



FFTs in Graphics and Vision

Alignment, Invariance
and Pattern Matching



Warning

From here on in, when considering the inner-product between vectors $\mathbf{a}, \mathbf{b} \in \mathbb{C}^n$, we will drop the subscript and simply write:

$$\langle \mathbf{a}, \mathbf{b} \rangle_{[0, 2\pi)} \equiv \langle \mathbf{a}, \mathbf{b} \rangle$$



Outline

Alignment

Shape Matching

Invariance

Pattern Matching



Recall

When working with periodic arrays, we have:

- $V = \mathbb{C}^n$ is the space of periodic arrays ($\mathbf{v}_i = \mathbf{v}_{i+n}$, for all $\mathbf{v} \in V$ and all $i \in \mathbb{Z}/n\mathbb{Z}$).
- $G = \mathbb{Z}/n\mathbb{Z}$ is the group of (periodic) shifts
- ρ_α is the representation that shifts the coefficients of the array by $\alpha \in G$ indices.

We have:

- The Fourier basis $\{\mathbf{z}^0, \dots, \mathbf{z}^{n-1}\} \subset V$, obtained by regularly sampling the complex exponentials (and normalizing).
- We have the scaling vectors $\{\mathbf{x}^0, \dots, \mathbf{x}^{n-1}\}$ with the property that:

$$\rho_\alpha(\mathbf{z}^k) = \mathbf{x}_\alpha^k \cdot \mathbf{z}^k$$



Shape Representation

For 2D shape matching/analysis, it is common to represent the geometry of a shape by a circular array of real values.




Shape Representation

Example:

The circular extent function represents the extent of the shape about the center of mass:


- The value at an angle θ is the distance to the last point of intersection of the ray from the origin, with angle θ , with the shape





Alignment


Since the shape of an object doesn't change when we rotate it, we would like to know if the two arrays are equivalent, up to rotation.





Alignment

Is there a rotation that will rotate the first array into the second?





Alignment

Is there a rotation that will rotate the first array into the second?

Given arrays $\mathbf{f}, \mathbf{g} \in \mathbb{R}^n$, is there an index $\alpha \in G$ such that:

$$\begin{aligned} \mathbf{g} &= \rho_{\alpha}(\mathbf{f}) \\ &\Updownarrow \\ \mathbf{g}_k &= \mathbf{f}_{k-\alpha} \quad \forall 0 \leq k < n \end{aligned}$$



Alignment

In a continuous setting, asking the binary question “are the arrays equal” is not very meaningful, since

- Sampling
- Noise
- Etc

can cause “equal” arrays to have different values.



Alignment

Is there a rotation that will rotate the first array
into the second?



For every rotation, how close is the rotation of the
first array to the second array?

For every rotation $\alpha \in G$, what is the value of:

$$D_{\mathbf{f},\mathbf{g}}^2(\alpha) = \|\rho_\alpha(\mathbf{f}) - \mathbf{g}\|^2$$



Alignment

At every α we would like to evaluate:

$$\begin{aligned} D_{\mathbf{f}, \mathbf{g}}^2(\alpha) &= \|\rho_\alpha(\mathbf{f}) - \mathbf{g}\|^2 \\ &= \langle \rho_\alpha(\mathbf{f}) - \mathbf{g}, \rho_\alpha(\mathbf{f}) - \mathbf{g} \rangle \\ &= \langle \rho_\alpha(\mathbf{f}), \rho_\alpha(\mathbf{f}) \rangle + \langle \mathbf{g}, \mathbf{g} \rangle - \langle \rho_\alpha(\mathbf{f}), \mathbf{g} \rangle - \overline{\langle \rho_\alpha(\mathbf{f}), \mathbf{g} \rangle} \end{aligned}$$

Since the Hermitian dot-product of two real-valued arrays is also a real value:

$$= \|\rho_\alpha(\mathbf{f})\|^2 + \|\mathbf{g}\|^2 - 2\langle \rho_\alpha(\mathbf{f}), \mathbf{g} \rangle$$

Since ρ_α is a unitary transformation:

$$= \|\mathbf{f}\|^2 + \|\mathbf{g}\|^2 - 2\langle \rho_\alpha(\mathbf{f}), \mathbf{g} \rangle$$



Alignment

$$D_{\mathbf{f},\mathbf{g}}^2(\alpha) = \|\mathbf{f}\|^2 + \|\mathbf{g}\|^2 - 2\langle \rho_\alpha(\mathbf{f}), \mathbf{g} \rangle$$

To compute the distance between a rotation of the circular array \mathbf{f} by α , and the circular array \mathbf{g} , we need to compute:

- The magnitude of \mathbf{f} : $\|\mathbf{f}\|^2$,
- The magnitude of \mathbf{g} : $\|\mathbf{g}\|^2$,
- The value of the correlation: $\langle \rho_\alpha(\mathbf{f}), \mathbf{g} \rangle$



Alignment


$$D_{\mathbf{f},\mathbf{g}}^2(\alpha) = \|\mathbf{f}\|^2 + \|\mathbf{g}\|^2 - 2\langle \rho_\alpha(\mathbf{f}), \mathbf{g} \rangle$$

- The magnitude of $\mathbf{f} \in \mathbb{R}^n$:
 - Constant independent of α : $O(n)$ time.
- The magnitude of $\mathbf{g} \in \mathbb{R}^n$:
 - Constant independent of α : $O(n)$ time.
- The value of the correlation:
 - With an FFT: $O(n \log n)$ time.

Alignment



$$D_{\mathbf{f},\mathbf{g}}^2(\alpha) = \|\mathbf{f}\|^2 + \|\mathbf{g}\|^2 - 2\langle \rho_\alpha(\mathbf{f}), \mathbf{g} \rangle$$






Alignment

$$D_{\mathbf{f},\mathbf{g}}^2(\alpha) = \|\mathbf{f}\|^2 + \|\mathbf{g}\|^2 - \boxed{2\langle \rho_\alpha(\mathbf{f}), \mathbf{g} \rangle}$$

Because the norms are constant, instead of looking for the minimum distance, we can look for the maximum dot-product.





Alignment

$$D_{\mathbf{f},\mathbf{g}}^2(\alpha) = \|\mathbf{f}\|^2 + \|\mathbf{g}\|^2 - \boxed{2\langle \rho_\alpha(\mathbf{f}), \mathbf{g} \rangle}$$

Because the norms are constant, instead of looking for the minimum distance, we can look for the maximum dot-product.

\mathbf{f}

\mathbf{g}

$(\mathbf{f} \star \mathbf{g})_\alpha$

The maximal dot-product:

- Lets us determine the best alignment
- Doesn't let us compare across shapes



Outline

Alignment

Shape Matching

Invariance

Pattern Matching

Shape Matching



In shape matching applications, we would like to find the shapes in a database that are most similar to a given query.

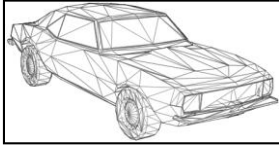
Princeton 3D Model Search Engine - Microsoft Internet Explorer provided by Verizon Online

















Princeton Shape Retrieval and Analysis Group
3D Model Search Engine

[Text & 2D Sketch](#) [Text & 3D Sketch](#) [File Compare](#) [Research](#) [Contact Us](#) [Links](#) [FAQ](#) [Main](#)

Search results in database [espona], 1000 models (click on a thumbnail for more information on that model)

[Next page \(17 - 32\)](#) search type: [similar shape], results: 100


Query

 1. karmann (esp) Find similar shape	 2. chevy (esp) Find similar shape	 3. prelude (esp) Find similar shape	 4. lanc b24 (esp) Find similar shape
 5. 850p (esp) Find similar shape	 6. skyline (esp) Find similar shape	 7. m300se (esp) Find similar shape	 8. mer300sl (esp) Find similar shape
 9. nsx (esp) Find similar shape	 10. jaguar (esp) Find similar shape	 11. citrxm (esp) Find similar shape	 12. mercedess600 (esp) Find similar shape
 13. astonm (esp) Find similar shape	 14. f250 (esp) Find similar shape	 15. bmw502 (esp) Find similar shape	 16. ch54 (esp) Find similar shape

[Next page \(17 - 32\)](#) Something didn't work? [Let us know!](#)

Opening <http://shape.cs.princeton.edu/search/sketch.cgi> Internet



Shape Matching

General approach:

Define a function that takes in two models and returns a measure of their proximity.

$$D\left(\begin{array}{|c|} \hline \text{Car Model } M_1 \\ \hline \end{array} \bigg/ \begin{array}{|c|} \hline \text{Car Model } M_2 \\ \hline \end{array}\right) \leq D\left(\begin{array}{|c|} \hline \text{Car Model } M_1 \\ \hline \end{array} \bigg/ \begin{array}{|c|} \hline \text{Wolf Model } M_3 \\ \hline \end{array}\right)$$

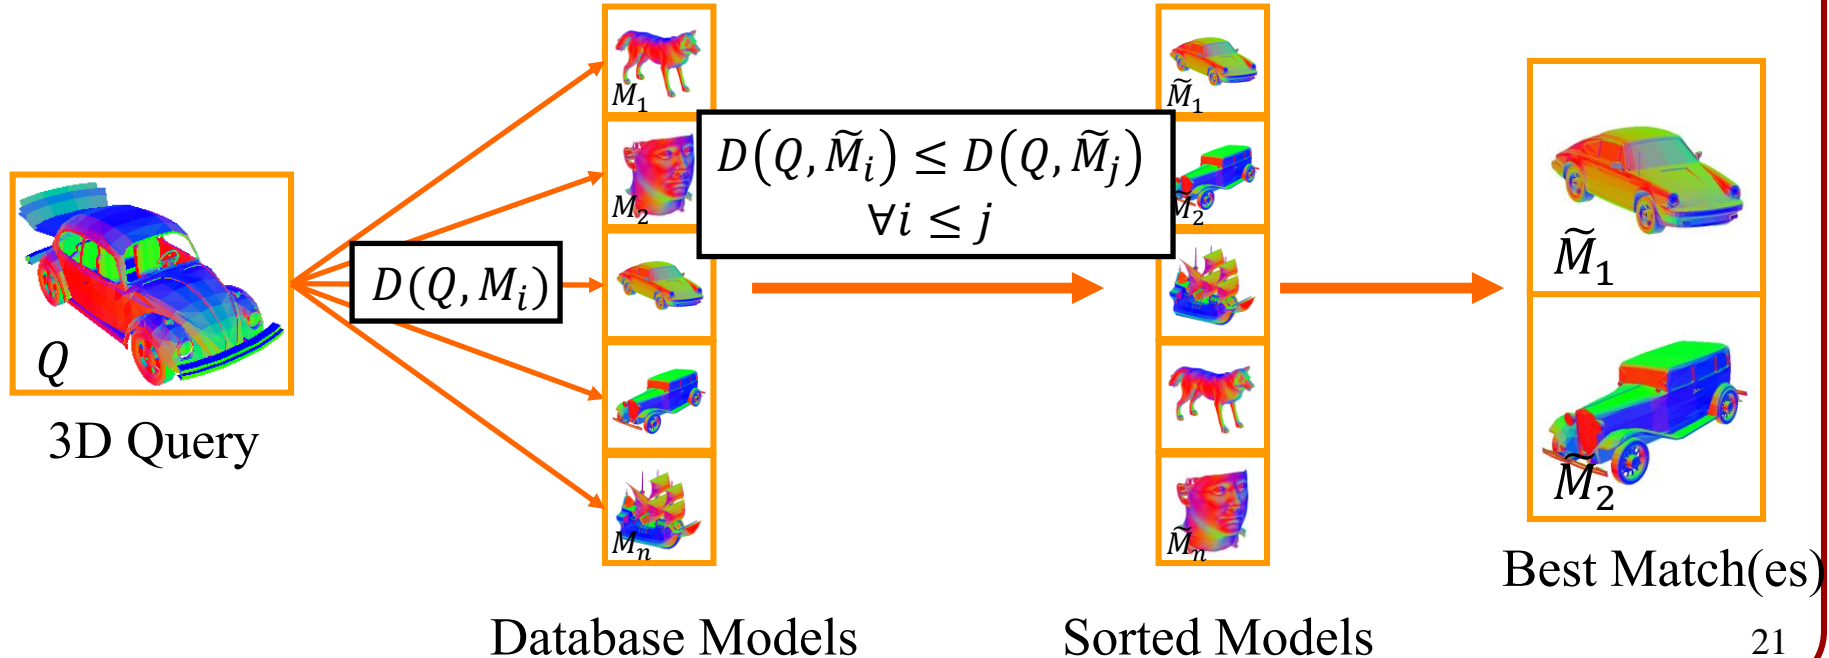


M_1 is closer to M_2 than it is to M_3



Database Retrieval

- Compute the distance from the query to each database model
- Sort the database models by proximity
- Return the closest matches





Shape Matching

To do this efficiently, models are often represented by *Shape Descriptors*:

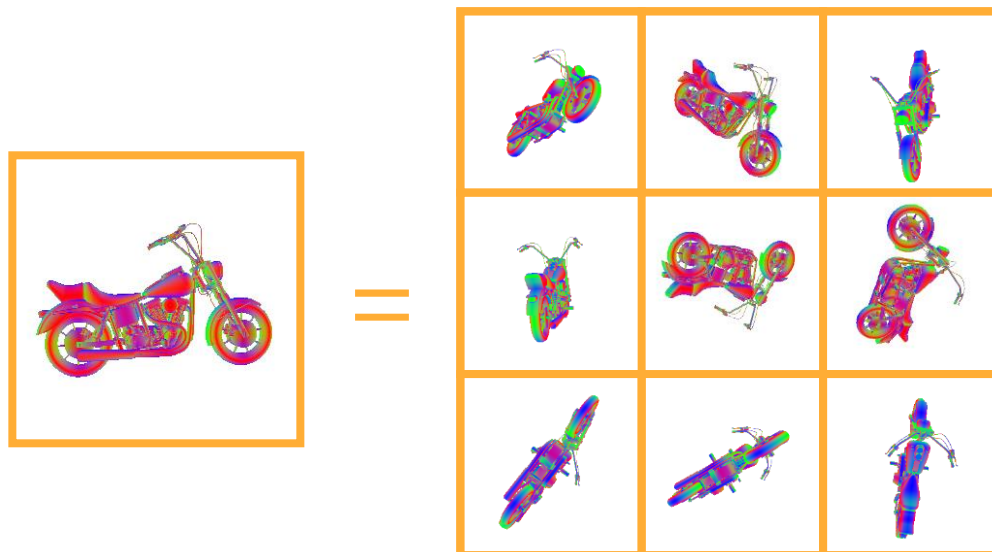
- Arrays of values encapsulating information about the shape of the model, such that
- The distance between the arrays gives a measure of proximity of the underlying shapes.



Shape Matching

Challenge:

Since the shape of the model doesn't change if we rotate it, we would like to match models across rotational poses.





Shape Matching

Challenge:

Since the shape of the model doesn't change if we rotate it, we would like to match models across rotational poses.

Solution 1:

Define the measure of similarity by using the FFT to find the distance between two models at the best possible alignment.



Shape Matching

Challenge:

Since the shape of the model doesn't change if we rotate it, we would like to match models across rotational poses.

Solu

This can be too slow for interactive applications that need to return the best match from very large databases.

Defi

to find the distance between two models at the best

Not quite true for 1D arrays, but becomes more true as the dimension increases.



Shape Matching

Challenge:

Since the shape of the model doesn't change if we rotate it, we would like to match models across rotational poses.

Solution 2:

Design a descriptor that is rotation-invariant:


- Instances of the same shape in different poses will give the same shape descriptor.



Invariance

Given an array $\mathbf{f} \in \mathbb{R}^n$, we would like to define a mapping $\Phi: \mathbb{R}^n \rightarrow \mathbb{R}^d$ taking \mathbf{f} to some other array, s.t.:

$$\Phi(\mathbf{f}) = \Phi(\rho_\alpha(\mathbf{f})) \quad \forall \alpha$$





Invariance

Given an array \mathbf{f} , we can express it in terms of its Fourier decomposition:

$$\mathbf{f} = \sum_{k=0}^{n-1} \hat{\mathbf{f}}_k \cdot \mathbf{z}^k$$

If we rotate \mathbf{f} by α we get:

$$\rho_{\alpha}(\mathbf{f}) = \sum_{k=0}^{n-1} \hat{\mathbf{f}}_k \cdot \rho_{\alpha}(\mathbf{z}^k)$$



Invariance

$$\rho_{\alpha}(\mathbf{f}) = \sum_{k=0}^{n-1} \hat{\mathbf{f}}_k \cdot \rho_{\alpha}(\mathbf{z}^k)$$

Since the $\{\mathbf{z}^0, \dots, \mathbf{z}^{n-1}\}$ are a basis for the one-dimensional irreducible representations:

$$\rho_{\alpha}(\mathbf{z}^k) = \mathbf{x}_{\alpha}^k \cdot \mathbf{z}^k$$

where \mathbf{x}_{α}^k is a complex number.

Since the representation is unitary, we know that \mathbf{x}_{α}^k has unit norm.



Invariance

$$\rho_{\alpha}(\mathbf{f}) = \sum_{k=0}^{n-1} \hat{\mathbf{f}}_k \cdot \mathbf{x}_{\alpha}^k \cdot \mathbf{z}^k \quad \text{w/ } \|\mathbf{x}_{\alpha}^k\| = 1$$

In particular, we have:

$$\|\widehat{\rho_{\alpha}(\mathbf{f})}_k\| = \|\mathbf{x}_{\alpha}^k \cdot \hat{\mathbf{f}}_k\| = \|\hat{\mathbf{f}}_k\| \quad \forall k$$

⇒ We can get a rotation invariant representation of \mathbf{f} by storing only the magnitudes of the Fourier coefficients (i.e. discarding phase):

$$\Phi(\mathbf{f}) = (\|\hat{\mathbf{f}}_0\|, \dots, \|\hat{\mathbf{f}}_{n-1}\|)$$



Invariance

What kind of information do we get when we compare just the amplitudes of the Fourier coefficients?



Invariance


Suppose we are given two arrays \mathbf{f} and \mathbf{g} with only one non-zero Fourier coefficient:

$$\mathbf{f} = \hat{\mathbf{f}}_k \cdot \mathbf{z}^k$$

$$\mathbf{g} = \hat{\mathbf{g}}_k \cdot \mathbf{z}^k$$

what is the measure of similarity at the optimal alignment?

If we rotate \mathbf{f} by α , this amounts to multiplying the k -th Fourier coefficient by $e^{-ik\alpha}$.



But this is just a rotation in the complex plane.



Invariance


Suppose we are given two arrays \mathbf{f} and \mathbf{g} with only one non-zero Fourier coefficient:

$$\mathbf{f} = \hat{\mathbf{f}}_k \cdot \mathbf{z}^k$$

$$\mathbf{g} = \hat{\mathbf{g}}_k \cdot \mathbf{z}^k$$

what is the measure of similarity at the optimal alignment?

At the optimal rotation, the Fourier coefficients are on the same line and the measure of similarity is the difference between the lengths.





Invariance

Storing only the amplitudes we get a shape representation $\Phi(\mathbf{f})$ that:

- Is invariant to rotations
- Provides a measure of similarity that is the distance between \mathbf{f} and \mathbf{g} if we could optimally align the frequency components independently.



- This is a lower bound for the distance between \mathbf{f} and \mathbf{g} at the optimal alignment.



Invariance

How good is the lower bound?

After discarding phase, what's left?


Experiment:

To test this, we can consider what happens when we take two arrays and swap the amplitudes of the Fourier coefficients:


$$\begin{aligned} \mathbf{f} &= \sum_{k=0}^{n-1} \boxed{r_k} e^{i\theta_k} \cdot \mathbf{z}^k & \mathbf{g} &= \sum_{k=0}^{n-1} s_k e^{i\boxed{\phi_k}} \cdot \mathbf{z}^k \\ \text{ASwap}(\mathbf{f}, \mathbf{g}) &= \sum_{k=0}^{n-1} \boxed{r_k} e^{i\boxed{\phi_k}} \cdot \mathbf{z}^k \end{aligned}$$

Diagram illustrating the ASwap operation: The amplitude r_k from \mathbf{f} and the phase ϕ_k from \mathbf{g} are combined to form the new coefficient $r_k e^{i\phi_k}$ in $\text{ASwap}(\mathbf{f}, \mathbf{g})$. Arrows indicate the flow of information from the boxed terms in the first two equations to the boxed terms in the third equation.


Invariance



Amplitude



Phase





Invariance

For human perception, dominant information occurs at image boundaries.

These discontinuities arise when the phases are lined up so the occurrence of boundaries is strongly phase dependent.

If the grid encodes other type of information (non-visual) amplitude can be more important.



Outline

Alignment

Shape Matching

Invariance

Pattern Matching



Notation

Given $\mathbf{f}, \mathbf{g} \in \mathbb{C}^n$, we can define the component-wise product $\mathbf{f} \cdot \mathbf{g}$:

$$(\mathbf{f} \cdot \mathbf{g})_j \equiv \mathbf{f}_j \cdot \mathbf{g}_j$$



Note 1

If $\mathbf{f} \in \mathbb{R}^n$ and $\mathbf{g}, \mathbf{h} \in \mathbb{C}^n$, then:

$$\begin{aligned}\langle \mathbf{f} \cdot \mathbf{g}, \mathbf{h} \rangle &= \frac{2\pi}{n} \sum_{k=0}^{n-1} (\mathbf{f} \cdot \mathbf{g})_k \cdot \bar{\mathbf{h}}_k \\ &= \frac{2\pi}{n} \sum_{k=0}^{n-1} \mathbf{f}_k \cdot \mathbf{g}_k \cdot \bar{\mathbf{h}}_k \\ &= \frac{2\pi}{n} \sum_{k=0}^{n-1} \mathbf{g}_k \cdot \bar{\mathbf{f}}_k \cdot \bar{\mathbf{h}}_k \\ &= \frac{2\pi}{n} \sum_{k=0}^{n-1} \mathbf{g}_k \cdot \overline{(\mathbf{f} \cdot \mathbf{h})_k} \\ &= \langle \mathbf{g}, \mathbf{f} \cdot \mathbf{h} \rangle\end{aligned}$$

That is, component-wise multiplication by $\mathbf{f} \in \mathbb{R}^n$ is a symmetric operator.



Note 2

ρ_α is the unitary representation that shifts an array by α indices:

$$\rho_\alpha(\mathbf{f})_k \equiv \mathbf{f}_{k-\alpha}$$

\Rightarrow For all $k \in \mathbb{Z}/n\mathbb{Z}$ we have:

$$\begin{aligned} (\rho_\alpha(\mathbf{f} \cdot \mathbf{g}))_k &= (\mathbf{f} \cdot \mathbf{g})_{k-\alpha} \\ &= \mathbf{f}_{k-\alpha} \cdot \mathbf{g}_{k-\alpha} \\ &= (\rho_\alpha(\mathbf{f}))_k \cdot (\rho_\alpha(\mathbf{g}))_k \end{aligned}$$

\Downarrow

$$\rho_\alpha(\mathbf{f} \cdot \mathbf{g}) = \rho_\alpha(\mathbf{f}) \cdot \rho_\alpha(\mathbf{g})$$




Pattern Matching

Given an instance of a pattern, find all occurrences of the pattern within a target image:

Pattern f




Target Image g





Pattern Matching

We could compute the correlation of the pattern with the image and look for local maxima:





Pattern Matching

The correlation has large values because the dot product of the image with the translated pattern instance is large.


What causes $\langle f, g \rangle$ to be large?



Pattern Matching

What causes $\langle f, g \rangle$ to be large?

- If the values of f and g are correlated
- If the values of g are large





Pattern Matching

We don't want to measure:

How correlated is the pattern instance with a region in the image?

What we want to measure is:


How similar is the pattern instance with a region in the image?



Pattern Matching

For every point in the image, we want to know how similar the region about the point is to the translated pattern.

Translated Pattern



Restricted Target







Pattern Matching

For every point in the image, we want to know how similar the region about the point is to the translated pattern.

Translated Pattern



Restricted Target







Pattern Matching

For every point in the image, we want to know how similar the region about the point is to the translated pattern.

Translated Pattern



Restricted Target





Pattern Matching

For every point in the image, we want to know how similar the region about the point is to the translated pattern.

Translated Pattern



Restricted Target



Pattern Matching

How do we express this formally?





Pattern Matching

How do we express this formally?

If we represent the pattern by \mathbf{f} , then the translation of \mathbf{f} to the point p is written as:

$$\rho_p(\mathbf{f})$$

We want to restrict \mathbf{g} to the region about p by zeroing out the part of \mathbf{g} away from p .



Pattern Matching

How do we express this formally?


Let \mathbf{m} be the masking function for the pattern:

$$\mathbf{m}_{jk} = \begin{cases} 1 & \text{if } \mathbf{f}_{jk} \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

\mathbf{f}



\mathbf{m}





Pattern Matching

How do we express this formally?

The restriction of \mathbf{g} to the region about p can be expressed as:

$$\rho_p(\mathbf{m}) \cdot \mathbf{g}$$

- $\rho_p(\mathbf{m})$ translates the mask so it's centered on p .
- Multiplying by $\rho_p(\mathbf{m})$ zeros out everything except for the region about p .







Pattern Matching

How do we express this formally?

The restriction of g to the region about p can be expressed as:

$$\rho_p(\mathbf{m}) \cdot g$$






Pattern Matching

How do we express this formally?

For every p , we would like to compute:

$$D_{\mathbf{f},\mathbf{g}}^2(p) = \|\rho_p(\mathbf{f}) - \rho_p(\mathbf{m}) \cdot \mathbf{g}\|^2$$





Pattern Matching

How do we express this formally?

For every p , we would like to compute:

$$D_{\mathbf{f},\mathbf{g}}^2(p) = \|\rho_p(\mathbf{f}) - \rho_p(\mathbf{m}) \cdot \mathbf{g}\|^2$$

Expanding in terms of dot-products gives:

$$D_{\mathbf{f},\mathbf{g}}^2(p) = \langle \rho_p(\mathbf{f}), \rho_p(\mathbf{f}) \rangle - 2\langle \rho_p(\mathbf{f}), \rho_p(\mathbf{m}) \cdot \mathbf{g} \rangle + \langle \rho_p(\mathbf{m}) \cdot \mathbf{g}, \rho_p(\mathbf{m}) \cdot \mathbf{g} \rangle$$



Pattern Matching

$$D_{\mathbf{f},\mathbf{g}}^2(p) = \boxed{\langle \rho_p(\mathbf{f}), \rho_p(\mathbf{f}) \rangle} - 2\langle \rho_p(\mathbf{f}), \rho_p(\mathbf{m}) \cdot \mathbf{g} \rangle + \langle \rho_p(\mathbf{m}) \cdot \mathbf{g}, \rho_p(\mathbf{m}) \cdot \mathbf{g} \rangle$$

Since the representation is unitary:

$$\langle \rho_p(\mathbf{f}), \rho_p(\mathbf{f}) \rangle = \|\mathbf{f}\|^2$$



Pattern Matching

$$D_{\mathbf{f},\mathbf{g}}^2(p) = \|\mathbf{f}\|^2 - \boxed{2\langle \rho_p(\mathbf{f}), \rho_p(\mathbf{m}) \cdot \mathbf{g} \rangle} + \langle \rho_p(\mathbf{m}) \cdot \mathbf{g}, \rho_p(\mathbf{m}) \cdot \mathbf{g} \rangle$$

Since \mathbf{m} is real-valued, we can move it to the other side of the dot-product:

$$\langle \rho_p(\mathbf{f}), \rho_p(\mathbf{m}) \cdot \mathbf{g} \rangle = \langle \rho_p(\mathbf{m}) \cdot \rho_p(\mathbf{f}), \mathbf{g} \rangle$$

Since the representation commutes with point-wise multiplication:

$$= \langle \rho_p(\mathbf{m} \cdot \mathbf{f}), \mathbf{g} \rangle$$

And since \mathbf{m} is one whenever \mathbf{f} is non-zero:

$$= \langle \rho_p(\mathbf{f}), \mathbf{g} \rangle$$



Pattern Matching

$$D_{\mathbf{f},\mathbf{g}}^2(p) = \|\mathbf{f}\|^2 - 2\langle \rho_p(\mathbf{f}), \mathbf{g} \rangle + \boxed{\langle \rho_p(\mathbf{m}) \cdot \mathbf{g}, \rho_p(\mathbf{m}) \cdot \mathbf{g} \rangle}$$

Since \mathbf{m} and \mathbf{g} are real-valued, we can move them to the other sides of the dot-product:

$$\begin{aligned}\langle \rho_p(\mathbf{m}) \cdot \mathbf{g}, \rho_p(\mathbf{m}) \cdot \mathbf{g} \rangle &= \langle \rho_p(\mathbf{m}) \cdot \rho_p(\mathbf{m}), \mathbf{g} \cdot \mathbf{g} \rangle \\ &= \langle \rho_p(\mathbf{m} \cdot \mathbf{m}), \mathbf{g} \cdot \mathbf{g} \rangle\end{aligned}$$

Since \mathbf{m} is strictly 0 or 1, we have:

$$\mathbf{m} = \mathbf{m} \cdot \mathbf{m}$$

So that:

$$\langle \rho_p(\mathbf{m}) \cdot \mathbf{g}, \rho_p(\mathbf{m}) \cdot \mathbf{g} \rangle = \langle \rho_p(\mathbf{m}), \mathbf{g} \cdot \mathbf{g} \rangle$$



Pattern Matching

$$D_{\mathbf{f},\mathbf{g}}^2(p) = \|\mathbf{f}\|^2 - 2\langle \rho_p(\mathbf{f}), \mathbf{g} \rangle + \langle \rho_p(\mathbf{m}), \mathbf{g}^2 \rangle$$

Or somewhat more cleanly:

$$D_{\mathbf{f},\mathbf{g}}^2(p) = \|\mathbf{f}\|^2 - 2\boxed{\mathbf{f} \star \mathbf{g}} + \boxed{\mathbf{m} \star \mathbf{g}^2}$$

The correlation

The windowed square norm



Pattern Matching

$$D_{f,g}^2(p) = \|f\|^2 - 2f \star g + m \star g^2$$

