Notes For Computational Photography and Imaging
Class of September 21, 2005

Topic: Dynamic Range and Focus

Basically, this was a re-cap over what the class format should be and also tips and instructions for the student note taker.

However, before we continued, we discussed an anecdote concerning the Renaissance. During the Renaissance, the highest level of art that one could achieve was “Photo realistic Painting”, which was in itself a challenge. However, once cameras became popular, this art form went into a decline throughout the 20th century. When asked if this meant if this meant that this meant for the art world, Pablo Picasso replied: “Now people will understand what art is NOT about”.

Curious note: Photography art was spawned as photo realistic painting declined. Will computational photography spawn a new art form?

Problem: Dynamic Range

1

1500

25,000

400,000

2,000,000,000

The real world is
High dynamic range

What does the eye sees?

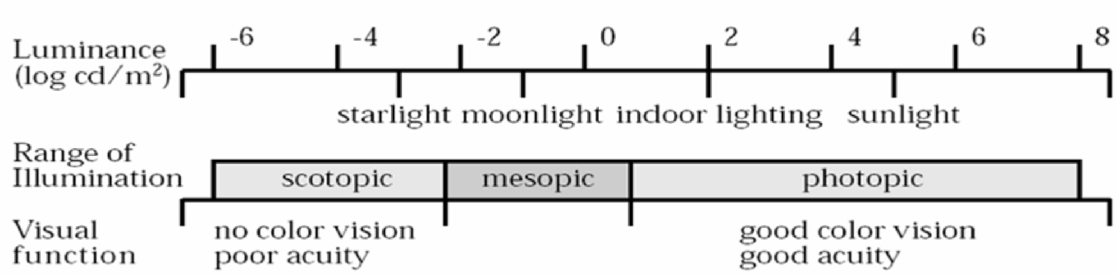

Figure 1: The range of luminances in the natural environment and associated visual parameters. After Hood (1986).

The eye has a huge dynamic range
Do we see a true radiance map?

How do we deal with the problem of Dynamic Range? The real world has a high high radiance to low radiance ration that it is impossible to capture in a single image.

Trick #1: Varying Exposure:

Here, we saw the “Belgian House” slide and saw how through different images with different exposures we can get a more “realistic” image.

But, how can WE, human beings, perceive such high radiance when our retinas can only capture 3 orders of magnitude?

Answer: Our eyes perform local variance and our eyes have local gain, whereas cameras have a global, indiscriminate gain.

Examples:

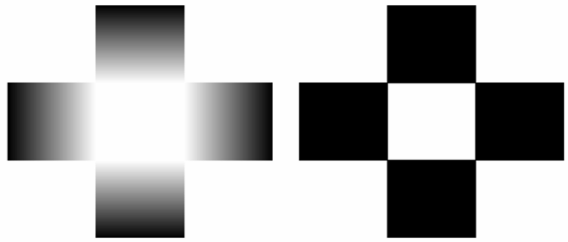

A typical photo taken with a 100ms exposure will only capture from 0-255 of World Intensity of a particular RGB value. If we were to reduce that exposure by a magnitude of 10, say 10 ms, we can capture 0 – 2550 of World Intensity

If we continue to decrease by a factor of ten:

1ms exposure = 0-25,500 World Intensity
.01 ms exposure = 0 – 255,000 World Intensity
ms exposure = 0 – 2.55 M World Intensity

Limitation: The Amount of light registered is very small as the exposure time decreases.

Typically, one could expect the World Intensity of the outdoor world to be around 10 ^ 6 while under a table one could register 10 of World Intensity. This means that a magnitude of 10^4 – 10^6 is a typical room environment

Second problem: Noise.

The Signal to Noise Ratio (SNR) characteristic is such that a signal captured between a 0 – 10 WI is useless and cannot be trusted. However, as more intensity is acquired, the SNR increases until we reach a ratio of 10 SNR (which is what we would expect when we capture something around 255).

So, how can we use this to our advantage? We can create a new image based on the good SNR's captured by each image with a different exposure.

Problem: If we take an image that has a good SNR from 180 – 255 with an exposure of 100ms, how should we decide to use this pixel instead of another taken from another image with a good SNR from 1800 – 2550 in an exposure of 10ms?

Before we answer this, we should note that the camera does not increase its SNR as we increase the intensity of an RGB value captured. Rather, after 240, our SNR decreases rapidly.

Solution: Created weighted average of pictures. Have each picture's pixel assigned a confidence value based upon the intensity registered and the exposure rate of when the picture was taken. Then, average these to create your new image.

Very curious note: All manufacturers will lie as to their SNR

Curious observation by a student: Why not have a comprehensive curve that could describe all of these things?

Curious fact: SMAL = logarithmic sensor that controls how every pixel is captured.

Now, to describe our assignment:
Our dataset's image exposures should vary by a factor of 2 and you must find a solution to overlapping exposures.

Simplest Solution:

Value = Pixel / Exposure weight (Note: There was some confusion as to what the formula was)

Not as simple solution:

Value = (Pixel value/ 255) / Exposure weight

Tougher solution:

Weight = Intensity Value / 255

Value = (Summation of Weights / Intensity) / Exposure weight) / Summation of Weights

Question posed by professor: With a scene of 0 – 1000 and a good camera that can capture from 0 – 100, how many pictures would you need to cover the range?

Also consider how noise affects the image capture process.

Note: Without knowing the SNR, it is impossible to answer the previous question.

Question posed by professor: Given two pictures, how do you find the exposure settings?

One: You can find the relative ratio between both pictures by comparing pixels over 240

Extra credit: Try to figure out the exposure settings of each picture.

Key Point: Final value of image is NOT a direct value of pixel value AND scene radiance MUST be preserved.

2nd Part of the class: Bilateral Filtering.

Overall Concept:

Take an HDR image

- Separate into low frequency and high frequency
- Reduce the quality of the low frequency image
- Multiply them together

Detail = Image / Base

'Unilateral' Filter

Traditional, linear, FIR filters

Key Idea: Convolution

- Output(x) = local weighted avg. of inputs.
- Weights vary within a 'window' of nearby x

Smoothes away details, **BUT** blurs result

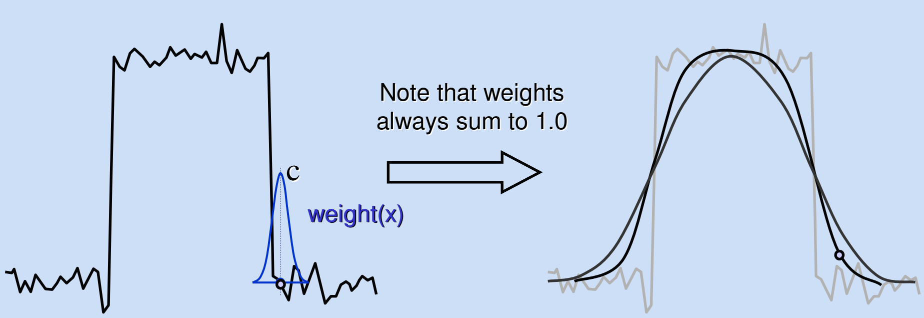

'Unilateral' Filter

Forces a Tradeoff:

-Broad window: better detail removal

-- OR --

-Narrow window: better large structure

But we want BOTH...

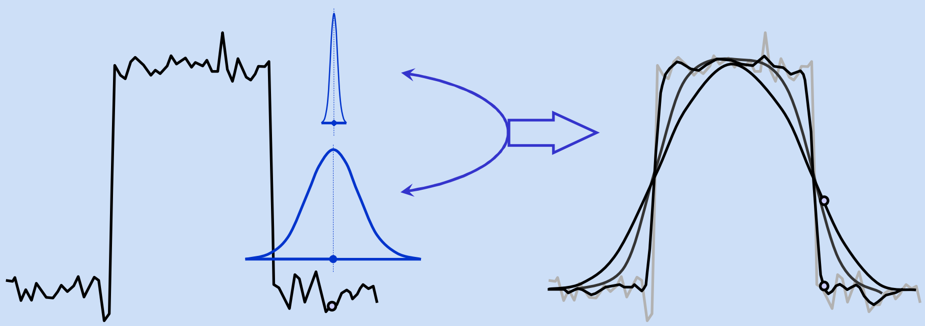

A bilateral filter is basically a double Gaussian Filter, in that it applies a regular Gaussian Filter while at the same time applying a second Gaussian Filter that is relative to a distance chosen as the high value of the image.

Bilateral Filter

A 2-D filter window: weights vary with intensity

[Tomasi&Manduchi1998]

Further Analysis: [Black99] [Elad02] [Durand&Dorsey02], ...

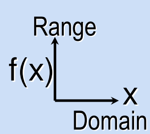

c: distance from input (***domain*** of input)

s: difference from input (***range*** of input)

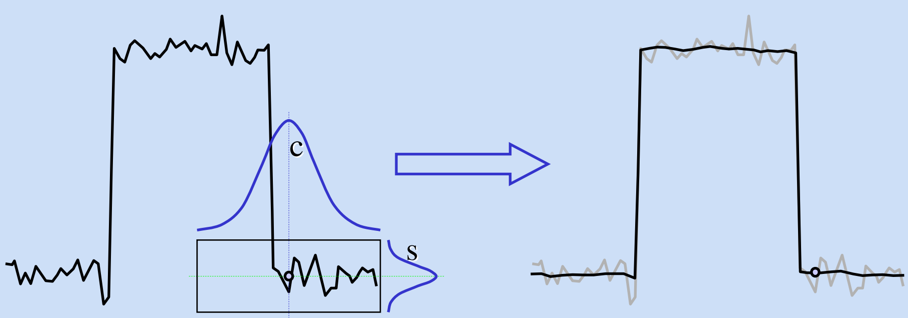

Bilateral Filter

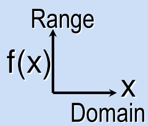

2 Gaussian Weights:
product =
ellisoidal footprint

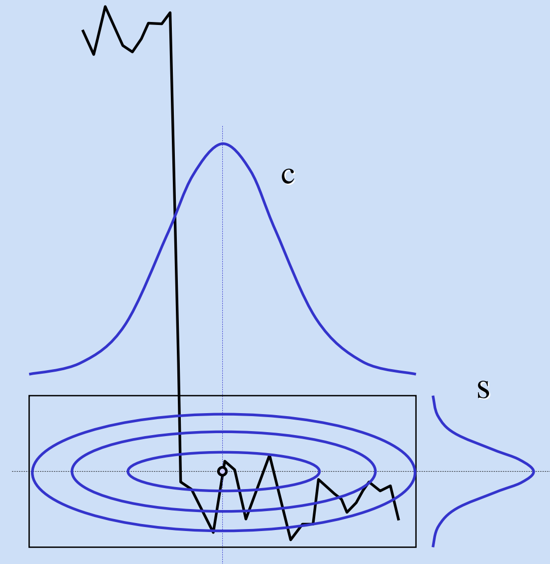

Bilateral Filter

Why it works: graceful segmentation

- Filtering in one region ignores filtering in another
- Gaussian **s** acts as a 'filtered region' finder

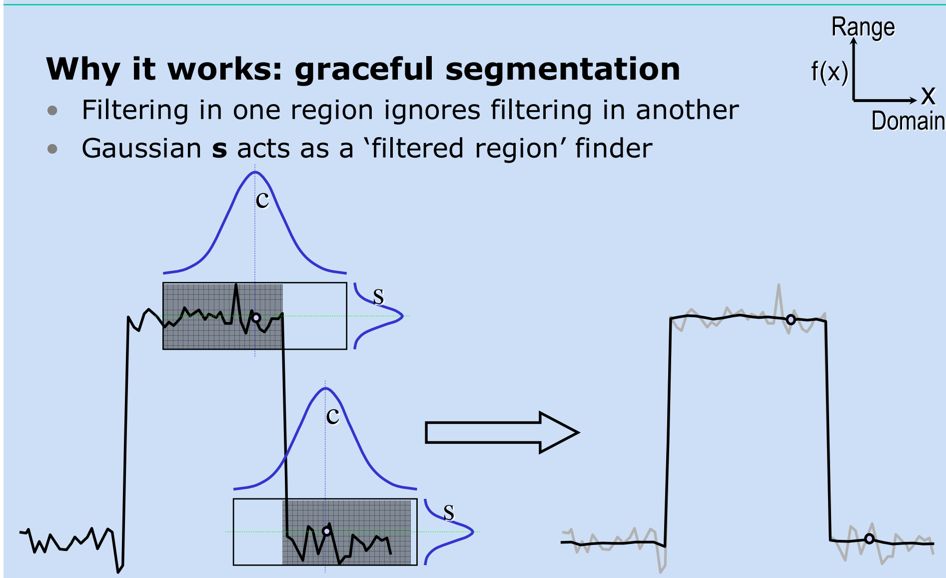

Bilateral Filter: Strengths

Piecewise smooth result

- averages local small details, ignores outliers
- preserves steps, large-scale ramps, and curves,...

- Equivalent to anisotropic diffusion and robust statistics [Black98,Elad02,Durand02]
- Simple & Fast (esp. w/ [Durand02] FFT-based speedup)

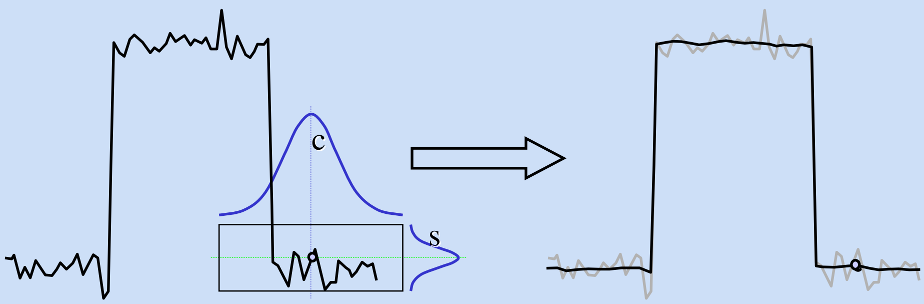

Suppose we have a signal that has a drastic range jump. On its left side, we see that this noisy signal ranges around 100 – 200 Intensity. However, on its right side, there is a drastic jump to a range of 1000 – 2000 Intensity. If we were to just apply a unilateral filter, we would get a smooth incline between these two sections of the signal, instead of two discrete, smooth sections.

A Bilateral filter is able to determine where the highest range of the image occurs, and then as our main unilateral filter approaches this region, it decides that only the left side should be filtered, leaving the right side intact. As the filter trespassed onto the right side of the signal, we apply a filter that only smoothes the right part of the signal, leaving the left side intact.