NUS National University of Singapore | School of Computing

CS4277 / CS5477
3D Computer Vision

Lecture 1:
2D and 1D Projective Geometry

Asst. Prof. Lee Gim Hee

AY 2020/21

Semester 2

Acknowledgements

- A lot of slides and content of this lecture are adopted from:

1. R. Hartley, and Andrew Zisserman: “Multiple view geometry in computer vision”, Chapter 2.

2. Y. Ma, S. Soatto, J. Kosecka, S. S. Sastry, “ An invitation to 3-D vision”, Chapter 2.

Projective Transformation

- The mapping of scene objects onto an image is an example of a projective transformation.

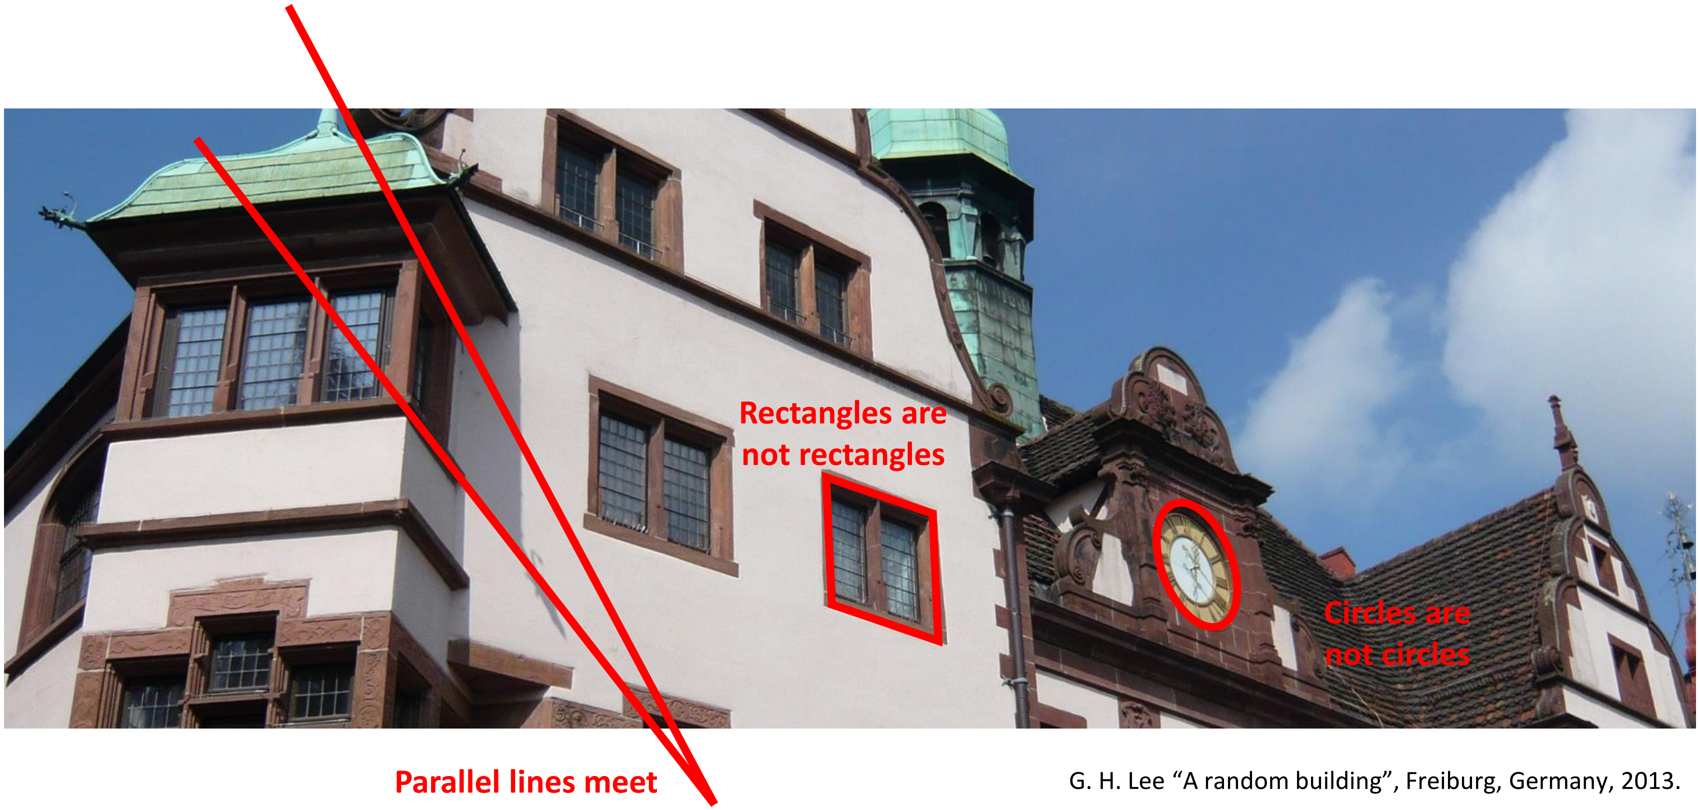

Parallel lines meet at a finite point

G. H. Lee "A random building", Freiburg, Germany, 2013.

What is Projective Geometry?

- We saw that certain geometric properties are not preserved by projective transformation, e.g.

  1. A circle may appear as an ellipse
  2. Parallel lines may meet at a finite point
  3. A rectangle may appear as a parallelogram

- In fact, angles, distance, ratios of distances – none of these are preserved!

What is Projective Geometry?

- A property that is preserved is straightness, which is the most general requirement on the mapping.
- **A thought**: we may define a projective transformation as any mapping that preserves straight lines.
- More generally, we study of geometric properties that are invariant with respect to projective transformations in projective geometry!

Euclidean vs Projective

- The familiar **Euclidean geometry** is an example of synthetic geometry.
- Use axiomatic method and its related tools, i.e. compass and straightedge to solve problems.
- **Projective geometry** uses coordinates and algebra – analytic geometry.
- We will see that one most important result is that geometry at infinity can now be nicely represented!

Raphel, “The School of Athens”, 1509-1511

Image source: https://en.wikipedia.org/wiki/Euclidean_geometry

Homogenous Coordinates

- A point in homogenous coordinates (kx, ky, k), corresponds to $\left(\frac{kx}{k}, \frac{ky}{k}\right) = (x, y)$ in Cartesian coordinates.
- (kx, ky, k) is equivalent for all k's.
- Now we can use (x, y, k), where $k = 0$ to represent the point at infinity, i.e. $\left(\frac{x}{0}, \frac{y}{0}\right)$ which is infinite.
- Generally, the $\mathbb{R}^n$ Euclidean space can be extended to a $\mathbb{P}^n$ projective space as homogeneous vectors.

Homogenous Coordinates

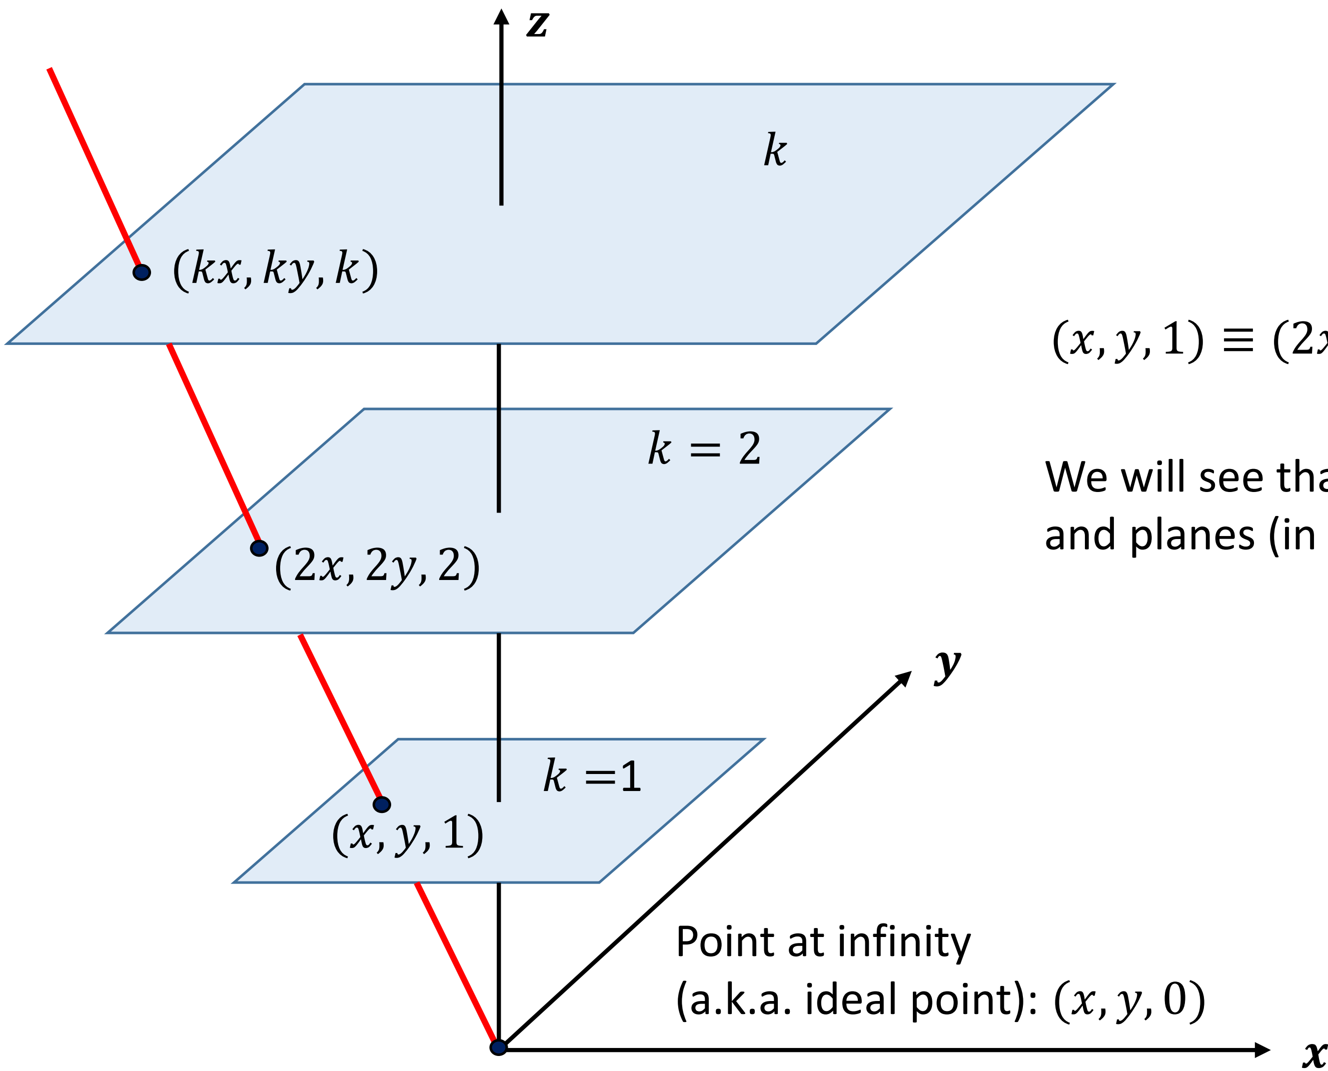

$(x, y, 1) \equiv (2x, 2y, 2) \equiv (kx, ky, k)$

We will see that this idea extends to lines and planes (in both 2D and 3D space)!

Point at infinity
(a.k.a. ideal point): $(x, y, 0)$

The 2D Projective Plane

- We will look at the homogeneous notation for lines $\mathbf{l}$ on a plane π, and the incidence relationship between points and lines.

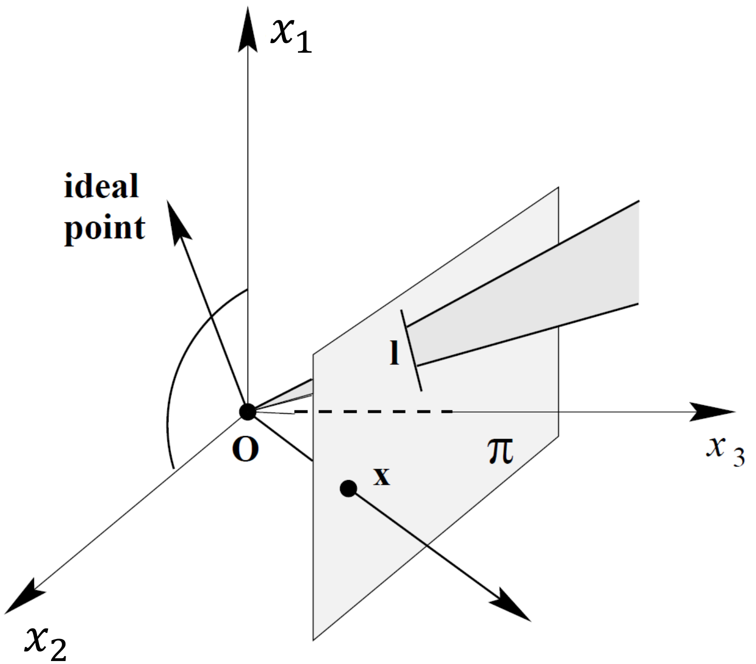

Image source: "Multiple View Geometry in Computer Vision", Richard Hartley and Andrew Zisserman

Homogeneous Representation of Lines and Points

- A line in the plane is represented by:

$$ax + by + c = 0$$

- Different choices of a, b and c giving rise to different lines.

- Thus, a line may naturally be represented by the vector $(a, b, c)^T$.

Homogeneous Representation of Lines and Points

- The correspondence between lines and vectors $(a, b, c)^T$ is not one-to-one.
- Since the lines $ax + by + c = 0$ and $(ka)x + (kb)y + (kc) = 0$ are the same, $\forall k \neq 0$.
- Thus $(a, b, c)^T$ and $k(a, b, c)^T$ represent the same line, for any non-zero k, i.e. equivalence class.
- **Note:** the vector $(0, 0, 0)^T$ does not correspond to any line.

Homogeneous Representation of Lines and Points

- A point $\mathbf{x} = (x, y)^T$ lies on the line $\mathbf{l} = (a, b, c)^T$ if and only if $ax + by + c = 0$, i.e.

$$(x, y, 1)(a, b, c)^T = (x, y, 1)\mathbf{l} = 0;$$

- Similarly, for any constant non-zero k,

$$(kx, ky, k)(a, b, c)^T = k(x, y, 1)\mathbf{l} = (x, y, 1)\mathbf{l} = 0.$$

Homogeneous Representation of Lines and Points

- Hence, $(kx, ky, k)^T \in \mathbb{P}^2$ for varying values of k to be a representation of the point $(x, y)^T \in \mathbb{R}^2$, i.e.

$$\mathbf{x} = (x_1, x_2, x_3)^T \in \mathbb{P}^2 \equiv \left(\frac{x_1}{x_3}, \frac{x_2}{x_3}\right)^T \in \mathbb{R}^2.$$

Homogeneous Representation of Lines and Points

- More formally: The point $\mathbf{x}$ lies on the line $\mathbf{l}$ if and only if $\mathbf{x}^T\mathbf{l} = 0$.
- Note that the expression $\mathbf{x}^T\mathbf{l}$ is just the inner or scalar product of the two vectors $\mathbf{l}$ and $\mathbf{x}$; the scalar product:

$$\mathbf{x}^T\mathbf{l} = \mathbf{l}^T\mathbf{x} = \mathbf{x}.\mathbf{l}$$

- Degree of freedom (dof): a point has 2 dof – x and y coordinates; a line also has 2 dof – two independent ratios $\{a : b : c\}$.

Intersection of Lines

- The intersection of two lines $\mathbf{l}$ and $\mathbf{l}'$ is the point $\mathbf{x} = \mathbf{l} \times \mathbf{l}'$.

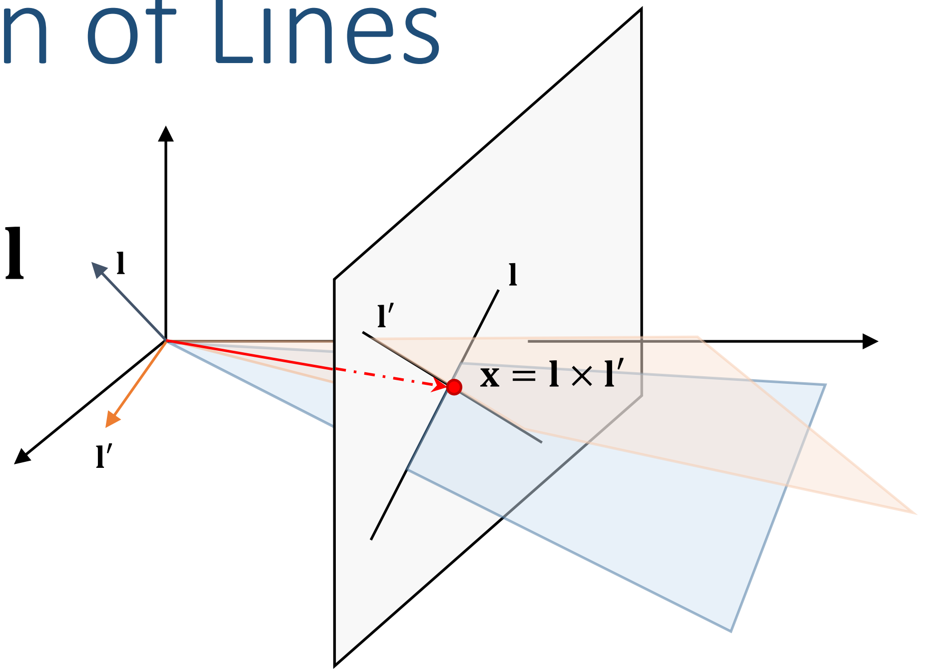

Proof:

Given two lines $\mathbf{l} = (a, b, c)^T$ and $\mathbf{l}' = (a', b', c')^T$, the triple scalar product identity gives $\mathbf{l}.(\mathbf{l} \times \mathbf{l}') = \mathbf{l}'.(\mathbf{l} \times \mathbf{l}') = 0$, which we rewrite as:

$$\mathbf{l}^{\mathsf{T}}\mathbf{x} = \mathbf{l}'^{\mathsf{T}}\mathbf{x} = 0.$$

If $\mathbf{x}$ is thought of as representing a point, then $\mathbf{x}$ lies on both lines $\mathbf{l}$ and $\mathbf{l}'$, and hence is the intersection of the two lines. □

Line Joining Points

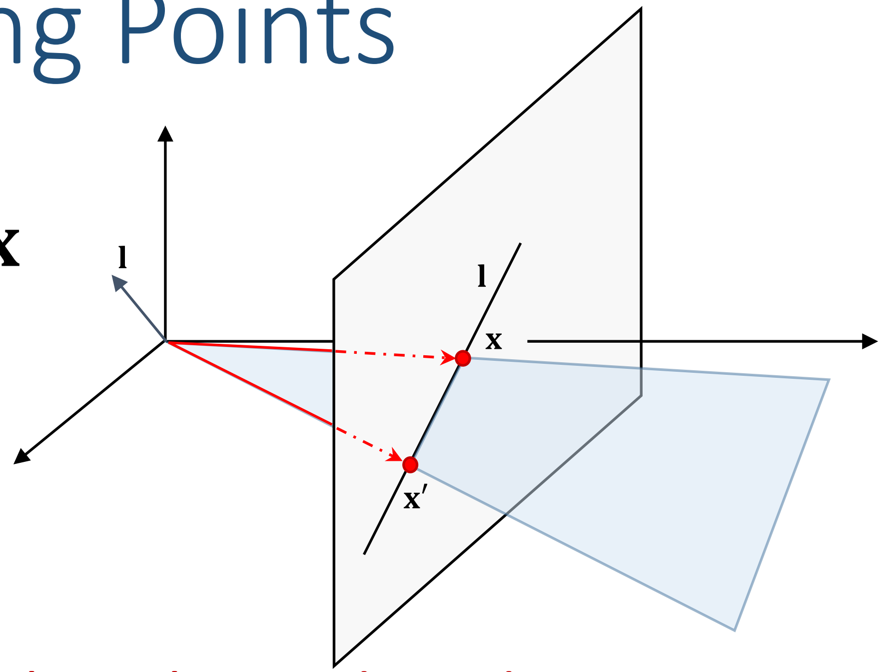

- The line through two points $\mathbf{x}$ and $\mathbf{x}'$ is $\mathbf{l} = \mathbf{x} \times \mathbf{x}'$.

Proof:

Given two points $\mathbf{x}$ and $\mathbf{x}'$, the triple scalar product identity gives $\mathbf{x}.(\mathbf{x} \times \mathbf{x}') = \mathbf{x}'.(\mathbf{x} \times \mathbf{x}') = 0$, which we rewrite as:

$$\mathbf{x}^{\mathrm{T}}\mathbf{l} = \mathbf{x}'^{\mathrm{T}}\mathbf{l} = 0.$$

If $\mathbf{l}$ is thought of as representing a line, then $\mathbf{l}$ contains both points $\mathbf{x}$ and $\mathbf{x}'$, and hence is the line joining the two points.

Ideal Points and the Line at Infinity

Intersection of parallel lines

- Consider two parallel lines $ax + by + c = 0$ and $ax + by + c' = 0$, i.e. $\mathbf{l} = (a, b, c)^T$ and $\mathbf{l}' = (a, b, c')^T$.

- The intersection is $\mathbf{l} \times \mathbf{l}' = (c' - c)(b, -a, 0)^T$, i.e $(b, -a, 0)^T$ ignoring the scale factor $(c' - c)$.

- $(b, -a, 0)^T$ is an infinite point and this implies that parallel lines meet at infinity.

Ideal Points and the Line at Infinity

Example:

Consider the two lines $x = 1$ and $x = 2$. Here the two lines are parallel, and consequently intersect "at infinity".

In homogeneous notation the lines are $\mathbf{l} = (-1, 0, 1)^T$, $\mathbf{l} = (-1, 0, 2)^T$, and their intersection point is:

$$\mathbf{x} = \mathbf{l} \times \mathbf{l}' = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ -1 & 0 & 1 \\ -1 & 0 & 2 \end{vmatrix} = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$$

which is the point at infinity in the direction of the *y*-axis.

Ideal Points and the Line at Infinity

- The points $\mathbf{x} = (x_1, x_2, x_3)^T$ with last coordinate $x_3 = 0$ are known as ideal points, or points at infinity.
- The set of all ideal points may be written $(x_1, x_2, 0)^T$, with a particular point specified by the ratio $x_1 : x_2$.
- Note that this set lies on a single line, the line at infinity, denoted by the vector $\mathbf{l}_\infty = (0, 0, 1)^T$.

Proof:

$$(0, 0, 1)(x_1, x_2, 0)^T = 0.$$

Ideal Points and the Line at Infinity

- The parallel lines $\mathbf{l} = (a, b, c)^T$ and $\mathbf{l}' = (a, b, c')^T$ intersects $\mathbf{l}_\infty$ at the ideal point $(b, -a, 0)^T$ for all c's.
- In inhomogeneous notation $(b, -a)^T$ is a vector tangent to the line, and orthogonal to the line normal (*a, b*), and so represents the line's *direction*.
- As the line's direction varies, the ideal point $(b, -a, 0)^T$ varies over $\mathbf{l}_\infty$.
- Hence, the line at infinity can be thought of as the set of directions of lines in the plane.

Duality principle

- Notice how the role of points and lines may be interchanged in:

1. Incidence equations, i.e. $\mathbf{l}^T\mathbf{x} = 0$ and $\mathbf{x}^T\mathbf{l} = 0$.
2. Intersection of two lines and the line through two points, i.e. $\mathbf{x} = \mathbf{l} \times \mathbf{l}'$ and $\mathbf{l} = \mathbf{x} \times \mathbf{x}'$.

- These observations lead to the duality principle.

Duality principle

- **Duality principle.** To any theorem of 2-dimensional projective geometry there corresponds a dual theorem, which may be derived by interchanging the roles of points and lines in the original theorem.
- Consequently, it not necessary to prove the dual of a given theorem once the original theorem has been proven.
- The proof of the dual theorem will be the dual of the proof of the original theorem.

Conics and Dual Conics

- A conic is a curve described by a second-degree equation in the plane.
- In Euclidean geometry, conics are of three main types: hyperbola, ellipse, and parabola.
- These three types of conic arise as conic sections generated by planes of differing orientation.
- **Note**: there are also degenerate conics, which we will define later.

Conics and Dual Conics

Types of conics:

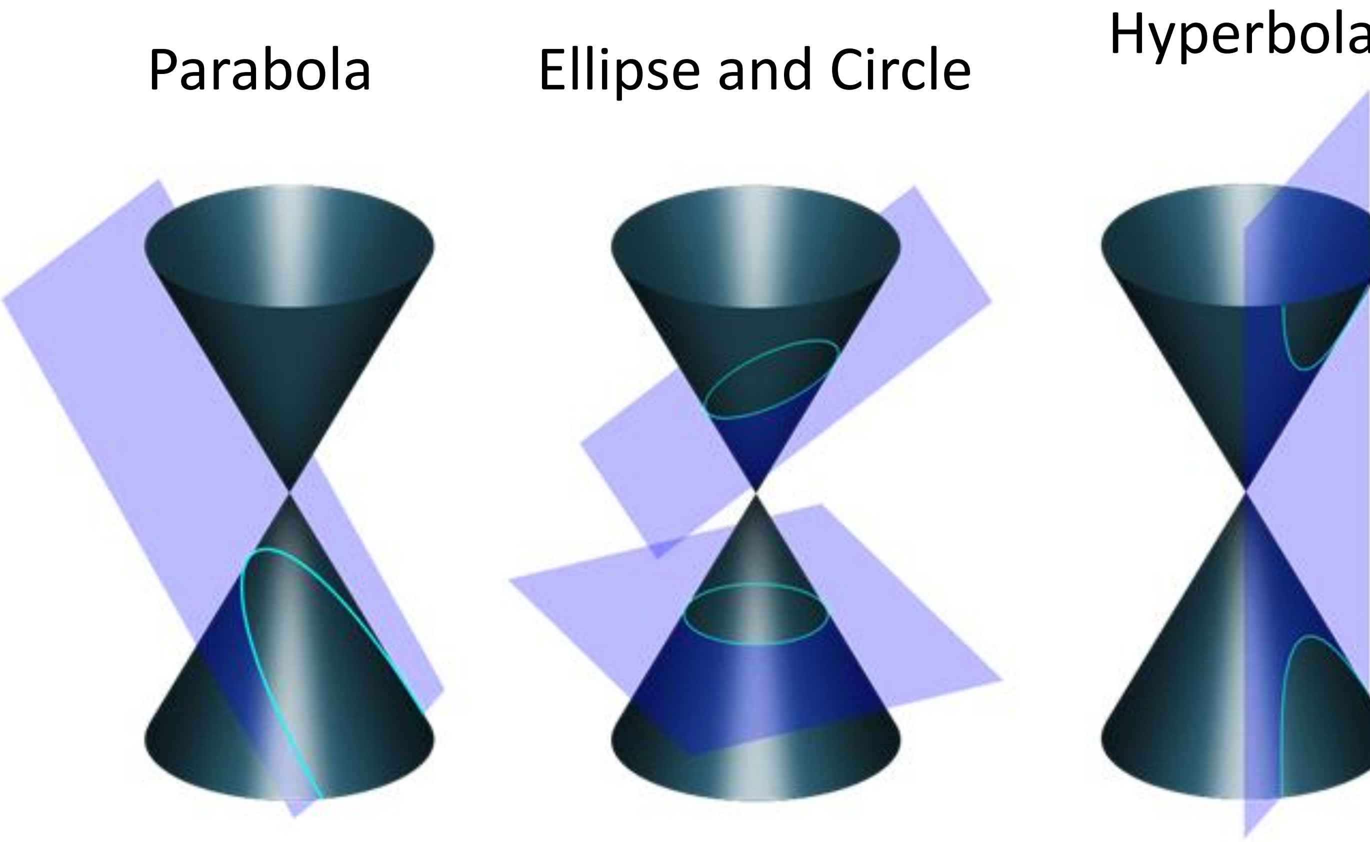

Image source: https://en.wikipedia.org/wiki/Conic_section

Conics and Dual Conics

- The equation of a conic in inhomogeneous coordinates is a polynomial of degree 2, i.e.

$$ax^2 + bxy + cy^2 + dx + ey + f = 0$$

- "Homogenizing" this by the replacements: $x \rightarrow \frac{x_1}{x_3}, y \rightarrow \frac{x_2}{x_3}$ gives

$$a{x_1}^2 + bx_1x_2 + c{x_2}^2 + dx_1x_3 + ex_2x_3 + f{x_3}^2 = 0$$

Conics and Dual Conics

- Or in matrix form: $\mathbf{x}^{\mathsf{T}}\mathtt{C}\mathbf{x} = 0$, where C is symmetric and given by:

$$\mathtt{C} = \begin{bmatrix} a & b/2 & d/2 \\ b/2 & c & e/2 \\ d/2 & e/2 & f \end{bmatrix}.$$

- C is a homogeneous representation of a conic.

- Only the ratios of the matrix elements are important, multiplying C by a non-zero scalar has no effect.

Conics and Dual Conics

$$\mathbf{x}^\mathsf{T}\mathtt{C}\mathbf{x} = 0, \qquad \mathtt{C} = \begin{bmatrix} a & b/2 & d/2 \\ b/2 & c & e/2 \\ d/2 & e/2 & f \end{bmatrix}.$$

- The conic has five degrees of freedom which can be thought of as the ratios $\{a : b : c : d : e : f\}$.
- Or equivalently the six elements of a symmetric matrix less one for scale.

Five Points Define a Conic

- Each point $\mathbf{x}_i = (x_i, y_i)$ places one constraint on the conic coefficients:

$$a{x_i}^2 + bx_iy_i + c{y_i}^2 + dx_i + ey_i + f = 0.$$

- This constraint can be written as:

$$\begin{pmatrix} x_i^2 & x_iy_i & y_i^2 & x_i & y_i & 1 \end{pmatrix} \mathbf{c} = 0$$

- where $\mathbf{c} = (a, b, c, d, e, f)^T$ is the conic C represented as a 6-vector.

Five Points Define a Conic

- Stacking the constraints from five points we obtain

$$\begin{bmatrix} x_1^2 & x_1y_1 & y_1^2 & x_1 & y_1 & 1 \\ x_2^2 & x_2y_2 & y_2^2 & x_2 & y_2 & 1 \\ x_3^2 & x_3y_3 & y_3^2 & x_3 & y_3 & 1 \\ x_4^2 & x_4y_4 & y_4^2 & x_4 & y_4 & 1 \\ x_5^2 & x_5y_5 & y_5^2 & x_5 & y_5 & 1 \end{bmatrix} \mathbf{c} = \mathbf{0}$$

- The conic is the null vector of this 5×6 matrix.
- This shows that a conic is determined uniquely (up to scale) by five points in general position.

Conics and Dual Conics

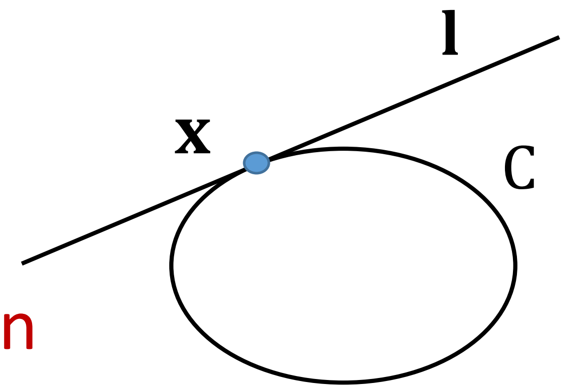

Tangent lines to conics:

The line $\mathbf{l}$ tangent to C at a point $\mathbf{x}$ on C is given by $\mathbf{l} = \mathrm{C}\mathbf{x}$.

Proof:

The line $\mathbf{l} = \mathrm{C}\mathbf{x}$ passes through $\mathbf{x}$, since $\mathbf{l}^T\mathbf{x} = \mathbf{x}^T\mathrm{C}\mathbf{x} = 0$. If $\mathbf{l}$ has one-point contact with the conic, then it is a tangent, and we are done.

$\square$

Conics and Dual Conics

- The conic C defined as far is more properly termed a point conic, as it defines an equation on points.
- There is also a dual (line) conic which defines an equation on lines denoted as C^*(3x3 matrix).
- A line $\mathbf{l}$ *tangent* to the conic C satisfies $\mathbf{l}^T \mathrm{C}^* \mathbf{l} = 0$.
- A dual conic has five degrees of freedom and can be computed from five lines.

Conics and Dual Conics

- For a non-singular symmetric matrix $C^* = C^{-1}$ (up to scale).

Proof:

A point $\mathbf{x}$ on C, the tangent is $\mathbf{l} = C\mathbf{x}$ and this implies $\mathbf{x} = C^{-1}\mathbf{l}$, i.e. $C^* = C^{-1}$ and $\mathbf{x} = C^*\mathbf{l}$.

Furthermore, since $\mathbf{x}$ satisfies $\mathbf{x}^T C\mathbf{x} = 0$, we obtain $(C^{-1}\mathbf{l})^T C(C^{-1}\mathbf{l}) = \mathbf{l}^T C^{-1}\mathbf{l} = 0$, where $C^{-T} = C^{-1}$; we can write as $\mathbf{l}^T C^*\mathbf{l} = 0$.

□

Conics and Dual Conics

- Dual conics are also known as conic envelopes:

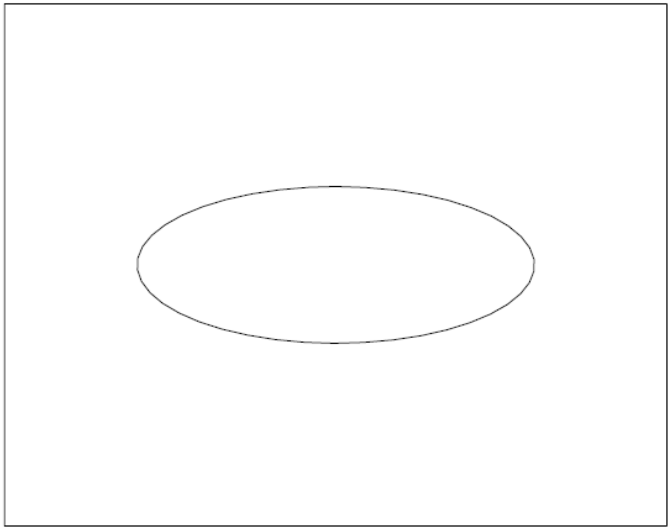

Points $\mathbf{x}$ satisfying $\mathbf{x}^T \mathrm{C} \mathbf{x} = 0$ lie on a point conic.

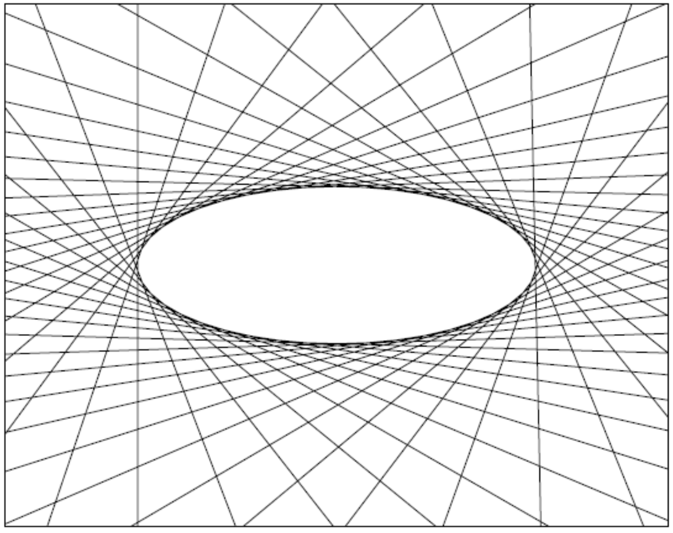

Lines $\mathbf{l}$ satisfying $\mathbf{l}^T \mathrm{C}^* \mathbf{l} = 0$ are tangent to the point conic C. The conic C is the envelope of the lines $\mathbf{l}$.

Image source: "Multiple View Geometry in Computer Vision", Richard Hartley and Andrew Zisserman

Degenerate Conics

- Suppose that $\mathbf{l}$ meets the conic in another point $\mathbf{y}$, then $\mathbf{y}^T \mathrm{C} \mathbf{y} = 0$ and $\mathbf{x}^T \mathrm{C} \mathbf{y} = \mathbf{l}^T \mathbf{y} = 0$.
- From this it follows that $(\mathbf{x} + \alpha \mathbf{y})^T \mathrm{C} (\mathbf{x} + \alpha \mathbf{y}) = 0$ for all α.
- This means that the whole line $\mathbf{l} = \mathrm{C}\mathbf{x}$ joining $\mathbf{x}$ and $\mathbf{y}$ lies on the conic C, which is therefore degenerate.

Degenerate Conics

$$(\mathbf{x} + \alpha\mathbf{y})^T \mathrm{C}(\mathbf{x} + \alpha\mathbf{y}) = 0$$

$$\mathbf{l} = \mathrm{C}\mathbf{x} = \mathrm{C}\mathbf{y} \text{ and } \mathbf{l}^{\mathrm{T}}\mathrm{y} = \mathbf{l}^{\mathrm{T}}\mathbf{x} = \mathbf{0};\ \mathrm{rank}(\mathrm{C}) < 3$$

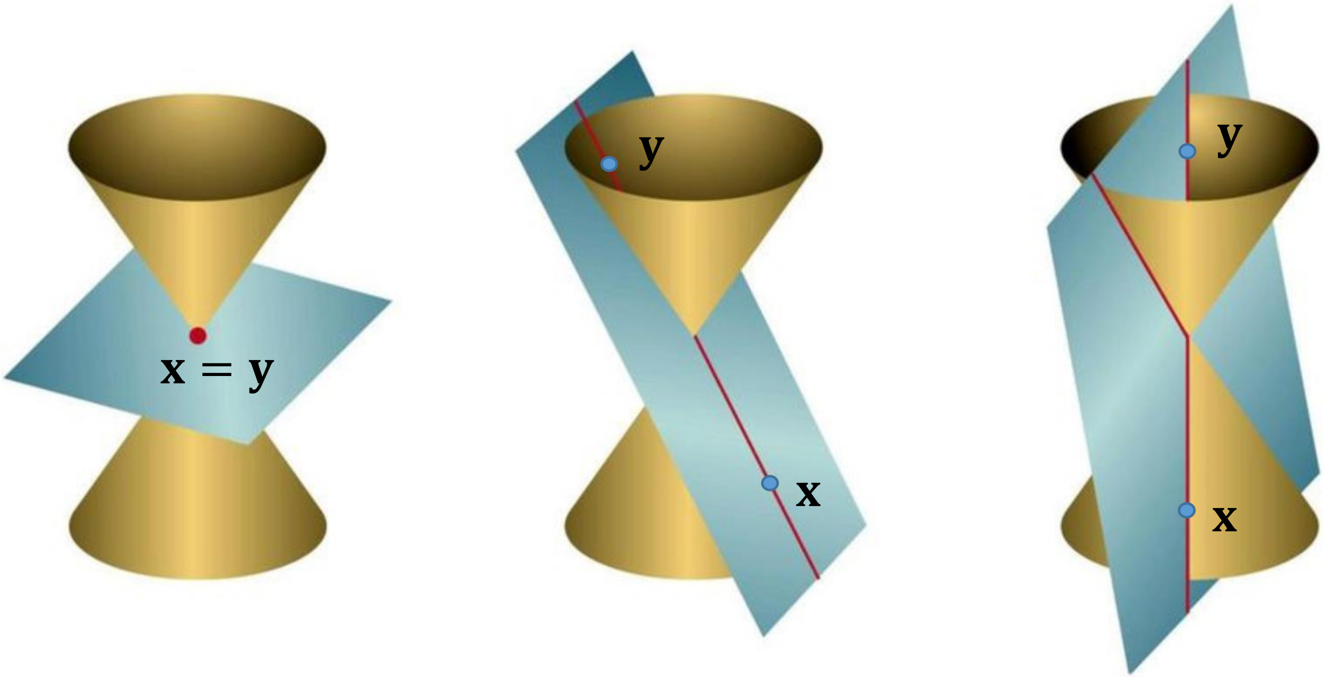

Image source: https://slideplayer.com/slide/12844330/

Planar Projective Transformations

- 2D projective geometry is the study of properties of the projective plane $\mathbb{P}^2$ that are invariant under a group of transformations known as projectivities.

- A projectivity is an invertible mapping h from $\mathbb{P}^2$ to itself such that three points $\mathbf{x}_1$, $\mathbf{x}_2$ and $\mathbf{x}_3$ lie on the same line if and only if $h(\mathbf{x}_1)$, $h(\mathbf{x}_2)$ and $h(\mathbf{x}_3)$ do.

- A projectivity is also called a collineation, a projective transformation or a homography.

Planar Projective Transformations

Theorem:

A mapping $h : \mathbb{P}^2 \rightarrow \mathbb{P}^2$ is a projectivity if and only if there exists a non-singular 3×3 matrix H such that for any point in $\mathbb{P}^2$ represented by a vector $\mathbf{x}$ it is true that $h(\mathbf{x}) = H\mathbf{x}$.

Partial Proof:

Let $\mathbf{x}_1$, $\mathbf{x}_2$ and $\mathbf{x}_3$ lie on a line **l**. Thus $\mathbf{l}^T\mathbf{x}_i = 0$ for $i = 1, \ldots, 3$. Let H be a non-singular 3×3 matrix.

We can verify that $\mathbf{l}^T H^{-1} H\mathbf{x}_i = 0$. Thus, the points $H\mathbf{x}_i$ all lie on the line $H^{-T}\mathbf{l}$, and hence collinearity is preserved by the transformation.

Note: We skip the converse which is harder to prove, i.e. each projectivity arises in this way.

Planar Projective Transformations

- We now define planar projective transformation as:

$$\begin{pmatrix} x'_1 \\ x'_2 \\ x'_3 \end{pmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \quad \text{or} \quad \mathbf{x}' = \mathrm{H}\mathbf{x}.$$

- Properties of H:

1. Non-singular 3×3 matrix;
2. Homogeneous matrix since only the ratio of the matrix elements is significant;
3. Eight degrees of freedom, i.e. eight independent ratios amongst the nine elements of H.

Planar Projective Transformations

- Central projection maps points on one plane to points on another plane.
- And represented by a linear mapping of homogeneous coordinates $\mathbf{x}' = H\mathbf{x}$.

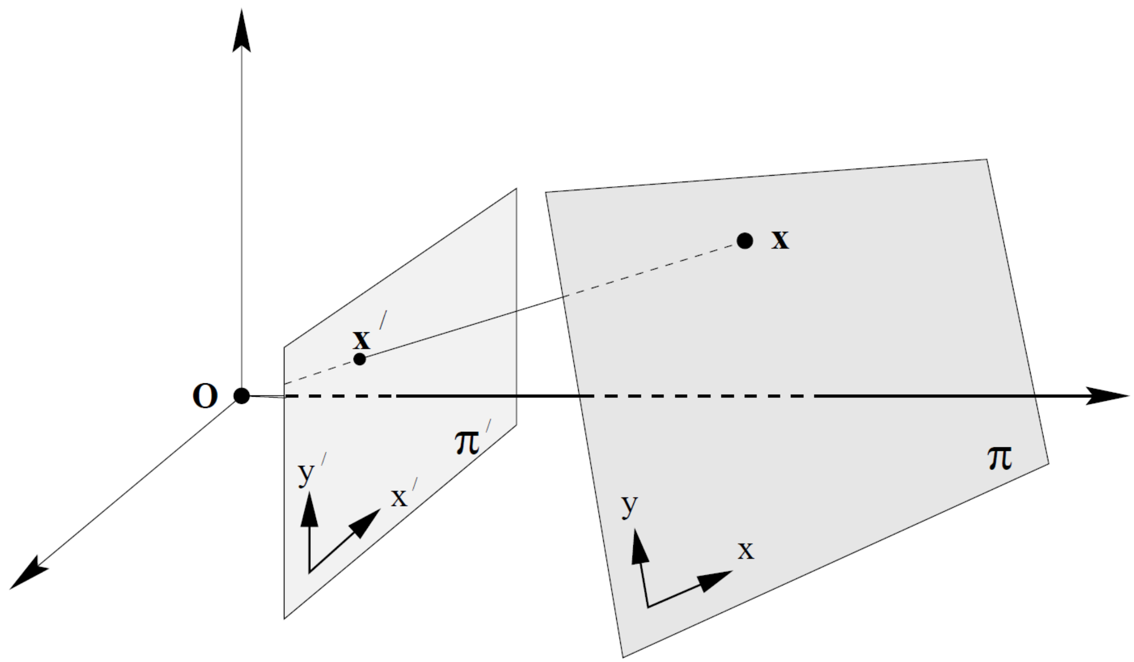

Image source: "Multiple View Geometry in Computer Vision", Richard Hartley and Andrew Zisserman

Planar Projective Transformations

- Examples of a projective transformation $\mathrm{x}' = H\mathrm{x}$, arising in perspective images.

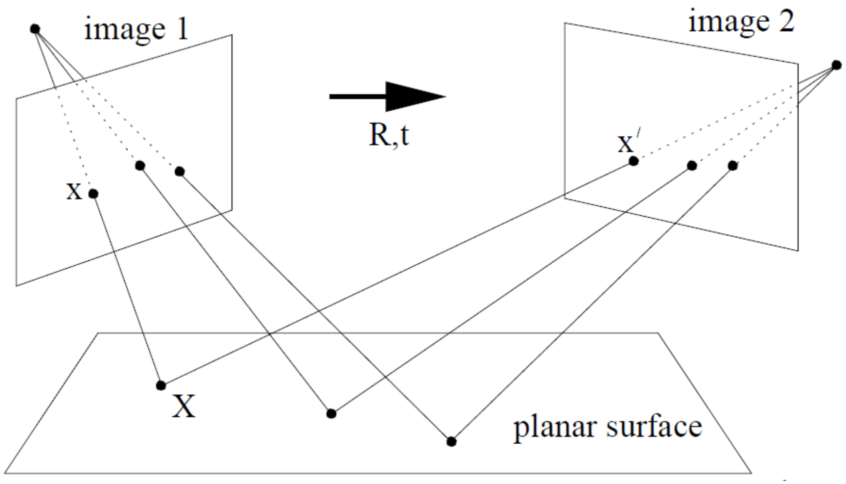

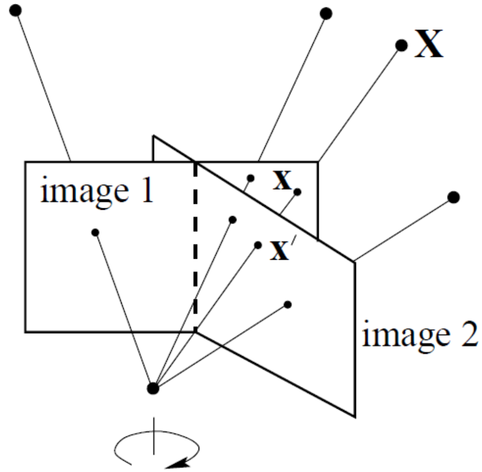

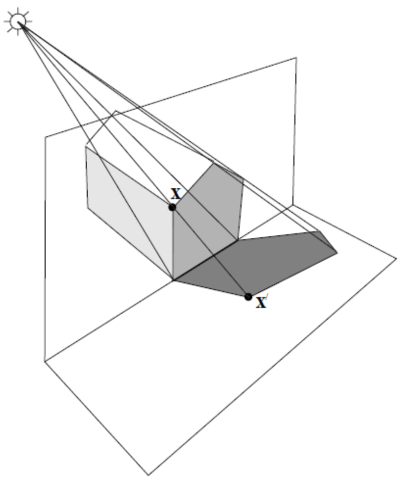

Image source: "Multiple View Geometry in Computer Vision", Richard Hartley and Andrew Zisserman

Transformations of Lines and Conics

- We have seen earlier that:

  If points $\mathbf{x}_i$ lie on a line $\mathbf{l}$, then the transformed points $\mathbf{x}_\mathbf{i}' = H\mathbf{x}_\mathbf{i}$ under a projective transformation lie on the line $\mathbf{l}' = H^{-T}\mathbf{l}$.

  In this way, incidence of points on lines is preserved, since $\mathbf{l}'^T\mathbf{x}_i' = \mathbf{l}^T H^{-1} H\mathbf{x}_i = 0$.

- This means that under the point transformation $\mathbf{x}' = H\mathbf{x}$, a line transforms as:

$$\mathbf{l}' = \mathrm{H}^{-\mathsf{T}}\mathbf{l}\text{, or}\qquad \mathbf{l}'^{\mathsf{T}} = \mathbf{l}^{\mathsf{T}}\mathrm{H}^{-1}.$$

Transformations of Lines and Conics

- Under a point transformation $\mathbf{x}' = H\mathbf{x}$, a conic C transforms to $\mathrm{C}' = H^{-T}CH^{-1}$.
- Under a point transformation $\mathbf{x}' = H\mathbf{x}$, a dual conic C^* transforms to $\mathrm{C}^{*\prime} = H\mathrm{C}^*H^T$.

Proof:

Under a point transformation $\mathbf{x}' = \mathrm{H}\mathbf{x}$,

$$\begin{aligned}\mathbf{x}^\top \mathrm{C}\mathbf{x} &= \mathbf{x}'^\top [\mathrm{H}^{-1}]^\top \mathrm{C}\mathrm{H}^{-1}\mathbf{x}' \\ &= \mathbf{x}'^\top \mathrm{H}^{-\top}\mathrm{C}\mathrm{H}^{-1}\mathbf{x}'\end{aligned}$$

which is a quadratic form $\mathbf{x}'^T\mathrm{C}'\mathbf{x}'$ with $\mathrm{C}' = H^{-T}CH^{-1}$.

Hierarchy of Transformations: Isometries

- Isometries are transformations of the plane $\mathbb{R}^2$ that preserve Euclidean distance, and represented as

$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{bmatrix} \epsilon\cos\theta & -\sin\theta & t_x \\ \epsilon\sin\theta & \cos\theta & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}, \text{ or}$$

$$\mathbf{x}' = \mathrm{H_E}\mathbf{x} = \begin{bmatrix} \mathrm{R} & \mathbf{t} \\ \mathbf{0}^\mathsf{T} & 1 \end{bmatrix} \mathbf{x}$$

- where $\epsilon = \pm 1$.

Hierarchy of Transformations: Isometries

- If ϵ =1, then the isometry is orientation-preserving and is a Euclidean transformation (rotation matrix R and translation t).

- If $\epsilon = -1$, then the isometry reverses orientation, e.g. reflection.

- **Invariants:** Length, angle and area.

Hierarchy of Transformations: Similarity

- Similarity transformation is an isometry composed with an isotropic scaling, and represented as

$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{bmatrix} s\cos\theta & -s\sin\theta & t_x \\ s\sin\theta & s\cos\theta & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}$$

, or

$$\mathbf{x}' = \mathrm{H}_{\mathrm{S}}\mathbf{x} = \begin{bmatrix} s\mathrm{R} & \mathbf{t} \\ \mathbf{0}^{\mathsf{T}} & 1 \end{bmatrix} \mathbf{x}$$

- where the scalar s represents the isotropic scaling.

Hierarchy of Transformations: Similarity

- A similarity transformation is also known as an *equi-form* transformation, because it preserves "shape" (form).

- H_s has four degrees of freedom (3 isometry + 1 scale) and can be computed from two point correspondences.

- **Invariants:** Angles, ratio of two lengths and ratio of areas.

Hierarchy of Transformations: Affinity

- Affine transformation is a non-singular linear transformation followed by a translation, and represented as

$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{bmatrix} a_{11} & a_{12} & t_x \\ a_{21} & a_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \quad \text{, or} \quad \mathbf{x}' = \mathrm{H}_{\mathrm{A}}\mathbf{x} = \begin{bmatrix} \mathrm{A} & \mathbf{t} \\ \mathbf{0}^{\mathsf{T}} & 1 \end{bmatrix} \mathbf{x}$$

- where A is a 2 × 2 non-singular matrix.

- H_A has six degrees of freedom and can be computed from three point correspondences.

- **Invariants:** parallel lines, ratio of lengths of parallel line segments and ratio of areas.

Hierarchy of Transformations: Affinity

- The affine matrix A can always be decomposed as:

$$\mathtt{A} \;=\; \mathrm{R}(\theta)\,\mathrm{R}(-\phi)\,\mathrm{D}\,\mathrm{R}(\phi)$$

- $R(\theta)$ and $R(\phi)$ are rotations by θ and ϕ respectively, and D is a diagonal matrix:

$$\mathtt{D} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}.$$

- This decomposition follows directly from the Singular Value Decomposition (SVD):

$$A = UDV^T = (UV^T)(VDV^T) = R(\theta)\big(R(-\phi)DR(\phi)\big).$$

- Since U and V are orthogonal matrices.

Hierarchy of Transformations: Affinity

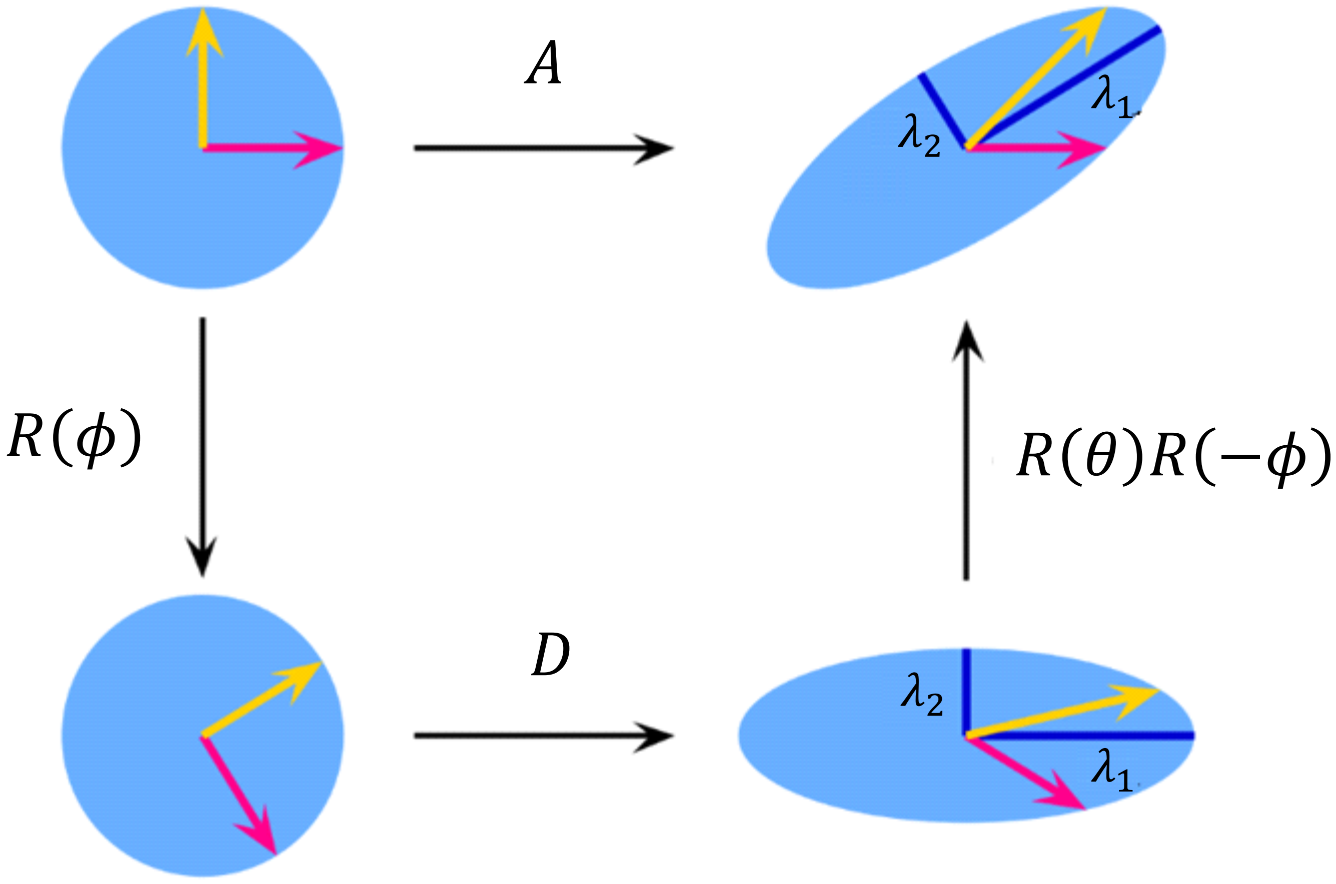

$$A = R(\theta)\big(R(-\phi)DR(\phi)\big)$$

Image modified from: https://en.wikipedia.org/wiki/Singular_value_decomposition

Hierarchy of Transformations: Projective

- Projective transformation is a general non-singular linear transformation of homogeneous coordinates, and represented as

$$\mathbf{x}' = \mathrm{H}_{\mathrm{P}}\mathbf{x} = \begin{bmatrix} \mathrm{A} & \mathbf{t} \\ \mathbf{v}^{\mathsf{T}} & v \end{bmatrix} \mathbf{x}$$

- where the vector $\mathbf{v} = (v_1, v_2)^T$ and v can be 0.

- H_p has nine elements with only their ratio significant, so the transformation has eight degrees of freedom.

Hierarchy of Transformations: Projective

- Note, it is not always possible to scale the matrix such that v is unity since v might be zero.
- A projective transformation between two planes can be computed from four point correspondences, with no three collinear on either plane.
- Not possible to distinguish between orientation preserving and orientation reversing projectivities in $\mathbb{P}^2$.
- **Invariants:** order of contact, tangency (2 pt contact) and cross ratio (details later).

Hierarchy of Transformations

Group	Matrix	Distortion	Invariant properties
Projective 8 dof	$\begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix}$	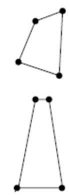	Concurrency, collinearity, **order of contact**: intersection (1 pt contact); tangency (2 pt contact); inflections (3 pt contact with line); tangent discontinuities and cusps. cross ratio (ratio of ratio of lengths).
Affine 6 dof	$\begin{bmatrix} a_{11} & a_{12} & t_x \\ a_{21} & a_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix}$	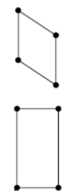	Parallelism, ratio of areas, ratio of lengths on collinear or parallel lines (e.g. midpoints), linear combinations of vectors (e.g. centroids). The line at infinity, $\mathbf{l}_\infty$ (more later).
Similarity 4 dof	$\begin{bmatrix} sr_{11} & sr_{12} & t_x \\ sr_{21} & sr_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix}$	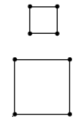	Ratio of lengths, angle. The circular points, **I**, **J** (more later).
Euclidean 3 dof	$\begin{bmatrix} r_{11} & r_{12} & t_x \\ r_{21} & r_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix}$	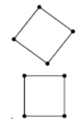	Length, area

Source: Page 44, Table 2.1, "Multiple View Geometry in Computer Vision", Richard Hartley and Andrew Zisserman

Decomposition of a Projective Transformation

- A projective transformation can be decomposed into a chain of transformations:

$$\mathtt{H} = \mathtt{H}_\mathrm{S}\,\mathtt{H}_\mathrm{A}\,\mathtt{H}_\mathrm{P} = \begin{bmatrix} s\mathtt{R} & \mathbf{t} \\ \mathbf{0}^\mathsf{T} & 1 \end{bmatrix} \begin{bmatrix} \mathtt{K} & \mathbf{0} \\ \mathbf{0}^\mathsf{T} & 1 \end{bmatrix} \begin{bmatrix} \mathtt{I} & \mathbf{0} \\ \mathbf{v}^\mathsf{T} & v \end{bmatrix} = \begin{bmatrix} \mathtt{A} & \mathbf{t} \\ \mathbf{v}^\mathsf{T} & v \end{bmatrix}$$

- A a non-singular matrix given by $A = s\mathrm{RK} + \mathbf{tv}^T$.
- K an upper-triangular matrix normalized as $\det(\mathrm{K}) = 1$.
- Decomposition is valid provided $v \neq 0$, and is unique if s is chosen positive.
- We will see that this decomposition preserves geometric properties of $\mathbf{l}_\infty$ and the circular points (Lecture 3).

Projective Geometry of 1D (Line), $\mathbb{P}^1$

- We denote a point on the line as the homogeneous coordinates $\bar{\mathbf{x}}' = (x_1, x_2)^T$.
- $x_2 = 0$ is an ideal point of the line.
- A projective transformation of a line is represented by a 2×2 homogeneous matrix,

$$\bar{\mathbf{x}}' = \mathrm{H}_{2\times2}\bar{\mathbf{x}}$$

- $H_{2\times2}$ has 3 dof corresponding to 4 elements less one for over scaling, and can be computed from 3 points.

Projective Geometry of 1D (Line), $\mathbb{P}^1$

The Cross Ratio

- The cross ratio is the basic projective invariant of $\mathbb{P}^1$. Given 4 points $\bar{\mathbf{x}}_i$ the *cross ratio* is defined as:

$$\text{Cross}(\bar{\mathbf{x}}_1, \bar{\mathbf{x}}_2, \bar{\mathbf{x}}_3, \bar{\mathbf{x}}_4) = \frac{|\bar{\mathbf{x}}_1 \bar{\mathbf{x}}_2||\bar{\mathbf{x}}_3 \bar{\mathbf{x}}_4|}{|\bar{\mathbf{x}}_1 \bar{\mathbf{x}}_3||\bar{\mathbf{x}}_2 \bar{\mathbf{x}}_4|} ,$$

where

$$|\bar{\mathbf{x}}_i \bar{\mathbf{x}}_j| = \det \begin{bmatrix} x_{i1} & x_{j1} \\ x_{i2} & x_{j2} \end{bmatrix} .$$

- If each point $\bar{\mathbf{x}}_i$ is a finite point and $x_2 = 1$, then $|\bar{\mathbf{x}}_i \bar{\mathbf{x}}_j|$ represents the signed distance from $\bar{\mathbf{x}}_i$ to $\bar{\mathbf{x}}_j$.

- Definition of the cross ratio is also valid if one of the points $\bar{\mathbf{x}}_i$ is an ideal point.

Projective Geometry of 1D (Line), $\mathbb{P}^1$

- The value of the cross ratio is invariant under any projective transformation of the line: if $\bar{\mathrm{x}}' = H_{2\times 2}\bar{\mathrm{x}}$ then

$$\mathrm{Cross}(\bar{\mathbf{x}}_1', \bar{\mathbf{x}}_2', \bar{\mathbf{x}}_3', \bar{\mathbf{x}}_4') = \mathrm{Cross}(\bar{\mathbf{x}}_1, \bar{\mathbf{x}}_2, \bar{\mathbf{x}}_3, \bar{\mathbf{x}}_4).$$

Exercise: Prove it!

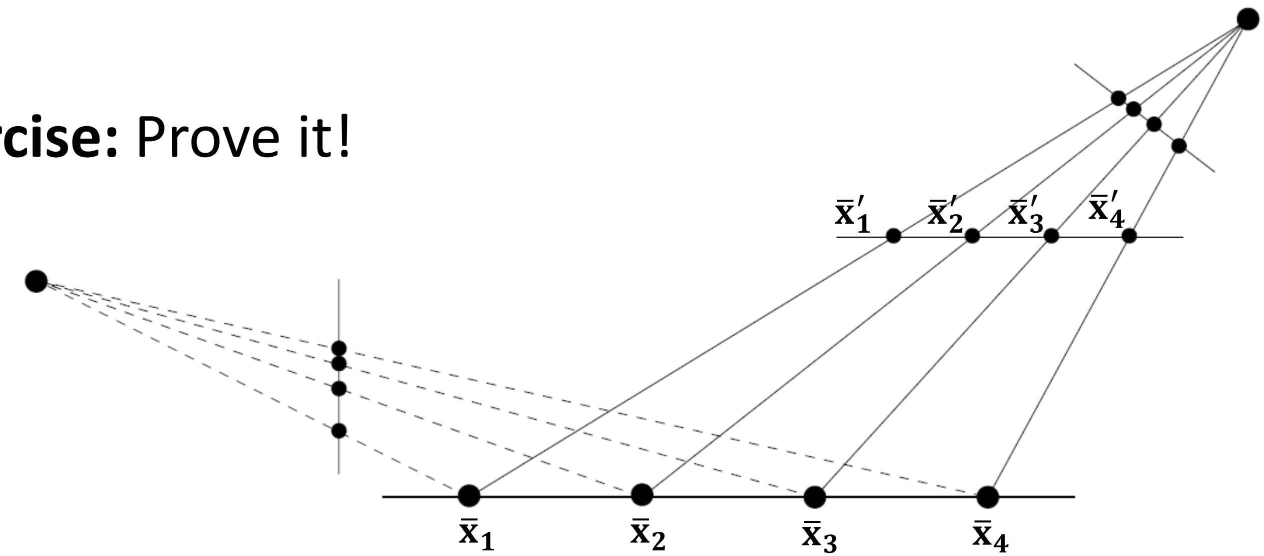

Image source: "Multiple View Geometry in Computer Vision", Richard Hartley and Andrew Zisserman

Projective Geometry of 1D (Line), $\mathbb{P}^1$

Concurrent Lines

- A configuration of concurrent lines is dual to collinear points on a line, i.e. concurrent lines on a plane are also in $\mathbb{P}^1$.
- Four concurrent lines $\mathbf{l}_i$ intersect the line $\mathbf{l}$ in the four points $\bar{\mathbf{x}}_i$.
- The cross ratio of these lines is an invariant to projective transformations of the plane.
- Its value is given by the cross ratio of the points, $\mathrm{Cross}(\bar{\mathbf{x}}_1, \bar{\mathbf{x}}_2, \bar{\mathbf{x}}_3, \bar{\mathbf{x}}_4)$.

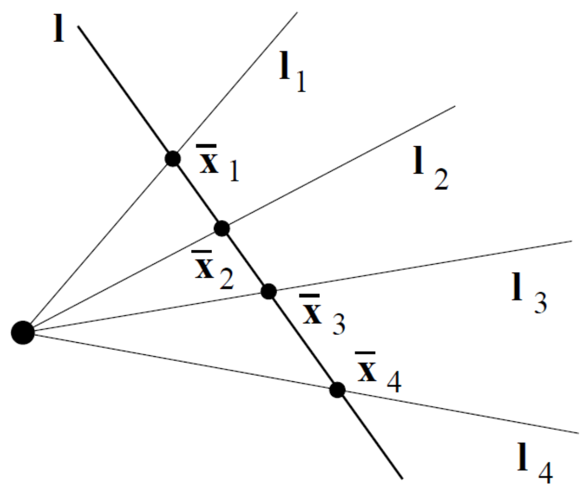

Image source: "Multiple View Geometry in Computer Vision", Richard Hartley and Andrew Zisserman

Projective Geometry of 1D (Line), $\mathbb{P}^1$

Concurrent Lines

- Coplanar points $\mathbf{x}_i$ are imaged onto a line $\mathbf{l}$ (also in the plane) by a projection with centre $\mathbf{C}$.
- May be thought of as representing projection of points in $\mathbb{P}^2$ into a 1-dimensional image.
- In particular, the line $\mathbf{l}$ represents an 1D analogue of the image plane.
- The cross ratio of the image points $\bar{\mathbf{x}}_\mathrm{i}$ is invariant to the position of the image line $\mathbf{l}$.

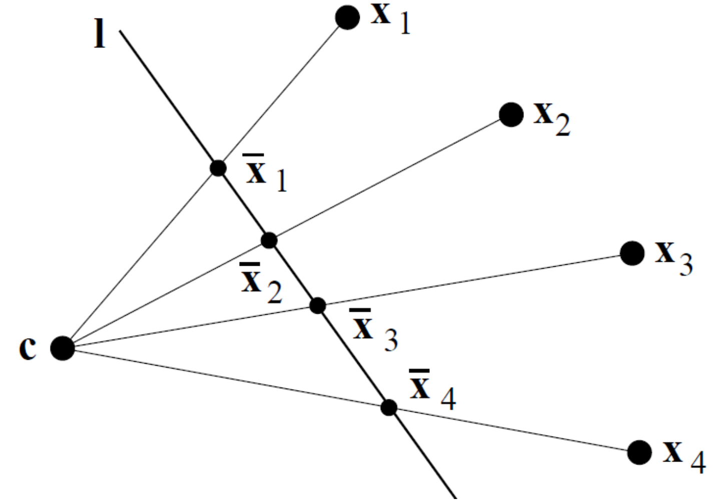

Image source: "Multiple View Geometry in Computer Vision", Richard Hartley and Andrew Zisserman