

# **Introduction to Graphics**

# Computer Science Tripos Part 1A Michaelmas Term 2025/2026

Department of Computer Science and Technology The Computer Laboratory

> William Gates Building 15 JJ Thomson Avenue Cambridge CB3 0FD

> > www.cst.cam.ac.uk

This handout includes copies of the slides that will be used in lectures. These notes do not constitute a complete transcript of all the lectures, and they are not a substitute for textbooks. They are intended to give a reasonable synopsis of the subjects discussed, but they give neither complete descriptions nor all the background material.

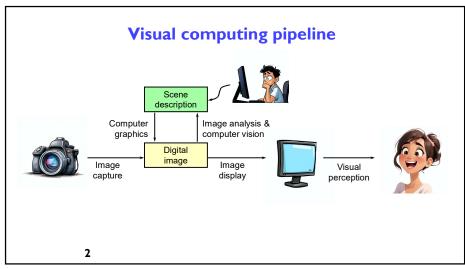
Selected slides contain a reference to the relevant section in the recommended textbook for this course: *Fundamentals of Computer Graphics* by Marschner & Shirley, CRC Press 2015 (4<sup>th</sup> or 5<sup>th</sup> edition). The references are in the format [FCG A.B/C.D], where A.B is the section number in the 4<sup>th</sup> edition and C.D is the section number in the 5<sup>th</sup> edition.

Material is copyright © Neil A Dodgson, Peter Robinson & Rafał Mantiuk, 1996-2025, except where otherwise noted.

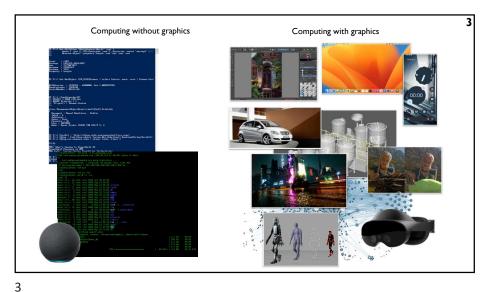
All other copyright material is made available under the University's licence. All rights reserved.

4

# **Introduction to Computer Graphics** Rafał Mantiuk www.cl.cam.ac.uk/~rkm38 Eight lectures & two practical tasks Part IA CST Two supervisions suggested Two exam questions on Paper 3



1 2



Why bother with CG? → All visual computer output depends on CG printed output (laser/ink jet/phototypesetter) monitor (CRT/LCD/OLED/DMD) • all visual computer output consists of real images generated by the computer from some internal digital image → Much other visual imagery depends on CG computer games TV & movie special effects & post-production most books, magazines, catalogues... VR/AR

6

8

## **Course Structure**

### **→ Background**

What is an image? Resolution and quantisation. Storage of images in memory. [1 lecture]

### Rendering

Perspective. Reflection of light from surfaces and shading. Geometric models. Ray tracing.
 [2 lectures]

## + Graphics pipeline

 Polygonal mesh models. Transformations using matrices in 2D and 3D. Homogeneous coordinates. Projection: orthographic and perspective. Rasterisation. [2 lectures]

### → Graphics hardware and OpenGL

GPU APIs. Vertex processing. Fragment processing. Working with meshes and textures.
 [I lecture]

### Human vision, colour and tone mapping

Colour perception. Colour spaces. Tone mapping [2 lectures]

# **Course books**

## → Fundamentals of Computer Graphics

- Shirley & Marschner
   CRC Press 2015 (4<sup>th</sup> or 5<sup>th</sup> edition)
- ◆ [FCG 8.1/9.1] reference to section 3.1 in the 4<sup>th</sup> edition, 9.1 in the 5<sup>th</sup> edition
- **★**Computer Graphics: Principles & Practice
  - Hughes, van Dam, McGuire, Sklar et al. Addison-Wesley 2013 (3<sup>rd</sup> edition)
- → OpenGL Programming Guide: The Official Guide to Learning OpenGL Version 4.5 with SPIR-V
  - Kessenich, Sellers & Shreiner
     Addison Wesley 2016 (7<sup>th</sup> edition and later)











5

# **Introduction to Computer Graphics**

# **→ Background**

- What is an image?
- Resolution and quantisation
- Storage of images in memory
- + Rendering
- **+** Graphics pipeline
- **→** Rasterisation

7

- + Graphics hardware and OpenGL
- → Human vision and colour & tone mapping

# What is a (digital) image?

- → A digital photograph? ("JPEG")
- → A snapshot of real-world lighting?

From computing perspective (discrete)

Image From mathematical perspective (continuous)

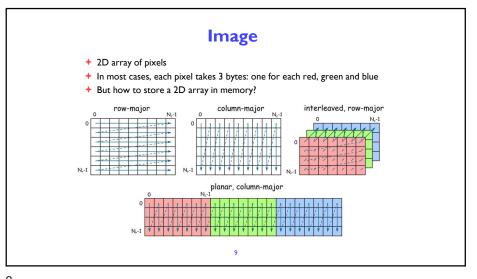
2D array of pixels

To express images in

- •To represent images in memory
- •To create image processing software

•To express image processing as a mathematical problem

•To develop (and understand) algorithms



**S**tride

→ Calculating the pixel component index in memory

For row-major order (grayscale)

$$i(x,y) = x + y \cdot n_{cols}$$

• For column-major order (grayscale)

$$i(x,y) = x \cdot n_{rows} + y$$

For interleaved row-major (colour)

$$i(x, y, c) = x \cdot 3 + y \cdot 3 \cdot n_{cols} + c$$

General case

$$i(x, y, c) = x \cdot s_x + y \cdot s_y + c \cdot s_c$$

where  $s_x$ ,  $s_v$  and  $s_c$  are the strides for the x, y and colour dimensions

10

9

10

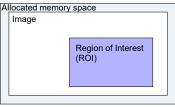
# Padded images and stride

- Sometimes it is desirable to "pad" image with extra pixels
   for example when using operators that need to access pixels outside the image border
- → Or to define a region of interest (ROI)

Allocated memory space
Image
Region of Interest
(ROI)

+ How to address pixels for such an image and the ROI?

Padded images and stride



 $i(x, y, c) = i_{first} + x \cdot s_x + y \cdot s_y + c \cdot s_c$ 

→ For row-major, interleaved, colour

- $i_{first} =$
- $\bullet$   $s_x =$
- $s_{\nu} =$
- $s_c =$

12

11

# **Pixel (PIcture ELement)**

- + Each pixel (usually) consist of three values describing the colour (red, green, blue)
- → For example

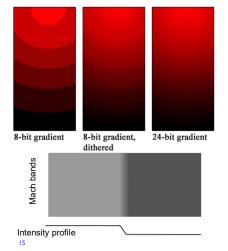
13

15

- (255, 255, 255) for white
- (0, 0, 0) for black
- (255, 0, 0) for red
- → Why are the values in the 0-255 range?
- → How many bytes are needed to store 5MPixel image? (uncompressed)

# **Colour banding**

- + If there are not enough bits to represent colour
- → Looks worse because of the Mach band or **Chevreul** illusion
- + Dithering (added noise) can reduce banding
  - Printers but also some LCD displays



Sample Length: 2 Channel Membership: None 19 18 17 16 15 14 13 12 11 1 R. G. B. A. X RGBAX 10.10.10.0.2 Sample Length Notation → But why? 14 What is a (computer) image? → A digital photograph? ("JPEG") → A snapshot of real-world lighting? From computing From mathematical Image perspective (discrete) (continuous)

16

2D function

•To express image processing as a mathematical problem

•To develop (and understand)

algorithms

Pixel formats, bits per pixel, bit-depth

5.6.5.0.0

→ Grayscale – single colour channel, 8 bits (1 byte)

→ Highcolor – 2<sup>16</sup>=65,536 colors (2 bytes)

RGBAX

→ Truecolor –  $2^{24}$  = 16,8 million colors (3 bytes)

→ Deepcolor – even more colors (>= 4 bytes)

2D array of pixels

•To represent images in

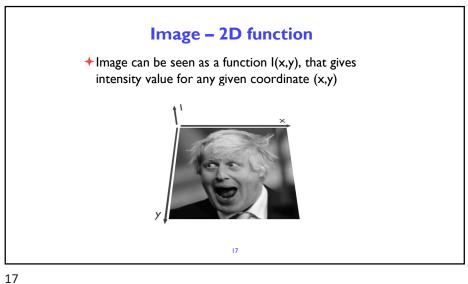
To create image processing

Sample Length Notation:

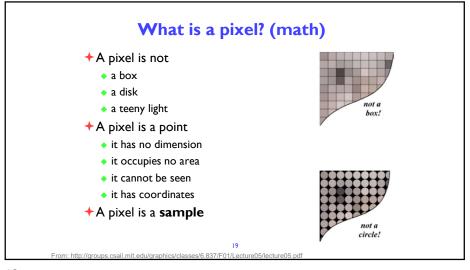
16

©1996–2025 Neil A. Dodgson, Peter Robinson & Rafał Mantiuk

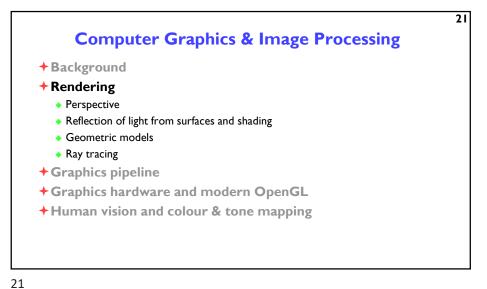
18

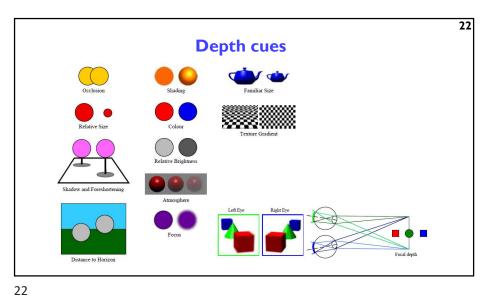


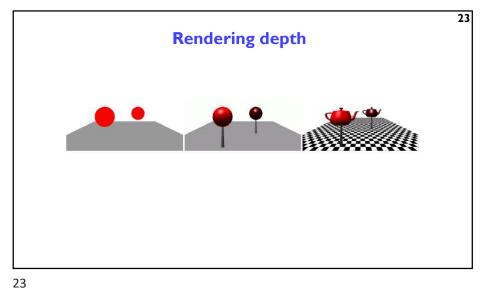
Sampling an image → The image can be sampled on a rectangular sampling grid to yield a set of samples. These samples are pixels.

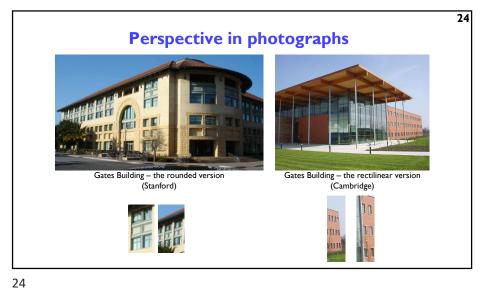


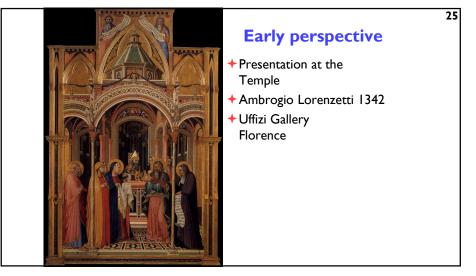
**Sampling and quantization** → Physical world is described in terms of continuous quantities → But computers work only with discrete numbers + Sampling - process of mapping continuous function to a → Quantization – process of mapping continuous variable to a discrete one

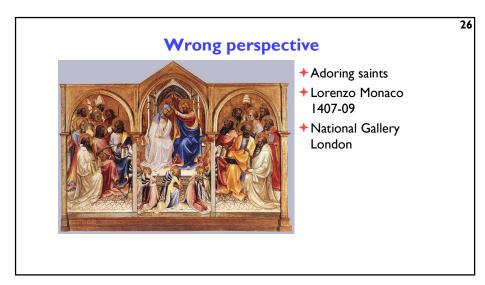




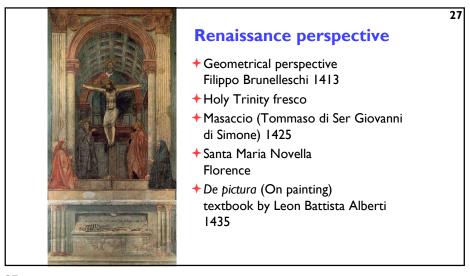


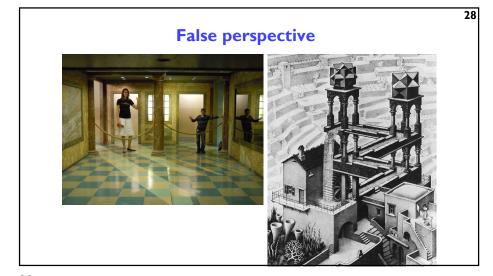


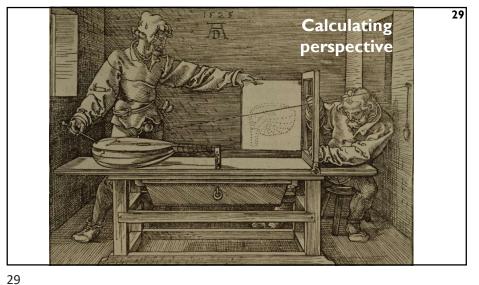


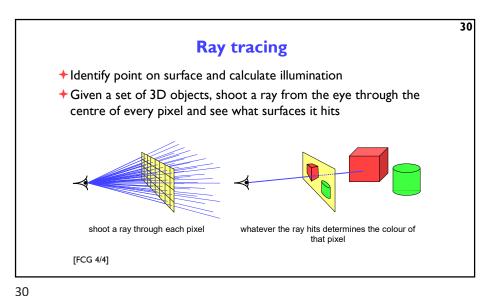


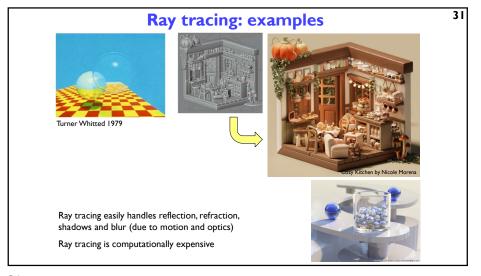
25 26



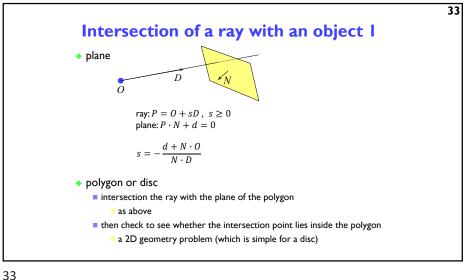


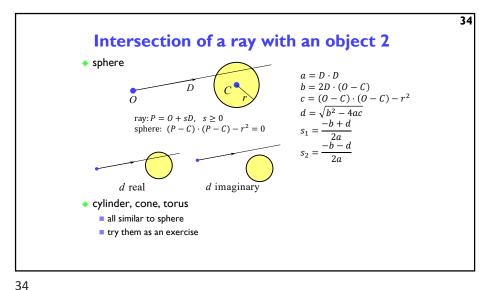


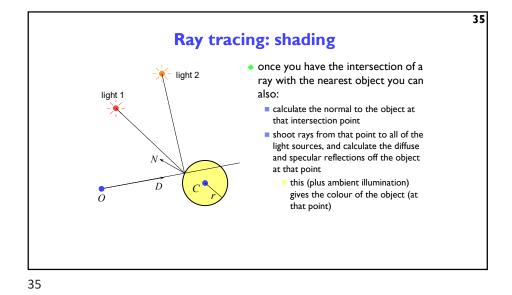


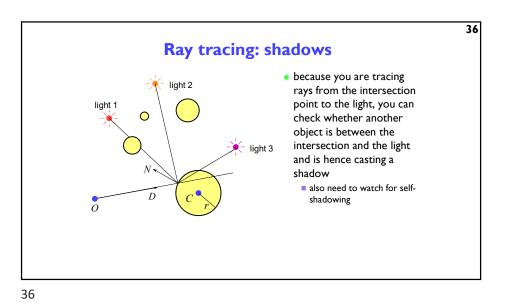


32 Ray tracing algorithm select an eye point and a screen plane FOR every pixel in the screen plane determine the ray from the eye through the pixel's centre FOR each object in the scene IF the object is intersected by the ray IF the intersection is the closest (so far) to the eye record intersection point and object END IF; END FOR; calculate colour for the closest intersection point (if any) END FOR;

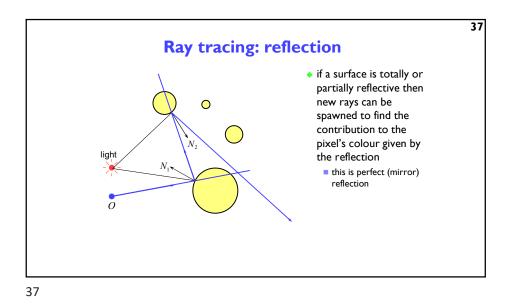


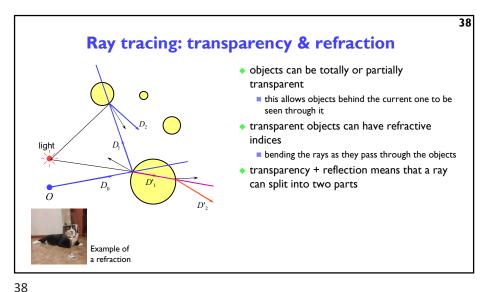






©1996–2025 Neil A. Dodgson, Peter Robinson & Rafał Mantiuk





Illumination and shading

Dürer's method allows us to calculate what part of the scene is visible in any pixel

But what colour should it be?

Depends on:

lighting

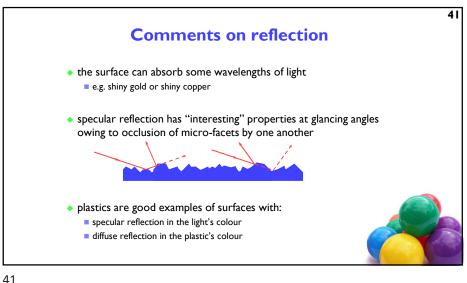
shadows

properties of surface material

How do surfaces reflect light?

perfect specular reflection (Lambertian reflection)

the surface of a specular reflector is facetted, each facet reflects perfectly but in a slightly different direction to the other facets



Calculating the shading of a surface

- gross assumptions:
  - there is only diffuse (Lambertian) reflection
  - all light falling on a surface comes directly from a light source
    - there is no interaction between objects
  - no object casts shadows on any other
    - so can treat each surface as if it were the only object in the scene
  - light sources are considered to be infinitely distant from the object
    - the vector to the light is the same across the whole surface
- observation:

42

• the colour of a flat surface will be uniform across it, dependent only on the colour & position of the object and the colour & position of the light sources

43 **Diffuse shading calculation** L is a normalised vector pointing in the direction of the light source N is the normal to the surface  $I_{i}$  is the intensity of the light source  $k_d$  is the proportion of light which is  $I = I_1 k_d \cos \theta$ diffusely reflected by the surface  $= I_1 k_d (N \cdot L)$ I is the intensity of the light reflected by the surface use this equation to calculate the colour of a pixel

**Diffuse shading: comments** 

- $\bullet$  can have different  $I_i$  and different  $k_d$  for different wavelengths (colours)
- watch out for  $\cos \theta < 0$ 
  - implies that the light is behind the polygon and so it cannot illuminate this side of
- do you use one-sided or two-sided surfaces?
  - one sided: only the side in the direction of the normal vector can be illuminated
    - if  $\cos \theta < 0$  then both sides are black
  - $\blacksquare$  two sided: the sign of  $\cos\theta$  determines which side of the polygon is illuminated
    - need to invert the sign of the intensity for the back side
- this is essentially a simple one-parameter ( $\theta$ ) BRDF
  - Bidirectional Reflectance Distribution Function

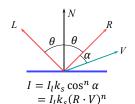
43 44

45

48

# Imperfect specular reflection

+ Phong developed an easy-to-calculate approximation to imperfect specular reflection



L is a normalised vector pointing in the direction of the light source

*R* is the vector of perfect reflection

N is the normal to the surface

V is a normalised vector pointing at the viewer

 $I_{l}$  is the intensity of the light source

 $k_s$  is the proportion of light which is specularly reflected by the surface

n is Phong's ad hoc "roughness" coefficient

I is the intensity of the specularly reflected light











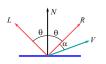
Phong Bui-Tuong, "Illumination for computer generated

pictures", CACM, 18(6), 1975, 311-7

# **Shading: overall equation**

• The overall shading equation can thus be considered to be the ambient illumination plus the diffuse and specular reflections from each light source

$$I = I_a k_a + \sum_i I_i k_d (L \cdot N) + \sum_i I_i k_s (R \cdot V)^n$$



- The equation above is computed for each colour channel (red, green and blue)
- The more lights there are in the scene, the longer this calculation will take

**Examples** specular reflection diffuse reflection

The gross assumptions revisited

- diffuse reflection
- approximate specular reflection
- no shadows
  - need to do ray tracing or shadow mapping to get shadows
- lights at infinity
  - can add local lights at the expense of more calculation
    - need to interpolate the L vector
- no interaction between surfaces
  - - assume that all light reflected off all other surfaces onto a given surface can be amalgamated into a single constant term: "ambient illumination", add this onto the diffuse and specular illumination

47

48

46

# Sampling • we have assumed so far that each ray passes through the centre of a pixel ■ i.e. the value for each pixel is the colour of the object which happens to lie exactly under the centre of the pixel • this leads to: ■ stair step (jagged) edges to objects ■ small objects being missed completely ■ thin objects being missed completely or split into small pieces

Anti-aliasing

These artefacts (and others) are jointly known as aliasing

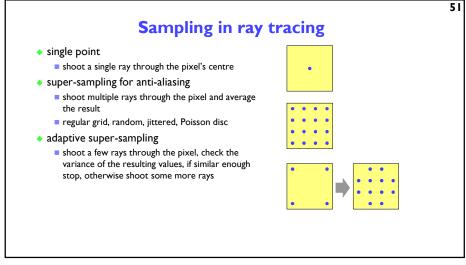
Methods of ameliorating the effects of aliasing are known as anti-aliasing

in signal processing aliasing is a precisely defined technical term for a particular kind of artefact

in computer graphics its meaning has expanded to include most undesirable effects that can occur in the image

this is because the same anti-aliasing techniques which ameliorate true aliasing artefacts also ameliorate most of the other artefacts

49 50



Types of super-sampling

• regular grid

• divide the pixel into a number of sub-pixels and shoot a ray through the centre of each

• problem: can still lead to noticeable aliasing unless a very high resolution sub-pixel grid is used

• random

• shoot N rays at random points in the pixel

• replaces aliasing artefacts with noise artefacts

• the eye is far less sensitive to noise than to aliasing

# Types of super-sampling 2 • Poisson disc • shoot N rays at random points in the pixel with the proviso that no two rays shall pass through the pixel closer than \$\epsilon\$ to one another • for N rays this produces a better looking image than pure random sampling • very hard to implement properly

Types of super-sampling 3

• Jittered (a.k.a. stratified sampling)

• divide pixel into N sub-pixels and shoot one ray at a random point in each sub-pixel

• an approximation to Poisson disc sampling

• for N rays it is better than pure random sampling

• easy to implement

• Poisson disc pure random

53

# More reasons for wanting to take multiple samples per pixel

- super-sampling is only one reason why we might want to take multiple samples per pixel
- many effects can be achieved by distributing the multiple samples over some range
  - called distributed ray tracing
    - N.B. distributed means distributed over a range of values
- can work in two ways
  - each of the multiple rays shot through a pixel is allocated a random value from the relevant distribution(s)
    - all effects can be achieved this way with sufficient rays per pixel
  - 2 each ray spawns multiple rays when it hits an object
    - this alternative can be used, for example, for area lights

**Examples of distributed ray tracing** 

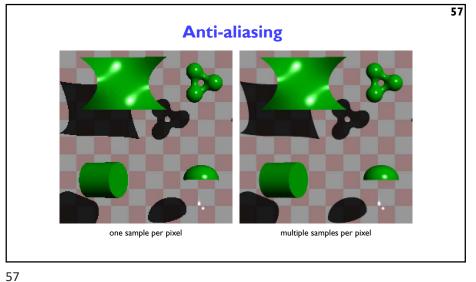
- distribute the samples for a pixel over the pixel area
  - get random (or jittered) super-sampling
  - used for anti-aliasing
- distribute the rays going to a light source over some area
  - allows area light sources in addition to point and directional light sources
  - produces soft shadows with penumbrae
- distribute the camera position over some area
  - allows simulation of a camera with a finite aperture lens
  - produces depth of field effects
- distribute the samples in time
  - produces motion blur effects on any moving objects

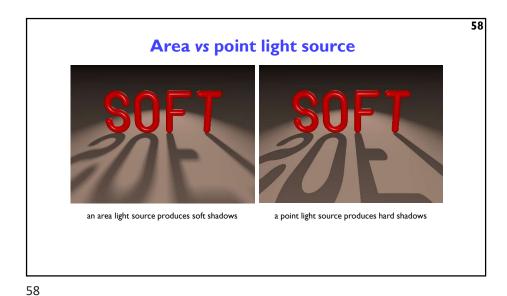
55

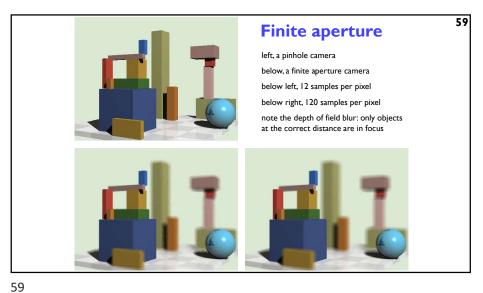
55

56

60







60 **Introduction to Computer Graphics +** Background **→** Rendering +Graphics pipeline Polygonal mesh models • Transformations using matrices in 2D and 3D Homogeneous coordinates Projection: orthographic and perspective **→** Rasterization **→** Graphics hardware and modern OpenGL → Human vision, colour and tone mapping

Introduction to Graphics

Unfortunately...

Ray tracing is computationally expensive

used for super-high visual quality

Video games and user interfaces need something faster

Most real-time applications rely on rasterisation

Model surfaces as polyhedra – meshes of polygons

Use composition to build scenes

Apply perspective transformation and project into the plane of the screen

Work out which surface was closest

Fill pixels with the colour of the nearest visible polygon

Graphics cards have hardware to support this

Ray tracing starts to appear in real-time rendering

The new generations of GPUs offer accelerated ray-tracing

But it is still not as efficient as rasterisation

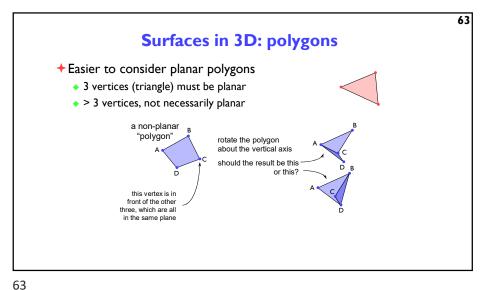
Three-dimensional objects

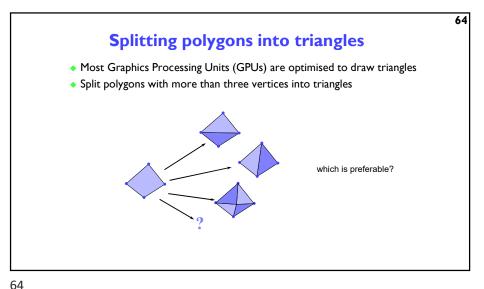
Polyhedral surfaces are made up from meshes of multiple connected polygons

Polygonal meshes
open or closed

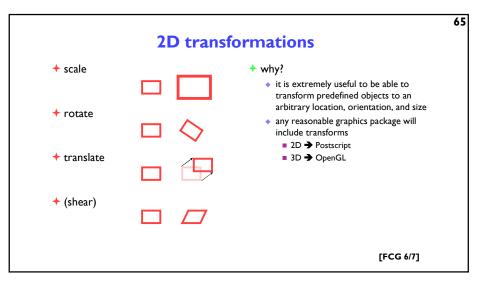
Curved surfaces
must be converted to polygons to be drawn

Michaelmas Term 2025/2026





68



66 **Basic 2D transformations** scale x' = mxabout origin y' = myby factor m rotate  $x' = x \cos \theta - y \sin \theta$ about origin by angle θ  $y' = x \sin \theta + y \cos \theta$ translate along vector  $(x_o, y_o)$  $x' = x + x_0$  $y' = y + y_0$ shear parallel to x axis x' = x + ay■ by factor *a* y' = y

5

65

Matrix representation of transformations

+ scale

• about origin, factor m  $\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} m & 0 \\ 0 & m \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$ + do nothing

• identity

• parallel to x axis, factor a  $\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$   $\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} 1 & a \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$ 

**Homogeneous 2D co-ordinates** 

 translations cannot be represented using simple 2D matrix multiplication on 2D vectors, so we switch to homogeneous co-ordinates

$$(x, y, w) \equiv \left(\frac{x}{w}, \frac{y}{w}\right)$$

- an infinite number of homogeneous coordinates maps to every 2D point
- w=0 represents a point at infinity
- usually take the inverse transform to be:

$$(x, y) \equiv (x, y, 1)$$

The symbol ≡ means equivalent

[FCG 6.3/7.3]

67

68

66

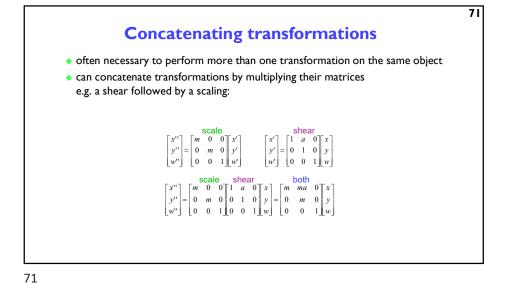
Matrices in homogeneous co-ordinates

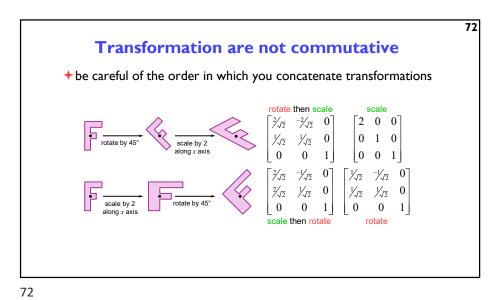
+ scale

+ about origin, factor m  $\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} m & 0 & 0 \\ 0 & m & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w' \end{bmatrix} = \begin{bmatrix} 1 & a & 0 \\ 0 & 0 &$ 

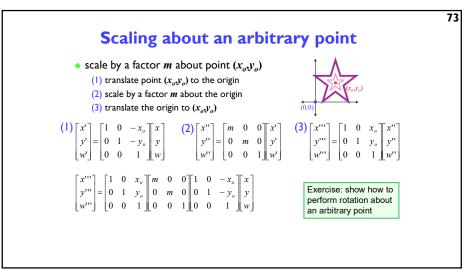
Translation by matrix algebra  $\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} 1 & 0 & x_o \\ 0 & 1 & y_0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$ In homogeneous coordinates  $x' = x + wx_o \qquad y' = y + wy_o \qquad w' = w$ In conventional coordinates  $\frac{x'}{w'} = \frac{x}{w} + x_0 \qquad \frac{y'}{w'} = \frac{y}{w} + y_0$ 

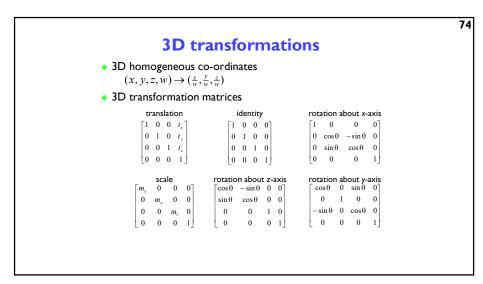
69 70

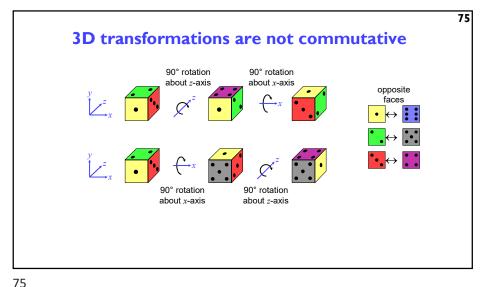


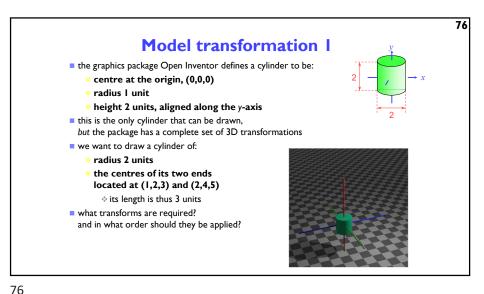


©1996–2025 Neil A. Dodgson, Peter Robinson & Rafał Mantiuk









80

**Model transformation 2** 

→ order is important:

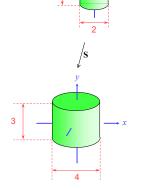
- scale first
- rotate
- translate last
- +scaling and translation are straightforward

$$\mathbf{S} = \begin{bmatrix} 2 & 0 & 0 & 0 \\ 0 & 1.5 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$
 scale from size (2,2,2)

to size (4,3,4)



translate centre of cylinder from (0,0,0) to halfway between (1,2,3) and (2,4,5)



**Model transformation 3** 

- → rotation is a multi-step process
  - break the rotation into steps, each of which is rotation about a principal axis
  - work these out by taking the desired orientation back to the original axisaligned position

**Model transformation 5** 

• then zero the x-coordinate by rotating about the z-axis • we now have the object's axis pointing along the y-axis

- the centres of its two ends located at (1,2,3) and (2,4,5)
- desired axis: (2,4,5)-(1,2,3)=(1,2,2)
- original axis: y-axis = (0,1,0)

77

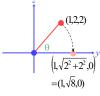
78

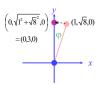
79

# **Model transformation 4**

- desired axis: (2,4,5)-(1,2,3) = (1,2,2)
- original axis: y-axis = (0,3,0)
- zero the z-coordinate by rotating about the x-axis

$$\mathbf{R}_{1} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta & 0 \\ 0 & \sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$
$$\theta = -\arcsin\frac{2}{\sqrt{2^{2} + 2^{2}}}$$





79

81

# Model transformation 6

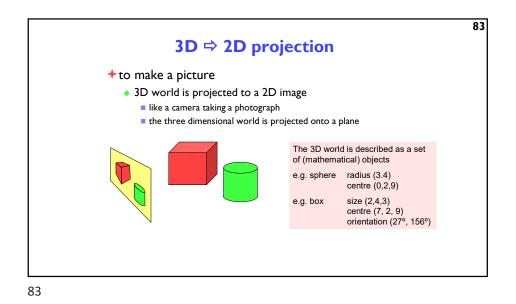
- +the overall transformation is:
  - first scale
  - then take the inverse of the rotation we just calculated
  - finally translate to the correct position

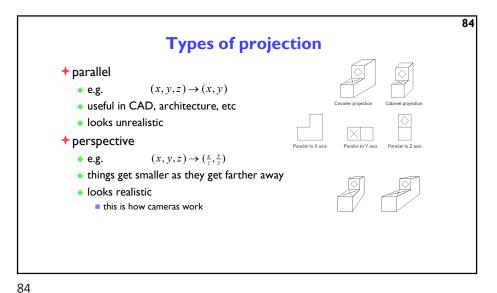
$$\begin{bmatrix} \begin{bmatrix} x' \\ y' \\ z' \\ w' \end{bmatrix} = \mathbf{T} \times \mathbf{R}_1^{-1} \times \mathbf{R}_2^{-1} \times \mathbf{S} \times \begin{bmatrix} x \\ y \\ z \\ w \end{bmatrix}$$

Application: display multiple instances

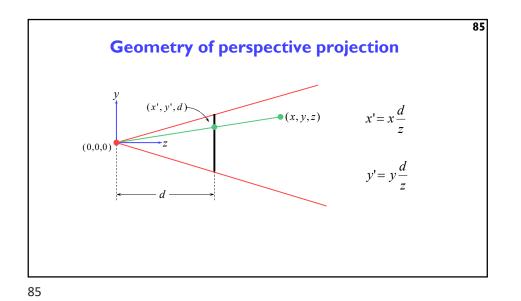
• transformations allow you to define an object at one location and then place multiple instances in your scene

82





©1996–2025 Neil A. Dodgson, Peter Robinson & Rafał Mantiuk



Projection as a matrix operation  $\begin{bmatrix} x \\ y \\ 1/d \\ z/d \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1/d \\ 0 & 0 & 1/d & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 0 \end{bmatrix} \qquad x' = x \frac{d}{z}$   $y' = y \frac{d}{z}$ This is useful in the z-buffer algorithm where we need to interpolate 1/z values rather than z values.

86

87

\_\_\_\_\_

Perspective projection with an arbitrary camera

we have assumed that:

screen centre at (0,0,d)

screen parallel to xy-plane

z-axis into screen

y-axis up and x-axis to the right

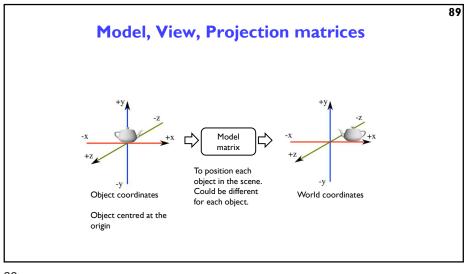
eye (camera) at origin (0,0,0)

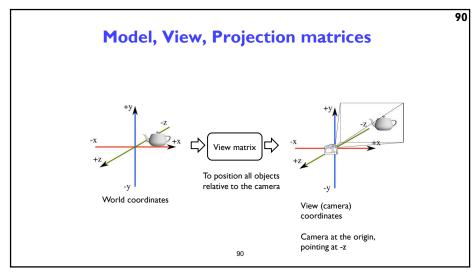
• for an arbitrary camera, we can either:

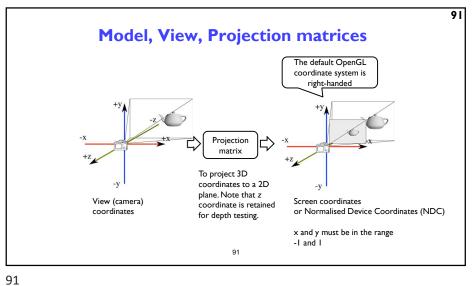
work out equations for projecting objects about an arbitrary point onto an arbitrary plane

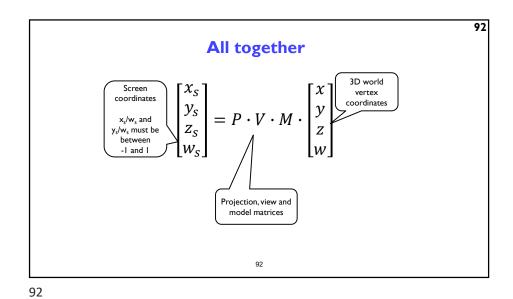
 transform all objects into our standard coordinate system (viewing coordinates) and use the above assumptions

88 A variety of transformations obiect in obiect in obiect in obiect in object world 2D screen co-ordinates co-ordinates co-ordinates co-ordinates modelling viewing transform transform the modelling transform and viewing transform can be multiplied together to produce a single matrix, taking an object directly from object coordinates into viewing coordinates either or both of the modelling transform and viewing transform matrices can be the identity e.g. objects can be specified directly in viewing co-ordinates, or directly in world co-ordinates this is a useful set of transforms, not a hard and fast model of how things should be done

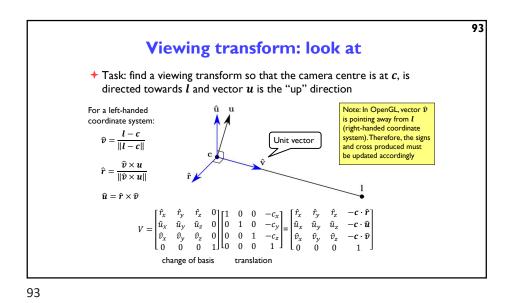








95



Transforming normal vectors

Transformation by a nonorthogonal matrix does not preserve angles

Since:  $N \cdot T = 0$ Normal transform  $N' \cdot T' = (GN) \cdot (MT) = 0$ Vertex position transform

Transformed normal and tangent vector

We can find that:  $G = (M^{-1})^T$ Derivation shown in the lecture

[FCG 6.2.2/7.2.2]

94

Scene construction

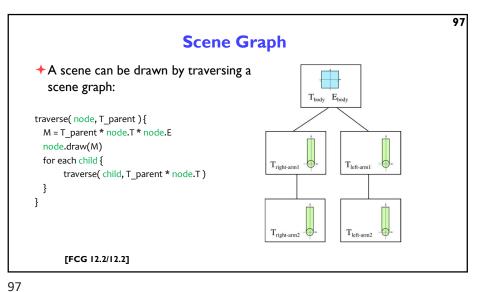
We will build a robot from basic parts
Body transformation  $M_{body} =$ Arm1 transformation  $M_{arm1} =$ Arm2 transformation  $M_{arm2} =$ 

Scene construction

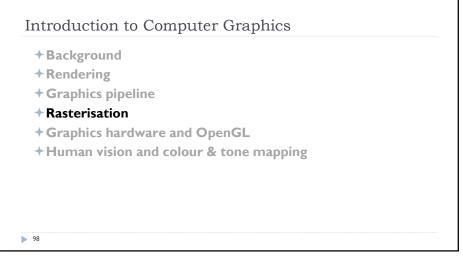
\*Body transformation  $E_{body} = scale \begin{bmatrix} 1\\2 \end{bmatrix}$   $T_{body} = translate \begin{bmatrix} x_0\\y_0 \end{bmatrix} \cdot rotate(30^\circ)$   $M_{body} = T_{body}E_{body}$ \*Arm1 transformation  $T_{arm1} = translate \begin{bmatrix} 1\\1.75 \end{bmatrix} \cdot rotate(-90^\circ)$   $M_{arm1} = T_{body}T_{arm1}$ \*Arm2 transformation  $T_{arm2} = translate \begin{bmatrix} 0\\2 \end{bmatrix} \cdot rotate(-90^\circ)$   $M_{arm2} = T_{body}T_{arm1}T_{arm2}$ Body

©1996–2025 Neil A. Dodgson, Peter Robinson & Rafał Mantiuk

Introduction to Graphics

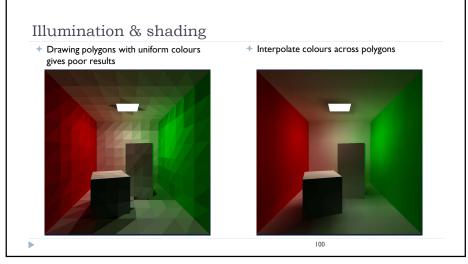


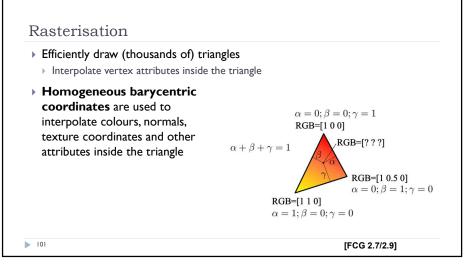
99

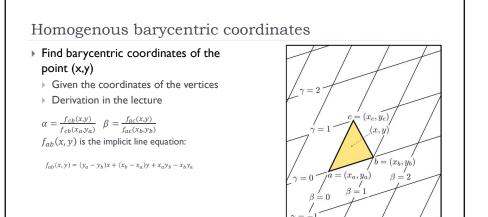


Rasterisation algorithm(\*) Set model, view and projection (MVP) transformations fragment - a candidate pixel in the triangle FOR every triangle in the scene transform its vertices using MVP matrices IF the **triangle** is within a view frustum clip the triangle to the screen border FOR each **fragment** in the triangle interpolate fragment position and attributes between vertices compute fragment colour IF the fragment is closer to the camera than any pixel drawn so far, update the screen pixel with the fragment colour END IF; END FOR; END IF; END FOR; (\*) simplified 99

98







Triangle rasterisation

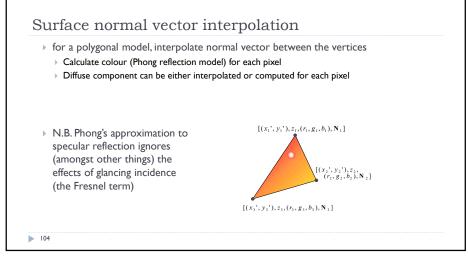
for  $y=y_{min}$  to  $y_{max}$  do

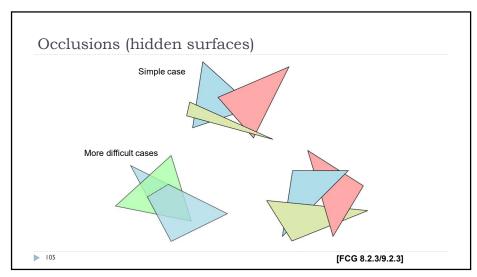
for  $x=x_{min}$  to  $x_{max}$  do  $\alpha = f_{cb}(x,y)/f_{cb}(x_a,y_a)$   $\beta = f_{ac}(x,y)/f_{ac}(x_b,y_b)$   $\gamma = 1 - \alpha - \beta$ if  $(\alpha > 0$  and  $\beta > 0$  and  $\gamma > 0$ ) then  $c = \alpha c_a + \beta c_b + \gamma c_c$ draw pixels (x,y) with colour c

Optimisation: the barycentric coordinates will change by the same amount when moving one pixel right (or one pixel down), regardless of the position

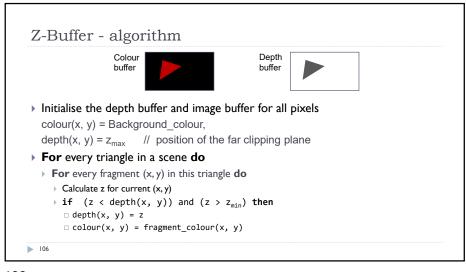
Precompute increments  $\Delta \alpha, \Delta \beta, \Delta \gamma$  and use them instead of computing barycentric coordinates when drawing pixels sequentially

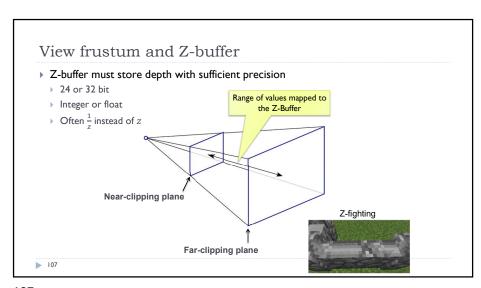
102



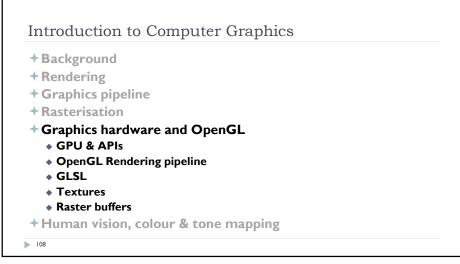


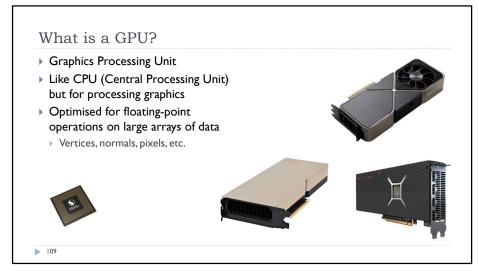
104





106





# What does a GPU do ▶ Performs all low-level tasks & a lot of high-level tasks ▶ Clipping, rasterisation, hidden surface removal, ... Essentially draws millions of triangles very efficiently Procedural shading, texturing, animation, simulation, ... ▶ Ray tracing (ray traversal, acceleration data structures) Video rendering, de- and encoding, ... Physics engines ▶ Full programmability at several pipeline stages fully programmable but optimized for massively parallel operations **110**

What makes GPU so fast?

- > 3D rendering can be very efficiently parallelized
- Millions of pixels
- ▶ Thousands of triangles
- Many operations executed independently at the same time
- ▶ This is why modern GPUs
- ▶ Contain between hundreds and thousands of SIMD processors
- Single Instruction Multiple Data operate on large arrays of data
- >>1000 GB/s memory access
  - This is much higher bandwidth than CPU
  - But peak performance can be expected for very specific operations

■ 111

111

110

# **GPU APIs** (Application Programming Interfaces)

# OpenGL

112



OpenGL.

- ▶ Multi-platform
- Open standard API
- ▶ Focus on general 3D applications
  - Dopen GL driver manages the
- No ray tracing extensions

### DirectX **DirectX**

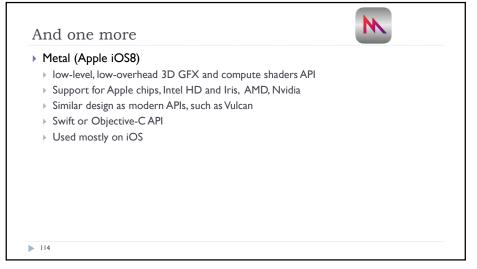
- Microsoft Windows / Xbox
- Proprietary API
- Focus on games
- Application manages resources

One more API



- ▶ Vulkan cross platform, open standard
- ▶ Low-overhead API for high performance 3D graphics
- Compared to OpenGL / DirectX
- ▶ Reduces CPU load
- ▶ Better support of multi-CPU-core architectures
- ▶ Finer control of GPU
- But
- The code for drawing a few primitives can take 1000s line of code
- Intended for game engines and code that must be very well optimized

113



GPGPU - general purpose computing

• OpenGL and DirectX are not meant to be used for general purpose computing

• Example: physical simulation, machine learning

• CUDA - Nvidia's architecture for parallel computing

• C-like programming language

• With special API for parallel instructions

• Requires Nvidia GPU

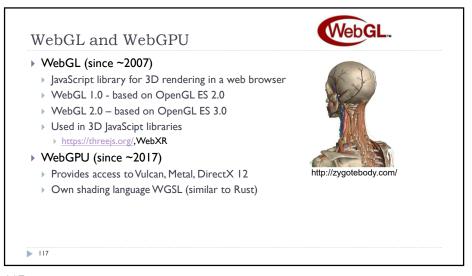
• OpenCL - Similar to CUDA, but open standard

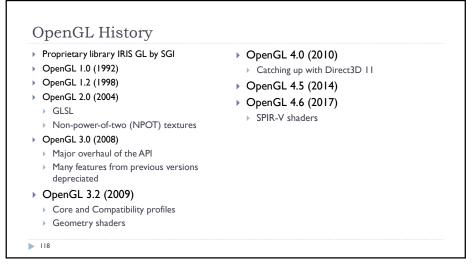
• Can run on both GPU and some CPUs

• Supported by AMD, Intel and NVidia, Qualcomm, Apple, ...

114







How to learn OpenGL?

Lectures – algorithms behind OpenGL, general principles

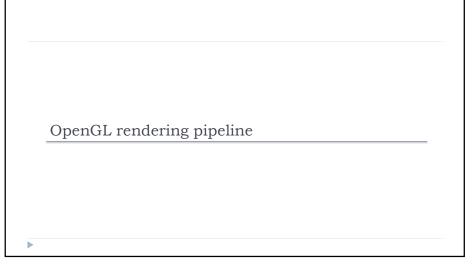
References

OpenGL Programming Guide: The Official Guide to Learning OpenGL, Version 4.5 with SPIR-V by John Kessenich, Graham Sellers, Dave Shreiner ISBN-10:0134495497

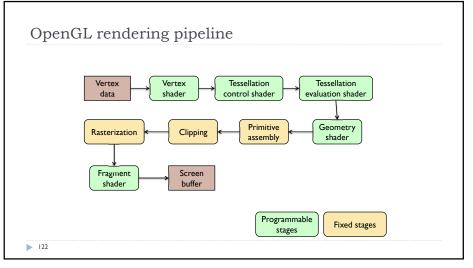
OpenGL quick reference guide https://www.opengl.org/documentation/glsl/

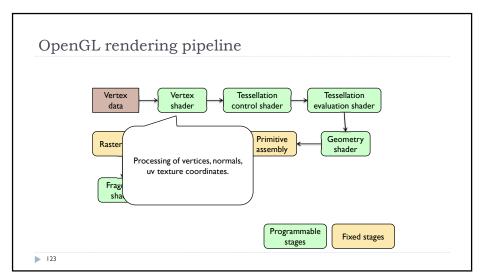
Google search: "man gl....."

118

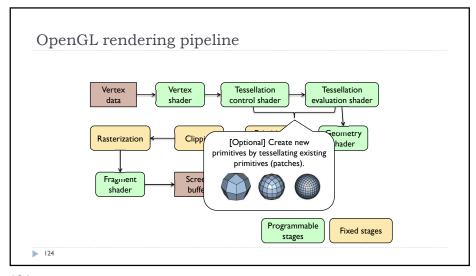


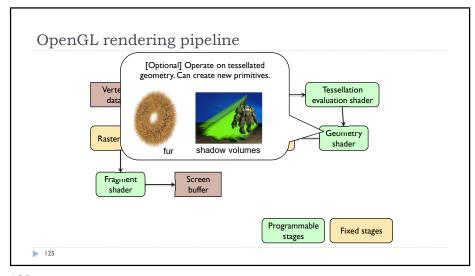
OpenGL programming model **CPU** code **GPU** code ▶ gl\* functions that Fragment shaders Create OpenGL objects Vertex shaders ▶ Copy data CPU<->GPU and other shaders Modify OpenGL state Written in GLSL ▶ Enqueue operations ▶ Similar to C Synchronize CPU & GPU From OpenGL 4.6 could be written in ▶ C99 library other language and compiled to SPIR-V Wrappers in most programming language ▶ 121

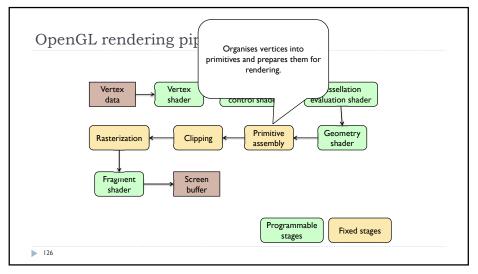


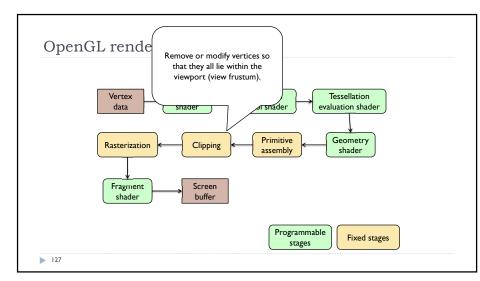


122

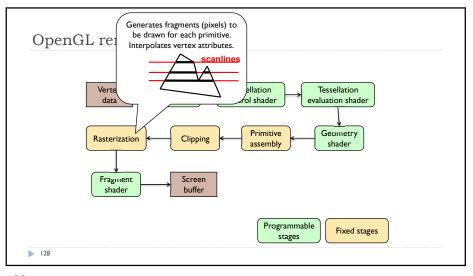


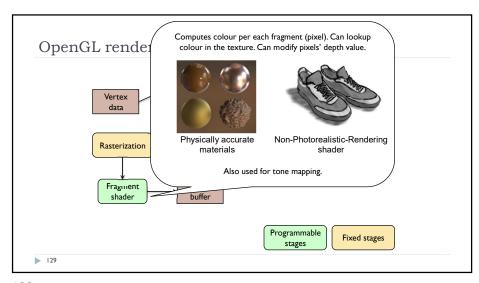


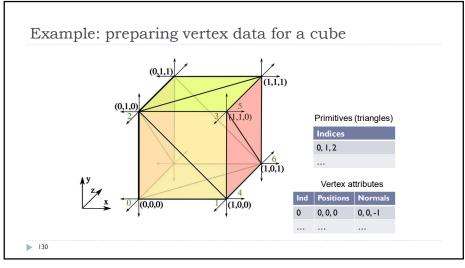


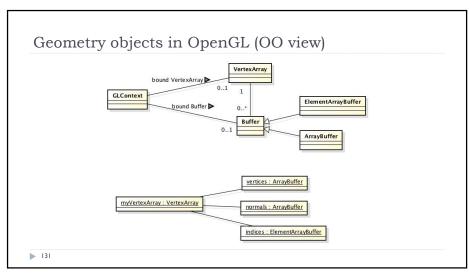


126









130

GLSL - fundamentals

Shaders

Shaders are small programs executed on a GPU

Executed for each vertex, each pixel (fragment), etc.

They are written in GLSL (OpenGL Shading Language)

Similar to C and Java

Primitive (int, float) and aggregate data types (ivec3, vec3)

Structures and arrays

Arithmetic operations on scalars, vectors and matrices

Flow control: if, switch, for, while

Functions

```
Example of a vertex shader
 #version 330
in vec3 position;
                              // vertex position in local space
in vec3 normal;
                              // vertex normal in local space
                              // fragment normal in world space
out vec3 frag_normal;
                              // model-view-projection matrix
uniform mat4 mvp matrix;
void main()
  // Typicaly normal is transformed by the model matrix
   // Since the model matrix is identity in our case, we do not modify normals
   frag_normal = normal;
   // The position is projected to the screen coordinates using mvp_matrix
   gl_Position = mvp_matrix * vec4(position, I.0);
                                             Why is this piece
                                             of code needed?
134
```

```
Data types
▶ Basic types
▶ float, double, int, uint, bool
▶ Aggregate types
▶ float: vec2, vec3, vec4; mat2, mat3, mat4
▶ double: dvec2, dvec3, dvec4; dmat2, dmat3, dmat4
▶ int: ivec2, ivec3, ivec4
▶ uint: uvec2, uvec3, uvec4
▶ bool: bvec2, bvec3, bvec4
vec3 V = vec3(1.0, 2.0, 3.0); mat3 M = mat3(1.0, 2.0, 3.0, 4.0, 5.0, 6.0, 7.0, 8.0, 9.0);
▶ 135
```

134

```
Indexing components in aggregate types

> Subscripts: rgba, xyzw, stpq (work exactly the same)

> float red = color.r;

> float v_y = velocity.y;

but also

> float red = color.x;

> float v_y = velocity.g;

> With 0-base index:

> float red = color[0];

> float m22 = M[1][1]; // second row and column

// of matrix M
```

```
Swizzling
You can select the elements of the aggregate type:
vec4 rgba_color( 1.0, 1.0, 0.0, 1.0 );
vec3 rgb_color = rgba_color.rgb;
vec3 bgr_color = rgba_color.bgr;
vec3 grayscale = rgba_color.ggg;
```

```
Arrays

> Similar to C
float lut[5] = float[5]( 1.0, 1.42, 1.73, 2.0, 2.23 );

> Size can be checked with "length()"
for( int i = 0; i < lut.length(); i++ ) {
    lut[i] *= 2;
}</pre>
```

```
Storage qualifiers

const – read-only, fixed at compile time
in – input to the shader

out – output from the shader

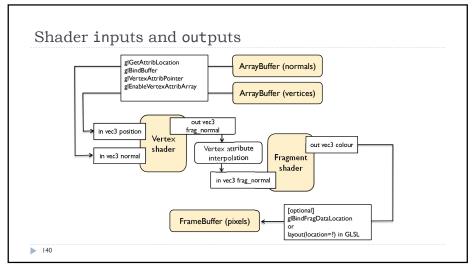
uniform – parameter passed from the application (Java), constant for the drawn geometry

buffer – GPU memory buffer (allocated by the application), both read and write access

shared – shared with a local work group (compute shaders only)

Example: const float pi=3.14;
```

138



```
GLSL Operators

Arithmetic: + - ++ --

Multiplication:

vec3 * vec3 - element-wise

mat4 * vec4 - matrix multiplication (with a column vector)

Bitwise (integer): <<, >>, &, |, ^

Logical (bool): &&, ||, ^^

Assignment:

float a=0;

a += 2.0; // Equivalent to a = a + 2.0

See the quick reference guide at: <a href="https://www.opengl.org/documentation/glsl/">https://www.opengl.org/documentation/glsl/</a>
```

```
GLSL Math

Trigonometric:
    radians( deg ), degrees( rad ), sin, cos, tan, asin, acos, atan, sinh, cosh, tanh, asinh, acosh, atanh

Exponential:
    pow, exp, log, exp2, log2, sqrt, inversesqrt

Common functions:
    abs, round, floor, ceil, min, max, clamp, ...

Graphics
    reflect, refract, inversesqrt

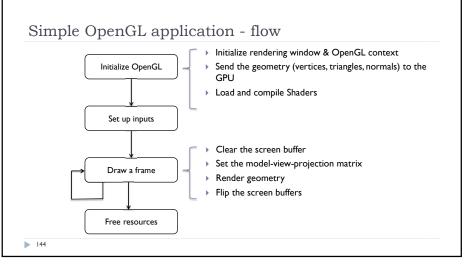
And many more

See the quick reference guide at: https://www.opengl.org/documentation/glsl/

142
```

```
GLSL flow control
 if( bool ) {
                                        for( int i = 0; i<10; i++ ) {
  // true
 } else {
  // false
                                        while( n < 10 ) {
 switch( int_value ) {
  case n:
    // statements
                                        do {
    break;
                                       } while ( n < 10 )
  case m:
    // statements
    break:
  default:
143
```

142



Rendering geometry

To render a single object with OpenGL

I.glUseProgram() – to activate vertex & fragment shaders

2.glVertexAttribPointer() – to indicate which Buffers with vertices and normals should be input to the vertex shader

3.glUniform\*() – to set uniforms (parameters of the fragment/vertex shader)

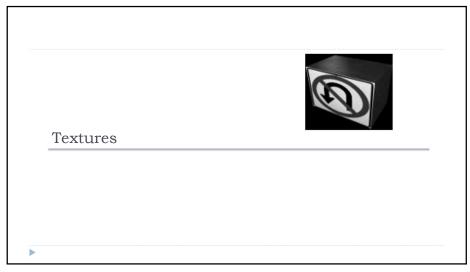
4.glBindTexture() – to bind the texture

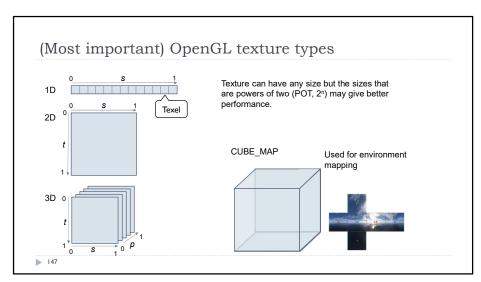
5.glBindVertexArray() – to bind the vertex array

6.glDrawElements() – to queue drawing the geometry

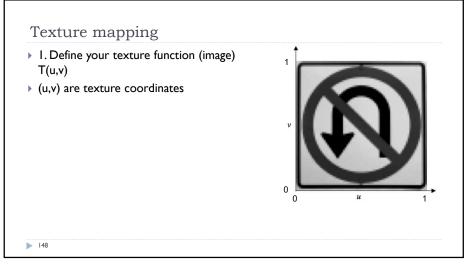
7. Unbind all objects

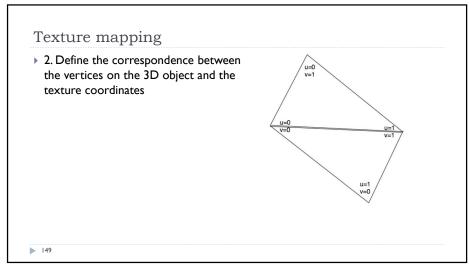
DenGL API is designed around the idea of a state-machine – set the state & queue drawing command

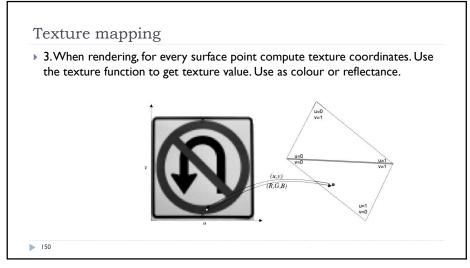


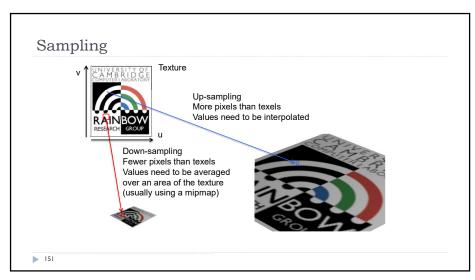


146

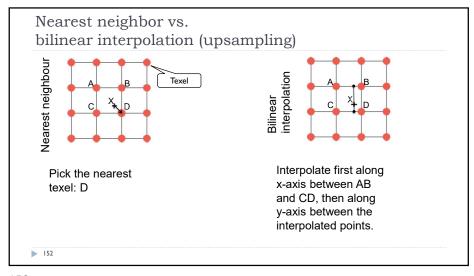


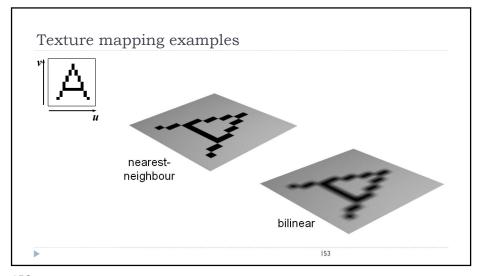


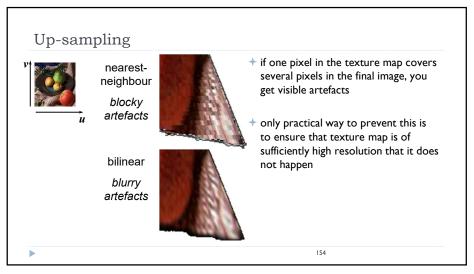


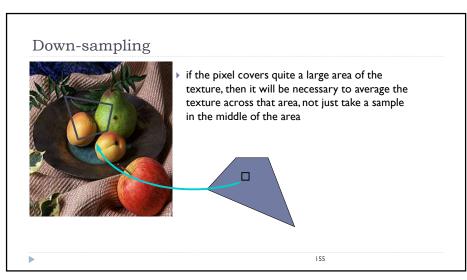


150

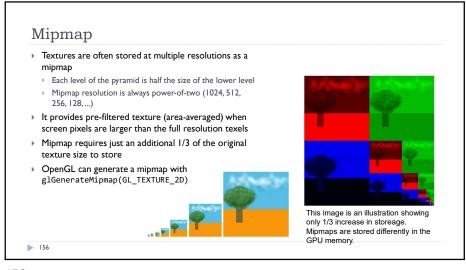


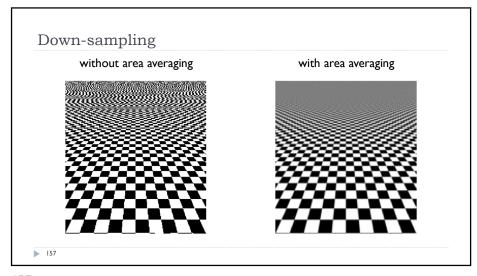


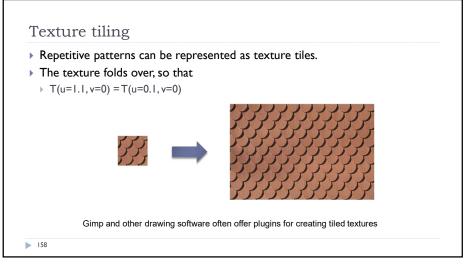


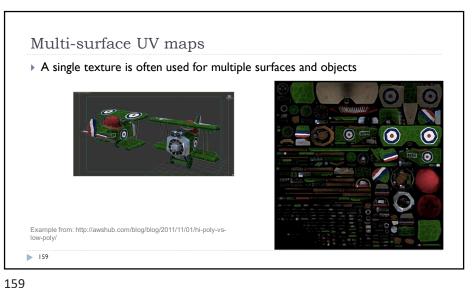


154

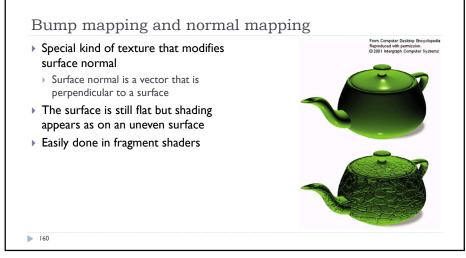


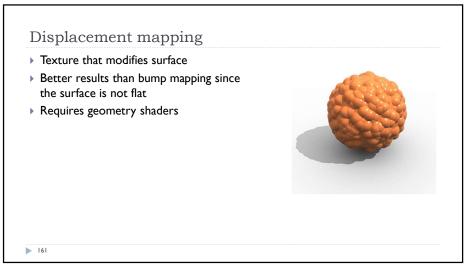


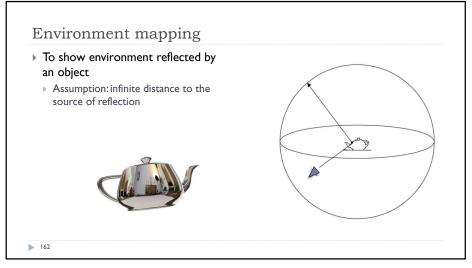


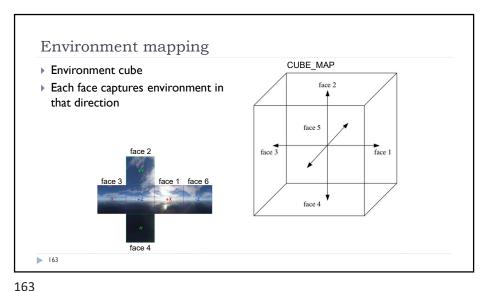


158

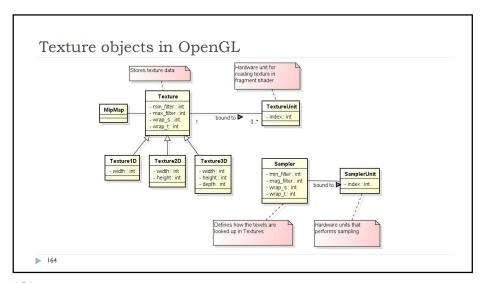








162



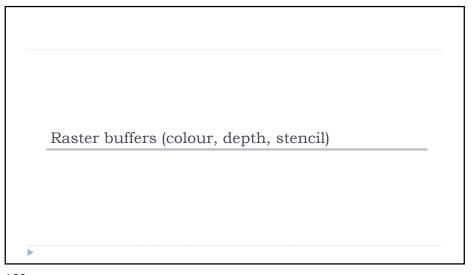
Texture parameters

//Setup filtering, i.e. how OpenGL will interpolate the pixels when scaling up or down glTexParameteri(GL\_TEXTURE\_2D, GL\_TEXTURE\_MAG\_FILTER, GL\_LINEAR); glTexParameteri(GL\_TEXTURE\_2D, GL\_TEXTURE\_MIN\_FILTER, GL\_LINEAR\_MIPMAP\_NEAREST);

How to interpolate between mipmap levels

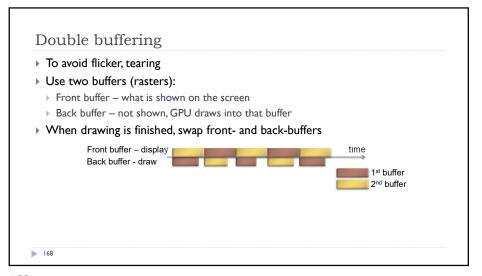
//Setup wrap mode, i.e. how OpenGL will handle pixels outside of the expected range glTexParameteri(GL\_TEXTURE\_2D, GL\_TEXTURE\_WRAP\_S, GL\_CLAMP\_TO\_EDGE); glTexParameteri(GL\_TEXTURE\_2D, GL\_TEXTURE\_WRAP\_T, GL\_CLAMP\_TO\_EDGE);

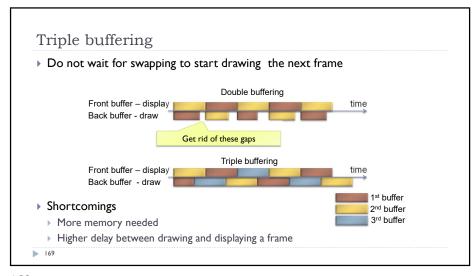
165

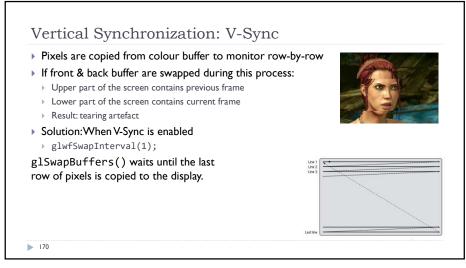


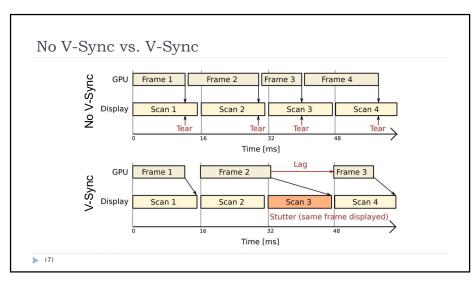
Render buffers in OpenGL Four components: GL\_BACK GL\_FRONT Colour: RGBA Typically 8 bits per component GL\_FRONT\_LEFT GL\_FRONT\_RIGHT In stereo: GL\_BACK\_LEFT GL\_BACK\_RIGHT To resolve occlusions (see Z-buffer algorithm) DEPTH Depth: Single component, usually >8 bits To block rendering selected pixels **STENCIL** Stencil: Single component, usually 8 bits. **167** 

166

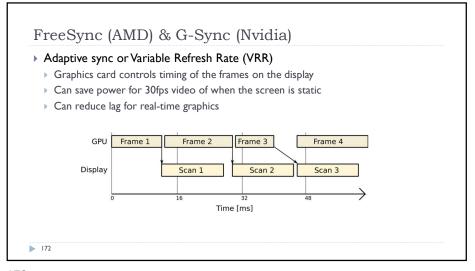


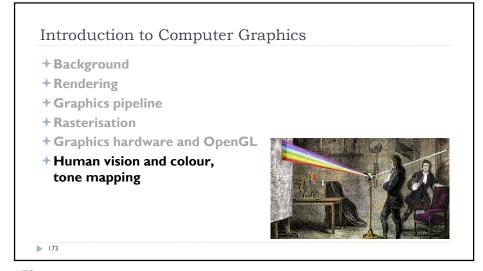


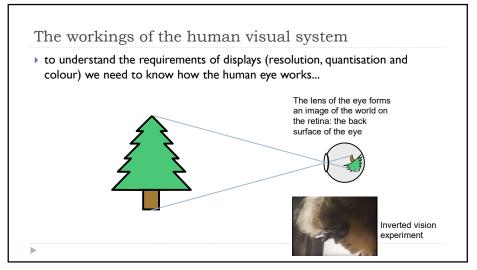




170







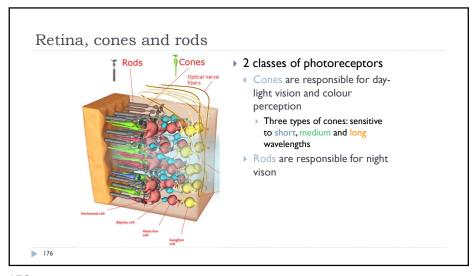
Structure of the human eye

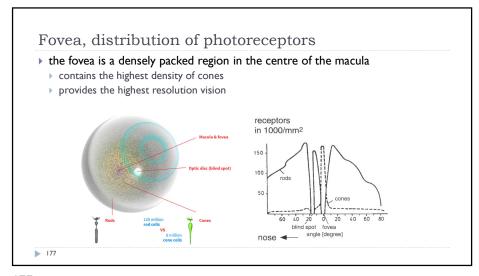
Cornea

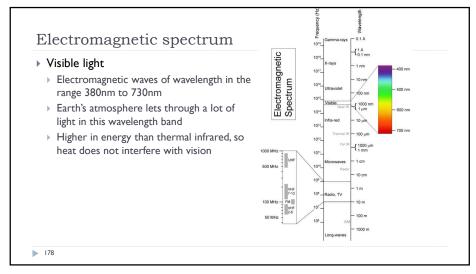
Iris
Pupil

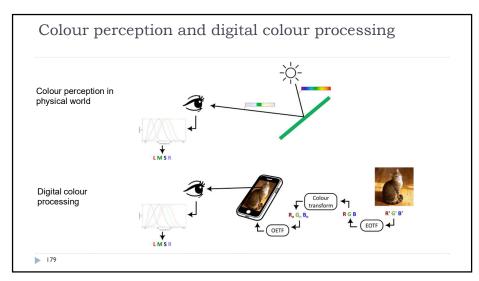
Optic
Optic
Optic
Optic
Optic
Optic
Optic
Optic
See Animagraffs web page for an animated visualization
https://animagraffs.com/human-eye/

174

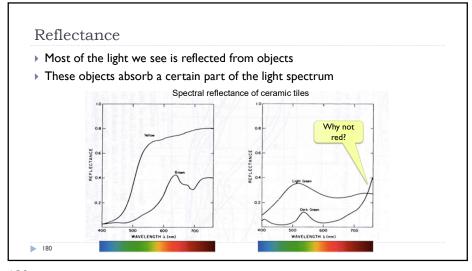


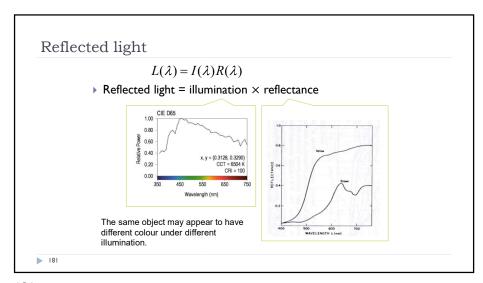


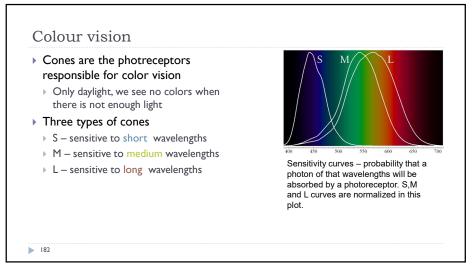


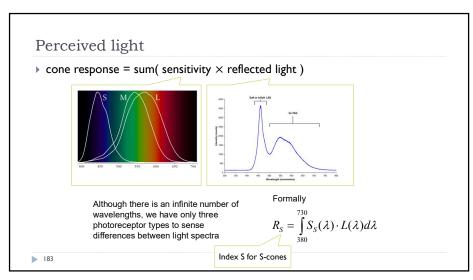


178

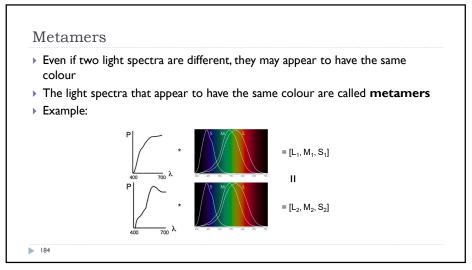


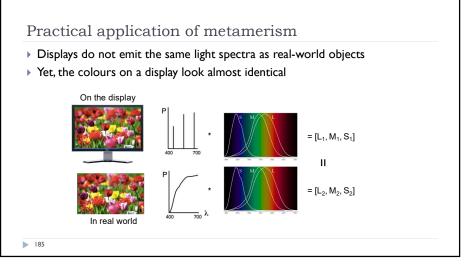


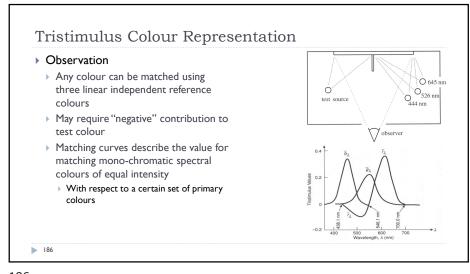




182







Standard Colour Space CIE-XYZ

- → CIE Experiments [Guild and Wright, 1931]
  - Colour matching experiments
  - ▶ Group ~12 people with normal colour vision
  - 2 degree visual field (fovea only)
  - ▶ Basis for CIE XYZ 1931 colour matching functions
- CIE 2006 XYZ
- Derived from LMS color matching functions by Stockman & Sharpe
- ▶ S-cone response differs the most from CIE 1931
- CIE-XYZ Colour Space
- ▶ Goals
- Abstract from concrete primaries used in experiment
- > All matching functions are positive
- Primary "Y" is roughly proportionally to light intensity (luminance)

187

189

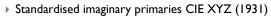
189

187

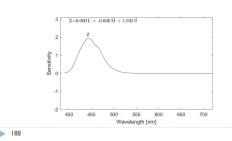
186

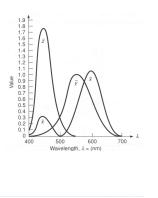
188

Standard Colour Space CIE-XYZ



- ▶ Could match all physically realizable colour stimuli
- Cone sensitivity curves can be obtained by a linear transformation of CIE XYZ



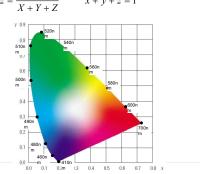


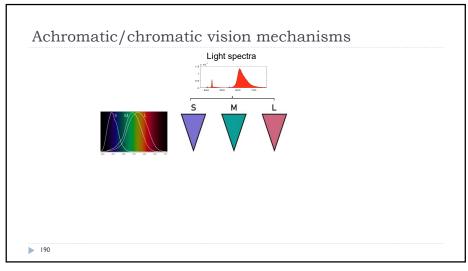
## CIE chromaticity diagram

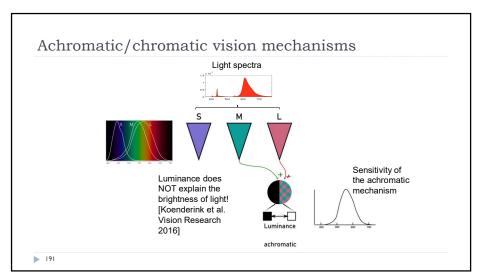
 $\blacktriangleright$  chromaticity values are defined in terms of x, y, z

$$x = \frac{X}{X+Y+Z}, \quad y = \frac{Y}{X+Y+Z}, \quad z = \frac{Z}{X+Y+Z} \qquad \quad x+y+z = \frac{Z}{X+Y+Z}$$

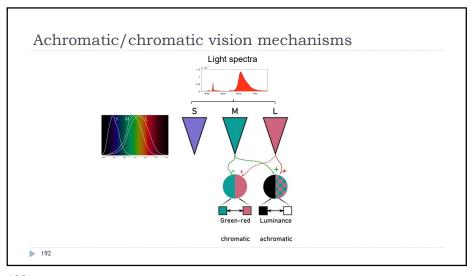
- ignores luminance
- > can be plotted as a 2D function
- pure colours (single wavelength) lie along the outer curve
- all other colours are a mix of pure colours and hence lie inside the curve
- points outside the curve do not exist as colours

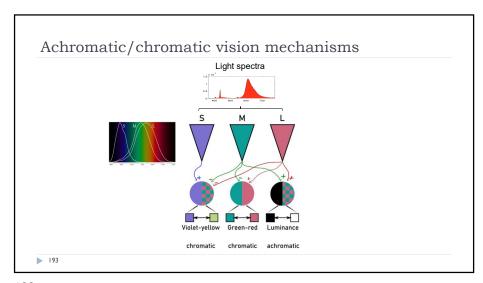


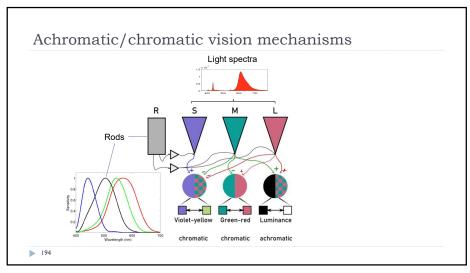


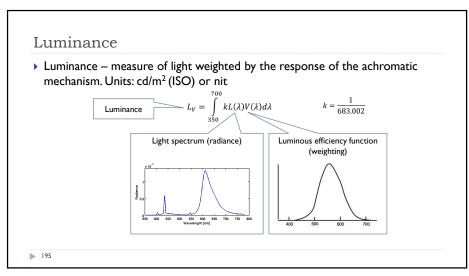


190









194

Visible vs. displayable colours

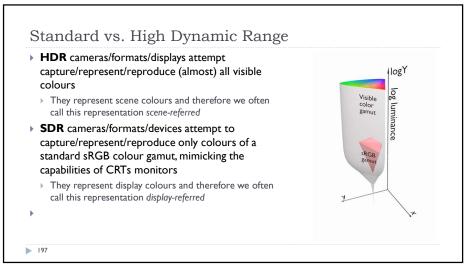
All physically possible and visible colours form a solid in XYZ space

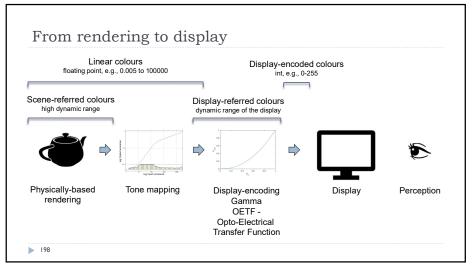
Each display device can reproduce a subspace of that space

A chromacity diagram is a slice taken from a 3D solid in XYZ space

Colour Gamut – the solid in a colour space

Usually defined in XYZ to be device-independent

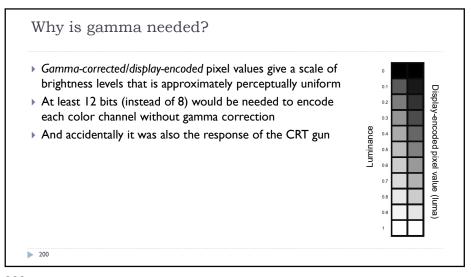


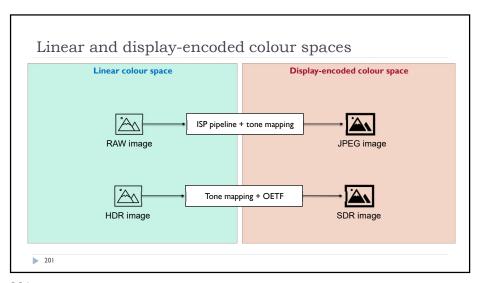


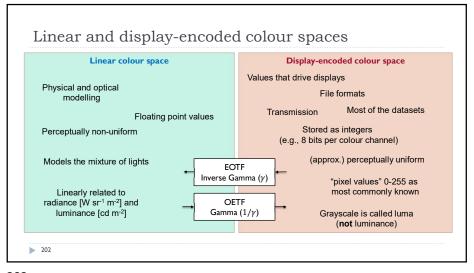
Display encoding (EOTF) for SDR: gamma correction

• Gamma correction is often used to encode luminance or tri-stimulus color values (RGB) in imaging systems (displays, printers, cameras, etc.)  $V_{out} = a \cdot V_{in}^{\gamma}$   $V_{in}^{(relative) Luminance}$   $V_{in}^{(relative) Luminance}$ 

198







Luma – gray-scale pixel value

• Luma - pixel brightness in gamma corrected units L' = 0.2126R' + 0.7152G' + 0.0722B'• R', G' and B' are gamma-corrected colour values

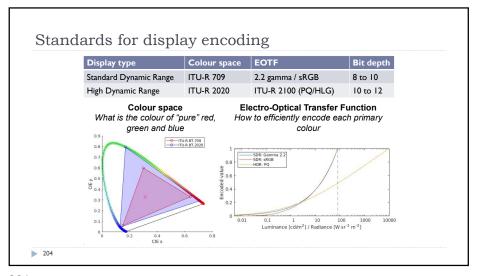
• Prime symbol denotes gamma corrected

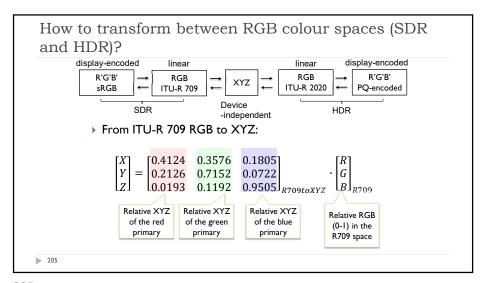
• Used in image/video coding

• Note that relative luminance if often approximated with  $L = 0.2126R + 0.7152G + 0.0722B = 0.2126(R')^{\gamma} + 0.7152(G')^{\gamma} + 0.0722(B')^{\gamma}$ • R, G, and B are linear colour values

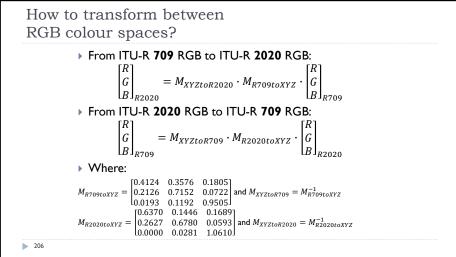
• Luma and luminace are different quantities despite similar formulas

202





206



Exercise: Map colour to a display

▶ 207

> 209

209

207

- Spectrum of the colour we want to reproduce: L (Nx1 vector)
- XYZ sensitivities: S<sub>XYZ</sub> (Nx3 matrix)
- ▶ Spectra of the RGB primaries: P<sub>RGB</sub> (Nx3 matrix)
- ightharpoonup Display gamma:  $\gamma = 2.2$
- We need to find display-encoded R'G'B' colour values
- ▶ Step I: Find XYZ of the colour

 $[X \quad Y \quad Z]^T = S_{XYZ}^T L$ 

- > Step 2: Find a linear combination of RGB primaries  $S_{XYZ}^T P_{RGB} = M_{RGB \to XYZ}$
- > Step 3: Convert and display-encode linear colour values

$$\begin{bmatrix} R & G & B \end{bmatrix}^{T} = M_{RGB \to XYZ}^{-1} \begin{bmatrix} X & Y & Z \end{bmatrix}^{T}$$

$$\begin{bmatrix} R' & G' & B' \end{bmatrix} = \begin{bmatrix} R^{1/\gamma} & G^{1/\gamma} & B^{1/\gamma} \end{bmatrix}$$

To obtain a metameric match. XYZ of the light emitted from the display and the XYZ of the spectrum L much be the same

### Representing colour

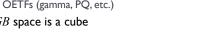
- > We need a mechanism which allows us to represent colour in the computer by some set of numbers
- A) preferably a small set of numbers which can be quantised to a fairly small number of bits each
- Display-encoded RGB, sRGB
- B) a set of numbers that are easy to interpret
- Munsell's artists' scheme
- HSV. HLS
- C) a set of numbers in a 3D space so that the (Euclidean) distance in that space corresponds to approximately perceptually uniform colour differences
  - CIE Lab. CIE Luv

208

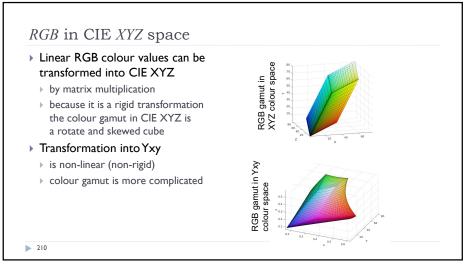
208

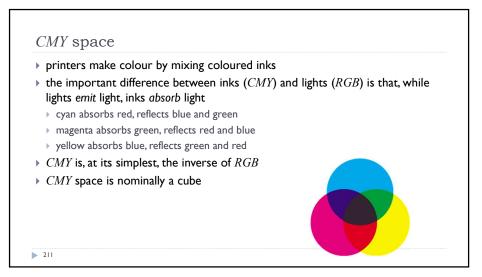
### RGB spaces

- Most display devices that output light mix red, green and blue lights to make colour
- televisions, CRT monitors, LCD screens
- ▶ RGB colour space
- ▶ Can be linear (RGB) or display-encoded (R'G'B')
- Can be scene-referred (HDR) or display-referred (SDR)
- ▶ There are multiple RGB colour spaces
  - ITU-R 709 (sRGB), ITU-R 2020, Adobe RGB, DCI-P3
  - Each using different primary colours
  - And different OETFs (gamma, PQ, etc.)
- Nominally, *RGB* space is a cube

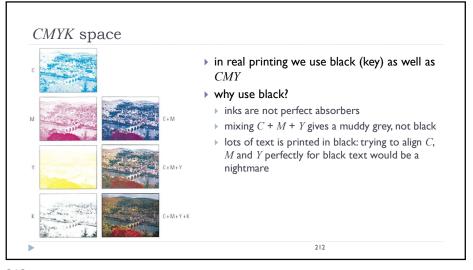


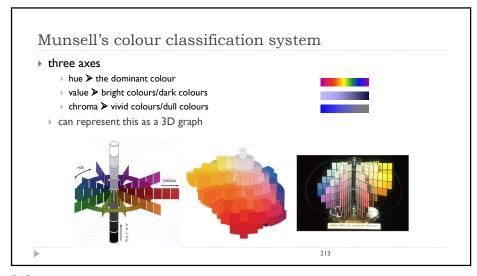


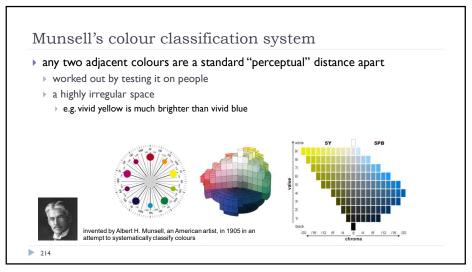




210 211



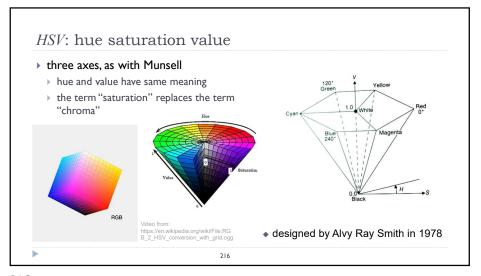




Colour spaces for user-interfaces

- ▶ *RGB* and *CMY* are based on the physical devices which produce the coloured output
- ▶ RGB and CMY are difficult for humans to use for selecting colours
- Munsell's colour system is much more intuitive:
- hue what is the principal colour?
- ▶ value how light or dark is it?
- ▶ chroma how vivid or dull is it?
- $\blacktriangleright$  computer interface designers have developed basic transformations of RGB which resemble Munsell's human-friendly system

214 215



### HLS: hue lightness saturation

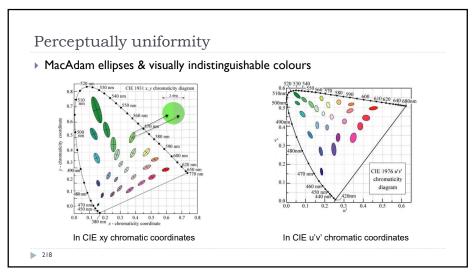
### a simple variation of ##SV

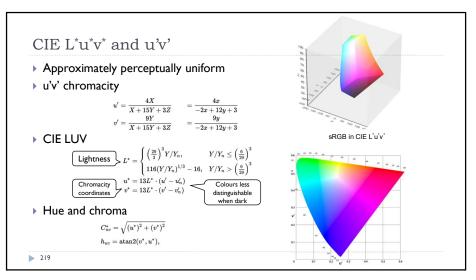
• hue and saturation have same meaning
• the term "lightness" replaces the term "value"

### designed to address the complaint that ##SV has all pure colours having the same lightness/value as white

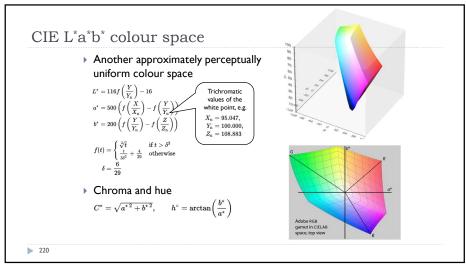
### Video from:
https://upload.wikimedia.org/wikipedia/commons
// ###/ ### Aldefres 2, HSL\_conversion\_with\_grid.ogg

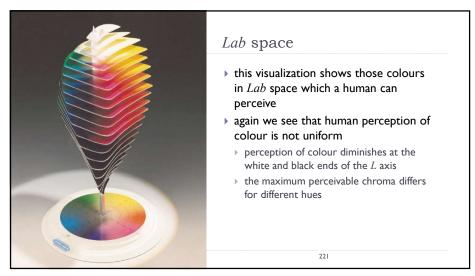
### designed by Metrick in 1979





218





#### Recap: Linear and display-encoded colour ▶ Linear colour spaces ▶ Examples: CIE XYZ, LMS cone responses, linear RGB Typically floating point numbers Directly related to the measurements of light (radiance and luminance) ▶ Perceptually non-uniform Transformation between linear colour spaces can be expressed as a matrix multiplication Display-encoded and non-linear colour spaces Examples: display-encoded (gamma-corrected, gamma-encoded) RGB, HVS, HLS, PQ-encoded RGB Typically integers, 8-12 bits per colour channel

Intended for efficient encoding, easier interpretation of colour, perceptual uniformity

Colour - references

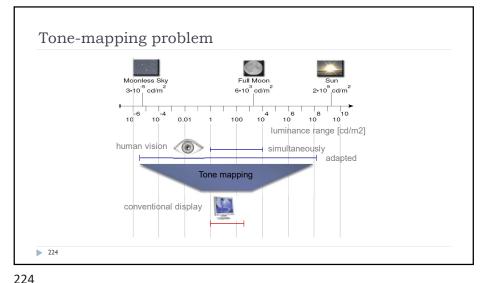
- ▶ Chapters "Light" and "Colour" in
- ▶ Shirley, P. & Marschner, S., Fundamentals of Computer Graphics
- ▶ Textbook on colour appearance
- Fairchild, M. D. (2005). Color Appearance Models (second.). John Wiley & Sons.

223

> 222

222

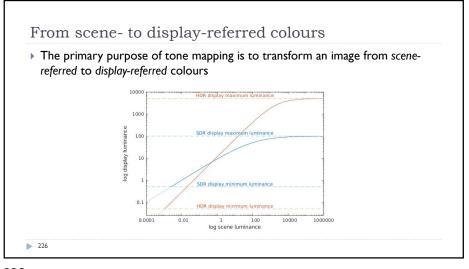
223



Why do we need tone mapping?

- ▶ To reduce dynamic range
- To customize the look (colour grading)
- To simulate human vision (for example night vision)
- To simulate a camera (for example motion blur)
- To adapt displayed images to a display and viewing conditions
- ▶ To make rendered images look more realistic
- ▶ To map from scene- to display-referred colours
- Different tone mapping operators achieve different combination of these goals

225

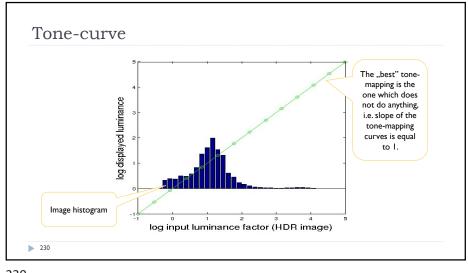


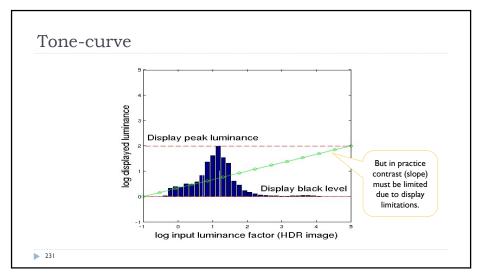
Tone mapping and display encoding Tone mapping is often combined with display encoding display-referred, scene-referred, display-referred, display-encoded, int linear, float linear, float Display encoding SDR raster Rendered HDR Tone mapping (inverse display buffer image model) Different for SDR and HDR displays Display encoding can model the display and account for Display contrast (dynamic range), brightness and ambient light levels > 227

226 227

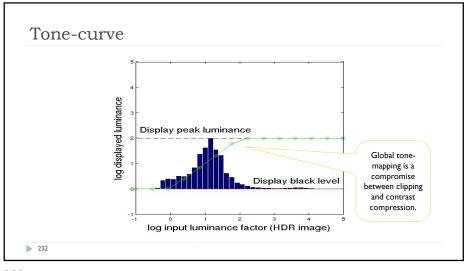
Basic tone-mapping and display coding ▶ The simplest form of tone-mapping is the exposure/brightness adjustment: Scene-referred Display-referred relative red value [0;1] Scene-referred luminance of white R for red, the same for green and blue No contrast compression, only for a moderate dynamic range ▶ The simplest form of display coding is the "gamma" Prime (') denotes a  $R' = (R_d)^{\overline{\gamma}}$ gamma-corrected value Typically  $\gamma = 2.2$ ▶ For SDR displays only 228

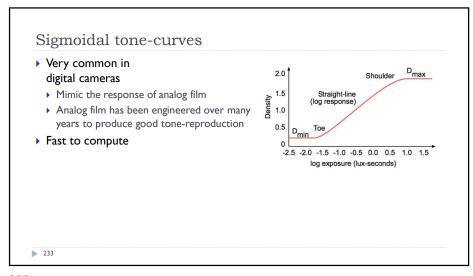
SRGB textures and display coding
 OpenGL offers sRGB textures to automate RGB to/from sRGB conversion
 SRGB textures store data in gamma-corrected space
 SRGB colour values are converted to (linear) RGB colour values on texture look-up (and filtering)
 Inverse display coding
 RGB to sRGB conversion when writing to sRGB texture
 with glEnable(GL\_FRAMEBUFFER\_SRGB)
 Forward display coding

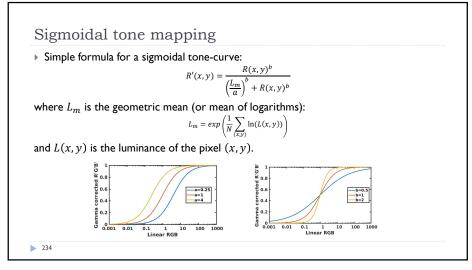


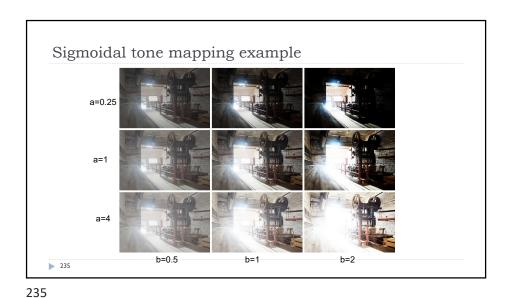


230 231









234

# Thank you for attending the lectures

- + Background
- → Rendering
- +Graphics pipeline
- **★ Rasterisation**
- + Graphics hardware and OpenGL
- +Human vision and colour & tone mapping

236