CSE455: Computer Vision

Geometric Primitives & Transformations

Simran Bagaria

Reference: Szeliski 2.1

What is the most popular topic at CVPR?

	Publication	<u>h5-index</u>	<u>h5-median</u>
1.	Nature	<u>467</u>	707
2.	The New England Journal of Medicine	<u>439</u>	876
3.	Science	<u>424</u>	665
4.	IEEE/CVF Conference on Computer Vision and Pattern Recognition	<u>422</u>	681
5.	The Lancet	<u>368</u>	688
6.	Nature Communications	<u>349</u>	456
7.	Advanced Materials	<u>326</u>	415
8.	Cell	<u>316</u>	503
9.	Neural Information Processing Systems	<u>309</u>	503
10.	International Conference on Learning Representations	<u>303</u>	563

h5-index: largest number h such that h articles published in the last 5 years have at least h citations each.

https://scholar.google.com/citations?view_op=top_venues&hl=en

CVPR 2023 by the Numbers

AUTHORS |

889 185

813 166

688 153

673 139 631 118

553 126

524 113 485 92

450 89

431 91

420 53

373 83

359 69

349 71

344 54

276 54

274 61

155

141 23

129 30

107 24 80 14

72 12

65 14

55 12

PAPERS

Selecting a category below changes the paper list on the right.

SELECT **J** Top 10 overall by number of authors

- 1 3D from multi-view and sensors
- 2 Image and video synthesis and generation
- 3 Humans: Face, body, pose, gesture, movement
- 4 Transfer, meta, low-shot, continual, or long-tail learning
- 5 Recognition: Categorization, detection, retrieval
- 6 Vision, language, and reasoning
- 7 Low-level vision
- 8 Segmentation, grouping and shape analysis
- 9 Deep learning architectures and techniques
- 10 Multi-modal learning
- 11 3D from single images
- 12 Medical and biological vision, cell microscopy
- 13 Video: Action and event understanding
- 14 Autonomous driving
- 15 Self-supervised or unsupervised representation learning
- 16 Datasets and evaluation
- 17 Scene analysis and understanding
- 18 Adversarial attack and defense
- 19 Efficient and scalable vision
- 20 Computational imaging
- 21 Video: Low-level analysis, motion, and tracking
- 22 Vision applications and systems
- 23 Vision + graphics
- 24 Robotics
- 25 Transparency, fairness, accountability, privacy, ethics in vision
- 26 Explainable computer vision
- 27 Embodied vision: Active agents, simulation
- 28 Document analysis and understanding
- 29 Machine learning (other than deep learning)
- 30 Physics-based vision and shape-from-X

Select PICK INSTITUTIONS Award Candidate (AII) • Highlight Paper **# AUTHORS # PAPERS** 1,000 800 600 400 200 3D from multi-view and sensors Award Candidate Highlight Paper NeuMap: Neural Coordinate Mapping by Auto-Transdecoder for Camera 33 Localization **Object Pose Estimation with Statistical Guarantees: Conformal Keypoint** 76 **Detection and Geometric Uncertainty Propagation** NeuralUDF: Learning Unsigned Distance Fields for Multi-view Reconstruction of 120 Surfaces with Arbitrary Topologies NEF: Neural Edge Fields for 3D Parametric Curve Reconstruction from 143 Multi-view Images Looking Through the Glass: Neural Surface Reconstruction Against High 330 Specular Reflections Multi-View Azimuth Stereo via Tangent Space Consistency 357 https://cvpr2023.thecvf.com/Conferences/2023/AcceptedPape

Why do we care about Geometry?

- Self-driving cars: navigation, collision avoidance
- Robots: navigation, manipulation
- Graphics & AR/VR: augment or generate images
- Photogrammetry (architecture, surveys)
- Pattern Recognition (web, medical imaging, etc)

Geometry is more useful now than ever!







Warped frame

Overview of Geometric Vision in CSE455

Geometric Image Formation

The Pinhole Camera model + Calibration

Multi-view Geometry

Structure-from-Motion

Reference textbooks: Szeliski, Hartley & Zisserman to go deeper

Slides credits: Fei-Fei Li, JC Niebles, J. Wu, K. Kitani, S. Lazebnik, S. Seitz, D. Fouhey, J.

What will we learn today?

- Why Geometric Vision Matters
- Geometric Primitives in 2D & 3D
- 2D & 3D Transformations

General Advice / Observations

- Fundamentals: need to (eventually) feel easy
- Try to do the math in parallel live in class!
- If not grokking this: practice later, ask on Ed, OH
- Lots of good (hard?) exercises in Szeliski's book

What will we learn today?

Why Geometric Vision Matters

Geometric Primitives in 2D & 3D

2D & 3D Transformations

Images are 2D projections of the 3D world

Simplified Image Formation



Figure: R. Szeliski

Perspective Projection





Figure: https://www.youtube.com/@huseyin_ozde...

Can we understand the 3D world from 2D images?



CV is an ill-posed inverse problem

2D Image

3D Scene



217	191	252	255	239	
102	80	200	146	138	
159	94	91	121	138	
179	106	136	85	41	
115	129	83	112	67	
94	114	105	111	89	

Objects Material Shape/Geometry Motion Semantics 3D Pose

Slide credit: Andreas Geiger

- 2020-: geometry + learning
- 2010s: deep learning

• . . .

- 2000s: local features, birth of benchmarks
- 1990s: digital camera, 3D reconstruction
- 1980s: epipolar geometry (stereo) [Longuet-Higgins]

• 1860s: first Computer Vision startup? [Willème]







Source: P. Sturm

- 1860s: first Computer Vision startup? [Willème]
- 1850s: birth of photogrammetry [Laussedat]
- 1840s: panoramic photography



Source: P. Sturm

- 1860s: first Computer Vision startup? [Willème]
- 1850s: birth of photogrammetry [Laussedat]
- 1840s: panoramic photography
- 1822-39: birth of photography [Niépce, Daguerre]
- 1773: general 3-point pose estimation [Lagrange]
- 1715: basic intrinsic calibration (pre-photography!) [Taylor]
- 1700's: topographic mapping from perspective drawings [Beautemps-Beaupré, Kappeler]



Niépce, "La Table Servie", 1822

• 15th century: start of mathematical treatment of 3D, <u>first AR app</u>?

Augmented reality invented by Filippo Brunelleschi (1377-1446)?

Tavoletta prospettica di Brunelleschi





• 5th century BC: principles of pinhole camera, a.k.a. camera obscura

- China: 5th century BC
- O Greece: 4th century BC
- O Egypt: 11th century
- Throughout Europe: from 11th century onwards

First mention ...

First camera?



Chinese philosopher Mozi (470 to 390 BC)



Greek philosopher Aristotle (384 to 322 BC)





Source: P. Sturm



What will we learn today?

Why Geometric Vision Matters

Geometric Primitives in 2D & 3D

2D & 3D Transformations

Points in Cartesian and Homogeneous Coordinates

2D points:
$$\mathbf{x} = (x, y) \in \mathcal{R}^2$$
 or column vect $\mathbf{x} = egin{bmatrix} x \ y \end{bmatrix}$

3D points: $\mathbf{x} = (x, y, z) \in \mathcal{R}^3$ (often noted X or P)

Homogeneous coordinates: append a 1

$$\mathbf{\bar{x}} = (x, y, 1)$$
 $\mathbf{\bar{x}} = (x, y, z, 1)$

Why?

Homogeneous coordinates in 2D

2D Projective Space: $\mathcal{P}^2 = \mathcal{R}^3 - (0,0,0)$ (same story in 3D with \mathcal{P}^3)

• heterogeneous \rightarrow homogeneous

$$\left[\begin{array}{c} x\\ y \end{array}\right] \Rightarrow \left[\begin{array}{c} x\\ y\\ 1 \end{array}\right]$$

 $\bullet \ homogeneous \rightarrow heterogeneous$

$$\left[\begin{array}{c} x\\ y\\ w \end{array}\right] \Rightarrow \left[\begin{array}{c} x/w\\ y/w \end{array}\right]$$

• points differing only by scale are *equivalent*: $(x, y, w) \sim \lambda (x, y, w)$

$$\tilde{\mathbf{x}} = (\tilde{x}, \tilde{y}, \tilde{w}) = \tilde{w}(x, y, 1) = \tilde{w}\bar{\mathbf{x}}$$

Homogeneous coordinates in 2D



In homogeneous coordinates, a point and its scaled versions are same

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = w \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} wx \\ wy \\ w \end{bmatrix} \quad w \neq 0$$

Figure: https://www.youtube.com/@huseyin_ozde...

Everything is easier in Projective Space

2D Lines: Representation: l = (a, b, c)Equation: ax + by + c = 0In homogeneous coordinates: $\bar{x}^T l = 0$

General idea: homogenous coordinates unlock the full power of linear algebra!

Everything is easier in Projective Space

2D Lines:

$$\mathbf{\tilde{x}}^{\mathrm{T}}\mathbf{l} = \mathbf{0}, \forall \mathbf{\tilde{x}} = (x, y, w) \in P^2$$

 $\mathbf{l} = (\hat{n}_x, \hat{n}_y, d) = (\mathbf{\hat{n}}, d) \text{ with } \|\mathbf{\hat{n}}\| = 1$



3D planes: same! $\tilde{\mathbf{x}}^{\mathrm{T}}\mathbf{m} = 0, \forall \tilde{\mathbf{x}} = (x, y, z, w) \in P^{3}$

$$\mathbf{m} = (\hat{n}_x, \hat{n}_y, \hat{n}_z, d) = (\mathbf{\hat{n}}, d) \text{ with } \|\mathbf{\hat{n}}\| = 1$$



Lines in 3D

Two-point parametrization:

 $\mathbf{r} = (1 - \lambda)\mathbf{p} + \lambda\mathbf{q}$ $\tilde{\mathbf{r}} = \mu\tilde{\mathbf{p}} + \lambda\tilde{\mathbf{q}}$

Two-plane parametrization:

coordinates $(x_0, y_0) \& (x_1, y_1)$ of intersection with planes at z = 0, 1 (or other planes)



Cross-product quick reminder $\mathbf{a} \times \mathbf{b} = \|\mathbf{a}\| \|\mathbf{b}\| \sin(\theta) \mathbf{n}$ a $\mathbf{a} imes \mathbf{b} = [\mathbf{a}]_{ imes} \mathbf{b} = egin{bmatrix} 0 & -a_3 & a_2 \ a_3 & 0 & -a_1 \ -a_2 & a_1 & 0 \end{bmatrix} egin{bmatrix} b_1 \ b_2 \ b_3 \end{bmatrix}$

Benefits of Homogeneous Coordinates

- Line Point duality:
 - line between two 2D points: $\tilde{\mathbf{l}} = \mathbf{ ilde{x}}_1 imes \mathbf{ ilde{x}}_2$
 - intersection of two 2D lines: $~{\bf \tilde{x}}={\bf \tilde{l}}_1\times{\bf \tilde{l}}_2$
- Representation of Infinity:
 - points at infinity: (x, y, 0); line at infinity: (0,0,1)
- Parallel & vertical lines are easy (take-home: intersect //)
- Makes 2D & 3D transformations linear!

Questions?

What will we learn today?

Why Geometric Vision Matters

Geometric Primitives in 2D & 3D

2D & 3D Transformations

The camera as a coordinate transformation



a 2D image

Source: K. Kitani

The camera as a coordinate transformation



a 2D image

2D to 2D transform (image warping)

Source: K.

Cameras and objects can move!



Figure 2.12 A point is projected into two images: (a) relationship between the 3D point coordinate (X, Y, Z, 1) and the 2D projected point (x, y, 1, d); (b) planar homography induced by points all lying on a common plane $\mathbf{\hat{n}}_0 \cdot \mathbf{p} + c_0 = 0$.
2D Transformations Zoo



Figure: R.

Transformation = Matrix Multiplication



Scaling



Rotation



x

Translation



$$\begin{aligned} x' &= x + t_x \\ y' &= y + t_y \end{aligned}$$

Slide: JC. Niebles

Translation with homogeneous coordinates



$$p' \rightarrow \begin{bmatrix} x + tx \\ y + ty \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} I & t \\ 0 & 1 \end{bmatrix} p = Tp$$

Slide: JC. Niebles

2D Transformations with homogeneous coordinates



Figure:

Questions?

2D Transformations Zoo



Figure: R.

Euclidean / Rigid Transformation

Euclidean (rigid): rotation + translation

$$\begin{bmatrix} \cos\theta & -\sin\theta & t_x \\ \sin\theta & \cos\theta & t_y \\ 0 & 0 & 1 \end{bmatrix}$$

How many degrees of freedom?



Similarity Transformation

Similarity: Scaling + rotation + translation

ion
$$\begin{bmatrix} a & -b & t_x \\ b & a & t_y \\ 0 & 0 & 1 \end{bmatrix}$$

How many degrees of freedom?



Similarity Transformation

Similarity: Scaling + rotation + translation
$$\begin{bmatrix} a & -b & t_x \\ b & a & t_y \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \alpha \cos \theta & -\alpha \sin \theta & b_0 \\ \alpha \sin \theta & \alpha \cos \theta & b_1 \\ 0 & 0 & 1 \end{bmatrix}$$

How many degrees of freedom?



Affine transformations are combinations of

• Arbitrary (4-DOF) linear transformations + translations

$$\begin{bmatrix} x'\\y' \end{bmatrix} = \begin{bmatrix} A_{00} & A_{01}\\A_{10} & A_{11} \end{bmatrix} \begin{bmatrix} x\\y \end{bmatrix} + \begin{bmatrix} b_0\\b_1 \end{bmatrix} \longrightarrow \begin{bmatrix} x'\\y'\\1 \end{bmatrix} = \begin{bmatrix} A_{00} & A_{01} & b_0\\A_{10} & A_{11} & b_1\\0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x\\y\\1 \end{bmatrix}$$

$$\begin{bmatrix} x\\y\\1 \end{bmatrix} \xrightarrow{} Cartesian coordinates} Homogeneous coordinates$$

Source: K. Kitapi

Affine transformations are combinations of

• Arbitrary (4-DOF) linear transformations + translations

coordinates

$$\begin{bmatrix} x'\\ y' \end{bmatrix} = \begin{bmatrix} A_{00} & A_{01}\\ A_{10} & A_{11} \end{bmatrix} \begin{bmatrix} x\\ y \end{bmatrix} + \begin{bmatrix} b_0\\ b_1 \end{bmatrix}$$
$$\begin{bmatrix} x\\ y \end{bmatrix} \begin{bmatrix} x\\ y' \end{bmatrix} \begin{bmatrix} x'\\ y' \end{bmatrix} \begin{bmatrix} x'\\ y' \end{bmatrix} \begin{bmatrix} x'\\ y' \end{bmatrix} \begin{bmatrix} x\\ y \end{bmatrix} \begin{bmatrix} x\\ y' \end{bmatrix} \begin{bmatrix} x\\ y' \end{bmatrix}$$
Cartesian Homogeneous

How many degrees of freedom?

 $\rightarrow \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} A_{00} & A_{01} & b_0 \\ A_{10} & A_{11} & b_1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$

Source: K.







matrix in 3D $\begin{bmatrix} x'\\y'\\1 \end{bmatrix} = \begin{bmatrix} A_{00} & A_{01} & b_0\\A_{10} & A_{11} & b_1\\0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x\\y\\1 \end{bmatrix}$ Then this column is the third basis vector of transformed vector space

This matrix is a linear transformation

And what b_0 and b_1 do is to change the orientation of that basis vector









Affine transformations are combinations of

• Arbitrary (4-DOF) linear transformations + translations

Properties of affine transformations:

- origin does not necessarily map to origin
- lines map to lines
- parallel lines map to parallel lines
- ratios are preserved



Source: K. Kitani

Projective transformations are combinations of

• Affine transformations + projective warps



How many degrees of freedom?



Source: K. Kitani













Original Image

$$\begin{bmatrix} x'\\y'\\z' \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0\\0 & 1 & 0\\h_0 & h_1 & 1 \end{bmatrix} \begin{bmatrix} x\\y\\1 \end{bmatrix}$$

 $h_0 > 0$, $h_1 < 0$ and $|h_0| > |h_1|$



Warped Image

Projective transformations are combinations of

• Affine transformations + projective warps



How many degrees of freedom?

Properties of projective transformations:

- origin does not necessarily map to origin
- lines map to lines
- parallel lines do not necessarily map to parallel lines
- ratios are not necessarily preserved



Source: K. Kitani

Questions?

Composing Transformations

Transformations = Matrices => Composition by Multiplication!

$$p' = R_2 R_1 S p$$

In the example above, the result is equivalent to

$$p' = R_2(R_1(Sp))$$

Equivalent to multiply the matrices into single transformation matrix:

$$p' = (R_2 R_1 S) p$$

Order Matters! Transformations from *right to left*.

Scaling & Translating != Translating & Scaling

$$p'' = TSp = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} s_x & 0 & t_x \\ 0 & s_y & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} s_x x + t_x \\ s_y y + t_y \\ 1 \end{bmatrix}$$

$$p''' = STp = \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} s_x & 0 & s_x t_x \\ 0 & s_y & s_y t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} s_x x + s_x t_x \\ s_y y + s_y t_y \\ 1 \end{bmatrix}$$

Similarity: Translation + Rotation + Scaling

p'= (T R S) p

$$p' = TRSp = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$
$$= \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$
$$= \begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} S & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} RS & t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$
This is the form of the general-purpose transformation matrix

2D Transforms = Matrix Multiplication

Transformation	Matrix	# DoF	Preserves	Icon
translation	$\begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix}_{2 imes 3}$	2	orientation	
rigid (Euclidean)	$\begin{bmatrix} \mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	3	lengths	\bigcirc
similarity	$\begin{bmatrix} s \mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	4	angles	\Diamond
affine	$\begin{bmatrix} \mathbf{A} \end{bmatrix}_{2 imes 3}$	6	parallelism	
projective	$\left[\mathbf{ ilde{H}} ight]_{3 imes 3}$	8	straight lines	

Table 2.1 Hierarchy of 2D coordinate transformations, listing the transformation name, its matrix form, the number of degrees of freedom, what geometric properties it preserves, and a mnemonic icon. Each transformation also preserves the properties listed in the rows below it, i.e., similarity preserves not only angles but also parallelism and straight lines. The 2×3 matrices are extended with a third $[0^T \ 1]$ row to form a full 3×3 matrix for homogeneous coordinate transformations.

Figure: R.

3D Transforms = Matrix Multiplication

Transformation	Matrix	# DoF	Preserves	Icon
translation	$\begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix}_{3 imes 4}$	3	orientation	
rigid (Euclidean)	$\begin{bmatrix} \mathbf{R} & \mathbf{t} \end{bmatrix}_{3 \times 4}$	6	lengths	\bigcirc
similarity	$\begin{bmatrix} s \mathbf{R} & \mathbf{t} \end{bmatrix}_{3 \times 4}$	7	angles	\Diamond
affine	$\left[\mathbf{A} ight]_{3 imes 4}$	12	parallelism	
projective	$\left[\mathbf{ ilde{H}} ight]_{4 imes 4}$	15	straight lines	

Table 2.2 Hierarchy of 3D coordinate transformations. Each transformation also preserves the properties listed in the rows below it, i.e., similarity preserves not only angles but also parallelism and straight lines. The 3×4 matrices are extended with a fourth $[\mathbf{0}^T \ 1]$ row to form a full 4×4 matrix for homogeneous coordinate transformations. The mnemonic icons are drawn in 2D but are meant to suggest transformations occurring in a full 3D cube.

Figure: R.

What did we learn today?

• Geometry is essential to Computer Vision!
What did we learn today?

- Geometry is essential to Computer Vision!
- Geometric Primitives in 2D & 3D
 - Homogeneous coordinates, points, lines, and planes in 2D & 3D

What did we learn today?

- Geometry is essential to Computer Vision!
- Geometric Primitives in 2D & 3D
 - Homogeneous coordinates, points, lines, and planes in 2D & 3D
- 2D & 3D Transformations
 - scaling, translation, rotation, rigid, similarity, affine, homography

Questions?

Appendix

3D Rotations: SO(3) representations



Intersecting Parallel Lines





2D planar transformations



Polar coordinates... $x = r \cos (\phi)$ $y = r \sin (\phi)$ $x' = r \cos (\phi + \theta)$ $y' = r \sin (\phi + \theta)$

Trigonometric Identity... $x' = r \cos(\phi) \cos(\theta) - r \sin(\phi)$ $\sin(\theta)$ $y' = r \sin(\phi) \cos(\theta) + r \cos(\phi)$ $\sin(\theta)$

Substitute... $x' = x \cos(\theta) - y \sin(\theta)$ $y' = x \sin(\theta) + y \cos(\theta) x$