## ECE 417, Lecture 6: Discrete Cosine Transform

Mark Hasegawa-Johnson

9/6/2019

## Outline

- DCT
- KNN
- How to draw the contour plots of a multivariate Gaussian pdf


## Discrete Cosine Transform

- Last time: PCA
- Why it's useful: PCs are uncorrelated with one another, so you can keep just the top-N (for $\mathrm{N} \ll \mathrm{D}$ ), and still get a pretty good nearestneighbor classifier.
- Why it's difficult: PCA can only be calculated when you've already collected the whole dataset.
- Question: can we estimate what the PCA will be in advance, before we have the whole dataset? For example, what are the PC axes for the set of "all natural images"?


## A model of natural images

Image made of random rectangles

1. Choose an object of a random color,
2. Make it a random size,
3. Position it at a random location in the image,
4. Repeat.


## Result: PCA = DFT!

Define the 2D DFT, $X\left[k_{1}, k_{1}\right]$, as

$$
\sum_{n_{1}=0}^{N_{1}-1} \sum_{n_{2}=0}^{N_{2}-1} x\left[n_{1}, n_{1}\right] e^{-j \frac{2 \pi k_{1} n_{1}}{N_{1}}} e^{-j \frac{2 \pi k_{2} n_{2}}{N_{2}}}
$$

It turns out that the pixels, $x\left[n_{1}, n_{1}\right]$, are highly correlated with one another (often exactly the same!)

But on average, as \# images $\rightarrow \infty$, the DFT coefficients $X\left[k_{1}, k_{1}\right.$ ] become uncorrelated with one another (because object sizes are drawn at random).

Image made of random rectangles


## 2D DFT as a vector transform

- Suppose we vectorize the image, for example, in raster-scan order, so that

$$
\vec{x}=\left[\begin{array}{c}
x[0,0] \\
x[0,1] \\
\vdots \\
x\left[N_{1}-1, N_{2}-1\right]
\end{array}\right]
$$

- ... and suppose we invent some mapping from $k$ to ( $k_{1}, k_{2}$ ), for example, it could be in diagonal order:
$0:(0,0), 1:(1,0), 2:(0,1), 3:(2,0), 4:(1,1), 5:(0,2), 6:(3,0), \cdots$. Then the features are $y_{k}=\vec{x}^{T} \vec{v}_{k}$, with basis vectors

$$
\vec{v}_{k}=\left[\begin{array}{c}
v_{0 k} \\
\vdots \\
v_{N_{1} N_{2}-1, k}
\end{array}\right], v_{n k}=e^{-j \frac{2 \pi k_{1} n_{1}}{N_{1}}} e^{-j \frac{2 \pi k_{2} n_{2}}{N_{2}}}
$$

## The problem with DFT...

... is that it's complex-valued! That makes it hard to do some types of statistical analysis and machine learning (some types of derivatives, for example, do not have a definition if the variable is complex-valued).

$$
\vec{v}_{k}=\left[\begin{array}{c}
v_{0 k} \\
\vdots \\
v_{N_{1} N_{2}-1, k}
\end{array}\right], v_{n k}=e^{-j \frac{2 \pi k_{1} n_{1}}{N_{1}}} e^{-j \frac{2 \pi k_{2} n_{2}}{N_{2}}}
$$

## How to make the DFT real

The DFT of a real symmetric sequence is real \& symmetric.

$$
\begin{aligned}
& x[n]=x^{*}[N-n] \leftrightarrow \operatorname{Im}\{X[k]\}=0 \\
& \operatorname{Im}\{x[n]\}=0 \leftrightarrow X[k]=X^{*}[N-k]
\end{aligned}
$$

## How to make the DFT real

- Most natural images are realvalued.
- Let's also make it symmetric: pretend that the observed image is just $1 / 4$ of a larger, mirrored image.



## Discrete Cosine Transform

Define $s[m]=\left\{\begin{array}{r}\frac{1}{2} x\left[m-\frac{1}{2}\right], m=\frac{1}{2}, \frac{3}{2}, \cdots, N-\frac{1}{2} \\ \frac{1}{2} x\left[2 N-m-\frac{1}{2}\right], m=N+\frac{1}{2}, N+\frac{3}{2}, \cdots, 2 N-\frac{1}{2}\end{array}\right.$

Then:

## 2D DCT as a vector transform

Assume that you have some reasonable mapping from $n$ to $\left(n_{1}, n_{2}\right)$, and from $k$ to $\left(k_{1}, k_{2}\right)$. Then $\vec{y}=V^{T} \vec{x}$, where $V=\left[\vec{v}_{0}, \cdots, \vec{v}_{N_{1} N_{2}-1}\right]$, and

$$
\vec{v}_{k}=\left[\begin{array}{c}
v_{0 k} \\
\vdots \\
v_{N_{1} N_{2}-1, k}
\end{array}\right], v_{n k}=\cos \left(\frac{\pi k_{1}\left(n_{1}+\frac{1}{2}\right)}{N_{1}}\right) \cos \left(\frac{\pi k_{2}\left(n_{2}+\frac{1}{2}\right)}{N_{2}}\right)
$$

## Basis Images: 9th_order 2D DCT

$$
\cos \left(\frac{\pi k_{1}\left(n_{1}+\frac{1}{2}\right)}{N_{1}}\right) \cos \left(\frac{\pi k_{2}\left(n_{2}+\frac{1}{2}\right)}{N_{2}}\right)
$$

- The $k 1=0, k 2=0$ case represents the average intensity of all pixels in the image.
- The k1=1 or k2=1 basis vectors capture the brightness gradient from top to bottom, or from left to right, respectively.
- The $k 1=2$ or $k 2=2$ basis vectors capture the difference in pixel intensity between the center vs. the edges of the image.


Basis image 1


Basis image 2


## Nearest neighbors: $9^{\text {th }}$-order 2D DCT

This image shows the four nearest neighbors of "Image 0" (Arnold Schwarzenegger) and "Image 47" (Jiang Zemin), calculated using a $9^{\text {th }}$ order 2D DCT.
Neighbors of "Image 0" are dark on the right-hand-side, and in the lowerleft corner.
Neighbors of "Image 47" are darker on the bottom than the top.
Neither of these features captures person identity very well...


## Basis Images: 36th-order 2D DCT

$$
\cos \left(\frac{\pi k_{1}\left(n_{1}+\frac{1}{2}\right)}{N_{1}}\right) \cos \left(\frac{\pi k_{2}\left(n_{2}+\frac{1}{2}\right)}{N_{2}}\right)
$$

With a $36^{\text {th }}$ order DCT (up to $k 1=5, k 2=5$ ), we can get a bit more detail about the image.


Basis image 1


Basis image 2


## Nearest neighbors: $36^{\text {th }}$-order 2D DCT

The 36 order DCT is, at least, capturing the face orientation: most of the images considered "similar" are at least looking in the same way.

Jiang Zemin seems to be correctly identified ( 2 of the 4 neighbors are the same person), but Arnold Schwarzenegger isn't (each of the 4 "similar" images shows a different person!)


## PCA vs. DCT

PCA is like DCT in some ways. In this example, $\vec{v}_{0}$ might be measuring average brightness; $\vec{v}_{1}$ is left-to-right gradient; $\vec{v}_{2}$ is measuring center-vs-edges.

But PCA can also learn what's important to represent sample covariance of the given data. For example, eyeglasses ( $\vec{v}_{5}, \vec{v}_{6}$ ), short vs. long nose ( $\vec{v}_{5}$ ), narrow vs. wide chin ( $\vec{v}_{7}$ ).


Basis image 1


Basis image 2


## Nearest neighbors: $9^{\text {th }}$-order PCA

For these two test images, $9^{\text {th }}$-order PCA has managed to identify both people.
Two of the four neighbors of "Image 0 " are Arnold Schwarzenegger.
Three of the four neighbors of "Image 47 " are Jiang Zemin.


## High-order PCA might be just noise!

It is not always true that PCA outperforms DCT. Especially for higherdimension feature vectors, PCA might just learn random variation in the training dataset, which might not be useful for identifying person identity.


## Summary

- As M $\rightarrow \infty$, PCA of randomly generated images $\rightarrow$ DFT
- DCT = half of the real symmetric DFT of a real mirrored image.
- As order of the DCT grows, details of the image start to affect its nearest neighbor calculations, allowing it capture more about person identity.
- PCA can pick out some details with smaller feature vectors than DCT, because it models the particular problem under study (human faces) rather than a theoretical model of all natural images.
- With larger feature vectors, PCA tends to learn quirks of the given dataset, which are usually not useful for person identification. DCT is a bit more robust (maybe because it's like using $\mathrm{M} \rightarrow \infty$ ).


## Outline

- DCT
- KNN
- How to draw the contour plots of a multivariate Gaussian pdf


## K-Nearest Neighbors (KNN) Classifier

1. To classify each test token, find the $K$ training tokens that are closest.
2. Look up the reference labels (known true person IDs) of those K neighbors. Let them vote. If there is a winner, then use that person ID as the hypothesis for the test token.

- If there is no winner, then fall back to 1 NN .


## Confusion Matrix

## Hypothesis



## Accuracy, Recall, and Precision

Accuracy:

$$
A=\frac{\sum_{r=0}^{3} C(r \mid r)}{\sum_{r=0}^{3} \sum_{h=0}^{3} C(h \mid r)}=\frac{\# \text { correct }}{\# \text { data }}
$$

Recall:

$$
R=\frac{1}{4} \sum_{r=0}^{3} \frac{C(r \mid r)}{\sum_{h=0}^{3} C(h \mid r)}=\frac{1}{4} \sum_{r=0}^{3} \frac{\# \text { times } r \text { correctly recognized }}{\# \text { times } r \text { presented }}
$$

Precision:

$$
P=\frac{1}{4} \sum_{h=0}^{3} \frac{C(h \mid h)}{\sum_{r=0}^{3} C(h \mid r)}=\frac{1}{4} \sum_{h=0}^{3} \frac{\# \text { times } h \text { correctly recognized }}{\# \text { times } h \text { guessed }}
$$

## Outline

- DCT
- KNN
- How to draw the contour plots of a multivariate Gaussian pdf


## The Multivariate Gaussian probability density function

If the dimensions of $\vec{x}$ are jointly Gaussian, then we can write their joint probability density function (pdf) as

$$
f_{\vec{X}}(\vec{x})=\mathcal{N}(\vec{x} ; \vec{\mu}, R)=\frac{1}{|2 \pi R|^{1 / 2}} e^{\frac{1}{2}(\vec{x}-\vec{\mu})^{T} R^{-1}(\vec{x}-\vec{\mu})}
$$

The exponent is sometimes called the Mahalanobis distance (with weight matrix $R$ ) between $\vec{x}$ and $\vec{\mu}$ (named after Prasanta Chandra Mahalanobis, 1893-1972):

$$
d_{R}^{2}(\vec{x}, \vec{\mu})=(\vec{x}-\vec{\mu})^{T} R^{-1}(\vec{x}-\vec{\mu})
$$

## Contour lines of a Gaussian pdf

The contour lines of a Gaussian pdf are the lines of constant Mahalanobis distance between $\vec{x}$ and $\vec{\mu}$. For example:
$\frac{f_{\vec{X}}(\vec{x})}{f_{\vec{X}}(\vec{\mu})}=e^{-\frac{1}{2}}$ when $1=d_{R}^{2}(\vec{x}, \vec{\mu})$ which happens when

$$
1=d_{R}^{2}(\vec{x}, \vec{\mu})=(\vec{x}-\vec{\mu})^{T} R^{-1}(\vec{x}-\vec{\mu})
$$

$\frac{f_{\vec{X}}(\vec{x})}{f_{\vec{X}}(\vec{\mu})}=e^{-2}$ when $4=d_{R}^{2}(\vec{x}, \vec{\mu})$ which happens when

$$
4=d_{R}^{2}(\vec{x}, \vec{\mu})=(\vec{x}-\vec{\mu})^{T} R^{-1}(\vec{x}-\vec{\mu})
$$



## Inverse of a positive definite matrix

The inverse of a positive definite matrix is:

$$
R^{-1}=V \Lambda^{-1} V^{T}
$$

Proof:

$$
R R^{-1}=V \Lambda V^{T} V \Lambda^{-1} V^{T}=V \Lambda \Lambda^{-1} V^{T}=V V^{T}=I
$$

So

$$
\begin{gathered}
d_{R}^{2}(\vec{x}, \vec{\mu})=(\vec{x}-\vec{\mu})^{T} R^{-1}(\vec{x}-\vec{\mu})=(\vec{x}-\vec{\mu})^{T} V \Lambda^{-1} V^{T}(\vec{x}-\vec{\mu}) \\
=\vec{y}^{T} \Lambda^{-1} \vec{y}
\end{gathered}
$$

## Facts about ellipses

The formula

$$
1=\vec{y}^{T} \Lambda^{-1} \vec{y}, \text { where } \vec{y}=V^{T}(\vec{x}-\vec{\mu})
$$

... or equivalently

$$
1=\frac{y_{0}{ }^{2}}{\lambda_{0}}+\cdots+\frac{y_{D-1}{ }^{2}}{\lambda_{D-1}} \quad \text { where } y_{d}=\vec{v}_{d}{ }^{T}(\vec{x}-\vec{\mu})
$$


... is the formula for an ellipsoid (in 2D, an ellipse).

- The principal axes are the vectors $\vec{v}_{0}, \vec{v}_{1}, \ldots$
- The radius of the ellipse in the $\vec{v}_{d}$ direction is $\sqrt{\lambda_{d}}$
- If $\lambda_{0}=\lambda_{1}=\cdots$, then it's a circle.


## Example

Suppose that $x_{1}$ and $x_{2}$ are linearly correlated Gaussians with means 1 and -1 , respectively, and with variances 1 and 4 , and covariance 1 .

$$
\vec{\mu}=\left[\begin{array}{c}
1 \\
-1
\end{array}\right]
$$

Remember the definitions of variance and covariance:

$$
\begin{gathered}
\sigma_{1}^{2}=E\left[\left(x_{1}-\mu_{1}\right)^{2}\right]=1 \\
\sigma_{2}^{2}=E\left[\left(x_{2}-\mu_{2}\right)^{2}\right]=4 \\
\rho_{12}=\rho_{21}=E\left[\left(x_{1}-\mu_{1}\right)\left(x_{2}-\mu_{2}\right)\right]=1
\end{gathered}
$$

$$
R=\left[\begin{array}{ll}
1 & 1 \\
1 & 4
\end{array}\right]
$$

## Example

We have that

$$
R=\left[\begin{array}{ll}
1 & 1 \\
1 & 4
\end{array}\right]
$$

We get the eigenvalues from the determinant equation: $|R-\lambda I|=$ $(1-\lambda)(4-\lambda)-1=\lambda^{2}-5 \lambda+3$ which equals zero for $\lambda=\frac{5 \pm \sqrt{13}}{2}$.
We get the eigenvectors by solving $\lambda \vec{v}=R \vec{v}$, which gives

$$
\vec{v}_{1} \propto\left[\begin{array}{c}
1 \\
\frac{3+\sqrt{13}}{2}
\end{array}\right], \vec{v}_{2} \propto\left[\begin{array}{c}
1 \\
\frac{3-\sqrt{13}}{2}
\end{array}\right]
$$

## Example

So the principal axes of the ellipse are in the directions

$$
\begin{aligned}
& \vec{v}_{1} \propto\left[\begin{array}{c}
1 \\
\frac{3+\sqrt{13}}{2}
\end{array}\right], \\
& \vec{v}_{2} \propto\left[\frac{3-\sqrt{13}}{2}\right]
\end{aligned}
$$



## Summary

In fact, it's useful to talk about $R$ in this way:

- The first principal component, $y_{1}$, is the part of $(\vec{x}-\vec{\mu})$ that's in the $\vec{v}_{1}$ direction. It has a variance of $\lambda_{1}$.
- The second principal component, $y_{2}$, is the part of $(\vec{x}-\vec{\mu})$ that's in the $\vec{v}_{2}$ direction. It has a variance of $\lambda_{2}$.
- The principal components are uncorrelated with each other.
- If $\vec{x}$ is Gaussian, then $y_{1}$ and $y_{2}$ are
 independent Gaussian random variables.

