BE/Bi 205: Lecture 2

David Van Valen MD, PhD

04/06/2023

Today's Lecture

- Images as data
- Sequences as data
- Structures as data

What is an image?

- A visual representation in the form of a function f(x, y) where f describes brightness (or color) at a point (x, y)
- Continuous in amplitude and space
- Most images are defined over a rectangle (but they don't have to be)







Copyright 2023 Caltech

Athanasius Kircher, Dark Chamber with Lenses, 1946

What is a digital image?

- Discrete samples f[x, y] representing the continuous image f(x, y)
- Each element in this 2D array f[x, y] is called a pixel (from picture element)



What is a digital image?

• Ultimately a digital image is an array of numbers



Useful digital image properties

- Spatial resolution
 - Width and height in pixels/length (inches, cm, etc.)
- Intensity
- Number of channels
 - RGB is 3 channels, grayscale is 1
 - Alpha channel is used to indicate transparency
- Color spaces
 - RGB, CMYK, HSV



Color components

 The link between array values in single channel images and color images are look up tables (LUTs)



'PLASMA' LUT

'GRAY' LUT

Copyright 2023 Caltech

Jonas Hartmann, Computational Image Analysis, 2017

Color comoponents



Digital images are multi-dimensional arrays

• Multi-channel images, time series images, and image stacks are just multi-dimensional arrays



Digital hangover: data types

• Digital images vary in the precision used to store the array values

| int8 | Byte (-128 to 127) | Used in image processing |
|---------|--|--------------------------|
| int16 | Integer (-32768 to 32767) | Produced by microscope |
| int32 | Integer (-2147483648 to 2147483647) | |
| int64 | Integer (-9223372036854775808 to 922337203685 | 4775807) |
| uint8 | Unsigned integer (0 to 255) | |
| uint16 | Unsigned integer (0 to 65535) | |
| uint32 | Unsigned integer (0 to 4294967295) | |
| uint64 | Unsigned integer (0 to 18446744073709551615) | |
| float16 | Half precision float: sign bit, 5 bits exponent, 10 bits | mantissa |
| float32 | Single precision float: sign bit, 8 bits exponent, 23 b | its mantissa |
| float64 | Double precision float: sign bit, 11 bits exponent, 52 | bits mantissa |

Digital hangover: data types

• A cautionary tale – image subtraction



| 224 | 77 | 224 | 167 | 55 |
|-----|-------|--------------|--------|-----|
| 206 | 194 | 122 | 240 | 103 |
| 21 | 210 | 113 | 34 | 207 |
| 222 | 161 | 47 | 14 | 20 |
| 236 | 29 | 3 | 34 | 91 |
| Α | – B = | C (n) | p.uint | :8) |
| | | | | |

| -32 | 77 | -32 | -89 | 55 |
|-----|------|------|-----|------|
| -50 | -62 | -134 | -16 | -153 |
| 21 | -46 | 113 | 34 | -49 |
| -34 | 161 | -209 | 14 | 20 |
| -20 | -227 | 3 | 34 | -165 |

A.astype(np.int16) – B.astype(np.int16) = C (np.int16)

Digital hangover: data types

• A cautionary tale - rescaling



Rescales an array between 0 and x:

$$\frac{A - \min(A)}{\max(A) - \min(A)} \cdot x = C$$

| 173 | 181 | 222 | 247 | 205 |
|-----|-----|-----|-----|-----|
| 36 | 158 | 207 | 164 | 181 |
| 255 | 120 | 0 | 130 | 56 |
| 206 | 21 | 72 | 218 | 160 |
| 207 | 39 | 152 | 14 | 87 |

192

27

249 219

79 112

83

57

(((A - A.min()) / (A.max() - A.min()))*255).astype(pp)uint8) = C (np.uint8)

Digital image processing: Contrast adjustment

- Contrast adjustment is the process of changing pixel intensities to take advantage of the full range of values
- The goal is to maximize the intensity differences between different objects or parts of the image, making them easier to see







Copyright 2023 Caltech https://www.mathworks.com/help/images/contrast-adjustment.html

Digital image processing: Contrast adjustment



- A convolution is a mathematical operation that produces a new version of an existing function using selections of that function and a second function
- In digital image processing this second function is called a kernel
- With the right kernel, convolutions can perform tasks and extract important features
 - Blurring Gaussian kernel
 - Edge finding
 - Corner finding





Note: Behavior at edges is undefined. Default in *scipy* is `reflect`.





• Convolutions are a mathematical operation that "respect" translation



Copyright 2023 Caltech https://chriswolfvision.medium.com/what-is-translation-equivariance-and-why-do-we-use-convolutions-to-get-it-6f18139d4c59

- The mathematical term for this "respect" is translational equivariance
- Equivariant mappings preserve the algebraic structure of transformations
 - The mapping commutes with the symmetry operation



Digital image processing: Affine transforms

- Affine transformations transform a shape while preserving
 - Straight lines
 - Relative distances of points on straight lines
- We use affine transformations for image augmentation
- Image augmentation boosts the diversity of our training data by performing augmentation operations that preserve the labels



Digital image processing: Affine transforms

- When used for augmentation with pixel-level predictions, affine transformations need to be applied to both the image and the labels
 - Applying the transformation to just one corrupts the learning signal
- Affine transformations typically involve interpolation of some kind
 - Pay attention to how this happens!
 - Image precision also matters!



Digital image processing: Foreground detection

RAW



THRESHOLDED



- Foreground detection splits the background and foreground of the image
 - The end result is a binary True/False array
 - This is also called a mask
- Thresholding is an essential part of foreground detection
 - Manual vs. automated
 - Uniform vs. adaptive
- Morphological operations can improve the quality of the final mask

Digital image processing: Thresholding

- There are many thresholding methods: Otsu's method is one that you should be aware of
- Otsu's method
 - Assume the image contains 2 classes of pixels following a bi-modal histogram
 - Search for the threshold that minimizes the intra-class variance
 - Minimizing the intra-class variance is the same as maximizing the inter-class variance

$$\sigma_b^2(t) \; = \omega_0(t) \omega_1(t) [\mu_0(t) - \mu_1(t)]^2$$

$$egin{aligned} &\omega_0\mu_0+\omega_1\mu_1=\mu_T\ &\omega_0+\omega_1=1 \end{aligned}$$

Algorithm

- 1. Compute histogram and probabilities of each intensity level
- 2. Set up initial $\omega_i(0)$ and $\mu_i(0)$
- 3. Step through all possible thresholds $t=1,\ldots$ maximum intensity
 - 1. Update ω_i and μ_i
 - 2. Compute $\sigma_b^2(t)$

Copyright 2023 Caltech 4. Desired threshold corresponds to the maximum $\sigma_b^2(t)$

Digital image processing: Thresholding

 Masks from thresholded images can be used to operate on other arrays/images



THRESHOLDED MORPHOLOGICALLY PROCESSED 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 0 0 0 0 1 1 1 1 0 0 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 1 1 0 0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 1 1 1 1 1 1 0 0 1 1 1 1 1 1 1 1 0 1 1 1 1 0 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 1 1 1 1 1 0 0 0 0 1 1 1 1 0 0 0 0 0 0 1 0 0 0 0 0 0 0 1 1 1 1 1 1 0 1 0 0 0 0 0 0 1 1 1 1 1 1 1 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 0 1 0 1 1 0 0 0 1 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 1 1 0 0 0 1 0 0 0 0 1 0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 1 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 1 1 1 1 11000000 0 0 0

- Very similar in principle to filtering
 - "Structuring elements" are used in a similar manner as kernels
- Common operations
 - Erosion and dilation
 - Opening and closing
 - Hole filling



DISC-SHAPED SE

• Example: Dilation in 1D





- Multiple operations and fine tuning may be required to get the desired result
 - This fine-tuning often introduces "brittleness" into the pipeline it works well for data it was finetuned on and poorly for other sets
 - Bypassing the need for this kind of fine-tuning is a big advantage of deep learning
 - Classical methods like morphological operations can still have a role in building deep learning-enabled systems

Digital image processing: Connected components



Sequences as data

- Typical examples: amino acid sequences (proteins), nucleotide sequences
 - Nucleotides: A, C, (T/U), G
 - Amino acids: A, R, N, D, C, Q, E, G, H, I, L, K, M, F, P, S, T, W, Y, V
 - Time series: EKG, EEG, etc.
- Sequences often need a numerical encoding to make them amenable to machine learning
 - Time series sequences are an exception
 - Even with a numerical encoding there are additional issues to make sequences ML ready (e.g., uneven length, padding, etc.)

One hot encoding



- Each token of a vocabulary gets labeled with all 0's and a single 1
- Pros:
 - Simple, fast, inexpensive

One hot encoding



- Cons:
 - Sparse, memory inefficient, and high dimensional
 - Low information content no notion of similarity between sequence or structural elements
 - Dimensionality is fixed

Learned embeddings



Map a token to a vector and require this mapping to have certain properties $X \rightarrow \vec{k}; k \in \mathbb{R}^n$

- Learned embeddings seek to map tokens to a vector space in a particular way to ensure that the embedding has specific properties
 - Similar tokens are close together, while dissimilar tokens are far apart
 - Embeddings have information about the sequence context
- There have been a variety of methods developed to create useful embeddings
 - word2vec, doc2vec
 - Unsupervised k-mer sequence embeddings
 - ESM Fold, AlphaFold, etc.

Learned embeddings

Unsupervised learning





Masked pre-training is a common technique with language models and has been successful on protein sequences

Learned embeddings



Embeddings learned by ESM clustered amino acids by their side chain properties

- Structural data captures the spatial locations of the component atoms in a molecule
- This data can be represented in a variety of ways so it can be presented to deep learning models
- Each representation lends itself to a different set of deep learning methods





- Voxel representation: Structural information is captured by assigning features to a 3D grid
 - Features can be occupancy, chemical properties, etc.
 - Voxel representations effectively convert structures to 3D images, which makes them amenable to vision models



• SMILES

- Simplified molecular input line entry system
- ASCII representation for chemical structures
- Typically used for small molecules
- Amenable to both language models and vision models
- One-hot encoding on top of SMILES representations has seen some use



One-hot encoding on top of SMILES representations has seen some use



- Graph representation
 - Atoms are nodes
 - Bonds are edges
 - The graph can have node and edge features that are captured by their respective matrices
 - The adjacency matrix captures the atom connectivity
 - The graph representation is amenable to graph neural networks



- Molecular point cloud
 - Version of the graph representation where the 3D coordinates of atoms are captured
 - Different variants of this idea exist
 - Coordinates as node features
 - Encoding location via torsion angles
 - Different levels of resolution also exist
 - e.g., only capture the location of the Cα carbon for amino acids
 - Amenable to graph neural networks