Other Neural Networks

How tasks affect architectures?

Unidimensional data

Artificial Neural Networks (ANNs)

- Tabular data

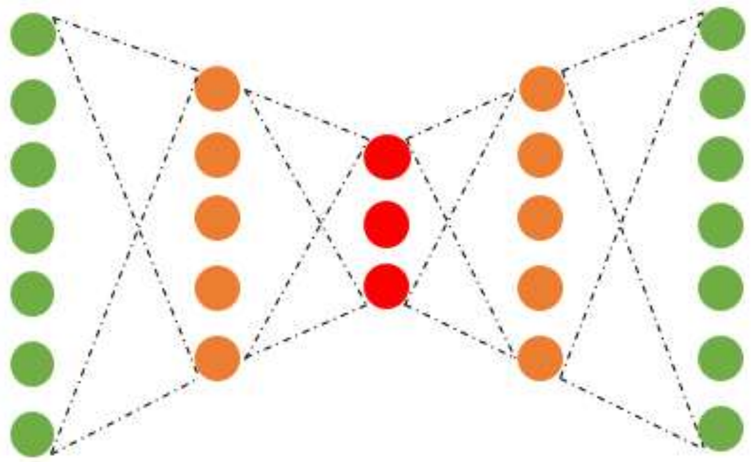

Autoencoder

- Reconstructs the input

Sequence-based Architecture

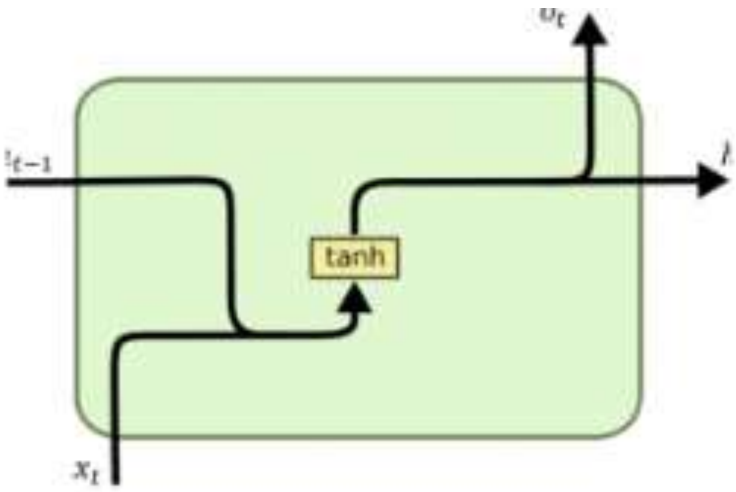

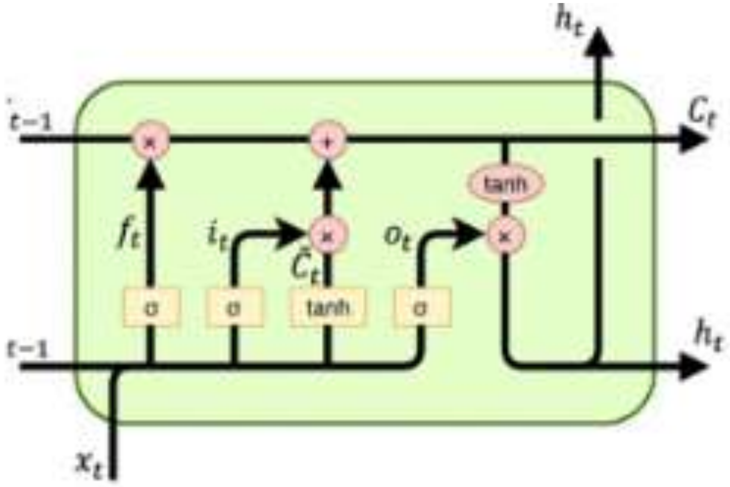

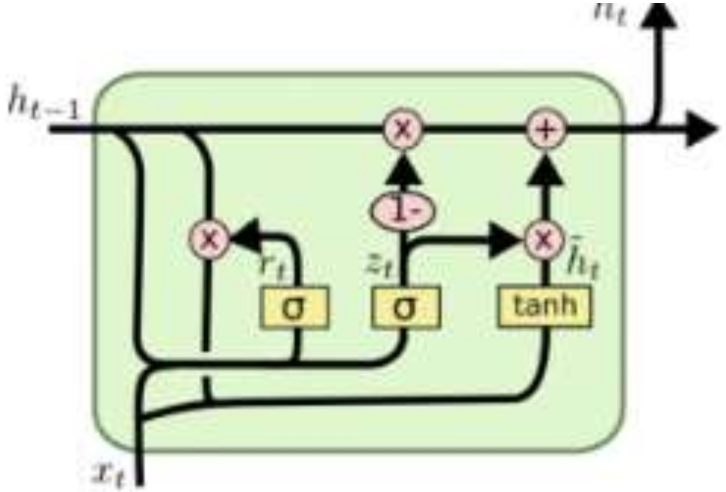

Recurrent Neural Network (RNN)

- Hidden state as memory
- Shallow architecture
- Difficult to model long-term relationships

Long Short-Term Memory (LSTM)

- Three gates: a) input, b) forget, c)output
- Mitigates vanishing gradients

Gated Recurrent Unit (GRU)

- Two gates: a) reset, b) update
- Efficient & simplified iteration of LSTM
- Comparable performance with better speed

Attention-based Architecture

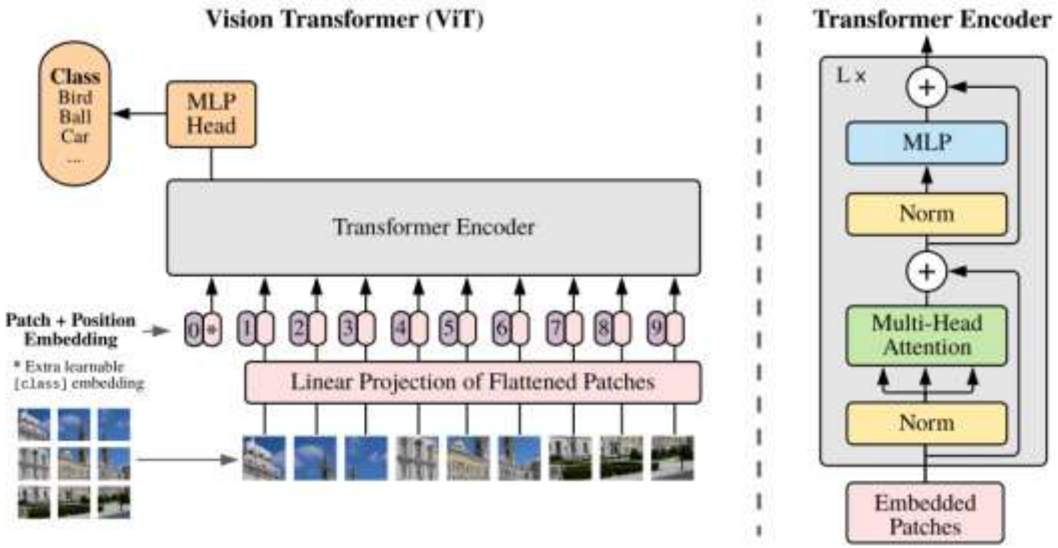

GPT, Bert, T5! ViT:

1. Extract patches
2. Flatten layer
3. Linear Projection layer
4. Add class token (CLS)
5. Add positional encoding
6. Transformer Encoder
7. Extract CLS
8. Classification Head
9. Softmax

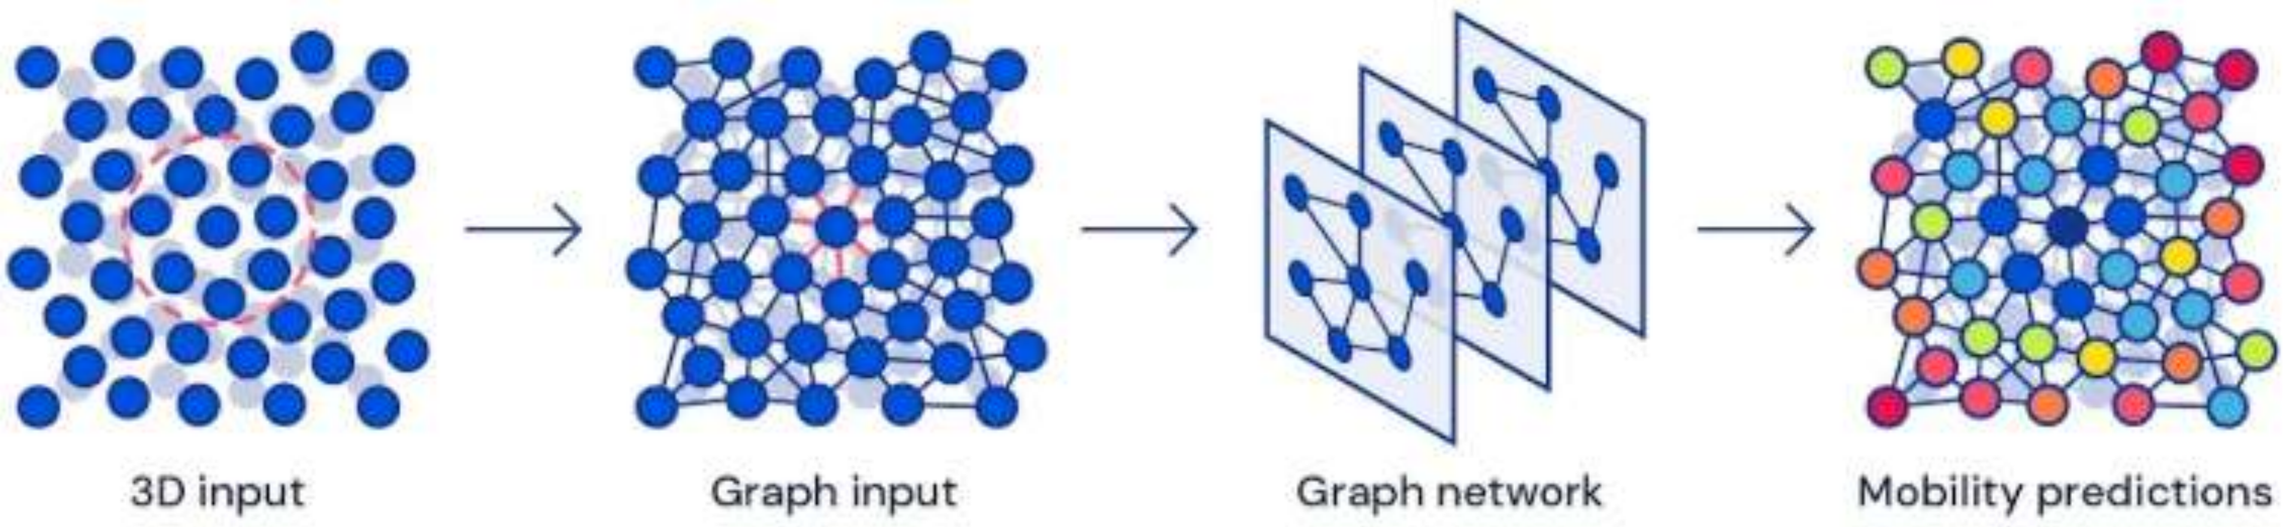

Graph-based Architectures

Architectures

- Graph Neural Networks (GNNs)
- Graph Convolutional Network (GCN)
- Graph Attention Network (GAT)

Fundamental Mechanism: **message passing**

- **Aggregate**, each node collects **feature information** from neighbors
- **Update**, combines its own **current** information with the **aggregated** neighbor information to create **new** information
- **Iterate**, a node eventually "learns" about parts of the graph (**local information**)

Generative Architectures

- Generative Adversarial Models
- Diffusion Models
- Variational Autoencoders (VAEs)

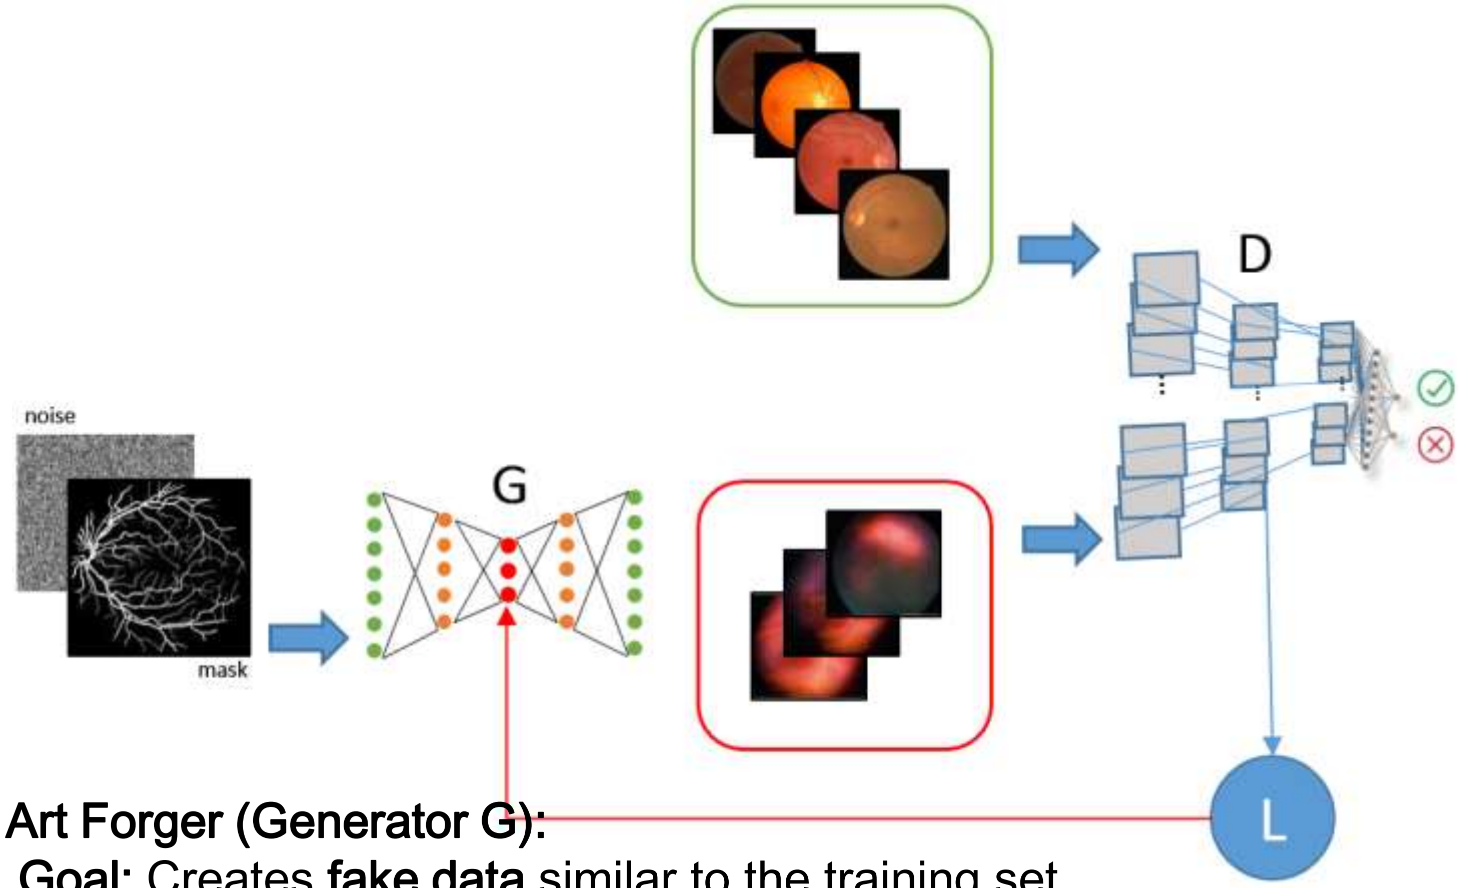

1. **The Art Forger (Generator G):**
   - **Goal:** Creates **fake data** similar to the training set
   - **Process:** Initial state – random noise or specific (multi-class) masks, Iterations – converge into realistic images
2. **The Art Critic (Discriminator D):**
   - **Goal:** Distinguishes between "real" data and "fake" data
   - **Process:** Outputs a probability of whether it is **forged**

Other Architectures

- Restricted Boltzmann Machines (RBMs)
  - Unsupervised feature learning
  - Dimensionality reduction
- Spiking Neural Networks
  - Low-power, real-time sensing
- Physics-Informed Neural Networks
  - Scientific computing
- Capsule Networks
  - Linke CNNs
- Neural Fields (NeRF)
  - Graphics / simulation
- Neural Turing Machines
  - Algorithmic reasoning
- Mixture of Experts (MoE)
  - Large-scale models like LLMs and multimodal systems
- Agentic AI
  - Autonomous task execution
  - Planning, tool use, multi-step reasoning

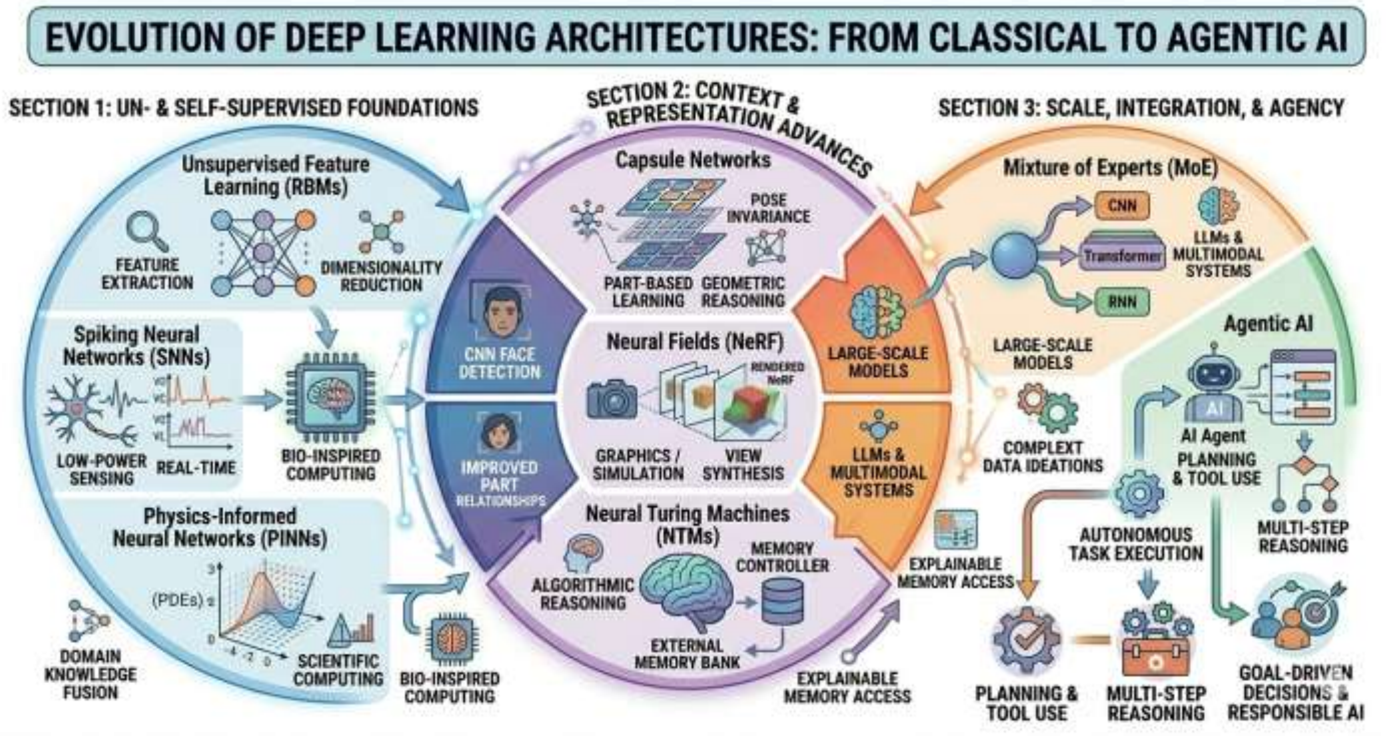

Comparison of DL Architectures

Architecture	Primary Use	Core Strength
RBM	Unsupervised learning	Probabilistic latent structure
SNN	Edge / real-time systems	Energy-efficient temporal processing
PINN	Scientific computing	Embeds physical laws
Capsule Net	Structured vision	Part–whole relationships
NeRF	3D reconstruction	Continuous spatial representation
NTM	Algorithmic tasks	External memory
MoE	Scalable AI systems	Sparse computation
Agentic AI	Autonomous systems	Goal-directed behavior

The AI Project Lifecycle

Hyper-Parameter Optimization & Model Evaluation

Pre-Processing Pipeline

Data preprocessing

- Import DICOM to Numpy
- Denoising
- Motion correction
- Normalization
- Harmonization
- Segmentation

Data Augmentation

- Deformation
- Flipping
- Rotation
- Mirroring

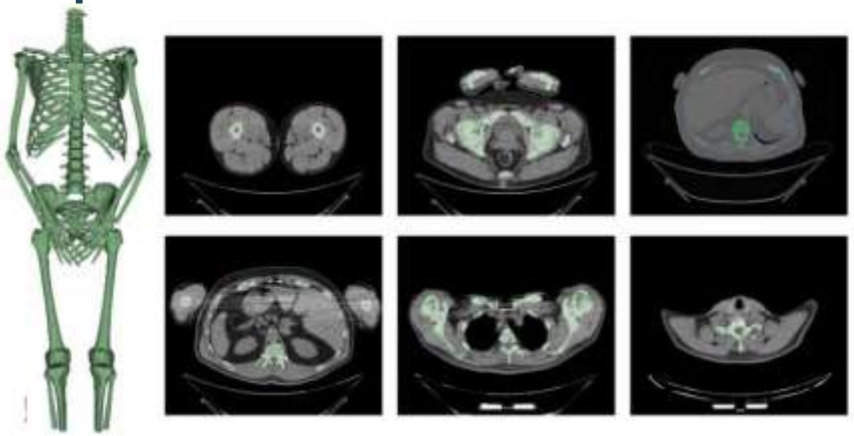

Segmentation

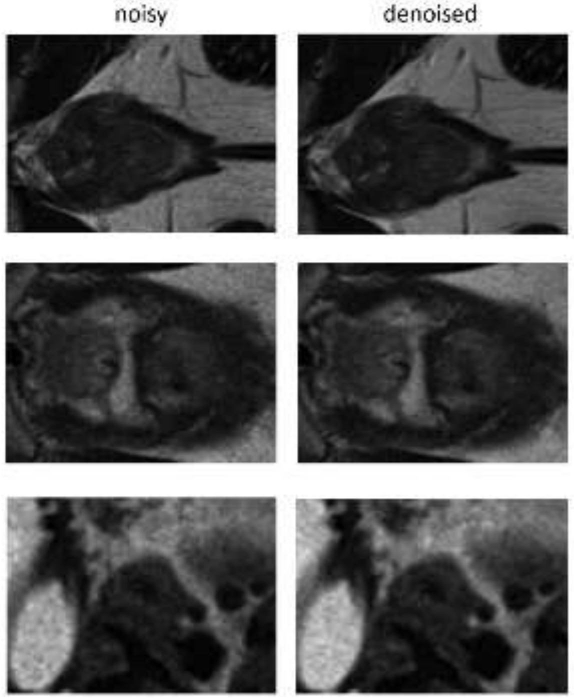

Denoising

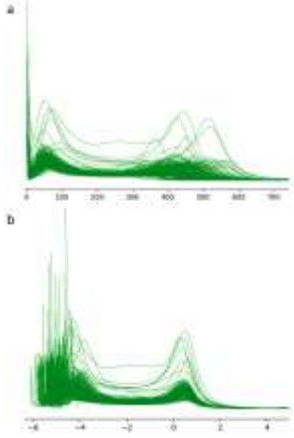

Harmonization

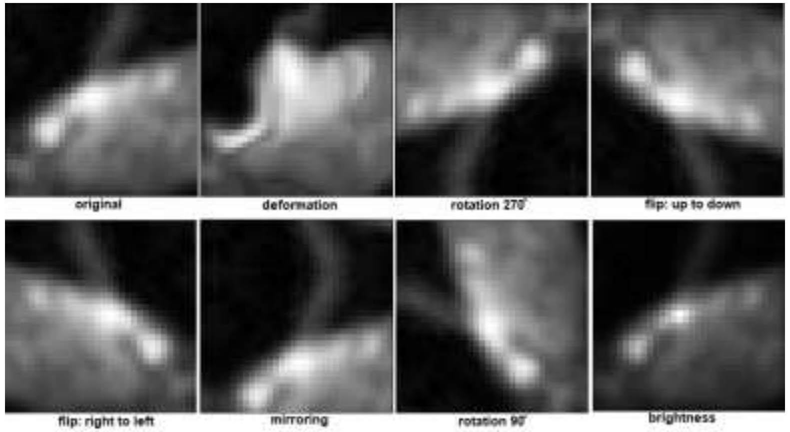

Data Augmentation

Dataset Stratification

Best Practices [1]

1. Cross-validation
2. Independent sets (training, validation, testing)
3. External validation sets

Stages

1. Model Convergence
   - Weight update
   - Gradient Descent
   - Backpropagation
2. Model Optimization
   - Early stopping
   - Assess overfitting
3. Model Evaluation
   - Independent patient cohort

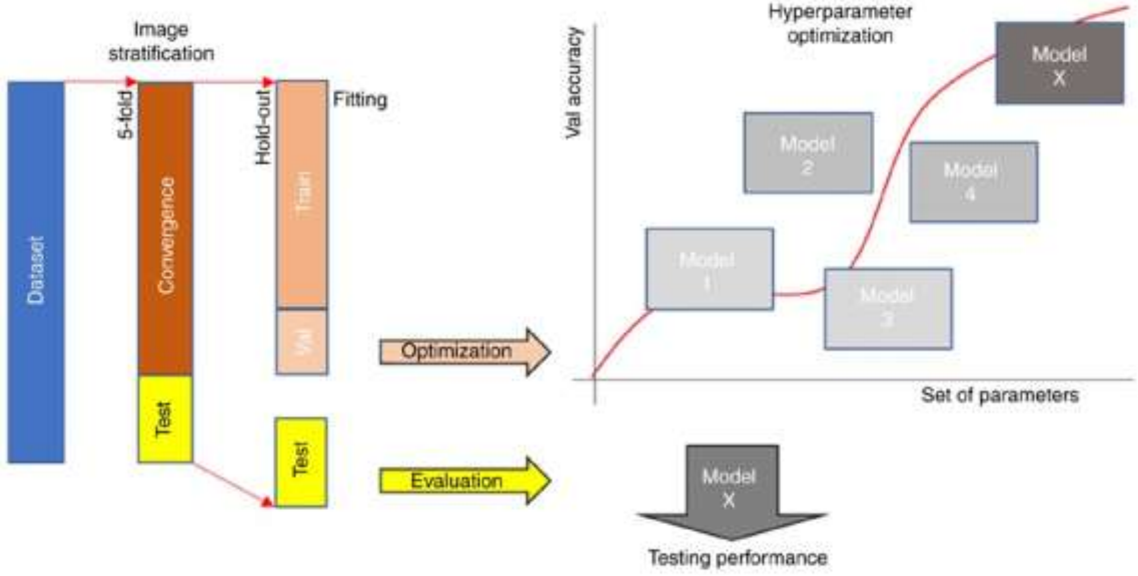

[1] Wujek, B., Hall, P. & Güneș, F. Best Practices for Machine Learning Applications. 1–23 (SAS Inst. Inc. 2016).

Architecture Type

Define Architecture

- Data driven
- Patch-wise 2D
- End-to-End 3D
- Pre-Trained Feature Extraction

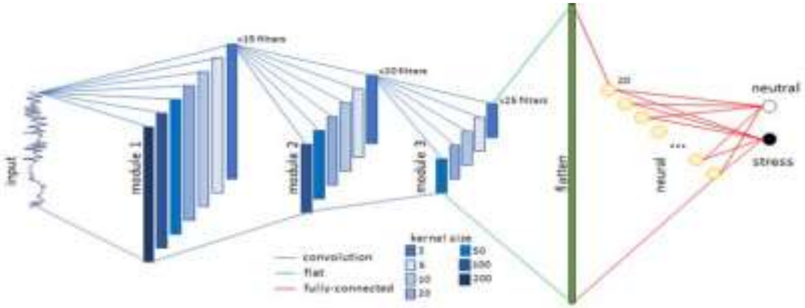

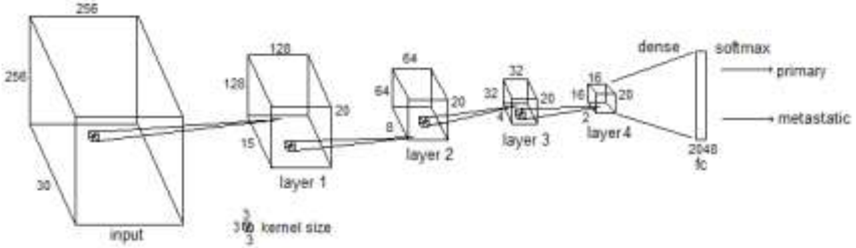

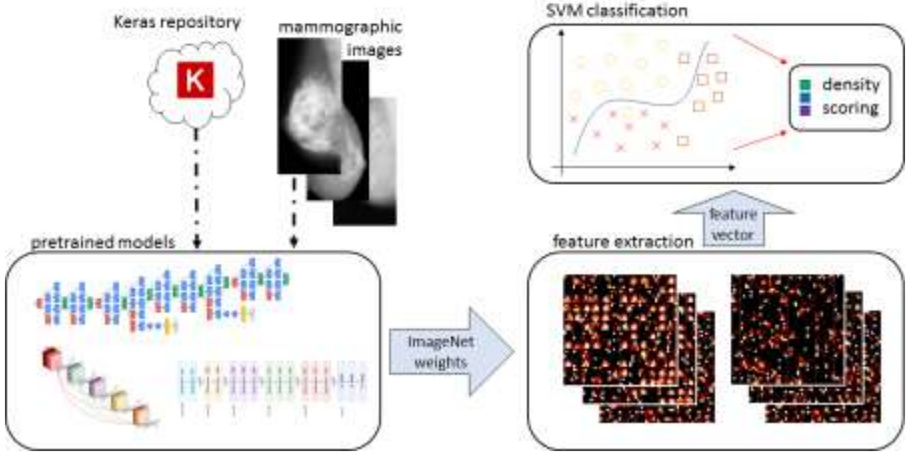

[1] Giannakakis G., Trivizakis E., Tsiknakis M, and Marias K., 2019. A novel multi-kernel 1D convolutional neural network for stress recognition from ECG. Eighth International Conference on Affective Computing and Intelligence Interaction (ACII), In press: IEEE Xplore
[2] Trivizakis, E., Ioannidis, G.S., Melissianos, V.D., Papadakis, G.Z., Tsatsakis, A., Spandidos, D.A. and Marias, K., 2019. A novel deep learning architecture outperforming 'off-the-shelf' transfer learning and feature-based methods in the automated assessment of mammographic breast density. Oncology reports, 42(5), pp.2009-2015.

Transfer Learning

1. easy deployment
2. p re-trained models or foundation models
3. use in other domains besides natural images
4. “off-the-shelf” for limited data
5. Fine-tune for best results
6. Repurpose for other imaging tasks like segmentation
7. open databases can be used for tuning or analysis
8. Ideal for medical applications

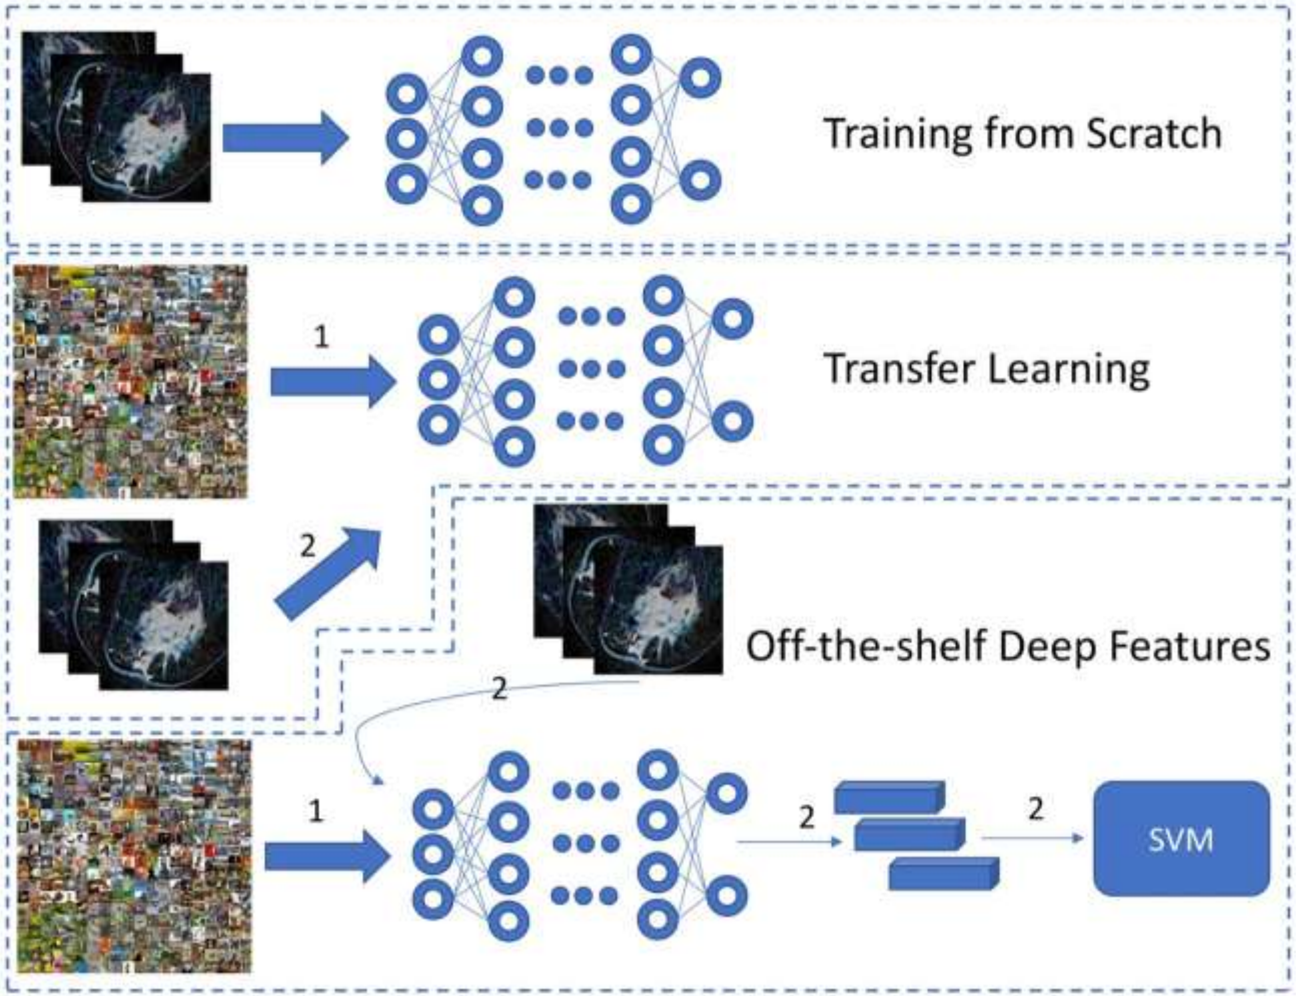

Zhu, Zhe, et al. "Deep Learning for identifying radiogenomic associations in breast cancer." Computers in biology and medicine 109 (2019): 85-90.

Hyperparameter Optimization

- Identify hypes
  - Learning rate
  - Number of filters
  - Number of neurons
  - Activation functions
- Define optimization method
  - Grid search
  - Random search
- Document performance changes
  - Exclusively on validation metrics
  - Change hypes accordingly
- Remember!
  - Testing set MUST remain **unseen**

Model Evaluation

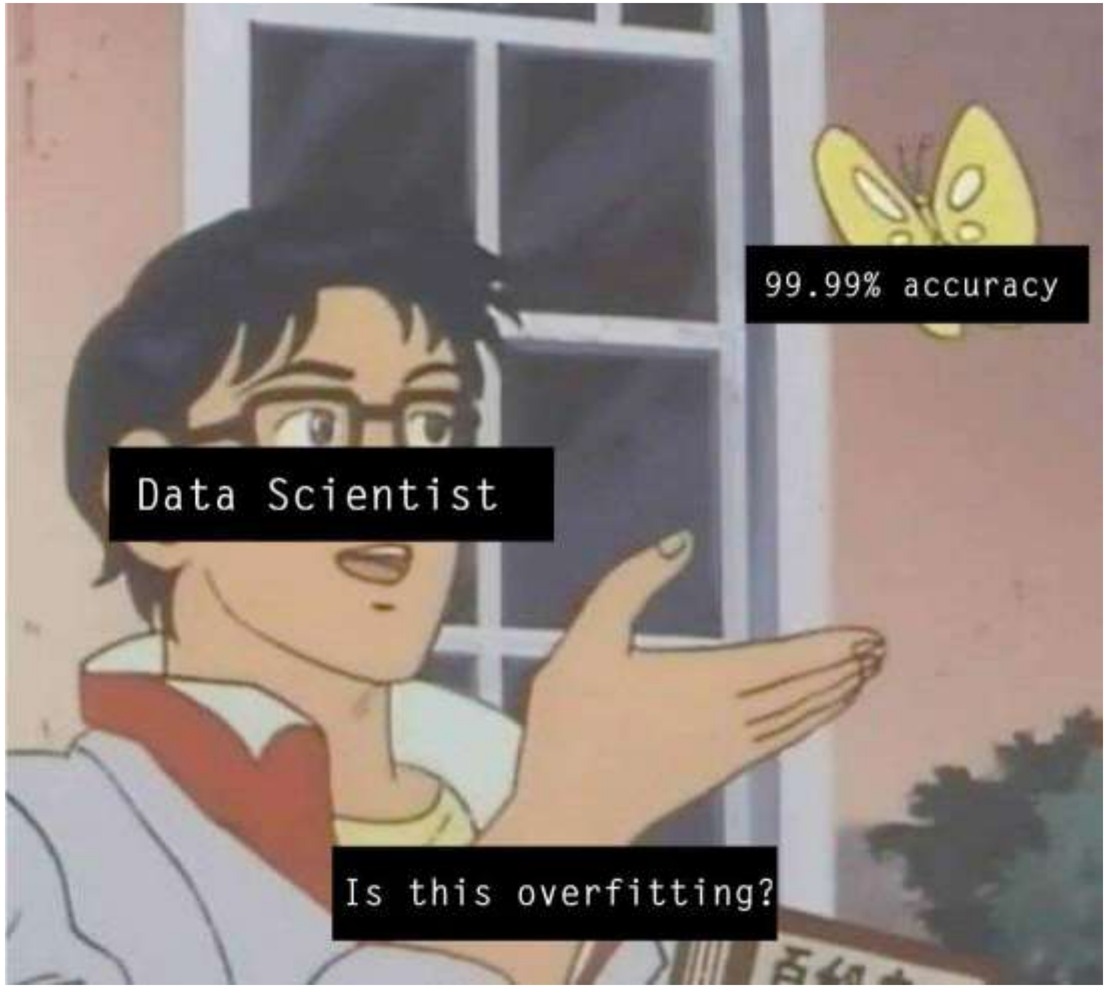

- **During model fitting**
  - Learning Curves
  - Training vs Validation Metrics
  - The gap should be small and, preferably, with a shrinking trend
  - Early-stopping prior to divergence
- **After convergence**
  - External validation set
  - Performance on par with validation error
  - Should be **unseen!**

Learning Curves – Patterns

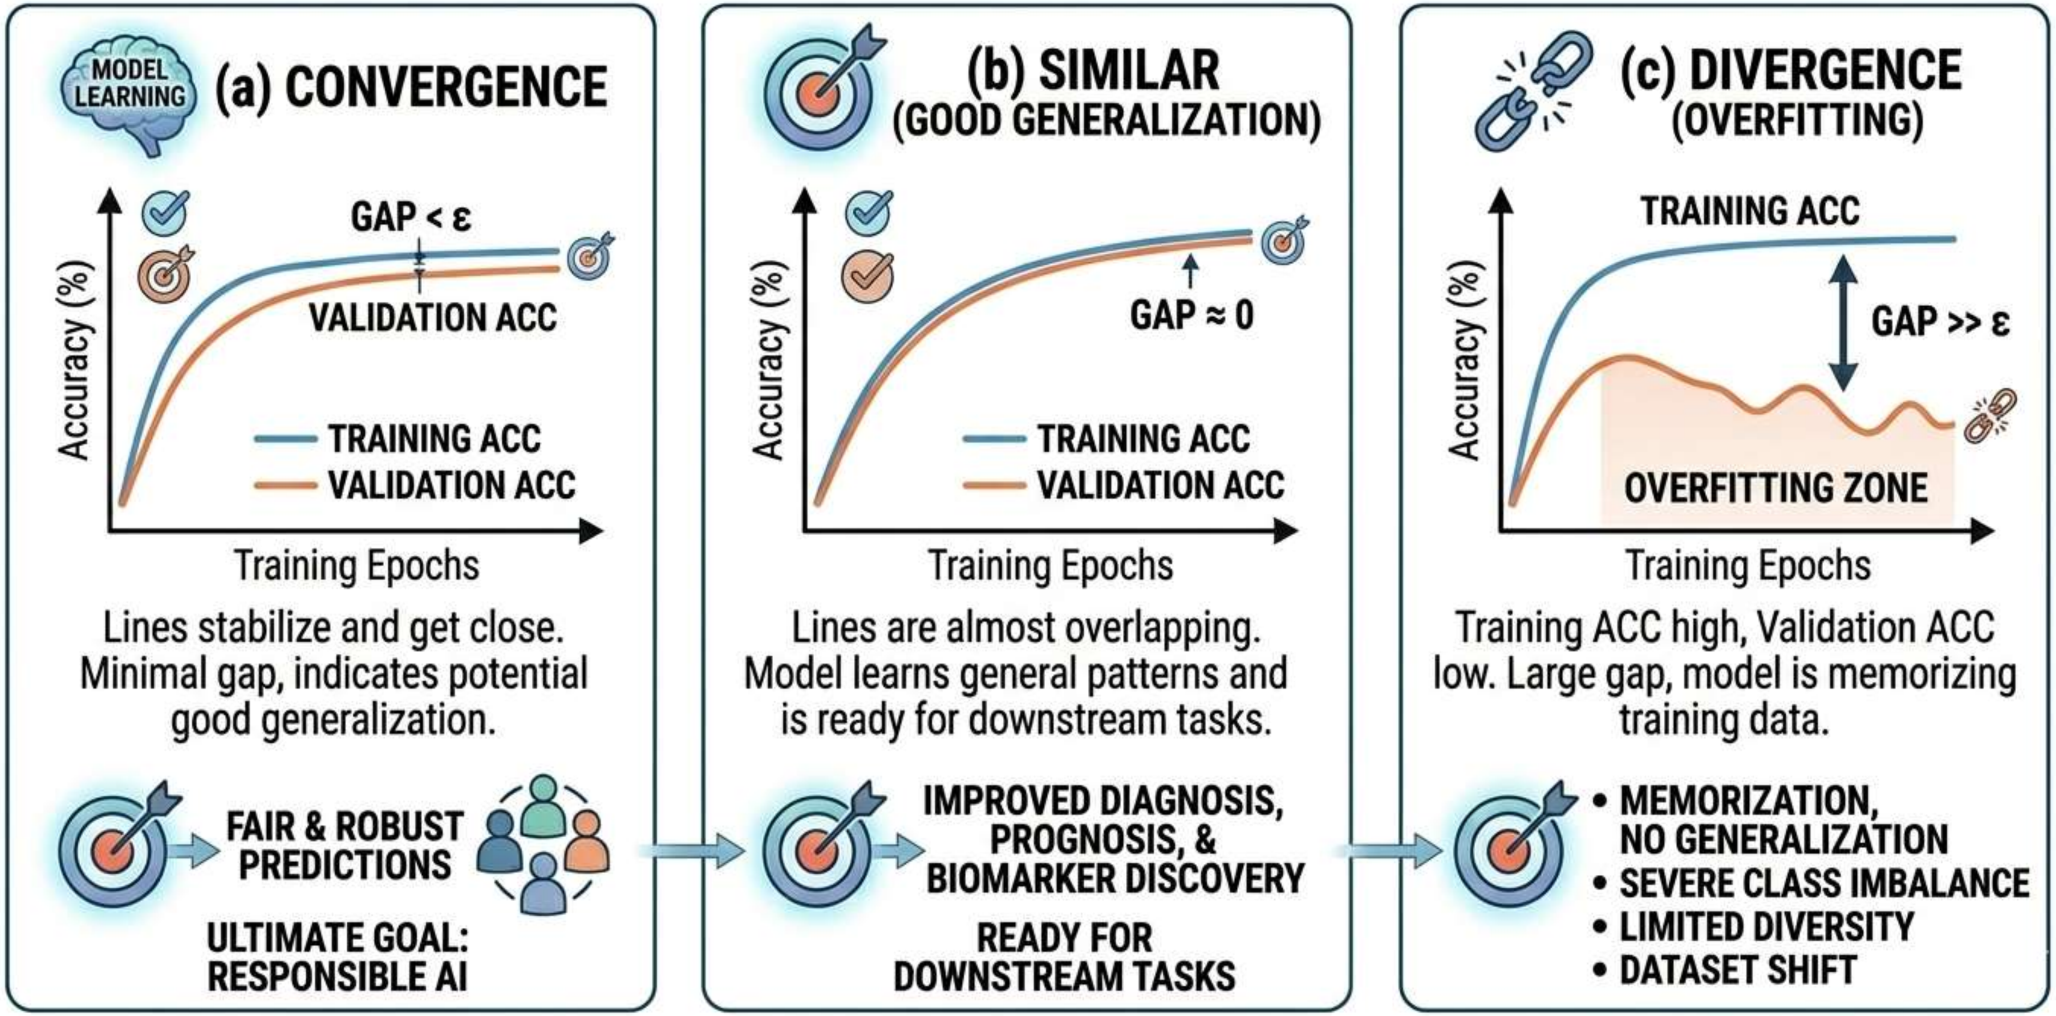

Learning Curves – Examples

Loss – Structural Differences Index

Denoising Model – Learning Curves
- **Converges** until epoch ~95
- Diverges after epoch >100 – **val_loss > tr_loss**

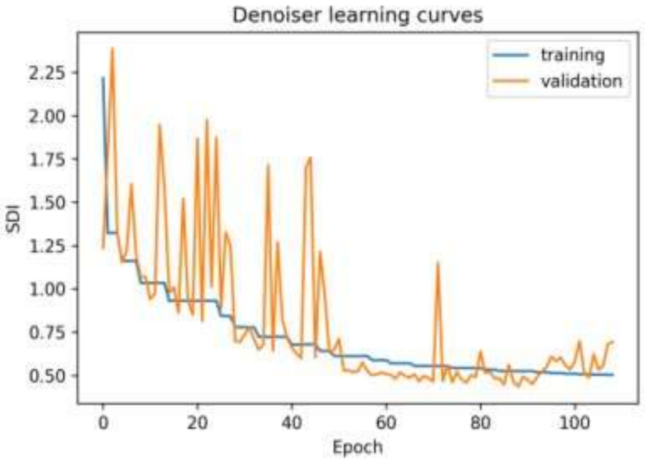

The learning curves of "real-world" models can be messy and noisy!

Training vs Validation Metrics

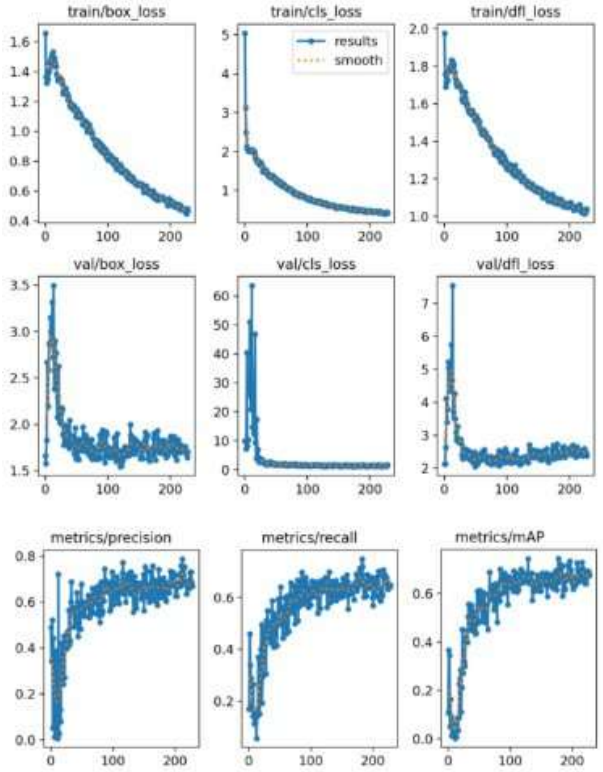

Detection Model – Learning Curves
- **3 losses:** bounding box , classification, distributed focal loss
- **Converges** until epoch ~170
- Diverges after epoch >180 – **val_loss > tr_loss**
- **Early-stopping** activated when mAP maxed!

Software Stack – Complex Dependencies

OS	Ubuntu 18.04 LTS
Drivers	Nvidia GPU 388.90
Cuda Libraries	Cuda Tools, CuDNN
Deep Learning Libraries	Tensorflow. Keras, Pytorch
Python Libraries for Image Processing	Numpy, Pydicom, OpenCV, skimage, simpleITK
Applications	Tissue Classification

prone to dependency errors

Learning Paradigms

HowTo learn?

Model Convergence Methods

- Supervised Learning
  - Learn mapping: x→y
  - Requires labeled data
- Unsupervised Learning
  - Unstructured data
- Self-Supervised Learning
  - e.g., masked prediction, contrastive learning
- Reinforcement Learning (RL)
  - Learn via interactions
    - state → action → reward
- Generative Learning
  - Reconstruct data from noise or lower dimensionality data
- Federated Learning
  - Model convergence from remote nodes
- Transfer Learning
  - Adaptation from another domain

Weakly Supervised: Breast Cancer Classification

- MRI
  - Breast
    - Malignancy prediction
    - BI-RADS based
  - Image analysis technique for ROI extraction (top)
  - 3D MRI input
  - Weakly annotated with Breast ROI only
- 3D Densenet architecture with global average pooling (GAP)
- Deep model's attention maps (b)
- Performance
  - ACC .837 [79.1%, 87.4%]
  - SN .908 [86.0%, 94.1%]
  - AUC .859

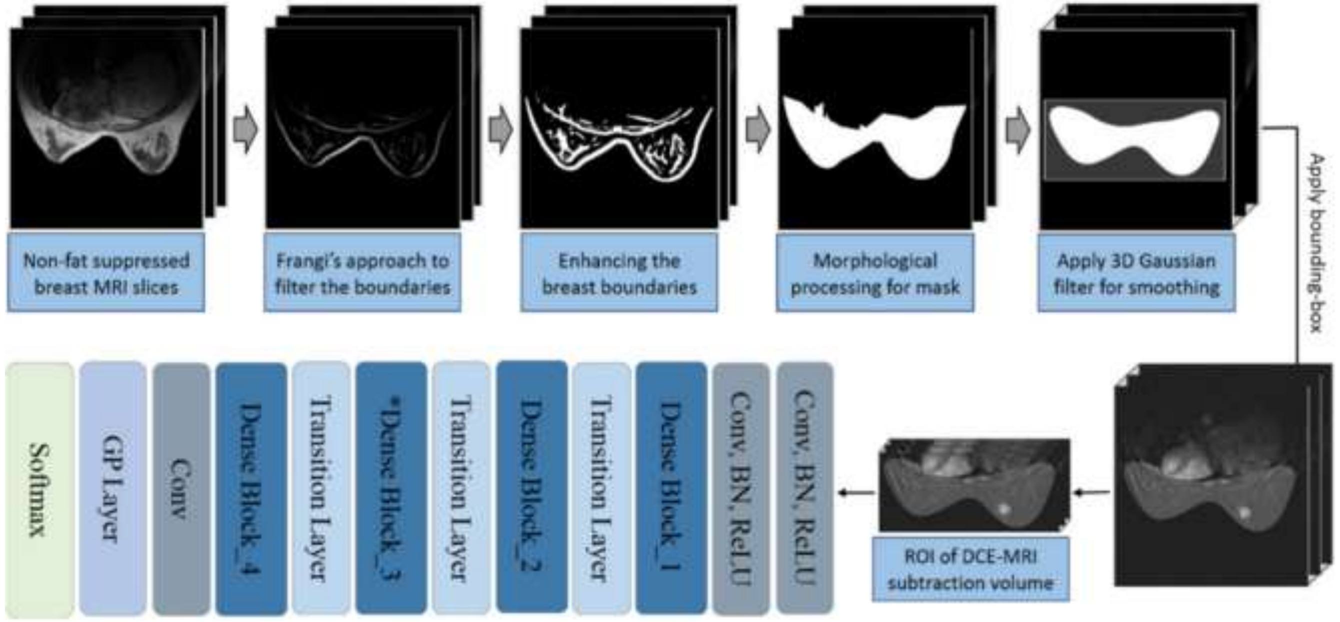

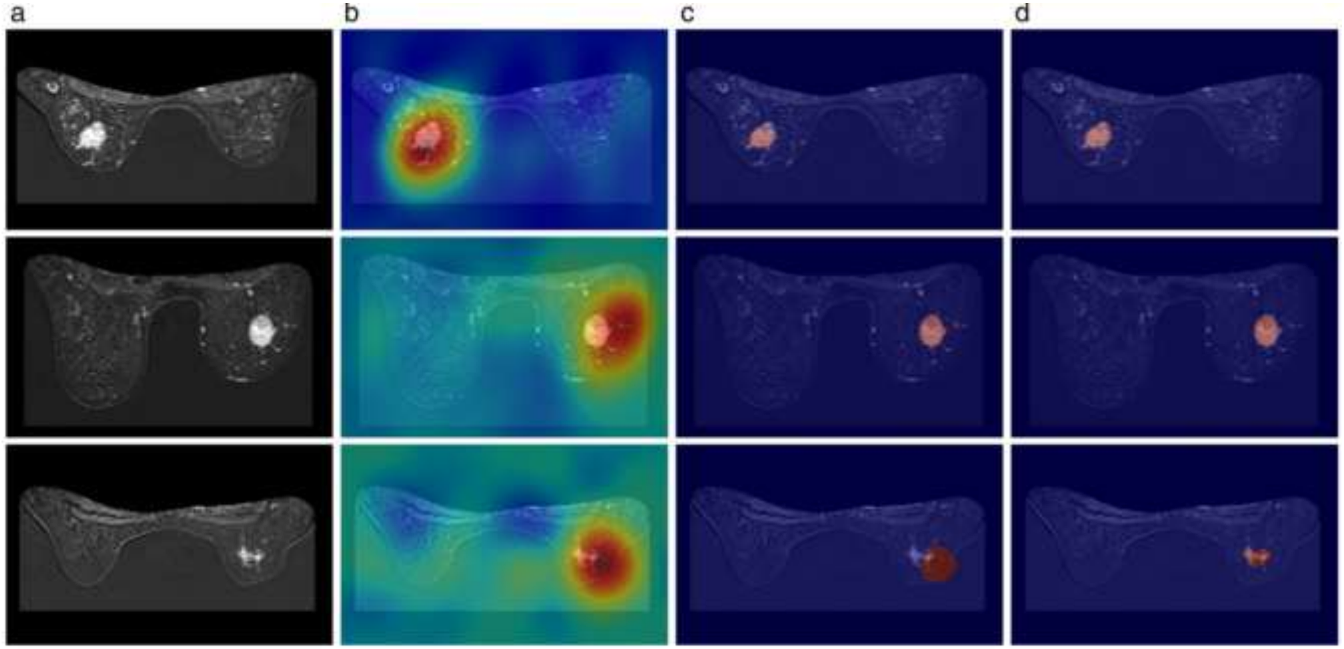

Zhou, Juan, et al. "Weakly supervised 3D deep learning for breast cancer classification and localization of the lesions in MR images." **Journal of Magnetic Resonance Imaging** 50.4 (2019): 1144-1151.

Transfer Learning: Microscopy Imaging

- Nucleus-based classification
- Custom preprocessing protocol
- Harmonization among different image magnification levels
- Fully-automated analysis
- Evaluated on unseen testing sets
- AUC 0.962 ± 0.04

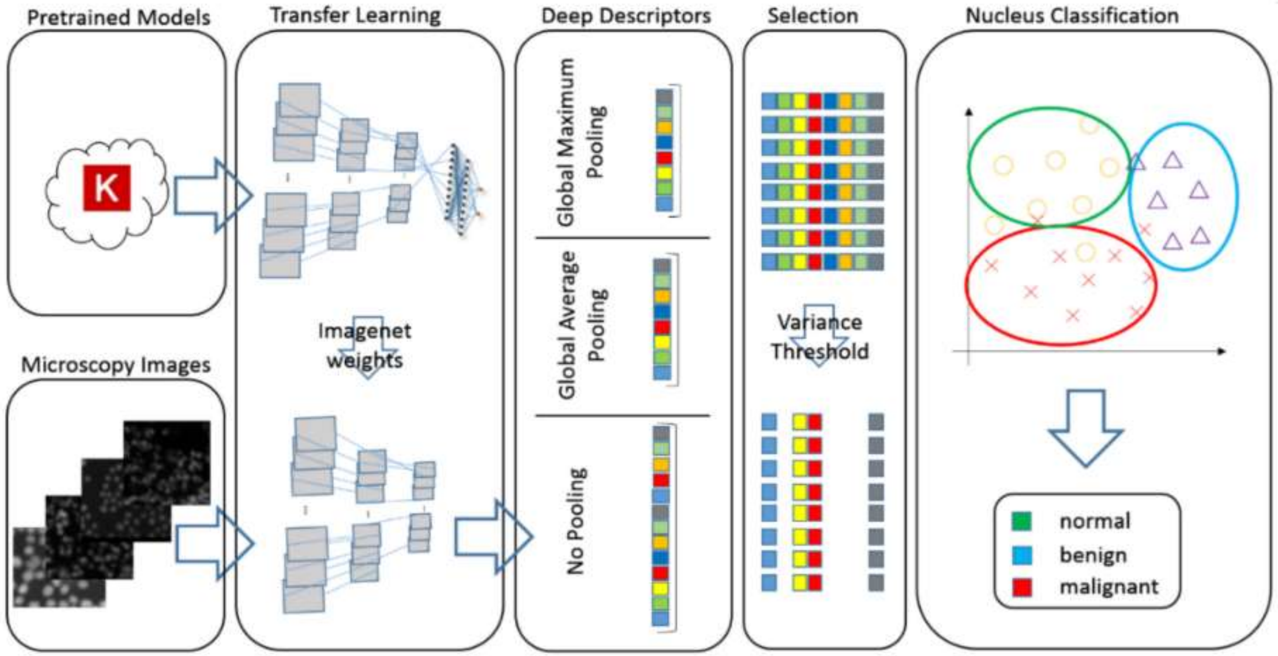

Ioannidis, G. S., Trivizakis, E., Metzakis, I., Papagiannakis, S., Lagoudaki, E., & Marias, K. (2021). Pathomics and Deep Learning Classification of a Heterogeneous Fluorescence Histology Image Dataset. Applied Sciences, 11(9), 3796.

Denoising – Prostate MRI

- Prostate T2WI
- Up-to 18% improvement in terms of SSIM
- Supervised Learning
  - Training Data: Synthetical Noise in HQ* T2WI images
  - Label: HQ Image
- Fully Convolutional
  - No pooling layers
  - Noise estimation

*High Quality

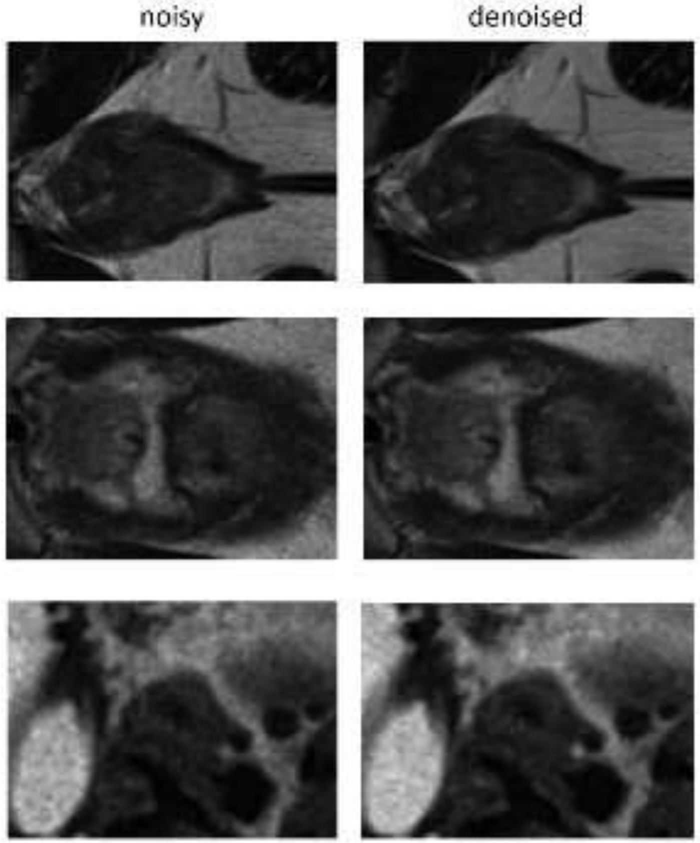

Trivizakis, E., Mylona, E., Nikiforaki, K., & Zaridis, D. I. (2025). Texture Preserving Deep Learning-based Noise Reduction for Anatomical Magnetic Resonance Images and Its Impact on Imaging.

Organ Segmentation – Lung CT

- Custom U-Net architecture
- Lung parenchyma segmentation
- Model trained on CT slices from the LIDC data
- Segmentation was performed on compressed CT found online
  - No Hounsfield Units available
  - Unknown CT windowing methodology
  - Multicentric data
  - Multiple resolutions (148 by 61 to 1,637 by 1,225 pixels)
- Performance Dice Similarity Coefficient of 0.996
- Segmented slices improved COVID-19 CT model
  - AUC by 4.8%

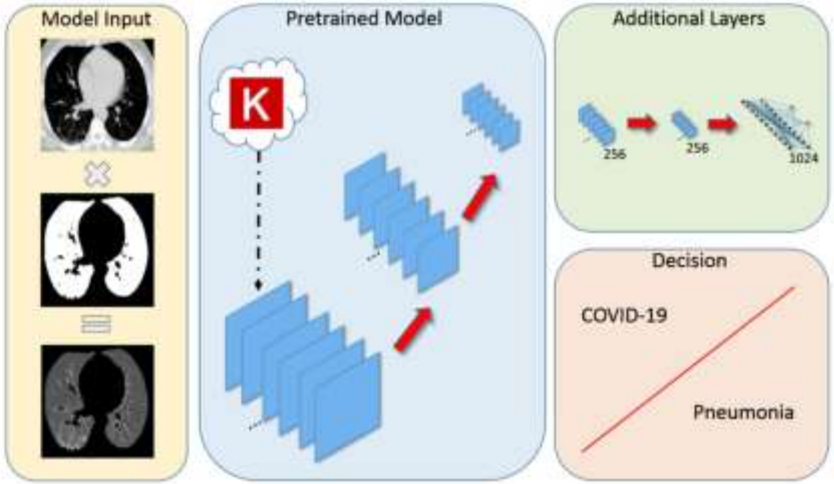

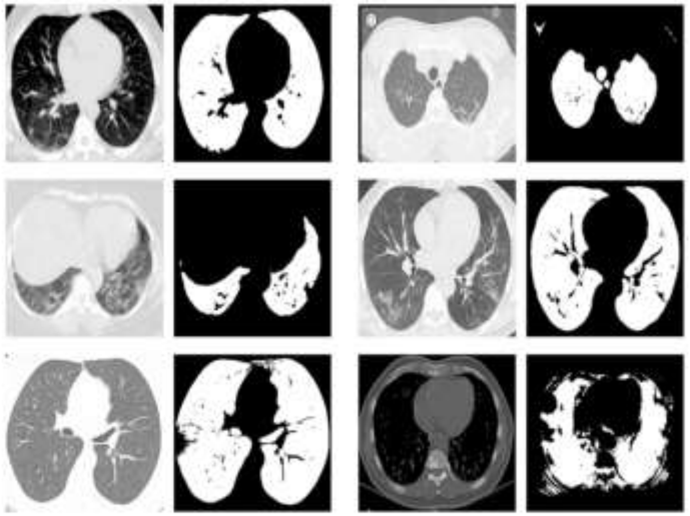

Trivizakis, E., Tsiknakis, N., Vassalou, E. E., Papadakis, G. Z., Spandidos, D. A., Sarigiannis, D., ... & Marias, K. (2020). Advancing COVID-19 differentiation with a robust preprocessing and integration of multi-institutional open-repository computer tomography datasets for deep learning analysis. *Experimental and therapeutic medicine*, *20*(5), 1-1.

Cross-Modality Translation - CT to PET

- Generative Adversarial Networks
  - Multi-channel architecture
  - Generation from binary mask and CT slice

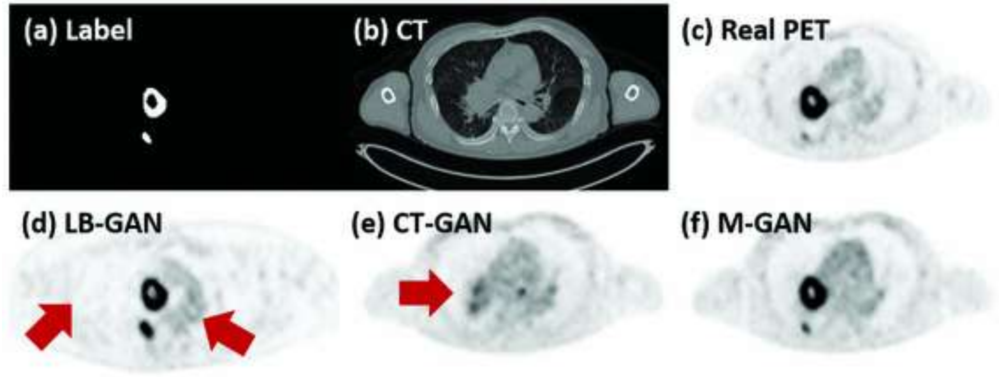

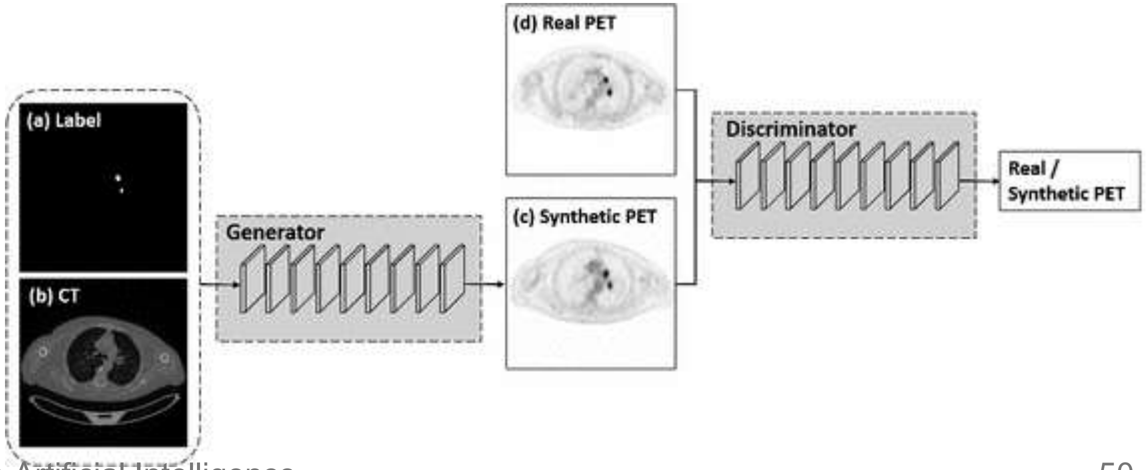

Bi, L., Kim, J., Kumar, A., Feng, D., & Fulham, M. (2017). Synthesis of positron emission tomography (PET) images via multi-channel generative adversarial networks (GANs). In molecular imaging, reconstruction and analysis of moving body organs, and stroke imaging and treatment (pp. 43-51). Springer, Cham.

Ethics & Bias

Challenges of Data-Driven Methods

Bias

- Data Acquisition
  - Sex
  - Ethnicity
  - Location
  - Country
  - Representation
  - Methodology
- Labeling
  - Inter-observer variability
  - Timing
  - Documented variables
  - Descriptions

Bias in GenAI Models

"Generate a **character** that is a cute orange fish with **human-like** facial features"

"Generate a blue fish **character** with big eyes and cute facial features"

"Generate a shark SETTINGS: floating, **studio flat lighting**, soft shadows"

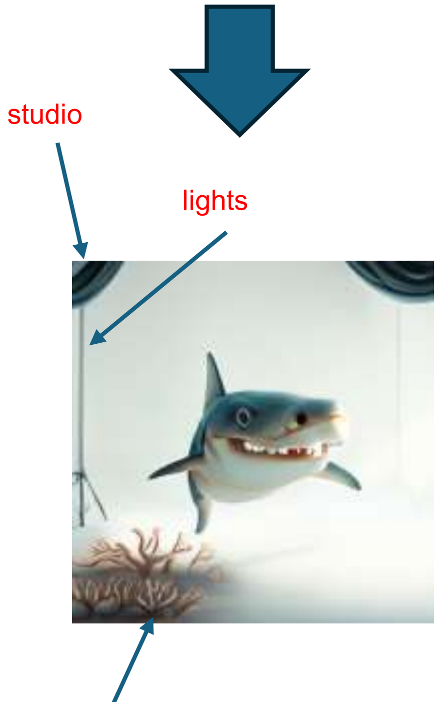

Not in prompt, relevant to a shark scene in a studio

Main AI Challenges

Limitations

- Still a **black box**!!!
  - **XAI** can be susceptible to **bias**
- Data biases (ethnicity, age, income, location, etc)
- Observer biases
- **Legal** and **ethical** frameworks
- **Accountability** in case of importand errors

Challenges

- **Quality** of data
- **Volume** of data
- Imaging **technology**
- **Temporal progression**

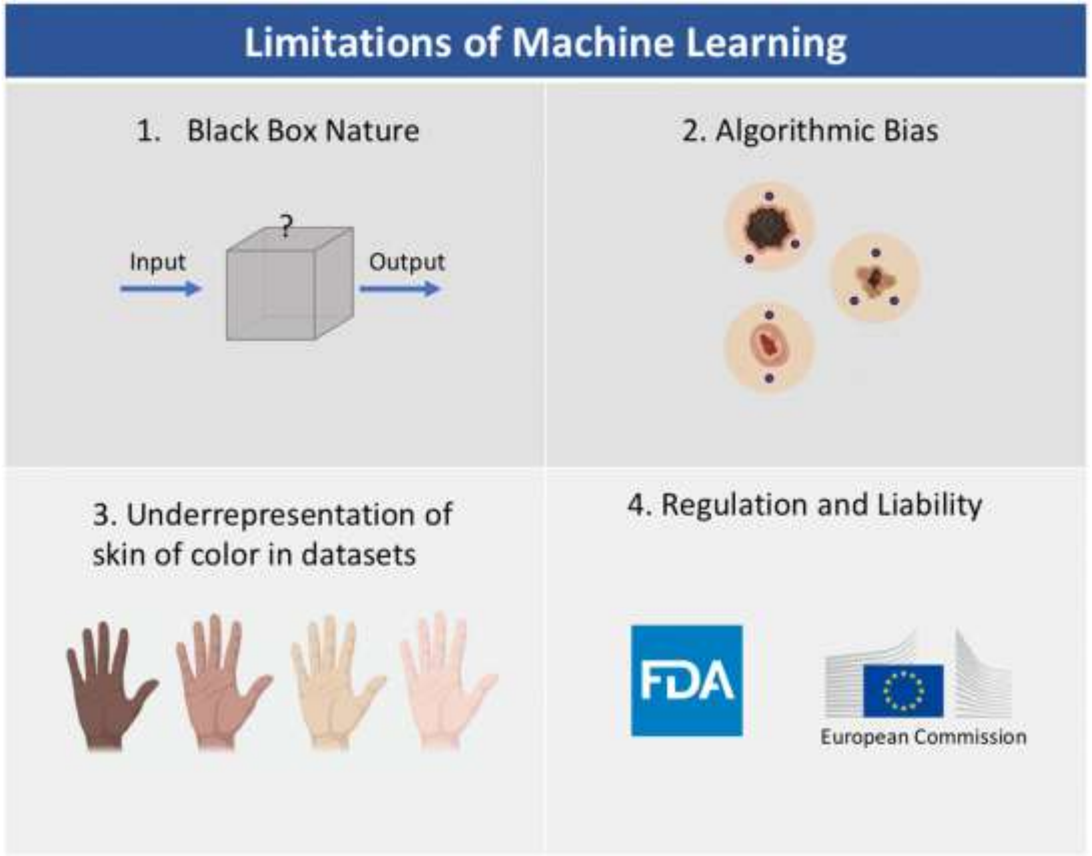

Chan, Stephanie, et al. "Machine learning in dermatology: current applications, opportunities, and limitations." *Dermatology and therapy* 10 (2020): 365-386.

Current Trends

- **Prognosis** & personalized treatment
  - Predict a future variable from current data
- Improved **objectivity** in lesion monitoring
  - Temporal data
- Multi-Modal AI
  - Integrate multi-source data
- Conversational AI
  - Agents and Dynamic Knowledge Bases

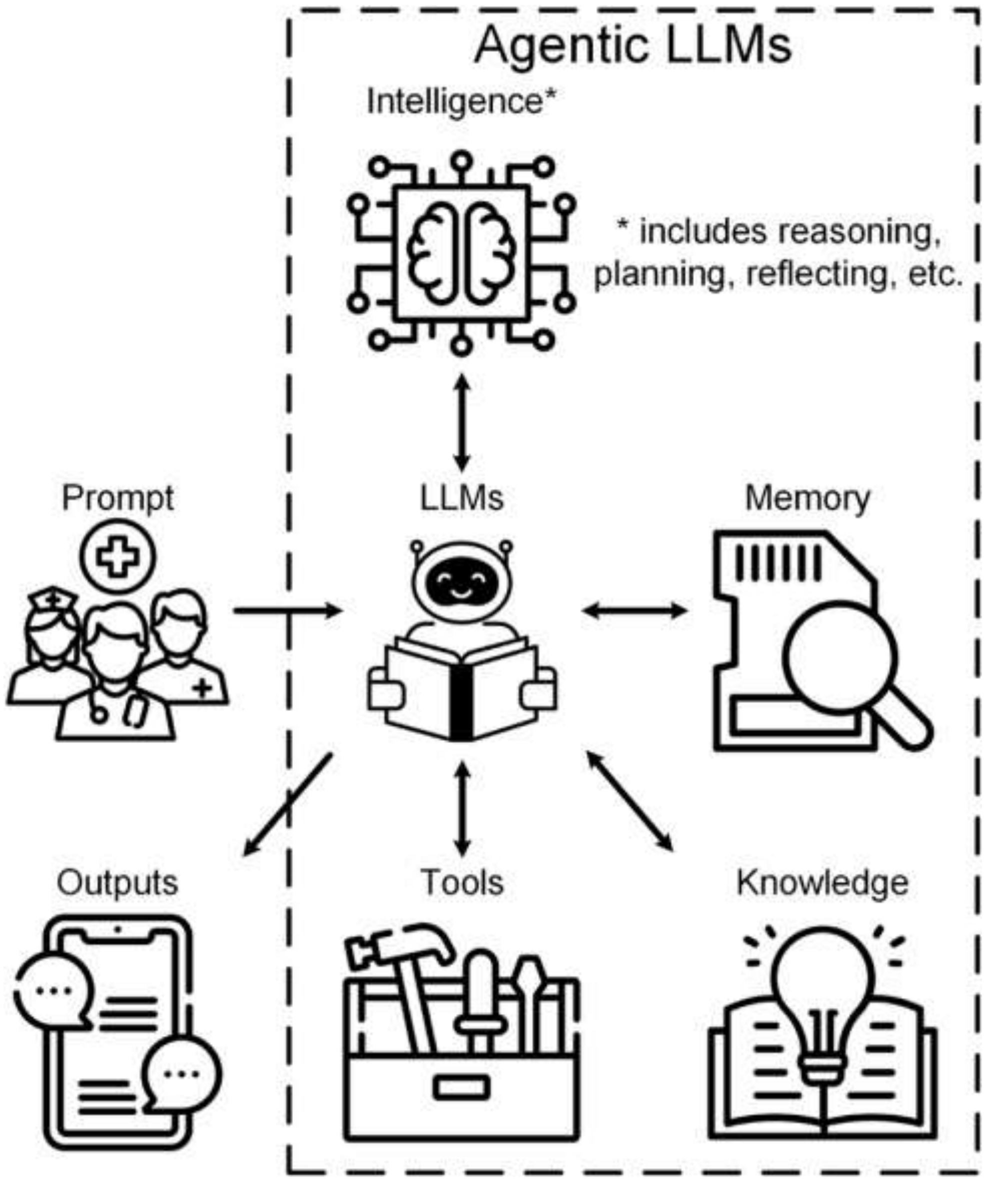

Thanks!

Any question?