



FFTs in Graphics and Vision

Representing Rotations



Outline

- Math Review
 - Polynomials
 - Eigenvectors
 - Orthogonal Transformations
 - Classifying the 2D Orthogonal Transformations
- Representing 3D Rotations



Math Review

Polynomials:

Let $P(x)$ be a polynomial of degree d :

$$P(x) = a_0 + a_1 \cdot x + \cdots + a_d \cdot x^d$$

Claim:

If d is odd, the polynomial $P(x)$ must have at least one real root.



Math Review

Polynomials:


Let $P(x)$ be a polynomial of degree d :

$$P(x) = a_0 + a_1 \cdot x + \cdots + a_d \cdot x^d$$

Proof:

Consider the sign of a_d :

- If a_d is positive:
 - » As $x \rightarrow -\infty$: $P(x) \rightarrow -\infty$
 - » As $x \rightarrow +\infty$: $P(x) \rightarrow +\infty$





Math Review

Polynomials:


Let $P(x)$ be a polynomial of degree d :

$$P(x) = a_0 + a_1 \cdot x + \cdots + a_d \cdot x^d$$

Proof:

Consider the sign of a_d :

- If a_d is positive:
 - » As $x \rightarrow -\infty$: $P(x) \rightarrow -\infty$
 - » As $x \rightarrow +\infty$: $P(x) \rightarrow +\infty$
- If a_d is negative:
 - » As $x \rightarrow -\infty$: $P(x) \rightarrow +\infty$
 - » As $x \rightarrow +\infty$: $P(x) \rightarrow -\infty$





Math Review

Polynomials:

Let $P(x)$ be a polynomial of degree d :

$$P(x) = a_0 + a_1 \cdot x + \cdots + a_d \cdot x^d$$

Proof:

In either case, the value of $P(x)$ changes signs so it must have a zero-crossing somewhere.



Math Review

Eigenvectors:

Given a vector space V and invertible linear op. $A: V \rightarrow V$, if v is an e.vector of A with e.value λ then v is an e.vector of A^{-1} with e.value $1/\lambda$.

Since $A^{-1} \cdot A$ is the identity, we have:

$$\begin{aligned} v &= A^{-1}(Av) \\ &= A^{-1}(\lambda v) \\ &= \lambda \cdot A^{-1}v \end{aligned}$$

\Downarrow

$$\frac{1}{\lambda} \cdot v = A^{-1}v$$



Math Review

Orthogonal Transformations:

For a real inner-product space V , a linear map R is orthogonal if for any $v, w \in V$, we have:

$$\langle v, w \rangle = \langle Rv, Rw \rangle$$

If the determinant of R is 1, the transformation is called a rotation.



Math Review

Orthogonal Transformations (Property 1):

The set of orthogonal transformations is a group.



Math Review

Orthogonal Transformations (Property 1):

The set of orthogonal transformations is a group.

To show this we need to show that if R and S are orthogonal transformations then:

- RS is orthogonal
- R^{-1} is orthogonal



Math Review

Orthogonal Transformations (Property 1):

If R and S are orthogonal transformations, then so is the transformation $R \cdot S$.

Since R is orthogonal:

$$\langle RSv, RS w \rangle = \langle Sv, Sw \rangle$$

Since S is orthogonal:

$$\langle Sv, Sw \rangle = \langle v, w \rangle$$

Thus, as desired, we get:

$$\langle RSv, RS w \rangle = \langle v, w \rangle$$



Math Review

Orthogonal Transformations (Property 1):

If R is an orthogonal transformation, then so is the transformation R^{-1} .

Starting with the identity:

$$\langle v, w \rangle = \langle RR^{-1}v, RR^{-1}w \rangle$$

Since R is orthogonal we get:

$$\langle RR^{-1}v, RR^{-1}w \rangle = \langle R^{-1}v, R^{-1}w \rangle$$

Thus, as desired, we get:

$$\langle v, w \rangle = \langle R^{-1}v, R^{-1}w \rangle$$



Math Review

Orthogonal Transformations (Property 2):

If R is an orthogonal transformation and v is an eigenvector of R with eigenvalue λ , then $\lambda = \pm 1$.

Since R orthogonal, we have:

$$\begin{aligned}\langle v, v \rangle &= \langle Rv, Rv \rangle \\ &= \langle \lambda v, \lambda v \rangle \\ &= \lambda^2 \langle v, v \rangle \\ &\Downarrow \\ \lambda^2 &= 1\end{aligned}$$

Equivalently, if λ , is an eigenvalue of (orthogonal) R , then:

$$\frac{1}{\lambda} = \lambda.$$



Math Review

Orthogonal Transformations (Property 3):

If R is an orthogonal transformation and v is an eigenvector of R , then if w is a vector perpendicular to v , Rw is also perpendicular to v .

Since R^{-1} is also an orthogonal transformation:

$$\begin{aligned}\langle v, Rw \rangle &= \langle R^{-1}v, R^{-1}Rw \rangle \\ &= \langle R^{-1}v, w \rangle \\ &= \frac{1}{\lambda} \langle v, w \rangle \\ &= 0\end{aligned}$$



Math Review

Orthogonal Transformations (Property 4):

If R is an orthogonal transformation and v_1 and v_2 are eigenvectors of R with eigenvalues λ_1 and λ_2 , then if $\lambda_1 \neq \lambda_2$, v_1 and v_2 must be perpendicular.

Since R^{-1} is orthogonal we have:

$$\begin{aligned}\langle Rv_1, v_2 \rangle &= \langle R^{-1}Rv_1, R^{-1}v_2 \rangle \\ &= \langle v_1, R^{-1}v_2 \rangle\end{aligned}$$

$$\Downarrow$$

$$\begin{aligned}\lambda_1 \langle v_1, v_2 \rangle &= 1/\lambda_2 \langle v_1, v_2 \rangle \\ &= \lambda_2 \langle v_1, v_2 \rangle\end{aligned}$$

$$\Downarrow$$

$$\langle v_1, v_2 \rangle = 0$$



Math Review

Classifying the 2D Orthogonal Transformations:

Let V be the space of 2D arrays with the standard basis:

$$\{\mathbf{e}_1 = (1,0), \mathbf{e}_2 = (0,1)\}$$

with the standard inner product:

$$\langle \mathbf{e}_i, \mathbf{e}_j \rangle = \delta_{ij}$$



Math Review

Classifying the 2D Orthogonal Transformations:

In the basis $\{\mathbf{e}_1, \mathbf{e}_2\}$ we can express a linear operator R as a matrix:

$$\mathbf{R} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

\mathbf{R} is orthogonal if $\mathbf{R}^\top \cdot \mathbf{R}$ is the identity:

$$\begin{aligned} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} &= \begin{pmatrix} a & c \\ b & d \end{pmatrix} \cdot \begin{pmatrix} a & b \\ c & d \end{pmatrix} \\ &= \begin{pmatrix} a^2 + c^2 & ab + cd \\ ab + cd & b^2 + d^2 \end{pmatrix} \end{aligned}$$



Math Review

Classifying the 2D Orthogonal Transformations:

$$\begin{pmatrix} \boxed{1} & 0 \\ 0 & \boxed{1} \end{pmatrix} = \begin{pmatrix} \boxed{a^2 + c^2} & ab + cd \\ ab + cd & \boxed{b^2 + d^2} \end{pmatrix}$$

The diagonal entries give rise to the equations:

$$1 = a^2 + c^2$$

$$1 = b^2 + d^2$$

For appropriate θ and ϕ , this gives:

$$a = \cos \theta \quad c = \sin \theta$$

$$b = \cos \phi \quad d = \sin \phi$$



Math Review

Classifying the 2D Orthogonal Transformations:

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} a^2 + c^2 & ab + cd \\ ab + cd & b^2 + d^2 \end{pmatrix}$$

The other equations then become:

$$0 = \cos \theta \cdot \cos \phi + \sin \theta \cdot \sin \phi$$

Or equivalently:

$$0 = \cos(\theta - \phi)$$

Which implies that:

$$\phi = \theta + k\pi + \pi/2$$



Math Review

Classifying the 2D Orthogonal Transformations:


$$\phi = \theta + k\pi + \pi/2$$

If \mathbf{R} is an orthogonal transformation, then in the basis $\{\mathbf{e}_1, \mathbf{e}_2\}$ we have one of two cases:

- k is even:

$$\mathbf{R} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$$

The determinant is 1 \Rightarrow this is a rotation.





Math Review

Classifying the 2D Orthogonal Transformations:

$$\phi = \theta + k\pi + \pi/2$$

If \mathbf{R} is an orthogonal transformation, then in the basis $\{\mathbf{e}_1, \mathbf{e}_2\}$ we have one of two cases:

- k is even:


$$\mathbf{R} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$$

The determinant is 1 \Rightarrow this is a rotation.

- k is odd:

$$\mathbf{R} = \begin{pmatrix} \cos \theta & \sin \theta \\ \sin \theta & -\cos \theta \end{pmatrix}$$

The determinant is $-1 \Rightarrow$ this is a reflection.






Math Review

Classifying the 2D Orthogonal Transformations:

Claim:

In the case that k is odd, the orthogonal transformation has eigenvalues 1 and -1 .





Math Review


Classifying the 2D Orthogonal Transformations:

Claim:

In the case that k is odd, the orthogonal transformation has eigenvalues 1 and -1 .

To compute the eigenvalues, we need to solve for the roots of the polynomial:

$$P_R(\lambda) = \det(R - \lambda \cdot \text{Id})$$





Math Review


Classifying the 2D Orthogonal Transformations:

Claim:

In the case that k is odd, the orthogonal transformation has eigenvalues 1 and -1 .

To compute the eigenvalues, we need to solve for the roots of the polynomial:

$$\begin{aligned} P_R(\lambda) &= \det \begin{pmatrix} \cos \theta - \lambda & \sin \theta \\ \sin \theta & -\cos \theta - \lambda \end{pmatrix} \\ &= \lambda^2 - \cos^2 \theta - \sin^2 \theta \\ &= \lambda^2 - 1 \end{aligned}$$





Math Review

Classifying the 2D Orthogonal Transformations:


Claim:

In the case that k is odd, the orthogonal transformation has eigenvalues 1 and -1 .

To compute the eigenvalues, we need to solve for the roots of the polynomial:

$$P_R(\lambda) = \det(R - \lambda \cdot \text{Id})$$

This polynomial has two roots, $\lambda = \pm 1$.





Outline


- Math Review
- Representing 3D Rotations
 - Axis-Angle
 - Euler Angles

Representing 3D Rotations (Axis-Angle)



We will show that any rotation R can be thought of as a rotation about some axis.

In particular, we need to show that every rotation R fixes some vector v and acts as a rotation in the plane P perpendicular to v .



Representing 3D Rotations (Axis-Angle)



Let V be the space of 3D arrays with the standard basis:

$$\{\mathbf{e}_1 = (1,0,0), \mathbf{e}_2 = (0,1,0), \mathbf{e}_3 = (0,0,1)\}$$

with the standard inner product:

$$\langle \mathbf{e}_i, \mathbf{e}_j \rangle = \delta_{ij}$$

Representing 3D Rotations (Axis-Angle)



In the basis $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ we can express the linear operator R as a matrix:

$$\mathbf{R} = \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix}$$

We can compute the eigenvalues of R by finding the roots of the determinant:

$$P_R(\lambda) = \det \begin{pmatrix} a - \lambda & b & c \\ d & e - \lambda & f \\ g & h & i - \lambda \end{pmatrix}$$

Representing 3D Rotations (Axis-Angle)



Since $P_R(\lambda)$ has odd degree ($d = 3$), it must have at least one root.



Thus, R has an eigenvector v with eigenvalue λ .

Since R is orthogonal $\lambda = \pm 1$.

Representing 3D Rotations (Axis-Angle)



Thus, for every orthogonal transformation R , there must exist a vector v that is either fixed by R or mapped to its antipode.




Representing 3D Rotations (Axis-Angle)



What happens to the plane P that is orthogonal to the eigenvector v ?

Since R maps the line spanned by v back into itself, and since R is orthogonal, R must map the plane P back into itself.

Since R preserves the inner-product and maps P to itself, the restriction of R to P is a 2D orthogonal operator.






Representing 3D Rotations (Axis-Angle)

Letting $\{w_1, w_2\}$ be an orthonormal basis for the plane P , with respect to the basis $\{v, w_1, w_2\}$, we can express R in matrix form as either:

$$\mathbf{R} = \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{pmatrix}$$

or:

$$\mathbf{R} = \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \cos \theta & \sin \theta \\ 0 & \sin \theta & -\cos \theta \end{pmatrix}$$





Representing 3D Rotations (Axis-Angle)


What happens in the case that R is a rotation?

If R is a rotation, then in addition to being orthogonal, it must have determinant 1.

For the two representations of R we get:

$$\det \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{pmatrix} = \lambda$$

$$\det \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \cos \theta & \sin \theta \\ 0 & \sin \theta & -\cos \theta \end{pmatrix} = -\lambda$$






Representing 3D Rotations (Axis-Angle)

What happens in the case that R is a rotation?

If $\lambda = 1$, we have:

$$\mathbf{R} = \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{pmatrix}$$

and R is a rotation in the plane P by angle θ .



Representing 3D Rotations (Axis-Angle)




What happens in the case that R is a rotation?

If $\lambda = -1$, we get:

$$\mathbf{R} = \begin{pmatrix} -1 & 0 & 0 \\ 0 & \cos \theta & \sin \theta \\ 0 & \sin \theta & -\cos \theta \end{pmatrix}$$

and R is the composition of a reflection in the plane P and a flip about the line spanned by v .



Representing 3D Rotations (Axis-Angle)




What happens in the case that R is a rotation?

Restricting R to the plane P , in the basis $\{w_1, w_2\}$ we get:

$$\mathbf{R}|_P = \begin{pmatrix} \cos \theta & \sin \theta \\ \sin \theta & -\cos \theta \end{pmatrix}$$

which has eigenvalues -1 and 1 .




Representing 3D Rotations (Axis-Angle)



What happens in the case that R is a rotation?

In particular, if we set w_1 and w_2 to be the corresponding eigenvectors, we get:

$$\mathbf{R} = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$






Representing 3D Rotations (Axis-Angle)

What happens in the case that R is a rotation?

$$\mathbf{R} = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$\Rightarrow R$ is a rotation by 180° in the plane spanned by v and w_1 .




Representing 3D Rotations (Axis-Angle)



What happens in the case that R is a rotation?

So, in both cases, the rotation R can be realized as a rotation by some angle θ about an axis v .

That is, R sends the vector v back into itself and rotates vectors in the plane that is perpendicular to v by the angle θ .



Representing 3D Rotations (Axis-Angle)



What happens in the case that R is a rotation?

This motivates a representation of rotations by specifying the axis about which the rotation occurs and the angle of the 2D rotation.



Outline

- Math Review
- Representing 3D Rotations
 - Axis-Angle
 - Euler Angles


Representing 3D Rotations (Euler)



We will consider a representation of rotations that describe what the rotation does to $(0,1,0)$.

Given a rotation R , if we know that R maps the North pole to the point p , is that enough information to define R ?


No. There can be many different rotations that all send $(0,1,0)$ to the point p .



Representing 3D Rotations (Euler)



In particular, if R_p is some (fixed) rotation taking the North pole to p and S is any rotation taking the North pole to p , then $R_p^{-1} \cdot S$ must map the North pole back to itself.





Representing 3D Rotations (Euler)

Denote by $R_y(\psi)$ the rotation about the y -axis (North pole) by ψ degrees:


$$\mathbf{R}_y(\psi) = \begin{pmatrix} \cos \psi & 0 & -\sin \psi \\ 0 & 1 & 0 \\ \sin \psi & 0 & \cos \psi \end{pmatrix}$$

Then:

$$R_p^{-1} \cdot S = R_y(\psi)$$



$$S = R_p \cdot R_y(\psi)$$



Representing 3D Rotations (Euler)




In order to represent all rotations, we need to find an expression for R_p -- some rotation that sends the North pole to the point p .

Let (θ, ϕ) be the spherical coordinates of p :

$$p = (\cos \theta \cdot \sin \phi, \cos \phi, \sin \theta \cdot \sin \phi)_y$$

The point p must lie on the circle about the y -axis with height $\cos \phi$.

We can get $(0,1,0)$ to this circle with a rotation by an angle of ϕ about the z -axis.



Representing 3D Rotations (Euler)




In order to represent all rotations, we need to find an expression for R_p -- some rotation that sends the North pole to the point p .

Let (θ, ϕ) be the spherical coordinates of p :

$$p = (\cos \theta \cdot \sin \phi, \cos \phi, \sin \theta \cdot \sin \phi)_y$$

We also know that the point p makes an angle of θ with the xz -plane.

We can get the rotation of $(0,1,0)$ to p by rotating by an angle of θ about the y -axis.




Representing 3D Rotations (Euler)



Thus, when the spherical coordinates of the point p are (θ, ϕ) , we can rotate $(0,1,0)$ to p by:

- First rotating by ϕ degrees about the z -axis, and
- Then rotating by θ degrees about the y -axis.



Representing 3D Rotations (Euler)



Since every rotation R can be described by a rotation about the y -axis, followed by a rotation that maps $(0,1,0)$ to $p = \Phi(\theta, \phi)$, we have:

$$R = R_y(\theta) \cdot R_z(\phi) \cdot R_y(\psi)$$

where $R_y(\alpha)$ is the rotation about the y -axis by α , and $R_z(\beta)$ is the rotation about the z -axis by β .

Representing 3D Rotations (Euler)



In matrix form, the triplet of angles (θ, ϕ, ψ) represents the rotation:

$$\mathbf{R}(\theta, \phi, \psi) = \underbrace{\begin{pmatrix} \cos \theta & 0 & -\sin \theta \\ 0 & 1 & 0 \\ \sin \theta & 0 & \cos \theta \end{pmatrix} \begin{pmatrix} \cos \phi & -\sin \phi & 0 \\ \sin \phi & \cos \phi & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{Rotation sending } (0,1,0) \rightarrow p = \Phi(\theta, \phi)} \underbrace{\begin{pmatrix} \cos \psi & 0 & -\sin \psi \\ 0 & 1 & 0 \\ \sin \psi & 0 & \cos \psi \end{pmatrix}}_{\text{Rotation about the } y\text{-axis by } \psi}$$

This is the Euler Angle parameterization of 3D rotations.