Image As Multi-Dimensional Data Structure

George Legrady © 2022
Experimental Visualization Lab
Media Arts & Technology
University of California, Santa Barbara

What is an image?

In computing, an image is a data structure

- a two-dimensional array of pixel values positioned in a grid of horizontal and vertical pixels

Red, green, blue, horizontal loc, vertical loc, alpha channel

The red, green, blue layers

Red, green, blue, horizontal loc, vertical loc, alpha channel

Inside the Digital Camera

How does it detect light?

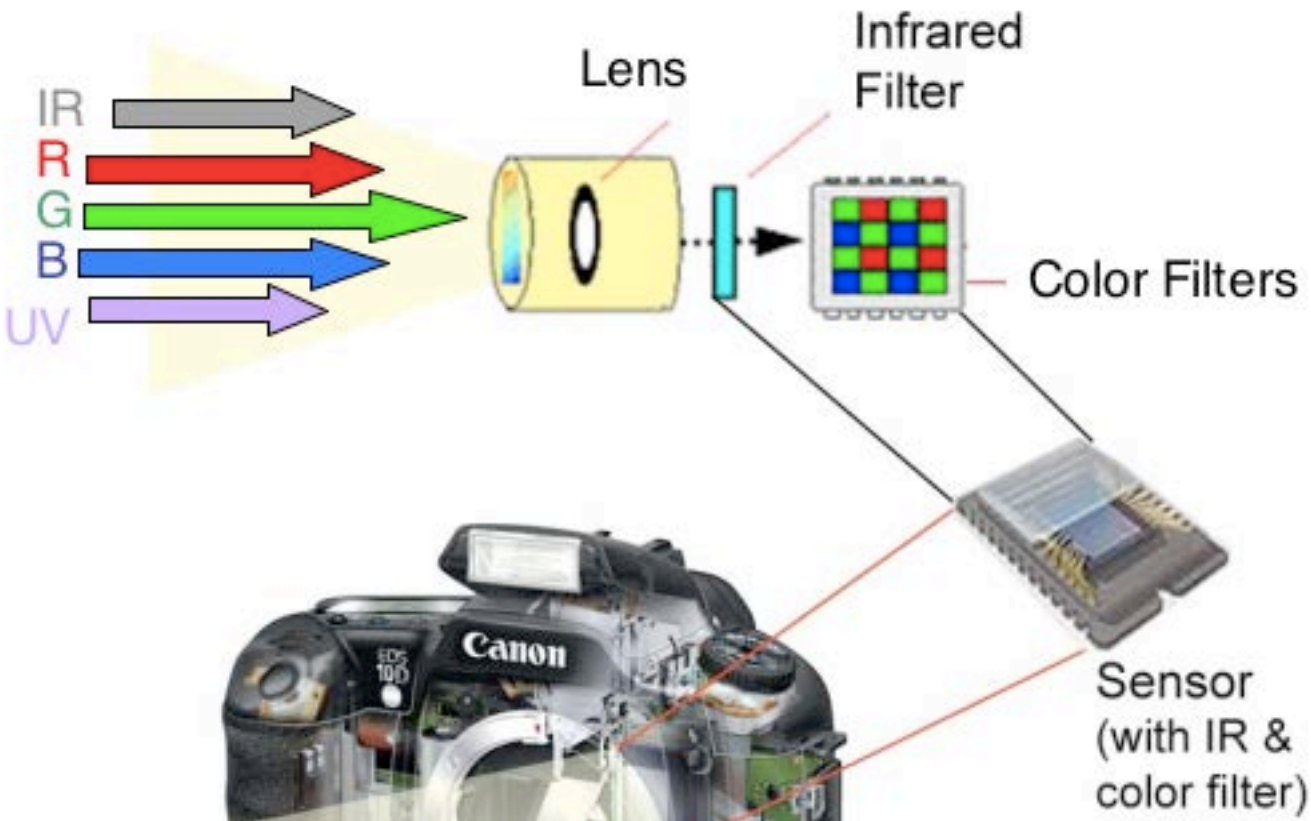

Note: Glass in camera lens blocks most UV from reaching sensor

The Digital image as Signal

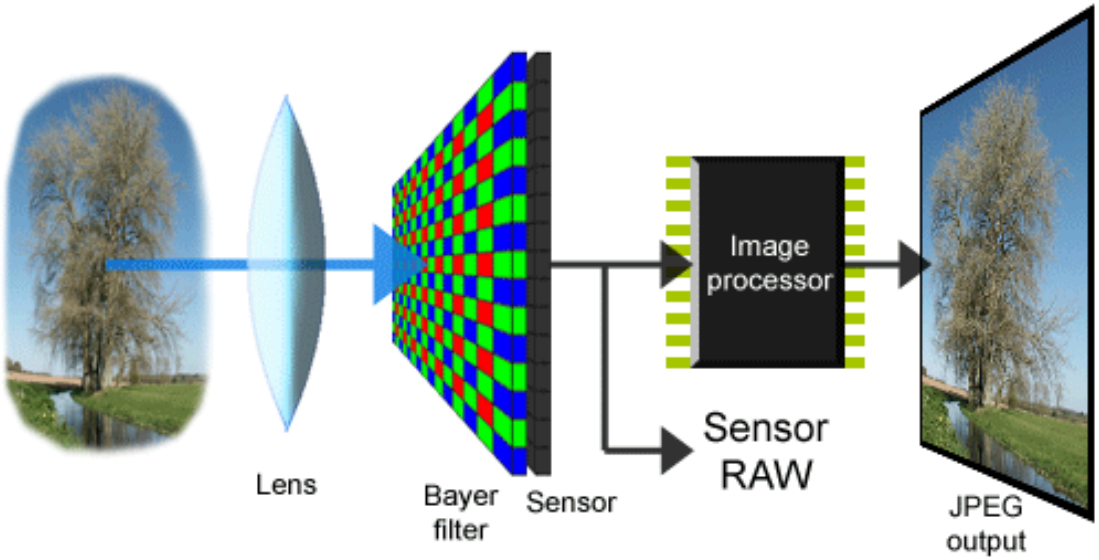

- Light captured by sensors in the camera is converted into digital form through sampling and quantization

Digital image *made up of pixels* is a multi-dimensional data structure

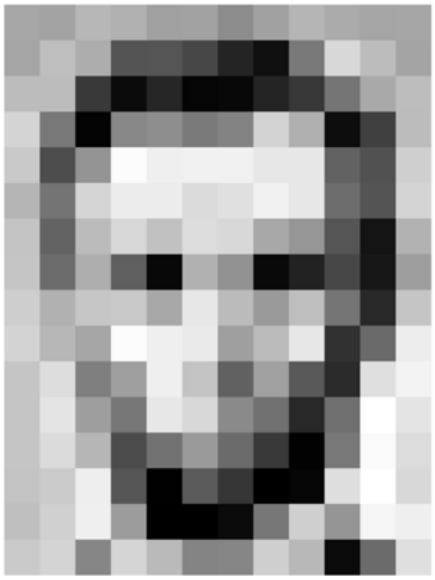

157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	33	17	110	210	180	154
180	180	50	14	34	6	10	33	48	106	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	106	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	86	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	33	17	110	210	180	154
180	180	50	14	34	6	10	33	48	106	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	106	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	86	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

- Pixel ***Horizontal*** location
- Pixel ***Vertical*** location
- Pixel ***Red*** color value
- Pixel ***Green*** color value
- Pixel ***Blue*** color value
- Pixel ***Alpha*** (transparency) value
- The whole image has a ***BitDepth*** resolution (2bit, 16bit, etc.)

The pixel is a discreet sample, Alvy Ray Smith (Pixar co-founder)

accurately IF the reconstruction is done correctly, and the Sampling Theorem tells how this is done. In other words, IF ONE IS CAREFUL, a discrete set of point samples is equivalent to a continuous infinity of points. Otherwise none of our digital displays would work. We think we are looking at a continuum when we watch digital light. To a large degree the history of computer graphics is the learning of how to cross back and forth across the discrete/continuous border defined by the Sampling Theorem. It is so fundamentally important that I think it should be taught in every school to every child. It defines the modern world.

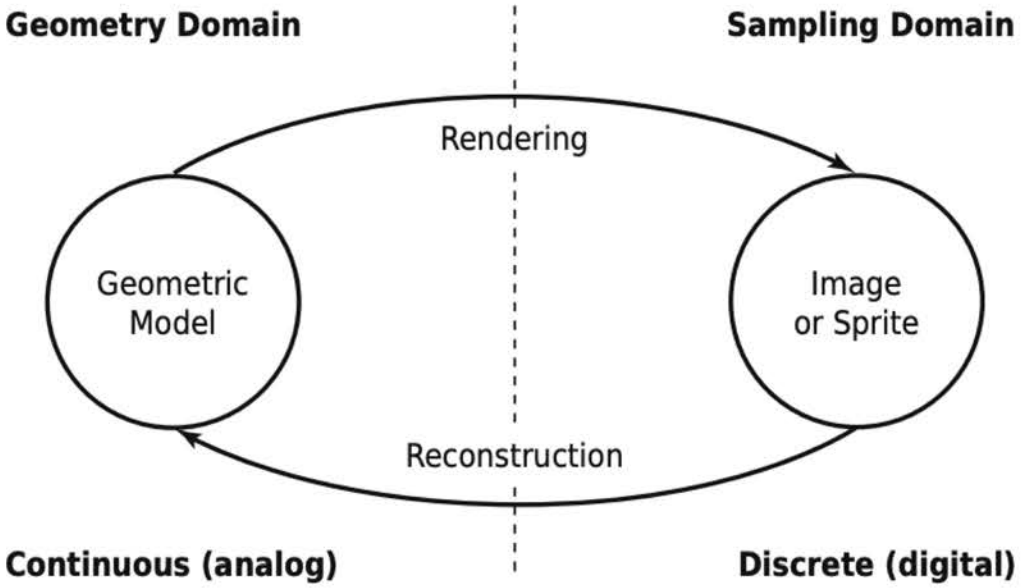

Figure 2. Geometry vs. sampling. Courtesy of Alvy Ray Smith.

The computer animation of Pixar is geometry-based. The sets and characters are defined with geometrical elements, assumed to move continuously through time. But consider digital photography. There is no geometry at all involved. The 'real world' is sampled with an array of sensors on a rectilinear grid.

Sampling rates: pixels per inch: 25%, 10%, 5%, 2%, 1%

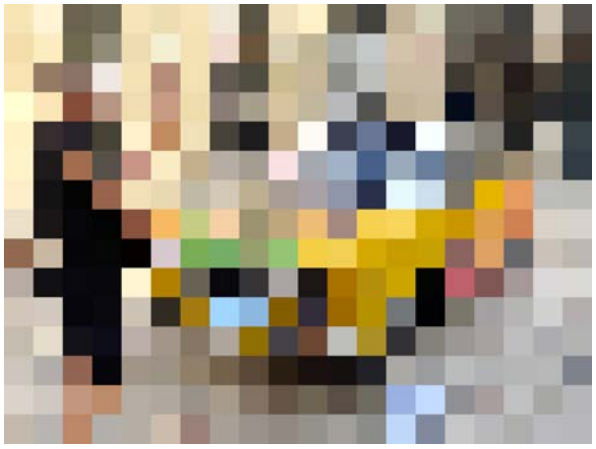

Finite Variations:

- If an image is 640 x 480 pixels, then the max variation of an image at this scale is 640 * 480 * 640 * 480 * 640 * 480 =
- or 307200 * 307200 * 307200 =
- or 2.8991E + 16

Digital image *made up of pixels* is a multi-dimensional data structure

- *Pixel* ***Horizontal*** location: **2560**
- Pixel ***Vertical*** location: **1920**
- **Each pixel has R,G,B values between 0 to 255**
- **Total bytes: 1,678,364 (1.7MB)**

Steganography: Compression allows for hiding data inside an image

- Steganography is the concealment of information within computer files
- When images are compressed, for instance if adjoining pixels have the same colors, this can be stored in shorthand as “3 x 245,23,67”, instead of “245,23,67”, “245,23,67”, “245,23,67” saving space
- Free space can be used to store other data which is then hidden

EXIF Data (*Digital cameras embed into the image how it was created*)

Sampling at different resolutions – DPI resolution (Dots per inch)

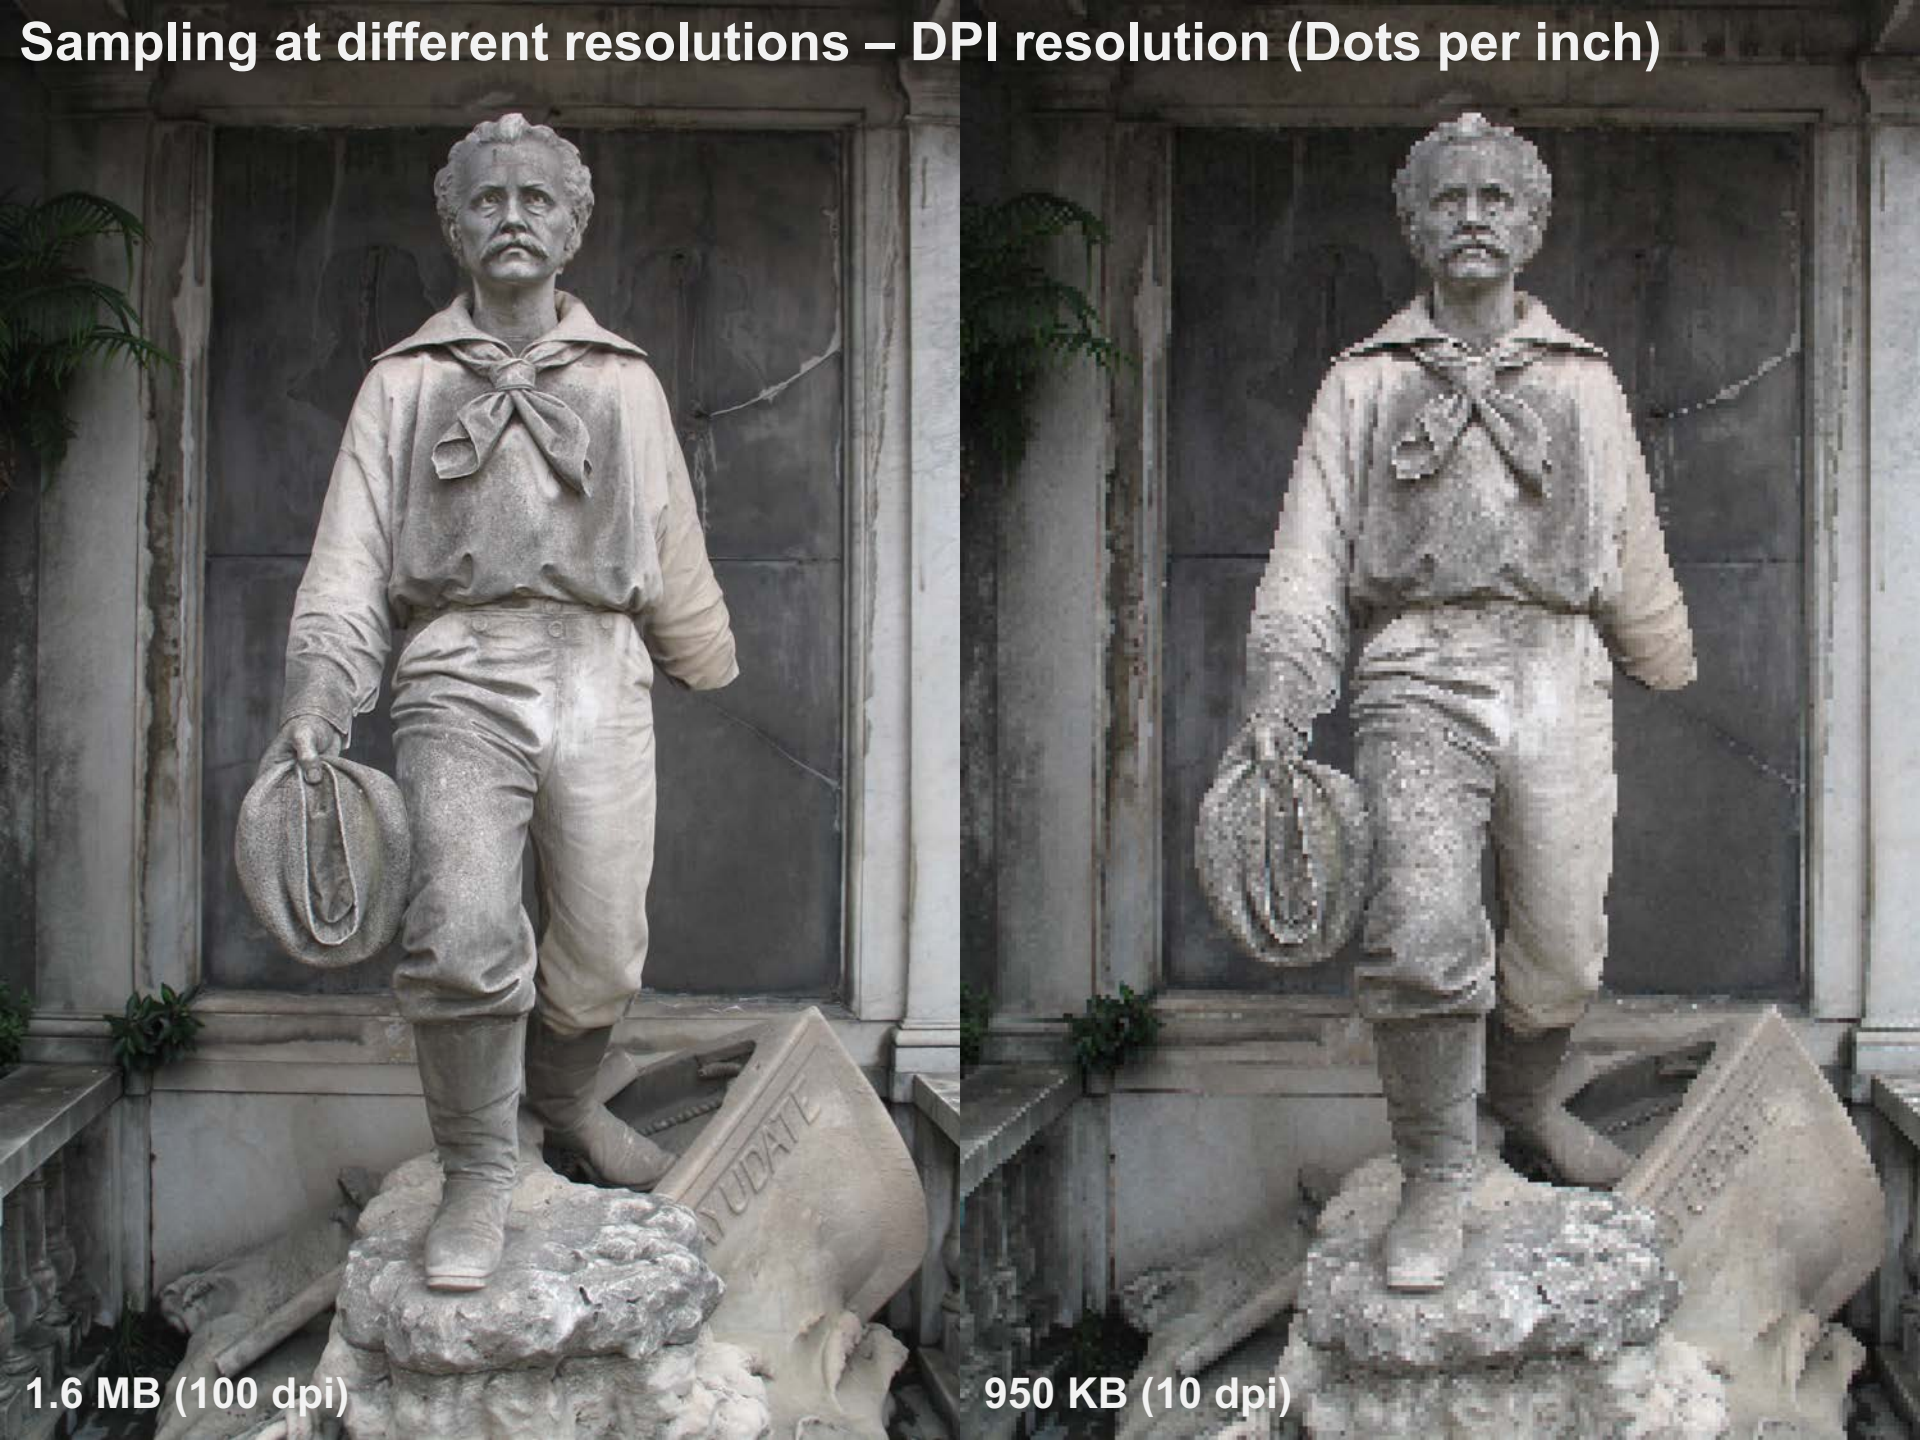

1.6 MB (100 dpi)

950 KB (10 dpi)

Two examples of different resolutions: grey scale (left), 2 bit dither (right)

1.2 MB

338 KB


Sharpen
[1,1,1]
[1,9,1]
[1,1,1]
1.6 MB
Maximum Saturation
4.9 MB

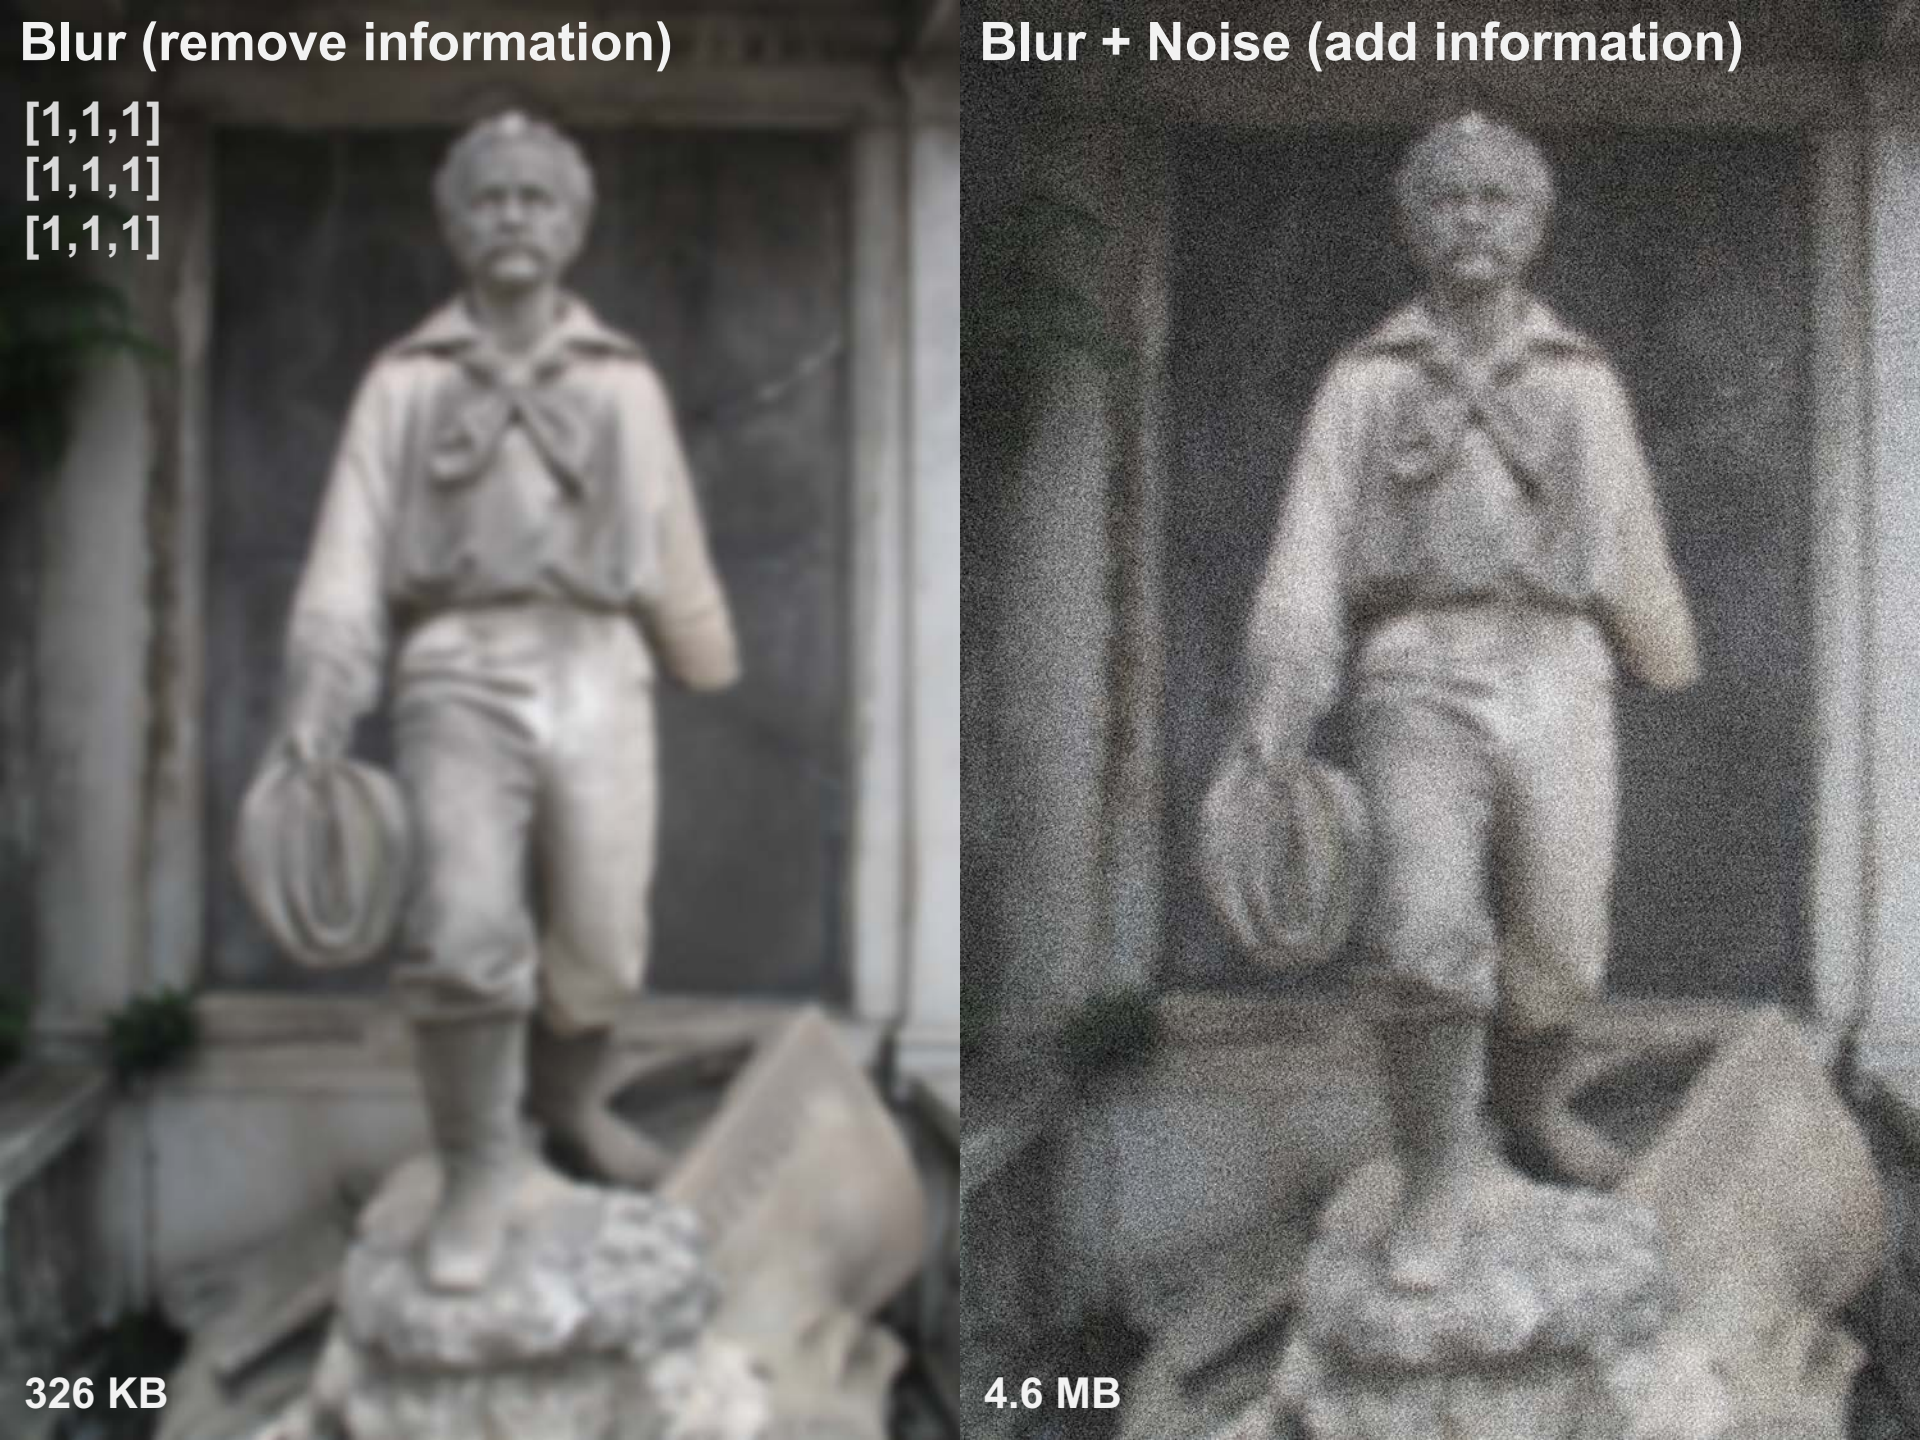
Blur (remove information)
[1,1,1]
[1,1,1]
[1,1,1]
326 KB
Blur + Noise (add information)
4.6 MB

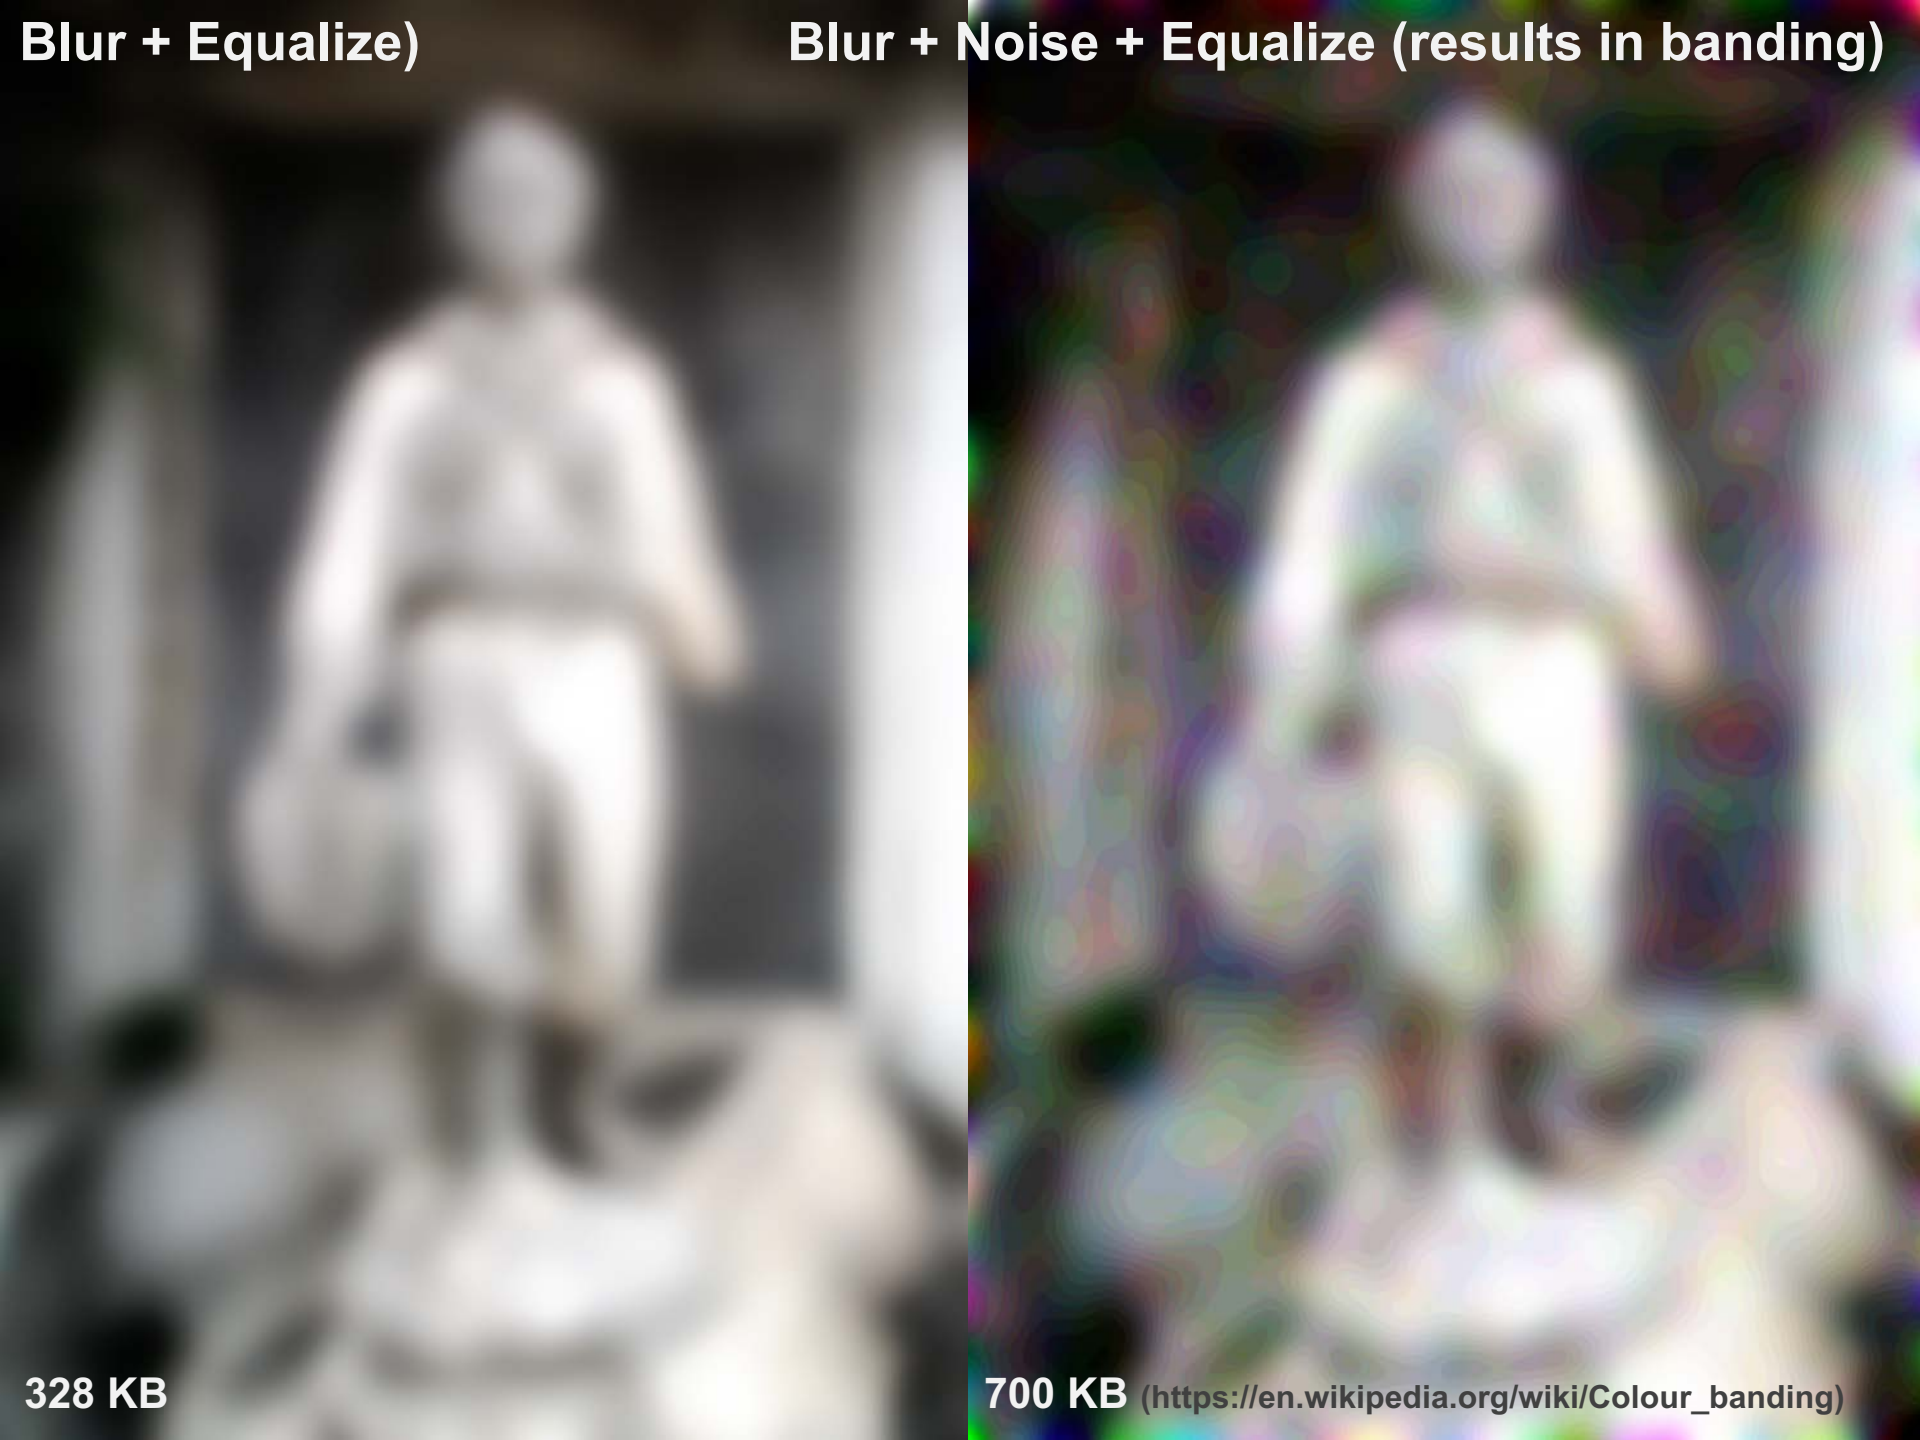
Blur + Equalize)
Blur + Noise + Equalize (results in banding)
328 KB
700 KB (https://en.wikipedia.org/wiki/Colour_banding)

Horizontal Edge Detection

[-1,-1,-1]
[9, 9, 9]
[-1,-1,-1]

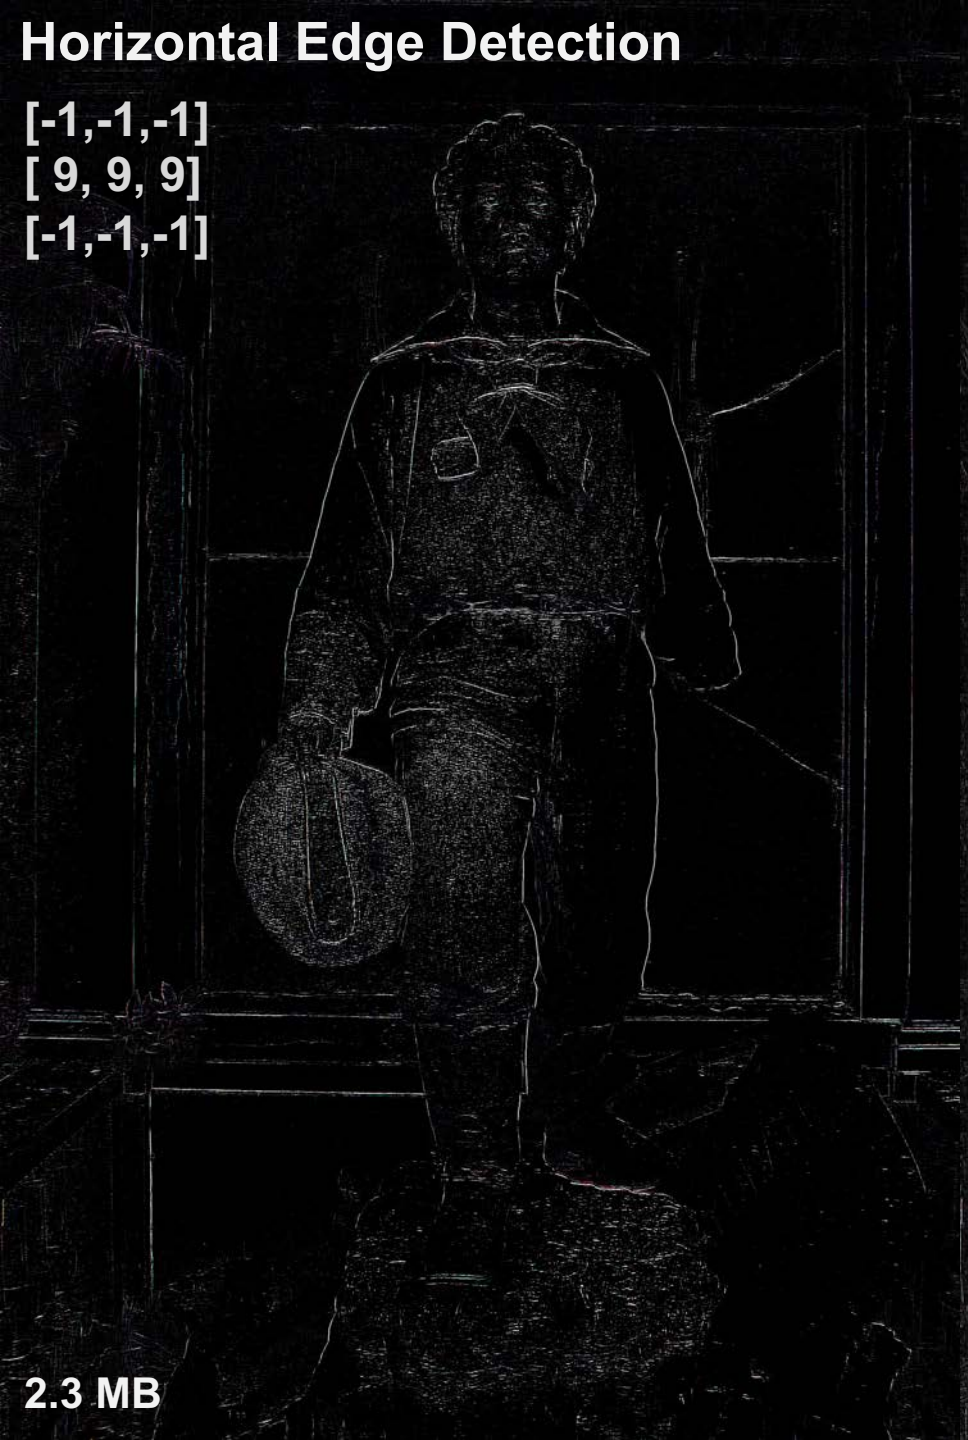

2.3 MB

Vertical Edge Detection

[-1,9,-1]
[-1,9,-1]
[-1,9,-1]

4.4 MB

Combining variations of 2D convolutions = Convolutional Neural Network

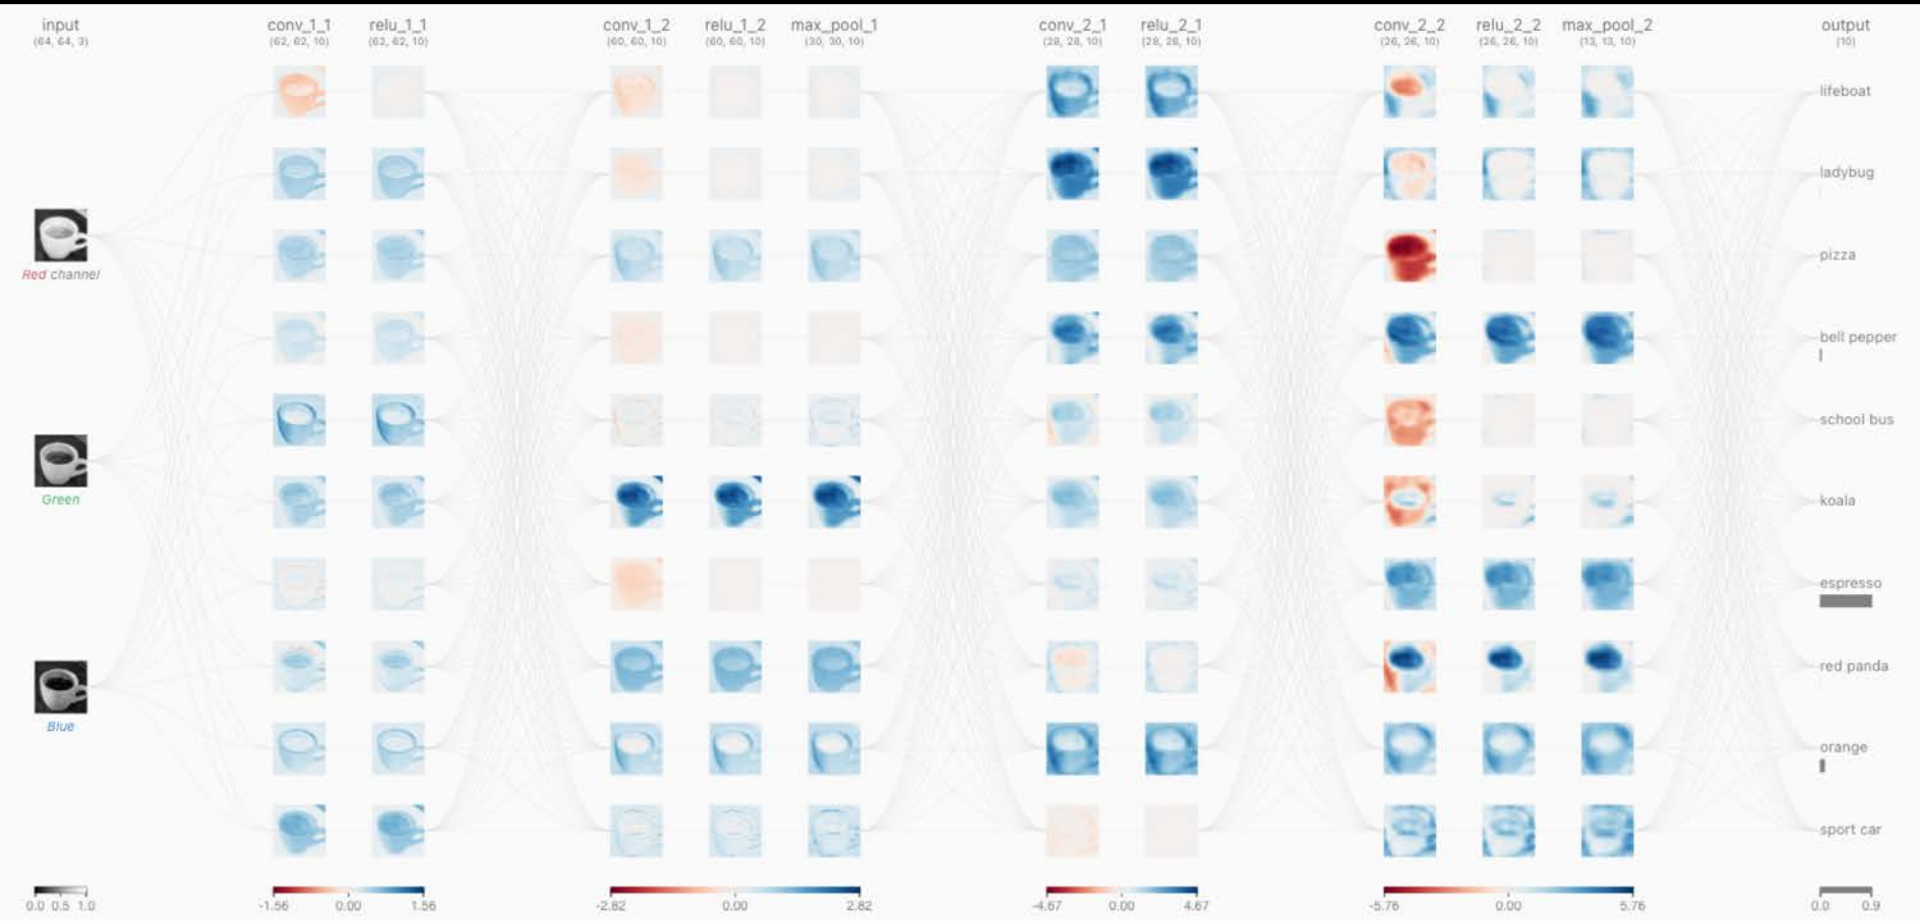

What is a Convolutional Neural Network?

In machine learning, a classifier assigns a class label to a data point. For example, an *image classifier* produces a class label (e.g, bird, plane) for what objects exist within an image. A *convolutional neural network*, or CNN for short, is a type of classifier, which excels at solving this problem!

Brig on the Water, Gustave Le Gray (1856) – two negatives exposed

https://artmuseum.princeton.edu/collections/objects/15941

Ansel Adams Zone System developed in collaboration with Fred Archer

1	2	3	4	5	6	7	8	9

Zone system chart for gamma = 2.2 (PC's, sRGB color space)>

Zone system chart for gamma = 1.8 (Macintosh)

Note 1. To display these tables correctly in Netscape, the Always use my colors, overriding document box must be unchecked. Click Edit, Preferences, Appearance, Colors) In Firefox, click Tools, Options, General, Fonts & Colors. To print in Internet Explorer 5, Click on Tools, Internet Options..., Advanced. Scroll down and check the box, "Print background colors and images." You might want to uncheck it afterwards.

Note 2. The best way to print these charts, which are HTML tables, ***not*** image files, is the following. (1) Adjust the width of the window for proportions you like. (2) Copy the window into the clipboard by pressing Ctrl-PrintScreen on your keyboard. (3) Paste the image into your image editor. (4) Crop it and otherwise adjust it in the editor. (5) Print it from the editor.

High-Dynamic Range: Emulating the Human Vision System

2 ev

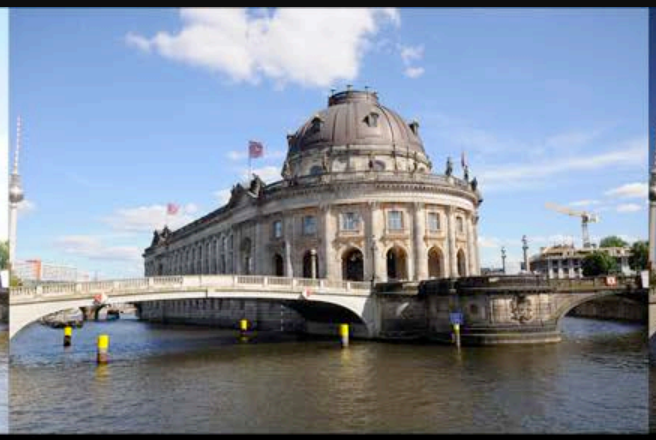

0 ev

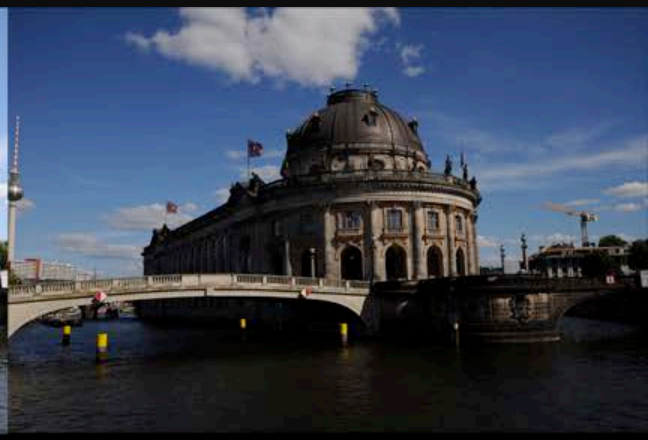

-2 ev

Tone-mapped HDR

Final image after post-processing

http://pristinekk.blogspot.com/2016/02/hdr-high-dynamic-range-hdr-photo-is.html

“jpeg ny02”, Thomas Ruff (2004)

https://www.davidzwirner.com/artists/thomas-ruff/survey#/jpeg-ny02--artwork-77CD3288-624F-46D1-8082-6DF161AEBBE0/Artwork

Lossy JPEG Compression

- In information technology, **lossy compression** or **irreversible compression** is the class of data encoding methods that uses inexact approximations and partial data discarding to represent the content. These techniques are used to reduce data size for storing, handling, and transmitting content.

- This is opposed to lossless data compression (reversible data compression) which does not degrade the data. The amount of data reduction possible using lossy compression is much higher than through lossless techniques.

- Lossy compression is most commonly used to compress multimedia data (audio, video, and images), especially in applications such as streaming media and internet telephony. By contrast, lossless compression is typically required for text and data files, such as bank records and text articles.

To be continued…