Lecture 4: Camera Imaging Pipeline

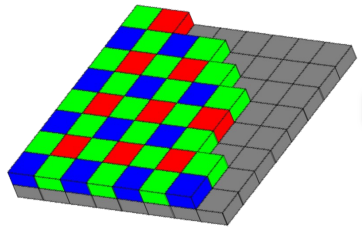
COMP 590/776: Computer Vision

Instructor: Soumyadip (Roni) Sengupta

TA: Mykhailo (Misha) Shvets

A typical color imaging pipeline





Sensor with color filter array (CFA)
(CCD/CMOS)

ISO gain and raw-image processing

RGB Demoasicing

Noise Reduction

White-Balance & Color Space Transform (CIE XYZ)

Color Manipulation (Photo-finishing)

Mapping to sRGB output

JPEG Compression

Save to file

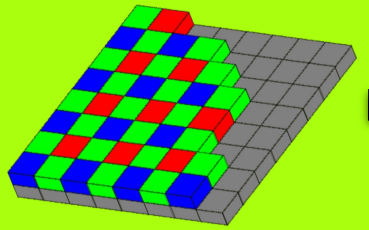
NOTE: This diagram represents the steps applied on a typical consumer camera pipeline. ISPs may apply these steps in a different order or combine them in various ways. A modern camera ISP will undoubtedly be more complex, but will almost certainly implement these steps in some manner.

Integrated signal processor (ISP)

- You will hear the term "ISP" associated with camera pipelines
- An ISP is dedicated hardware used to process the sensor image to produce the final output (JPEG image) that is saved on your device
- The ISP is usually integrated as part of a system on a chip (SoC) that has other modules
- Companies such as Qualcomm, HiSilicon, Intel (and more) sell ISP chips
  - An ISP can be customized by the customer (Samsung, Huawei, LG, Apple, etc)
- Note that it is also possible to perform operations common on an ISP on your device's CPU and GPU

A typical color imaging pipeline





Sensor with color filter array (CFA)
(CCD/CMOS)

ISO gain and raw-image processing

RGB Demoasicing

Noise Reduction

White-Balance & Color Space Transform (CIE XYZ)

Color Manipulation (Photo-finishing)

Mapping to sRGB output

JPEG Compression

Save to file

Camera sensor



CMOS sensor

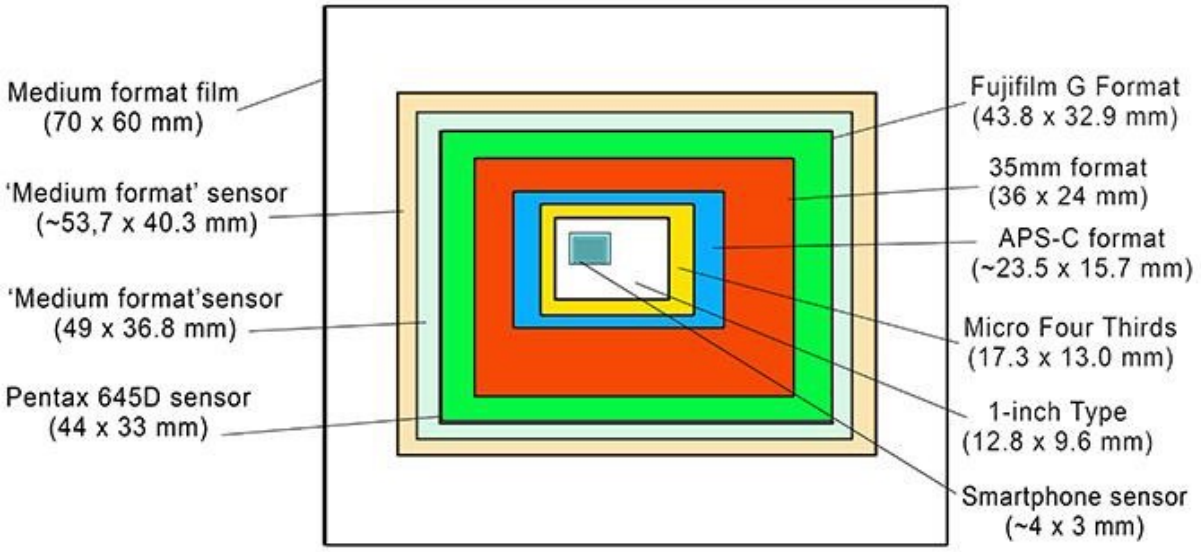
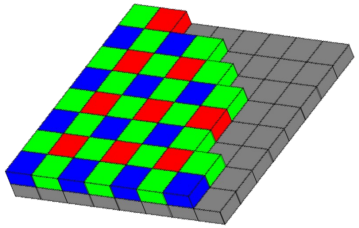


Figure from Photo Review website.

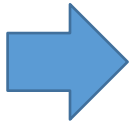
Almost all consumer camera sensors are based on complementary metal-oxide semiconductor (CMOS) technology.

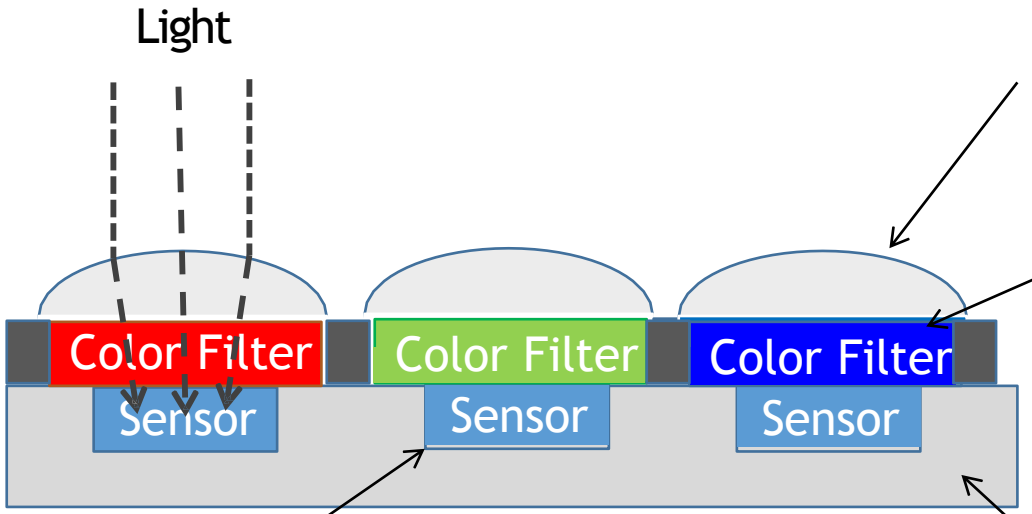
- We generally describe sensors in terms of number of pixels and size.
- The larger the sensor, the better the noise performance as more light can fall on each pixel.
- Smart phones have small sensors!

Camera sensor RGB values



Color filter array
or "Bayer" pattern.





Micro-lenses are placed over the diode to help increase light collection on the sensor

Color filters place over the sensor. This forms a Color Filter Array (CFA) also called a "Bayer Pattern" after inventor Bryce Bayer.

Silicon/Circuitry

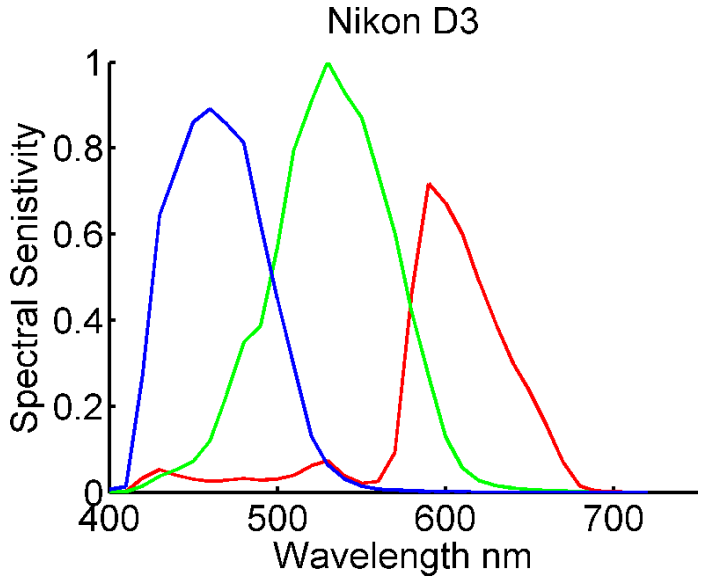
Photodiode
Photons hit the diode and force out electrons.
This design is similar to a solar cell!

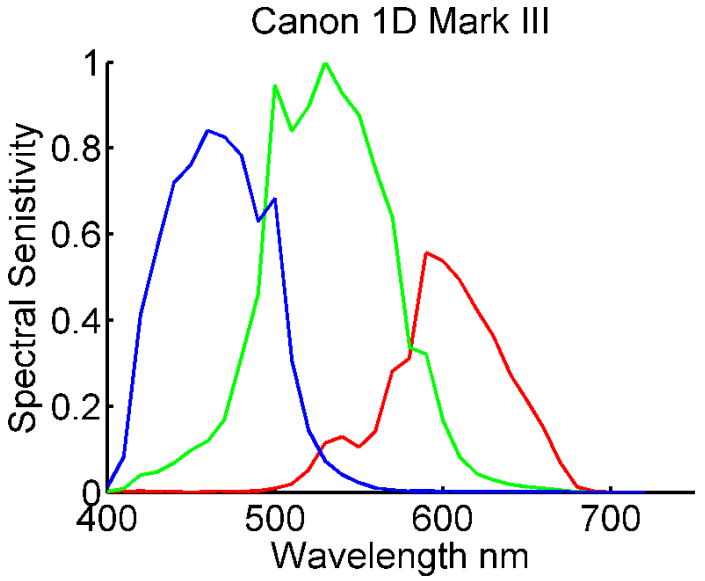


Bryce Bayer

Camera RGB sensitivity

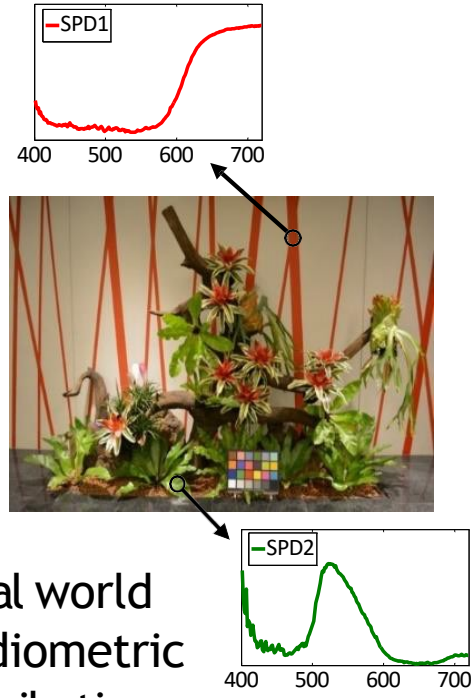
- The color filter array (CFA) on the camera filters the light into three *sensor-specific* RGB primaries



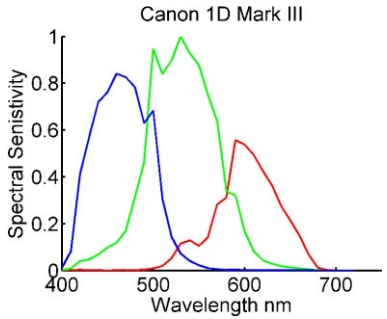


Plotted from camera sensitivity database by Dr. Jinwei Gu from Rochester Institute of Technology (RIT). Dr. Gu is now at SenseTime (USA). http://www.cis.rit.edu/jwgu/research/camspec/

Sensor raw-RGB image



Remember: physical world is measured by radiometric spectral power distributions.



Color Matching Functions CIE XYZ

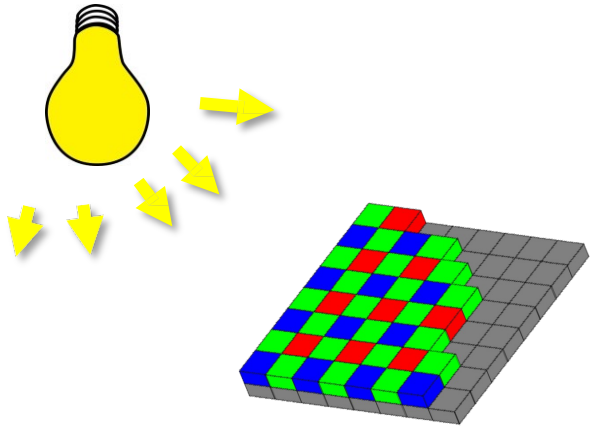
Your camera sensor RGB filter is sensitive to different regions of the incoming SPD.



raw-RGB represents the physical world's SPD "projected" onto the sensor's spectral filters.

Sensors are linear to irradiance

- Camera sensors are decent light measuring devices
- If you double the amount of light hitting a sensor's pixel, the digital value output of that pixel will double

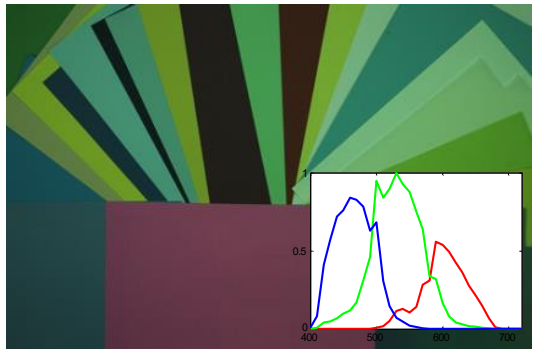


Sensor output is linear with respect to irradiance falling over the sensor over a certain amount of time.

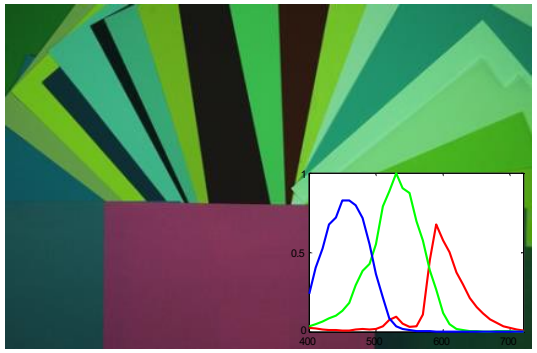
$$I = i * t$$

Digital value I is a linear function of irradiate i and exposure t.

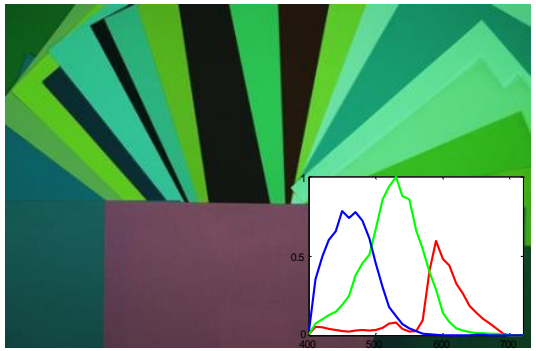
IMPORTANT: raw-RGB sensor images are not in a standard color space



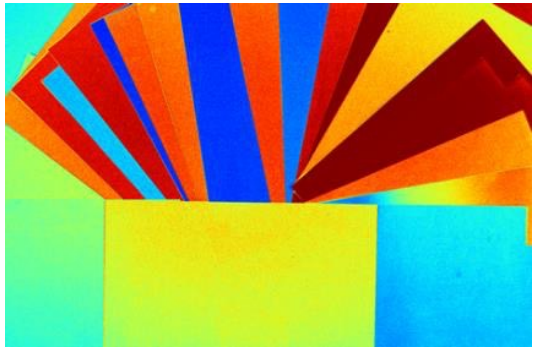
Canon 1D



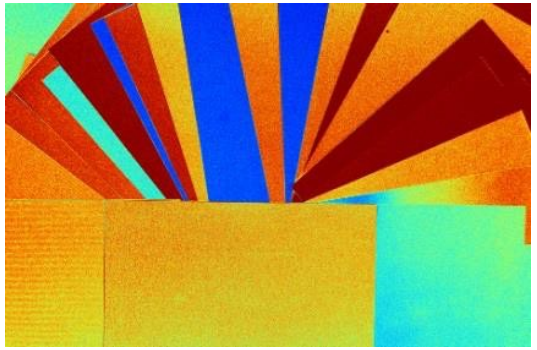
Nikon D40



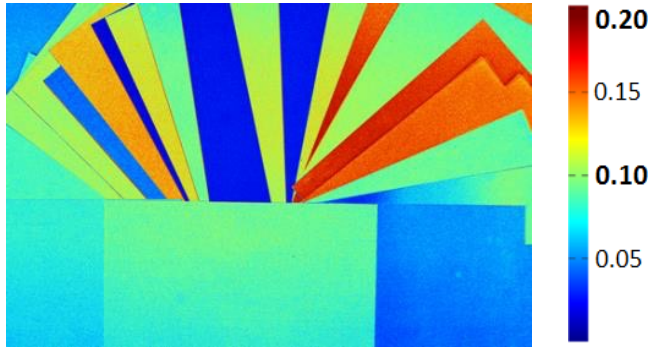
Sony α57



$||$Canon 1D – Nikon D40$||_2$



$||$Canon 1D - Sony α57$||_2$



$||$Nikon D40 - Sony α57$||_2$

Color plots show L2 distance between the raw-RGB values with different cameras.

Displaying raw-RGB images

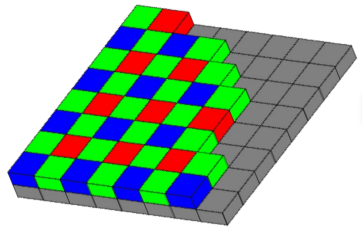
- Inserting a raw-RGB image in your slides, research paper, etc will result in strange colors.
- Why? Our devices (computers, printers, etc) expect the image to be in a standard color space like sRGB.



This is a raw-RGB image. Why does it look bad?
Because the RGB values are not sRGB values.
Knowing your color space is important!

A typical color imaging pipeline





Sensor with color filter array (CFA)
(CCD/CMOS)

ISO gain and raw-image processing

RGB Demoasicing

Noise Reduction

White-Balance & Color Space Transform (CIE XYZ)

Color Manipulation (Photo-finishing)

Mapping to sRGB output

JPEG Compression

Save to file

ISO signal amplification (gain)

- Imaging sensor signal is amplified and digitized
- Amplification to assist A/D conversion
  - Need to get the voltage to the range required to the desired digital output
- This gain is used to accommodate camera ISO settings
  - Gain to signal applied on sensor
  - Note – gaining the signal also gains image noise

Different ISO settings (note: the exposure will be shorter for higher ISO)

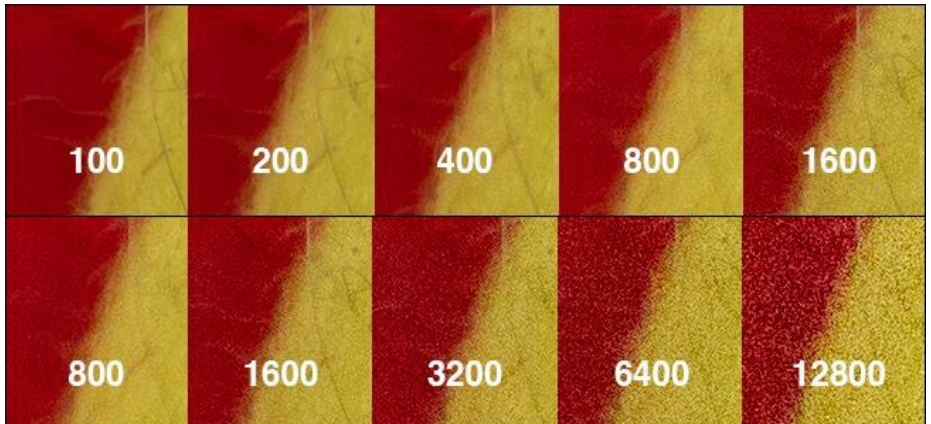


Image: Harry Guinness

ISO gain and raw-image processing

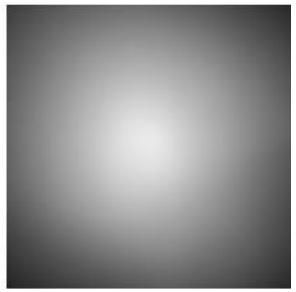
Pixel "intensity"

- We often talk about a pixel's intensity, however, a pixel's numerical value has no *unit*

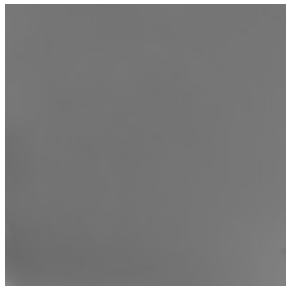
- The digital value of a pixel is based on several factors
  - Exposure (which is a function of both shutter speed and exposure)
  - Gain (ISO setting on the camera)
  - Camera hardware that digitizes the signal

- We typically rely on the ***relative*** digital values in the image and not the absolute digital values

Flat-field correction



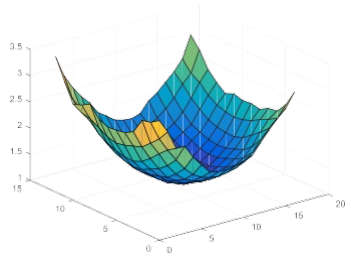
Uniform light falling on the sensor may not appear uniform in the raw-RGB image. This can be caused by the lens, sensor position in the camera housing, etc.



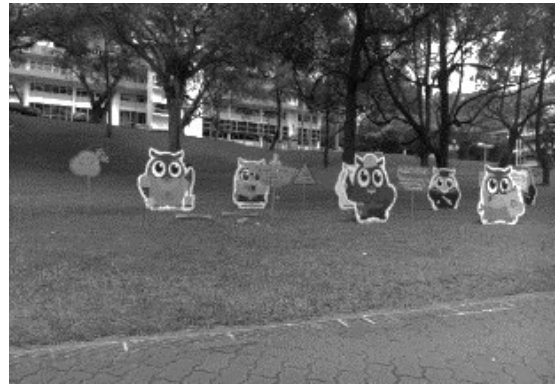
We want to correct this problem such that we get a "flat" output.



Before correction



Apply a correction gain over the sensor values.



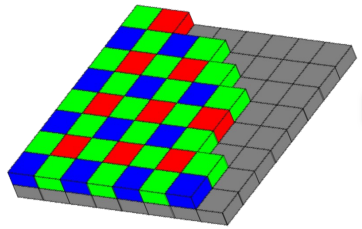
After correction

ISO gain and raw-image processing

This operation can also be called "lens shading" correction.

A typical color imaging pipeline





Sensor with color filter array (CFA)
(CCD/CMOS)

ISO gain and raw-image processing

RGB Demoasicing

Noise Reduction

White-Balance & Color Space Transform (CIE XYZ)

Color Manipulation (Photo-finishing)

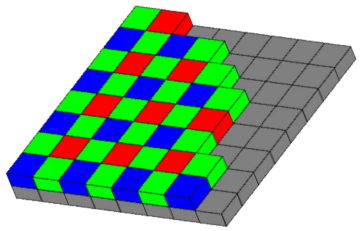
Mapping to sRGB output

JPEG Compression

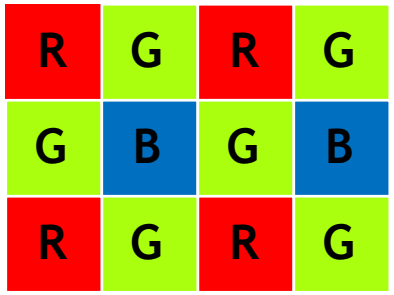
Save to file

CFA/Bayer pattern demosaicing

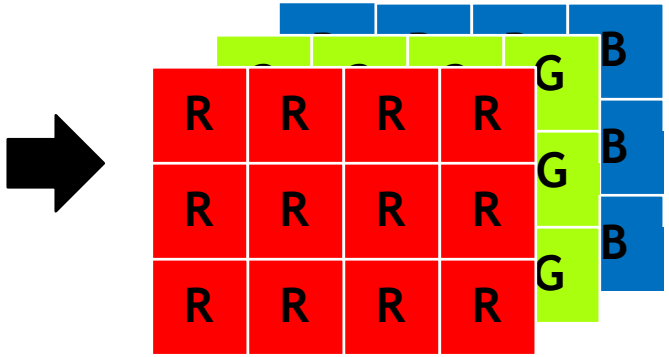
- Color filter array/Bayer pattern placed over pixel sensors
- We want an RGB value at each pixel, so we need to perform interpolation



Sensor with color filter array (CMOS)



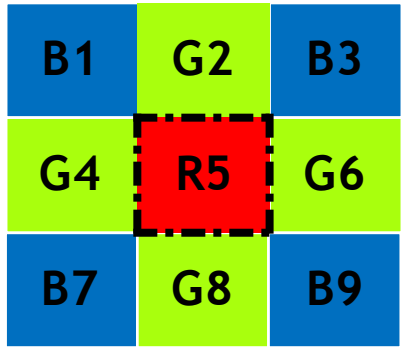
Sensor RGB layout



Desired output with RGB per pixel.

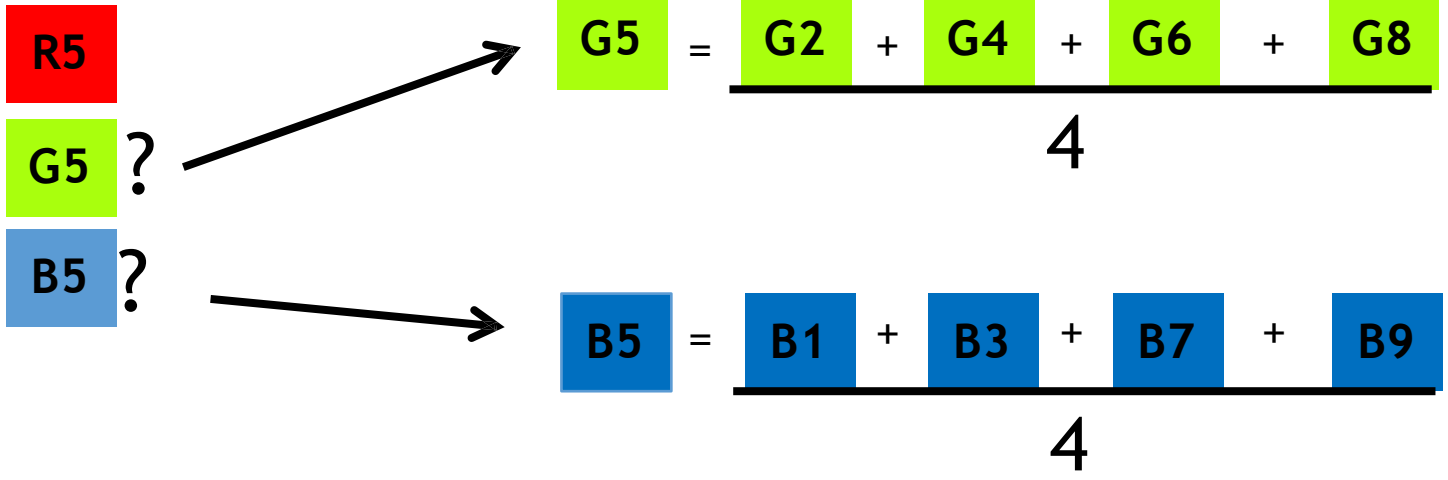
Simple interpolation

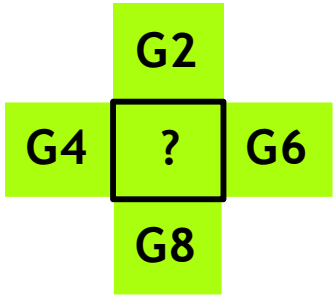
This is a zoomed up version of the Bayer pattern.

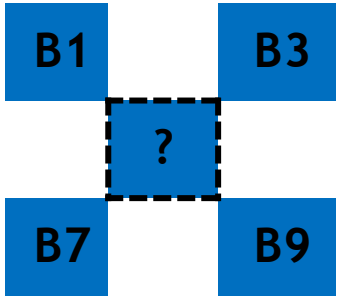


At location R5, we have a red pixel value, but no Green or Blue pixel.

We need to estimate the G5 & B5 values at location R5.

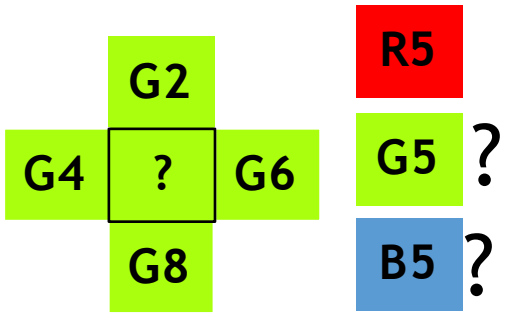






Simple "edge aware" interpolation





If (|(G2–G8) && |(G4-G8)| both < Thres):

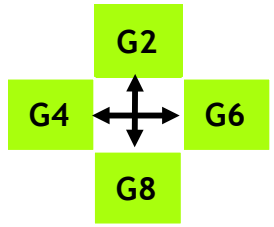
$$G5 = \frac{G2 + G4 + G6 + G8}{4}$$

elseif (|G2-G8| > Thres):

$$G5 = \frac{G4 + G6}{2}$$

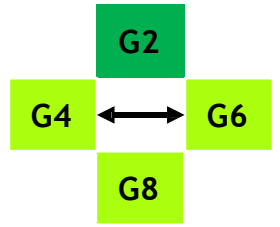
else:

$$G5 = \frac{G2 + G8}{2}$$



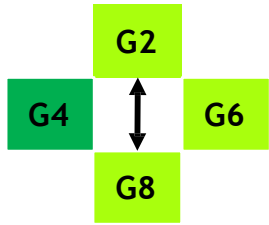
Case 1

All about the same.



Case 2

G2 and G8 differ - ignore them in the interpolation



Case 3

G4 and G6 differ - ignore them in the interpolation

Do this procedure also for the blue pixel, B5.

Demosaicing in practice

- The prior examples are illustrative algorithms only
- Camera IPSs use more complex and proprietary algorithms.
- Demosaicing can be combined with additional processing
  - Highlight clipping
  - Sharpening
  - Noise reduction

D N N for demosaicing (and denoising)

Gharbi et al SIGGRAPH Asia, 2016

Deep Joint Demosaicking and Denoising

Michaël Gharbi MIT CSAIL
Gaurav Chaurasia MIT CSAIL
Sylvain Paris Adobe
Frédo Durand MIT CSAIL



Figure 1: *We propose a data-driven approach for jointly solving denoising and demosaicking. By carefully designing a dataset made of rare but challenging image features, we train a neural network that outperforms both the state-of-the-art and commercial solutions on demosaicking alone (group of images on the left, insets show error maps), and on joint denoising–demosaicking (on the right, insets show close-ups). The benefit of our method is most noticeable on difficult image structures that lead to* moiré *or* zippering *of the edges.*

Abstract

Demosaicking and denoising are the key first stages of the digital imaging pipeline but they are also a severely ill-posed problem that infers three color values per pixel from a single noisy measurement.

1 Introduction

Demosaicking and denoising are simultaneously the crucial first steps of most digital camera pipelines. They are quintessentially ill-posed reconstruction problems: at least two-thirds of the data is

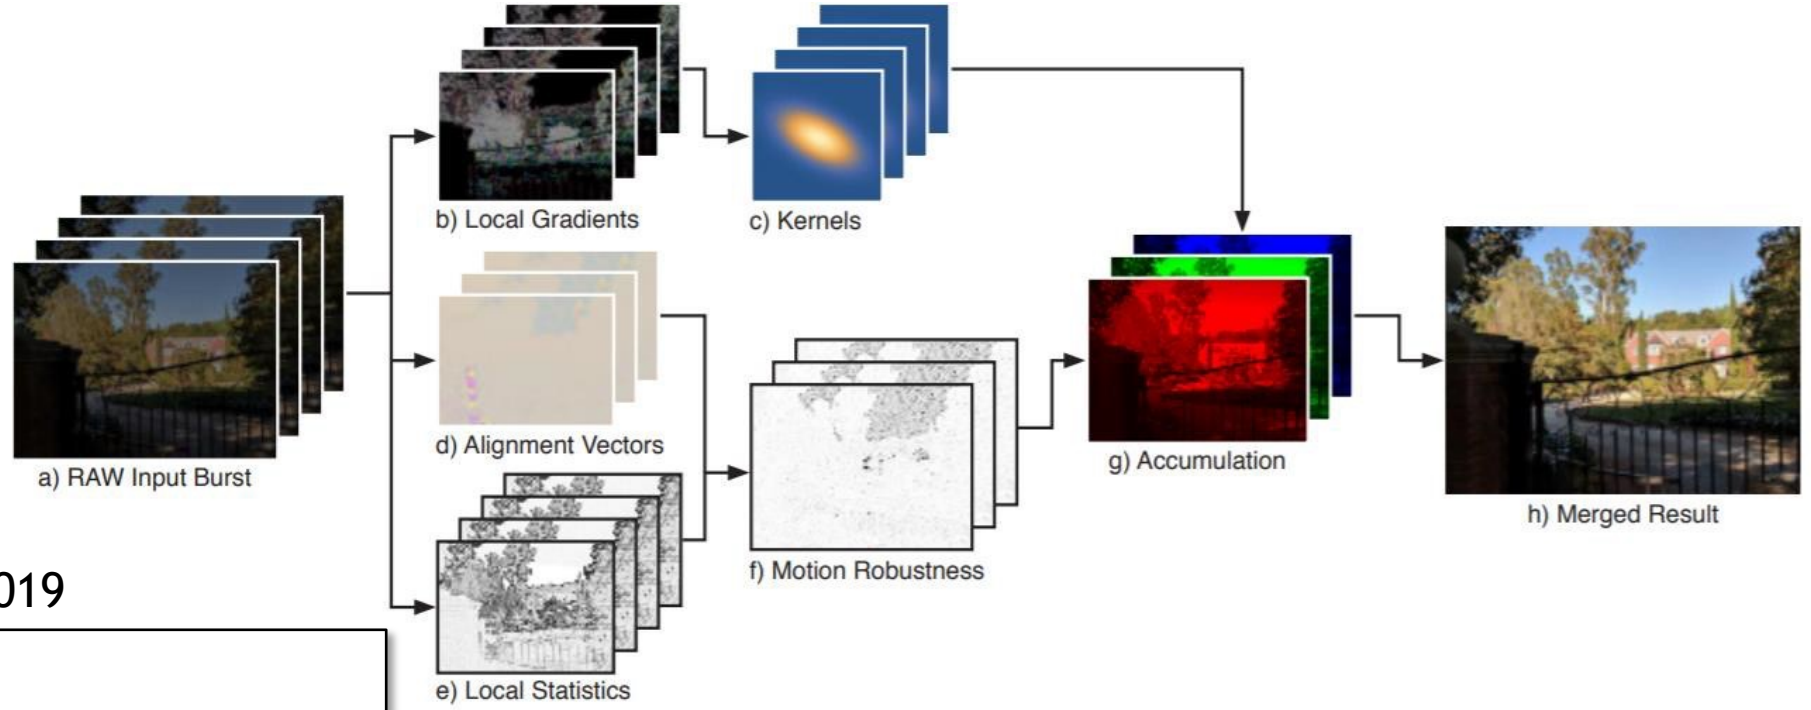
This is a very interesting paper that examines using a D N N to perform demoasicing.

Interestingly, the act of demosaicing to a clean ground truth image implicitly performs denoising.

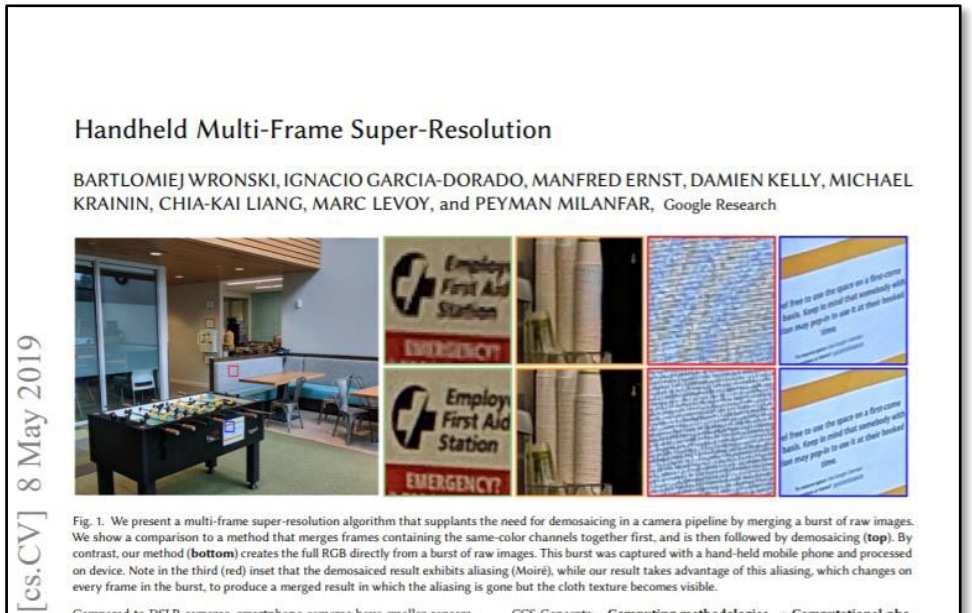
The only drawback of this paper is that training/testing demosiaced images were synthetically generated by sampling sRGB images and adding noise.

A lack of a raw-RGB dataset is a problem and would be a nice addition in this area.

Non-learning improvement for demosaicing



Wronski et al SIGGRAPH 2019

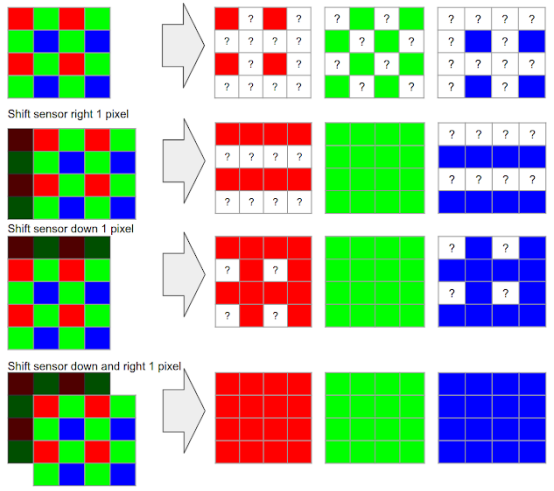


Handheld Multi-Frame Super-Resolution

BARTLOMIEJ WRONSKI, IGNACIO GARCIA-DORADO, MANFRED ERNST, DAMIEN KELLY, MICHAEL KRAININ, CHIA-KAI LIANG, MARC LEVOY, and PEYMAN MILANFAR, Google Research

[cs.CV] 8 May 2019

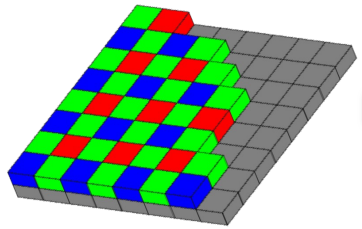
Fig. 1. We present a multi-frame super-resolution algorithm that supplants the need for demosaicing in a camera pipeline by merging a burst of raw images. We show a comparison to a method that merges frames containing the same-color channels together first, and is then followed by demosaicing (**top**). By contrast, our method (**bottom**) creates the full RGB directly from a burst of raw images. This burst was captured with a hand-held mobile phone and processed on device. Note in the third (red) inset that the demosaiced result exhibits aliasing (Moiré), while our result takes advantage of this aliasing, which changes on every frame in the burst, to produce a merged result in which the aliasing is gone but the cloth texture becomes visible.



This paper uses multiple frames and very small camera motion (from hand tremors) to perform demosaicing and super-resolution. By exploiting motion, they can fill in missing Bayer data.

A typical color imaging pipeline





Sensor with color filter array (CFA)
(CCD/CMOS)

ISO gain and raw-image processing

RGB Demoasicing

Noise Reduction

White-Balance & Color Space Transform (CIE XYZ)

Color Manipulation (Photo-finishing)

Mapping to sRGB output

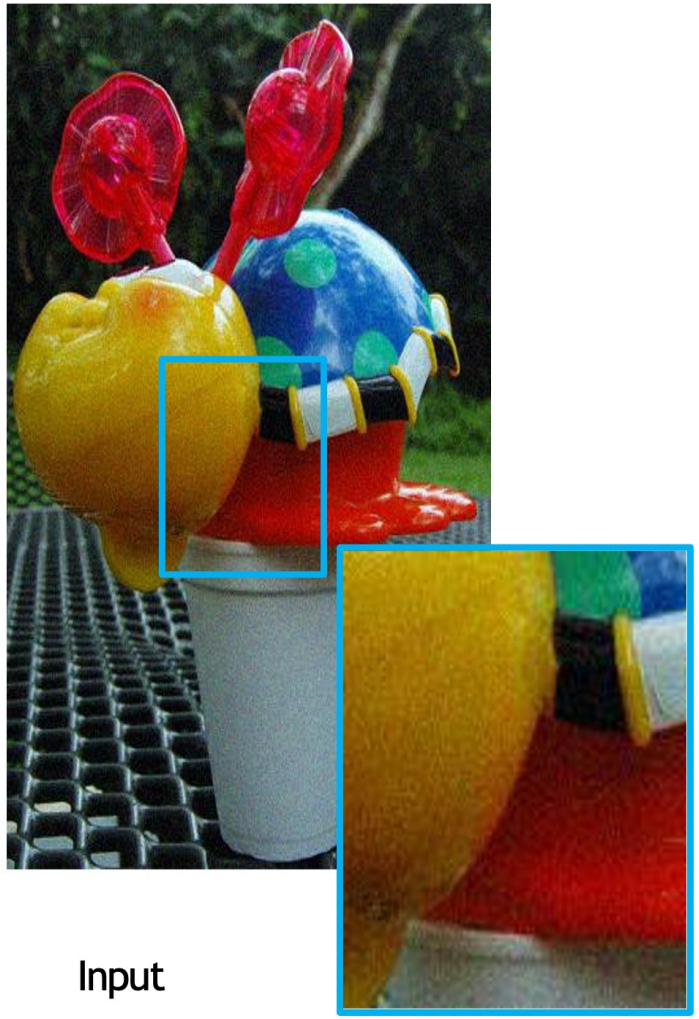
JPEG Compression

Save to file

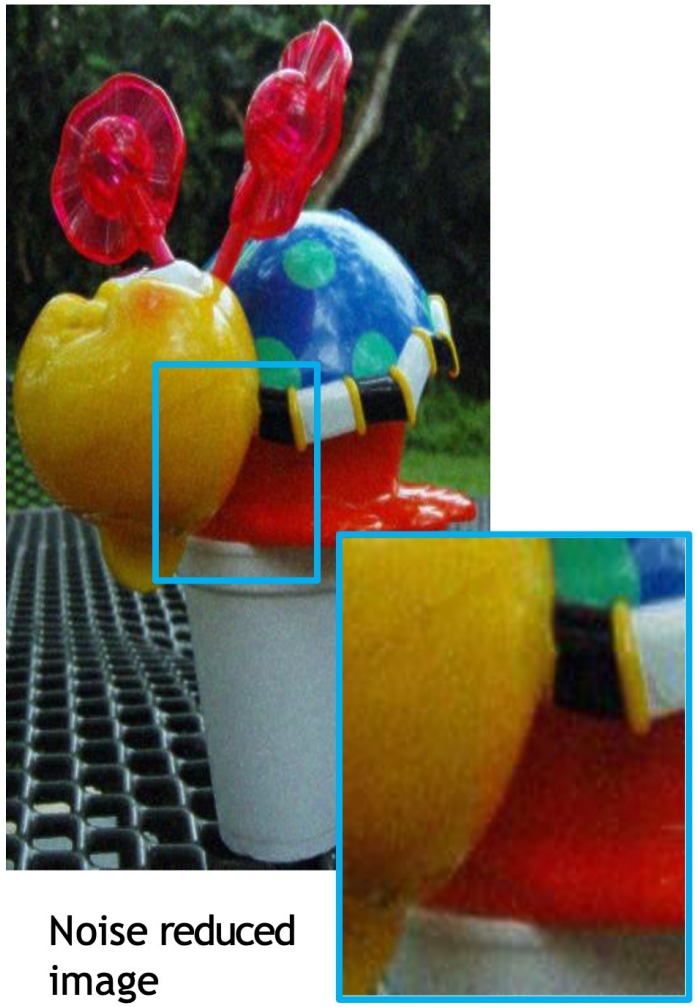
* Note that steps can be optional (e.g. noise reduction) or applied in slightly different order.

Noise reduction (NR)

- All sensors inherently have noise
- Most cameras apply additional NR after A/D conversion
- A simple method is described in the next slide
- For high-end cameras, it is likely that cameras apply different strategies depending on the ISO settings, e.g. high ISO will result in more noise, so a more aggressive NR could be used
- Smartphone cameras, because the sensor is small, apply aggressive noise reduction.



Input



Noise reduced image