Sensing Color

Ron Parr
CPS 1/296

Sensors are Photon Counters

- CCD and CMOS sensors register incident photons
- Are photons alike to sensors?
  - Some penetrate deeper into the substrate
  - Some may have different probabilities of detection
- With a few exceptions, sensors are fundamentally B&W devices

Filters

- For our purposes, filters
  - Allow certain wavelengths to pass
  - Reflect or absorb others
- Typical filters absorb/pass
- Dichroic filters reflect/pass
- See laser/filter demo in class

Filters for Color Sensing

- Idea 1: Pick color filters that are identical to human cones
- Idea 2: Pick color filters with sensitivities that are a linear transformation of human cone sensitivity (why is this sufficient?)
- Idea 3: Pick color filters that are "close" to a linear transformation of human cone sensitivity and then optimize "average" performance – try to minimize potential for metameric failure

Implementing Color Filters

- Use 3 sensors with dichroic prisms
- See, e.g., http://www.bockoptronics.ca/pdfs/Jai/CV-L107CL_datasheet_web.pdf
- Each prism:
  - Deflects range of wavelengths to 1 sensor
  - Allows rest to pass
- "3CCD" label seen in high end video cameras

3CCD Pros and Cons

- Pros:
  - Every channel captured at full resolution
- Cons:
  - Size
  - Weight
  - Expense
    - Duplication of sensors
    - Dichroic prisms
  - Alignment
  - Some limitations in filter design

Some Quirks about Human Vision

- Our perception of luminance is biased towards green wavelengths
- We are most sensitive to high frequency spatial variations in luminance
- We are surprisingly insensitive to high frequency spatial variations at the extreme ends of the spectrum
- See subsampling experiments

The Bayer Pattern

- Idea: If luminance (green) is more salient than more extreme wavelengths, allocate more resources to green
- Bayer pattern sensor uses a color filter array on a single sensor

GRGRGRGRGRGR
BGBGBGBGBGBG
GRGRGRGRGRGR
BGBGBGBGBGBG

N.B.: Patent and concept are more general than just RGB

Features of the Bayer Pattern

- Green sampled at 50% B&W sensor resolution
- Red, Blue sampled at 25% B&W sensor resolution
- Single chip color sensor solution
- Most widely used color sensing solution

Pros and Cons of the Bayer Pattern

- Pros:
  - Surprisingly easy to manufacture
  - Good flexibility in filter design
  - Good bang for the buck
- Cons:
  - Must estimate missing color channel components
    - Computation requirements
    - Possibility (reality) of errors in estimation
  - Potential for image degradation

Interpolated Color

- Basic problem:
  - Each pixel sees just one of R,G,B
  - Must estimate other two from nearby pixels
- Many, many, many approaches to this problem (see reading)
- No silver bullet

Bilinear Interpolation

- Simplest interpolation algorithm
- Likely similar to or exactly what your webcam does
- Missing components are averaged from nearest 4 (or two) pixels of the corresponding color
- Can lead to interesting color hallucinations, AKA, aliasing (see example on board)
- Once you are aware of this, you will start noticing it on your TV when guests wear shirts with high frequency patterns!

Color Interpolation Fallout

- No matter how clever the algorithm, there will be some cases where it will make mistakes
- Is there a solution?
- Use an anti-aliasing (low pass) filter in front of the sensor
- Blurs high frequency detail
- Problems:
  - Can be expensive
  - Softens image (effectively reduces resolution)
  - Typically combined with a software "sharpening" filter
- Trade off between optimizing for average case or worst case
  - Average case: Usually sharp images, but some aliasing
  - Worst case: Usually soft images, but no aliasing

Implications for Robotics

- Can't rely upon sensing small red or blue objects, e.g., red laser dots from laser pointers
- Edges of objects may have artifacts
- Effective sensor resolution can be lower than stated resolution

- In general, a X MP Bayer pattern sensor will behave about as well as a X/2 MP 3CCD solution, and possibly as bad as X/4 MP 3CCD solution in the worst case.
- When moving from B&W to color, plan on doubling pixel count

Sanity Check

- Q: Suppose your image is pure B&W, will a Bayer pattern sensor perform comparably to a B&W sensor?

- A: Only if the interpolation algorithm knows that the image is pure B&W

- Why: Interpolation mistakes in color channels imply interpolation mistakes in brightness

Single Chip Alternatives

- Some have proposed exploiting a quirk of photon absorption in the silicon substrate
- P(absorption) = f(wavelength)

- Idea: Layer photodiodes on top of each other
- Patented by Foveon

Idealized X3 Sensor

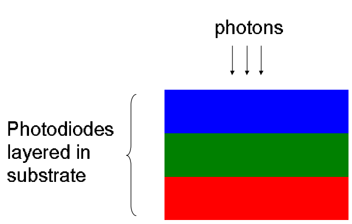

X3 in Practice

- Less control over properties of filters - limited by physics of substrate (see paper)
- New technology, still evolving

- For a given sensor size, Foveon tends to offer ~½ the pixels of Bayer pattern sensors
- Recall that Bayer pattern sensors tend to perform like a ½ resolution 3CCD system for most images

- Difference in practice not yet as large as one might hope

X3 Potential for Robotics

- Many camera systems optimized to produce visually satisfactory illusions
- Image processing often a black box
- Potential to use information that other cameras, and even our own visual system discards

Other Approaches

- Why limit ourselves to 3 channels?
- Uses for more filter types:
  - Compensate for imperfect matches between cone absorption and filter absorption
  - Create multi/hyperspectral sensors – a useful form of metameric failure
- Uses for multi/hyperspectral sensors
  - Infer chemical composition
  - See through camoflauge