IMAGE PROCESSING (RRY025)

One of the Exams in 2012/2013

1 NOISE REMOVAL / IMAGE PRE-COMPRESSION [15 points]

(a) [4p] What is additive white Gaussian noise? How can you remove it from an image?

(b) [8p] What is Poissonian noise? What are the fundamental differences between this type of noise and additive white Gaussian noise? When does Poissonian noise occur? How can you remove it from an image?

(c) [3p] How would you pre-compress a noisy image in a rigorous way? The noise present in the image is of no interest, and pre-compression should produce as few artifacts as possible! Does your answer depend on the type of noise, or not?

2 MISCELLANEA [15 points]

(a) [5p] A dear friend of yours invites you to her new office. As soon as you enter, you get surprised: far away from the door, there is a nice photo of you! You get flattered and ask: 'Why do you have a photo of me in your office?' She answers: 'Go close to the photo and you will see that you are wrong!' You go close to the photo and get even more surprised: it is a photo of her!! How is it possible?? (Your friend is not a magician, but an expert image processer :-)

(b) [5p] Suppose that you throw a Swedish krona in the air 100 times, and that you write down the result of this experiment in the form of a binary image of size 10×10: when you get 'head', you write 0; when you get 'crown', you write 1. What will be the (approximate) value of the single-pixel entropy for this image? If a person tells you that she has done a similar experiment using another coin, and that the resulting single-pixel entropy is much less than 1, what can you conclude about the coin (or the person)? Help: $p \log_2 p \to 0$ as $p \to 0$.

(c) [2p] You have two images, $A(x, y)$ and $B(x, y)$, that have the same size but represent different objects. Suppose that you do the following. You low-pass filter the first image; call it $A_{\mathrm{LP}}(x, y)$. You high-pass filter the second image; call it $B_{\mathrm{HP}}(x, y)$. You then add $B_{\mathrm{HP}}(x, y)$ to $A_{\mathrm{LP}}(x, y)$. What does the resulting image look like when you watch it from a short distance? And when you watch it from a long distance? Perhaps, I am helping you too much :-)

(d) [3p] In your opinion, what is the most interesting topic of the course? Explain how important this topic is in the context of image processing, and how important it is for your studies/job.

HELP TO SOME OF THE QUESTIONS

1 Noise Removal / Image Pre-Compression

(c) What does pre-compression mean?

What is the idea behind data de-noising?

So data de-noising is a rigorous way to pre-compress noisy data! Isn't it??

→ What do you conclude?

2 Miscellanea

(a) Perhaps, I have helped you too much :-)

See question **(c)**!

But I am sure you want to know more about

HYBRID IMAGES

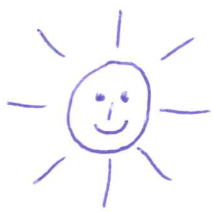

- http://cvcl.mit.edu/publications/Talk-Hybrid_Siggraph06.pdf

  These are presentation slides!

- http://cvcl.mit.edu/publications/Hybrid.mp4

  This is a movie presentation!

- http://cvcl.mit.edu/publications/OlivaTorralb_Hybrid_Siggraph06.pdf

  This is the original paper by Oliva, Torralba & Schyns (2006)

- http://cvcl.mit.edu/hybridimage.htm

  This is their home page with other useful links

HAVE FUN!

(b) Not as simple as it seems .....

- Single-pixel entropy:

$$H_1 = - \sum_{i=0}^{1} p_i \log_2 p_i$$

- What about p_i ?

- The expected probabilities are equal:

$$P_0 = P_1 = \frac{1}{2}$$

- But the observed probabilities can be significantly different, for example $p_0 \ll p_1$!

- Let us quantify this point .....

- What is the probability of getting n_0 heads when you throw the coin N times ?

$$P_B(\underset{\text{variable}}{n_0}, \underset{\text{parameter}}{N}) = \frac{N!}{n_0!\,(N-n_0)!} \left(\frac{1}{2}\right)^{n_0} \left(\frac{1}{2}\right)^{N-n_0}$$

$$= \frac{N!}{n_0!\,(N-n_0)!} \left(\frac{1}{2}\right)^{N}$$

This is the binomial distribution!

- What is the mean of n_0?

$$\mu_0 = \frac{1}{2} N$$

- What is the standard deviation of n_0?

$$\sigma_0 = \frac{1}{2}\sqrt{N}$$

➔ The observed n_0 can be estimated as

$$\mu_0 \pm \sigma_0 = \frac{1}{2} N \left(1 \pm \frac{1}{\sqrt{N}}\right)$$

expected value

uncertainty

- What about the observed p_0?

$$p_0 = \frac{n_0}{N} \approx \frac{\mu_0 \pm \sigma_0}{N} = \frac{1}{2}\left(1 \pm \frac{1}{\sqrt{N}}\right)$$

- And what about the observed p_1?

$$p_1 \approx \frac{1}{2}\left(1 \mp \frac{1}{\sqrt{N}}\right)$$

Why?

$$\rightarrow H_1 \approx -\frac{1}{2}\left(1 \pm \frac{1}{\sqrt{N}}\right) \log_2 \left[\frac{1}{2}\left(1 \pm \frac{1}{\sqrt{N}}\right)\right] - \frac{1}{2}\left(1 \mp \frac{1}{\sqrt{N}}\right) \log_2 \left[\frac{1}{2}\left(1 \mp \frac{1}{\sqrt{N}}\right)\right]$$

$\approx$ ..... compute it by yourself, if you don't believe me ☺ .....

$$\approx 1 - \frac{\log_2 e}{N}$$

- Summarizing:

$$\left|\frac{\Delta p_i}{p_i}\right| \approx \frac{1}{\sqrt{N}} \quad \underline{\text{versus}} \quad \left|\frac{\Delta H_1}{H_1}\right| \approx \frac{1}{N}$$

- What does this mean?

* $N \sim 100 \Rightarrow$ $\begin{cases} \text{uncertainty in } p_i \sim 10\% \\ \text{uncertainty in } H_1 \sim 1\% \end{cases}$

* $N \sim 10 \Rightarrow$ $\begin{cases} \text{uncertainty in } p_i \sim 30\% \\ \text{uncertainty in } H_1 \sim 10\% \end{cases}$

Do you get the point?

$\rightarrow$ H_1 is much more robust than p_i! Even when N is as small as ~ 10, H_1 is very close to 1 !!

- Why is H_1 so robust ?

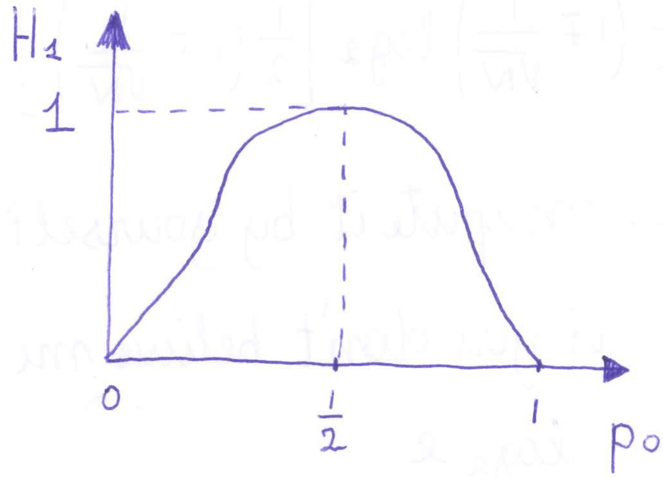

and so ??

→ What do YOU conclude ?