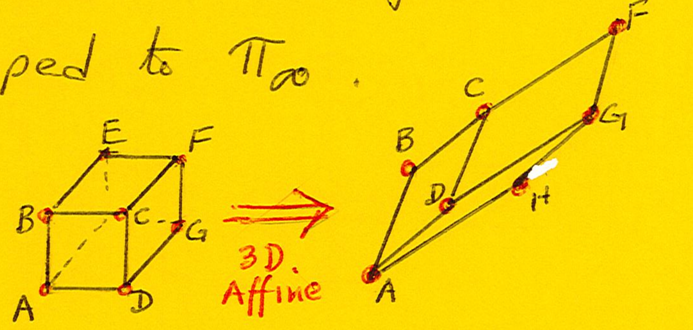

- We obtain the Similarity Subgroup of the Affine Group by imposing the additional constraint that $A^TA = \lambda^2 I$ where I is a 3x3 identity matrix. Using reasoning similar to what you saw in Lecture 4, we can show that a similarity transform preserves angles — in other words, it is shape preserving in 3D. That implies that A can only do two things: rotate a 3D object and scale it isotropically. It takes 3 parameters to specify a rigid-body rotation in 3D (two for defining the direction of the axis of rotation and one for the angle of rotation around the axis). In addition, we need one parameter for the isotropic scaling, and 3 for $\vec{t}$. So, in all, a similarity transformation has 7 DoF.
- We obtain the Euclidean Subgroup by further restricting A to $A^TA = I$. That means, A must now be a pure rigid-body rotation. Reasoning as above, we see that a Euclidean transformation has 6 DoF.

The Absolute Conic and The Absolute Dual Quadric

- The role played by the two circular points for the case of 2D projectivity is played by The Absolute Conic for 3D projectivity. (Lecture 4 coverd Circular Points) The Absolute Conic is the intersection of any sphere with π_∞, just as the Circular Points were the intersection of any circle with l_∞.
- Representing a world 3D point by $X = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}$ in homogeneous coordinates, the general equation of a sphere is given by $\boxed{x_1^2 + x_2^2 + x_3^2 + dx_1x_4 + ex_2x_4 + fx_3x_4 + gx_4^2 = 0}$. Since π_∞ is defined by $x_4 = 0$, the intersection of any sphere with π_∞ is given by $\boxed{x_1^2 + x_2^2 + x_3^2 = 0}$ which can be expressed as $\boxed{(x_1\ x_2\ x_3)\, I_{3\times3} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = 0}$ which looks like the equation of a conic. However, don't forget this conic resides in the plane π_∞. This is our Absolute Conic $\Omega_\infty = I_{3\times3}$. The conic consists only of imaginary points.
- The dual of the Absolute Conic Ω_∞ is called the Absolute Dual Quadric and denoted Ω^*_∞. Algebraically, we have $\boxed{\Omega^*_\infty = \begin{bmatrix} I_{3\times3} & \vec{0} \\ \vec{0}^T & 0 \end{bmatrix}}$. Geometrically, Ω^*_∞ consists of all planes tangent to the point conic Ω_∞. We can show that all planes $\pi = \begin{pmatrix} \vec{v} \\ k \end{pmatrix}$ that satisfy $\vec{v}^T\vec{v} = 0$ are tangent to the point conic Ω_∞ and that at the same time satisfy $\boxed{\pi^T \Omega^*_\infty \pi = 0}$. The second part of this proof is trivial. The first part requires work.