MACHINE LEARNING

Course Recap.

Last Update: 21st December 2022

Prof. Dr. Shadi Albarqouni

Director of Computational Imaging Research Lab. (Albarqouni Lab.)

University Hopsital Bonn | University of Bonn | Helmholtz Munich

STRUCTURE

OVERVIEW.

WHAT IS MACHINE LEARNING (ML)?

Definition (Tom Mitchell)

A computer program is said to learn from experience E with respect to some class of tasks T, and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.

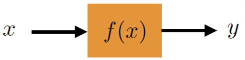

$$f(x) = \int_0^{\frac{\pi}{2}} cos(x)e^{-2\gamma x}dx$$

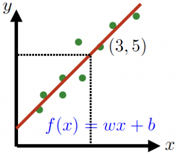

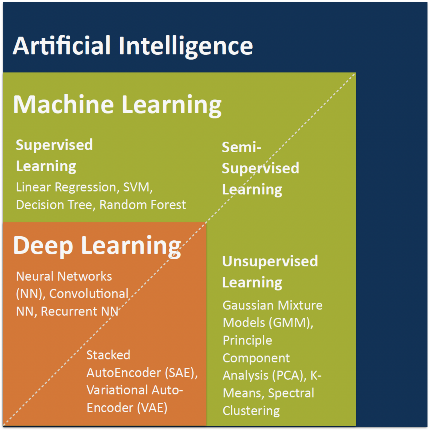

TASKS - CLASSIFICATION

$f($ 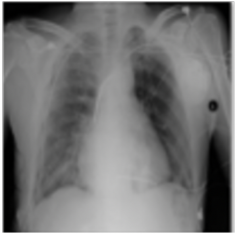 $;\theta)$

$$I_q \in \mathbb{R}^{m \times n}$$

T$_1$: Classification

$$y_q \in \{c_1, c_2, \ldots, c_K\}$$

Tuberculosis Atelectasis

TASKS - ANOMALY DETECTION

$f($ 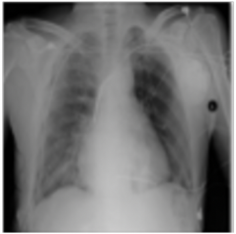 $;\theta)$

$$I_q \in \mathbb{R}^{m \times n}$$

T₂: Anomaly Detection

$$y_q \in \{c_N, c_A\}$$

Normal **Abnormal**

TASKS - REGRESSION

$f($ 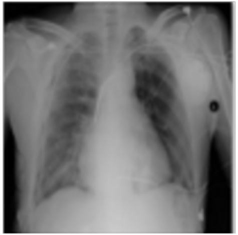 $; \theta)$

$$I_q \in \mathbb{R}^{m \times n}$$

T_3: Regression

$y_q \in \mathbb{R}$

45 yrs

TASKS - RETRIEVAL

$f($ 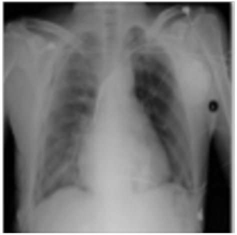 $;\theta)$

$$I_q \in \mathbb{R}^{m \times n}$$

T$_4$: Retrieval

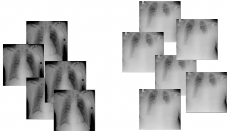

TASKS - SEGMENTATION

$f($ 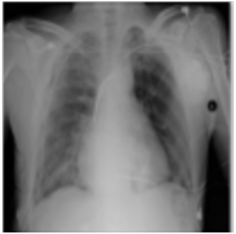 $;\theta)$

$I_q \in \mathbb{R}^{m \times n}$

T_5: Segmentation

$y_q \in \mathbb{R}^{m \times n \times K}$

ML EXPERIENCE

EXPERIENCE

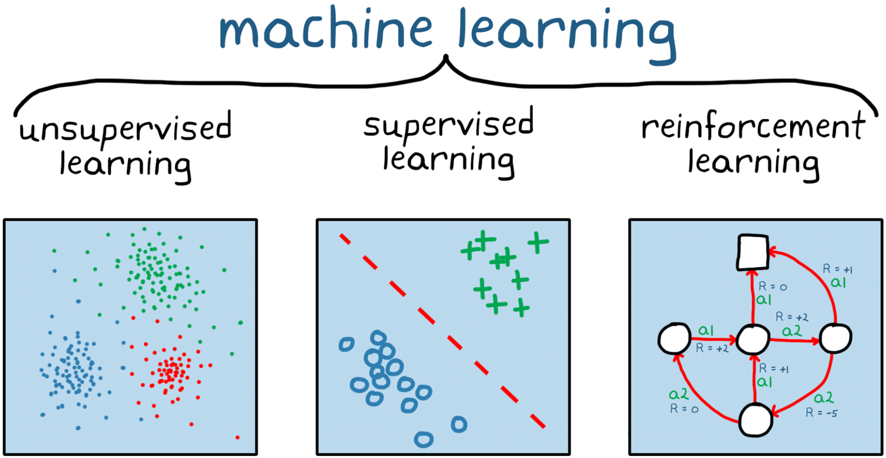

Source: https://www.mathworks.com/discovery/reinforcement-learning.html

EXPERIENCE

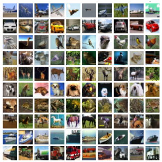

Training Set

EXPERIENCE


Pre-Processing
Feature Extraction
Dimensionality Reduction
Training
Loss
Validation
Evaluation Metrics
Testing

FEATURE EXTRACTION

FEATURE EXTRACTION

PRE-PROCESSING -- MOTIVATION

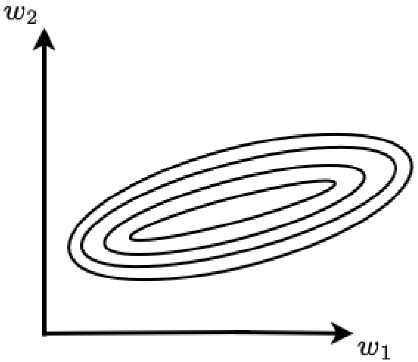

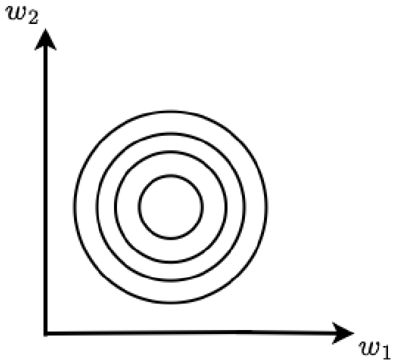

PRE-PROCESSING -- FEATURE SCALING

Normalization: It is the process of rescaling the values of all features to a range between 0 and 1.

$$z_i = \frac{x_i - \min(x)}{\max(x) - \min(x)}$$

Image Source:
https://mkang32.github.io/python/2020/12/27/feature-scaling.html

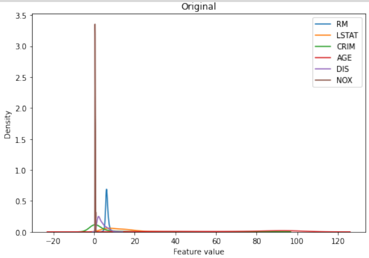

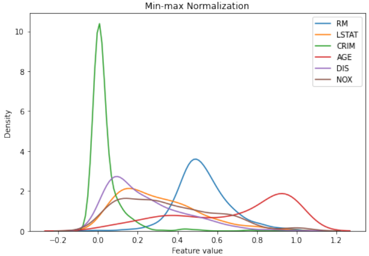

PRE-PROCESSING -- FEATURE SCALING

Standardization: It is the process of representing the data as a Normal distribution with a 0 mean and a unit (1) standard deviation.

$$z_i = \frac{x_i - \mu_x}{\sigma_x}$$

Image Source:
https://mkang32.github.io/python/2020/12/27/feature-scaling.html

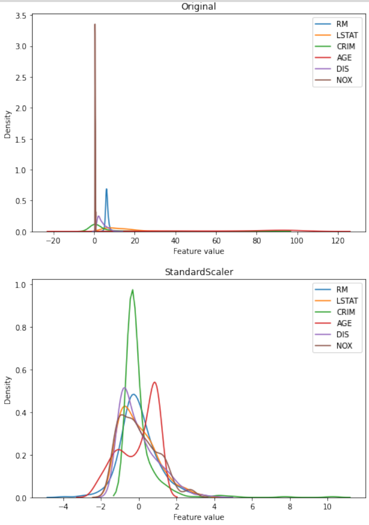

FEATURE EXTRACTION

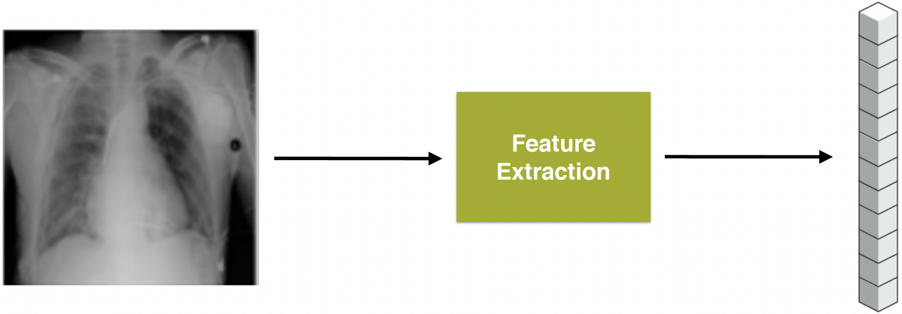

$$I_q \in \mathbb{R}^{m \times n}$$

$$x_q \in \mathbb{R}^D$$

FEATURE EXTRACTION

- Local Pixel Features (Binary, Spectra, e.g., SIFT, SURF, HoG, ...etc.)
- Global Pixel Features (Texture, SDM, ...etc.)
- Shape of Pixel Regions (Area, Perimeter, Centroids ...etc.)
- Basis sets (Haarlike, Bag of words, ...etc.)

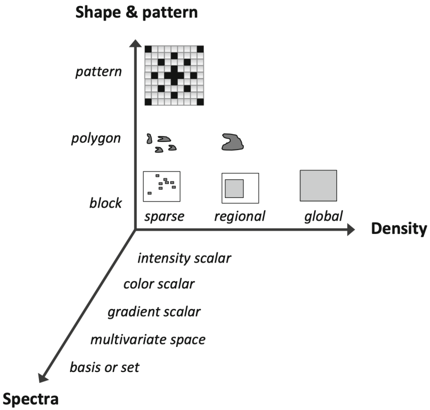

Adopted from Fig.5.1 in Krig, S., 2014. Computer vision metrics: Survey, taxonomy, and analysis (p. 508). Springer nature.

FEATURE EXTRACTION

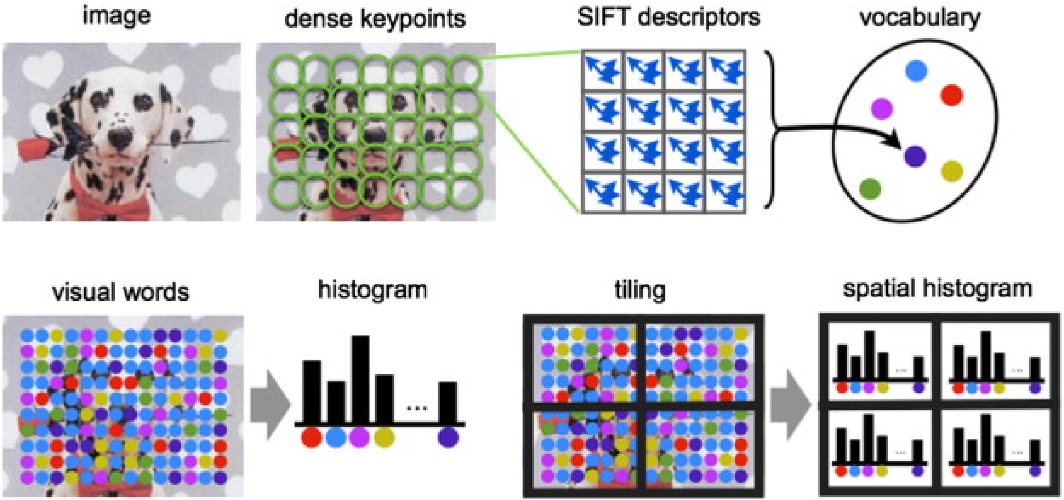

Adopted from Fig.1 in El-Gayar, M. M., and H. Soliman. "A comparative study of image low level feature extraction algorithms." Egyptian Informatics Journal 14.2 (2013): 175-181.

DIMENSIONALITY REDUCTION

DIMENSIONALITY REDUCTION

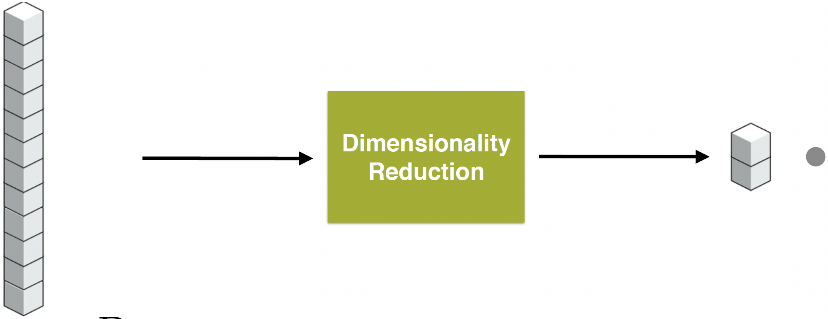

$x_q \in \mathbb{R}^D$ $d \lll D$ $x_q \in \mathbb{R}^d$

DIMENSIONALITY REDUCTION

Principal Compnent Analysis (PCA)

- is a statistical technique for reducing the dimensionality of a dataset
- linearly transform the data into a new coordinate system where (most of) the variation in the data can be described with fewer dimensions than the initial data
- The new coordinate system components are called Principal Components (PCs)

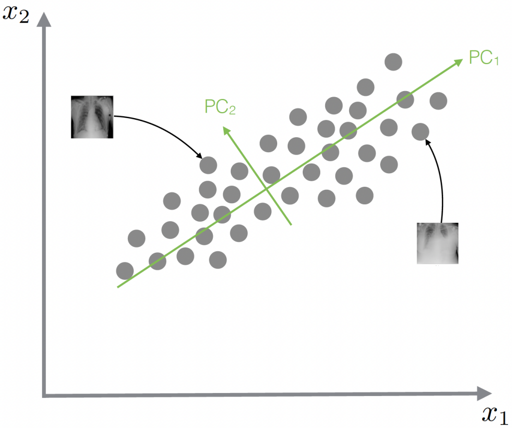

DIMENSIONALITY REDUCTION

Principal Component Analysis (PCA)

- is a statistical technique for reducing the dimensionality of a dataset
- linearly transform the data into a new coordinate system where (most of) the variation in the data can be described with fewer dimensions than the initial data
- The new coordinate system components are called Principal Components (PCs)

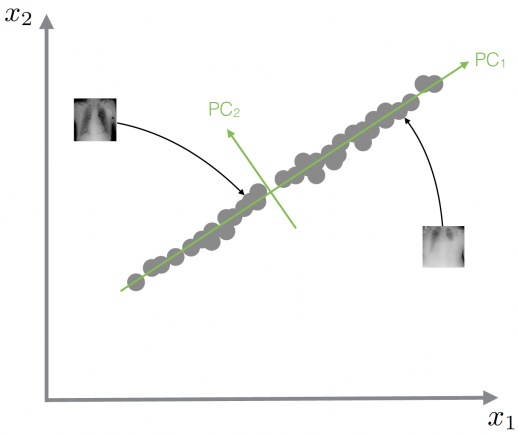

DIMENSIONALITY REDUCTION

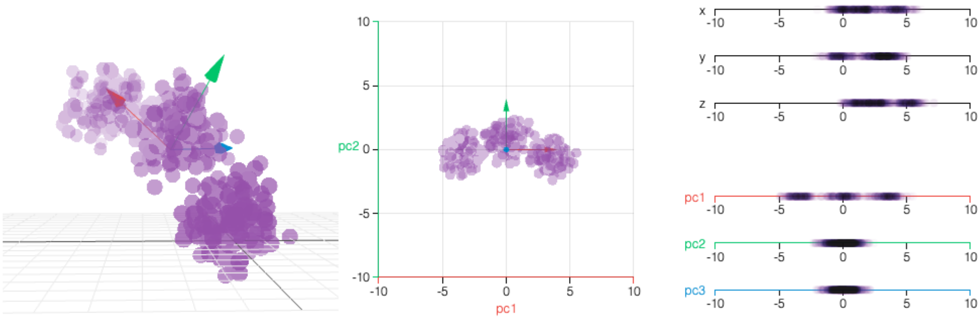

Source: **https://setosa.io/ev/principal-component-analysis/**

DIMENSIONALITY REDUCTION

Few example images
Each picture has 19x19 pixels (i.e. the feature space has 361 dims)

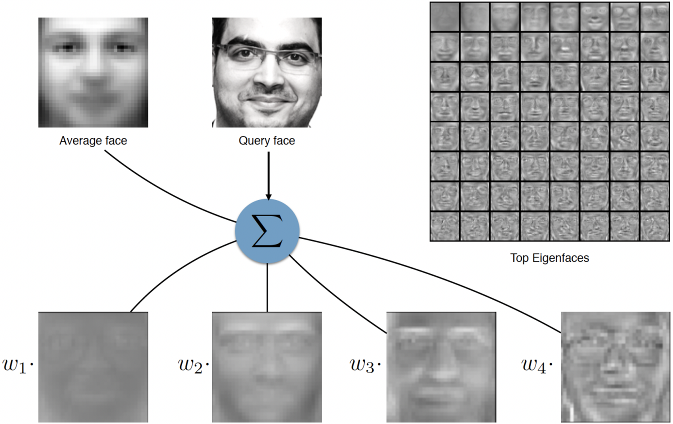

TRAINING

TRAINING


Pre-Processing
Feature Extraction
Dimensionality Reduction
Training
Loss
Validation
Evaluation Metrics
Testing

TRAINING

- **Supervised Learning**
  - Derive general rules from labeled examples
- Unsupervised Learning
  - Discover similarities within unlabelled data. Estimate their distribution
- Semi-Supervised Learning
  - Make use of both labeled and unlabelled data
- Reinforcement Learning
  - Make right decisions from the past experience

labeled examples:

$$\mathcal{D} = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), ..., (\mathbf{x}_N, y_N)\}$$

Input feature:

$$\mathbf{x} = (x_1, x_2, ..., x_d) \in \mathbb{R}^d$$

Predicted output:

$$\underbrace{y \in \{c_1, c_2, ..., c_K\}}_{classification} \quad \text{or} \quad \underbrace{y \in \mathbb{R}^k}_{regression}$$

Conditional distribution:

$$P(y|\mathbf{x}) \qquad f(x) = arg\,max_x P(y|\mathbf{x})$$

Joint distribution:

$$P(\mathbf{x}, y) \qquad P(y|\mathbf{x}) = \frac{P(\mathbf{x}, y)}{\sum_y P(\mathbf{x}, y)}$$

Complexity

TRAINING

$$\mathcal{D} = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), ..., (\mathbf{x}_N, y_N)\}$$

$$y_q \in \{c_1, c_2, ..., c_K\}$$

TRAINING

- **Supervised Learning**
  - Derive general rules from labeled examples
- **Unsupervised Learning**
  - Discover similarities within unlabelled data. Estimate their distribution
- Semi-Supervised Learning
  - Make use of both labeled and unlabelled data
- Reinforcement Learning
  - Make right decisions from the past experience

Unlabelled examples:

$$\mathcal{D} = \{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_N\}$$

Input feature:

$$\mathbf{x} = (x_1, x_2, ..., x_d) \in \mathbb{R}^d$$

Output (clusters):

$$\underbrace{\nabla = \{C_1, C_2, ..., C_K\}}_{clusters}$$

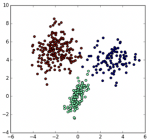

Mixture density:

$$f(\mathbf{x}) = \sum_{k=1}^{K} \pi_k f_k(\mathbf{x})$$

$$P(\theta|\mathbf{x}) = \sum_{k=1}^{K} \pi_k \mathcal{N}(\mathbf{x}; \mu_i, \Sigma_i)$$

TRAINING

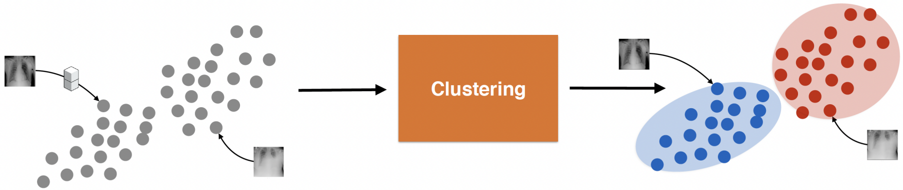

$$x_q \in \mathbb{R}^d$$

$$y_q \in \{c_1, c_2, \dots, c_K\}$$

TRAINING

- **Supervised Learning**
  - Derive general rules from labeled examples
- **Unsupervised Learning**
  - Discover similarities within unlabelled data. Estimate their distribution
- **Semi-Supervised Learning**
  - Make use of both labeled and unlabelled data
- Reinforcement Learning
  - Make right decisions from the past experience

Labeled & Unlabelled examples:

$$\mathcal{D} = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), ..., (\mathbf{x}_L, y_L), x_{L+1}, x_{L+2}, ..., x_{L+U}\}$$

Input feature:

$$\mathbf{x} = (x_1, x_2, ..., x_d) \in \mathbb{R}^d$$

Predicted Output:

$$y \in \{c_1, c_2, ..., c_K\}$$

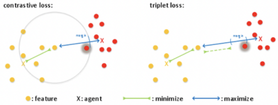

TRAINING

$$\mathcal{D} = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_L, y_L), x_{L+1}, x_{L+2}, \dots, x_{L+U}\}$$

$$y_q \in \{c_1, c_2, \dots, c_K\}$$

TRAINING

- **Supervised Learning**
  - Derive general rules from labeled examples
- **Unsupervised Learning**
  - Discover similarities within unlabelled data. Estimate their distribution
- **Semi-Supervised Learning**
  - Make use of both labeled and unlabelled data
- **Reinforcement Learning**
  - Make right decisions from the past experience

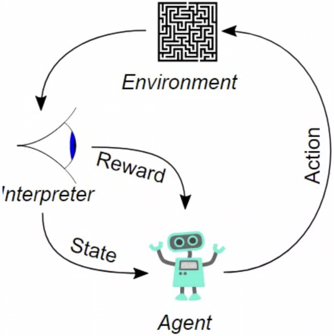

TRAINING

- **Bias**
  - Model Complexity
  - More data?
- **Variance**
  - Data variance
  - Weight decay
  - Dropout

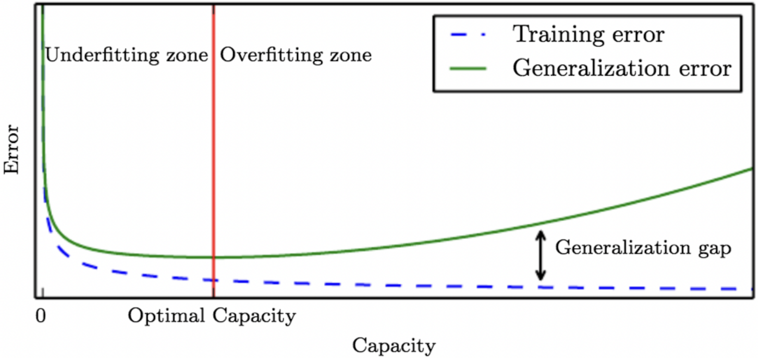

Source: **http://www.deeplearningbook.org/contents/ml.html**

VALIDATION

VALIDATION


Pre-Processing
Feature Extraction
Dimensionality Reduction
Training
Loss
Validation
Evaluation Metrics
Testing

VALIDATION

Non-Exhaustive Cross Validation

- Holdout method
- k-fold Cross Validation (k-fold CV)

Exhaustive Cross Validation

- Leave-one-out Cross Validation (LOOCV)
- Leave-p-out Cross Validation (LpOCV)

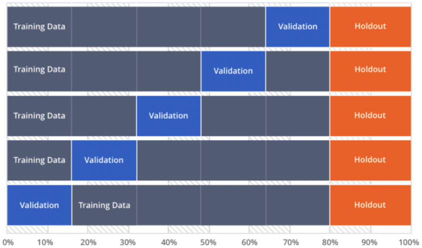

EVALUATION METRICS

EVALUATION METRICS

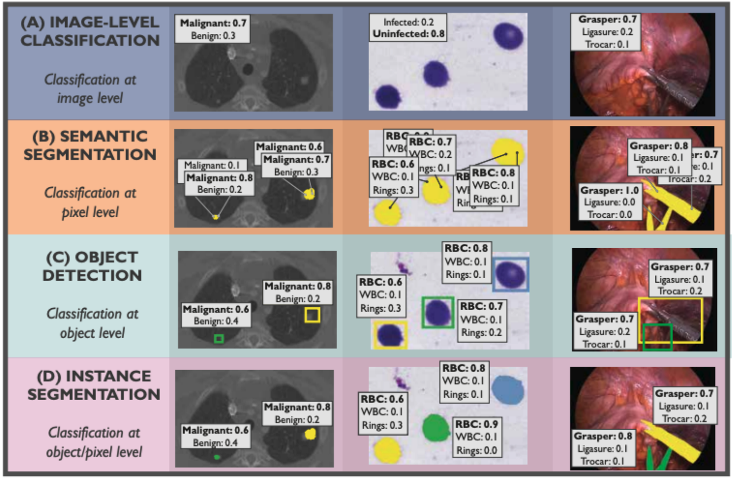

Source: Reinke, Annika, et al. "Common limitations of image processing metrics: A picture story." arXiv preprint arXiv:2104.05642 (2021).

EVALUATION METRICS

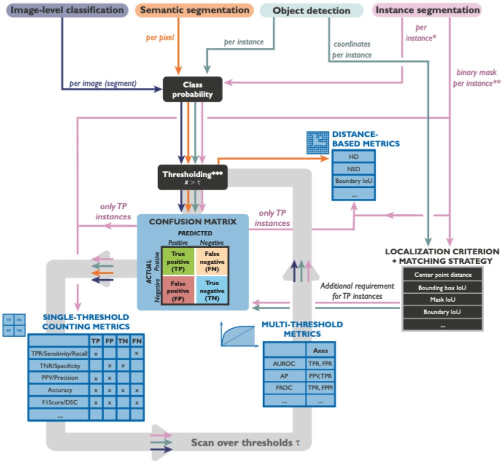

Source: Reinke, Annika, et al. "Common limitations of image processing metrics: A picture story." arXiv preprint arXiv:2104.05642 (2021).

EVALUATION METRICS -- SINGLE THRESHOLD

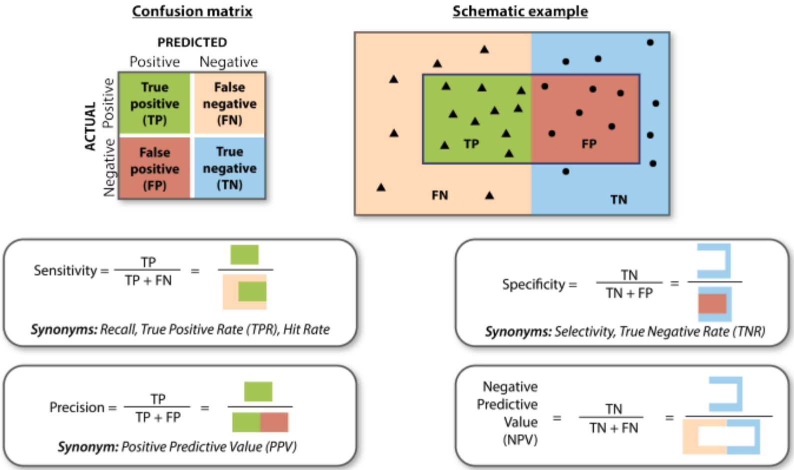

Source: Reinke, Annika, et al. "Common limitations of image processing metrics: A picture story." arXiv preprint arXiv:2104.05642 (2021).

EVALUATION METRICS -- SINGLE THRESHOLD

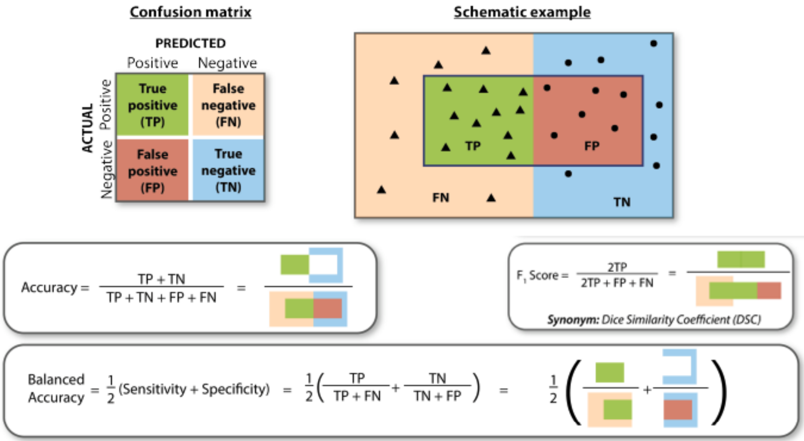

Source: Reinke, Annika, et al. "Common limitations of image processing metrics: A picture story." arXiv preprint arXiv:2104.05642 (2021).

EVALUATION METRICS -- MULTI-THRESHOLDS

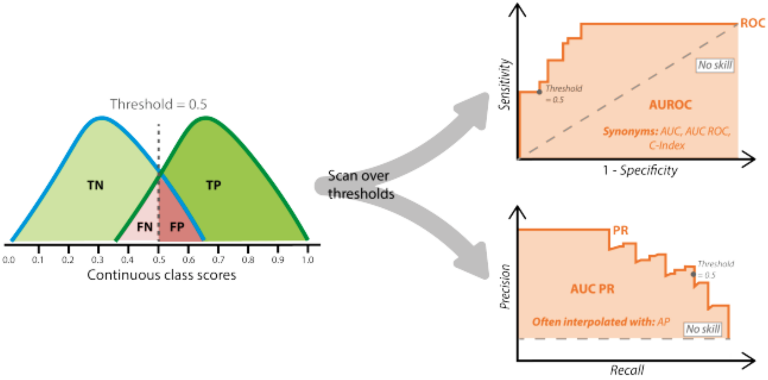

Source: Reinke, Annika, et al. "Common limitations of image processing metrics: A picture story." arXiv preprint arXiv:2104.05642 (2021).

EVALUATION METRICS -- SEGMENTATION

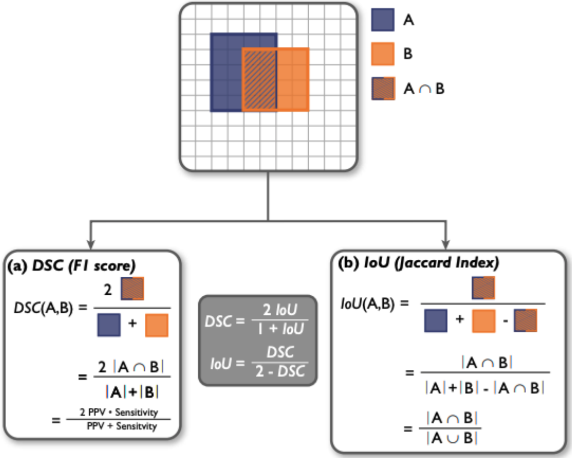

Source: Reinke, Annika, et al. "Common limitations of image processing metrics: A picture story." arXiv preprint arXiv:2104.05642 (2021).

EVALUATION METRICS -- OTHER METRICS

- **Classification**
  - Accuracy (ACC)
  - Error Rate (top 1%, top 5%)
  - Precision
  - Recall
  - F-Score
  - Area Under ROC Curve
  - Area Under PR Curve
- **Segmentation**
  - Dice Coefficient (DICE)
  - Jaccard index
- **Regression**
  - Mean Absolute Error (MAE)
  - Mean Square Error (MSE)
  - Normalized Cross Correlation (NCC)
- **Synthesis/Denoising**
  - Mean Square Error (MSE)
  - Peake Signal to Noise Ratio (PSNR)
  - Structural Similarity Image Measure (SSIM)
  - Contrast to Noise Ratio (CNR)
- **Clustering**
  - Davies-Bouldin index
  - Purity
  - Normalized Mutual Information (NMI)

DEEP LEARNING

DEEP LEARNING

WHAT'S NEXT?

NEXT COURSE?

Machine Learning II

- Neural Networks for Sequences (Ch15)
- Kernel Methods – Support Vector Machine (Ch 17)
- Trees, Forests, Bagging, and Boosting – Boosting (Ch 18)
- Beyond Supervised Learning – Learning with Fewer Labeled Examples (Ch 19)
- Beyond Supervised Learning – Recommender Systems
- Beyond Supervised Learning – Graph Embeddings

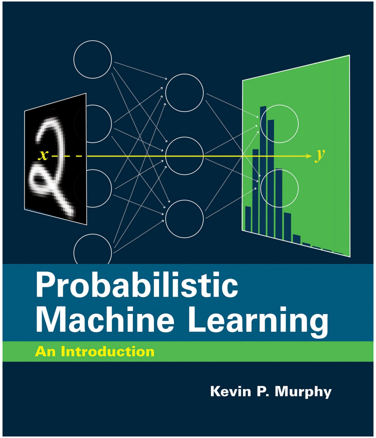

Questions