FFTs in Graphics and Vision

Correlation of Spherical Functions

Outline

- Math Review
- Spherical Correlation

Review

Dimensionality:

Given an n-dimensional array $a[\,]$ representing regular samples of a function on the circle, we can express the array in terms of its Fourier decomposition:

$$a[\,] = \sum_k \hat{a}[k]\, e_k[\,]$$

Where the $e_k[\,]$ are regular samples of the (normalized) complex exponentials.

Review

Dimensionality:

How many complex exponentials do we use?

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:

$$a[\,] = \sum_k \hat{a}[k] e_k[\,]$$

↓

$$a[\,] = \sum_{k=-n/2}^{n/2} \hat{a}[k] e_k[\,]$$

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.

$$a[\,] = \sum_{k=-n/2}^{n/2} \hat{a}[k] e_k[\,]$$

Review

<u>Dimensionality</u>:

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

$$f[j][k] = f\left(\cos\theta_j \sin\phi_k, \cos\phi_k, \sin\theta_j \sin\phi_k\right)$$

Where θ_j and ϕ_k are the angles:

$$\theta_j = \frac{2\pi j}{n}$$

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

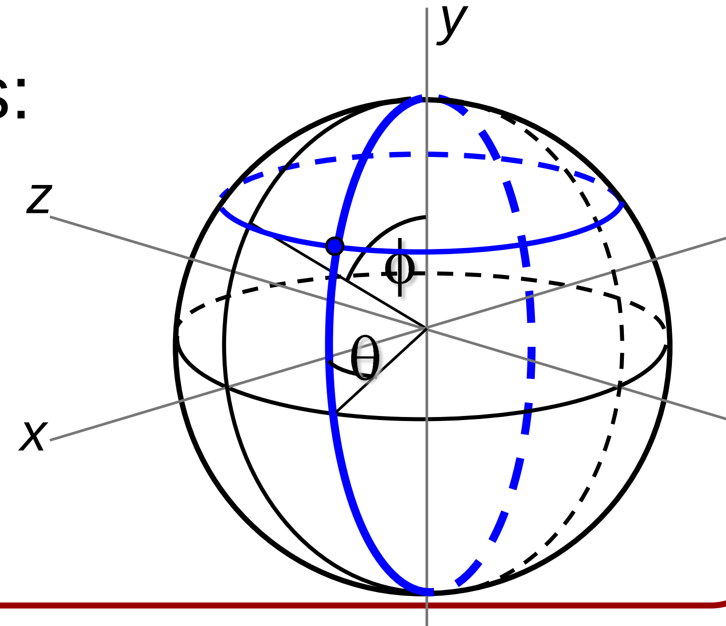

Review

<u>Dimensionality</u>:

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

$$f[\ ][\ ] = \sum_{l} \sum_{m=-l}^{l} \hat{f}[l][m] Y_l^m[\ ][\ ]$$

Review

Dimensionality:

How many frequencies should we use?

Review

<u>Dimensionality</u>:

How many frequencies should we use?

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

$$f[\;][\;] = \sum_{l} \sum_{m=-l}^{l} \hat{f}[l][m] Y_l^m[\;][\;]$$

$$f[\;][\;] = \sum_{l=0}^{n/2-1} \sum_{m=-l}^{l} \hat{f}[l][m] Y_l^m[\;][\;]$$

Review

Dimensionality:

$$f[\;][\;]=\sum_{l=0}^{n/2-1}\sum_{m=-l}^{l}\hat{f}[l][m]Y_l^m[\;][\;]$$

In this case, the number of coefficients is only:

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

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:

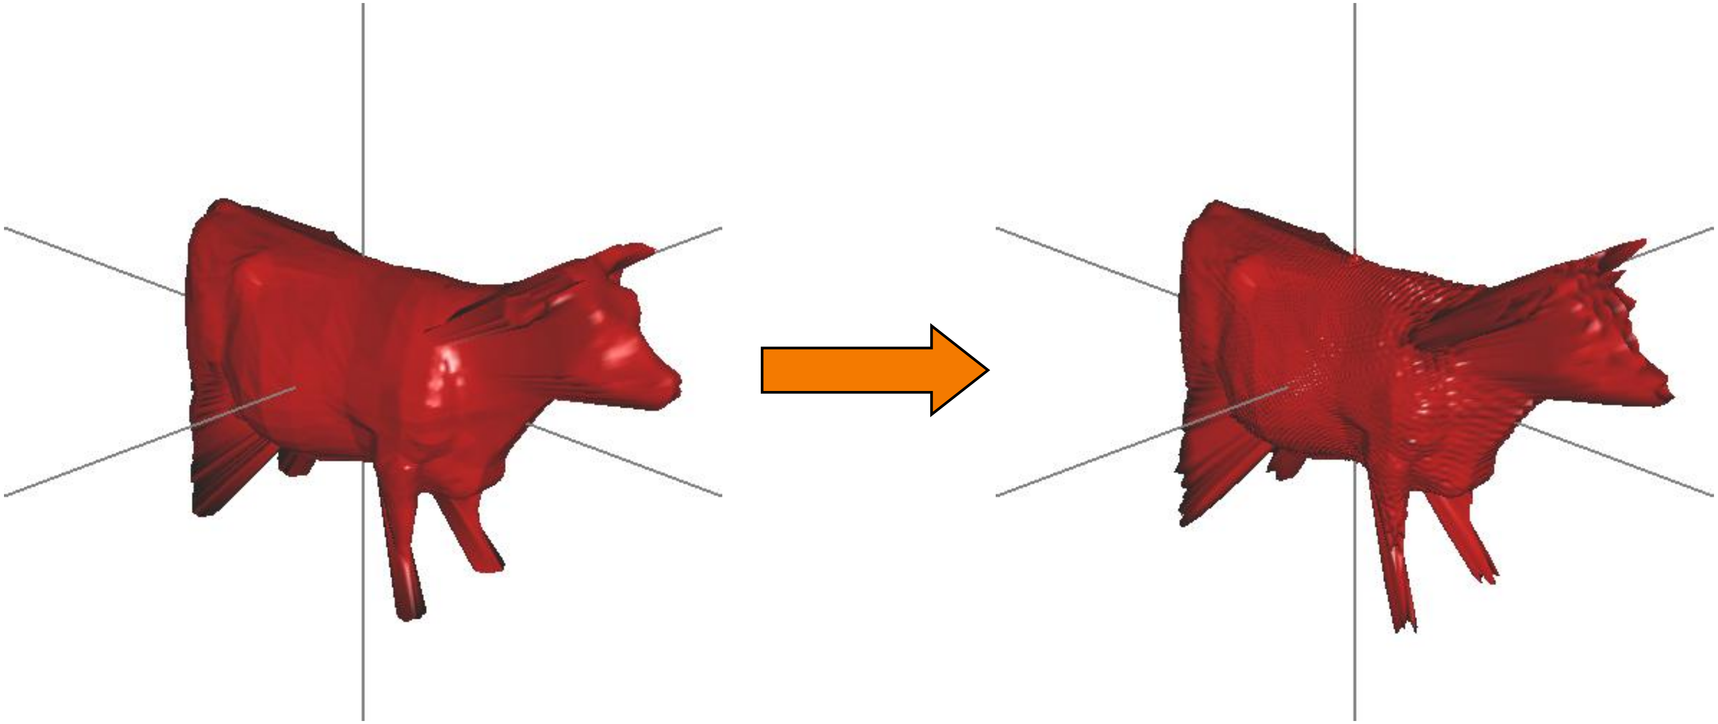

Outline

- Math Review
- Spherical Correlation

Goal

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

$$R = \underset{R \in \text{Rotations}}{\arg\min} \left\| f - \rho_R g \right\|^2$$

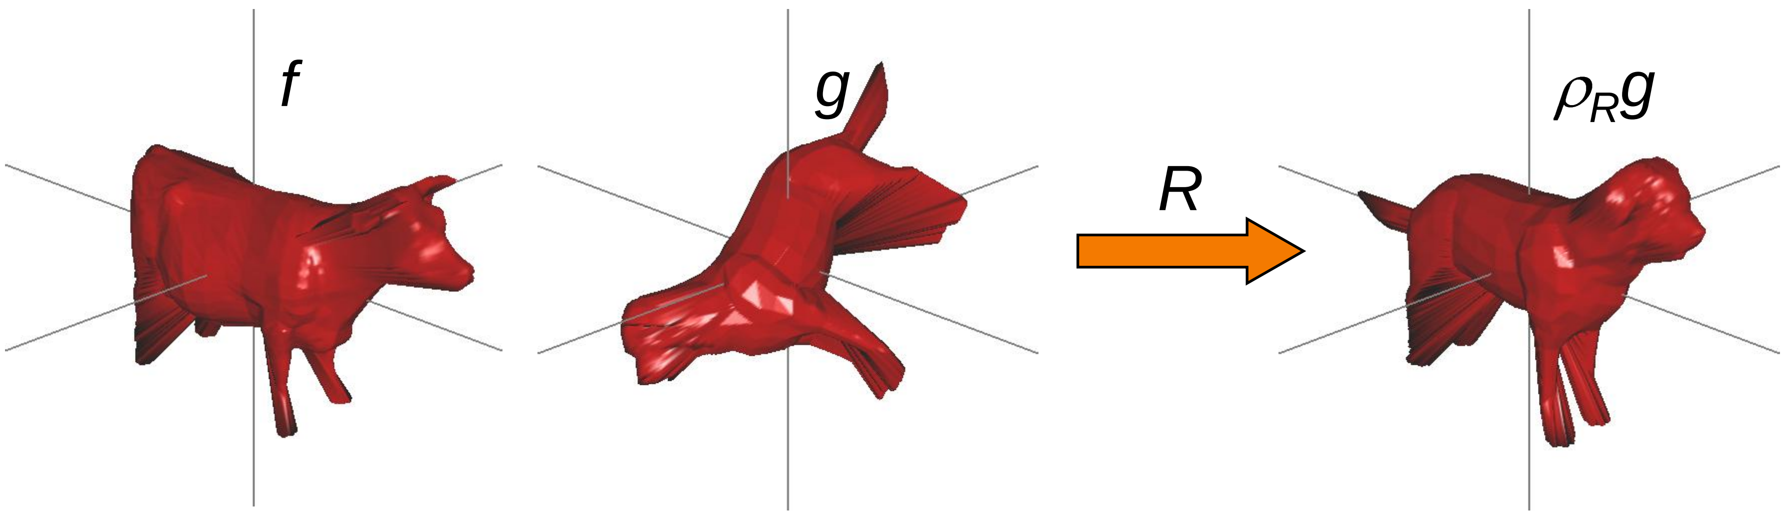

Recall

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

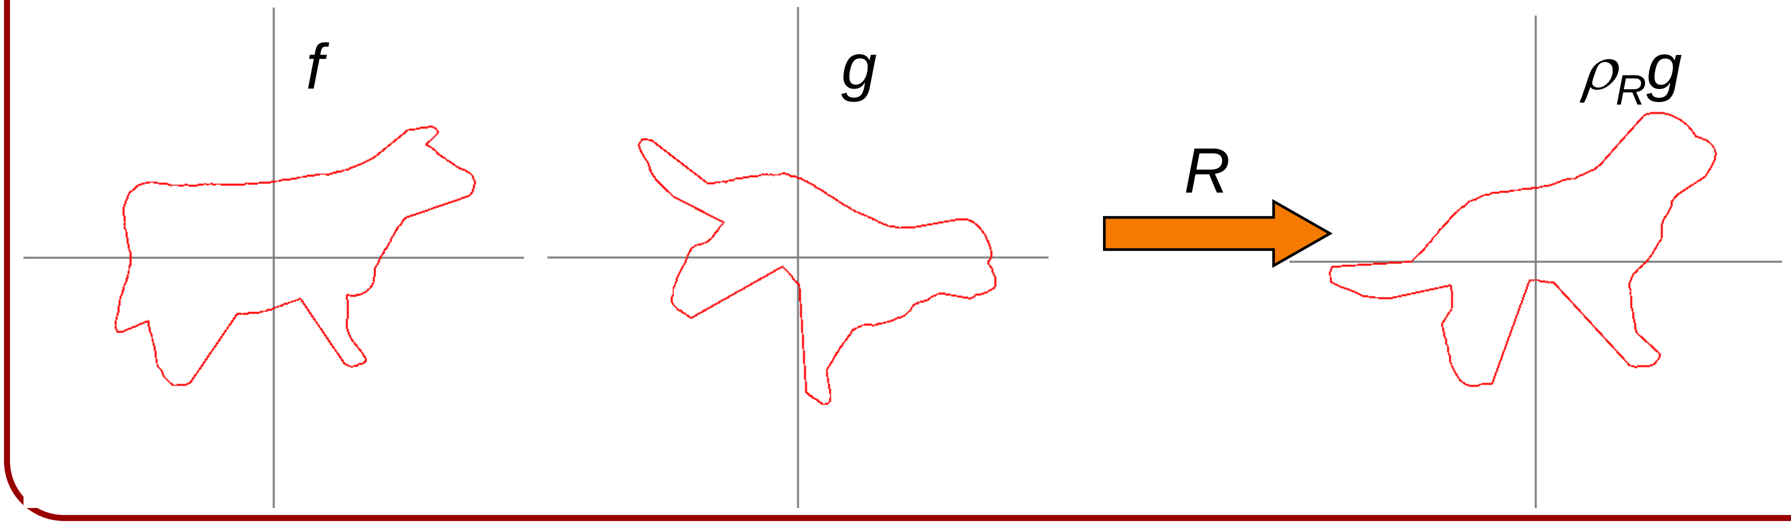

Reduction to a Moving Dot-Product

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

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

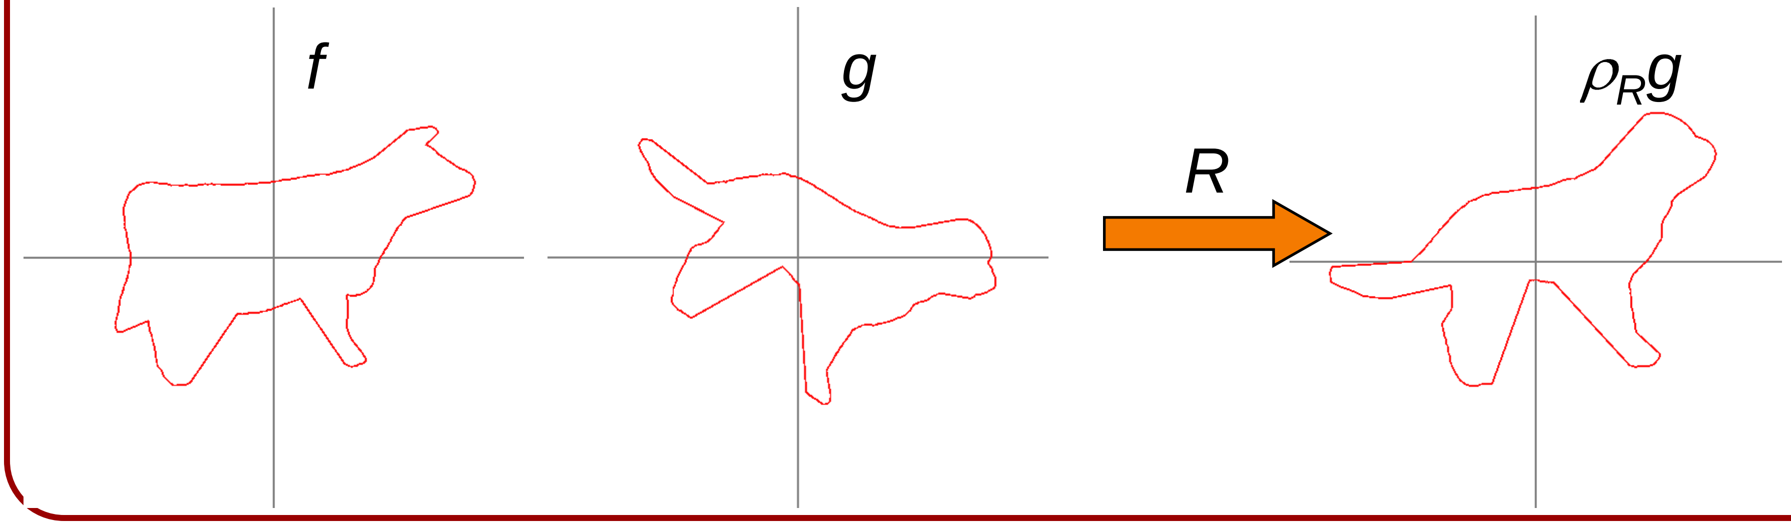

Reduction to a Moving Dot-Product

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

$$\left\| f - \rho_R g \right\|^2 = \left\langle f - \rho_R g, f - \rho_R g \right\rangle$$
$$= \left\langle f, f \right\rangle + \left\langle \rho_R g, \rho_R g \right\rangle - \left\langle f, \rho_R g \right\rangle - \left\langle \rho_R g, f \right\rangle$$

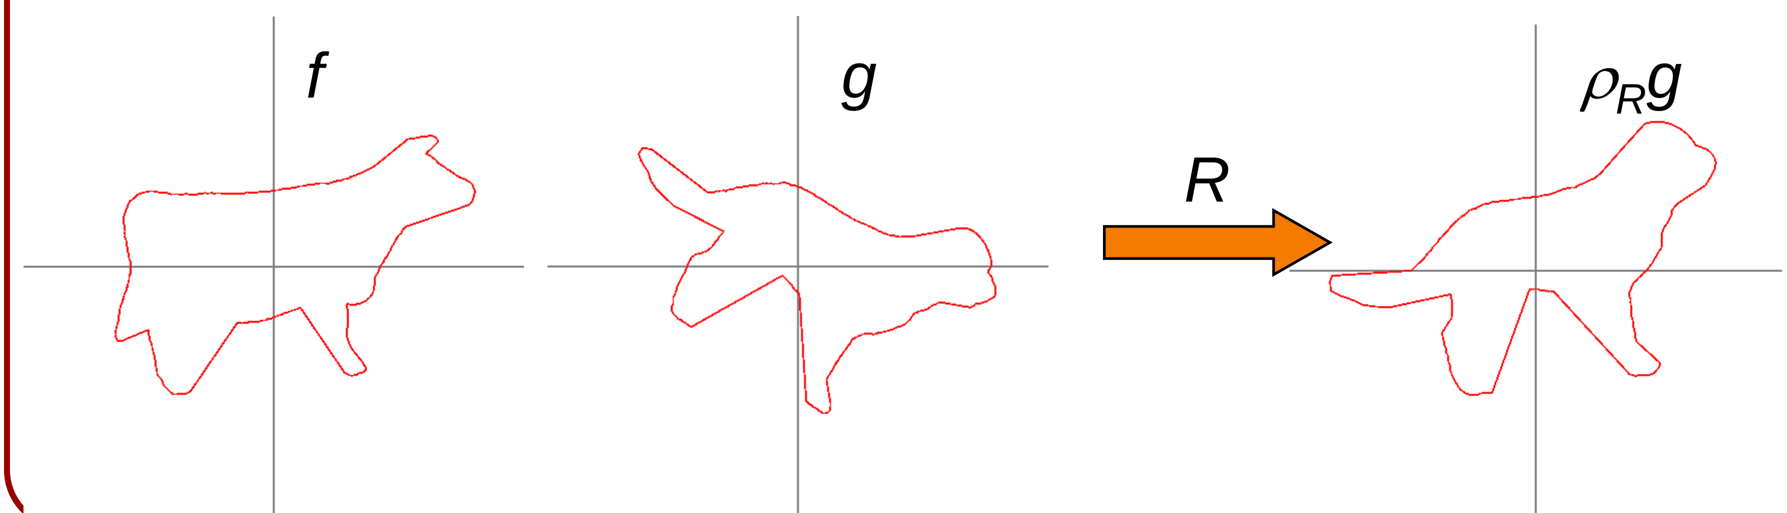

Reduction to a Moving Dot-Product

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

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

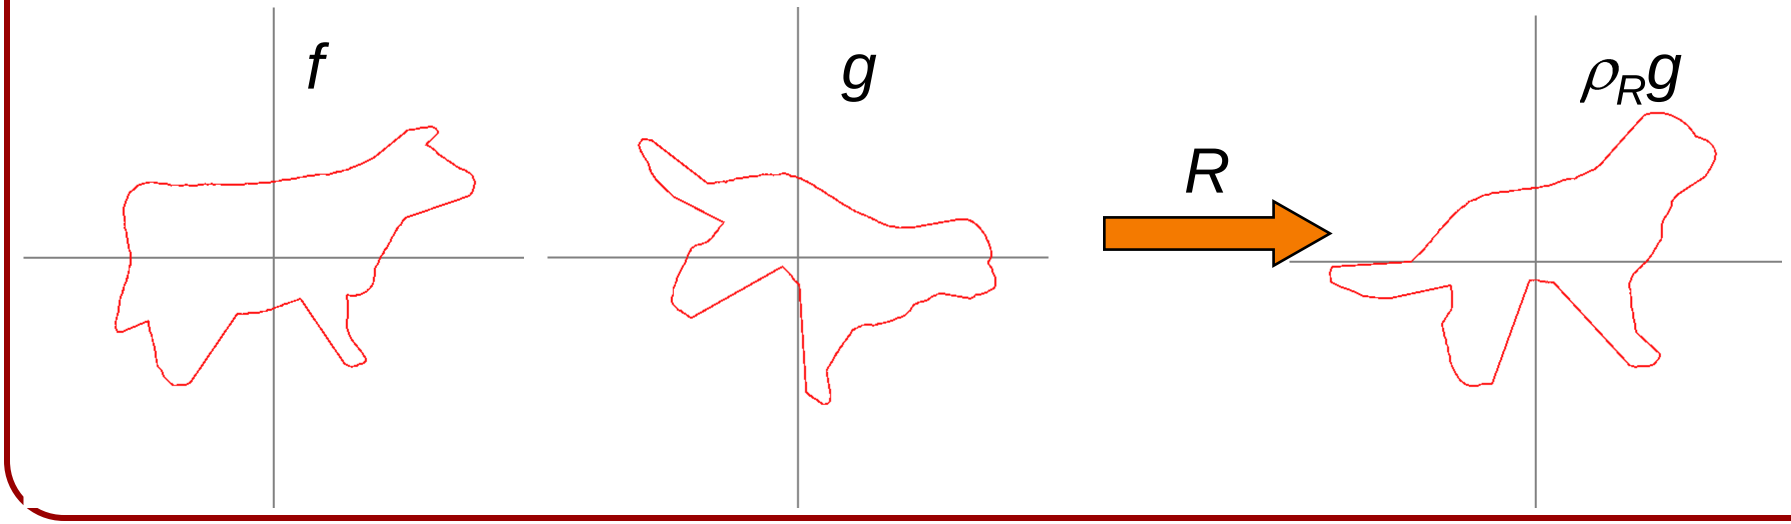

Reduction to a Moving Dot-Product

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

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

$\Rightarrow$ Finding the rotation minimizing the norm is equivalent to finding the rotation maximizing the dot-product.

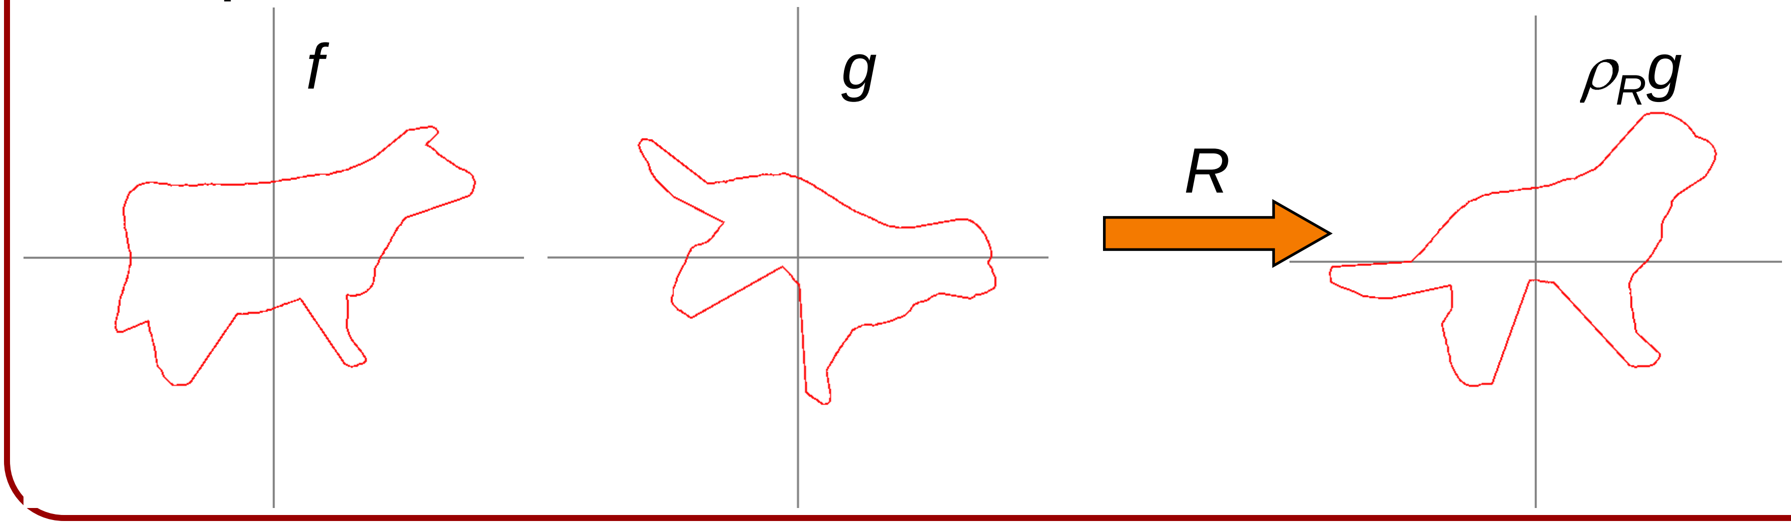

General Approach

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

$$\text{Dot}_{f,g}(\alpha) = \langle f, \rho_\alpha g \rangle$$

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

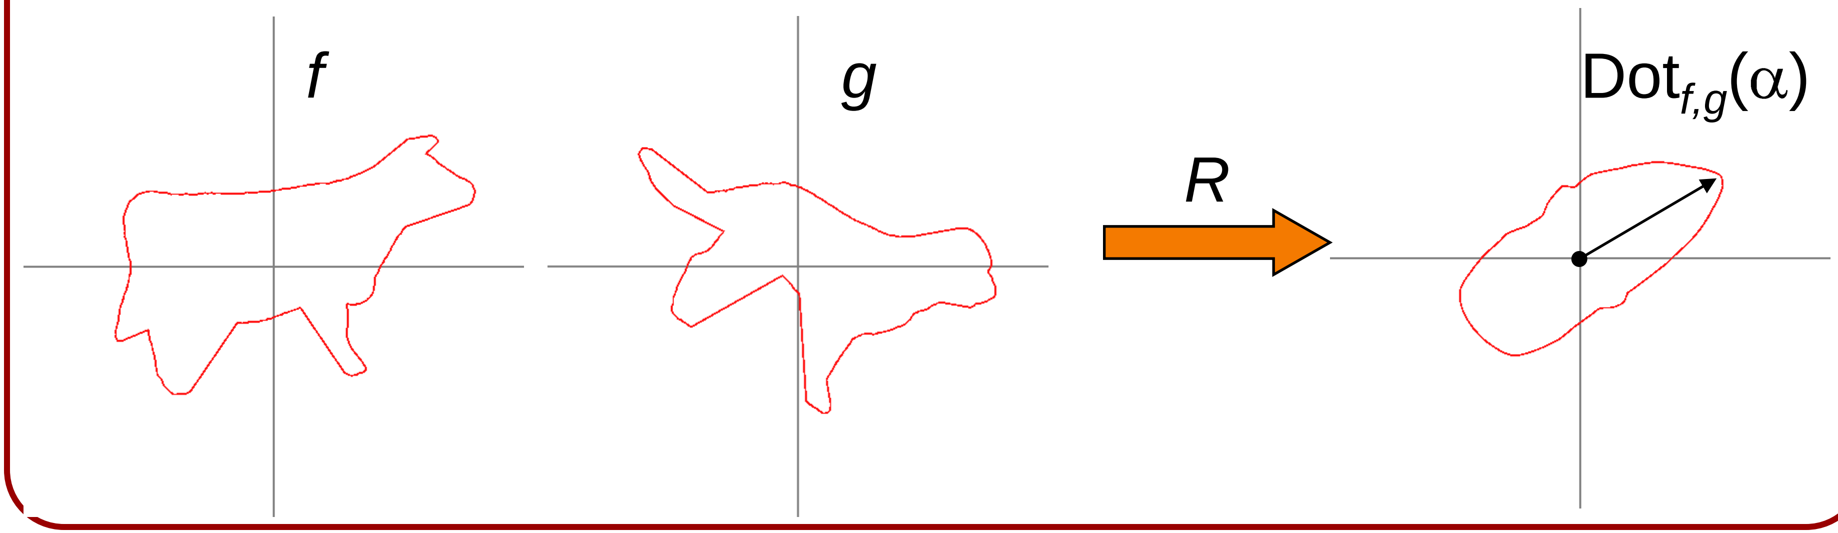

Brute-Force

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

Brute-Force

To compute $\text{Dot}_{f,g}(\alpha)$, 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:

$$f(\theta) = \sum_k \hat{f}[k] e^{ik\theta}$$

$$g(\theta) = \sum_k \hat{g}[k] e^{ik\theta}$$

Irreducible Representations

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

$$\text{Dot}_{f,g}(\alpha) = \left\langle \sum_k \hat{f}(k)e^{ik\theta}, \rho_\alpha\left(\sum_{k'} \hat{g}(k')e^{ik'\theta} \right) \right\rangle$$

Irreducible Representations

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

$$\begin{aligned}\text{Dot}_{f,g}(\alpha) &= \left\langle \sum_k \hat{f}(k)e^{ik\theta}, \rho_\alpha\left(\sum_{k'} \hat{g}(k')e^{ik'\theta}\right)\right\rangle \\ &= \left\langle \sum_k \hat{f}(k)e^{ik\theta}, \sum_{k'} \hat{g}(k')\rho_\alpha\left(e^{ik'\theta}\right)\right\rangle\end{aligned}$$

Irreducible Representations

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

$$\text{Dot}_{f,g}(\alpha) = \left\langle \sum_k \hat{f}(k) e^{ik\theta}, \rho_\alpha \left(\sum_{k'} \hat{g}(k') e^{ik'\theta} \right) \right\rangle$$

$$= \left\langle \sum_k \hat{f}(k) e^{ik\theta}, \sum_{k'} \hat{g}(k') \rho_\alpha \left(e^{ik'\theta} \right) \right\rangle$$

$$= \sum_{k,k'} \hat{f}(k) \overline{\hat{g}(k')} \left\langle e^{ik\theta}, \rho_\alpha \left(e^{ik'\theta} \right) \right\rangle$$

Irreducible Representations

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

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

Irreducible Representations

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

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

Then the equation for the dot-product becomes:

$$\mathrm{Dot}_{f,g}(\alpha) = \sum_{k,k'} \hat{f}(k)\overline{\hat{g}(k')}\left\langle e^{ik\theta}, \rho_\alpha\left(e^{ik'\theta}\right)\right\rangle$$

$$\mathrm{Dot}_{f,g}(\alpha) = \sum_{k,k'} \hat{f}(k)\overline{\hat{g}(k')}D_{k,k'}(\alpha)$$

Irreducible Representations

$$\text{Dot}_{f,g}(\alpha) = \sum_{k,k'} \hat{f}(k)\overline{\hat{g}(k')} D_{k,k'}(\alpha)$$

Thus, 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 $D_{k,k'}(\alpha)$

Irreducible Representations

$$\text{Dot}_{f,g}(\alpha) = \sum_{k,k'} \hat{f}(k)\overline{\hat{g}(k')} D_{k,k'}(\alpha)$$

Thus, 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

The advantage of using the Fourier decomposition, is that 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 $k \neq k'$)

Irreducible Representations

The advantage of using the Fourier decomposition, is that 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:

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

Irreducible Representations

This means that the expression for the correlation becomes:

$$\mathrm{Dot}_{f,g}(\alpha) = \sum_k \hat{f}(k)\overline{\hat{g}(k)}D_k(\alpha)$$

Irreducible Representations

This means that the expression for the correlation becomes:

$$\mathrm{Dot}_{f,g}(\alpha)=\sum_k \hat{f}(k)\overline{\hat{g}(k)}D_k(\alpha)$$

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

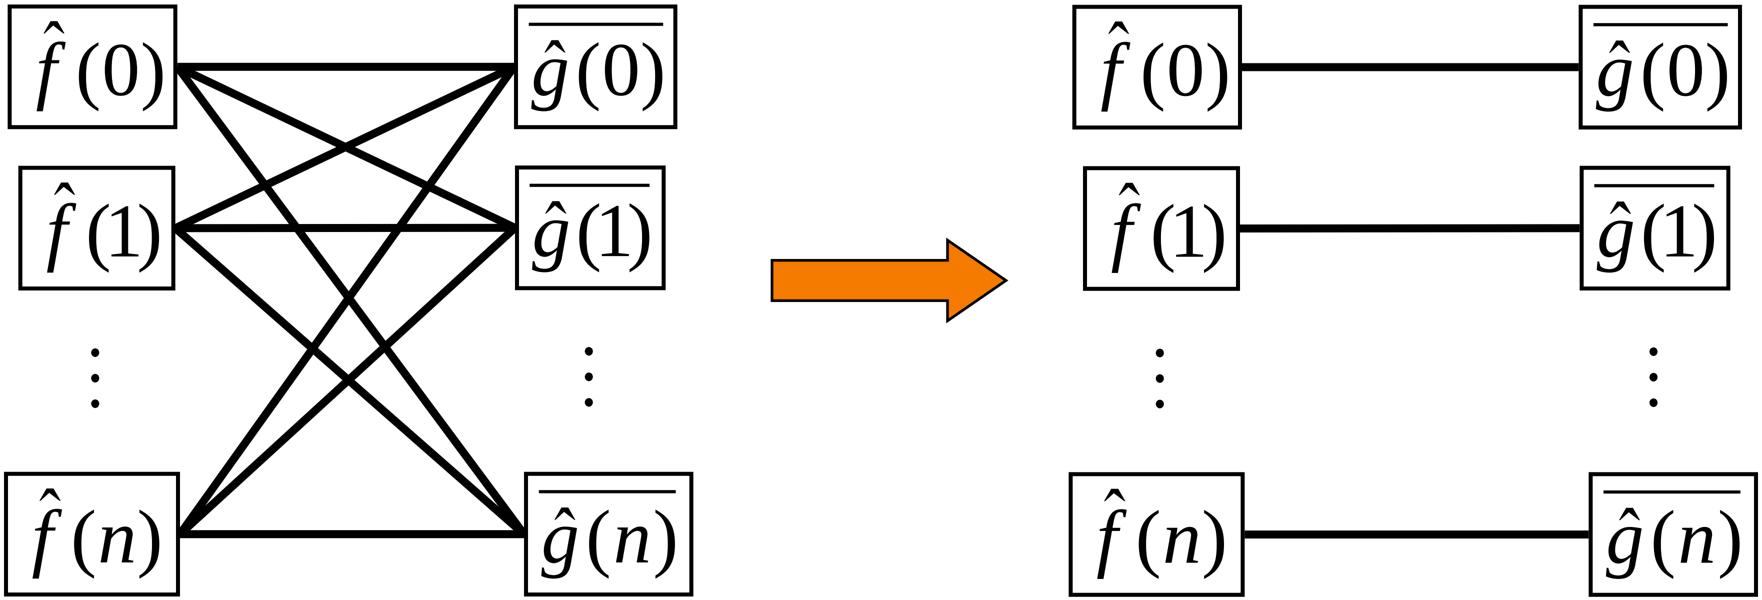

Change of Basis

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

$$\text{Dot}_{f,g}(\alpha) = \sum_k \hat{f}(k)\overline{\hat{g}(k)}D_k(\alpha)$$

Change of Basis

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

$$\text{Dot}_{f,g}(\alpha) = \sum_k \hat{f}(k)\overline{\hat{g}(k)}D_k(\alpha)$$

This is not quite good enough, because to evaluate the correlation at α, we need to get the value of each of the $D_k(\alpha)$, and then take the linear combination, using the weights $\hat{f}(k)\overline{\hat{g}(k)}$.

Change of Basis

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

$$\mathrm{Dot}_{f,g}(\alpha) = \sum_k \hat{f}(k)\overline{\hat{g}(k)}D_k(\alpha)$$

This is not quite good enough, because to evaluate the correlation at α, we need to get the value of each of the $D_k(\alpha)$, and then take the linear combination, using the weights $\hat{f}(k)\overline{\hat{g}(k)}$.

That, is evaluating the correlation at any single angle would require $O(n)$ computations.

Evaluating at all angles would take $O(n^2)$

Change of Basis

If we let *c*[ ] be the *n*-dimensional array of coefficients:

$$c[k] = \hat{f}(k)\overline{\hat{g}(k)}$$

and we let *a*[ ] be the *n*-dimensional array of correlation values:

$$a[k] = \mathrm{Dot}_{f,g}\left(\frac{2k\pi}{n}\right)$$

Change of Basis

If we let *c*[ ] be the *n*-dimensional array of coefficients:

$$c[k] = \hat{f}(k)\overline{\hat{g}(k)}$$

and we let *a*[ ] be the *n*-dimensional array of correlation values:

$$a[k] = \mathrm{Dot}_{f,g}\left(\frac{2k\pi}{n}\right)$$

We get:

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

Change of Basis

If we let *c*[ ] be the *n*-dimensional array of coefficients:

$$c[k] = \hat{f}(k)\overline{\hat{g}(k)}$$

and we let *a*[ ] be the *n*-dimensional array of correlation values:

$$\left(\frac{2k\pi}{\cdots}\right)$$

We [illegible]

To get the desired expression for the correlation, we need to perform a change of basis!

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

Change of Basis

It turns out that computing this change of basis amounts to computing the Inverse Fourier Transform.

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 *g* to *f*:

$$R = \underset{R \in \text{Rotations}}{\arg\min} \left\| f - \rho_R g \right\|^2$$

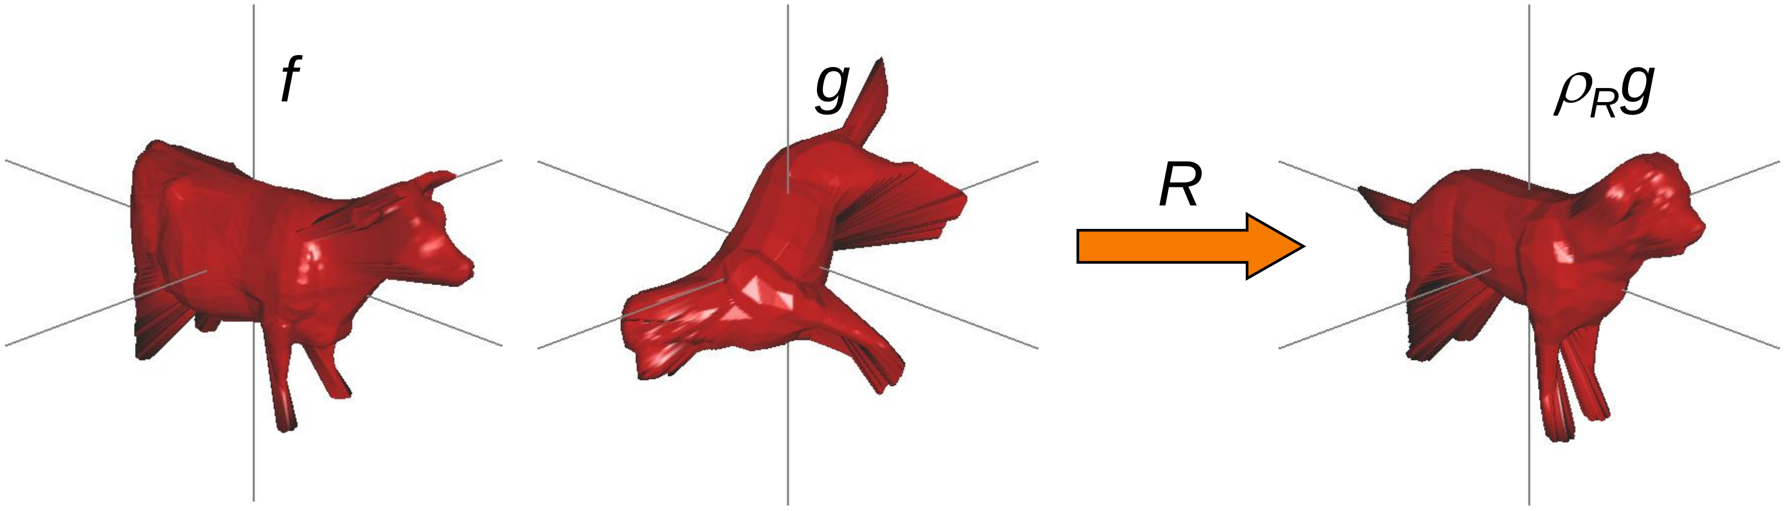

Expanding the Norm

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

$$R = \underset{R \in \mathrm{Rotations}}{\arg\min} \left\| f - \rho_R g \right\|^2$$

Expanding the norm, we get:

$$\left\| f - \rho_R g \right\|^2 = \left\| f \right\|^2 + \left\| g \right\|^2 - 2\left\langle \rho_R g, f \right\rangle$$

Expanding the Norm

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

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

$$\text{Dot}_{f,g}(R) = \langle \rho_R g, f\rangle$$

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 $\rho_R g$ with f.

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 $\rho_R g$ with f.

If we think of n as the resolution of the spherical function, the "size" of a spherical function is $O(n^2)$ and the "size" of the space of rotations in $O(n^3)$.

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 $\rho_R g$ with f.

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

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:

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

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

Irreducible Representations

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

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

Irreducible Representations

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

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

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

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

Irreducible Representations

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

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

Using the conjugate-linearity of the inner product, we can pull the linear summations outside of the inner product:

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

Irreducible Representations

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

Using the facts 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*≠*l'*)

we get:

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

Change of Basis

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

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

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

Change of Basis

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

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

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

We get:

$$\text{Dot}_{f,g}(R)=\sum_{l=0}^{b}\sum_{m,m'=-l}^{l}\hat{g}(l,m)\overline{\hat{f}(l,m')}D_l^{m',m}(R)$$

Change of Basis

$$\text{Dot}_{f,g}(R) = \sum_{l=0}^{b} \sum_{m,m'=-l}^{l} \hat{g}(l,m)\overline{\hat{f}(l,m')} D_l^{m',m}(R)$$

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

$$\mathrm{Dot}_{f,g}(R) = \sum_{l=0}^{b} \sum_{m,m'=-l}^{l} \hat{g}(l,m)\,\overline{\hat{f}(l,m')}\,D_l^{m',m}(R)$$

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 $-l \le m, m' \le l$,
- And then take the linear sum weighted by the product of the harmonic coefficients

Change of Basis

$$\text{Dot}_{f,g}(R) = \sum_{l=0}^{b} \sum_{m,m'=-l}^{l} \hat{g}(l,m)\overline{\hat{f}(l,m')} D_l^{m',m}(R)$$

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.

Change of Basis

$$\mathrm{Dot}_{f,g}(R) = \sum_{l=0}^{b} \sum_{m,m'=-l}^{l} \hat{g}(l,m)\overline{\hat{f}(l,m')} D_l^{m',m}(R)$$

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 the brute force method since it requires O(n^6) operations to get the values of the correlation at every rotation, while the brute force method requires only O(n^5) operations.

Change of Basis

What is that we really want to do?

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.

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^3x n^3= n^6.

Change of Basis

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

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

what do the functions:

$$D_l^{m',m}\left(R(\theta,\phi,\psi)\right)$$

look like?

Change of Basis

Recall that the spherical harmonics can be expressed as a complex exponential in θ times a "polynomial" in cosϕ:

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

Change of Basis

Recall that the spherical harmonics can be expressed as a complex exponential in θ times a "polynomial" in cosφ:

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

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

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

Change of Basis

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

$$D_l^{m',m}(\theta,\phi,\psi) = \left\langle \rho_{R_y(\theta)}\rho_{R_z(\phi)}\rho_{R_y(\psi)}\left(Y_l^m\right), Y_l^{m'} \right\rangle$$

Change of Basis

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

$$D_l^{m',m}(\theta,\phi,\psi) = \left\langle \rho_{R_y(\theta)}\rho_{R_z(\phi)}\rho_{R_y(\psi)}\left(Y_l^m\right), Y_l^{m'} \right\rangle$$

$$= \left\langle \rho_{R_z(\phi)}\rho_{R_y(\psi)}\left(Y_l^m\right), \rho_{R_y(-\theta)} Y_l^{m'} \right\rangle$$

Change of Basis

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

$$D_l^{m',m}(\theta,\phi,\psi) = \left\langle \rho_{R_y(\theta)}\rho_{R_z(\phi)}\rho_{R_y(\psi)}\big(Y_l^m\big), Y_l^{m'} \right\rangle$$

$$= \left\langle \rho_{R_z(\phi)}\rho_{R_y(\psi)}\big(Y_l^m\big), \rho_{R_y(-\theta)} Y_l^{m'} \right\rangle$$

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

Change of Basis

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

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

Change of Basis

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

Thus, denoting:

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

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

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

Change of Basis

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

The advantage of this representation is that now a significant part of the basis is expressed in terms of the complex exponentials, so we can use the inverse FFT to help us perform the change of basis efficiently.

Change of Basis

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

We can think of the sampled correlation function as an *n*x*n*x*n* grid, whose (*p*,*q*,*r*)-th entry corresponds to the value of the correlation at the Euler angle (θ_p,ϕ_q,ψ_r)

$$\theta_p = \frac{2\pi p}{n}$$

$$\phi_q = \frac{\pi(2q+1)}{2n}$$

$$\psi_r = \frac{2\pi r}{n}$$

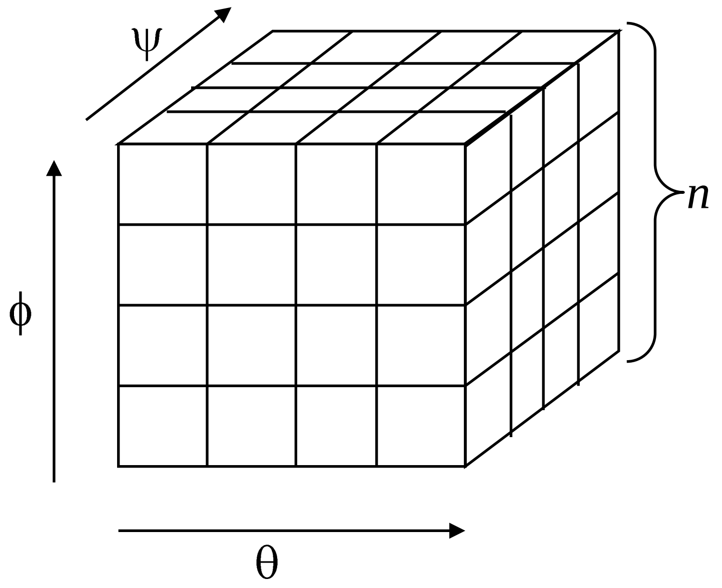

Change of Basis

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

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

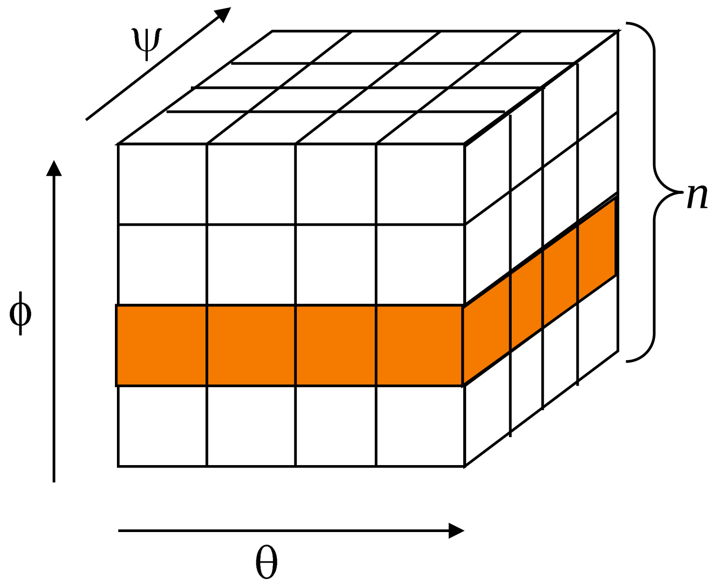

Change of Basis

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

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:

$$\Omega_{f,g,\phi}(\theta,\psi) = \sum_{l=0}^{b} \sum_{m,m'=-l}^{l} \hat{g}(l,m)\overline{\hat{f}(l,m')} D_l^{m',m}(\theta,\phi,\psi)$$

Change of Basis

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

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:

$$\Omega_{f,g,\phi}(\theta,\psi) = \sum_{l=0}^{b} \sum_{m,m'=-l}^{l} \hat{g}(l,m)\overline{\hat{f}(l,m')} D_l^{m',m}(\theta,\phi,\psi)$$

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

Change of Basis

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

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

$$\hat{\Omega}_{f,g,\phi}(m,m')=\sum_{l=\max(|m|,|m'|)}^{b}\hat{g}(l,-m)\overline{\hat{f}(l,-m')}d_l^{-m,-m'}(\phi)$$

Change of Basis

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

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

$$\hat{\Omega}_{f,g,\phi}(m,m')=\sum_{l=\max(|m|,|m'|)}^{b}\hat{g}(l,-m)\overline{\hat{f}(l,-m')}d_l^{-m,-m'}(\phi)$$

So we can get the values in this 2D slice by running the 2D inverse FFT.

Change of Basis

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

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

Change of Basis

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

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

For every sampled value of ϕ:

- We compute the Fourier coefficients:

$$\hat{\Omega}_{f,g,\phi}(m,m') = \sum_{l=\max(|m|,|m'|)}^{b} \hat{g}(l,-m)\overline{\hat{f}(l,-m')} d_l^{-m,-m'}(\phi)$$

- 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.

Change of Basis

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

For each slice there are *n*x*n* different Fourier coefficients, requiring O(*n*) calculations per coefficient:

$$\hat{\Omega}_{f,g,\phi}(m,m') = \sum_{l=\max(|m|,|m'|)}^{b} \hat{g}(l,-m)\overline{\hat{f}(l,-m')}\,d_l^{-m,-m'}(\phi)$$

Change of Basis

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

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

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

Change of Basis

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

For each slice there are nxn 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 nxn different Fourier coefficients, requiring O(n) calculations per co

And, in particular, we can do much better than the brute force algorithm

An

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.

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 $n^2 \mathrm{x} n^2$ matrix multiplication in two of the variables by an $O(n^2 \log n)$ inverse FFT.

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^2xn^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 nxn 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 $O(n \log^2 n)$.

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)|≤1) we would like to compute the distance between the functions at every rotation:

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

Correlating 3D Functions

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

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

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

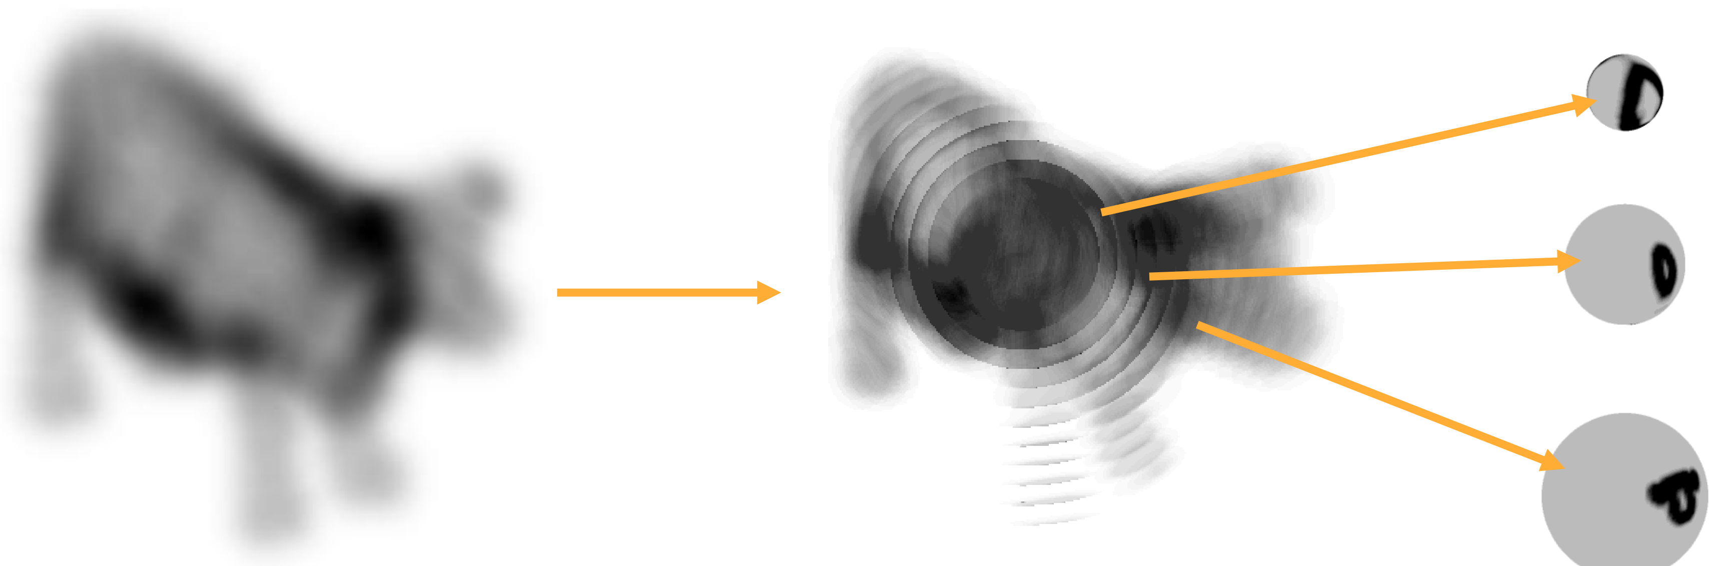

Correlating 3D Functions

The value of the correlation then becomes:

$$\langle \rho_R G, F \rangle = \int_0^1 \langle \rho_R G_r, F_r \rangle r^2 dr$$

Correlating 3D Functions

The value of the correlation then becomes:

$$\langle \rho_R G, F \rangle = \int_0^1 \langle \rho_R G_r, F_r \rangle r^2 dr$$

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

$$F_r = \sum_{l=0}^{b} \sum_{m=-l}^{l} \hat{F}_r(l,m) Y_l^m$$

$$G_r = \sum_{l=0}^{b} \sum_{m=-l}^{l} \hat{G}_r(l,m) Y_l^m$$

Correlating 3D Functions

The value of the correlation then becomes:

$$\langle \rho_R G, F \rangle = \int_0^1 \langle \rho_R G_r, F_r \rangle r^2 dr$$

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

$$F_r = \sum_{l=0}^{b} \sum_{m=-l}^{l} \hat{F}_r(l,m) Y_l^m$$

$$G_r = \sum_{l=0}^{b} \sum_{m=-l}^{l} \hat{G}_r(l,m) Y_l^m$$

we get:

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

Correlating 3D Functions

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

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

Correlating 3D Functions

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

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 given an algorithm with complexity O(n^5) / O($n^4 \log^2 n$).

Correlating 3D Functions

We can do better.

Correlating 3D Functions

We can do better.

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

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

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

Correlating 3D Functions

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

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

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

Algorithm: (Assuming O(n) radial samples)

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

Correlating 3D Functions

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

Algorithm: (Assuming O(n) radial samples)

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

Correlating 3D Functions

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

Algorithm: (Assuming O(n) radial samples)

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

Correlating 3D Functions

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

Algorithm: (Assuming O(n) radial samples)

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

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