Computational Vision U. Minn. Psy 5036 Daniel Kersten Lecture 10: Image processing

Initialize

Read in Add-in packages:

```
In[8]= Off[General::"spell1"];
SetOptions[ArrayPlot, ColorFunction → "GrayTones",
DataReversed → True, Frame → False, AspectRatio → Automatic,
Mesh → False, PixelConstrained → True, ImageSize → Small];
```

The input 64x64 image: face

Cell below is closed to conserve screen space.

```
In[11]:= width = Dimensions[face] [[1]]; size = width;
hsize = width
2;
hfwidth = hsize;
height = Dimensions[face] [[2]];
face;
gface = ArrayPlot[face]
```



Upcoming dates:

Mid-term: Oct. 26th. Study-guide available later this week. Final project outlines due: Nov. 18th.

Final projects

Format

Should be written like a scientific paper. Might require most of the code to be put in appendices. Can use modules you find elsewhere, but preserve copyrights, and reference Your "audience" will be your class peers.

Previously: Multiresolution, spatial fiter models of early visual processing

Single-channel spatial filtering

One assumes a convolution filter of a fixed shape (e.g. the weights of a DOG filter) that gets applied across the whole image to predict an output "neural image".

images can be represented in the fourier domain in terms of: amplitude and phase spectra.

convolution of an image with a filter in the space domain is mathematically equivalent to: InverseFourier[-Fourier[filter]*Fourier[image]].

convolution is used to model: blur, neural images, the linear stage in neural networks.

Multiple spatial frequency channels

Many neurons in primary visual cortex, in particular "simple cells", have receptive fields whose sizes are limited, and are selective for particular spatial frequencies and orientations. This suggests that neurons in primary visual cortex can be grouped into "channels" consisting of neurons at different retinotopic locations, but with similar frequency and orientation tuning. (See classic experiment of Campbell and Robson, 1968)

To model these, one assumes a set of convolution filters, each set has a fixed shape (e.g. the weights of gabor filters describe a fixed template) that get applied across the whole image to predict a set of output "neural images", where each neural image represents information at a particular spatial scale and orientation. The collection of filters, identical except for position, is called a channel.

```
Global vs. local:
```

Sinusoidal basis functions are global filters. Global filtering is not a good model because they give up spatial localization, at the expense of spatial frequency information (the filters would have to extend across the whole space implying large receptive fields).

Multiresolution analysis with local filters

Instead of sinusoidal basis functions, we can filter with localized "gabor function" filters:

```
In[13]:= Clear[Grating, GratingPatch, kern];
```

```
In[14]= Grating[x_,y_,fx_,fy_,phase_] := Cos[(2.0 Pi (fx x + fy y) + phase)];
GratingPatch[x_,y_,fx_,fy_,sig_,phase_] := Exp[-((x)^2 + (y)^2)/(2*sig^2)]*Grating[x,y,fx,
kern[fx_, fy_, sig_,phase_] :=
Table[GratingPatch[x, y, fx, fy, sig,phase], {x, -1, 1, .1}, {y, -1, 1, .1}];
```

The following demo illustrates a single neural image for various choices of spatial frequency and orientation. In addition, the envelope width manipulates "bandwidth"--i.e. more cycles under the gabor means narrow spatial frequency tuning and thus narrow bandwidth. The phase slider moves one from "bar" to "edge" type receptive fields.

```
In[17]:= Manipulate[
GraphicsRow[{
ArrayPlot[kern[fr * Cos[theta], fr * Sin[theta], sig, phase]], ArrayPlot[
ListConvolve[kern[fr * Cos[theta], fr * Sin[theta], sig, phase], face]]}],
{{fr, 1, "radial frequency"}, .1, 2}, {{theta, .4, "orientation"}, 0, Pi},
{{sig, .4, "envelope width"}, .001, 1}, {{phase, 0, "phase"}, .0, Pi/2}]
```



If you wanted to design a simple template to detect (rather than identify) faces based on a convogabor filter, what parameters might you use?

Self-similarity

When the filters have the same shape except for a change of scale $(x \rightarrow \alpha x)$, they are called self-similar. The self - similar idea is important to vision because of the importance of scale sensitive and scale invariant processing. Further, the self - similar aspect of neural filter models bears a close resemblance to the emerging mathematical field of wavelet analysis.

Human statistical efficiency for detecting gabor patches

Burgess, Wagner, Jennings and Barlow (1981) combined the SKE observer and spatial frequency analysis of human vision to find out how efficiently humans detected patterns. They showed in a 1981 Science article that narrowly windowed sinusoids were detected with high efficiency (>70%) when added to static visual noise. Further, these targets were detected more efficiently than disks of light.

You have all the tools to replicate the experiment of Burgess et al. You can compute d' for the ideal observer for signal-known-exactly patterns. And you can generate Gaussian-windowed sinusoids and add them to gaussian white noise. If you measure the percent correct, and convert that to d' for the human observer, you can calculate the absolute efficiency for human detection--and contribute to

answer the question of what the eye sees best.

Watson, Barlow & Robson (1983) found that that a 7 c/deg grating drifting at 4 Hz, (with a narrow gaussian envelope in space and time) was detected more efficiently than other patterns. Further, the quantum efficiency was very low (<0.05%).

Bottom-line: image coding in terms of scale and orientation: A model for human spatial image representation

At each spatial location, project the image onto a collection of basis vectors (i.e. compute the dot product) that span a range of spatial scales and orientations:



In general, these neural models of basis functions may be over-complete, and non-orthogonal. And there may be a range of phases. Above we show only the "sine-phase" or "edge-detectors" of Hubel and Wiesel.

How linear are VI neurons?

There are kinds of non-linearities that have been introduced to refine, or model other types of neural populations: 1) rectification--neuronal firing rate is by definition non-negative. The combination of oncenter, off-center responses can be treated as a theoretical unit to represent negative and positive signal values. 2) adding squared outputs of sine-phase and cosine-phase filter to produce "contrast energy" filters; 3) Contrast normalization, where the output of an otherwise linear neuron is normalized by the "energy" ouputs of ones nearby in space. We'll discuss some of these in later lectures

What is a multiresolution scale/orientation representation good for?

What is the computational significance of a wavelet-like decomposition? Efficient coding?

- -> savings in neurons, or metabolic requirements?
- -> representations for efficient learning, subsequent coding?
- -> analysis of natural image statistics

Analysis of what vision needs to recognize objects, etc..

-> Edge detection?

-> Edge detection at different spatial scales. Combining over spatial scale

Image processing tools

Mathematica has a library of built-in functions for doing image manipulations

See the Basic Image Manipulation, Image Processing & Analysis, and the Image Filtering & Neighborhood Processing guides. Here are some examples:

here are some examples

In[18]:= Image[Reverse[face]]



E.

We've seen a simple blur using BoxMatrix:

 $\begin{array}{l} \text{Out[20]=} & \left\{ \left\{ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9} \right\}, \ \left\{ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9} \right\}, \\ & \left\{ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9} \right\}, \ \left\{ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9} \right\}, \ \left\{ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9} \right\}, \ \left\{ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9}, \ \frac{1}{9} \right\}, \end{array}$

And the built-in Laplacian of a Gaussian filter which can be used to model center-surround, lateral inhibitory filtering:

In[21]:=





In[22]:= ImageAdjust[ifiltered]



What does ImageAdjust do? Try ListPlot[ImageData[ifiltered][[32]]] with and without using ImageAdjust

We can quickly generate the "neural images" for on- and off-center responses:

In[23]:= GraphicsRow[{ifiltered, ImageMultiply[Binarize[ifiltered], ifiltered], ImageMultiply[ColorNegate[Binarize[ifiltered]], ifiltered]}]



The above used a default threshold to separate above mean from below mean responses. You could use: threshold = Mean [Mean [ImageData []]] as the second argument of Binarize[].

Point operations

Various contrast definitions

For simple stimuli, contrast can be defined as:

(Imax - Imin)/(Imax + Imin): Called "Michelsen contrast". Particularly appropriate for gratings, or stimuli with primary luminance peak and valleys in the image.

 $\Delta I/I_{background}$: Often used in psychophysics of small points/disks against a larger background, or temporal increments/decrements (ΔI) relative to a base level ($I_{background}$).

 $\Delta I/I_{mean}$: Gives same number as Michelsen for gratings, when ΔI corresponds with the amplitude of the grating.

For complex stimuli, we could represent contrast in terms of standard statistical measures on luminance. Let the mean luminance: $I_{\text{mean}} = \sum_{x,y} I(x,y)/N$, where N is the number of pixels. The variance of the luminance is: $\frac{\sum_{x,y} (I(x,y)-I_{\text{mean}})^2}{N}$. Then the standard deviation could be used to provide an overall measure of contrast:

$$\sqrt{\frac{\sum_{\mathbf{x},\mathbf{y}} (\mathbf{I}(\mathbf{x},\mathbf{y}) - \mathbf{I}_{mean})^2}{N}}$$

(The formula for the unbiased variance estimate could also used, where N is replaced by N-1).

However, the above standard deviation definition depends on the units (e.g. candelas/meter^2). Further, it is useful to have contrast definitions that come closer to capturing the perceptual aspects of contrast in an image. We begin by defining contrast at a point--called local contrast.

The visual system is relatively insensitive to the mean luminance, suggesting that a useful measure of local contrast is luminance divided by the mean, whose ratio is dimensionless. In addition, the subjective, qualitative difference between "bright" and "dark" suggests positive and negative contrast, respectively. Thus we define local contrast at a point (x,y) as:

$$c(x,y) = \frac{I(x,y) - I_{mean}}{I_{mean}}$$

The mean of c(x,y) is zero by definition, with positive and negative values corresponding to bright and dark relative to the mean. (Note that local contrast can be defined as function of time too, c(x,y,t).) The definition of local contrast at a point raises the question: over what range should I_{mean} be calculated? This question is relevant to the problems of local adaption to light level, and to tone mapping. The default for us will be to take the mean over the whole image.

Given c(x,y), we now calculate a summary measure of contrast. The variance of c(x,y) is the same as the variance of I(x,y)/ I_{mean} and is given by: $\frac{\sum_{x,y} c^2(x,y)}{N}$. The root mean square (r.m.s.) of set of measurements m_i^2 is by definition: $\sqrt{\frac{\sum_i m_i^2}{N}}$. Thus one can define r.m.s. contrast as the square root of the variance of c:

$$\sqrt{\frac{\sum_{x,y} c^2(x,y)}{N}}$$

This definition of r.m.s. contrast provides us with a useful summary measure of overall "contrastiness" for a complex image.

By analogy with physics (substituting space for time) the square of r.m.s contrast is sometimes called "contrast power". And then "contrast energy" is defined as: contrast power x area. In psychophysics, "area" is often measured in squared degrees of visual angle. Contrast energy is a useful measure when one is concerned about how the size (or duration) of an image patch affects its visibility. If x and y are measured in degrees of visual angle, then, $\int c^2(x, y) dxdy$ is the contrast energy, which in the discrete case can be estimated as:

 $\sum_{x,y} c^2(x, y) \times \left(\frac{\text{area}}{N}\right) = \sum_{x,y} c^2(x, y) \Delta x \Delta y$, where Δx and Δy are the dimensions of a pixel in degrees.

we extend the above contrast defir

If we extend the above contrast definitions to time, in terms of general appearance, the apparent contrast of an image doesn't increase with time, so contrast energy (contrast power x duration) is irrelevant. However, over short durations (less than 100 msec or so), human vision approximately integrates contrast power, so contrast energy is a better predictor of threshold. But for longer durations, contrast power is the better predictor.

Side note on the relation of contrast energy of an image to its fourier representation: Parseval's theorem says that the contrast energy of an image is equal to the integral of the square of the amplitude spectrum (called the power spectrum) over spatial frequency. Contrast energy doesn't depend on the phase relationships.

Contrast manipulations

Adjusting contrast (gain=1 leaves image unchanged, gain=0 reduces it to a uniform field):

```
In[24]:= \mu = Mean[Flatten[face]]; gain = 0.045;
Manipulate[
ArrayPlot[gain (face - \mu) + \mu, Mesh \rightarrow False,
Frame \rightarrow False, PlotRange \rightarrow \{Min[face], Max[face]\}],
\{\{gain, 0.045\}, 0, 1\}]
```



Psychophysics and contrast

When measuring human visual sensitivity, it is important to carefully measure and calibrate the image stimuli. Because standard 8-bit computer displays resolve 256 graylevels, it can be useful to convert the stimuli into a range going from 0 to 255. Scale so values are represented as graylevels between 0 and 255:

```
\ln[26]:= \alpha = 255 / (Max[face] - Min[face]);

\beta = -\alpha Min[face];
```

```
ln[28]:= face256 = \alpha face + \beta;
```

- Exercise: Normalize "face" so that it has a mean level of zero, and an r.m.s. contrast of 1. Use ListPlot and Flatten[] to show a scatter plot of the values before and after the scaling.
- Exercise: Given that human contrast sensitivity for a sinewave grating can be as high as 500, could you get a good measure of it using a typical computer graphics screen? (answer: for fine measurements of human contrast sensitivity, 8 bits of intensity is too coarse. 10 to 12 bits is better.)
- Exercise: Produce a negative image by reversing the contrast

Gamma correction

Computer operating systems allow one to adjust for various non-linearities between displays. A typical screen has a non-linear relationship between measured screen intensity and the voltage supplied. A traditional way of summarizing the non-linear relationship is in terms of "gamma": intensity = a x volt-age^gamma. Let's assume that we are using intensity units that range from 0 to 255 and voltage units also going from 0 to 255. intensity = 255^(1-gamma) x voltage^gamma:

```
In[29]:= output[input_, gamma_] := 255^{1-gamma} input^{gamma};
Plot[output[x, 2], {x, 0, 255}, Frame <math>\rightarrow True, ImageSize \rightarrow Small]
Out[30]= 100 \\ 100 \\ 50 \\ 0 \\ 50 \\ 0 \\ 50 \\ 100 \\ 50 \\ 100 \\ 150 \\ 200 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 250 \\ 2
```

The computer's display card has a look-up-table (LUT) that can be loaded with the inverse gamma function to linearize the display. The LUT remaps input values to output voltages so that there is a linear relationship between internal values and screen luminance.

```
In[31]:= Solve[output == 255<sup>(1-gamma)</sup> input<sup>gamma</sup>, input]
```

Solve::ifun:

 $Inverse \, functions are \, being \, used \, by \, Solve, \, so \, some \, solutions may \, not \, be \, found \, \, use \, Reduce \, for \, complete solution information \, \gg \, complete \, solution \, so$

```
\mathsf{Out}[31]= \left\{ \left\{ \texttt{input} \rightarrow \left(255^{-1+\texttt{gamma}} \texttt{output}\right)^{\frac{1}{\texttt{gamma}}} \right\} \right\}
```

One can do silly things too, like a many-to-one mapping:



 $ln[36]:= ArrayPlot[output2[face256], PlotRange \rightarrow \{0, 255\}]$



..but maybe this isn't so silly. What does the fact that you can still see a recognizable form in the above picture tell you about human vision?

Using gamma to do point operations

You can also use the gamma transformation to do non-linear point operations on an image to concentrate some intensity values more than others.



Sigmoidal contrast manipulation

Here is a gain function (called the "logistic function") that manipulates contrast smoothly--a "soft" threshold:

```
 [n[38]:= squash[x_, \mu_, \gamma_] := N[\frac{1}{1 + e^{-\gamma (x-\mu)}}]; 
 Plot[squash[x, 128, 1], \{x, 0, 255\}, Frame \rightarrow True, ImageSize \rightarrow Small] 
 Out[38]= \begin{array}{c} 1.0 \\ 0.8 \\ 0.6 \\ 0.4 \\ 0.2 \\ 0.0 \\ 0 \\ 50 \\ 100 \\ 150 \\ 200 \\ 250 \end{array}
```

```
 \begin{split} & \text{In}[39]:= \text{gain} = 0.045^{;} \ \mu 0 = \text{Mean}[\text{Flatten}[\text{face256}]]; \\ & \text{Manipulate}[\text{GraphicsRow}[ \\ & \left\{ \text{Plot}[\text{squash}[x, \mu, \gamma], \{x, 0, 255\}, \text{Frame} \rightarrow \text{True}, \text{PlotRange} \rightarrow \{\{0, 255\}, \{0, 1\}\}\}, \\ & \text{ArrayPlot}[\text{squash}[(\text{face256} - \mu 0) + \mu 0, \mu, \gamma], \text{PlotRange} \rightarrow \{\text{Min}[\text{face}], \text{Max}[\text{face}]\}] \right\}], \\ & \left\{ \{\mu, 128\}, 0, 255\}, \{\{\gamma, .5\}, 0, 1\} \right\} \end{split}
```



As γ grows, the sigmoid approaches a hard-threshold. Images that are forced to have only two values are called "Mooney images".

We can make Mooney images more directly using a function that takes an image and sets pixels bigger than τ to 255, and if less than (or equal to) τ , to 0:

```
In[41]:= Mooney[image_, τ_] := Map[If[# > τ, 255, 0] &, image, {2}];
Image[Mooney[face256, 32]]
```

Out[42]=

See also: Binarize[], ColorQuantize[]. And the famous Dalmation dog in perception books.

See Moore and Engel, 2001, and Hegdé J & Kersten D. (2007) for applications of Mooney images to studying vision.

Gray levels are effectively quantized at low light levels. Mooney images mimic this effect but at high light levels.

Exercise: Write a function that quantizes an image to a set of gray levels specified by a set of thresholds:

 $\tau_1, \tau_2, \tau_3, ..., \tau_{N-1}$. Set N=3. (Try using Which[]).

Simple statistics

First-order, i.e. don't take into account relations between pixels

Mean, variance, r.m.s. contrast

```
In[43]:= \mu = Mean[Flatten[face]];

\sigma = Sqrt[Variance[Flatten[face]]];
```

r.m.s. contrast can be calculated as:

```
In[45]:= Sqrt[Variance[Flatten[face]]] / Mean[Flatten[face]]
```

Out[45]= 0.600059

Histograms

For images:



You can tell that the image is quantized at a coarse level (less than 4 bits).

Alternatively, you could calculate the histogram with more basic functions. To do the pattern match below, the floating point numbers are first converted to integers using Round[]:

```
In[48]:= domain = Range[0, 255];
```

```
Freq = Map[Count[Round[Flatten[face256]], #] &, domain];
```

If we normalize the histogram so that the sum is one, then we have a probability:



Getting regions of images

Three ways to pull out a region:

```
In[51]:= ArrayPlot[Take[face256, {1, 32}, {1, 64}]]
```



Or to get rows 1 to 32 and columns 1 to 64 you can use:

```
In[52]:= ArrayPlot[face256[[1;; 32, 1;; 64]]]
```



The above method is particularly useful to learn because it has parallels in both python/numpy and matlab.

We can also use the built-in image function ImageTake[] to get rows 32 through 64:



 Use Manipulate[] and ImageTake[] to make an interactive selection tool for picking out rectangular regions of arbitrary size

Getting coordinates by hand

ln[54]:= ArrayPlot[face256, PixelConstrained \rightarrow {1, 1}]





In[55]:=

Out[55]=



Click on the image above to select it. Now bring up the Drawing Tools, under the Graphics menu. Now use the "cross hairs" tool to select the first the lower left point $\{x0,y0\}$, and then the upper right point $\{x1,y1\}$ as the corners of the rectangular patch that you want. Copy and paste in the argument of the Round[] cell below. Here are coordinates for diagonal points for the eyes.

In[56]:=

```
Round[{{6, 36.4}, {52.4, 48.4}}];
Reverse[Transpose[%]];
ArrayPlot[Take[face256, %[[1]], %[[2]]]]
```



And...finally

Click out the output of Image[]. This will bring up a whole set of image processing tools





Geometric image manipulations using function interpolation

Compare the plots with InterpolationOrder \rightarrow 0 and InterpolationOrder \rightarrow 1.

```
In[60]:= faceFunction =
```

```
ListInterpolation[Transpose[face], \{\{-1, 1\}, \{-1, 1\}\}, InterpolationOrder \rightarrow 1];
\texttt{DensityPlot[faceFunction[x, y], \{x, -1, 1\}, \{y, -1, 1\}, \texttt{PlotPoints} \rightarrow 256,}
 \texttt{Mesh} \rightarrow \texttt{False}, \texttt{AspectRatio} \rightarrow \texttt{Automatic}, \texttt{Frame} \rightarrow \texttt{None}, \texttt{ColorFunction} \rightarrow \texttt{"GrayTones"} ]
```





Now you can calcualte intensities "between pixels":

```
In[62]:= faceFunction[.1, .23]
```

Out[62]= 0.437402

Morphing

```
In[63]:= faceFunction = ListInterpolation[Transpose[face], {{-1, 1}, {-1, 1}}];
```

```
\begin{aligned} & \text{In}_{[64]:=} \text{ DensityPlot} \Big[ \text{faceFunction} \Big[ \text{Sign} [x] \ x^2, \ \text{Sign} [y] \ y^2 \Big], \ \{x, -1, 1\}, \\ & \{y, -1, 1\}, \ \text{PlotPoints} \rightarrow 100, \ \text{Mesh} \rightarrow \text{False}, \ \text{AspectRatio} \rightarrow \text{Automatic}, \\ & \text{Frame} \rightarrow \text{None}, \ \text{ColorFunction} \rightarrow \text{"GrayTones"}, \ \text{ImageSize} \rightarrow \text{Small} \Big] \end{aligned}
```



- FINAL project idea: How does our ability to recognize objects degrade with geometrical deformations?
- FINAL project idea: Use the log polar model of the retinotopy of primary visual cortex to show how the foveal representation gets expanded on cortex.

More filtering: Calculating the spatial gradient of an image using function interpolation

As we will see later, both first and second spatial derivatives have been used to model edge detection/amplification by neural processes in primary visual cortex.

Hubel and Wiesel's "edge detector" can be interpreted as an approximation to a first order derivative ("gradient" in 2D, see ∇ f), and their "bar detector" as a second derivative (see ∇ ²G below).

Here we show how to use symbolic derivatives to find the gradient.

The gradient of an image intensity function f, ∇f , has a maximum value in the direction of greatest change.b

⊽f =	$\left(\frac{\partial \mathbf{f}}{\partial \mathbf{x}},\right)$	$\frac{\partial \mathbf{f}}{\partial \mathbf{y}} \bigg)$	= \	$\left(\frac{\partial \mathbf{f}}{\partial \mathbf{x}}\right)^2,$	$\left(\frac{\partial \mathbf{f}}{\partial \mathbf{y}}\right)^{2}$
--------	---	--	-----	---	--

Let filterface =f, where we've blurred out face a little to reduce quantization artifacts:

```
In[66]:= faceFunction = ListInterpolation[Transpose[filterface], {{-1, 1}, {-1, 1}}];
```

```
In[67]:= Clear[x, y]
```

```
In[68]:= nx[x_, y_] := Evaluate[D[faceFunction[x, y], x]];
ny[x_, y_] := Evaluate[D[faceFunction[x, y], y]];
ImageGradient[x_, y_] := Evaluate[Sqrt[D[nx[x, y], x]^2 + D[ny[x, y], y]^2]];
```

Plot the rate of change in the x-direction:

```
In[71]:= temp = Table[nx[x, y], {x, -1, 1, .005}, {y, -1, 1, .005}];
ArrayPlot[Transpose[temp]]
```



You can use VectorPlot to visualize the gradient field:

```
[n[73]:= VectorPlot[{nx[x, y], ny[x, y]}, {x, -1, 1}, {y, -1, 1}, ImageSize \rightarrow Small]
```



Plot the magnitude of the gradient to highlight regions of the image where contrast is changing the most rapidly:

```
ln[74]:= DensityPlot[ImageGradient[x, y], \{x, -1, 1\}, \{y, -1, 1\}, PlotPoints \rightarrow width, Mesh \rightarrow False, Frame \rightarrow False, ColorFunction \rightarrow "Rainbow", ImageSize \rightarrow Small]
```



Out[74]=

Discrete spatial derivative filtering with built-in image functions

Use can also use image specific filters. GradientFilter[image,r] gives an image corresponding to the magnitude of the gradient of image, using discrete derivatives of a Gaussian of pixel radius r.



Out[75]=



The $\nabla^2 G$ operator (one of the forms used to represent a "mexican hat" filter) in the previous lecture first convolves the image with a Gaussian blur kernel, and then takes the Laplacian ∇^2 . This filter effectively blurs an image before taking the derivatives.

Later we'll see why when we study edge detection.

In[76]:= LaplacianGaussianFilter



2] // ImageAdjust



Manipulating color images

```
In[77]:= image = ExampleData[{"TestImage", "Mandrill"}]
```



- In[78]:= RGBvalues = ImageData[image];
 Dimensions[RGBvalues]
- Out[79]= $\{512, 512, 3\}$
- $In[80]:= \{512, 512, 3\}$
- Out[80]= $\{512, 512, 3\}$
- In[81]:= **{512, 512, 3**}
- Out[81]= $\{512, 512, 3\}$
- In[82]:= **{512, 512, 3**}
- Out[82]= $\{512, 512, 3\}$

- In[83]:=
 {512, 512, 3}

 Out[83]=
 {512, 512, 3}

 In[84]:=
 {512, 512, 3}

 Out[84]=
 {512, 512, 3}
- In[85]:= reds = Map[#[[1]] &, N[RGBvalues], {2}];
 greens = Map[#[[2]] &, N[RGBvalues], {2}];
 blues = Map[#[[3]] &, N[RGBvalues], {2}];
- In[88]:= GraphicsRow[{Image[reds], Image[greens], Image[blues]}]



You can also use the function:

In[90]:= ColorSeparate



A weighted average of RGB values to produce a luminance image:

 $\ln[91] = \text{grayscale} = \text{Map}\left[\frac{0.3 \# [[1]] + 0.59 \# [[2]] + 0.11 \# [[3]]}{255} \&, \text{N}[\text{RGBvalues}], \{2\}\right];$

Image[grayscale] // ImageAdjust ;

Note the weights above are arbitrary, and the chosen values will depend on the color calibration.

Putting the R, G, B images back together:

```
In[93]= r = reds;
g = greens;
b = blues;
temp2 =
Partition[Transpose[{Flatten[r], Flatten[g], Flatten[b]}], Dimensions[r][[2]]];
Image[
temp2]
```



Next time

Efficient coding

References

Adelson, E. H., Simoncelli, E., & Hingorani, R. (1987). Orthogonal Pyramid Transforms for Image Coding. Paper presented at the Proc. SPIE - Visual Communication & Image Proc. II, Cambridge, MA. Barlow, H. B., & Olshausen, B. A. (2004). Convergent evidence for the visual analysis of optic flow through anisotropic attenuation of high spatial frequencies. J Vis, 4(6), 415-426.

Daugman, J. G. (1988). An information-theoretic view of analog representation in striate cortex, Computational Neuroscience. Cambridge, Massachusetts: M.I.T. Press.

Engel, S. A., Glover, G. H., & Wandell, B. A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. Cereb Cortex, 7(2), 181-192.

Gold, J. M., Murray, R. F., Bennett, P. J., & Sekuler, A. B. (2000). Deriving behavioural receptive fields for visually completed contours. Curr Biol, 10(11), 663-666.

Hegdé, J., & Kersten, D. (2010). A Link between Visual Disambiguation and Visual Memory. Journal of Neuroscience, 30(45), 15124–15133. doi:10.1523/JNEUROSCI.4415-09.2010

Konishi, S. M., Yuille, A. L., Coughlan, J. M., & Zhu, S. C. (2003). Statistical edge detection: Learning and evaluating edge cues. IEEE Transactions on Pattern Analysis and Machine Intelligence, 25(1), 57-74.

Olman, C. A., & Kersten, D. (2004). Classification objects, ideal observers & generative models. Cognitive Science, 28, 227-239.

Moore, C., & Engel, S. A. (2001). Neural response to perception of volume in the lateral occipital complex. Neuron, 29(1), 277-286.

Schwartz, E. L. (1980). A quantitative model of the functional architecture of human striate cortex with application to visual illusion and cortical texture analysis. Biol Cybern, 37(2), 63-76.

http://library.wolfram.com/howtos/images/#histograms

© 2008, 2010, 2013, 2015 Daniel Kersten, Computational Vision Lab, Department of Psychology, University of Minnesota. kersten.org