



Deep learning and bioimage analysis

Data Driven Life Sciences (DDLS) course

KTH - 2023

Estibaliz Gómez de Mariscal, EMBO Postdoctoral Fellow

Optical Cell Biology Group (Prof. Ricardo Henriques)

Instituto Gulbenkian de Ciência





@gomez_mariscal



egomez@igc.gulbenkian.pt

Today's lecture


- BioImage analysis: definition
- Deep learning for image processing
- Segmentation
- Considerations about DL



Bioimages contain lot of information




How many cells can we count?




Computational image processing:

- Precise
- Reproducible
- Transferable
- Automatic → FAST



Digital images


A **digital image** is a mapping of intensities from a 2D grid of (uniformly spaced) discrete points, into a set of numerical values. The grid elements are called **pixels**.



BioImages: biological information given by numbers

Raw data

Codifies the information contained in the image



0	1	2	3	4	5	6	7
0	0.912062	0.894625	0.21853	0.094814	0.745898	0.284247	0.980677
1	0.593781	0.881641	0.283428	0.249029	0.375144	0.394311	0.251872
2	0.936699	0.396945	0.937976	0.990739	0.651616	0.369892	0.806337
3	0.466779	0.841376	0.495952	0.670456	0.922456	0.826899	0.712388
4	0.282393	0.422493	0.812124	0.756462	0.798258	0.181772	0.1563
5	0.234833	0.935021	0.623741	0.186595	0.297894	0.789515	0.62341
6	0.821762	0.494741	0.246734	0.253439	0.127357	0.541123	0.134337
7	0.548848	0.498073	0.985419	0.129993	0.268893	0.856691	0.51205
8	0.218218	0.225782	0.459472	0.812888	0.288139	0.181945	0.807689
9	0.525016	0.259971	0.137972	0.426034	0.87146	0.24694	0.976575
10	0.231879	0.338959	0.391136	0.876129	0.558661	0.189497	0.341358
11	0.348801	0.418849	0.748876	0.256852	0.524362	0.957343	0.712228
12	0.875186	0.392295	0.88925881	0.958719	0.544945	0.237546	0.76039
13	0.518258	0.907361	0.915734	0.8792741	0.117125	0.813719	0.192075
14	0.989263	0.523213	0.839298	0.271368	0.498379	0.169455	0.985422
15	0.929697	0.178869	0.954913	0.852831	0.516879	0.629484	0.251428
16	0.572548	0.620854	0.984113	0.352866	0.993376	0.888374	0.248578
17	0.8328735	0.147678	0.356826	0.996781	0.173752	0.192928	0.989551
18	0.793844	0.471979	0.884959	0.648271	0.894491	0.213725	0.382664
19	0.188833	0.8898789	0.281579	0.168874	0.768456	0.15985	0.811869
20	0.24135	0.318843	0.52667	0.838611	0.523253	0.988714	0.64341
21	0.631555	0.197771	0.8288182	0.279186	0.419381	0.852642	0.581324
22	0.429872	0.726135	0.661441	0.646851	0.899366	0.883813	0.458745
23	0.488012	0.182647	0.143884	0.358812	0.392898	0.721293	0.885897
24	0.283615	0.349788	0.918145	0.544887	0.498818	0.394286	0.379541
25	0.88235189	0.18338	0.28423	0.982319	0.378447	0.888439	0.217819
26	0.573237	0.982540	0.129727	0.578382	0.38266	0.759982	0.278411
27	0.154159	0.387383	0.539944	0.8361467	0.213495	0.116879	0.453791
28	0.591882	0.858113	0.178423	0.295442	0.783162	0.474241	0.687441
29	0.237285	0.114423	0.533374	0.656558	0.463769	0.463523	0.162685

Metadata

Set of text data providing additional information about the image.

- Imaging modality
- Objective
- Magnification
- Resolution or Pixel/Voxel size (microns, mm)
- Number of channels
- Excitation spectrum
- Information about the patient

Key	Value
BitsPerPixel	12
DimensionOrder	XYCZT
IsDeleteaved	false
IsRGB	false
LibreEndian	true
PixelType	uint16
Series 0 Name	punto004.nd2 (series 1)
SizeC	3
SizeT	1
SizeX	1024
SizeY	1024
SizeZ	62
Device	TI.ZDrive
Sep	0.1 Dm
Average	1
AverageToQuality	0
CH1 (Laser Wavelength) #1	488.0 (Laser Power): 100.0
CH1 (Laser Wavelength) #2	488.0 (Laser Power): 100.0
CH1 (Laser Wavelength) #3	488.0 (Laser Power): 100.0
CH1 (Laser Wavelength) #4	488.0 (Laser Power): 100.0
CH1CLEMnBrightness	4
CH1CLEMnEnoughSignalThreshold	800
CH1CLEMnNoSignalThreshold	200
CH1ChannelColor	16711680
CH1ChannelCyeName	C491
CH1ChannelLaserIndex	0
CH1LaserPower	100
CH1LaserSimulatorPower1	0
CH1LaserSimulatorPower2	0
CH1LaserSimulatorPower3	0
CH1PMTHighVoltage	121
CH1PMTOnset	9
CH2 (Laser Wavelength) #1	488.0 (Laser Power): 85.6
CH2 (Laser Wavelength) #2	488.0 (Laser Power): 85.6
CH2 (Laser Wavelength) #3	488.0 (Laser Power): 85.6
CH2 (Laser Wavelength) #4	488.0 (Laser Power): 85.6
CH2CLEMnBrightness	4
CH2CLEMnEnoughSignalThreshold	800
CH2CLEMnNoSignalThreshold	200
CH2ChannelColor	65280
CH2ChannelCyeName	FITC
CH2ChannelLaserIndex	3
CH2LaserPower	85.6625
CH2LaserSimulatorPower1	0
CH2LaserSimulatorPower2	0
CH2LaserSimulatorPower3	0
CH2PMTHighVoltage	78



The solution for reading proprietary microscopy image data and metadata

Digital (bio)-image analysis





Image processing

Image Processing is any form of data processing for which the input is an image – the output is not necessarily an image.





Input: Image - Output: Coordinates (bounding boxes)



Input: Image
Output: Image




Input: Image
Output: Image











Input: Image
Output: Label

Non infected



Infected


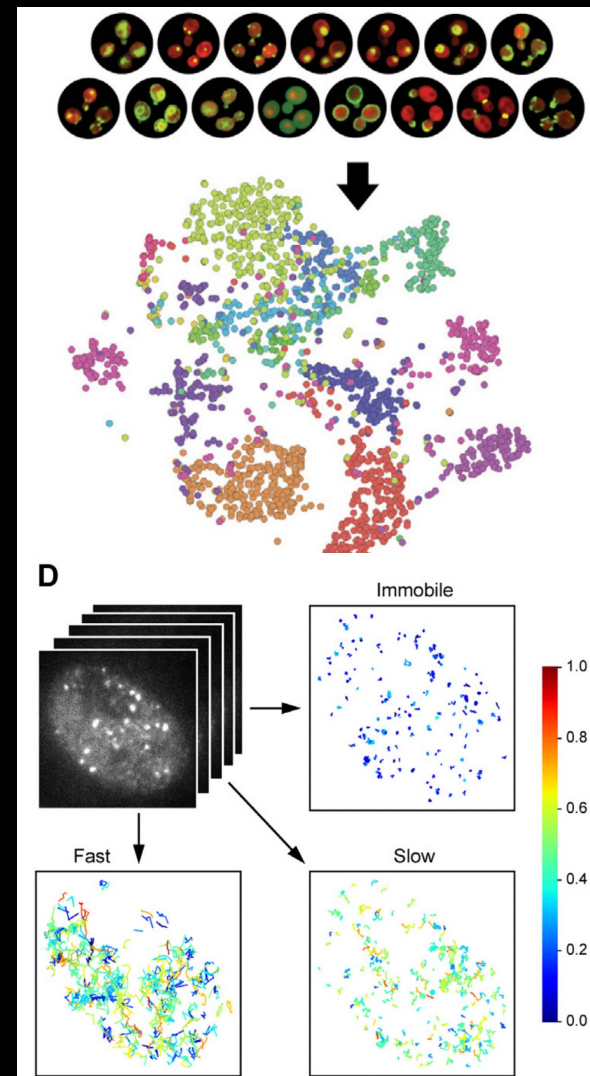

(Classical) image processing tasks



Christoph Spahn, et al., bioRxiv, 2021
Elnaz Fazeli, et al., F1000Research 2020
Ashley Kiemen, et al., bioRxiv 2020

BiImage analysis


Collection of image processing techniques to extract numerical information from scientific images



Deep learning:
an extremely hot topic in the field 🔥

The deep learning landscape for microscopy imaging


Segmentation



F. Lux & P. Matula, arXiv, 2020


Naert et al., Development 2021

Artificial labelling



L. Von Chamier, ..., R. Henriques, Nat Comms 2021


Resolution enhancement & restoration



C. Qiao et al., bioRxiv 2023

C. Qiao et al., Nat Biotech 2022

SMLM



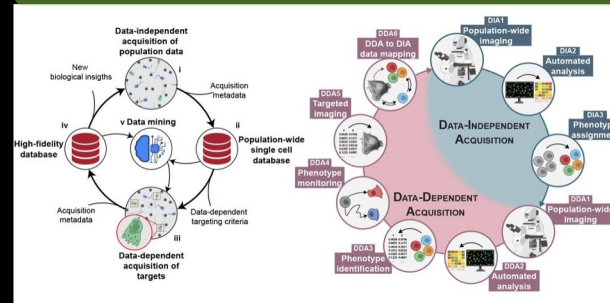
E. Nehme et al., Optica, 2018

Detection

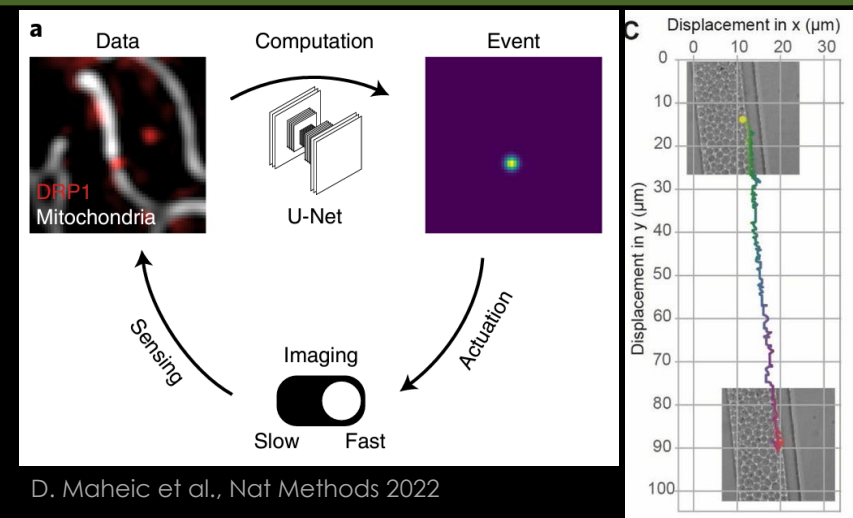


C. Spahn et al., Comm Biology 2022

Data driven microscopy



O. André et al., Cell reports 2023




D. Maheic et al., Nat Methods 2022

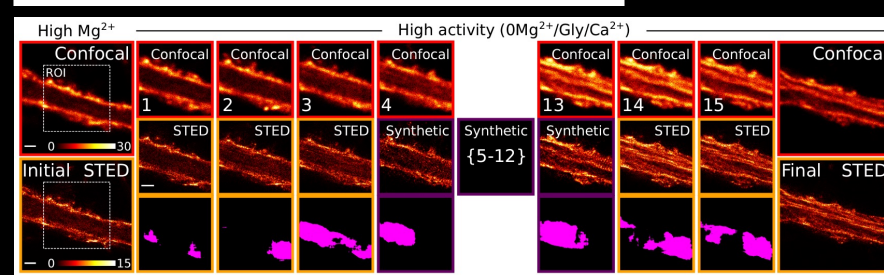
C. Bouchard et al., bioRxiv 2023

L. Chiron et al., Sci Reports 2022

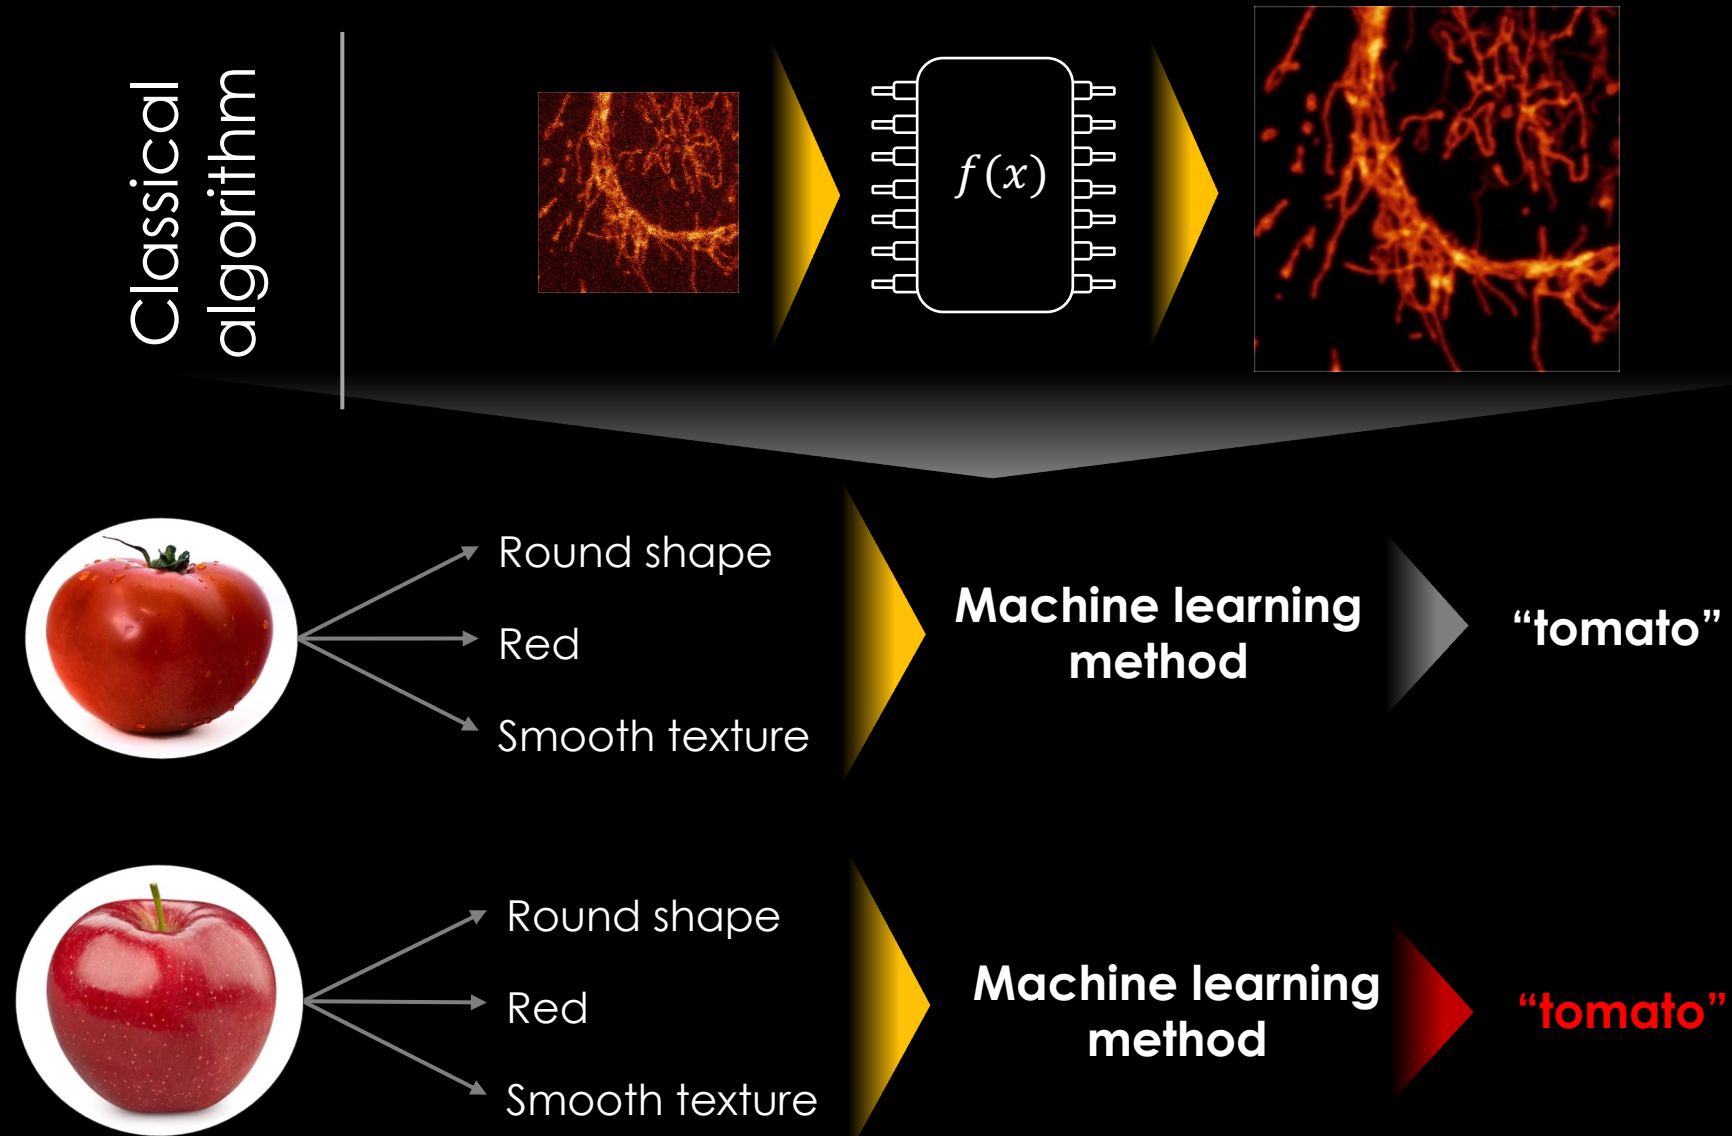
Denoising




L. Von Chamier, ..., R. Henriques, Nature Communications 2021



Why does deep learning pose a new paradigm?

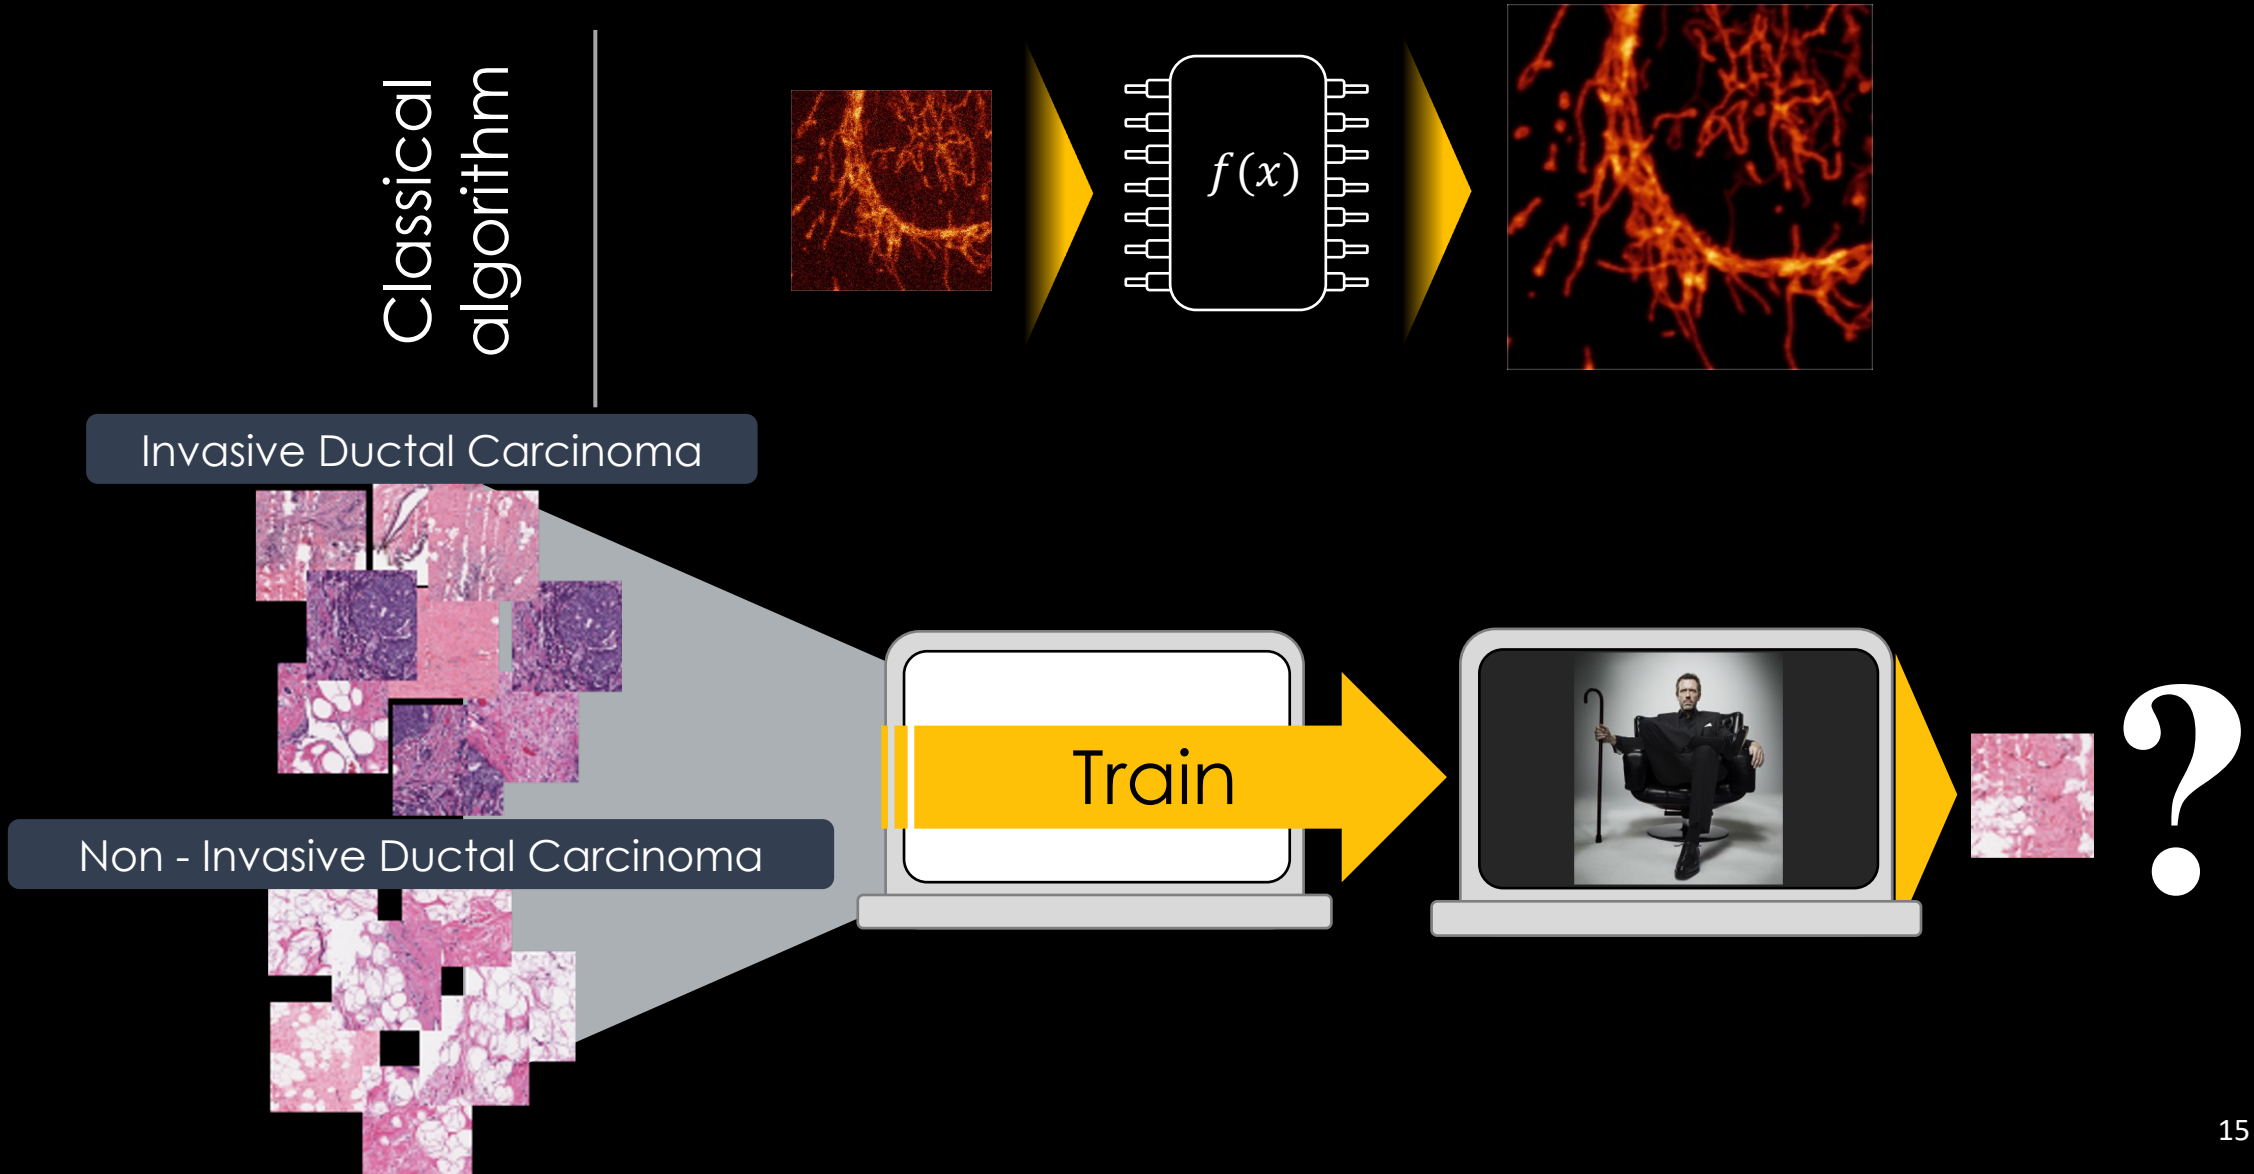


Why does deep learning pose a new paradigm?






What if the system could learn
automatically from the **data**?

Why does deep learning pose a new paradigm?









Why does deep learning pose a new paradigm?

Classical
algorithm



Deep Learning
algorithm



Trained (convolutional
neural) network

What are (convolutional) neural networks?


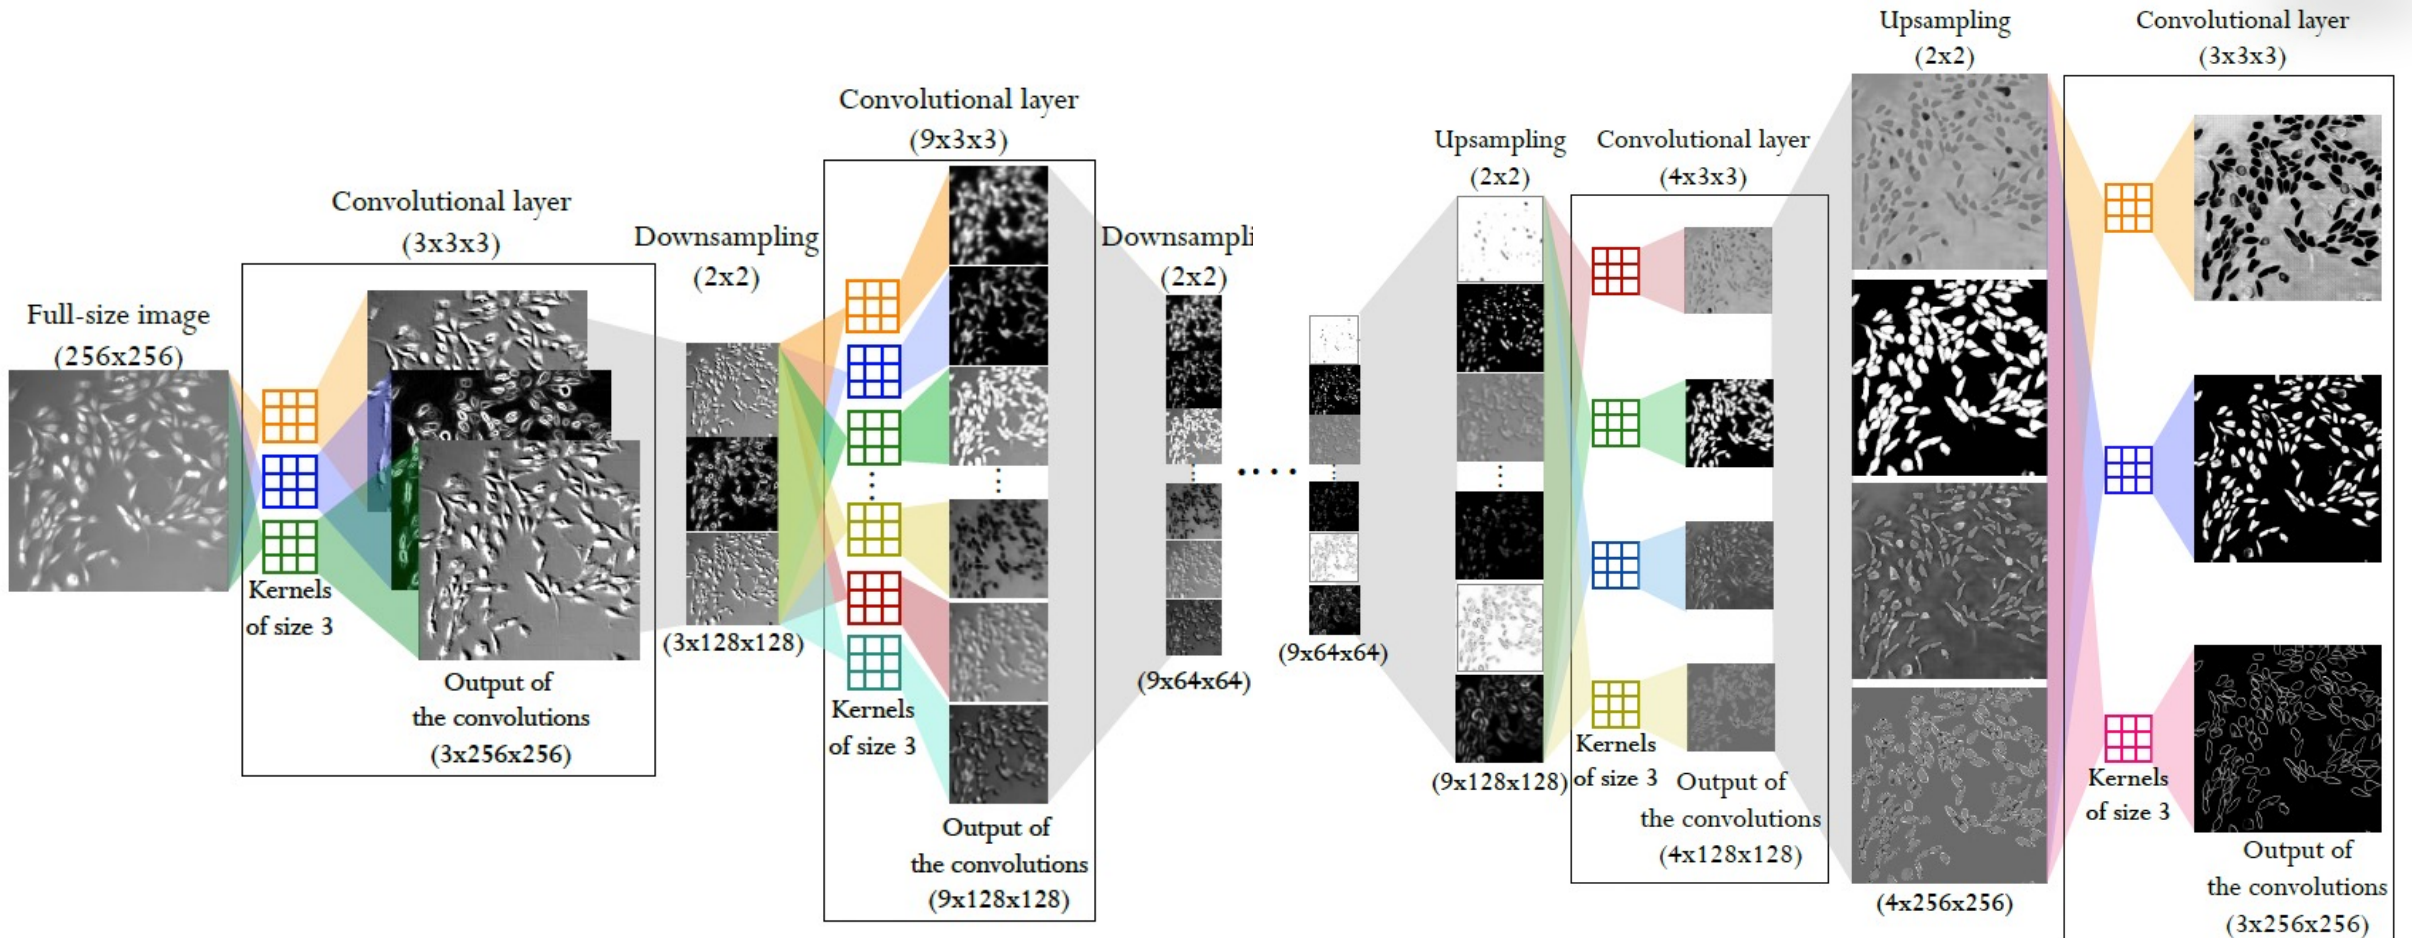


Image segmentation convolutional neural network architecture (2D U-Net)



Convolutions

1	0.5	0.5	0.5	1
0.5	1	0.5	1	0.5
0.5	0.5	1	0.5	0.5
0.5	1	0.5	1	0.5
1	0.5	0.5	0.5	1



Mean

	0.66	0.66		
		0.77		

1	0.5	0.5	0.5	1
0.5	1	0.5	1	0.5
0.5	0.5	1	0.5	0.5
0.5	1	0.5	1	0.5
1	0.5	0.5	0.5	1

0.66	0.66	0.66
0.66	0.77	0.66
0.66	0.66	0.66

Convolutions




Contextual information around the pixel







Spatial filtering

Convolutional kernel: determines de feature to enhance

Convolutions




Capacity to quantify and enhance features of interest in the image



Filtered images and image filters can be combined in multiple ways

What if the system could learn the optimal combinations **automatically** from the **data**?

What are (convolutional) neural networks?



Trainable filter (convolutional layer)

W_{00}	$W_{01}b$	W_{02}
W_{10}	W_{11}	W_{12}
W_{20}	W_{21}	W_{22}

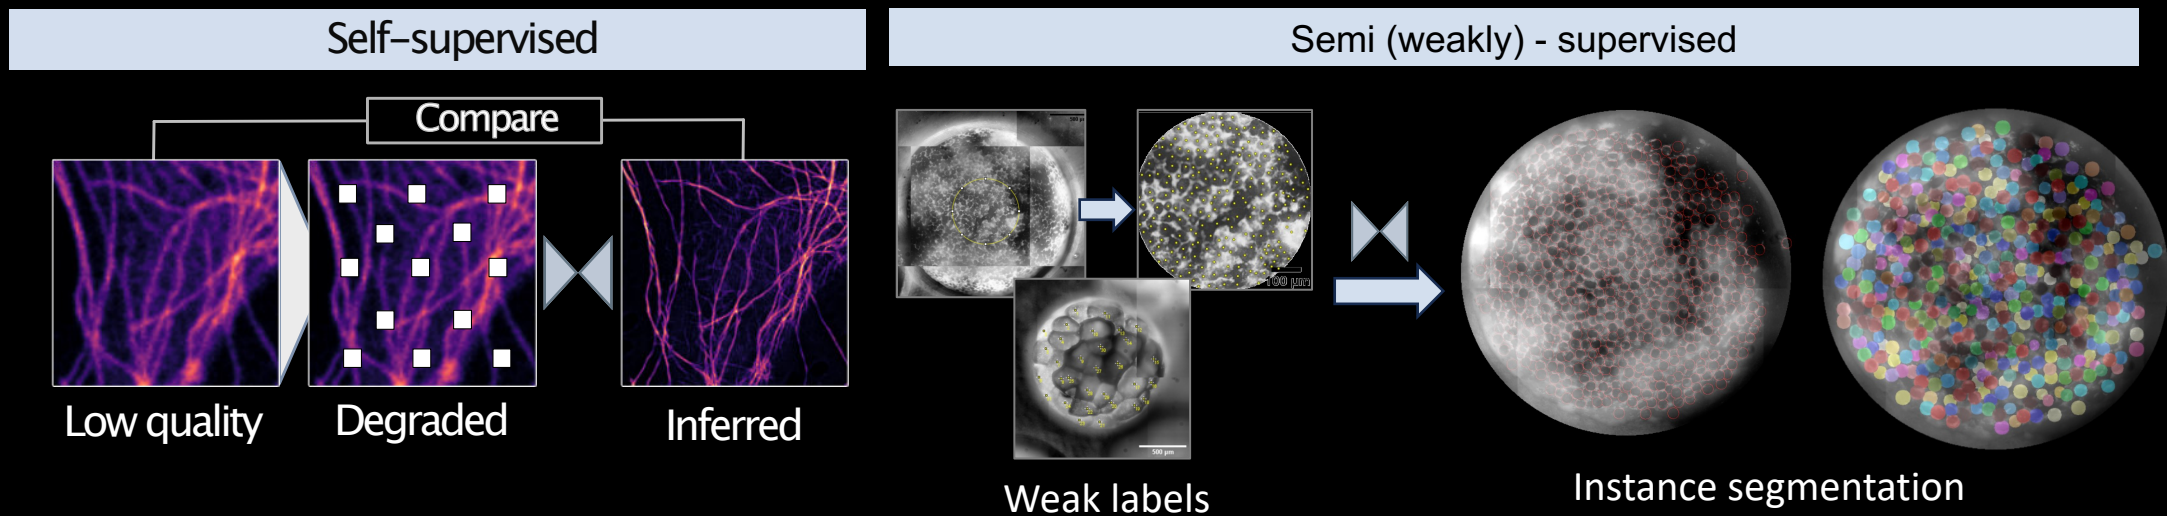

Convolutional kernel: determines de feature to enhance

Note: Each unit in the kernel will have its weight, but each convolutional filter will also have a bias:


$$C(x) = W \otimes x + b,$$

where x is an input image


The learning process




Supervised CNNs training



Supervised CNNs training



Supervised CNNs training




Input image

*

W_{00}	W_{01}	W_{02}
W_{10}	W_{11}	W_{12}
W_{20}	W_{21}	W_{22}

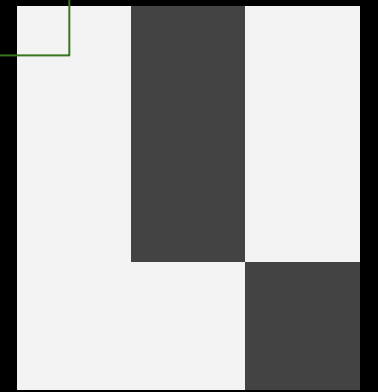
Trainable filter
(convolutional layer)

+ b =



Output

Stop training




Ground truth

Supervised CNNs training: backpropagation

- Optimization → Gradient descent

- Gradients computation → Backpropagation (use the chain rule for derivatives):

After each forward pass through the network, a backward pass is performed to adjust the model's parameters (weights and biases) according to the error made by the output of the network.





Loss function: quantitative measure of the error

Learning rate: proportion used to update the parameters on each pass

Most used loss functions:

- Mean Squared Error (MSE or L1)
- Mean Absolute Error (MAE or L2)
- Binary Cross Entropy (Categorical cross entropy)


Training in progress...




Introduced in 1960s
Popularized by Rumelhart, Hinton and Williams in "Learning representations by back-propagating errors", 1989

Supervised CNNs training

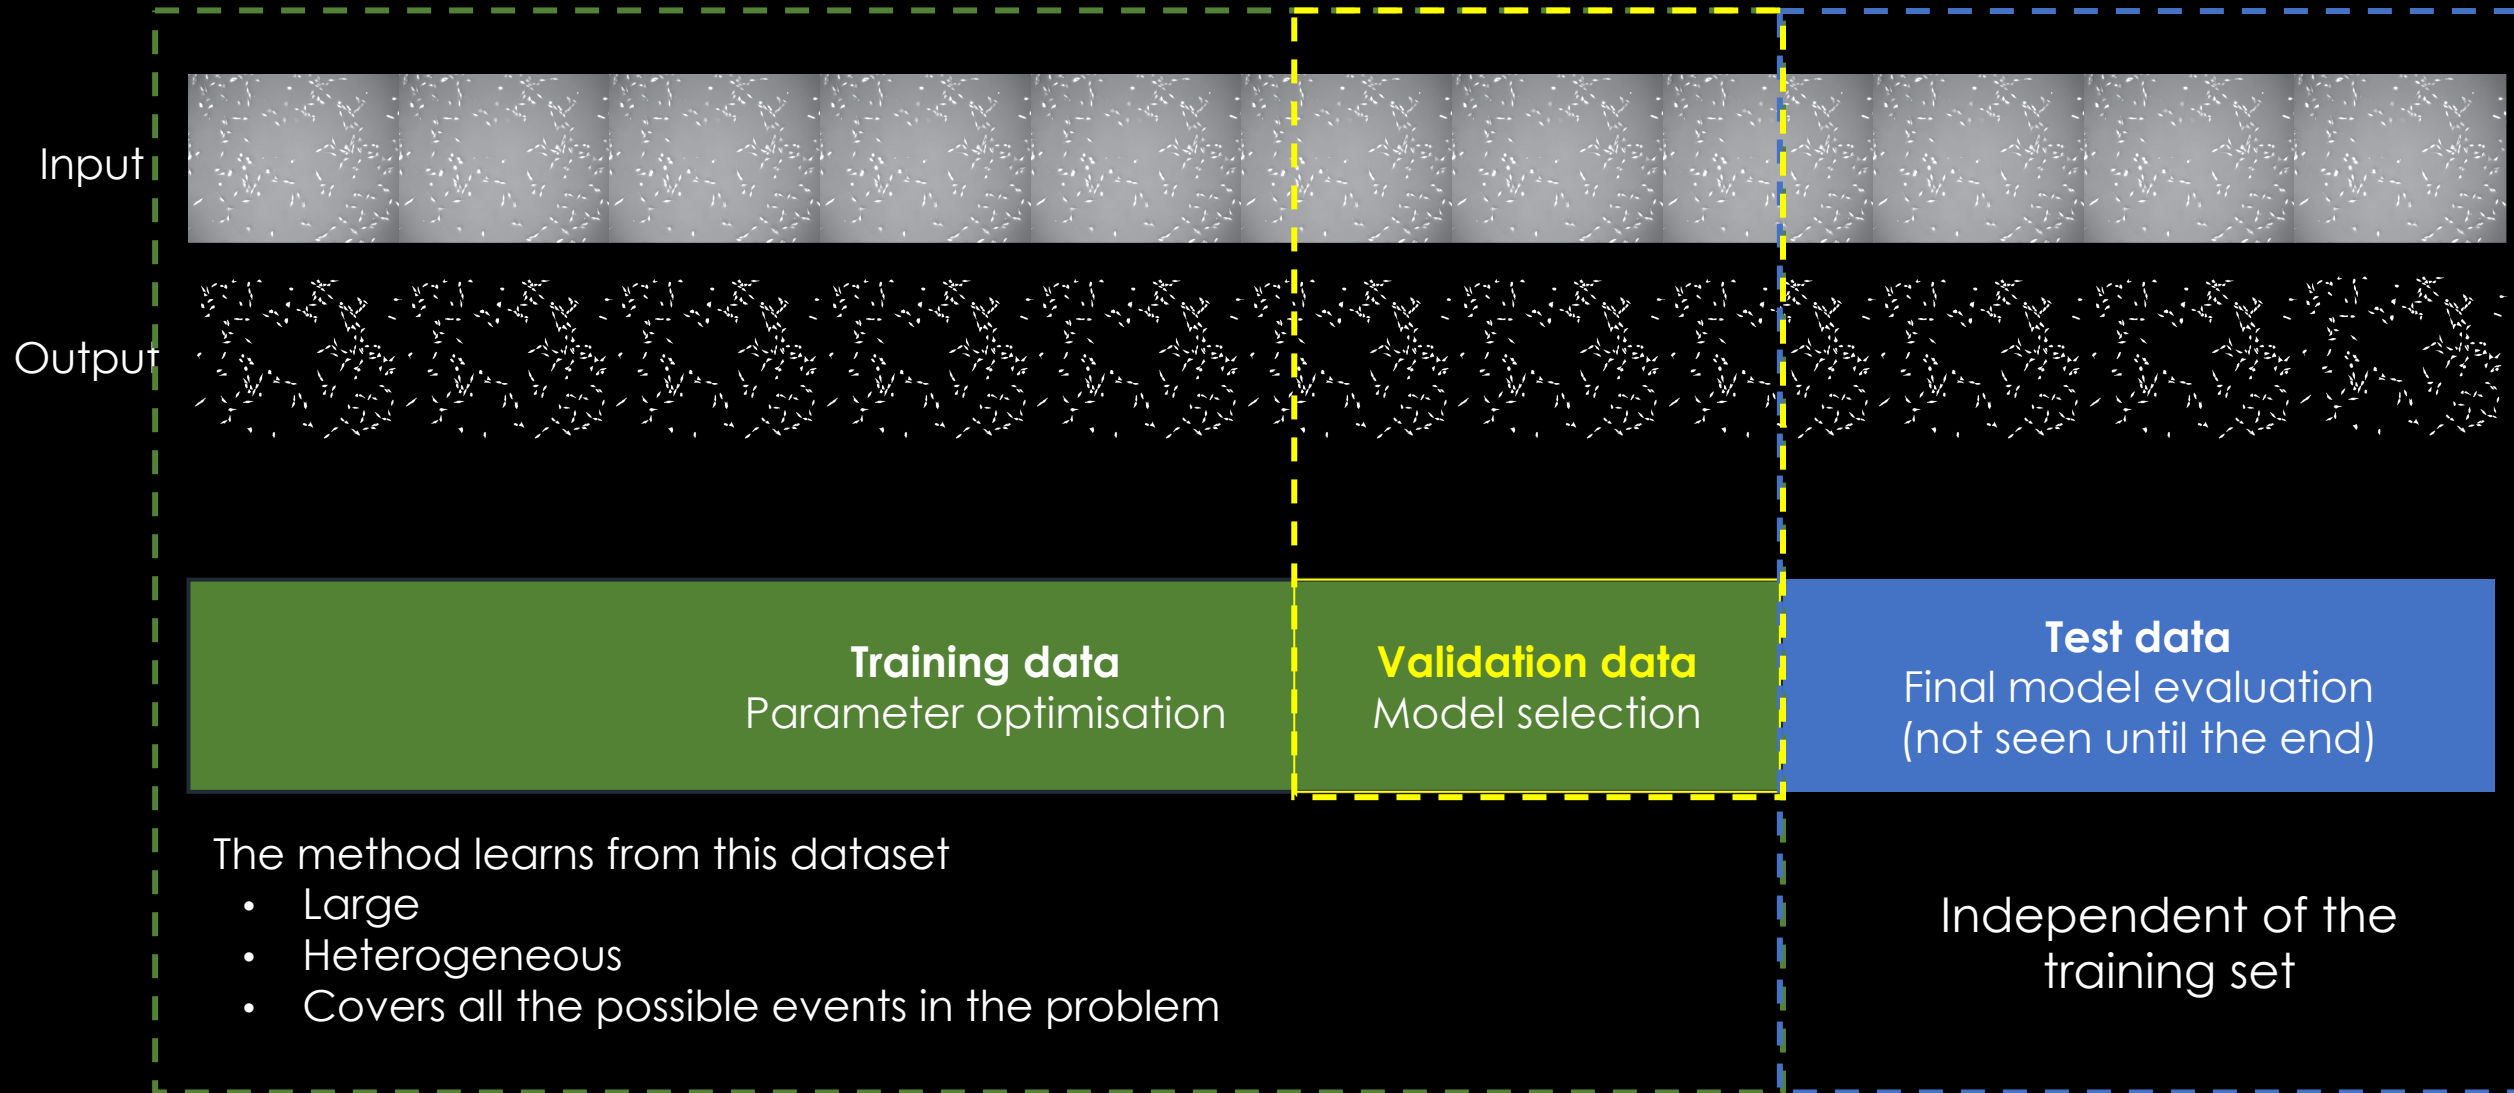
Pairs of inputs and desired outputs (i.e., ground truth)




Training a neural network



Training a neural network: data




Training a neural network: data




Training a neural network: data

If the method fails in the validation data, then it is called **overfitting**.



If the test processing is as good as in the training data, then we say that it can

generalize.





You: how do I prevent overfitting?

Esti & Wei: more data 😊

You: 🌀 🤔


Training a neural network: data

The bigger the data the better
→ Cover a real scenario




Data augmentation


All of them are cats, indeed, the same cat




No doubt, it is a fox or an airplane



It is always the number 6



Most common strategies to augment data in image classification



Training a neural network: data augmentation in microscopy


- I. The **ground truth** also needs to be augmented with **the same transformations**.
- II. Image transformations need to **preserve the meaning and biophysical properties** of the data.

Geometrical transformations

Linear transformations (preserve shape)

- Rotation
- Translation

Linear transformations



Original patch

Rotation + Shift


Rotation



Non-linear (elastic) transformations (shape changes)

- Zooming
- Shearing

Non-linear transformations



Original patch

Zoom


Shearing



Signal artifacts:

- Noise
- Contrast
- Blurring

Adding noise

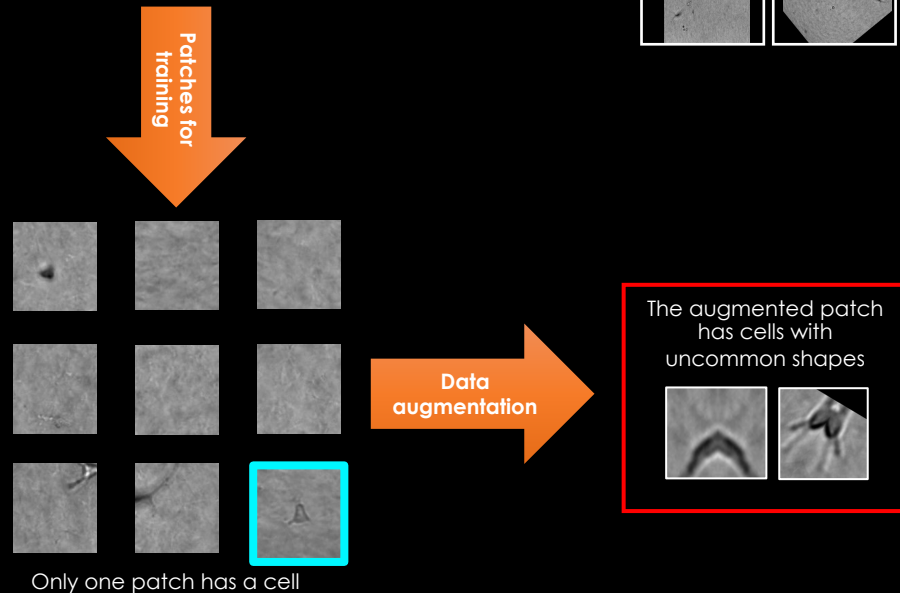



Original image

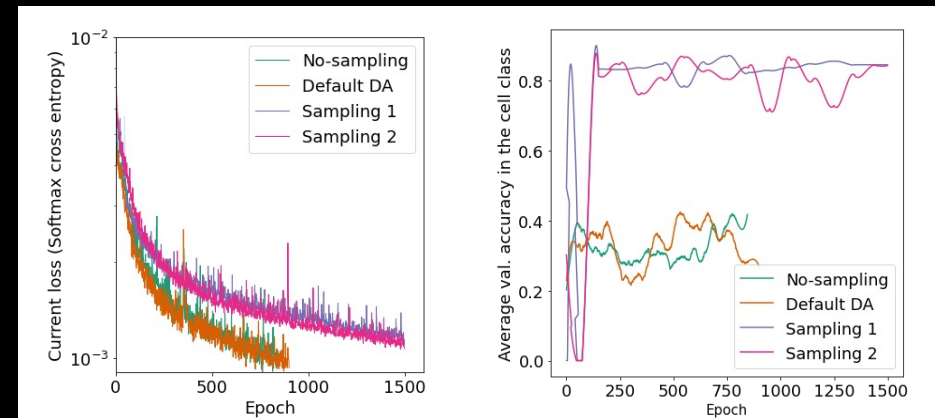
Noisy

Segmentation with CNNs: patches and data augmentation

Make sure that artifacts are not introduced when augmenting the patching



Results on the learning process with different strategies for data augmentation



Training a neural network: data

Data collection (&curation) is expensive.






Small datasets for bioimage analysis

Fine tuning

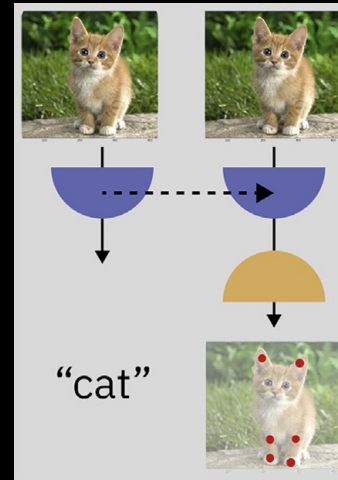
Pretrained model for boundary segmentation



Same task, similar features, different data distribution



Transfer learning



Classification



Segmentation

Classification



Pose estimation

Deep learning systems

Model architecture

- U-Net
- ResNet
- MobileNetV2
- (cycle)GANs
- DenseNets

Loss function

- Mean Squared Error
- Mean Absolute Error
- (binary) cross-entropy
- Focal loss
- Dice loss

Optimisation

- Stochastic gradient descent
- ADAM

Data (curation)

- Labelled masks
- Keypoints&landmarks
- Paired images (high&low SNR)
- Bounding boxes
- Tracks

Influence the final performance and behaviour of your system

Deep learning systems

Model architecture

Loss function

Optimisation

Data



Each configuration affects the inference speed, training data requirements, memory demands.



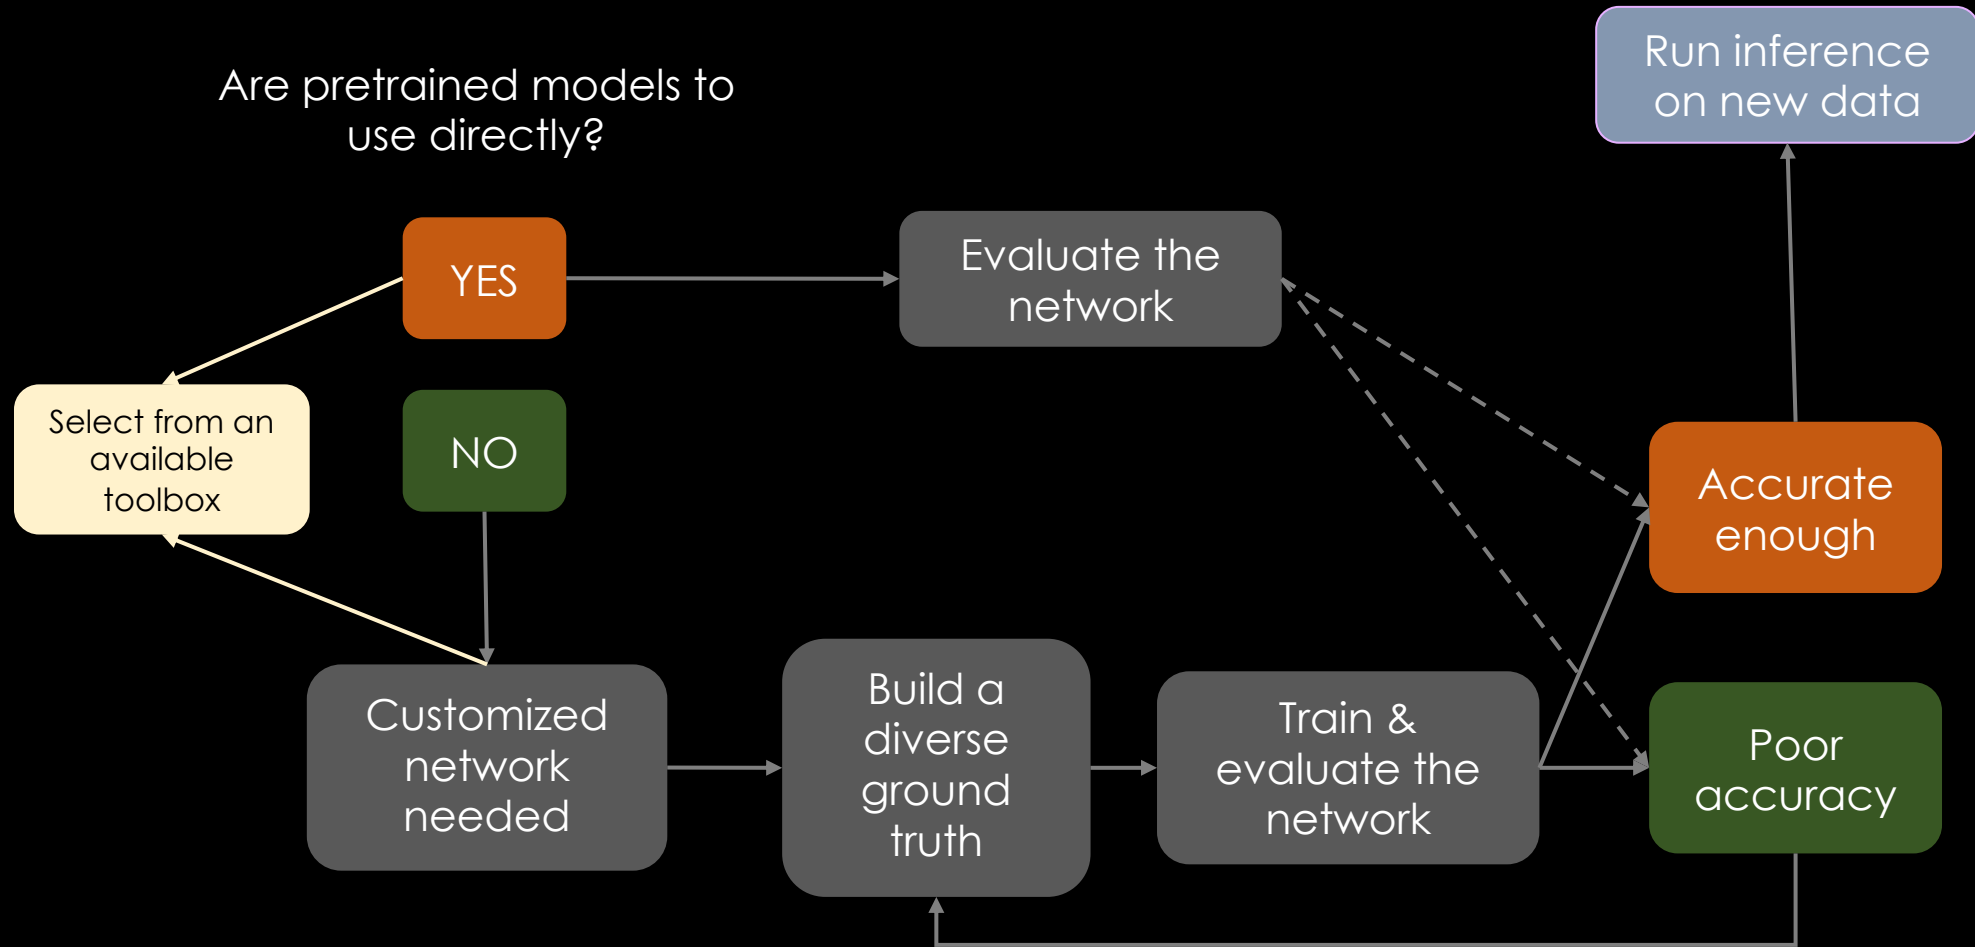
Task specific:

- 2D, 3D, time lapse, multichannel.
- Image-to-image vs. image-to-vector processing.
- Criterion needed to learn: regression vs classification.

Deep learning systems

Identify the sample type(s) and features you need to analyse

Questions to consider:
What features give me the info I need?
Live or fixed?
Highly accurate?
Multichannel?





Segmentation

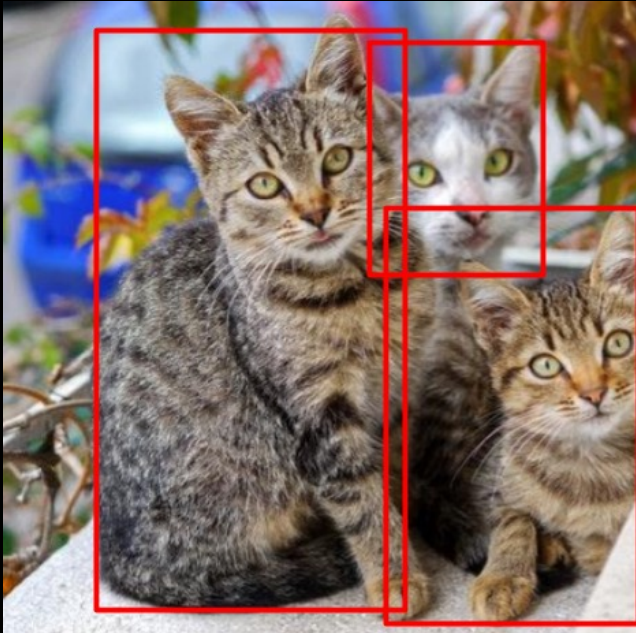
Segmentation

Formal definition: partitioning of the image domain Ω into several (usually disjoint) regions Ω_i


$$\Omega = \bigcup_i \Omega_i, \quad \Omega_i \cap \Omega_j = \emptyset, \forall i \neq j$$




Segmentation




Detection




Binary segmentation



Instance segmentation
(Detection + segmentation)



Semantic segmentation




Panoptic segmentation
(Instance + Semantic segmentation)

Segmentation


The information is partitioned in different segments to simplify its representation into something that is easier to analyze

Cell counting




F. Lux & P. Matula, arXiv, 2020

Morphology assessment




Data: Cell Tracking Challenge (Ulman, V., et al., Nat Methods 2017), Training: João Luis Soares Lopes (EPFL)

Determine anatomical regions (telencephalons) to measure cell activity (proliferating pHH3+ cells)



Thomas Naert, Development, 2021


Segmentation: Thresholding is the most basic form of obtaining binary images



Segmentation with CNNs: U-Net encoder-decoder for binary segmentations


Skip connections:

Take the output of each level in the encoder path and copy it with the input of the decoder path. It helps preserving high resolution details during decoder process.



Segmentation: alternative strategies



Weighted loss functions



$$w(\mathbf{x}) = w_0 \cdot \exp\left(-\frac{(d_1(\mathbf{x}) + d_2(\mathbf{x}))^2}{2\sigma^2}\right)$$

Olaf Ronneberger, et al., arXiv 2015

Use different labels



Wolny et al., eLife 2020






Fig. 4: Two segmentation results ($\tau = 0.5$) for DSB2018. See Fig. 2 caption for legend.

https://github.com/maweigert/neubias_academy_stardist
Schmidt, Weigert et al 2018

Instance segmentation with ROI classification and segmentation



K. He et al., arXiv 2018




Limitations of binary image segmentation

- Will not work with dense, packed or clustered objects.
- \rightarrow Additional labels to split independent objects
- Overlapping objects cannot be represented in one single mask
- ROIs need to be predefined and do not scale well to cellular shapes


Segmentation: learn deterministic features rather than discrete labels

StarDist



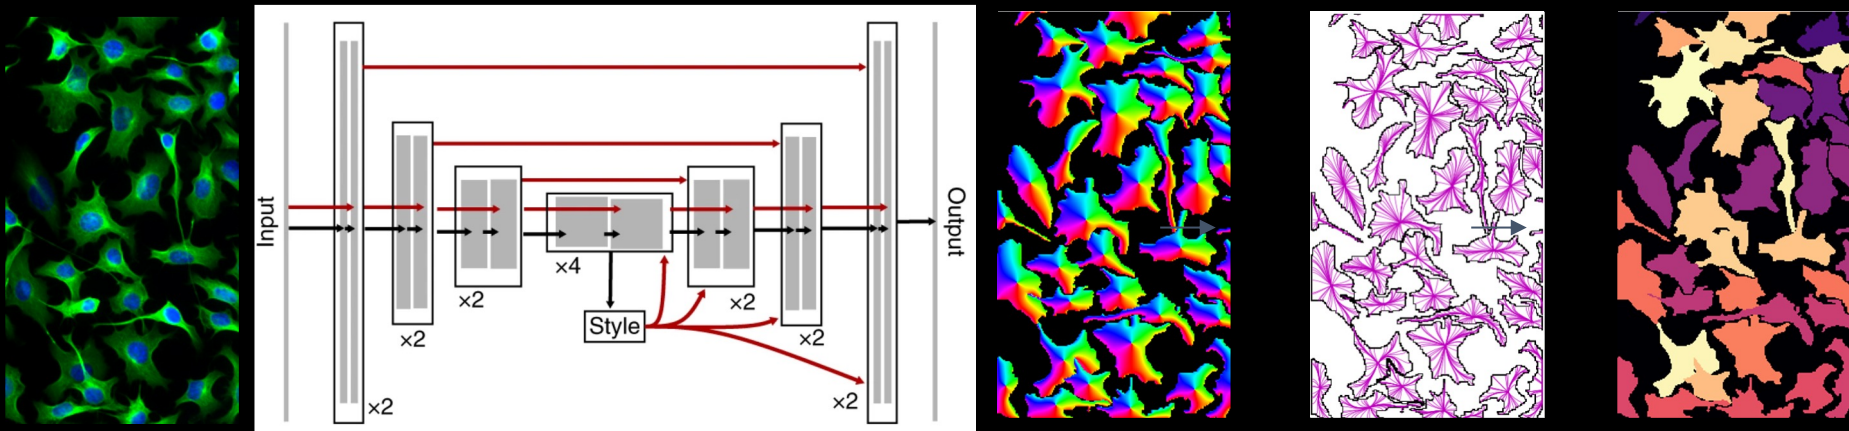
https://github.com/maweigert/neubias_academy_stardist
Schmidt, Weigert et al 2018

Cellpose



<https://www.cellpose.org>

Stringer, Wang,
Michaelos, Pachitariu,
Nature Methods, 2021


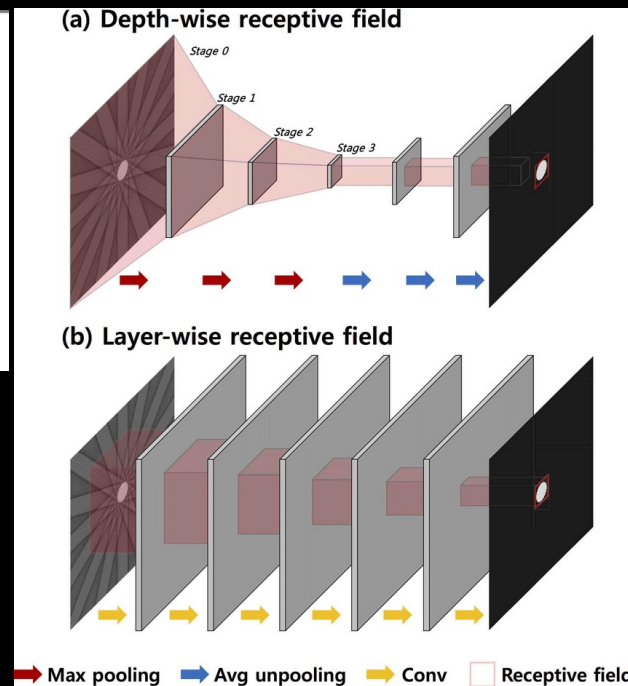





Segmentation with CNNs: image preparation and features

Biomedical images can get really large (up to TB for electron microscopy) → GPU memory a major limitation

Divide images into patches → increase the training data variability

How big? → Receptive field of the network → it needs to have enough information to learn and discriminate





TissueNet,
Greenwald, Miller et al 2021


Segmentation with CNNs: image preparation and features

Image processing task: segment cells in mitosis



Result: empty masks



Training patch size for StarDist



Receptive field of StarDist: "What the network sees to determine the value of one pixel"




Original videos: 4x4 field of views (x63)

Pixel size = 0.108 $\mu\text{m}/\text{pixel}$

Schmidt, Weigert et al 2018



Segmentation with CNNs: image preparation and features

A way to understand could be... asking what is the perfect distance to decipher the scenes of Claude Monet's art
→ enough as to get the context with still meaningful details



Segmentation with CNNs: image preparation and features

A way to understand could be... asking what is the perfect distance to decipher the scenes of Claude Monet's art
→ enough as to get the context with still meaningful details




Garden path at Giverny



The house among roses

Segmentation with CNNs: image resolution and its effect in segmentation



Segmentation with CNNs: image resolution and its effect in segmentation

Basic definition: Number of pixels in an image



Pixel size: physical size (length and width) covered by one pixel

16 cm

16 cm



8x8



64x64




128x128

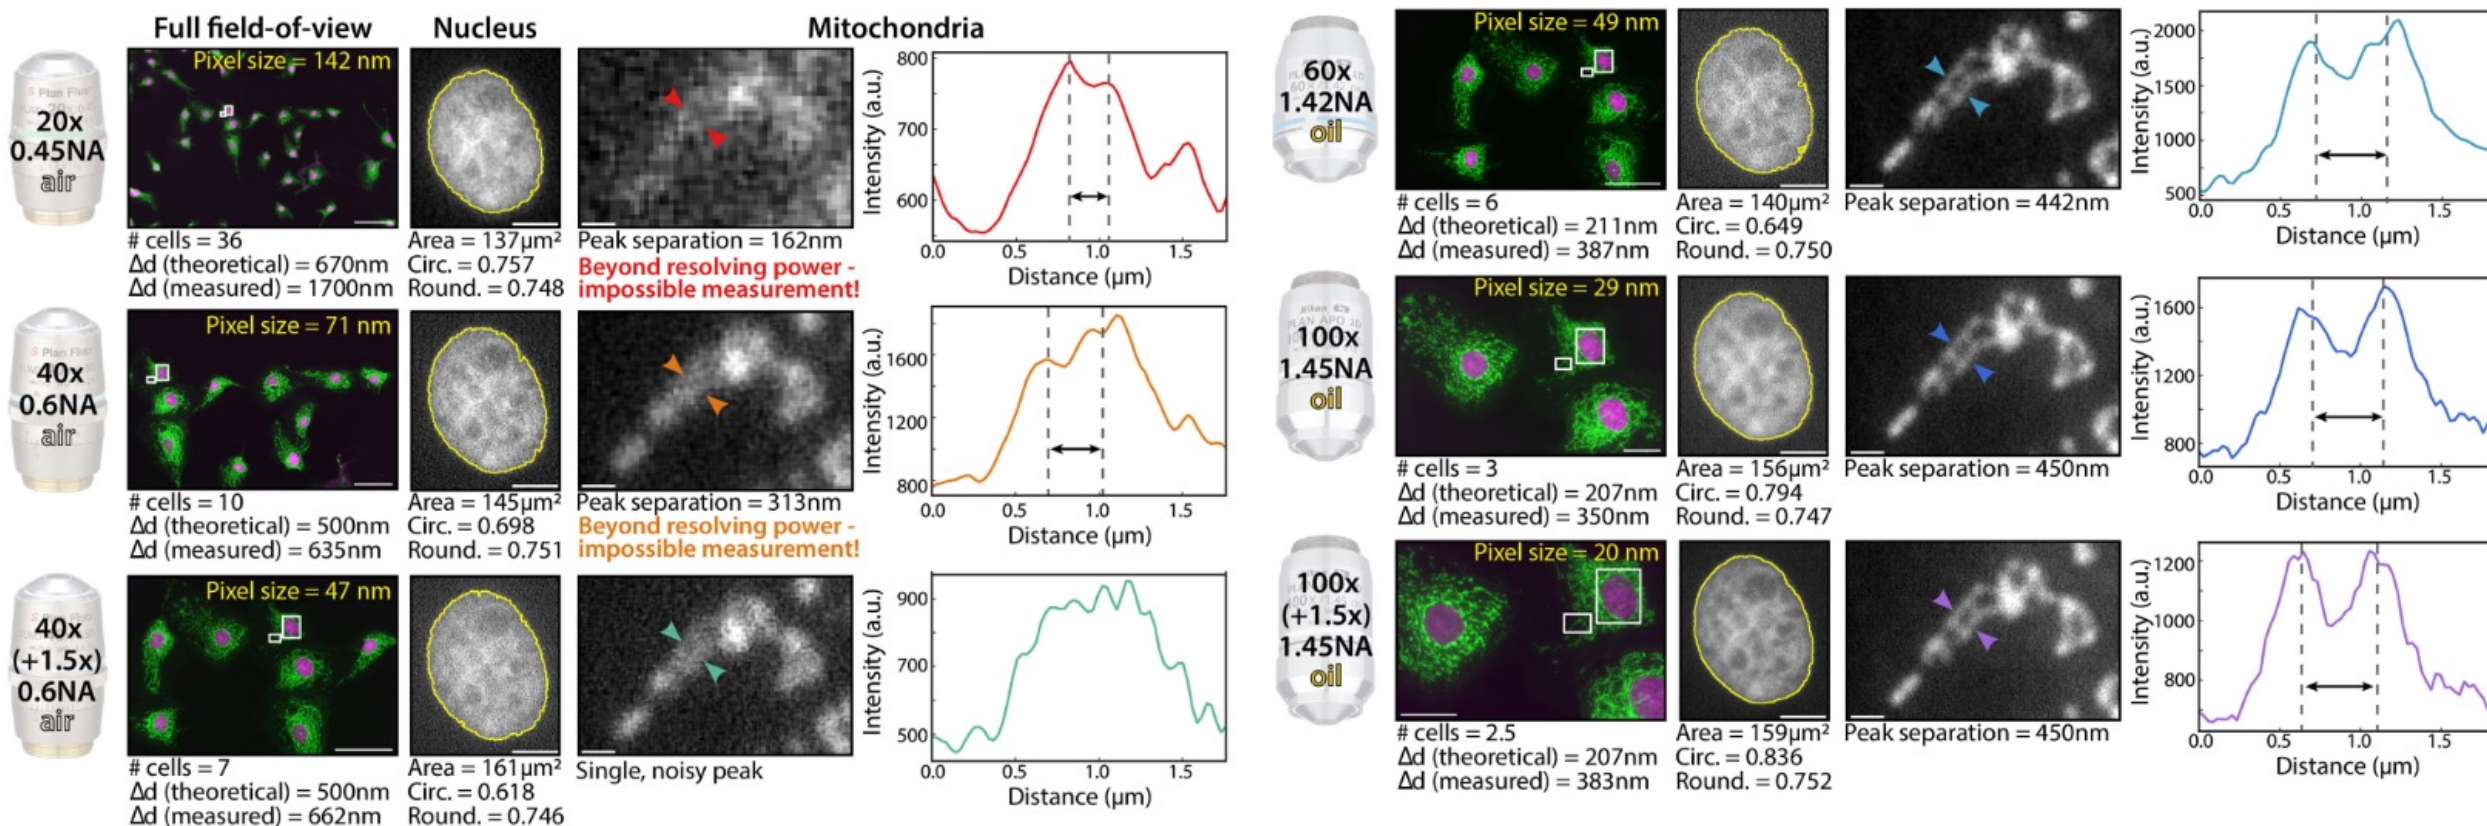


512x512

- The pixel size influence the amount of details for a given field of view.
- Detectors also limit the size of the finest detail that we can acquire




Segmentation with CNNs: image resolution and its effect in quantification



Segmentation with CNNs: image preparation and features

Intensity values vary with the physical properties of the data, the calibration of imaging devices or the natural variability of the sample


Non-normalized images




MNIST data: black and white pixels

0 1 2 3 4 5 6 7 8 9
0 1 2 3 4 5 6 7 8 9
0 1 2 3 4 5 6 7 8 9
0 1 2 3 4 5 6 7 8 9
0 1 2 3 4 5 6 7 8 9
0 1 2 3 4 5 6 7 8 9
0 1 2 3 4 5 6 7 8 9

Before normalization



After percentile normalization

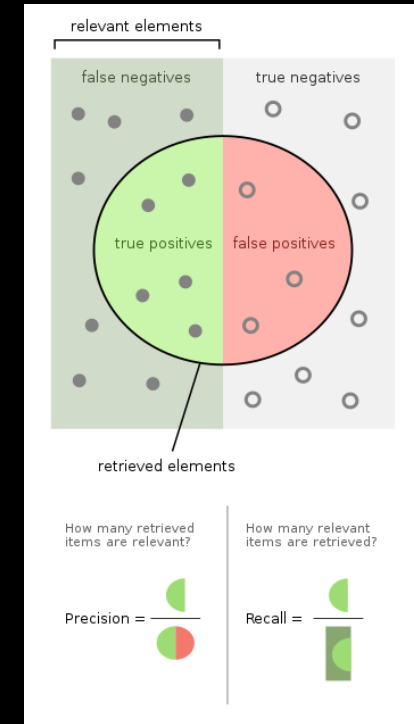
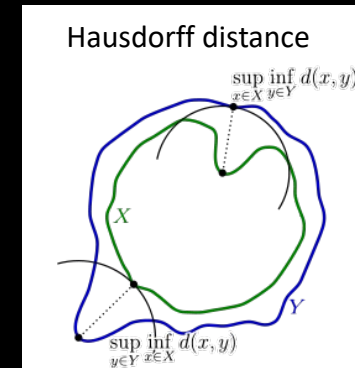
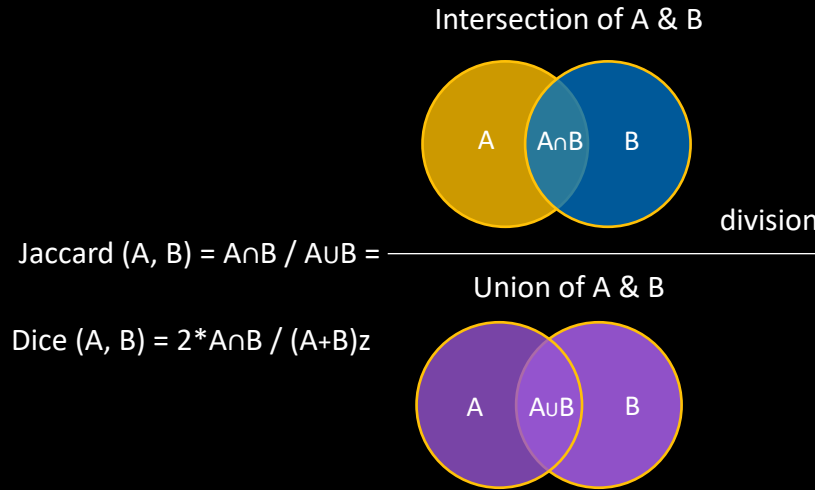
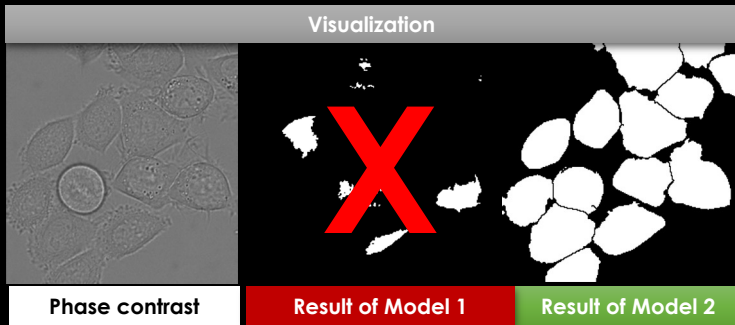


Common strategies:

- Intensity projection: Clip the dynamic range of values to the $[0, 1]$ range
- Standardization with the mean and standard deviation
- Percentile projection (common in fluorescence): remove outliers (i.e., noise and artifacts) from the intensity distribution (extremes in the tails) and clip to the $[0, 1]$ range.
- Normalize w.r.t. the entire population (training data)


Evaluation of the model performance → Accuracy

You need to verify that the model is doing precisely what you want



Quantify the accuracy

- Precision, recall, F1 (= $(2 * \text{precision} * \text{recall}) / (\text{precision} + \text{recall})$)
- Jaccard index / Dice coefficient
- Hausdorff distance
- Mean Squared Error (L2)
- Structural similarity index (SSIM)
- Biologically relevant measures (cell densities, fluorescence intensities, diameters)



https://github.com/maweigert/neubias_academy_stardist
Schmidt, Weigert et al 2018

Biological relevance of the segmentation results

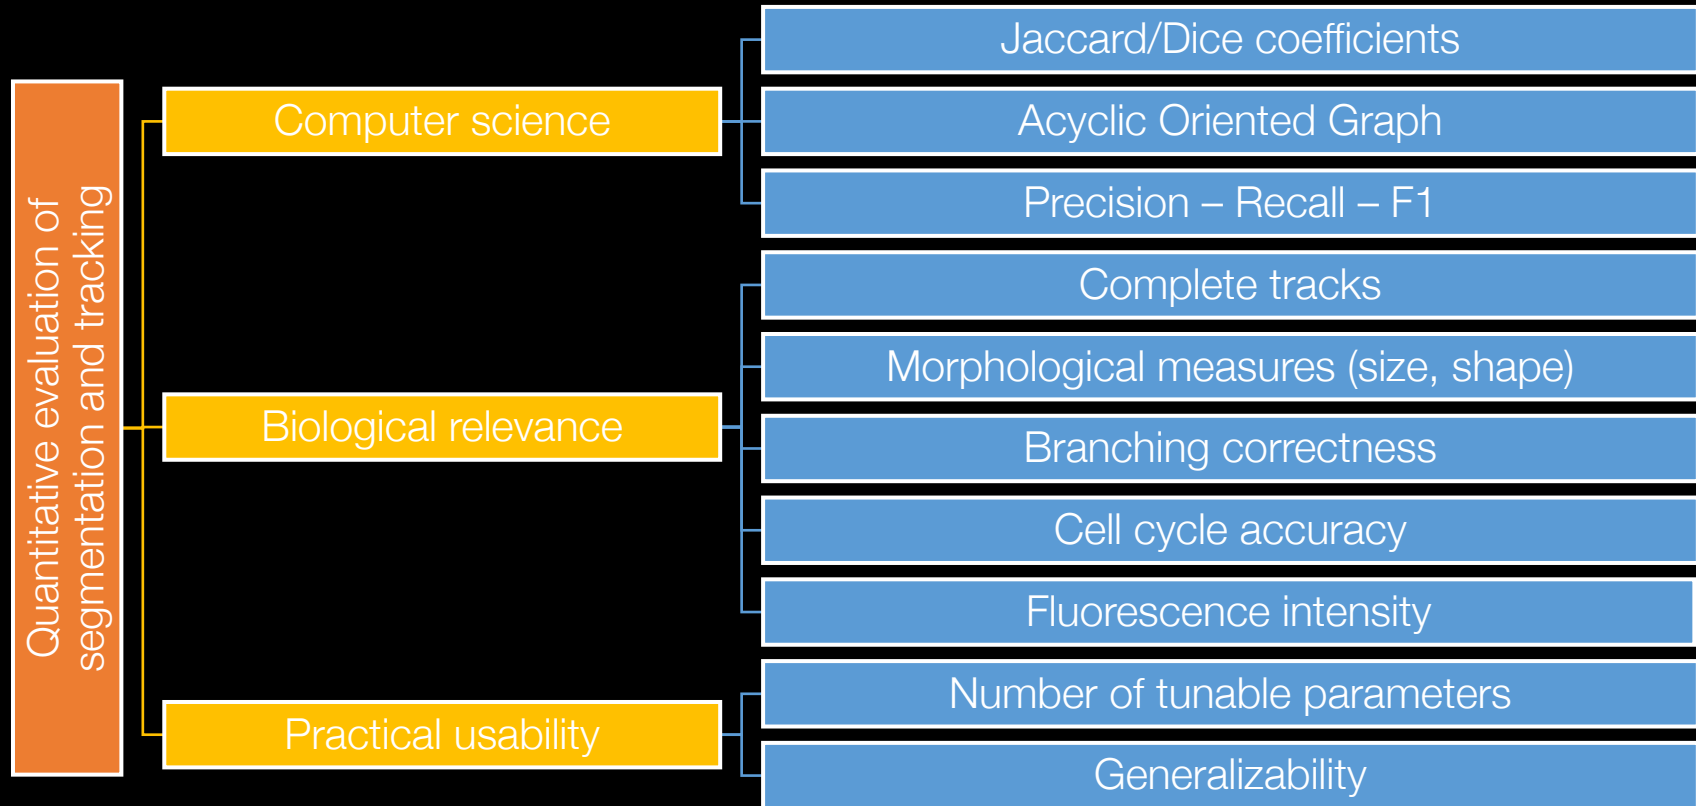


 Segmentation result

 Real object

 Accuracy (overlap)

 Length, curvature, diameter



Metrics reloaded: Recommendations for image analysis validation

Lena Maier-Hein, Annika Reinke, Patrick Godau, Minu D. Tizabi, Florian Büttner, Evangelia Christodoulou, Ben Glocker, Fabian Isensee, Jens Kleesiek, Michal Kozubek, Mauricio Reyes, Michael A. Riegler, Manuel Wiesenfarth, A. Emre Kavur, Carole H. Sudre, Michael Baumgartner, Matthias Eisenmann, Doreen Heckmann-Nötzel, A. Tim Radsch, Laura Acion, Michela Antonelli, Tal Arbel, Spyridon Bakas, Arriel Benis, Matthew Blaschko, M. Jorge Cardoso, Veronika Cheplygina, Beth A. Cimini, Gary S. Collins, Keyvan Farahani, Luciana Ferrer, Adrian Galdran, Bram van Ginneken, Robert Haase, Daniel A. Hashimoto, Michael M. Hoffman, Merel Huisman, Pierre Jannin, Charles E. Kahn, Dagmar Kainmueller, Bernhard Kainz, Alexandros Karargyris, Alan Karthikesalingam, Hannes Kenngott, Florian Kofler, Annette Kopp-Schneider, Anna Kreshuk, Tahsin Kurc, Bennett A. Landman, Geert Litjens, Amin Madani, Klaus Maier-Hein, Anne L. Martel, Peter Mattson, Erik Meijering, Bjoern Menze, Karel G.M. Moons, Henning Müller, Brennan Nihyporuk, Felix Nickel, Jens Petersen, Nasir Rajpoot, Nicola Rieke, Julio Saez-Rodriguez, Clara I. Sánchez, Shrayya Shetty, Maarten van Smeden, Ronald M. Summers, Abdel A. Taha, Aleksei Tulpin, Sotirios A. Tsafaris, Ben Van Calster, Gaël Varoquaux, Paul F. Jäger

(a) VARIOUS PITFALLS RELATED TO CHOICE OF VALIDATION METRIC

INAPPROPRIATE CHOICE OF THE PROBLEM CATEGORY

1 object detected ✗ 3 objects detected ✓

DSC = 0.92 >> DSC = 0.79

Example - object detection confused with semantic segmentation: DSC is strongly biased towards single objects and is therefore not appropriate for measuring the detection of multiple objects.

POOR METRIC SELECTION

DSC = 0.80 >> DSC = 0.50

Example - neglecting the small size of structures: Single-pixel differences can hugely impact the metric scores, which is especially relevant given high inter-rater variability and the non-deterministic nature of AI algorithms.

POOR METRIC APPLICATION

Example - inappropriate aggregation scheme: Hierarchical data structure is often neglected when aggregating metric values, which is especially important for different numbers of cases per variable.

DSC_{patient} = 0.6 ✓

(b) ADDRESSED BY PROBLEM-DRIVEN METRICS RELOADED FRAMEWORK

(1) PROBLEM FINGERPRINTING ENABLES MODALITY-INDEPENDENT METRIC SELECTION

Driving biomedical problem → Problem fingerprint → Metric selection

Problem fingerprints encapsulate relevant properties of a driving problem in a structured manner. ✓

Users are educated on pitfalls while being guided through the process of metric selection. ✓

(2) APPLICATION TO COMMON USE CASES DEMONSTRATES BROAD APPLICABILITY

Example input images:

Example outputs:

Inclusion criterion: classification at image, object or pixel level ✓

(3) ONLINE TOOL GUIDES THE USER

User-centric design. ✓

Lena Maier-Hein, Metrics reloaded: Recommendations for image analysis validation, <https://arxiv.org/abs/2206.01653>, arXiv 2022 (last update June 2023)

Fingerprint name	Fingerprint illustration	Fingerprint description
Image processing category identified by category mapping		Semantic segmentation (SemS): assignment of one or multiple category labels to each pixel.
Domain interest-related properties (selection)		
Particular importance of structure boundaries		The biomedical application requires exact structure boundaries. Example: segmentation for radiotherapy planning; knowledge of exact structure boundaries is crucial to destroy the tumor while sparing healthy tissue.
Particular importance of structure center (e.g., in cells, vessels)		Important: Overlap-based metrics do not measure shape agreement. In the case of complex shapes (high boundary-to-volume ratio) it is therefore typically advisable to set this property to TRUE.
Compensation for annotation imprecisions requested		The biomedical application requires accurate knowledge of structure centers. Example: cell centers are subsequently used for cell tracking and cell motion characterization, so false center movement should be suppressed.
...
Target structure-related properties (selection)		
Small size of structures relative to pixel size		Structures of the provided class are only a few pixels in size. Example: multiple sclerosis lesions in magnetic resonance imaging (MRI) scans.
High variability of structure sizes (within an image and/or across images)		The target structures vary substantially in size, such that some structures are several times the size of others. Example: polyps in colonoscopy screening, where some polyps are several times the size of others. Counterexample: large organs, such as the liver or the kidneys, which are relatively comparable in size across individuals.
...
Data set-related properties (selection)		
Presence of class imbalance		The class prevalences differ substantially. Example: In a screening application, the positive class (e.g., cancer) may occur extremely rarely. In this case, prevalence-dependent metrics, such as Accuracy, may be extremely misleading.
Non-independence of test cases		The test cases are hierarchically structured, indicating non-independence of test cases. Examples: multiple images of the same patient, hospital or video.
...
Algorithm output-related properties (selection)		
Possibility of algorithm output not containing the target structure(s)		The algorithm may yield outputs in which not all classes are present.
...


Considerations: The objective

Accuracy versus validity

Example:

High segmentation accuracy but
poor temporal consistency

→ Limit object tracking



Mistaken objective

Example:


Diagnosis of diabetic retinopathy.

Issues:

1. Discrepancy among doctors and non-valid majority voting
2. Hidden real objective → *“Should this patient see a doctor?”*

Considerations: Generalizability


DL models are extremely sensitive to pixel sizes (object size), imaging modalities, morphologies, cell types, fluorescence channels...



FluoC3DLMDA231



PhC-C2DL-PSC



Usiigaci

Generalizability is still an active and hot field of study with important open questions:

- Is it because we lack enough training data?
- Should we get deeper models?
- Is it possible to have one single model for a specific task regardless the data?

