GRAPHICS & VISION SUMMER SCHOOL

Computer Graphics

Jenser **Henrik Wa**

```
light:
 shape: box [-0.5,0.5] ...
 intensity: rgb (10,10,10)
sphere:
 center: (-0.5, 0.5, -0.5)
 radius: 0.5
 material: mirror
sphere:
 center: (0.5, 0.5, 0.5)
 radius: 0.5
 material: glass
wall:
 x: -1
 colour: rgb (0.8,0.2,0.2)
\bullet \bullet \bullet
```





Art, design, architecture







1000

-401

C. .

1000

-

- St.

- 27

() 臣 :

and the second sec

......

0

(A)

Scientific visualization and training





Inverse graphics











Actually, computer graphics is **omnipresent** in how all of us interact with computers today!

...Wait, how?

Graphical user interfaces, typography





Computing without graphics



ENIAC (1945)



Punched card from a Fortran program



[root@localhost ~]# ping -q fa.wikipedia.org PING text.pmtpa.wikimedia.org (208.80.152.2) 56(84) bytes of data. ^С

--- text.pmtpa.wikimedia.org ping statistics ---

[root@localhost ~]# cd /var [root@localhost var]# ls -la total 72 drwxr-xr-x. 18 root root 4096 Jul 30 22:43 . drwxr-xr-x. 23 root root 4096 Sep 14 20:42 ... drwxr-xr-x. 2 root root 4096 May 14 00:15 account drwxr-xr-x. 11 root root 4096 Jul 31 22:26 cache drwxr-xr-x. 3 root root 4096 May 18 16:03 db drwxr-xr-x. 3 root root 4096 May 18 16:03 empty drwxr-xr-x. 2 root root 4096 May 18 16:03 games drwxrwx--T. 2 root gdm 4096 Jun 2 18:39 gdm drwxr-xr-x. 38 root root 4096 May 18 16:03 lib drwxr-xr-x. 2 root root 4096 May 18 16:03 local lrwxrwxrwx. 1 root root 11 May 14 00:12 lock -> ../run/lock drwxr-xr-x. 14 root root 4096 Sep 14 20:42 log lrwxrwxrwx. 1 root root 10 Jul 30 22:43 mail -> spool/mail drwxr-xr-x. 2 root root 4096 May 18 16:03 nis drwxr-xr-x. 2 root root 4096 May 18 16:03 opt drwxr-xr-x. 2 root root 4096 May 18 16:03 preserve drwxr-xr-x. 2 root root 4096 Jul 1 22:11 report lrwxrwxrwx. 1 root root 6 May 14 00:12 run -> ../run drwxr-xr-x. 14 root root 4096 May 18 16:03 spool drwxrwxrwt. 4 root root 4096 Sep 12 23:50 tmp drwxr-xr-x. 2 root root 4096 May 18 16:03 yp [root@localhost var]# yum search wiki Loaded plugins: langpacks, presto, refresh-packageki rpmfusion-free-updates rpmfusion-free-updates/primary_db rpmfusion-nonfree-updates updates/metalink updates updates/primary_db 73% [=====

Complete rensmitted, 1 receive, 0% packet loss, time 0ms Complete a period of 04.28 042 0 12 0 h CS

t, remove-with-leaves			
		2.7 kB	00:00
		206 kB	00:04
		2.7 kB	00:00
		5.9 kB	00:00
		4.7 kB	00:00
	1 62 kB/s	s 2.6 MB	00:15 ETA

Aspects of computer graphics



Modelling





Rendering

Animation



Modelling

How to work with geometry?

- Representation
- Manipulation and editing
- Geometric queries









Rendering

- Quantifying light and materials
- Computing light transport in a scene
- Real-time approximations







Animation

- Character animation
- Physics-based animation











Rasterization

To display any 2D or 3D shape on a pixel display, it needs to be rasterized!

Input: Geometrical "primitives" (usually triangles) with attributes (e.g. colour)

Output: Raster image approximating the given shape

Usually this is performed by the **graphics processing unit** (GPU)







•
•
×
×
×
×
×



Preview: The (real-time) graphics pipeline

Even for 3D graphics,

1. first we project the vertices of each 3D triangle to their 2D locations on the screen

2. then we rasterize the 2D triangle!

So it makes sense to study rasterization of 2D graphics first.

How to draw an arbitrary triangle on a pixel grid?

For now, let's pick a sample point at the center of each pixel, and colour the pixel if the sample point lies inside the triangle.



How to check whether a point is inside a triangle?



A point **p** is inside triangle **abc** if:

- **p** is to the left of edge **ab**, and
- **p** is to the left of edge **bc**, and
- **p** is to the left of edge **ca**.



Edge tangent vector:

 $\mathbf{t} = \mathbf{b} - \mathbf{a} = (b_x - a_x, b_y - a_y)$ Edge "normal" vector: $\mathbf{n} = \operatorname{perp}(\mathbf{t}) = (-t_y, t_x)$ $\mathbf{p} \text{ is to the left of } \mathbf{ab} \text{ if}$

 $\mathbf{n} \cdot (\mathbf{p} - \mathbf{a}) \ge 0$



Points to the left of **ab**



Points to the left of **ab** and to the left of **bc**



Points to the left of **ab** and to the left of **bc** and to the left of **ca**

Would this still work if the vertices were given in clockwise order instead?



Easy to fix: First check if **c** is to the left or the right of **ab**. But, better if you ensure all triangles are anticlockwise in the first place.



So, here's what our rasterization algorithm looks like so far. drawTriangle(triangle, colour): for $x = 0 \dots$ imageWidth: for y = 0 ... imageHeight: if isInside(x, y, triangle): image[x, y] = colour

Is this an efficient algorithm?

How can we make it faster?

Better to only check the pixels in the **bounding box** of the triangle.

What are the coordinates of this box?

 $x_{\min} = \min \{a_x, b_x, c_x\},\$

• • •

Are there any cases where this is also terribly inefficient?









What about more complex shapes?

Split them up into triangles, then draw each triangle

for each triangle:
 for each sample (x, y) it covers:
 image[x, y] = triangle.colour(x, y)

Something to think about: What if we did this instead?

for each sample (x, y):
 for each shape that covers it:
 image[x, y] = shape.colour(x, y)





Transformations



Transformations









Translation

Rotation



Applications: Instancing

Star Wars: Episode II – Attack of the Clones (2002)



Applications: Posing





Transformation matrices

As you probably know, we can represent many transformations by matrices:

$$\mathbf{v} = \begin{bmatrix} v_x \\ v_y \end{bmatrix} \qquad \qquad \mathbf{A} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$$

and similarly in 3D:

$$\mathbf{v} = \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix} \qquad \mathbf{A} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}$$

$$\mathbf{Av} = \begin{bmatrix} a_{11}v_x + a_{12}v_y \\ a_{21}v_x + a_{22}v_y \end{bmatrix}$$

$$a_{13}$$

 a_{23}
 a_{33}

$$\mathbf{Av} = \begin{bmatrix} a_{11}v_x + a_{12}v_y + a_{13}v_z \\ a_{21}v_x + a_{22}v_y + a_{23}v_z \\ a_{31}v_x + a_{32}v_y + a_{33}v_y \end{bmatrix}$$

What are the matrices for these transformations?

What can't matrices do?



$\mathbf{v}_{new} \neq \mathbf{A}\mathbf{v}_{old}$



Nonlinear deformation



Transformations

A transformation is just a function that map points to points

Now: linear transformations (easy to represent with matrices)

Later: affine transformations (linear transformations + translation)

 $T: \mathbb{R}^n \to \mathbb{R}^n$





Linear algebra

Linear algebra

Linear algebra is not about little lists of numbers!



A vector only has coordinates once you make an (arbitrary) choice of basis



Outcomes of operations should be independent of arbitrary choices!

If $\mathbf{a} + 2\mathbf{b} = \mathbf{c}$ in my basis, then it should be true in your basis as well. Best to think in a **basis-independent** way as much as possible

Though, to compute anything we will always need a basis in the end...



What are vectors, really?

- A vector is an element of a vector space.
- A vector space over \mathbb{R} is any set V equipped with two operations:
- scalar multiplication: $\mathbb{R} \times V \rightarrow V$
- vector addition: $V \times V \rightarrow V$
- satisfying various identities, e.g. $\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$, $a(\mathbf{u} + \mathbf{v}) = a\mathbf{u} + a\mathbf{v}$, etc.

To do geometry, we also need a third operation:

• dot product / inner product: $V \times V \rightarrow \mathbb{R}$ satisfying identities like $\mathbf{u} \cdot \mathbf{v} = \mathbf{v} \cdot \mathbf{u}$, $(a\mathbf{u} + b\mathbf{v}) \cdot \mathbf{w} = a(\mathbf{u} \cdot \mathbf{w}) + b(\mathbf{v} \cdot \mathbf{w})$, etc.

and in any *n* dimensions!

(It will also work for other vector spaces: functions, images, etc. ...)



- Think of these three operations as the "public API" of the vector data type.
- Write your algorithm and code in terms of these, and it will work in 2D, in 3D,

Bases

A basis is just a set of vectors $\{e_1, e_2, ...\}$ such that any vector can be written **uniquely** as a linear combination of them.

$$\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \\ \vdots \end{bmatrix} \text{ in this basis } \Leftrightarrow \mathbf{v} = v_1 \mathbf{e}_1$$

What happens when you apply a matrix **A** to the basis vectors?

$$\mathbf{Ae}_{1} = \begin{bmatrix} a_{11} & a_{12} & \cdots \\ a_{21} & a_{22} & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ \vdots \end{bmatrix} = \begin{bmatrix} a_{11} \\ a_{21} \\ \vdots \end{bmatrix} = 1 \text{ st column of } \mathbf{A}$$





This determines the action of **A** on all other vectors!

 $Av = A(v_1e_1 + v_2e_2 + \cdots) = v_1(Ae_1) + v_2(Ae_2) + \cdots = v_1a_1 + v_2a_2 + \cdots$



- other vectors follow.
- Interpretation 2: Matrix-vector multiplication Av produces a linear



• Interpretation 1: A matrix transforms the basis vectors to its columns; all

combination of the columns of A, weighted by the components v_1 , v_2 , ...

Now, what is the matrix for this transformation?



 $\mathbf{a}_1 = \text{image of } \mathbf{e}_1 \approx \begin{bmatrix} -0.8 \\ 0.5 \end{bmatrix}$



8],
$$\mathbf{a}_2 = \text{image of } \mathbf{e}_2 \approx \begin{bmatrix} 1.1 \\ 0.2 \end{bmatrix}$$

 $\mathbf{A} \approx \begin{bmatrix} -0.8 & 1.1 \\ 0.5 & 0.2 \end{bmatrix}$





Composition of transformations

Apply transformation **A** then transformation **B**:

Column interpretation:

 $\mathbf{B}[\mathbf{a}_1 \ \mathbf{a}_2 \ \cdots] = \begin{bmatrix} \mathbf{B}\mathbf{a}_1 \ \mathbf{B}\mathbf{a}_2 \ \cdots \end{bmatrix}$

- $\mathbf{v} \rightarrow \mathbf{A}\mathbf{v} \rightarrow \mathbf{B}(\mathbf{A}\mathbf{v}) = (\mathbf{B}\mathbf{A})\mathbf{v}$



Often, want to apply a sequence of *n* transformations on millions of vertices. Just compute the product: then only 1 matrix-vector multiplication per vertex.



 $AB \neq BA$







Rotations in 3D

Rotations about the coordinate axes:



Are these all the possible rotations?

Rotations in 3D

Are these all possible rotations?

Not at all!

A rotation is any transformation which:

- preserves distances and angles
- preserves orientation
 ("F" can become "J" but not "\")

Equivalently, $\mathbf{R}^{\mathsf{T}}\mathbf{R} = \mathbf{I}$, and det $\mathbf{R} = 1$



Rodrigues' rotation formula

Rotation around an axis **n** by angle θ :

How? Hints:

- $[\mathbf{n}]_{\times}$ is the "cross-product matrix": $[\mathbf{n}]_{\times} \mathbf{v} = \mathbf{n} \times \mathbf{v}$
- Assume an orthogonal basis \mathbf{n} , \mathbf{e}_1 , \mathbf{e}_2 and see what \mathbf{R} does to it





Given unit vectors \mathbf{u} and \mathbf{v} , find a way maps \mathbf{u} to \mathbf{v} , i.e. $\mathbf{Ru} = \mathbf{v}$.

Is it unique, or are there many different such rotations?



Given unit vectors **u** and **v**, find a way to construct a rotation matrix **R** which

15 minute break

Translations

Move all points by a constant displacement

So a linear transformation followed by a translation will be of the form $T(\mathbf{p}) = \mathbf{A}\mathbf{p} + \mathbf{b}$

A bit tedious to compose:

 $T_2(T_1(\mathbf{p})) = \mathbf{A}_2(\mathbf{A}_1\mathbf{p} + \mathbf{b}_1) + \mathbf{b}_2 = (\mathbf{A}_2\mathbf{A}_1)\mathbf{p} + (\mathbf{A}_2\mathbf{b}_1 + \mathbf{b}_2)$









Suppose I have both points and directions/velocities/etc. to transform.



It seems translation should only affect some things, not others. But why?



Translation by (0, 0.5): $\mathbf{p} = (0.5, 1)$ v = (1, 0.5)?

Are points really vectors?



$p_1 + p_2 = ?$ $5p_3 = ?$

How about I just choose an origin and then add the displacement vectors?

Points vs. vectors

Points form an affine space A over the vector space V.

- Point-vector addition: $A \times V \rightarrow A$
- Point subtraction: $A \times A \rightarrow V$





with the obvious properties e.g. $(\mathbf{p} + \mathbf{u}) + \mathbf{v} = \mathbf{p} + (\mathbf{u} + \mathbf{v}), \mathbf{p} + (\mathbf{q} - \mathbf{p}) = \mathbf{q}$, etc.

Coordinate frames

To specify a vector numerically, we need a basis

 $p = p_1 e_1 + p_2 e_2 + \dots + o$ so r





To specify a point numerically, we need a coordinate frame: origin and basis

maybe
$$\mathbf{p} = \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ 1 \end{bmatrix}$$
?









e.g. $\begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} p_x \\ p_y \\ 1 \end{bmatrix} = \begin{bmatrix} s_x p_x \\ s_y p_y \\ 1 \end{bmatrix}$

Translation by a vector **t**: $\begin{vmatrix} \mathbf{I} & \mathbf{t} \\ \mathbf{0} & 1 \end{vmatrix}$, mapping $\mathbf{e}_i \rightarrow \mathbf{e}_i$ but $\mathbf{o} \rightarrow \mathbf{o} + \mathbf{t}$ e.g. $\begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix}$

$$\begin{bmatrix} p_x \\ p_y \\ 1 \end{bmatrix} = \begin{bmatrix} p_x + t_x \\ p_y + t_y \\ 1 \end{bmatrix}$$

What about vectors?

Apply a translation:

$\mathbf{v} = v_1 \mathbf{e}_1 + v_2 \mathbf{e}_2 + \dots + 0 \mathbf{o} \quad \Leftrightarrow \quad \mathbf{v} = \begin{vmatrix} v_1 \\ v_2 \\ \vdots \\ \vdots \end{vmatrix}$

 $\begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \\ 0 \end{bmatrix} = \begin{bmatrix} v_x \\ v_y \\ 0 \end{bmatrix}$





Homogeneous coordinates

Add an extra coordinate w at the end.

- Points: w = 1
- Vectors: w = 0

Transformations become $(n+1)\times(n+1)$ matrices

- Linear transformations: $\begin{bmatrix} A & 0 \\ 0 & 1 \end{bmatrix}$
- Translations:
 I t
 0 1

General affine transformation: $\begin{bmatrix} A & t \\ 0 & 1 \end{bmatrix}$

- Corresponds to linearly transforming basis vectors $\mathbf{e}_i \rightarrow \mathbf{A}\mathbf{e}_i$ and translating origin $\mathbf{o} \rightarrow \mathbf{o} + \mathbf{t}$
- Maps parallel lines to parallel lines, but does not preserve the origin
- Composition: just matrix multiplication again.

 $\begin{bmatrix} I \\ J \end{bmatrix}$

Example: Rotate by given angle θ about given point **p** (instead of about origin)





Translate by -**p**

 $\mathbf{M} = \mathbf{T}(\mathbf{p}) \ \mathbf{R}(\boldsymbol{\theta}) \ \mathbf{T}(-\mathbf{p})$





Rotate by θ about origin

Translate by **p**

Given coordinates of **p** in frame 1, what are its coordinates in frame 2?

 $p = p_1 e_1 + p_2 e_2 + \cdots + o$

Write coords of e_1 , e_2 , ... and o in frame 2:





Change of coordinates looks exactly like a transformation matrix!



Active transformation: Moves points to new locations in the same frame

Change of coordinates (passive transformation): Gives coordinates of the same point in a different frame

Matrices are the same but the meaning is different! You have to keep track.

e.g. world_driver = world_from_car * car_driver Vec3 Mat3x3



Vec3



So far we know:

- How to draw 2D shapes
- How to transform 2D and 3D shapes

Today: How to draw 3D shapes on a 2D screen?

Parallel projection

Easy way: Just drop one of the coordinates lol

- Useful for engineering drawings
- Doesn't match how eyes and cameras actually see things!



Perspective


Algorithmic drawing in the 1500s

A point is drawn where the ray from



Pinhole camera model



sensor / retina



Assume camera is at the origin, pointing in the direction -z.

Where is the point **p** projected to?

$$\frac{x}{z} = \frac{u}{d}$$
$$u = \frac{xd}{z}$$

Similarly v = yd/z

(W.l.o.g., let's take d = -1)

What if the camera is not at the origin and/or not looking along -z?



Just change to a coordinate system in which it is.

Viewing transformation



- center of projection c
- Construct orthonormal basis

Usually, user specifies:

- target point t or view vector $\mathbf{v} = (\mathbf{t} \mathbf{c})/\|\mathbf{t} \mathbf{c}\|$
- "up vector" u

$$\mathbf{e}_2 = (\mathbf{v} \times \mathbf{u}) / \|\mathbf{v} \times \mathbf{u}\|$$
$$\mathbf{e}_1 = \mathbf{v} \times \mathbf{e}_2$$
$$\mathbf{e}_3 = -\mathbf{v}$$



Camera \rightarrow world: $\mathbf{M} = [\mathbf{e}_1 \ \mathbf{e}_2 \ \mathbf{e}_3 \ \mathbf{c}]$ World \rightarrow camera: \mathbf{M}^{-1}

Once point is in camera space, project

$$ted point = \begin{bmatrix} xd/z \\ yd/z \end{bmatrix}$$



16mm











Canon EF Lens Work III























And finally, we can rasterize our triangles!





- Object space → world space
- World space \rightarrow camera space
- Camera space → projection plane (division by z)
- Projection plane → NDC
- NDC \rightarrow screen coordinates

Two problems:

- Every step is a matrix, except perspective division.
- Final result has lost depth information (the z coordinate): don't know which points are in front of which















Visibility a.k.a. hidden surface removal

Which surfaces are visible? Those that are not hidden by nearer surfaces.



Triangles drawn without considering depth / visibility



Correct result





LearnOpenGL.com



LearnOpenGL.com

Homogeneous coordinates revisited

Recall points vs. vectors: $\mathbf{p} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$, $\mathbf{v} = \begin{bmatrix} x \\ y \\ 0 \end{bmatrix}$

Let's generalize: points can have any $w \neq 0$



Any point in homo corresponds to the



e geneous coordinates
$$\hat{\mathbf{p}} = \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$
 with $w \neq 0$
a 2D point $\mathbf{p} = (x/w, y/w)$



The main idea: Points in 2D correspond to lines through the origin in 3D!

Linear and affine transformations still work as before!

All points $\hat{\mathbf{p}} = \begin{bmatrix} cx \\ cy \\ c \end{bmatrix}$ on a line represent the same point $\mathbf{p} = (x, y)$ where the line meets the plane w = 1

Analogy: Various tuples (2,4), (-1,-2), (5,10), ... all represent the same rational number ¹/₂



Perspective projection: $(x,y,z) \rightarrow (xd/z, yd/z)$

With homogeneous coordinates:



Hang on, we've still lost depth information.



To retain depth information, let's copy w into the z-coordinate:





$$\begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \xrightarrow{\rightarrow} \begin{bmatrix} x \\ y \\ 1/d \\ z/d \end{bmatrix} \sim \begin{bmatrix} xd/z \\ yd/z \\ 1/z \\ 1 \end{bmatrix}$$

The view frustum



Angel & Shreiner, Interactive Computer Graphics

In theory, horizontal and vertical angles of view define an infinite view cone

In practice, cut off at near and far "clipping planes": view frustum

Why?

- Exclude objects behind the camera
- Finite precision of depth coordinate (we'll see why later)







Marschner & Shirley, Fundamentals of Computer Graphics





$$\mathbf{M} = \begin{bmatrix} \frac{2|n|}{r-l} & 0 & \frac{r+l}{r-l} & 0 \\ 0 & \frac{2|n|}{t-b} & \frac{t+b}{t-b} & 0 \\ 0 & 0 & \frac{|n|+|f|}{|n|-|f|} & \frac{2|n||f|}{|n|-|f|} \\ 0 & 0 & -1 & 0 \end{bmatrix}$$

Normalized device coordinates (for real this time)



Clipping



- Discard triangles outside view frustum
- Clip triangles partially intersecting view frustum

Usually implemented in homogeneous coordinates (before division)

Keenan Crane