



FFTs in Graphics and Vision

Correlation of Spherical Functions



Assignment 2

- Only submit the .cpp files with the main function
- To declare an array, the size of the array needs to be a constant expression. The following code is not legal:

```
void function( int res )  
{  
    int values[res];  
    // do something  
}
```

Instead, do something like:

```
void function( int res )  
{  
    int *values = new int[res];  
    // do something  
    delete[] res;  
}
```

Outline

- Math Review
- Spherical Correlation





Review

Dimensionality:

Given a complex n -dimensional array $\mathbf{a} \in \mathbb{C}^n$ representing regular samples of a function on the circle, we can express the array in terms of its Fourier decomposition:

$$\mathbf{a} = \frac{1}{\sqrt{2\pi}} \sum_j \hat{\mathbf{a}}_j \cdot e^{-\frac{ik2\pi j}{n}}$$



Review

Dimensionality:

How many complex exponentials do we use?

Because the array is of dimension n , we need n Fourier coefficients to capture all the data:

$$\mathbf{a} = \frac{1}{\sqrt{2\pi}} \sum_j \hat{\mathbf{a}}_j \cdot e^{-\frac{ik2\pi j}{n}}$$

\Downarrow

$$\mathbf{a} = \frac{1}{\sqrt{2\pi}} \sum_{j=-n/2}^{n/2-1} \hat{\mathbf{a}}_j \cdot e^{-\frac{ik2\pi j}{n}}$$



Review

Dimensionality:

How many complex exponentials do we use?

Because the array is of dimension n , we need n Fourier coefficients to capture all the data:

The value of the largest frequency is often referred to as the *bandwidth* of the function.

$$\mathbf{a} = \frac{1}{\sqrt{2\pi}} \sum_{j=-n/2}^{n/2-1} \hat{\mathbf{a}}_j \cdot e^{-\frac{ik2\pi j}{n}}$$



Review


Dimensionality:

We represent a spherical function by an $n \times n$ grid whose entries are the regular samples of the function along the lines of latitude and longitude:

$$\mathbf{f}_{jk} = f(\cos \theta_j \cdot \sin \phi_k, \cos \phi_k, \sin \theta_j \cdot \sin \phi_k)$$

where θ_j and ϕ_k are the angles:

$$\theta_j = \frac{2\pi j}{n}$$
$$\phi_k = \frac{\pi(2k + 1)}{2n}$$





Review

Dimensionality:

We can express the spherical function as a sum of spherical harmonics:

$$\mathbf{f} = \sum_l \sum_{m=-l}^l \hat{\mathbf{f}}_{lm} \cdot Y_l^m$$



Review

Dimensionality:

How many frequencies should we use?

As in the case of functions on a circle, we use a bandwidth that is half the resolution:

$$\mathbf{f} = \sum_l \sum_{m=-l}^l \hat{\mathbf{f}}_{lm} \cdot Y_l^m$$
$$\Downarrow$$
$$\mathbf{f} = \sum_{l=0}^{n/2-1} \sum_{m=-l}^l \hat{\mathbf{f}}_{lm} \cdot Y_l^m$$



Review

Dimensionality:

$$\mathbf{f} = \sum_{l=0}^{n/2-1} \sum_{m=-l}^l \hat{\mathbf{f}}_{lm} \cdot Y_l^m$$

In this case, the number of coefficients is:


$$\sum_{l=0}^{n/2-1} (2l + 1) = \left(\frac{n}{2}\right)^2$$



Review

Dimensionality:

Since we go from n^2 spherical samples to $(n/2)^2$ spherical harmonic coefficients, there is a loss of information at the higher frequencies:





Review

Multi-dimensional FFT:

To compute a multi-dimensional FFT we leverage the fact that the d -dimensional basis vectors can be expressed as the product of 1D basis vectors:

$$\zeta^{l_1, \dots, l_d}(\theta_1, \dots, \theta_d) = \sqrt{\frac{1}{(2\pi)^d}} \cdot e^{il_1\theta_1} \dots e^{il_d\theta_d}$$



Review

Multi-dimensional FFT:

$$\zeta^{l_1, \dots, l_d}(\theta_1, \dots, \theta_d) = \sqrt{\frac{1}{(2\pi)^d}} \cdot e^{il_1\theta_1} \dots e^{il_d\theta_d}$$

Then, the (l_1, \dots, l_d) -th Fourier coefficient of a function $f(\theta_1, \dots, \theta_d)$ is:

$$\begin{aligned} \hat{f}_{l_1, \dots, l_d} &= \int \dots \int f(\theta_1, \dots, \theta_d) \cdot \overline{\zeta^{l_1, \dots, l_d}(\theta_1, \dots, \theta_d)} d\theta_d \dots d\theta_1 \\ &= \sqrt{\frac{1}{(2\pi)^d}} \int \dots \left(\int f(\theta_1, \dots, \theta_d) \cdot e^{-il_d\theta_d} d\theta_d \right) \dots e^{-il_1\theta_1} d\theta_1 \end{aligned}$$



Review

Multi-dimensional FFT:

$$\hat{\mathbf{f}}_{l_1, \dots, l_d} = \sqrt{\frac{1}{(2\pi)^d}} \int \dots \left(\int f(\theta_1, \dots, \theta_d) \cdot e^{-il_d \theta_d} d\theta_d \right) \dots e^{-il_1 \theta_1} d\theta_1$$

This can be computed by performing a successive set of 1D Fourier transforms.

Note that the only facts we used were that the basis functions are separable and orthonormal.

Outline

- Math Review
- Spherical Correlation






Goal

Given real-valued functions on the sphere f and g , find the rotation R that optimally aligns f to g :


$$R = \arg \min_{R \in SO(3)} \|\rho_R(f) - g\|^2$$





Recall

Given real-valued functions on the circle f and g , we would like to find the rotation R that optimally aligns f to g .






Reduction to a Moving Dot-Product

Expressing the norm in terms of the dot-product, we get:

$$\begin{aligned}\|\rho_R(f) - g\|^2 &= \langle \rho_R(f) - g, \rho_R(f) - g \rangle \\ &= \langle \rho_R(f), \rho_R(f) \rangle + \langle g, g \rangle - \langle g, \rho_R(f) \rangle - \langle \rho_R(f), g \rangle \\ &= \|f\|^2 + \|g\|^2 - 2\langle \rho_R(f), g \rangle\end{aligned}$$






Reduction to a Moving Dot-Product

Expressing the norm in terms of the dot-product, we get:

$$\|\rho_R(f) - g\|^2 = \|f\|^2 + \|g\|^2 - 2\langle \rho_R(f), g \rangle$$

⇒ Finding the rotation minimizing the norm is equivalent to finding the rotation maximizing the dot-product.






General Approach

If we define the function $D_{f,g}(\alpha)$ giving the dot-product of the rotation of f by angle α with g :

$$D_{f,g}(\alpha) = \langle \rho_\alpha(f), g \rangle$$

we can find the aligning rotation by finding the value of α maximizing $D_{f,g}(\alpha)$.





Brute-Force

To compute $D_{f,g}(\alpha)$, we could explicitly compute the value at each angle of rotation α .

If we represent a function on a circle by the values at n regular samples, this would give an algorithm whose complexity is $O(n^2)$



Fourier Transform

We can do better by using the Fourier transform:

- We can leverage the irreducible representations to minimize the number of multiplications that need to be performed.
- We can use the FFT to compute the Inverse Fourier Transform efficiently.



Irreducible Representations

Given the functions f and g on the circle, we can express the functions in terms of their Fourier decomposition:

$$f(\theta) = \frac{1}{\sqrt{2\pi}} \sum_k \hat{\mathbf{f}}_k e^{ik\theta}$$
$$g(\theta) = \frac{1}{\sqrt{2\pi}} \sum_k \hat{\mathbf{g}}_k e^{ik\theta}$$



Irreducible Representations

In terms of this decomposition, the expression for the dot-product becomes:

$$\begin{aligned} D_{f,g}(\alpha) &= \left\langle \rho_\alpha \left(\frac{1}{\sqrt{2\pi}} \sum_k \hat{\mathbf{f}}_k e^{ik\theta} \right), \frac{1}{\sqrt{2\pi}} \sum_{k'} \hat{\mathbf{g}}_{k'} e^{ik'\theta} \right\rangle \\ &= \left\langle \frac{1}{\sqrt{2\pi}} \sum_k \hat{\mathbf{f}}_k \cdot \rho_\alpha(e^{ik\theta}), \frac{1}{\sqrt{2\pi}} \sum_{k'} \hat{\mathbf{g}}_{k'} e^{ik'\theta} \right\rangle \\ &= \frac{1}{2\pi} \sum_{k,k'} \hat{\mathbf{f}}_k \overline{\hat{\mathbf{g}}_{k'}} \cdot \left\langle \rho_\alpha(e^{ik\theta}), e^{ik'\theta} \right\rangle \end{aligned}$$



Irreducible Representations

Let $D_{k,k'}(\alpha)$ be the function giving the dot-product of the rotation of the k -th complex exponential by an angle of α with the k' -th complex exponential:

$$D_{k,k'}(\alpha) = \left\langle \rho_\alpha(e^{ik\theta}), e^{ik'\theta} \right\rangle$$

Then the equation for the dot-product becomes:

$$D_{f,g}(\alpha) = \frac{1}{2\pi} \sum_{k,k'} \hat{\mathbf{f}}_k \overline{\hat{\mathbf{g}}_{k'}} \cdot \left\langle \rho_\alpha(e^{ik\theta}), e^{ik'\theta} \right\rangle$$

\Downarrow

$$D_{f,g}(\alpha) = \frac{1}{2\pi} \sum_{k,k'} \hat{\mathbf{f}}_k \overline{\hat{\mathbf{g}}_{k'}} \cdot D_{k,k'}(\alpha)$$



Irreducible Representations

$$D_{f,g}(\alpha) = \frac{1}{2\pi} \sum_{k,k'} \hat{\mathbf{f}}_k \overline{\hat{\mathbf{g}}_{k'}} \cdot D_{k,k'}(\alpha)$$

Up to this point, the algorithm looks like:

- Compute the Fourier coefficients of f and g .
- Cross-multiply the Fourier coefficients to get the coefficients of the correlation in terms of the functions $D_{k,k'}(\alpha)$

This doesn't seem particularly promising since it in the second step, we need to perform $O(n^2)$ multiplies – which is no better than brute force.



Irreducible Representations

But, using the Fourier decomposition, we know that the space of functions on a circle of frequency k are:

- Fixed by rotation (i.e. a sub-representation)
- Perpendicular to the space of functions of frequency k' (for $k \neq k'$)

Thus, for $k \neq k'$, we know that:

$$D_{k,k'}(\alpha) = \left\langle \rho_\alpha(e^{ik\theta}), e^{ik'\theta} \right\rangle = 0$$




Irreducible Representations

So the expression for the correlation becomes:

$$D_{f,g}(\alpha) = \frac{1}{2\pi} \sum_k \hat{\mathbf{f}}_k \overline{\hat{\mathbf{g}}_k} \cdot D_k(\alpha)$$

Reducing the number of cross-multiplications that need to be performed from $O(n^2)$ to $O(n)$:





Change of Basis

At this point, we have an expression for the correlation as a linear sum of the function $D_k(\alpha)$:

$$D_{f,g}(\alpha) = \frac{1}{2\pi} \sum_k \hat{\mathbf{f}}_k \overline{\hat{\mathbf{g}}_k} \cdot D_k(\alpha)$$

To evaluate the correlation at α we need to get the value of each of the $D_k(\alpha)$, and take the linear combination, using weights $\frac{1}{2\pi} \hat{\mathbf{f}}_k \overline{\hat{\mathbf{g}}_k}$.

That, is evaluating the correlation at any single angle requires $O(n)$ computations and evaluating at all angles would take $O(n^2)$.



Change of Basis

Set $\mathbf{c} \in \mathbb{C}^n$ to be the n -dimensional array:

$$\mathbf{c}_k = \frac{1}{2\pi} \hat{\mathbf{f}}_k \overline{\hat{\mathbf{g}}_k}$$

and set $\mathbf{a} \in \mathbb{C}^n$ to be the n -dimensional array:

$$\mathbf{a}_k = D_{f,g} \left(\frac{2k\pi}{n} \right)$$

Then we get:

$$\mathbf{a} = \begin{pmatrix} D_0(0) & \cdots & D_{n-1}(0) \\ \vdots & \ddots & \vdots \\ D_0\left(\frac{2(n-1)\pi}{n}\right) & \cdots & D_{n-1}\left(\frac{2(n-1)\pi}{n}\right) \end{pmatrix} \cdot \mathbf{c}$$



Change of Basis

Set $\mathbf{c} \in \mathbb{C}^n$ to be the n -dimensional array:

$$\mathbf{c}_k = \frac{1}{2\pi} \hat{\mathbf{f}}_k \overline{\hat{\mathbf{g}}_k}$$

and set $\mathbf{a} \in \mathbb{C}^n$ to be the n -dimensional array:

$$\mathbf{a}_k = D_c \left(\frac{2k\pi}{n} \right)$$

Th To get the desired expression for the correlation, we need to do a matrix vector multiply!

$$\mathbf{a} = \begin{pmatrix} \vdots & \ddots & \vdots \\ D_0 \left(\frac{2(n-1)\pi}{n} \right) & \cdots & D_{n-1} \left(\frac{2(n-1)\pi}{n} \right) \end{pmatrix} \cdot \mathbf{c}$$



Change of Basis

Computing this change of basis amounts to computing the Inverse Fourier Transform.

$$D_k(\alpha) = \langle e^{ik(\theta-\alpha)}, e^{ik\theta} \rangle = e^{-ik\alpha}$$



Algorithm for Circular Functions

In sum, we get an algorithm for computing the value of the correlation of f with g :

1. Compute the Fourier coefficients of f and g :

$$O(n \log n)$$

2. Cross-multiply the Fourier coefficients:

$$O(n)$$

3. Compute the inverse Fourier transform:


$$O(n \log n)$$



Goal

Given real-valued functions on the sphere f and g , find the rotation R that optimally aligns f to g :

$$R = \arg \min_{R \in SO(3)} \|\rho_R(f) - g\|^2$$





Expanding the Norm

Given real-valued functions on the sphere f and g , find the rotation R that optimally aligns f to g :

$$R = \arg \min_{R \in SO(3)} \|\rho_R(f) - g\|^2$$

Expanding the norm, we get:

$$\|\rho_R(f) - g\|^2 = \|f\|^2 + \|g\|^2 - 2\langle \rho_R(f), g \rangle$$



Expanding the Norm

$$\|\rho_R(f) - g\|^2 = \|f\|^2 + \|g\|^2 - 2\langle \rho_R(f), g \rangle$$

Thus, to find the rotation minimizing the norm of the difference, we need to find the rotation maximizing the dot-product:

$$D_{f,g}(R) = \langle \rho_R(f), g \rangle$$



Brute-Force

Again, we can try to compute the value of the dot-product using a brute force algorithm:

For each rotation R , we could compute the dot-product of the rotated function $\rho_R(f)$ with g .

If n is the resolution of the spherical function, the “size” of a spherical function is $O(n^2)$ and the “size” of the space of rotations is $O(n^3)$.

This means that a brute force algorithm would take on the order of $O(n^5)$ time.



Approach

As in the case of functions on a circle, we will take a two step approach:

1. We will use the irreducible representations to minimize the number of cross multiplications.
2. We will compute an efficient change of basis.



Irreducible Representations

Expanding the functions f and g in terms of their spherical harmonic decompositions, we get:

$$f(\theta, \phi) = \sum_{l=0}^b \sum_{m=-l}^l \hat{\mathbf{f}}_{lm} Y_l^m(\theta, \phi)$$
$$g(\theta, \phi) = \sum_{l=0}^b \sum_{m=-l}^l \hat{\mathbf{g}}_{lm} Y_l^m(\theta, \phi)$$



Irreducible Representations

Expanding the dot-product in terms of the spherical harmonics, we get:

$$D_{f,g}(R) = \left\langle \rho_R \left(\sum_{l=0}^b \sum_{m=-l}^l \hat{\mathbf{f}}_{lm} Y_l^m(\theta, \phi) \right), \sum_{l'=0}^b \sum_{m'=-l'}^{l'} \hat{\mathbf{g}}_{l'm'} Y_{l'}^{m'}(\theta, \phi) \right\rangle$$

Using the linearity of ρ_R , we can pull the linear summation outside of the rotation:

$$D_{f,g}(R) = \left\langle \sum_{l=0}^b \sum_{m=-l}^l \hat{\mathbf{f}}_{lm} \rho_R(Y_l^m(\theta, \phi)), \sum_{l'=0}^b \sum_{m'=-l'}^{l'} \hat{\mathbf{g}}_{l'm'} Y_{l'}^{m'}(\theta, \phi) \right\rangle$$



Irreducible Representations

Expanding the dot-product in terms of the spherical harmonics, we get:

$$D_{f,g}(R) = \left\langle \sum_{l=0}^b \sum_{m=-l}^l \hat{\mathbf{f}}_{lm} \rho_R(Y_l^m(\theta, \phi)), \sum_{l'=0}^b \sum_{m'=-l'}^{l'} \hat{\mathbf{g}}_{l'm'} Y_{l'}^{m'}(\theta, \phi) \right\rangle$$

Using the conjugate-linearity of the inner product, we can pull out the linear summation:

$$D_{f,g}(R) = \sum_{l,l'=0}^b \sum_{m=-l}^l \sum_{m'=-l'}^{l'} \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{l'm'}} \cdot \left\langle \rho_R(Y_l^m(\theta, \phi)), Y_{l'}^{m'}(\theta, \phi) \right\rangle$$



Irreducible Representations

$$D_{f,g}(R) = \sum_{l,l'=0}^b \sum_{m=-l}^l \sum_{m'=-l'}^{l'} \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{l'm'}} \cdot \left\langle \rho_R(Y_l^m(\theta, \phi)), Y_{l'}^{m'}(\theta, \phi) \right\rangle$$

Recall that:

1. Rotations of l -th frequency functions are l -th frequency functions
2. The space of l -th frequency functions is orthogonal to the space of l' -th frequency functions (for $l \neq l'$)

We get:

$$D_{f,g}(R) = \sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{lm'}} \cdot \left\langle \rho_R(Y_l^m(\theta, \phi)), Y_l^{m'}(\theta, \phi) \right\rangle$$



Change of Basis

$$D_{f,g}(R) = \sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{lm'}} \cdot \left\langle \rho_R(Y_l^m(\theta, \phi)), Y_l^{m'}(\theta, \phi) \right\rangle$$

Set $D_l^{m,m'}$ to be the functions on the space of rotations defined by:

$$D_l^{m,m'}(R) = \left\langle \rho_R(Y_l^m), Y_l^{m'} \right\rangle$$

This gives:

$$D_{f,g}(R) = \sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{lm'}} \cdot D_l^{m,m'}(R)$$

The $D_l^{m,m'}$ are called *Wigner-D* functions.



Change of Basis

$$D_{f,g}(R) = \sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{lm'}} \cdot D_l^{m,m'}(R)$$

Thus, given the spherical harmonic coefficients of f and g , we can express the correlation as a sum of the functions $D_l^{m,m'}$ by cross-multiplying the harmonic coefficients within each frequency.



Change of Basis

$$D_{f,g}(R) = \sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{lm'}} \cdot D_l^{m,m'}(R)$$

The problem is that this expression for the correlation is not easy to evaluate.

To compute the value at a particular rotation R , we need to:

- Evaluate $D_l^{m,m'}(R)$ at every frequency l and every pair of indices $-l \leq m, m' \leq l$,
- And then take the linear sum weighted by the product of the harmonic coefficients



Change of Basis

$$D_{f,g}(R) = \sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{lm'}} \cdot D_l^{m,m'}(R)$$

That is, for each of $O(n^3)$ rotations, we would need to evaluate:

$$\sum_{l=0}^{O(n)} (2l+1)^2 = O(n^3)$$

different functions.

This is worse than brute force method since it requires $O(n^6)$ while brute force requires $O(n^5)$.



Change of Basis

What is that we really want to do?

We would like to take a function expressed as a linear sum of the $D_l^{m,m'}$ and get an expression of the function, “regularly” sampled at n^3 rotations.

As in the case of circular correlation, this amounts to a change of basis. Only in the spherical case:

- The vectors themselves are of dimension n^3
- So the matrices are of $n^3 \times n^3 = n^6$.



Change of Basis

If we represent rotations in terms of the triplet of Euler angles $(\theta, \phi, \psi) \in [0, 2\pi) \times [0, \pi] \times [0, 2\pi)$:

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

Rotation sending
 $(0,1,0) \rightarrow p = \Phi(\theta, \phi)$

Rotation about
the y -axis by ψ

what do the function $D_l^{m,m'}(R(\theta, \phi, \psi))$ look like?



Change of Basis

Recall that the spherical harmonics can be expressed as a complex exponential in θ times a “polynomial” in $\cos \phi$:

$$Y_l^m(\theta, \phi) = P_l^m(\cos \phi) \cdot e^{im\theta}$$

So a rotation by an angle of α about the y -axis acts on the (l, m) -th spherical harmonics by:

$$\rho_{R_y(\alpha)}(Y_l^m) = e^{-im\alpha} \cdot Y_l^m$$



Change of Basis

Thus, writing out the functions $D_l^{m,m'}$ as functions of the Euler angles, we get:

$$\begin{aligned} D_l^{m,m'}(\theta, \phi, \psi) &= \left\langle \left(\rho_{R_y(\theta)} \circ \rho_{R_z(\phi)} \circ \rho_{R_y(\psi)} \right) (Y_l^m), Y_l^{m'} \right\rangle \\ &= \left\langle \rho_{R_z(\phi)} \left(\rho_{R_y(\psi)} (Y_l^m) \right), \rho_{R_y^{-1}(\theta)} (Y_l^{m'}) \right\rangle \\ &= \left\langle \rho_{R_z(\phi)} (e^{-im\psi} \cdot Y_l^m), e^{im'\theta} \cdot (Y_l^{m'}) \right\rangle \\ &= e^{-im\psi} \cdot e^{-im'\theta} \left\langle \rho_{R_z(\phi)} (Y_l^m), Y_l^{m'} \right\rangle \end{aligned}$$



Change of Basis

$$D_l^{m,m'}(\theta, \phi, \psi) = e^{-im\psi} \cdot e^{-im'\theta} \left\langle \rho_{R_Z(\phi)}(Y_l^m), Y_l^{m'} \right\rangle$$

Denoting:

$$d_l^{m,m'}(\phi) = \left\langle \rho_{R_Z(\phi)}(Y_l^m), Y_l^{m'} \right\rangle$$

we can express the functions $D_l^{m,m'}$ as:

$$D_l^{m,m'}(\theta, \phi, \psi) = e^{-im'\theta} \cdot d_l^{m,m'}(\phi) \cdot e^{-im\psi}$$

$$d_l^{m,m'}(\phi) = \sum_k (-1)^k \frac{\sqrt{(l+m)! \cdot (l-m)! \cdot (l+m')! \cdot (l-m')!}}{(l-m'-k)! \cdot (l+m-k)! \cdot k! \cdot (k+m'-m)!} \cos^{2l+m-m'-2k} \left(\frac{\phi}{2} \right) \cdot \sin^{2k+m'-m} \left(\frac{\phi}{2} \right)$$

The $d_l^{m,m'}$ are sometimes called
Wigner-small-d functions.



Change of Basis

$$D_l^{m,m'}(\theta, \phi, \psi) = e^{-im'\theta} \cdot d_l^{m,m'}(\phi) \cdot e^{-im\psi}$$

The advantage of this representation is that the basis functions are separable so instead of doing a single $O(n^3 \times n^3)$ we need to do $3n^2$ different $O(n \times n)$ matrix multiplies.

We can make things even faster noting that in two of the dimensions we are performing a Fourier transform.




Change of Basis

$$D_l^{m,m'}(\theta, \phi, \psi) = e^{-im'\theta} \cdot d_l^{m,m'}(\phi) \cdot e^{-im\psi}$$

We can think of the sampled correlation function as an $n \times n \times n$ grid, whose (p, q, r) -th entry corresponds to the value of the correlation at the Euler angle $(\theta_p, \phi_q, \psi_r)$

$$\begin{aligned}\theta_p &= \frac{2\pi p}{n} \\ \phi_q &= \frac{\pi(2q+1)}{2n} \\ \psi_r &= \frac{2\pi r}{n}\end{aligned}$$






Change of Basis

$$D_l^{m,m'}(\theta, \phi, \psi) = e^{-im'\theta} \cdot d_l^{m,m'}(\phi) \cdot e^{-im\psi}$$

Fixing the angle ϕ , we restrict ourselves to a 2D slice of the correlation values.





Change of Basis

$$D_l^{m,m'}(\theta, \phi, \psi) = e^{-im'\theta} \cdot d_l^{m,m'}(\phi) \cdot e^{-im\psi}$$

Fixing the angle ϕ , we restrict ourselves to a 2D slice of the correlation values.

On this 2D slice, the values of the correlation are:

$$\begin{aligned} D_{f,g}^{\phi}(\theta, \psi) &= \sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{lm'}} \cdot D_l^{m',m}(\theta, \phi, \psi) \\ &= \sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{lm'}} \cdot e^{-im'\theta} \cdot d_l^{m,m'}(\phi) \cdot e^{-im\psi} \end{aligned}$$



Change of Basis

$$D_{f,g}^{\phi}(\theta, \psi) = \sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{lm'}} \cdot e^{-im'\theta} \cdot d_l^{m,m'}(\phi) \cdot e^{-im\psi}$$

That is, for fixed ϕ we get a 2D function which is the sum of complex exponentials, with (m, m') Fourier coefficient defined by:

$$\left(\hat{\mathbf{D}}_{f,g}^{\phi}\right)_{mm'} = \sum_{l=\max(|m|, |m'|)}^b \hat{\mathbf{f}}_{l,-m} \cdot \overline{\hat{\mathbf{g}}_{l,-m'}} \cdot d_l^{-m,-m'}(\phi)$$

⇒ We can get the values in this 2D slice by running the 2D inverse FFT.



Change of Basis

$$D_{f,g}^{\phi}(\theta, \psi) = \sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{f}}_{lm} \cdot \overline{\hat{\mathbf{g}}_{lm'}} \cdot e^{-im'\theta} \cdot d_l^{m,m'}(\phi) \cdot e^{-im\psi}$$

This allows us evaluate the correlation on a slice by slice basis.

For every sampled value of ϕ :

- We compute the Fourier coefficients:

$$\left(\widehat{\mathbf{D}}_{f,g}^{\phi}\right)_{mm'} = \sum_{l=\max(|m|,|m'|)}^b \hat{\mathbf{f}}_{l,m} \cdot \overline{\hat{\mathbf{g}}_{l,-m'}} \cdot d_l^{-m,-m'}(\phi)$$

- And then we compute the 2D inverse FFT.



Change of Basis

Since we sample ϕ at n different values, there are n different 2D slices.

For each slice there are $n \times n$ different Fourier coefficients, requiring $O(n)$ calculations per coefficient:

$$\left(\widehat{\mathbf{D}}_{f,g}^{\phi}\right)_{mm'} = \sum_{l=\max(|m|,|m'|)}^b \hat{\mathbf{f}}_{l,m} \cdot \overline{\hat{\mathbf{g}}_{l,-m'}} \cdot d_l^{-m,-m'}(\phi)$$



Change of Basis

Since we sample ϕ at n different values, there are n different 2D slices.

For each slice there are $n \times n$ different Fourier coefficients, requiring $O(n)$ calculations per coefficient.

And each inverse FFT takes $O(n^2 \log n)$ time.

Thus, the computational complexity becomes:

- $O(n^4)$ for computing all the 2D slice Fourier coefficients
- $O(n^3 \log n)$ to compute all the 2D inverse FFTs.



Change of Basis

Since we sample ϕ at n different values, there are n different 2D slices.

For each slice there are $n \times n$ different Fourier coefficients, requiring $O(n)$ calculations per

co
Ar In particular, we can do much better than
the brute force algorithm

Thus, the computational complexity becomes:

- $O(n^4)$ for computing all the 2D slice Fourier coefficients
- $O(n^3 \log n)$ to compute all the 2D inverse FFTs.



General Overview

To make the computation of the correlation efficient, we used the fact that in two of the three coefficients – θ and ψ – the functions $D_l^{m,m'}$ could be expressed as complex exponentials.

This allowed us to replace the $n^2 \times n^2$ matrix multiplication in two of the variables by an $O(n^2 \log n)$ inverse FFT.

In the third variable – ϕ – we still end up doing a full $n \times n$ matrix multiplication:

$$n^3 \times n^3 \rightarrow \underbrace{n^2 \cdot (n \times n)}_{\phi} + \underbrace{(n^2 \log n) \cdot n}_{\theta, \psi}$$



General Overview

In practice, the change of basis in ϕ can also be performed using an FFT like approach, giving rise to an algorithm with complexity $O(n \log^2 n)$.

Thus, the total complexity of computing the correlation drops down to $O(n^3 \log^2 n)$.

Aligning 3D Functions



What kind of penalty hit do we pay for aligning functions defined in 3D?



Correlating 3D Functions

Given two functions F and G defined on the unit ball (i.e. (x, y, z) with $\|(x, y, z)\| \leq 1$) we would like to compute the distance between the functions at every rotation:

$$\|\rho_R(F) - G\|^2 = \|F\|^2 + \|G\|^2 - 2\langle \rho_R(F), G \rangle$$




Correlating 3D Functions

Using the fact that rotations fix spheres about the origin, we express the functions as a set of spherical functions:

$$F^r(\theta, \phi) = F(r \cdot \cos \theta \cdot \sin \phi, r \cdot \cos \phi, r \cdot \sin \theta \cdot \sin \phi)$$

$$G^r(\theta, \phi) = G(r \cdot \cos \theta \cdot \sin \phi, r \cdot \cos \phi, r \cdot \sin \theta \cdot \sin \phi)$$





Correlating 3D Functions

The value of the correlation then becomes:

$$\langle \rho_R(F), G \rangle = \int_0^1 \langle \rho_R(F^r), G^r \rangle \cdot r^2 dr$$

Thus, if we express each radial restriction in terms of its spherical harmonics:

$$F^r = \sum_{l=0}^b \sum_{m=-l}^l \hat{\mathbf{F}}_{lm}(r) \cdot Y_l^m \quad G^r = \sum_{l=0}^b \sum_{m=-l}^l \hat{\mathbf{G}}_{lm}(r) \cdot Y_l^m$$

we get:

$$\langle \rho_R(F), G \rangle = \int_0^1 \left(\sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{F}}_{lm}(r) \cdot \overline{\hat{\mathbf{G}}_{lm'}(r)} \cdot D_l^{m,m'}(R) \right) r^2 dr$$



Correlating 3D Functions

$$\langle \rho_R(F), G \rangle = \int_0^1 \left(\sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{F}}_{lm}(r) \cdot \overline{\hat{\mathbf{G}}_{lm'}(r)} \cdot D_l^{m,m'}(R) \right) r^2 dr$$

This implies that we can compute the correlation, by performing a correlation for each radial restriction and then take the (area weighted) sum.

Assuming that we sample the radius at $O(n)$ different values, this would give an algorithm with complexity $O(n^5) / O(n^4 \log^2 n)$.



Correlating 3D Functions

We can do better.

The functions $D_l^{m,m'}$ do not depend on the radius, so we can pull them out of the integral:

$$\langle \rho_R(F), G \rangle = \int_0^1 \left(\sum_{l=0}^b \sum_{m,m'=-l}^l \hat{\mathbf{F}}_{lm}(r) \cdot \overline{\hat{\mathbf{G}}_{lm'}(r)} \cdot D_l^{m,m'}(R) \right) r^2 dr$$

\Downarrow

$$\langle \rho_R(F), G \rangle = \sum_{l=0}^b \sum_{m,m'=-l}^l \left(\int_0^1 \hat{\mathbf{F}}_{lm}(r) \cdot \overline{\hat{\mathbf{G}}_{lm'}(r)} \cdot r^2 dr \right) \cdot D_l^{m,m'}(R)$$



Correlating 3D Functions

$$\langle \rho_R(F), G \rangle = \sum_{l=0}^b \sum_{m,m'=-l}^l \left(\int_0^1 \hat{\mathbf{F}}_{lm}(r) \cdot \overline{\hat{\mathbf{G}}_{lm'}(r)} \cdot r^2 dr \right) \cdot D_l^{m,m'}(R)$$

The advantage of this expression, is that by gathering values across different radii first, we only need to perform a single change of basis.



Correlating 3D Functions

$$\langle \rho_R(F), G \rangle = \sum_{l=0}^b \sum_{m, m'=-l}^l \left(\int_0^1 \boxed{\hat{\mathbf{F}}_{lm}(r)} \cdot \boxed{\hat{\mathbf{G}}_{lm'}(r)} \cdot r^2 dr \right) \cdot D_l^{m, m'}(R)$$

Algorithm: (Assuming $O(n)$ radial samples)

1. Compute the spherical harmonic transform of each radial restriction:
 $O(n) \cdot O(n^2 \log^2 n)$



Correlating 3D Functions

$$\langle \rho_R(F), G \rangle = \sum_{l=0}^b \sum_{m, m'=-l}^l \boxed{\left(\int_0^1 \hat{\mathbf{F}}_{lm}(r) \cdot \overline{\hat{\mathbf{G}}_{lm'}(r)} \cdot r^2 dr \right)} \cdot D_l^{m, m'}(R)$$

Algorithm: (Assuming $O(n)$ radial samples)

1. Compute the spherical harmonic transform of each radial restriction:
 $O(n) \cdot O(n^2 \log^2 n)$
2. Cross multiply intra-frequency harmonic coeffs. and sum over the radii:
 $O(n) \cdot O(n^3)$



Correlating 3D Functions

$$\langle \rho_R(F), G \rangle = \sum_{l=0}^b \sum_{m,m'=-l}^l \left(\int_0^1 \hat{\mathbf{F}}_{lm}(r) \cdot \overline{\hat{\mathbf{G}}_{lm'}(r)} \cdot r^2 dr \right) \cdot D_l^{m,m'}(R)$$

Algorithm: (Assuming $O(n)$ radial samples)

1. Compute the spherical harmonic transform of each radial restriction: $O(n) \cdot O(n^2 \log^2 n)$
2. Cross multiply intra-frequency harmonic coeffs. and sum over the radii: $O(n) \cdot O(n^3)$
3. Do the change of basis: $O(n^3 \log^2 n) / O(n^4)$

\Rightarrow In 3D, correlations can be done in $O(n^4)$ time.