Interpretable Multimodal Deep Learning for Objective Diagnosis, Prognosis and Biomarker Discovery

Faisal Mahmood, PhD
Postdoctoral Fellow
Department of Biomedical Engineering
Johns Hopkins University
faisalm@jhu.edu | http://faisal.ai

Deep Learning for Medical Imaging – Major Challenges

1. Limited Annotated Data

- Under representation of rare conditions.
- Limited experts available for annotation.
- Privacy Issues
- Cost

Transfer Learning

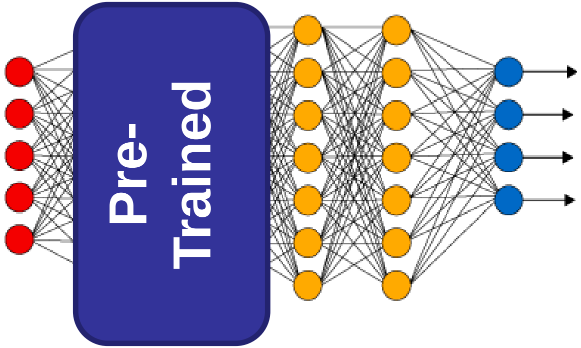

Unlabeled

Unstructured

Does not capture clinical diversity.

Johns Hopkins
Electronic Medical Records

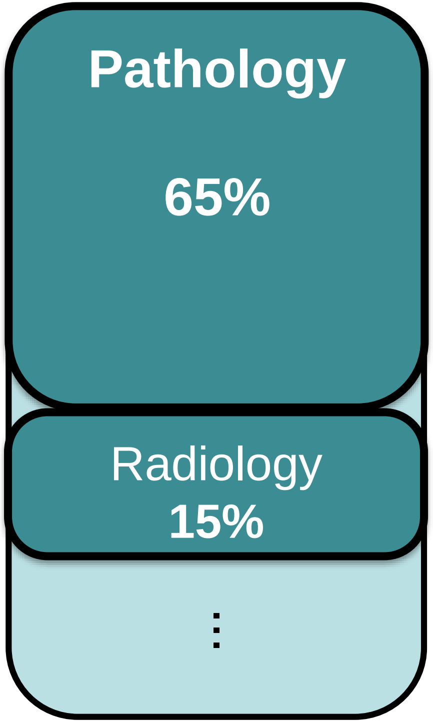

Faisal Mahmood, Nicholas J. Durr et al."Unsupervised Reverse Domain Adaptation for Synthetic Medical Images via Adversarial Training." ***IEEE Transactions on Medical Imaging*** (2018).

Deep Learning for Medical Imaging – Major Challenges

2. **Domain Adaptation**

- Diversity in data, different sensors, cites and patients.
- Patient specific texture and color information.

Train **Test**

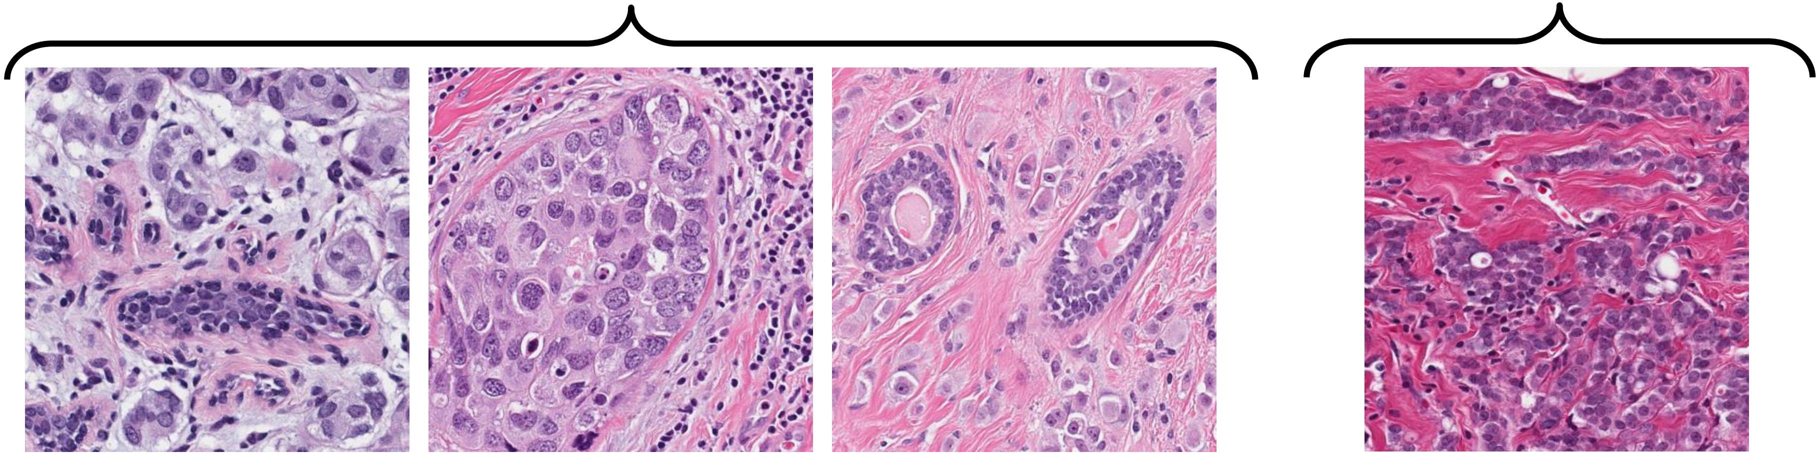

How can we train AI systems robust to variability in the data?

Deep Learning for Medical Imaging – Major Challenges

2. **Domain Adaptation**
- Diversity in data, different sensors, cites and patients.
- Patient specific texture and color information.

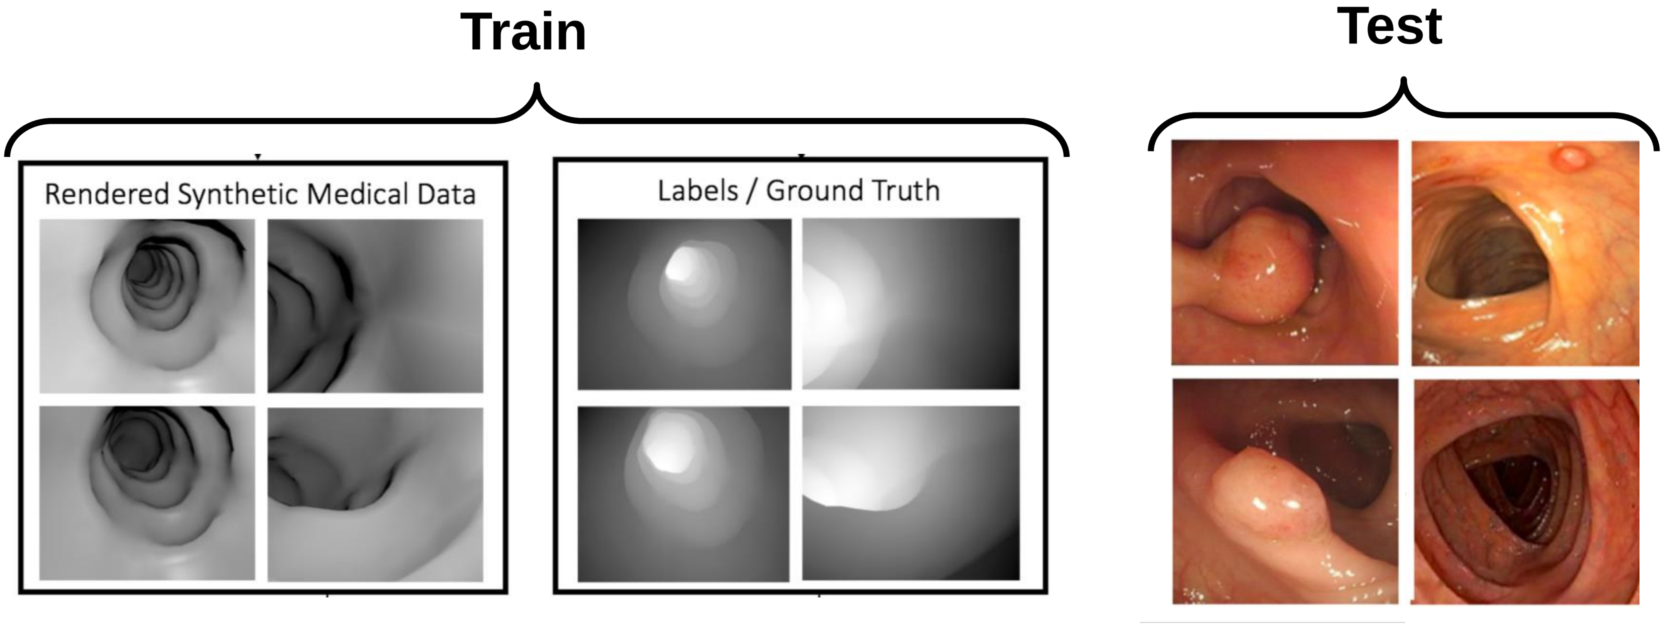

(F. Mahmood et al., 2018)

Deep Learning for Medical Imaging – Major Challenges

3. Structured Prediction

- Global vs Local features.

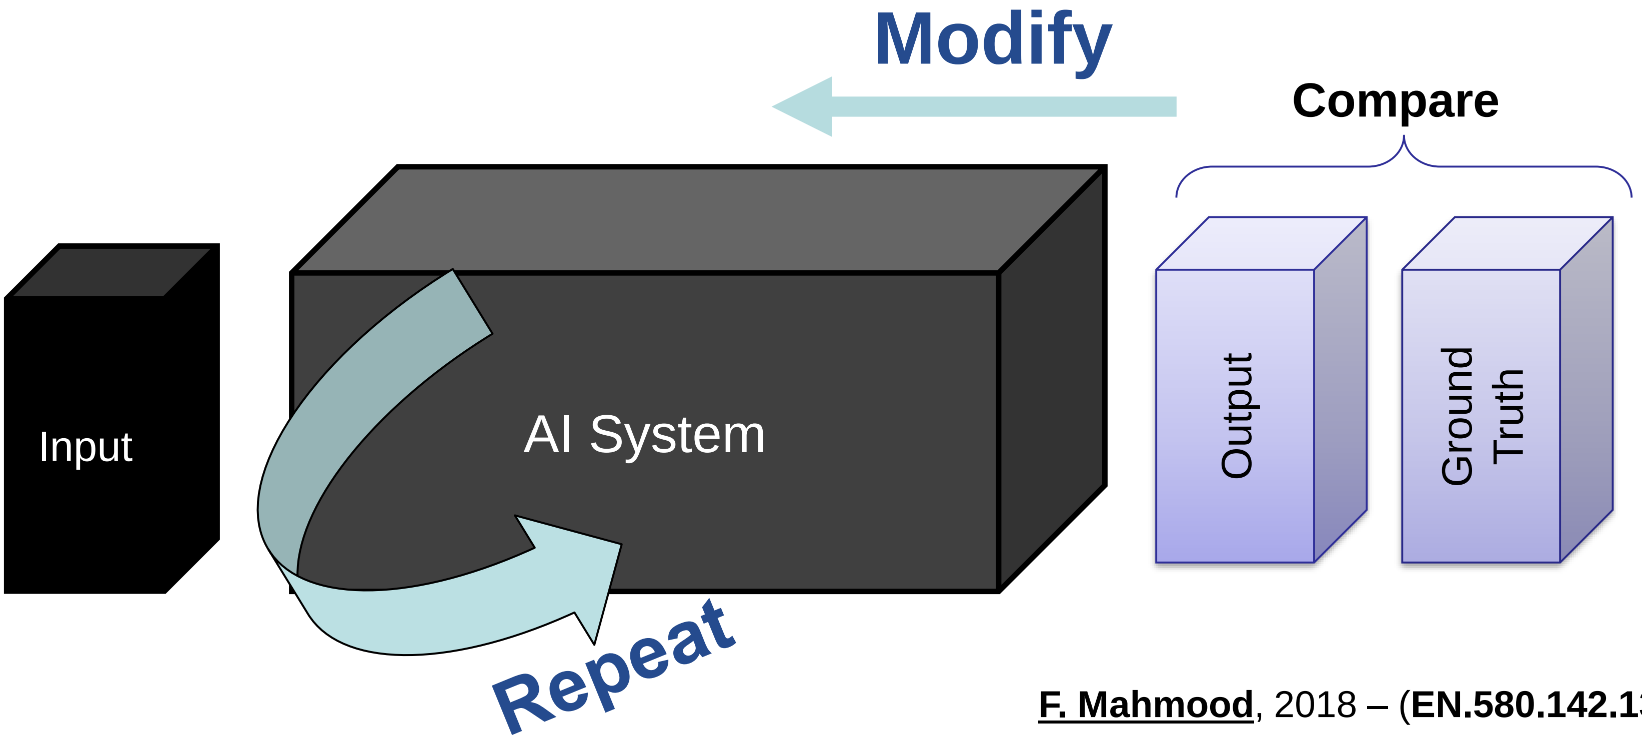

<u>F. Mahmood</u>, 2018 – (**EN.580.142.13**)

3. Structured Prediction

- Global vs Local features.

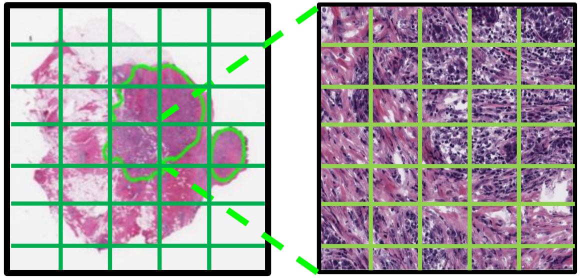

Per-pixel classification or regression is **unstructured**.

Each pixel is considered conditionally independent.

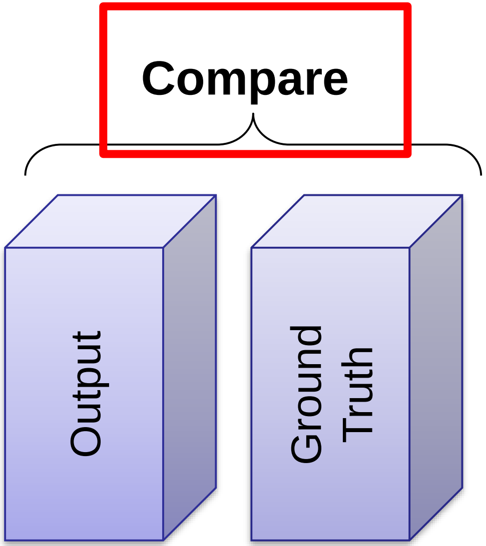

How can we develop conditionally dependent deep learning models?

Deep Learning for Medical Imaging – Major Challenges

JOHNS HOPKINS MEDICINE

4. Incorporating Multimodal Informatio

- Subjective diagnosis is multimodal

Imaging Data

Histopathology

Immunohistochemistry

Radiology

Endoscopy

…

…

How Can we Fuse Unregistered, Uncorrelated and Noisy Data?

How can we separate patient specific and general information?

Multi-omics Data

Genomics

Proteomics

Transcriptomics

miRNAomics

Metabolomics

…

Patient Specific Data

Patient History

Familial History

Clinical Phenotyping

Pharmacologic Data

…

Deep Learning for Medical Imaging – Major Challenges

JOHNS HOPKINS MEDICINE

1. **Limited Annotated Data**
   - Under representation of rare conditions.
   - Limited experts available for annotation.
   - Privacy Issues, Cost
2. **Domain Adaptation**
   - Diversity in data, different sensors, cites and patients.
   - Patient specific texture and color information.
3. **Structured Prediction**
   - Global vs Local features.
4. **Incorporating Multimodal Information**
   - Subjective diagnosis is multimodal.

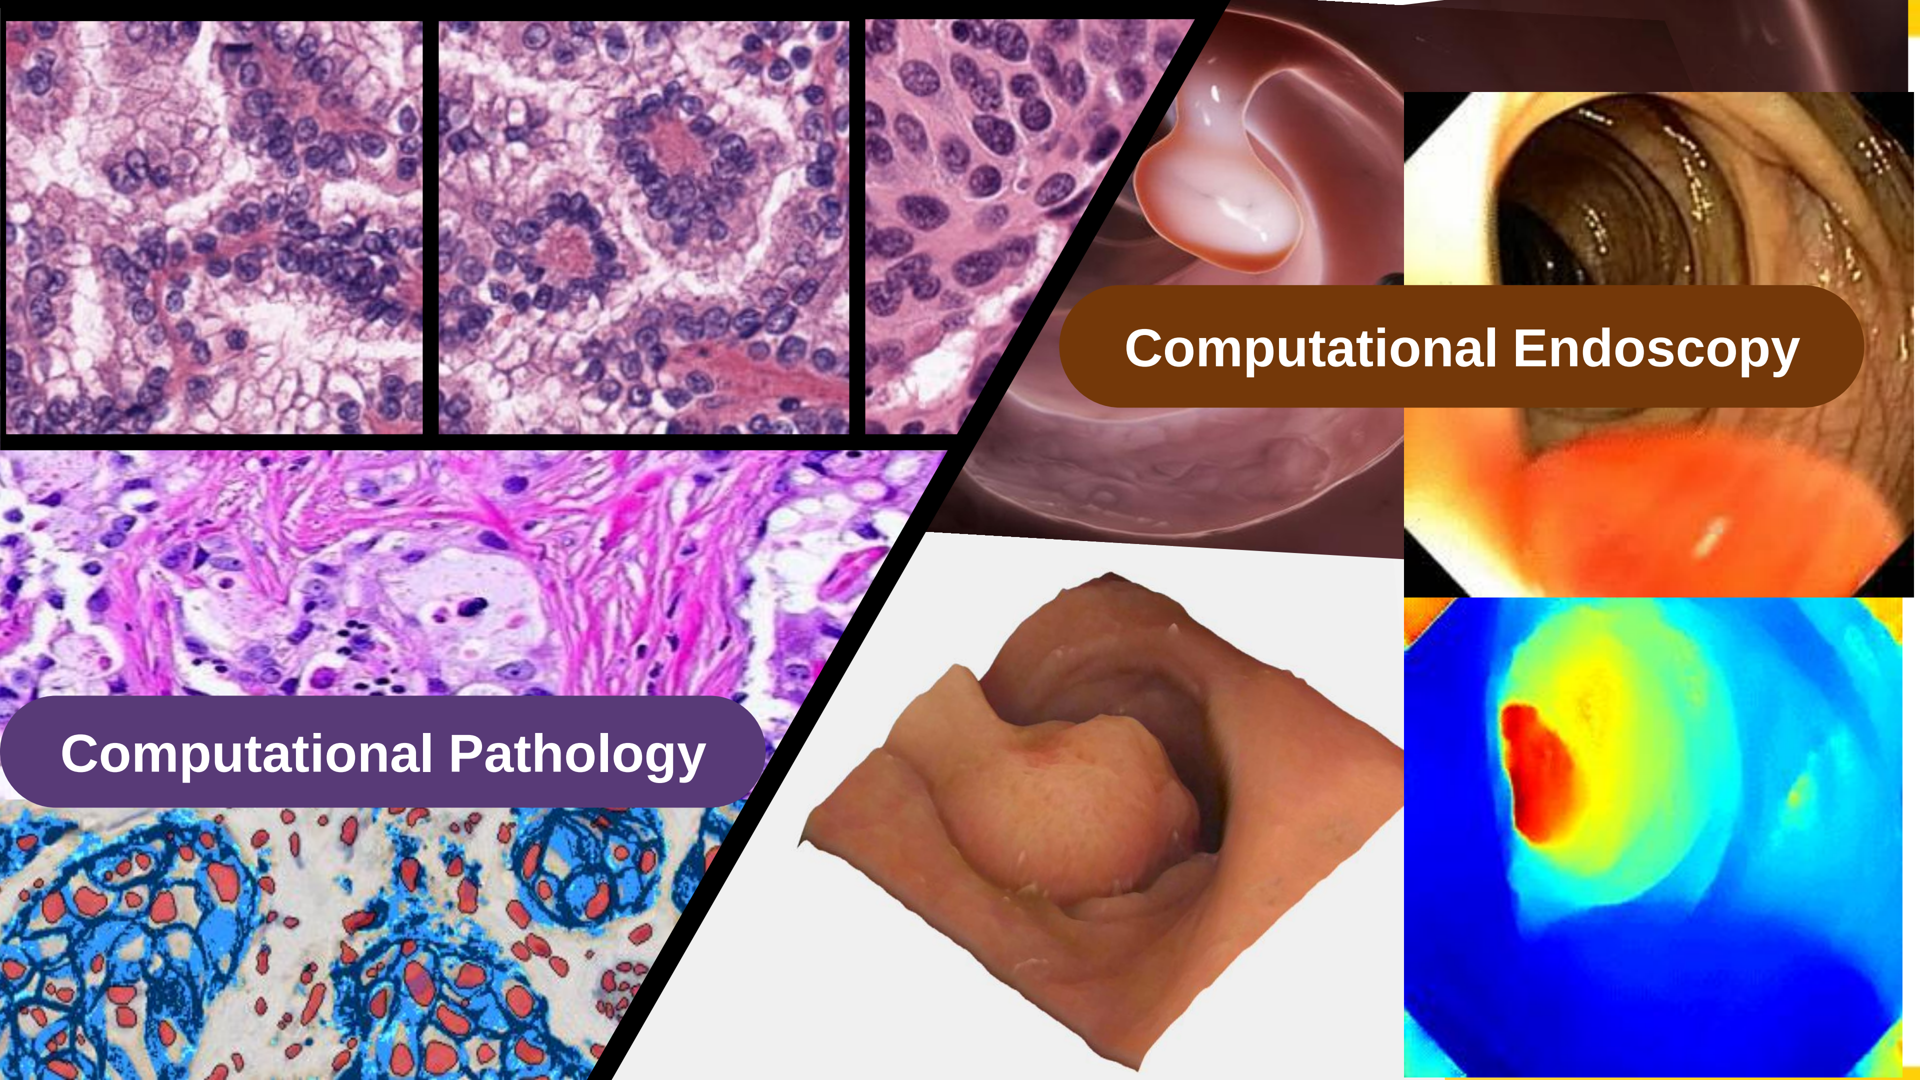
Computational Endoscopy
Computational Pathology

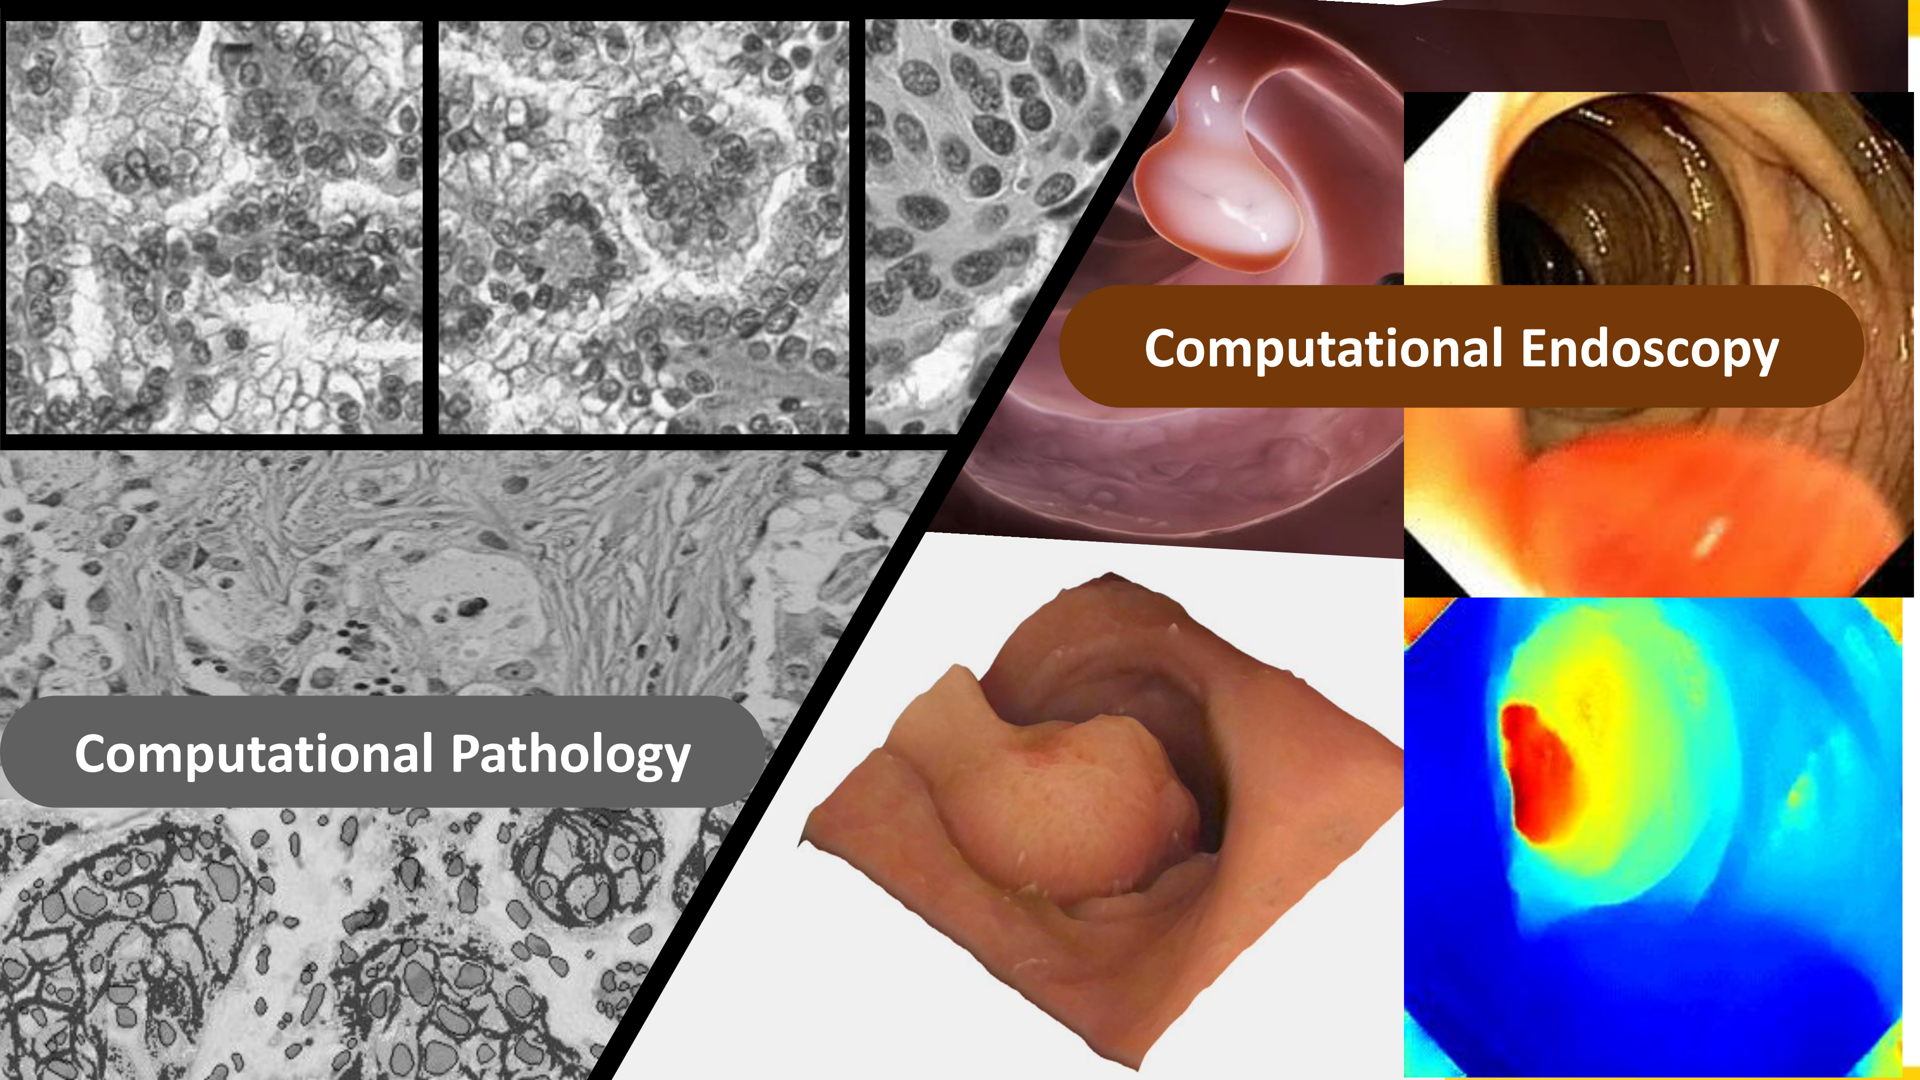
Computational Endoscopy
Computational Pathology

Endoscopic Depth and Topography

Application: Depth Estimation for Endoscopy
Purpose: Predict Topography from Monocular Images

Colonoscopy Gives 2D Images

Topography Matters

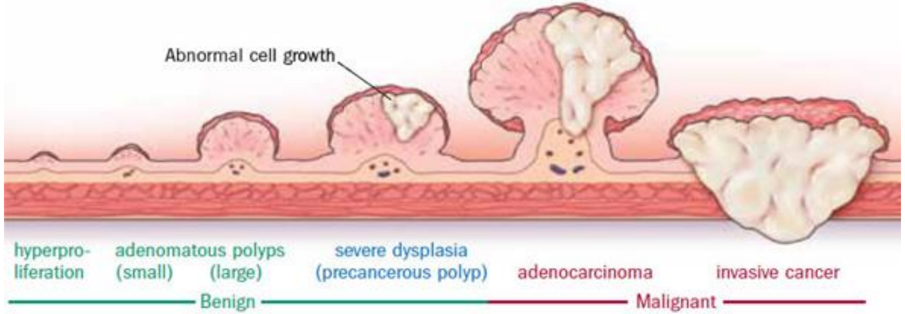

Faisal Mahmood, Nicholas J. Durr et al." Deep learning and conditional random fields-based depth estimation and topographical reconstruction from conventional endoscopy" ***Medical Image Analysis*** (2018).

Endoscopic Depth and Topography

Application: Depth Estimation for Endoscopy
Purpose: Predict Topography from Monocular Images

60% of colorectal cancer cases detected after optical colonoscopy are associated with missed lesions.

How do gastroenterologists predict the presence of a polyp?
Predict the size of the perforations.
Predict surface topography.

Faisal Mahmood, Nicholas J. Durr et al." Deep learning and conditional random fields-based depth estimation and topographical reconstruction from conventional endoscopy" ***Medical Image Analysis*** (2018).

Depth Estimation from Monocular Endoscopy Images

No Ground Truth Depth Data:

- Limited real estate on an endoscope.
- Regulatory approvals required to add depth sensor.

Solution: Generate Synthetic Endoscopy Data

Faisal Mahmood, Nicholas J. Durr et al." Deep learning and conditional random fields-based depth estimation and topographical reconstruction from conventional endoscopy" ***Medical Image Analysis*** (2018).

Generating Synthetic Endoscopy Data with GT Depth

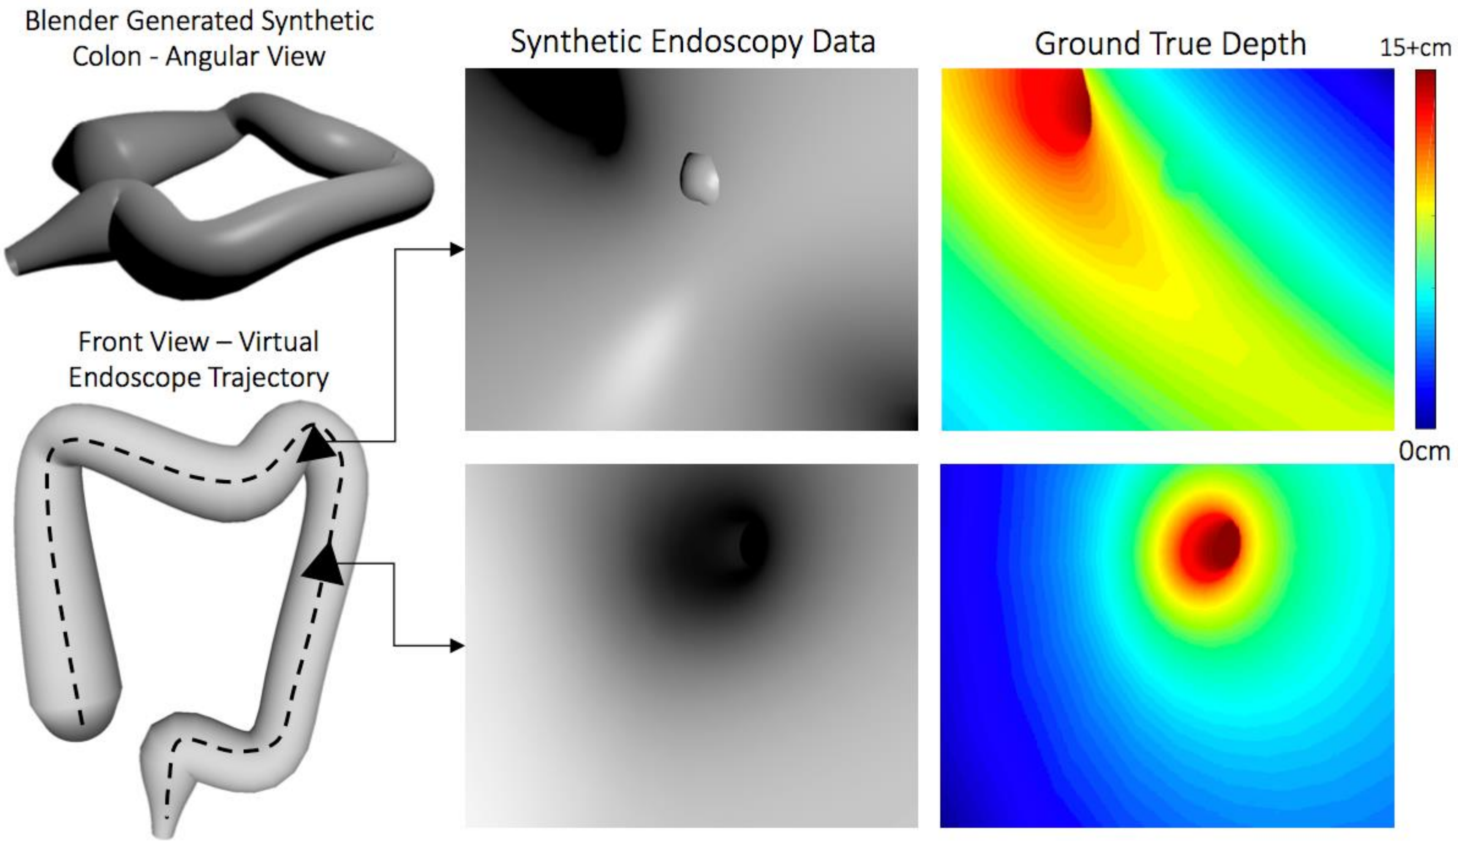

(F. Mahmood et al., 2018)

Generating Synthetic Endoscopy Data with GT Depth

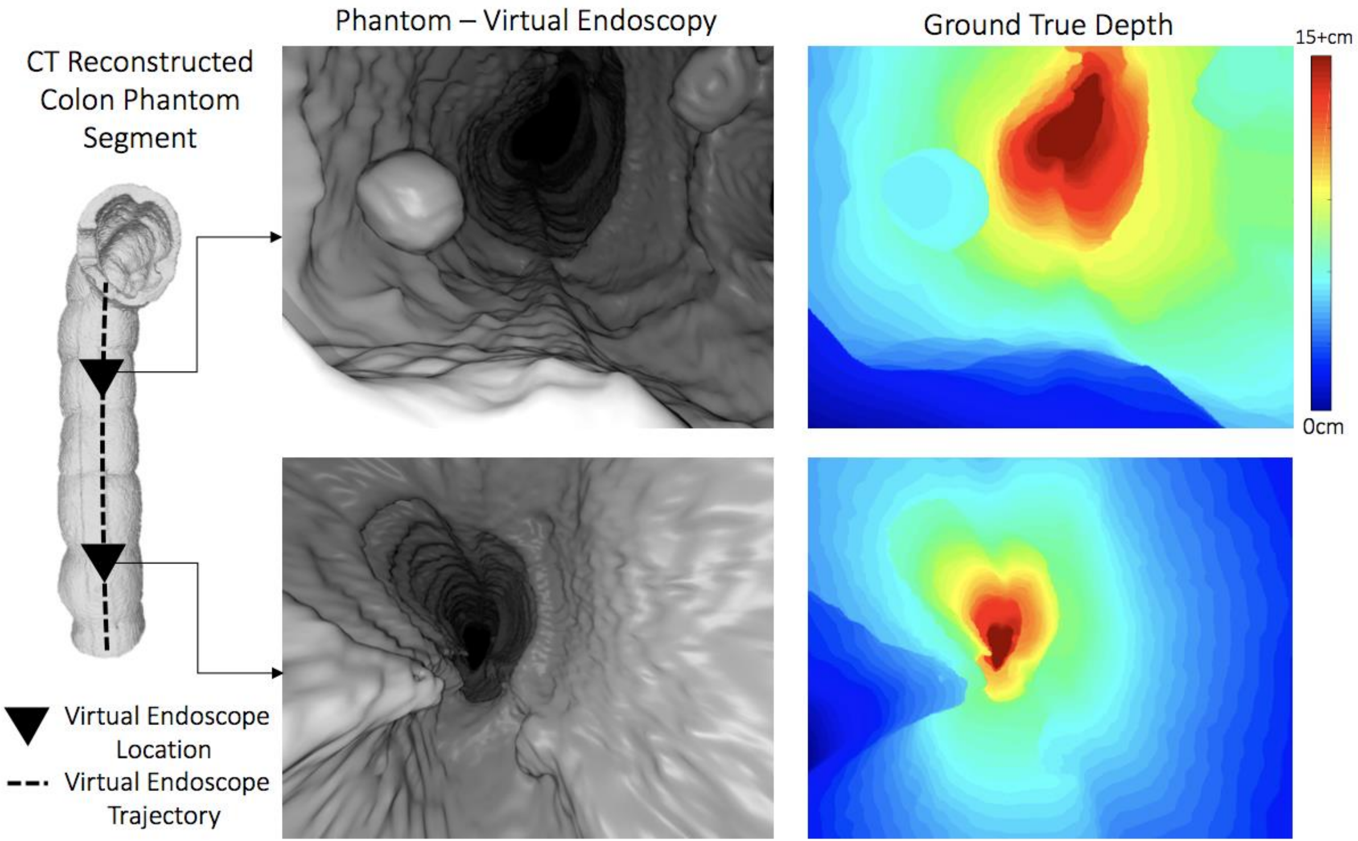

(F. Mahmood et al., 2018)

Deep Learning for Medical Imaging – Major Challenges

JOHNS HOPKINS MEDICINE

1. **Limited Annotated Data**
   - Under representation of rare conditions.
   - Limited experts available for annotation.
   - Privacy Issues
   - Cost
2. **Domain Adaptation**
   - Diversity in data, different sensors, cites and patients.
   - Patient specific texture and color information.
3. **Structured Prediction**
   - Global vs Local features.
4. **Incorporating Multimodal Information**
   - Subjective diagnosis is multimodal.

Training with Endoscopy Synthetic Data

Problem:

Standard Deep Learning Networks are not sufficiently context aware.

Solution:

Add non-local information using a joint CNN-Graphical Model Setup.

(F. Mahmood et al., 2018)

Solution: Joint CNN-CRF Model

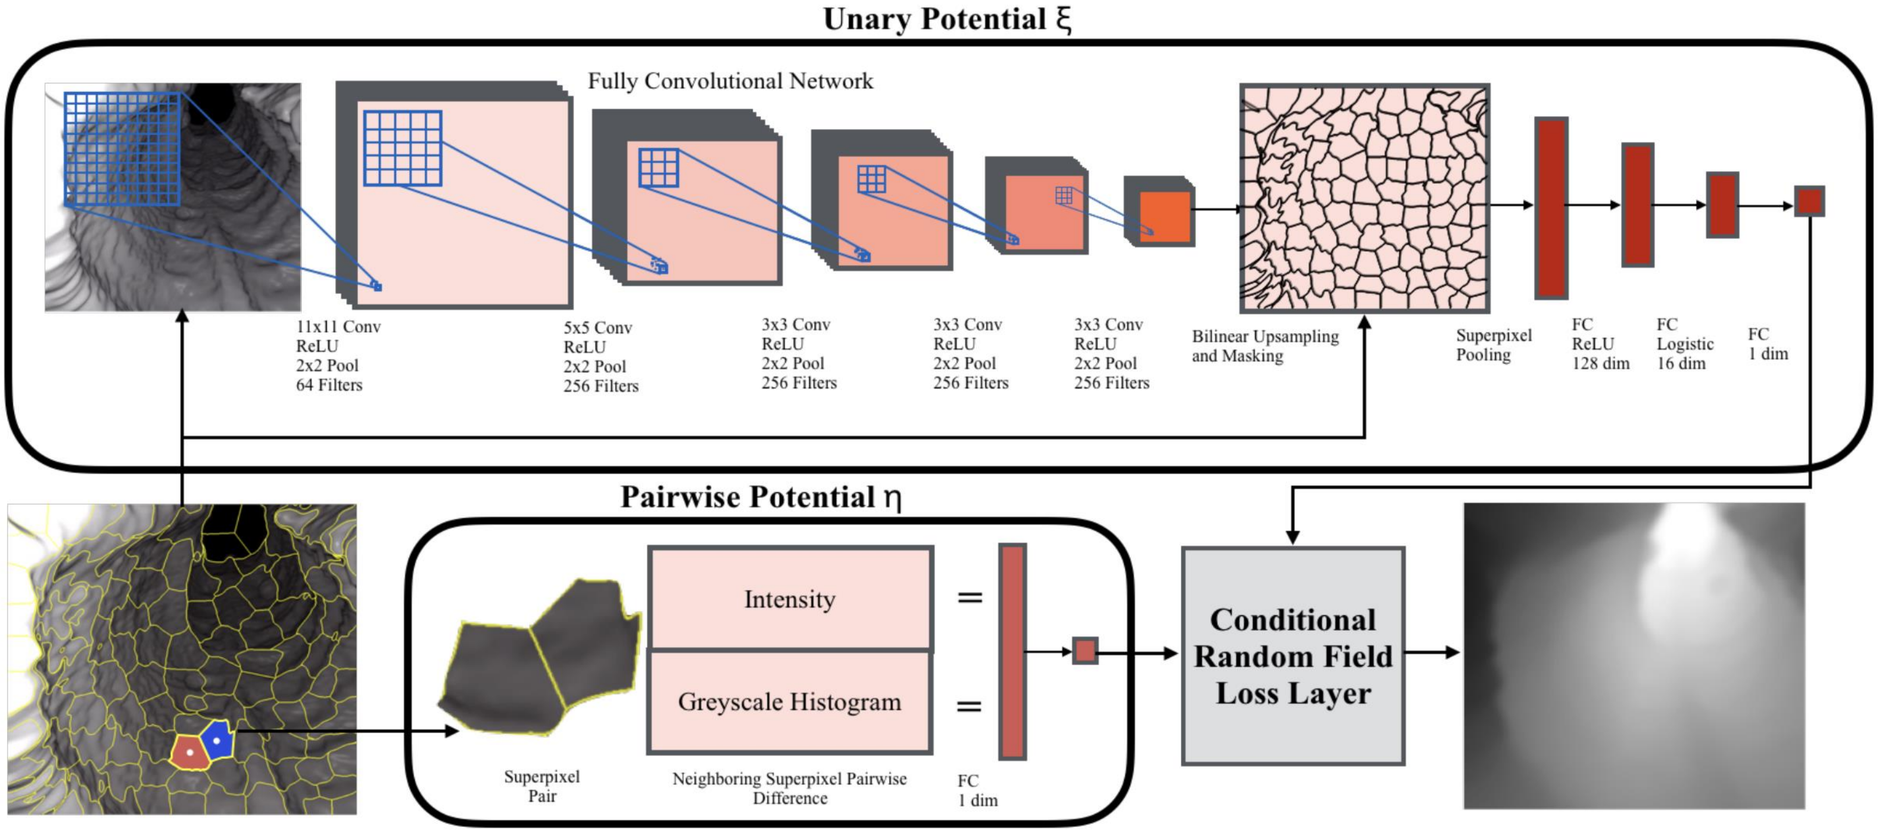

(F. Mahmood et al., 2018)

Solution: Joint CNN-CRF Model

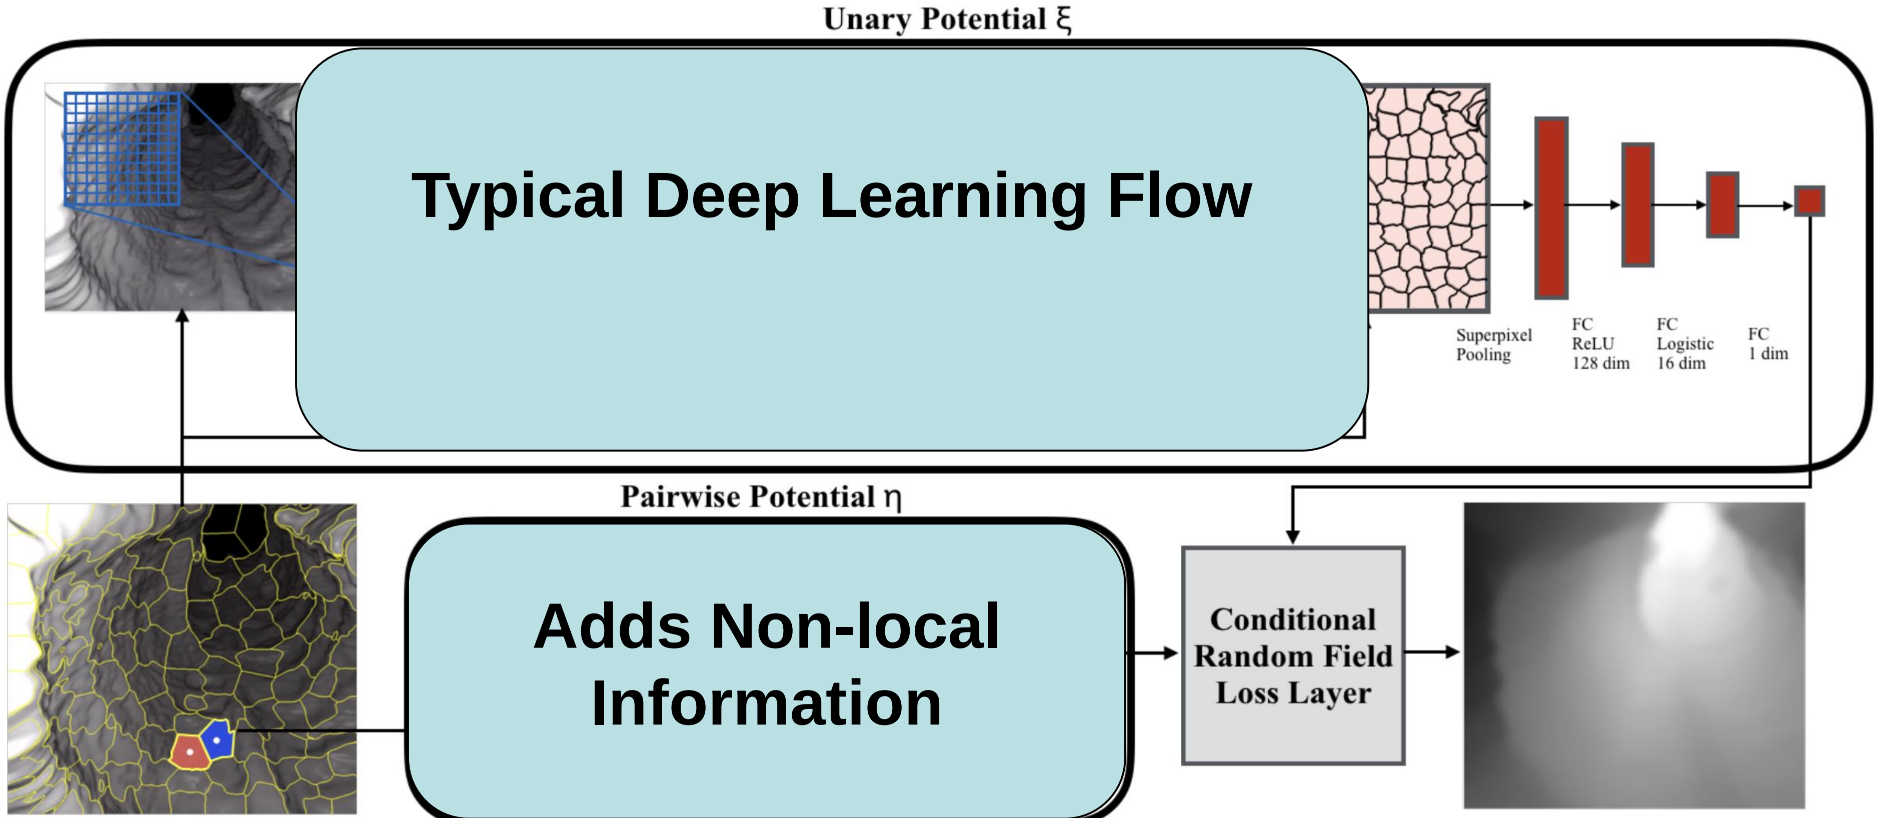

(F. Mahmood et al., 2018)

Deep Learning for Medical Imaging – Major Challenges

JOHNS HOPKINS MEDICINE

1. **Limited Annotated Data**
   - Under representation of rare conditions.
   - Limited experts available for annotation.
   - Privacy Issues
   - Cost
2. **Domain Adaptation**
   - Diversity in data, different sensors, cites and patients.
   - Patient specific texture and color information.
3. **Structured Prediction**
   - Global vs Local features.
4. **Incorporating Multimodal Information**
   - Subjective diagnosis is multimodal.

Adapting Synthetic Networks to Real Data

JOHNS HOPKINS MEDICINE

Problem: Network trained on synthetic data does not work with real data.

Solution: Adversarial Reverse Domain Adaptation.

Train

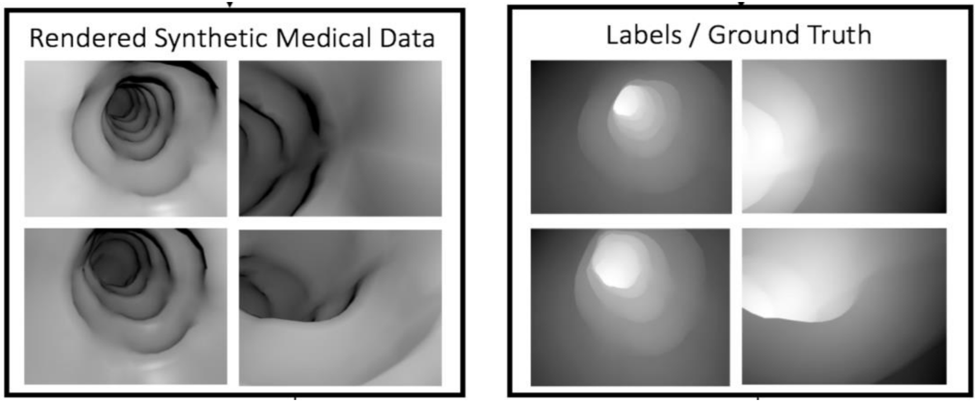

Test

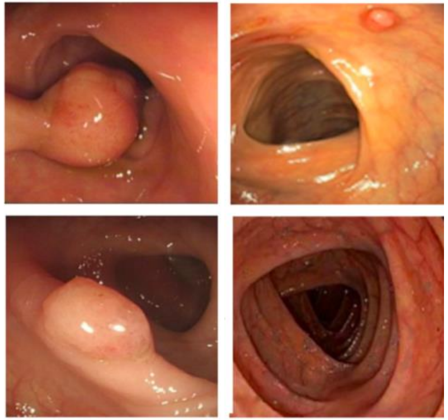

(Mahmood et al., 2018)

Adversarial Reverse Domain Adaptation

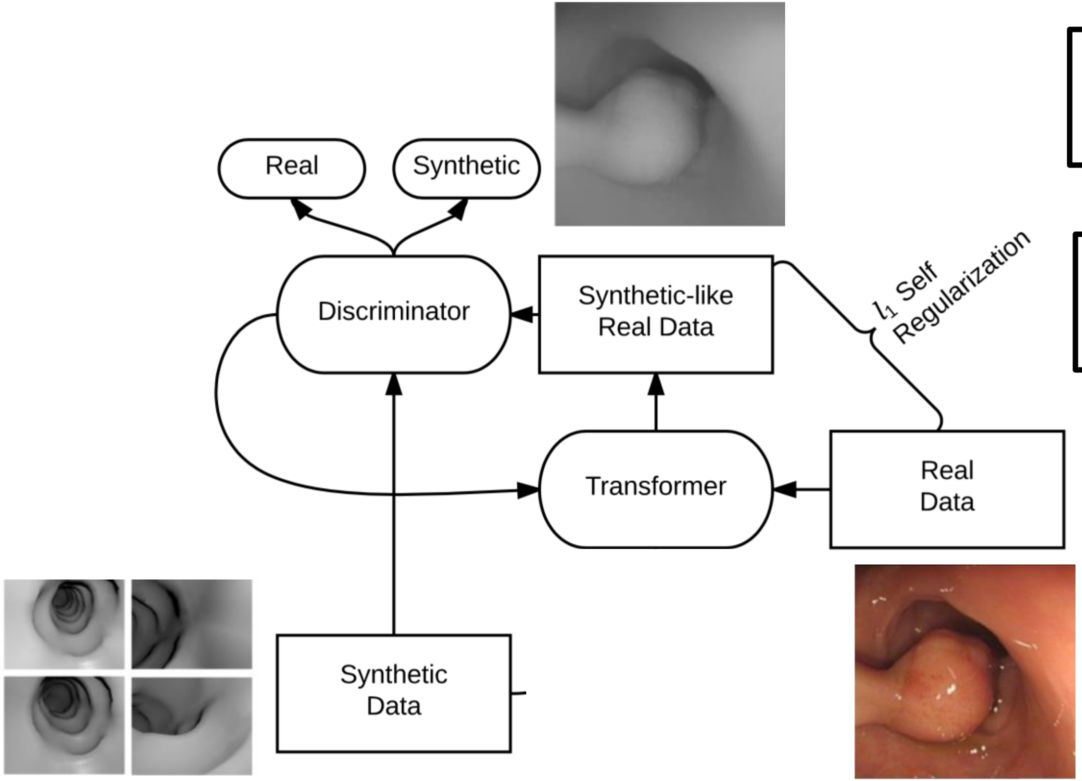

Typical Flow: Transfer Synthetic Data to Real-like Domain.

Proposed Flow: Transfer Real Data to Synthetic-like Domain.

Remove Patient Specific Features while preserving features necessary for depth estimation.

F. Mahmood, Nicholas J. Durr et al."Unsupervised Reverse Domain Adaptation for Synthetic Medical Images via Adversarial Training." ***IEEE Transactions on Medical Imaging*** (2018).

Adversarial Reverse Domain Adaptation

Endoscopy Images	Synthetic-like Representation

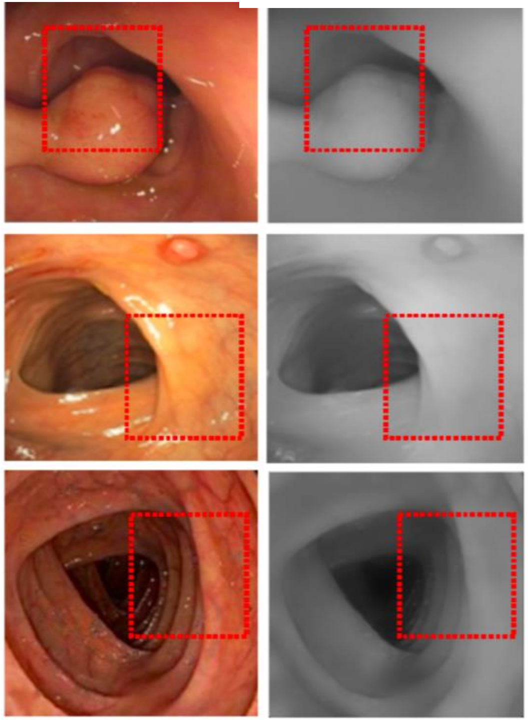

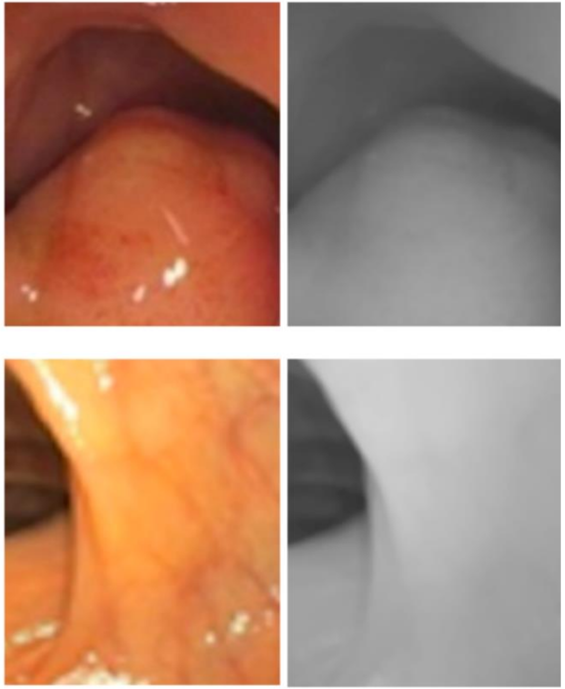

Shape, Shading, Intensity Preserved

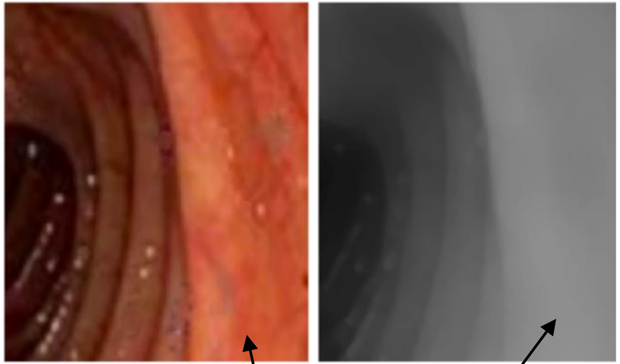

Patient Specific Details Removed

(Mahmood et al., 2018)

Endoscopy Depth Estimation

Colonoscopy Video

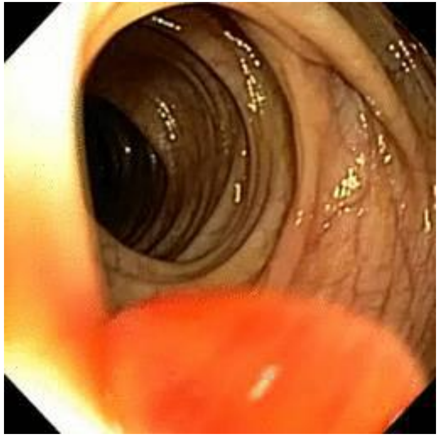

Depth Estimate

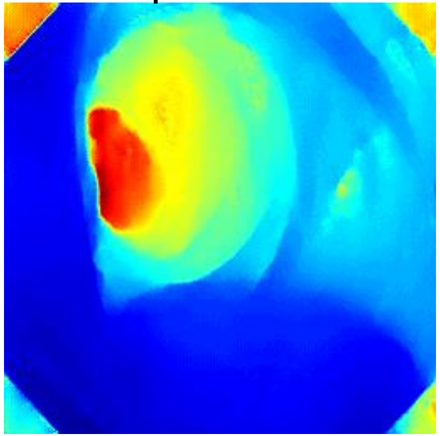

Colonoscopy Video

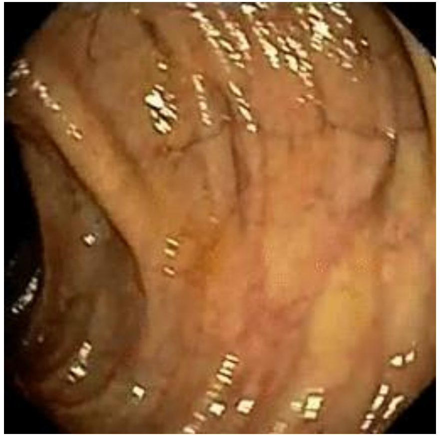

Depth Estimate

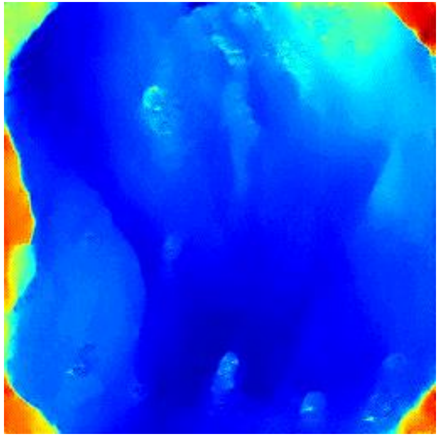

(Mahmood et al., 2018)

Validation – Endoscopy Depth Estimation

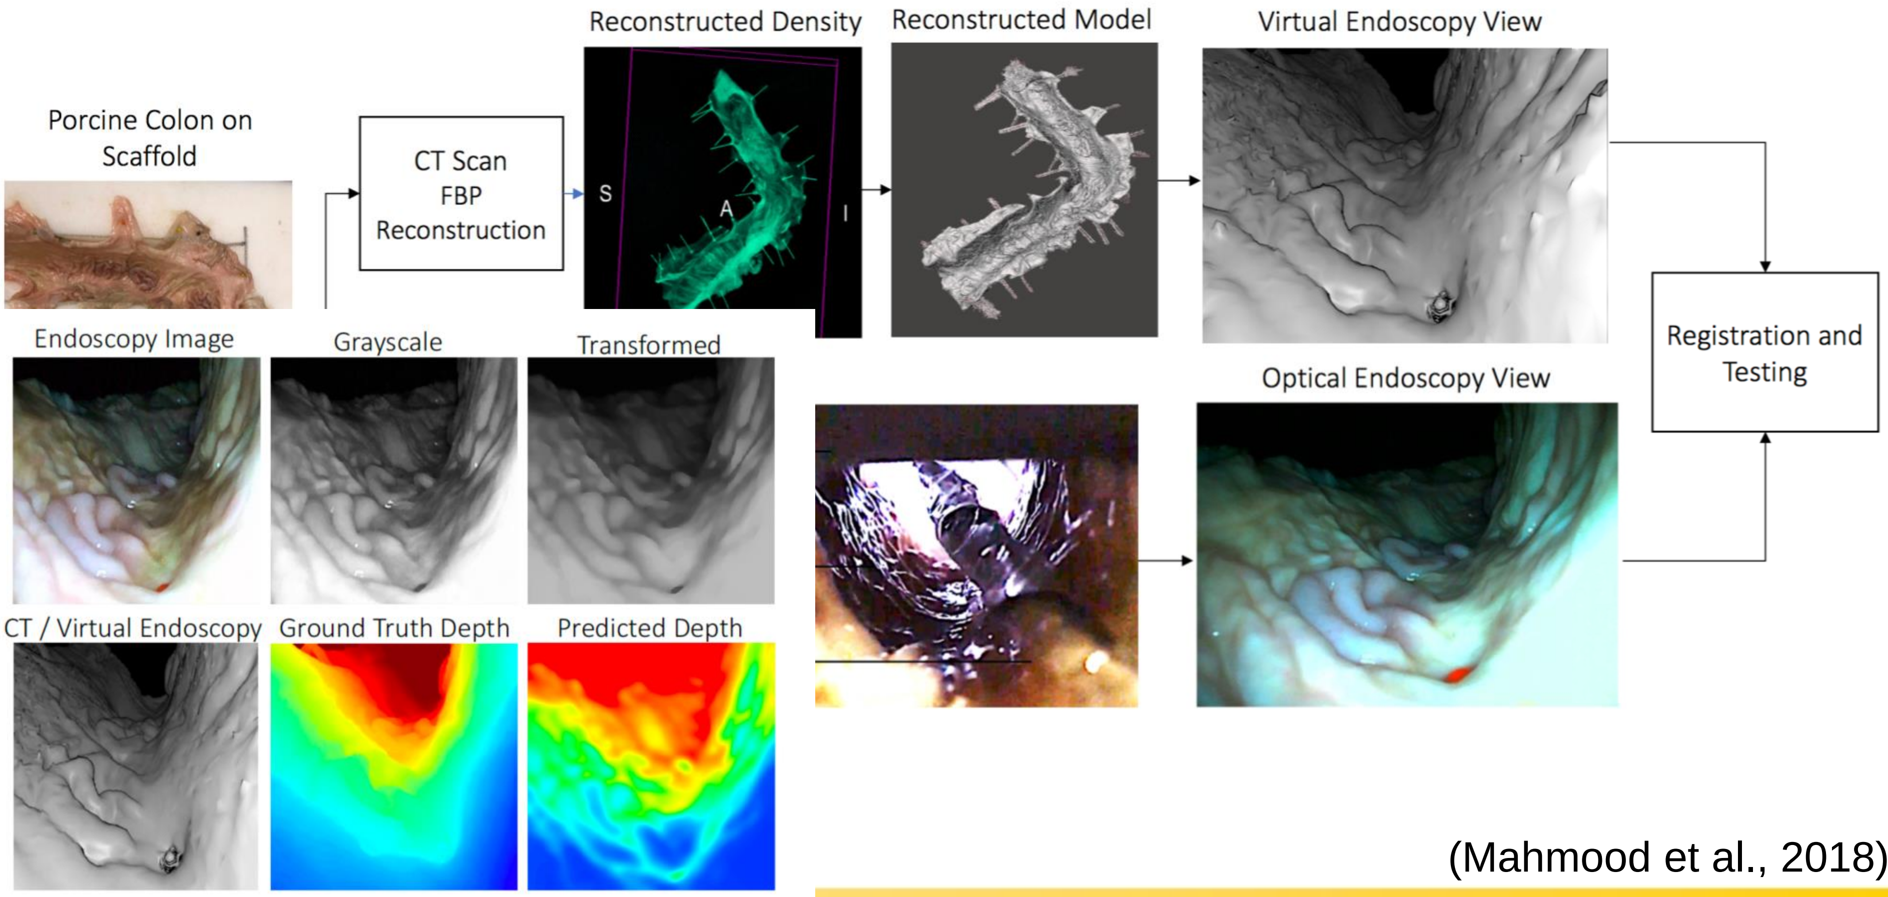

(Mahmood et al., 2018)

Estimated Depth to Topography

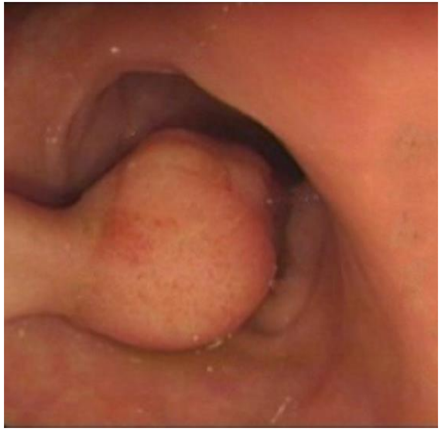

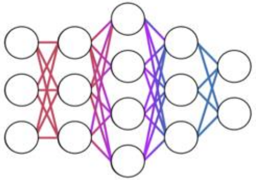

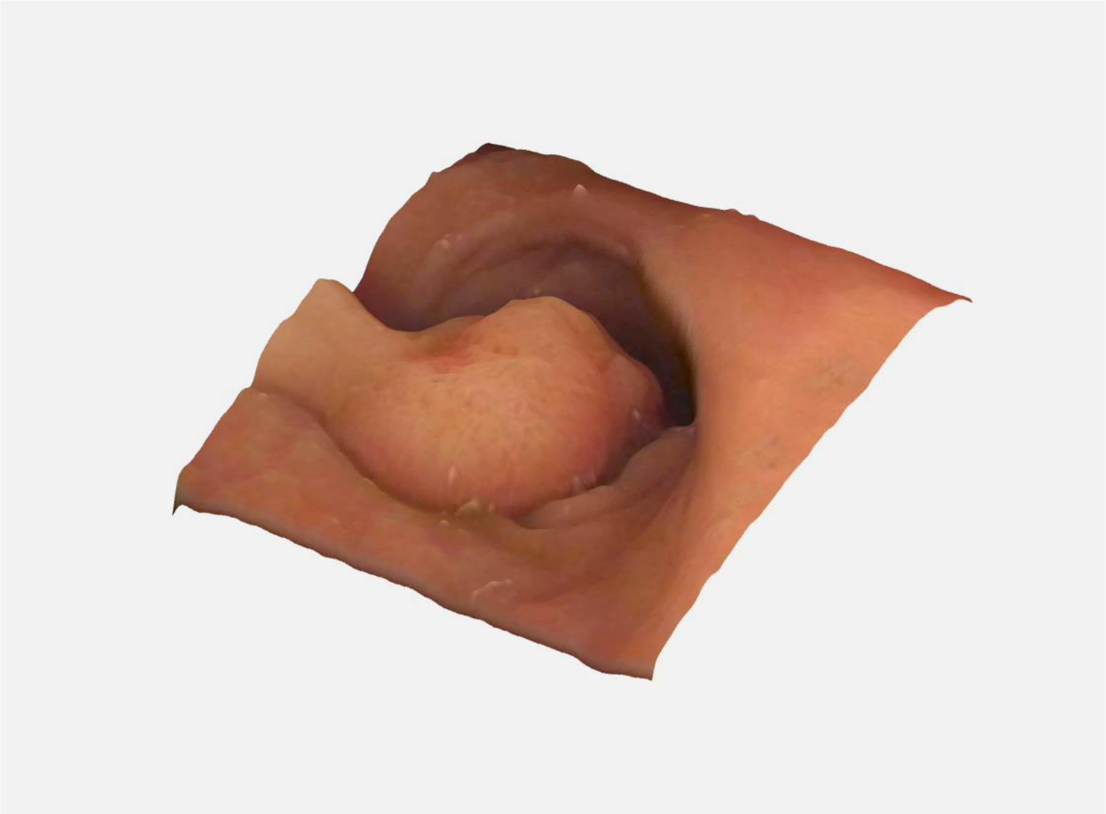

Faisal Mahmood, Nicholas J. Durr et al."Unsupervised Reverse Domain Adaptation for Synthetic Medical Images via Adversarial Training."
IEEE Transactions on Medical Imaging (2018).

Polyp Charcterization

Hyperplastic

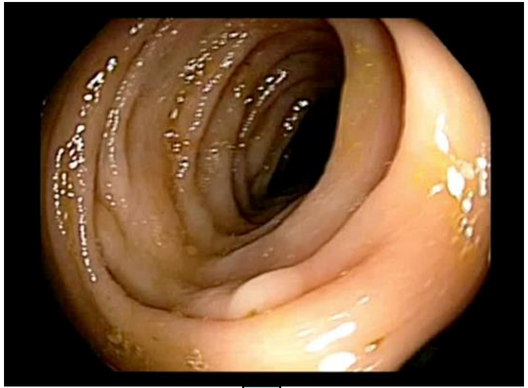

Serrated

Adenoma

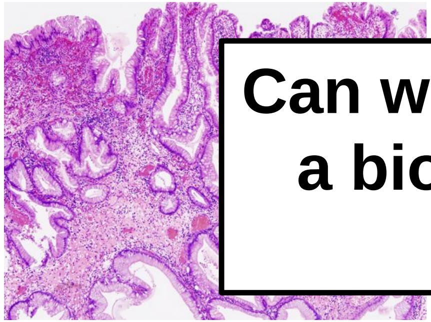

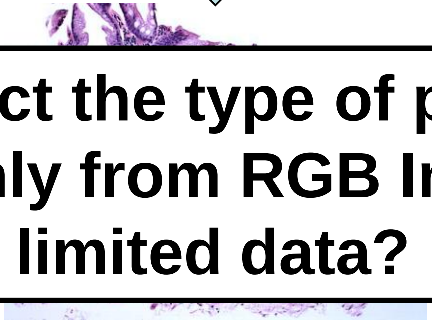

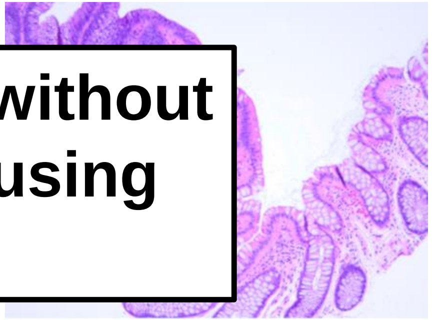

Can we predict the type of polyp without a biopsy only from RGB Image using limited data?

Polyp Charcterization

Hyperplastic

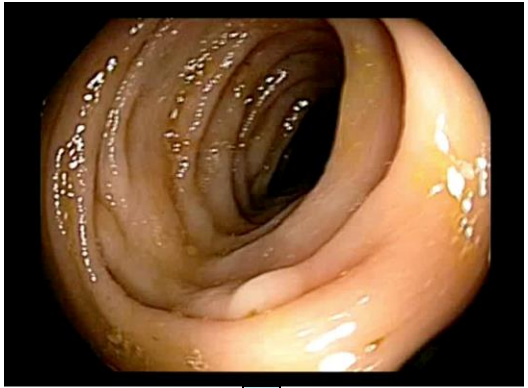

Serrated

Adenoma

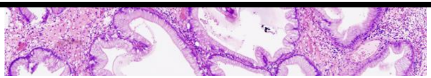

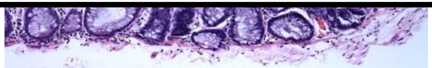

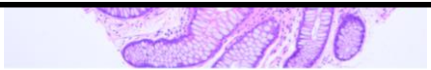

- 76 Videos
- All videos labeled by 4 Senior Gastroenterologists & 3 Fellows
- Average GI Accuracy: Senior: **63.4%** Fellow: **53.7%**

Deep Learning for Medical Imaging – Major Challenges

JOHNS HOPKINS MEDICINE

1. **Limited Annotated Data**
   - Under representation of rare conditions.
   - Limited experts available for annotation.
   - Privacy Issues
   - Cost
2. **Domain Adaptation**
   - Diversity in data, different sensors, cites and patients.
   - Patient specific texture and color information.
3. **Structured Prediction**
   - Global vs Local features.
4. **Incorporating Multimodal Information**
   - Subjective diagnosis is multimodal.

Multimodal Data Fusion

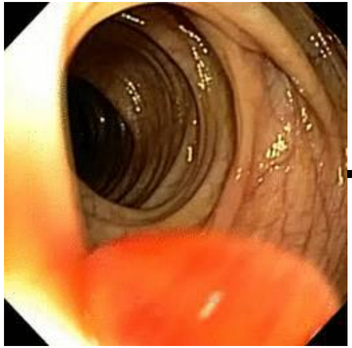

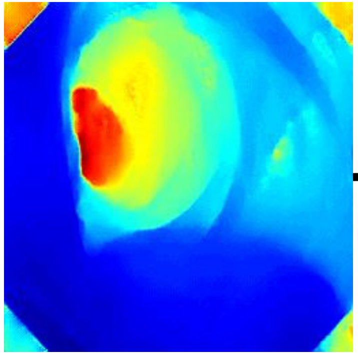

Fusion Network

Depth is an additional predicted modality.

Can we fuse depth and RGB to get better polyp classification results?

Does depth fusion help train a network that requires less data?

RGB-D Classification via Depth Fusion

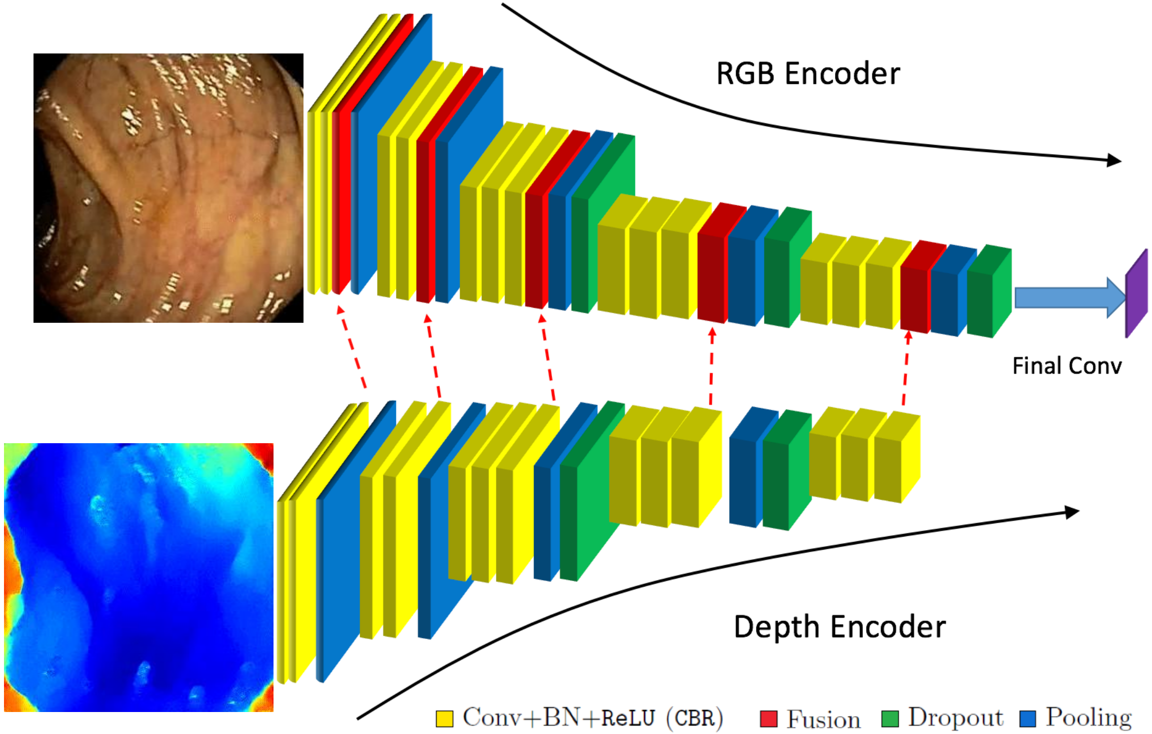

Data Fusion in Feature Space is better than Concatenation.

(Mahmood et al., 2018)

Multimodal Densenet

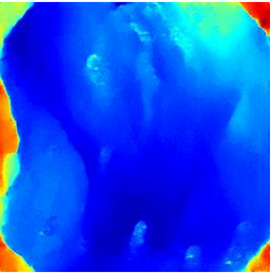

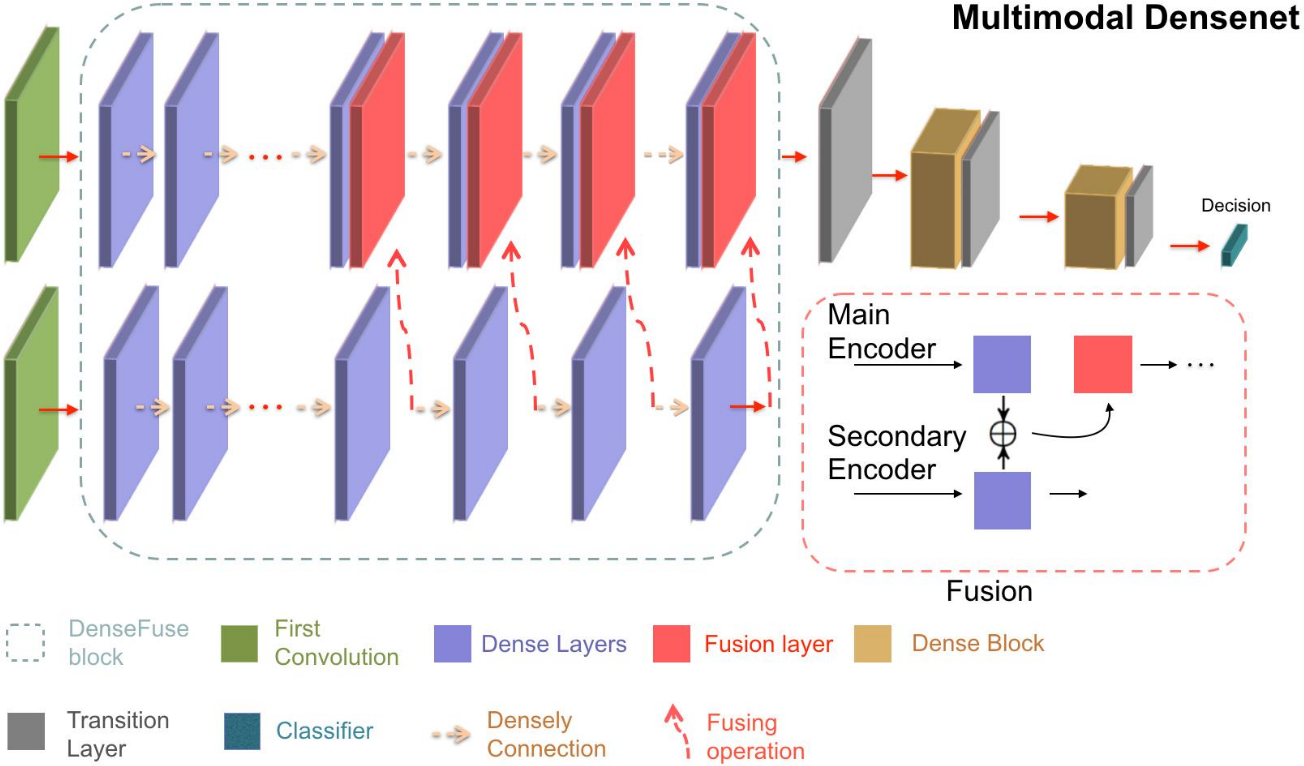

(F. Mahmood et al., 2018)

RGB-D Classification

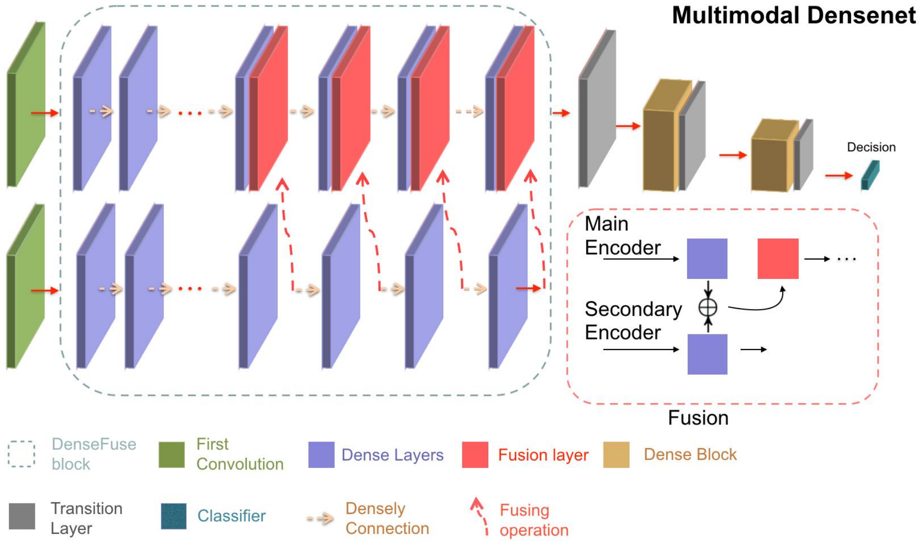

RGB vs RGB-D

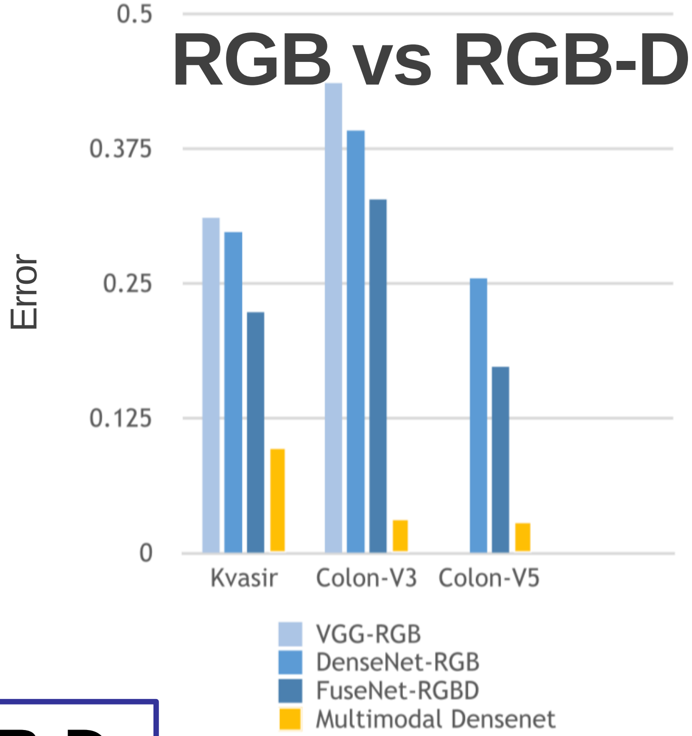

Average Accuracy = 93% using RGB-D

RGB vs RGB-D Classification

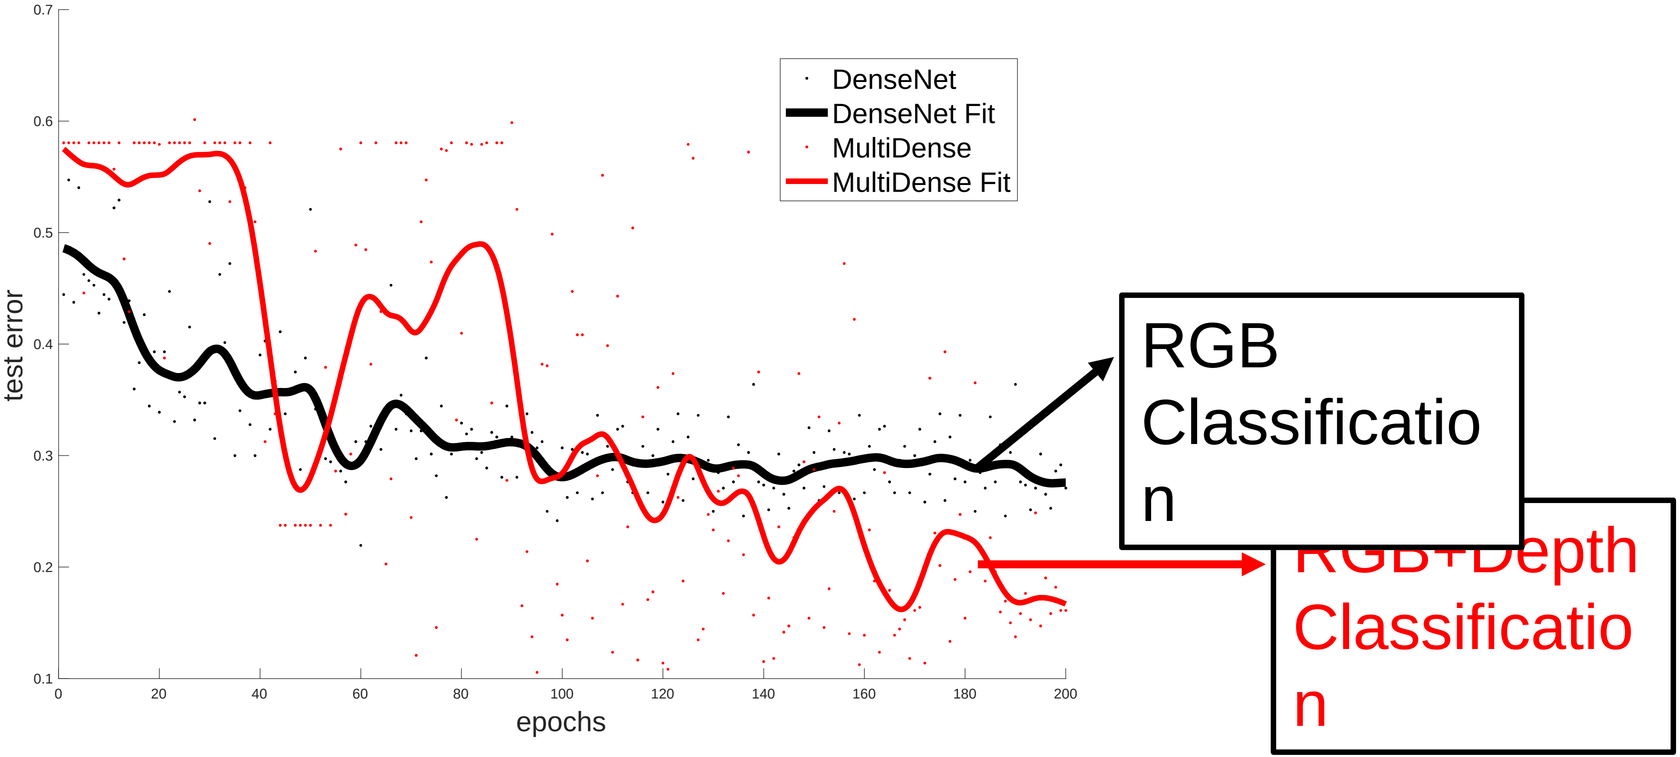

(F. Mahmood et al., 2018)

Gradient Class Activation Maps

Adenoma

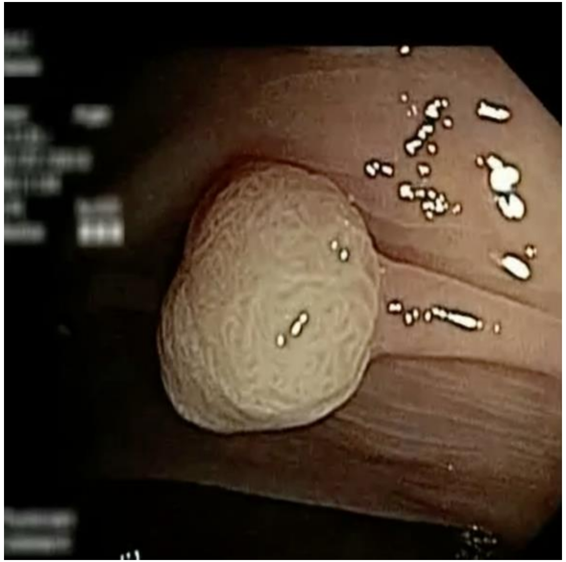

RGB
Classification

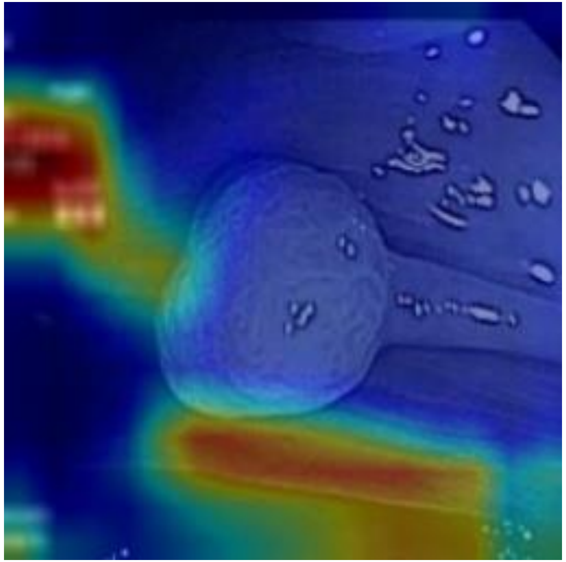

RGB-D
Classification

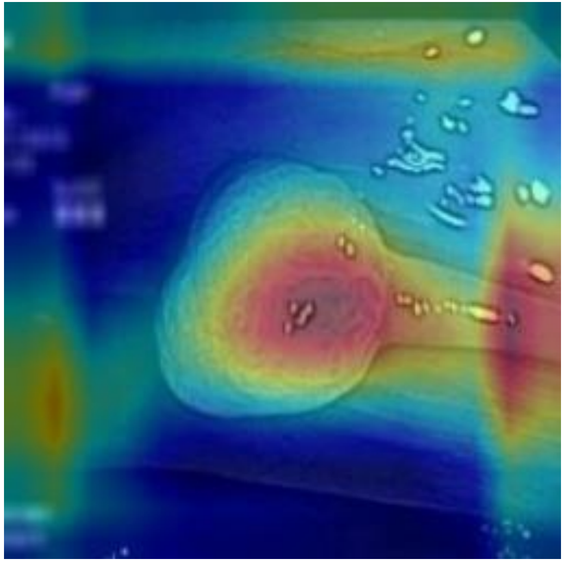

(F. Mahmood et al., 2018)

Gradient Class Activation Maps

Input

RGB Classification

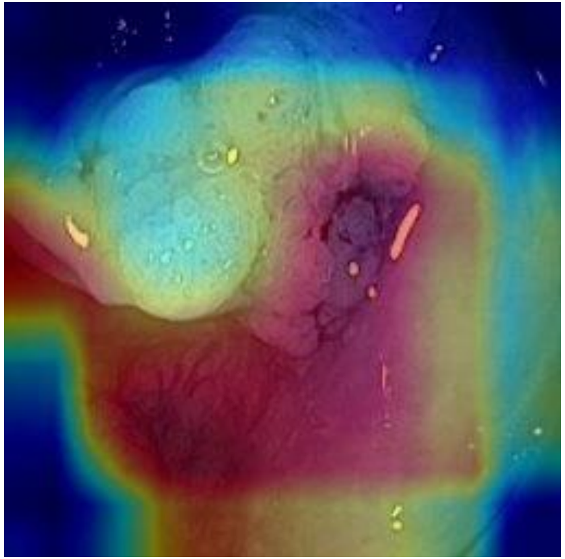

RGB-D Classification

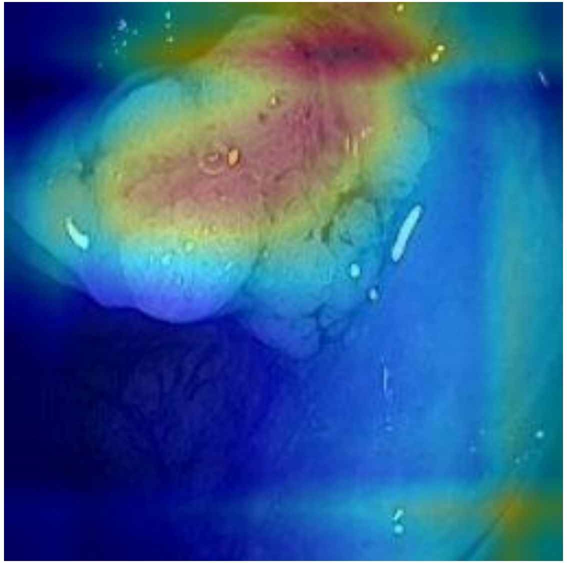

Using just 76 polyp videos with fused depth it is possible to build a classifier with AUC > 0.9

(F. Mahmood et al., 2018)

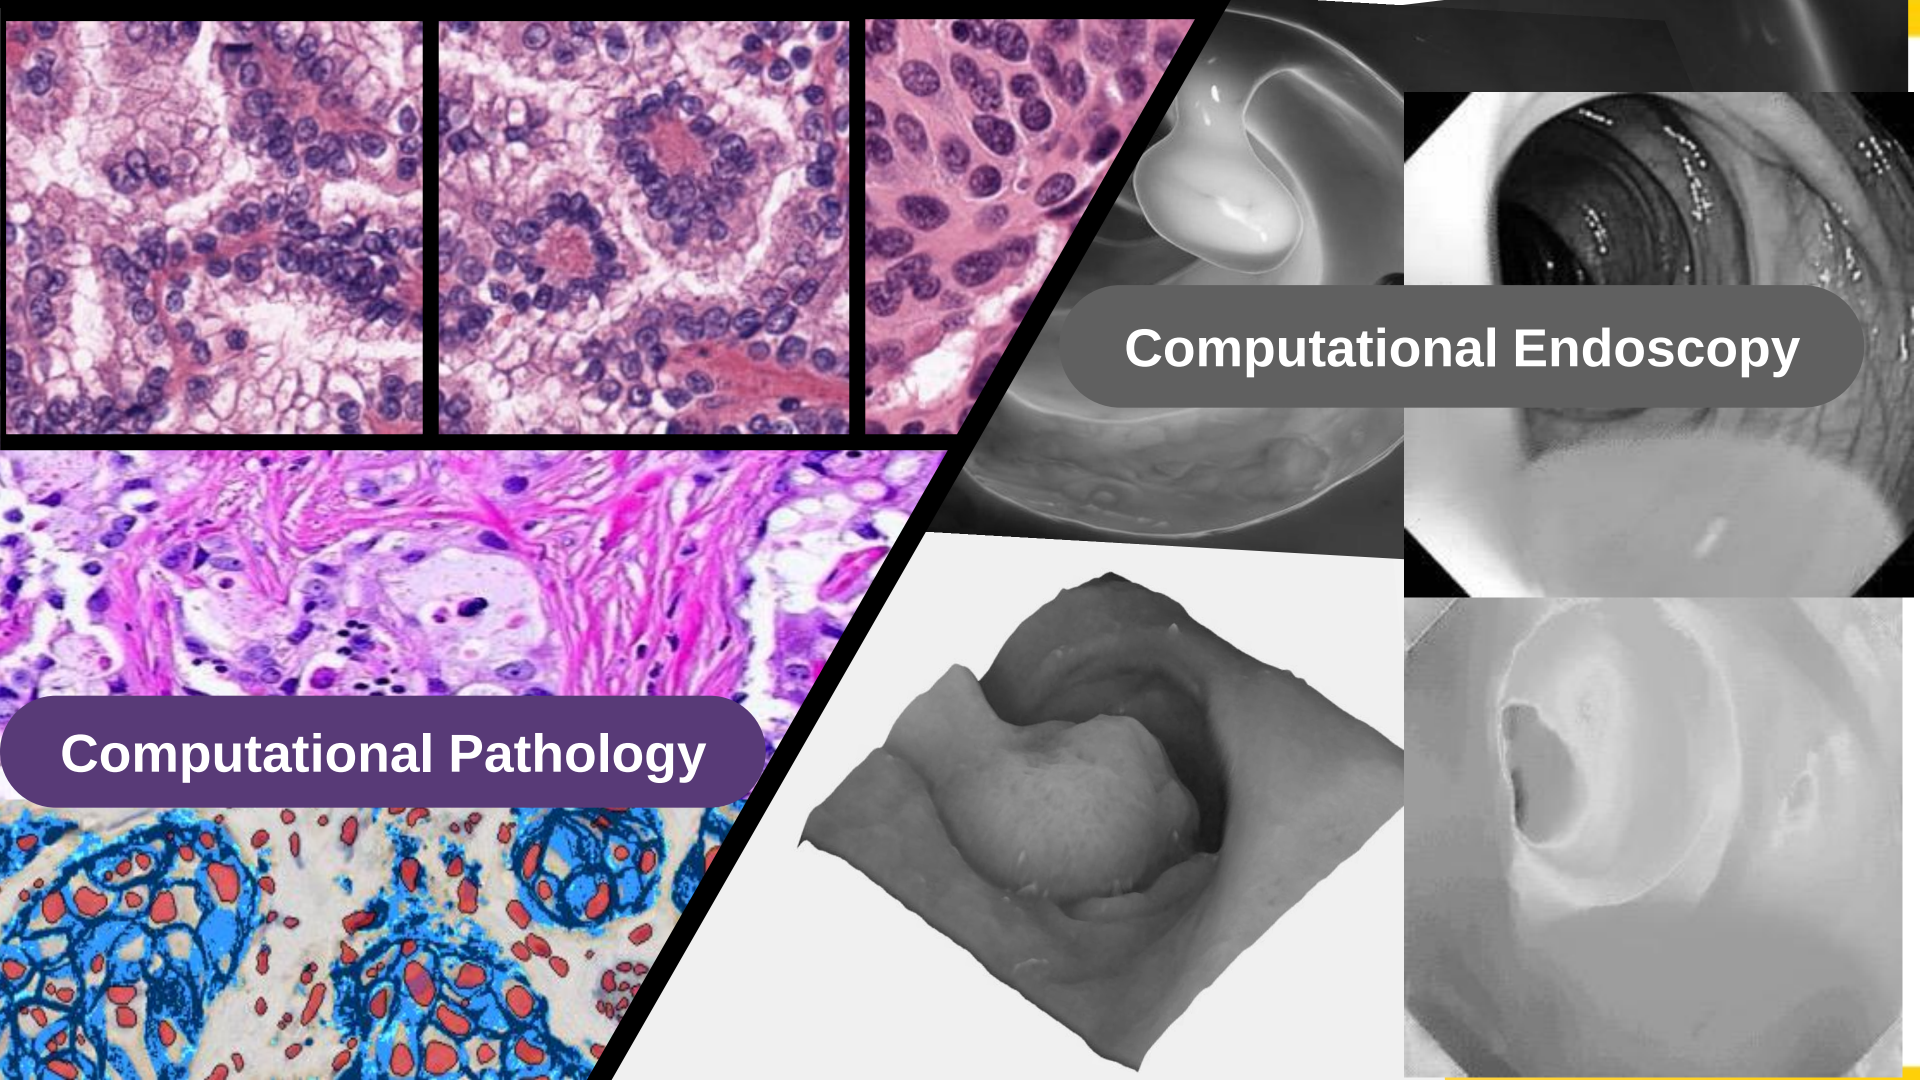
Computational Endoscopy
Computational Pathology

Automated Breast Cancer Grading

Nuclear Atypia

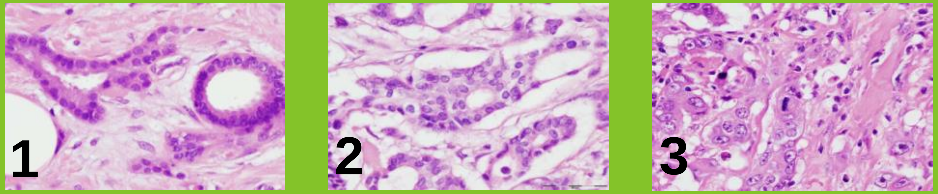

Tubule/Gland Formation

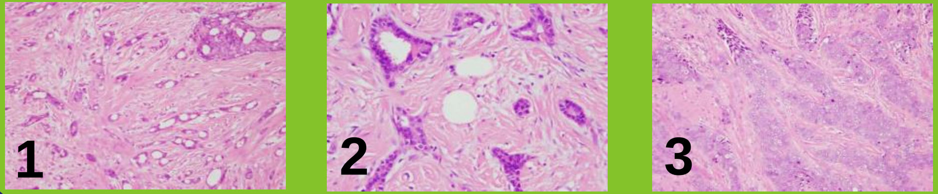

Mitotic Activity

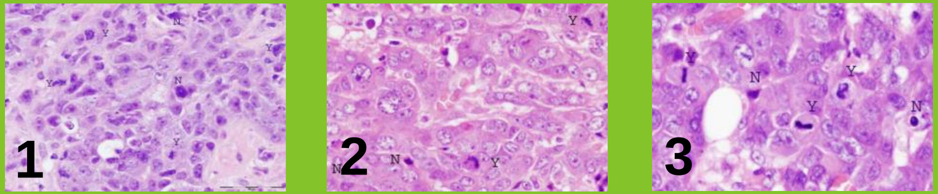

Typical AI for Pathology Flow

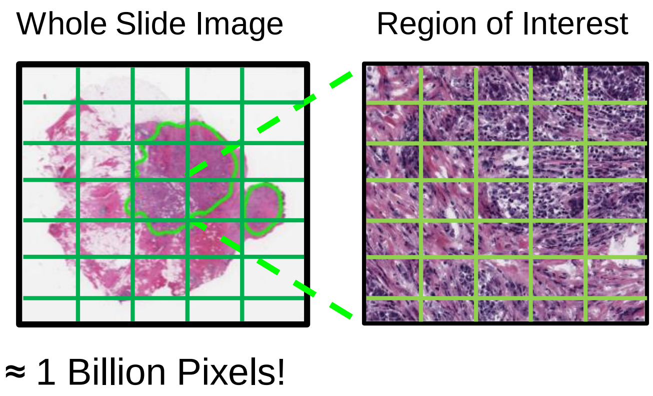

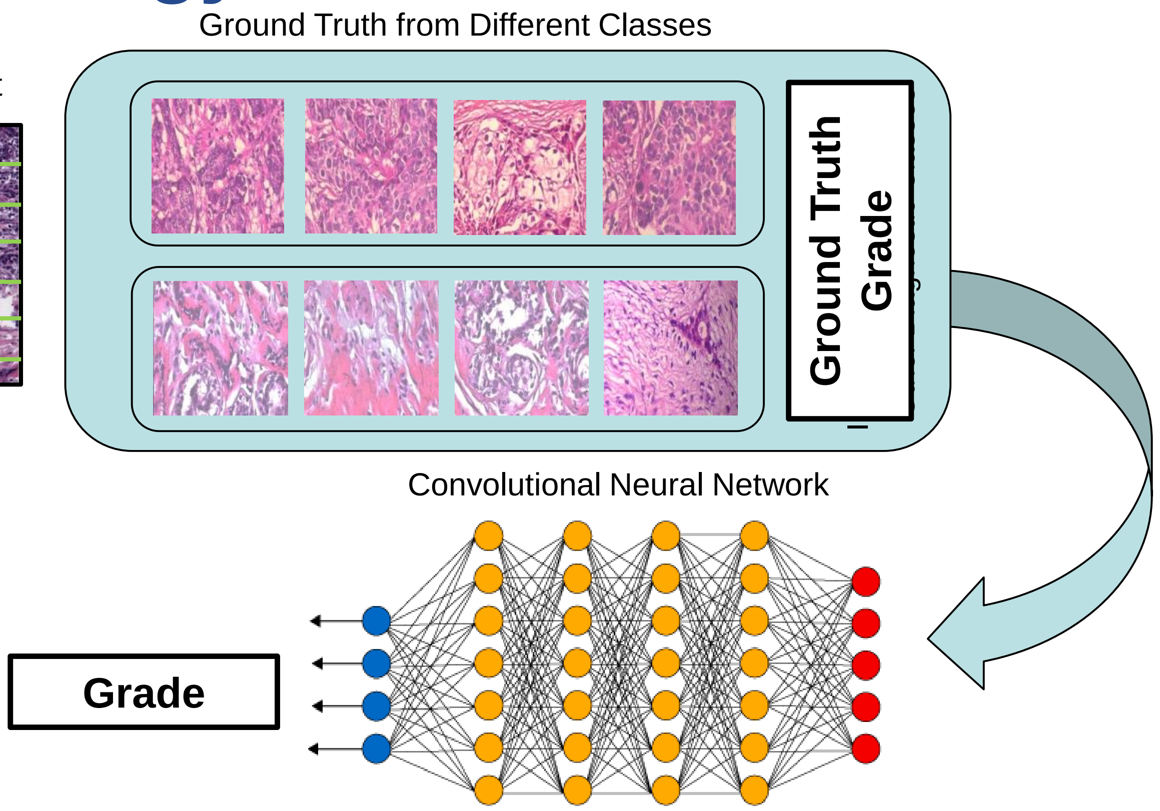

This Needs a Lot of Labeled Data!

Interobserver & Intraobserver Variability?

Automated Breast Cancer Grading

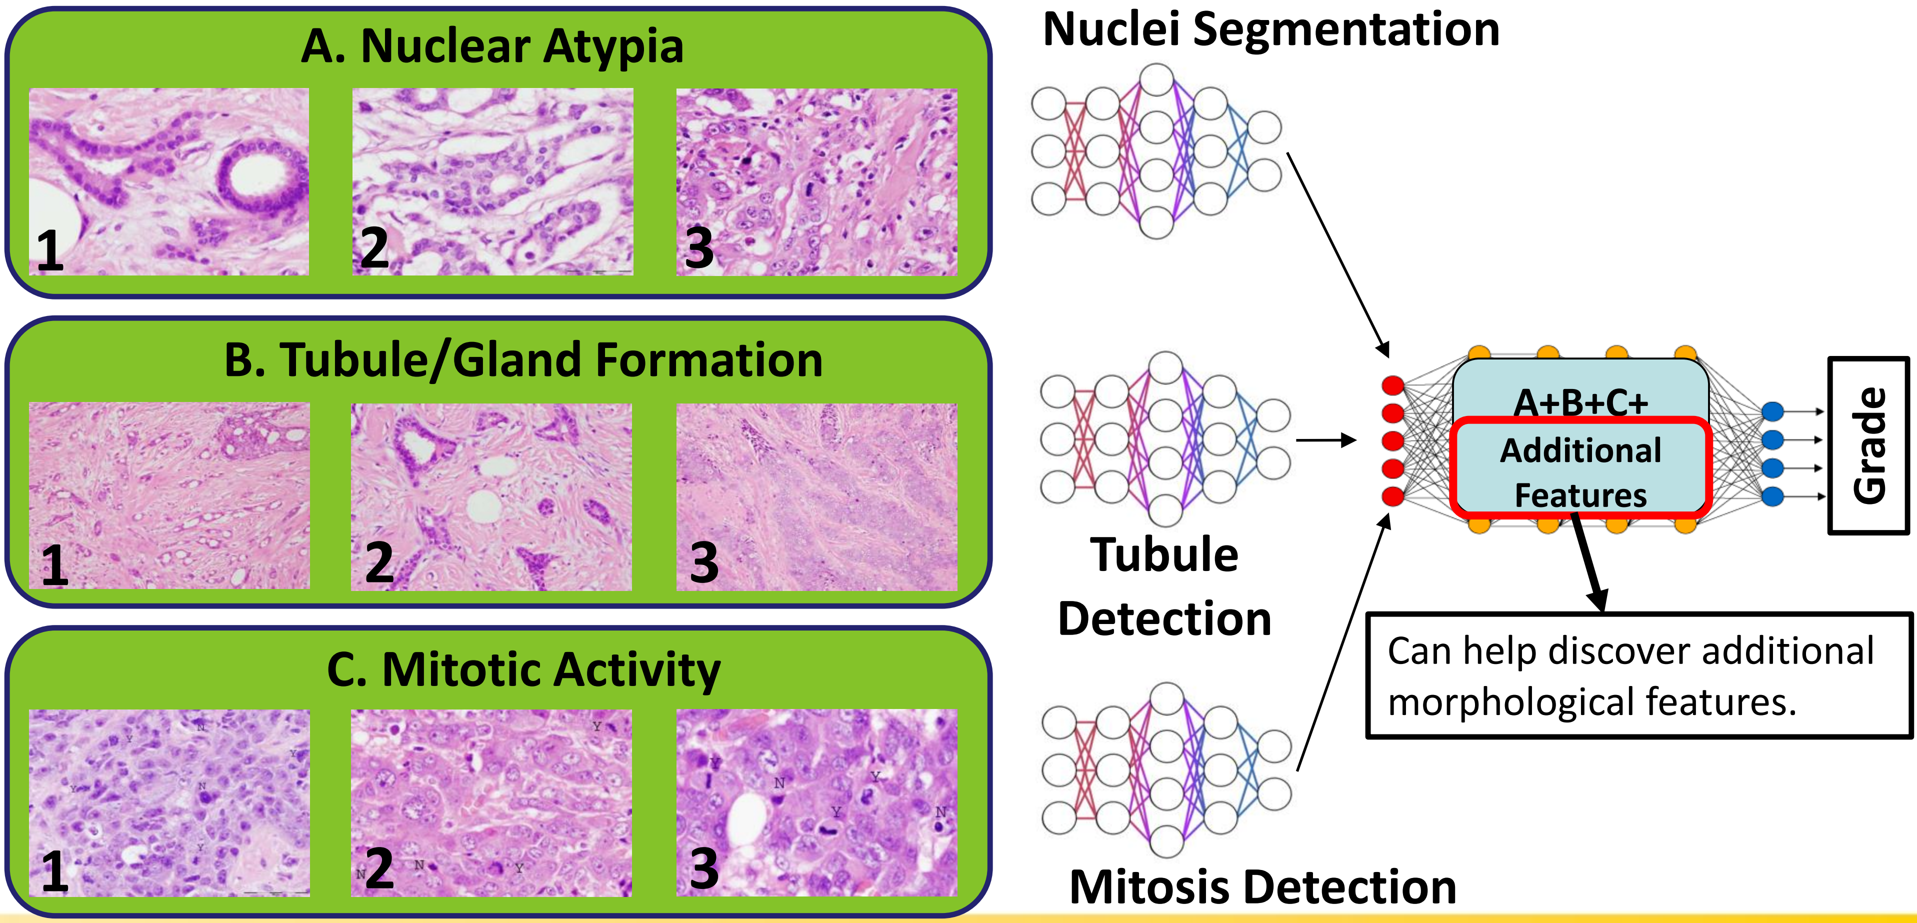

Automated Breast Cancer Grading

A. Nuclear Atypia

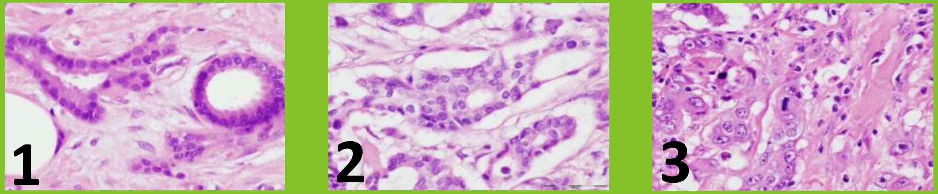

Nuclei Segmentation

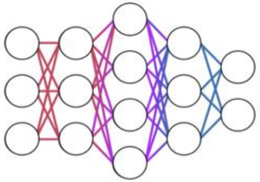

CellProfiler™
cell image analysis software

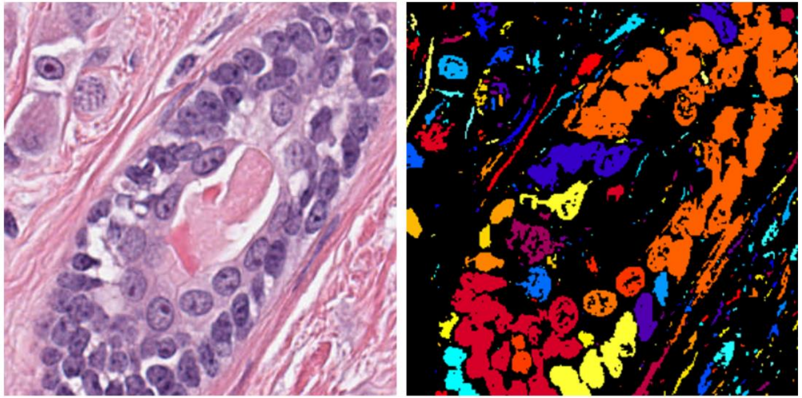

Can we build a single AI model that can segment nuclei from any H&E image regardless or organ?

Labeled Nuclei Segmentation Data

32 1000x1000 Slide Patches from 8 Different Organs

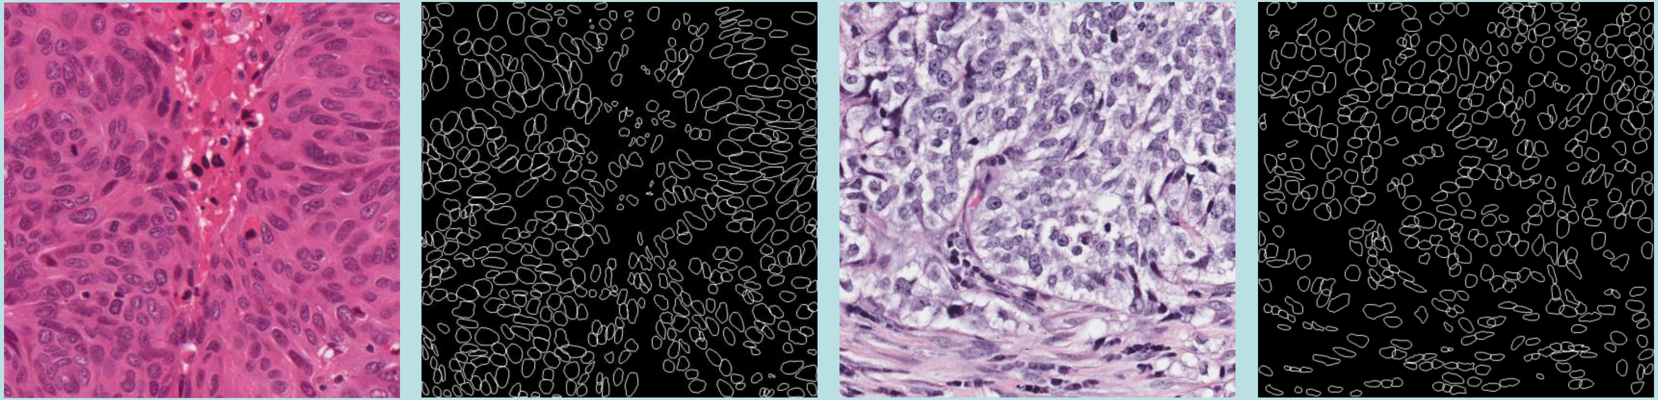

Small subjectively labeled datasets are not enough for capturing the diversity needed for a singular multi-organ nuclei segmentation network.

Sparse Stain Normalization

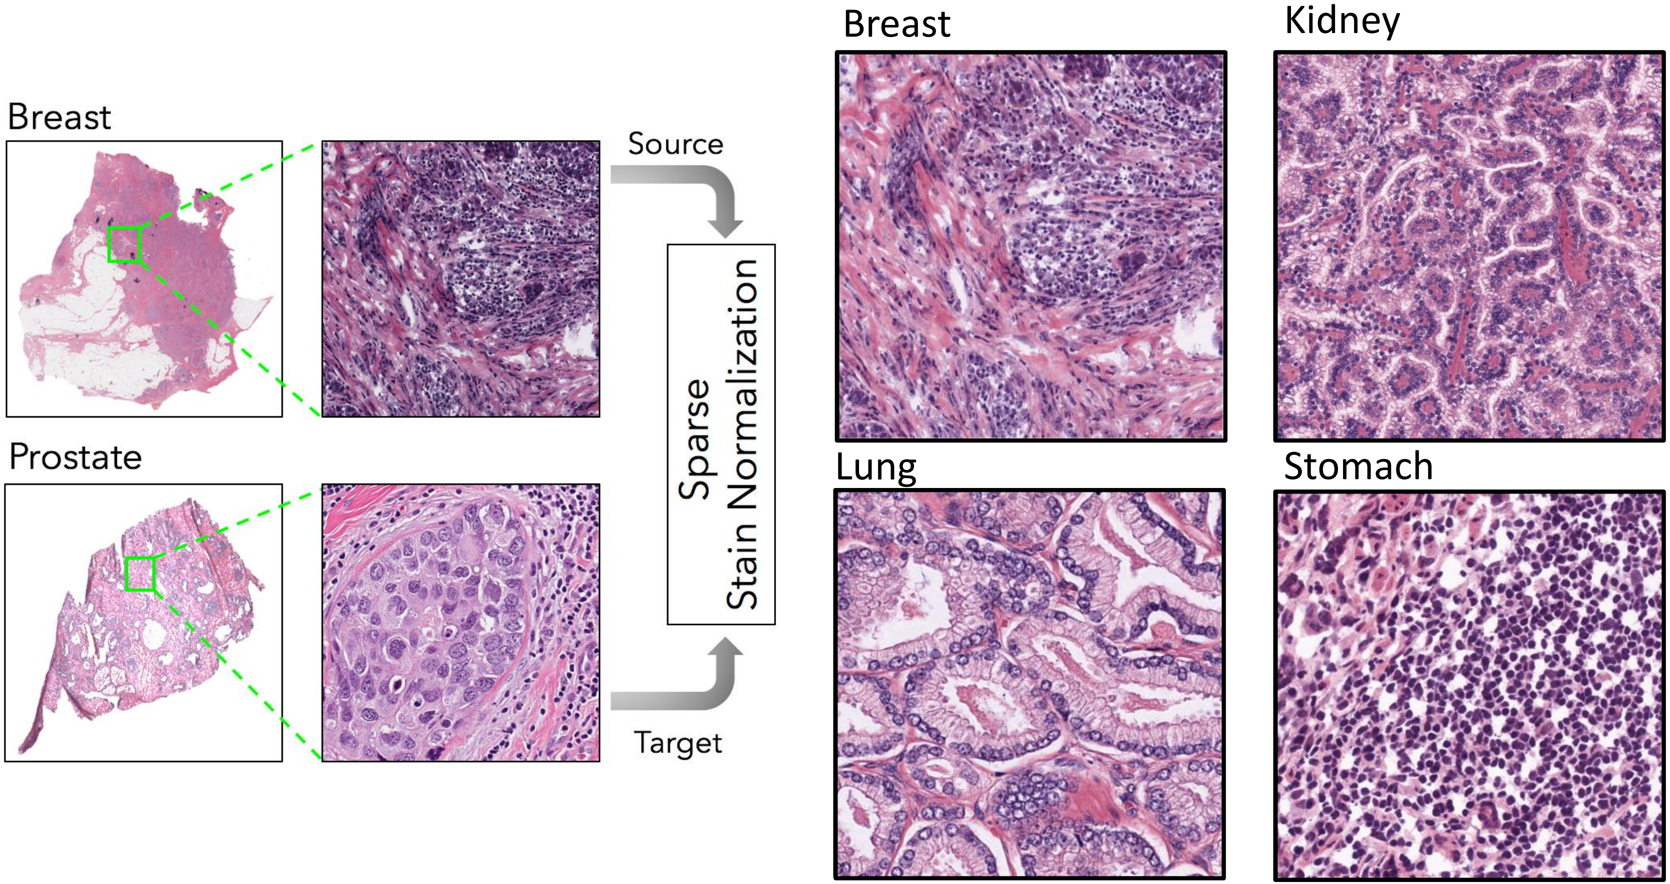

Synthetic Data Generation for Nuclei Segmentation

JOHNS HOPKINS MEDICINE

The variability in data can be captured using synthetic data.

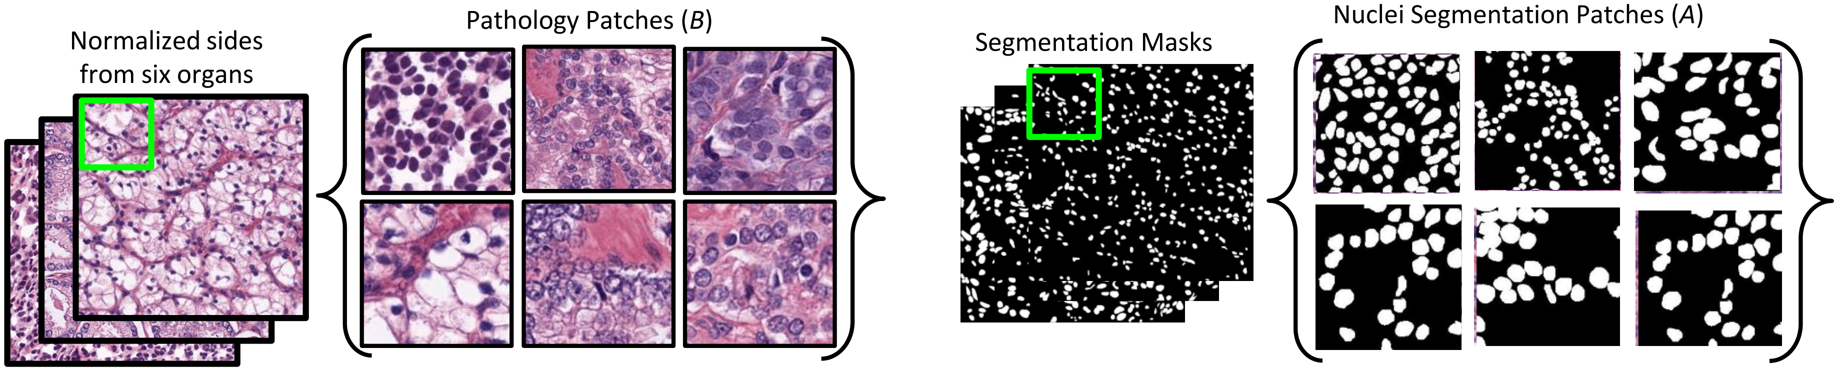

Unpaired Synthetic Data Generation

Unpaired Mapping between random polygons and synthetic H&E patches.

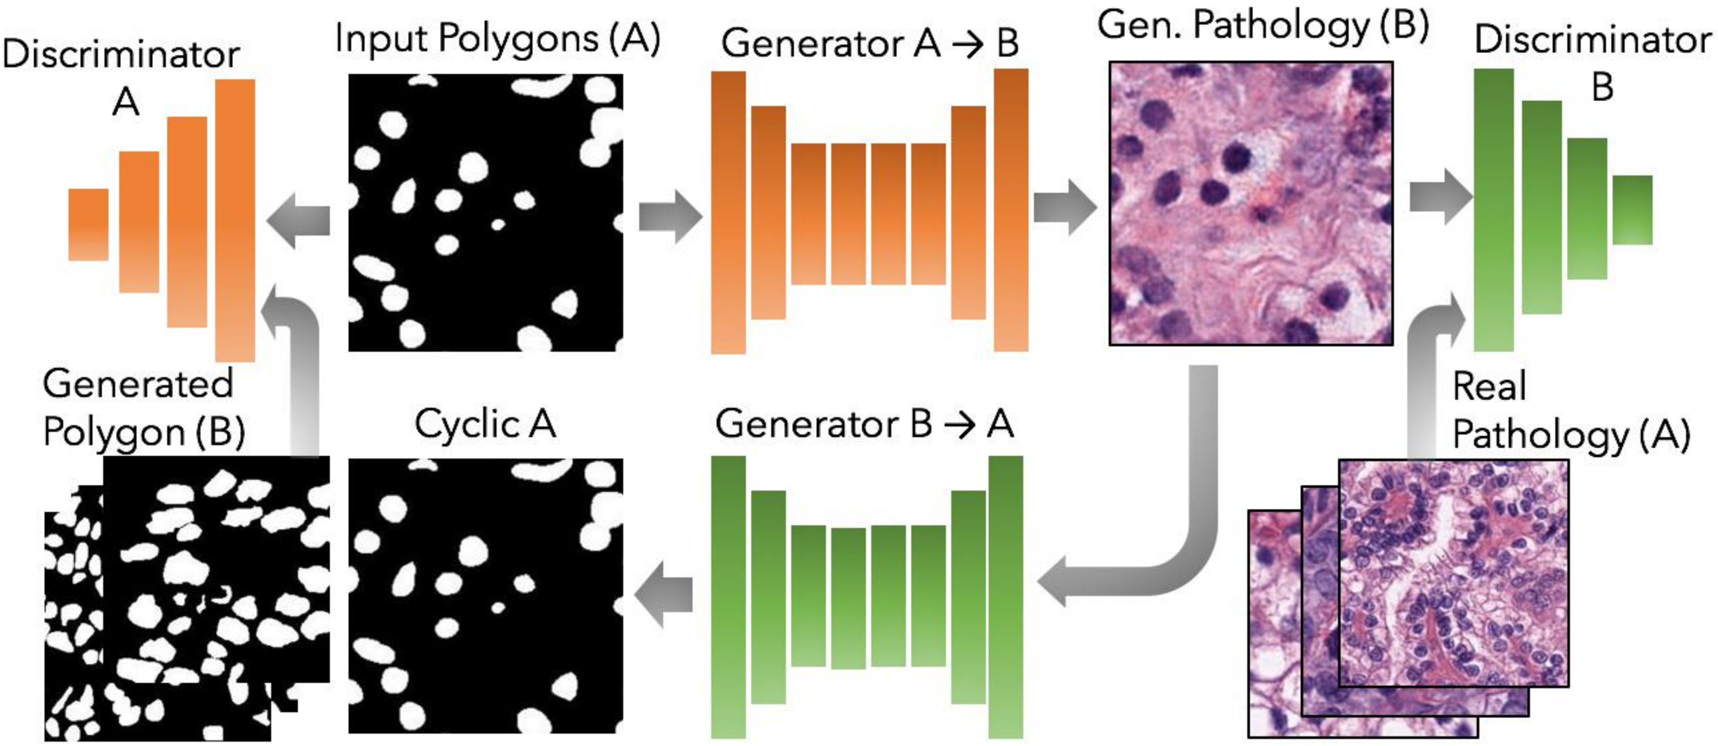

JOHNS HOPKINS
MEDICINE

Unpaired Mapping between random polygons and synthetic H&E patches.

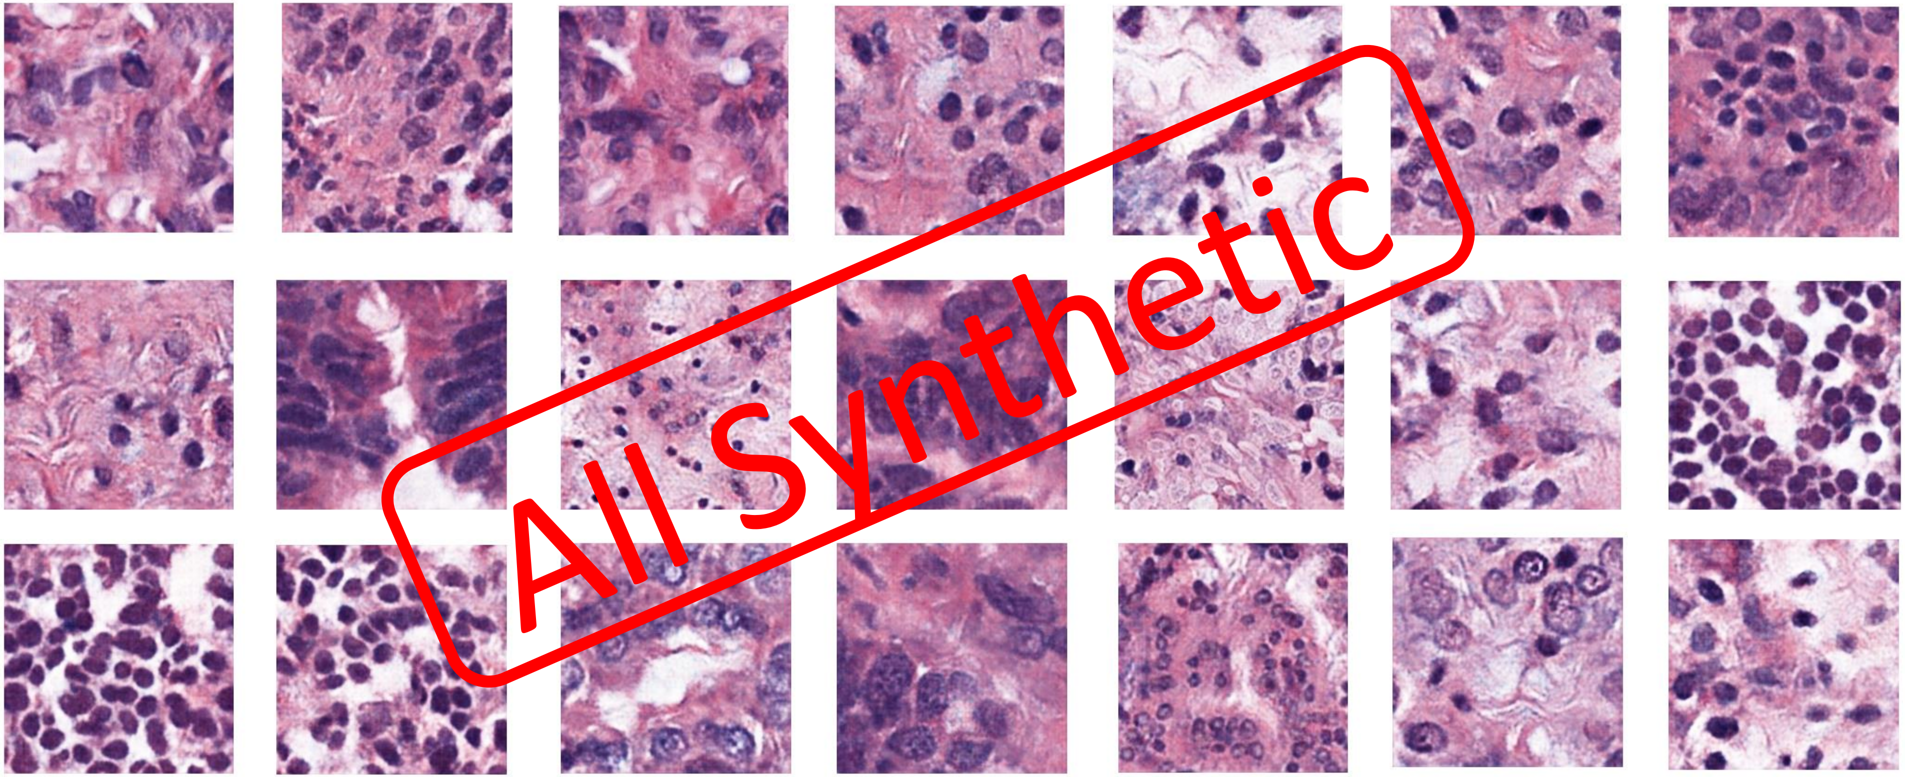

Deep Learning for Medical Imaging – Major Challenges

1. **Limited Annotated Data**
   - Under representation of rare conditions.
   - Limited experts available for annotation.
   - Privacy Issues
   - Cost
2. **Domain Adaptation**
   - Diversity in data, different sensors, cites and patients.
   - Patient specific texture and color information.
3. **Structured Prediction**
   - Global vs Local features.
4. **Incorporating Multimodal Information**
   - Subjective diagnosis is multimodal.

IEEE Transactions on Medical Imaging (2016).

Context-Aware Nuclei Segmentation with No CRF Post-processing Step

JOHNS HOPKINS MEDICINE

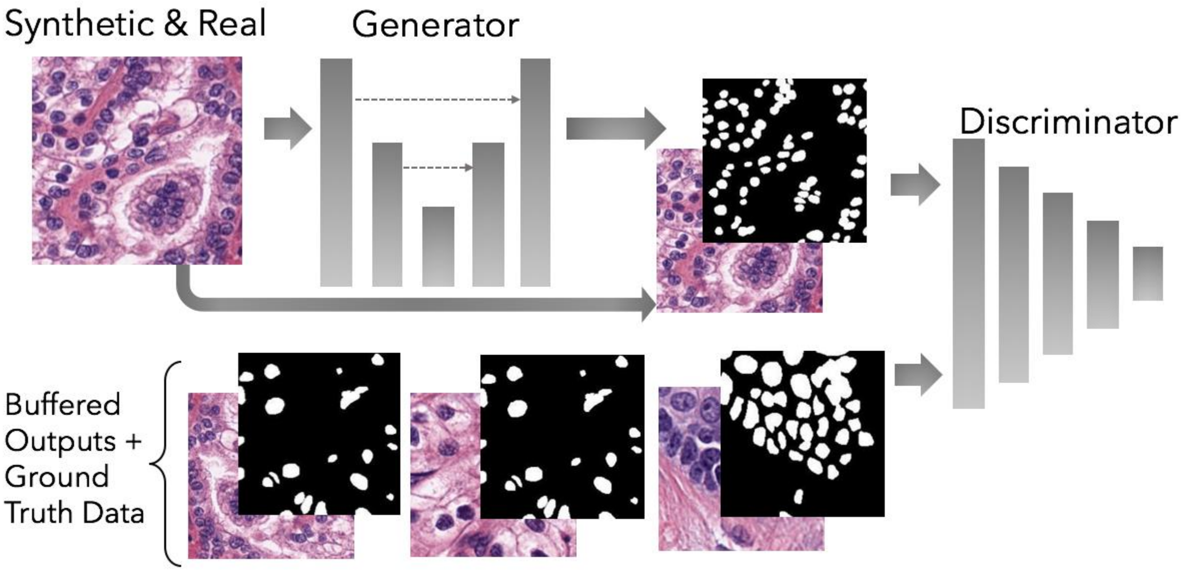

Pathology Patches with Adversarial Nuclei Segmentation

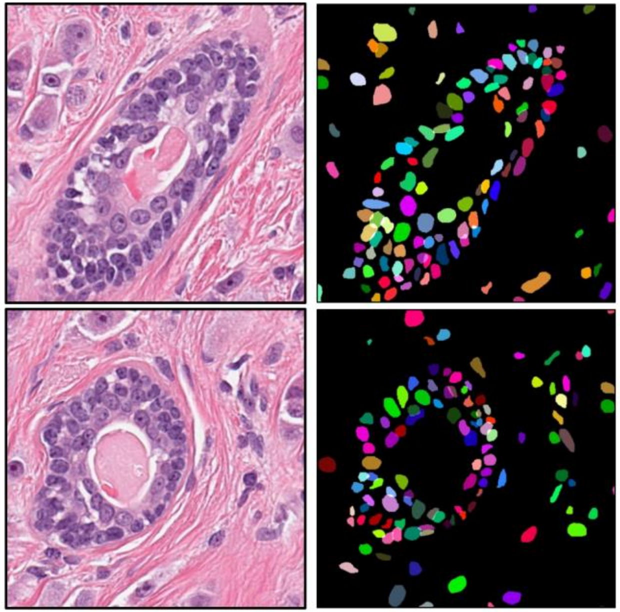

Overlapping Nuclei

Faisal Mahmood, Daniel Borders, et al. "Deep Adversarial Training for Multi-Organ Nuclei Segmentation in Histopathology Images." arXiv preprint arXiv:1810.00236 (2018).

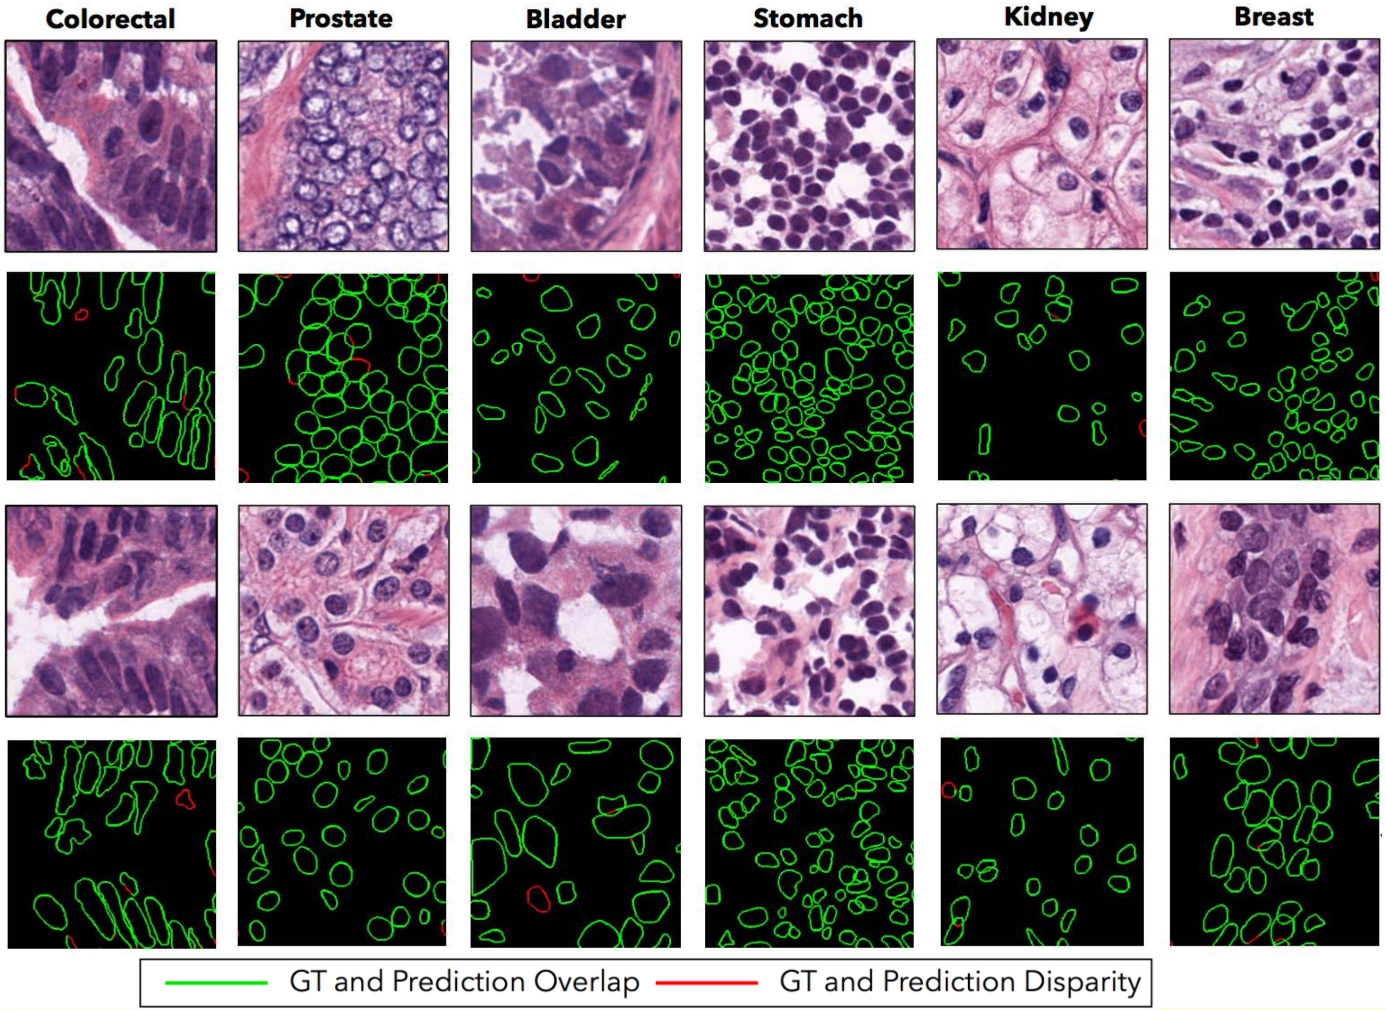

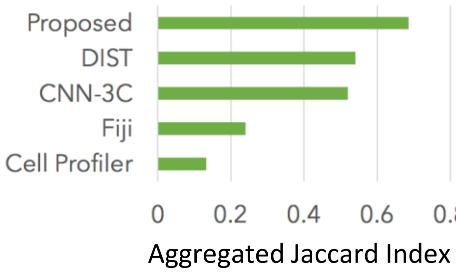

(F. Mahmood et al., 2018)

Extension to Mitotic Event Detection

Synthetic Data Generation

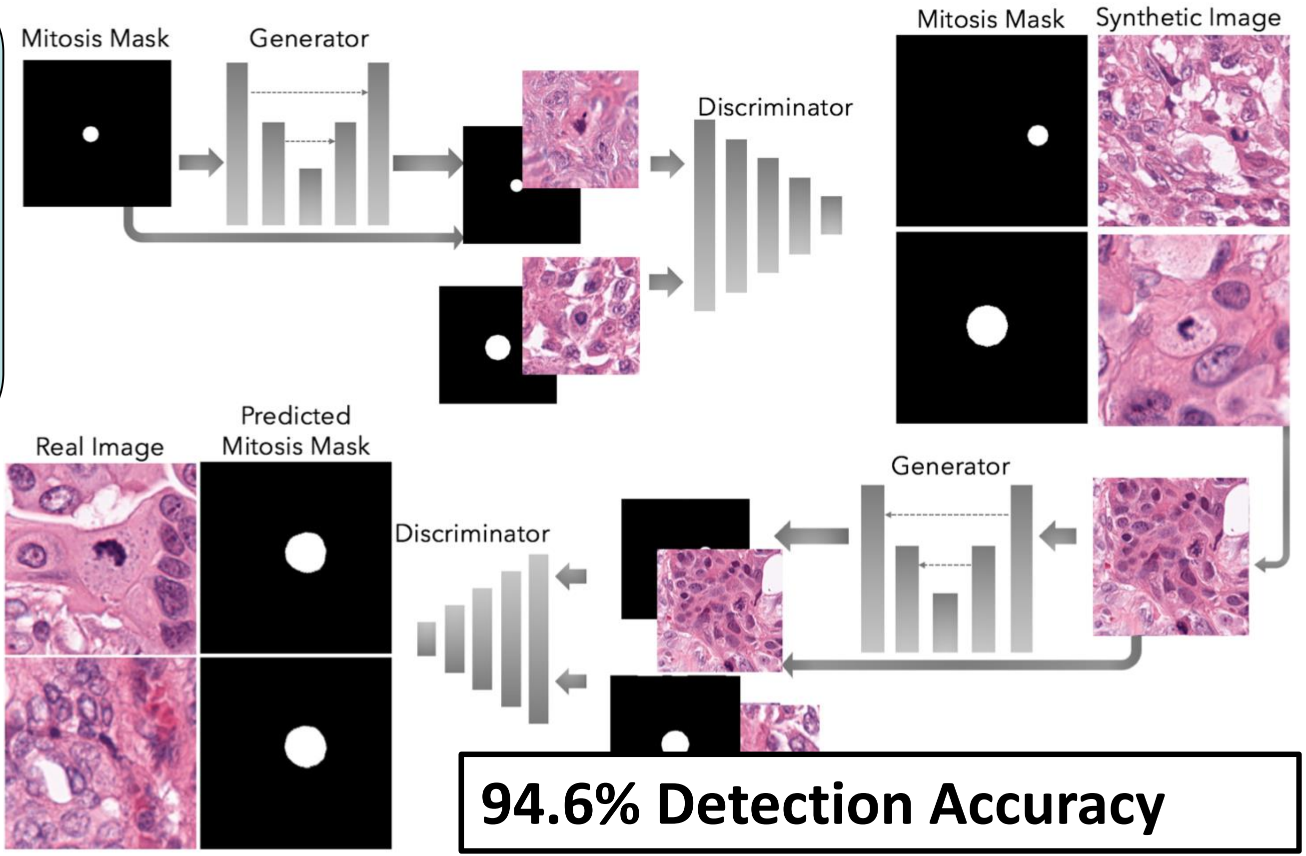

94.6% Detection Accuracy

Context-Aware Mitosis Detection

(F. Mahmood et al., 2018)

Extension to Epithelium Segmentation

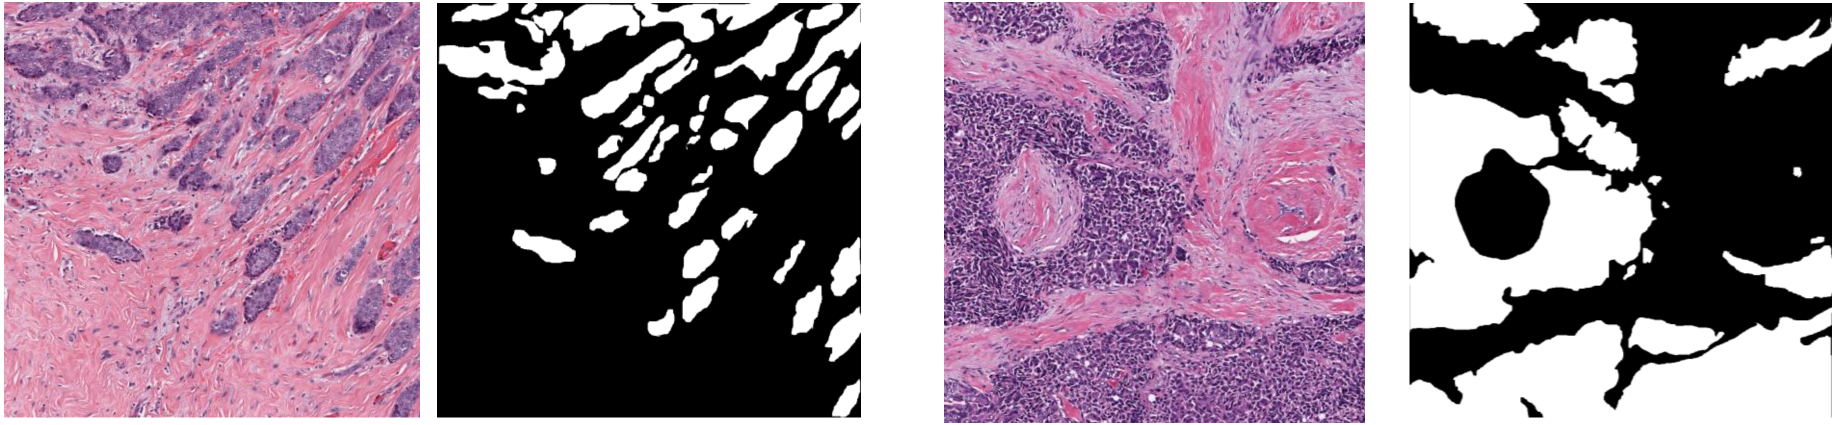

93.8% Segmentation Accuracy

(F. Mahmood et al., 2018)

Extension to Tubule Segmentation

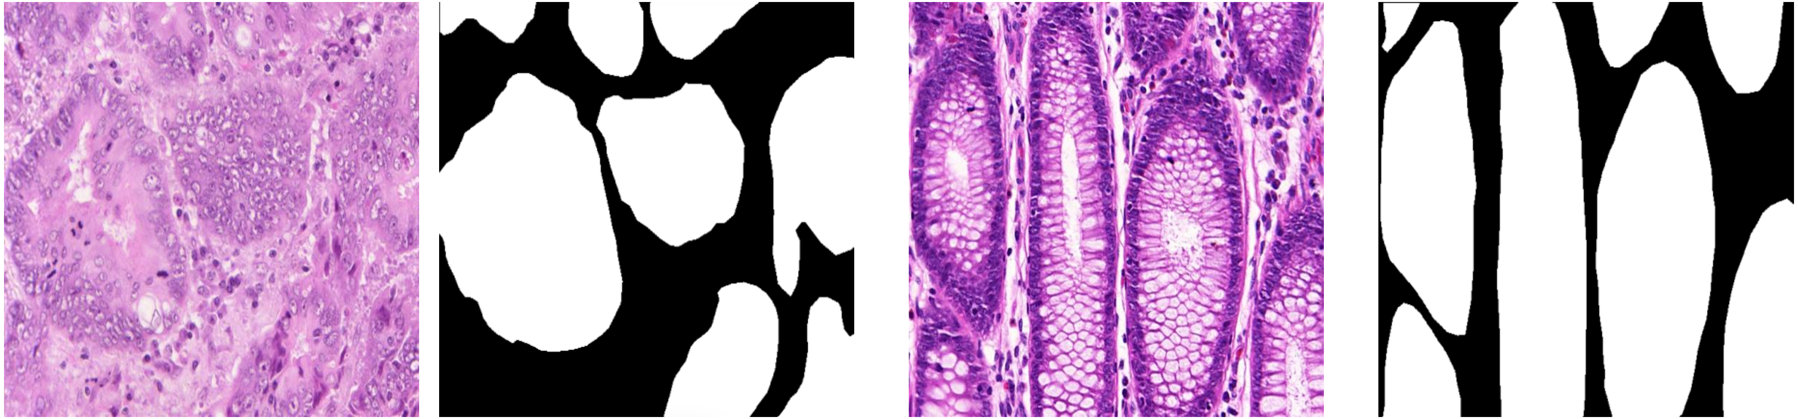

96.9% Segmentation Accuracy

(F. Mahmood et al., 2018)

Tissue Level Semantic Segmentation

JOHNS HOPKINS
MEDICINE

TMA

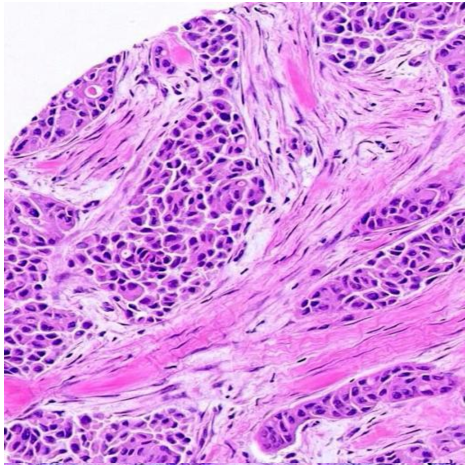

Prediction

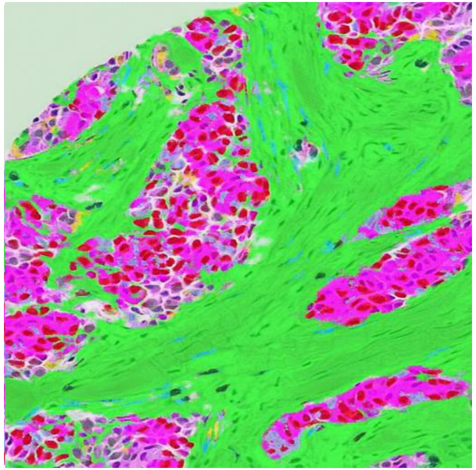

Ground Truth

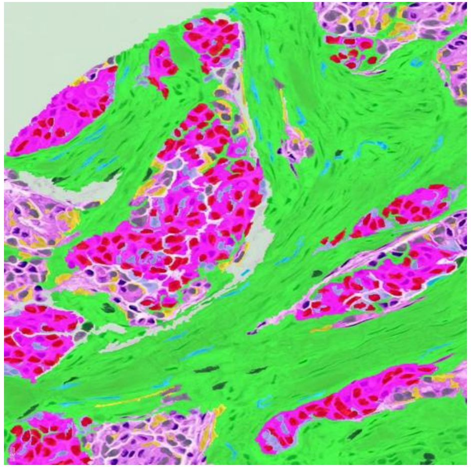

91.4% Segmentation Accuracy

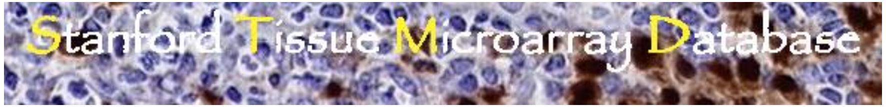

Predicted Feature Fusion

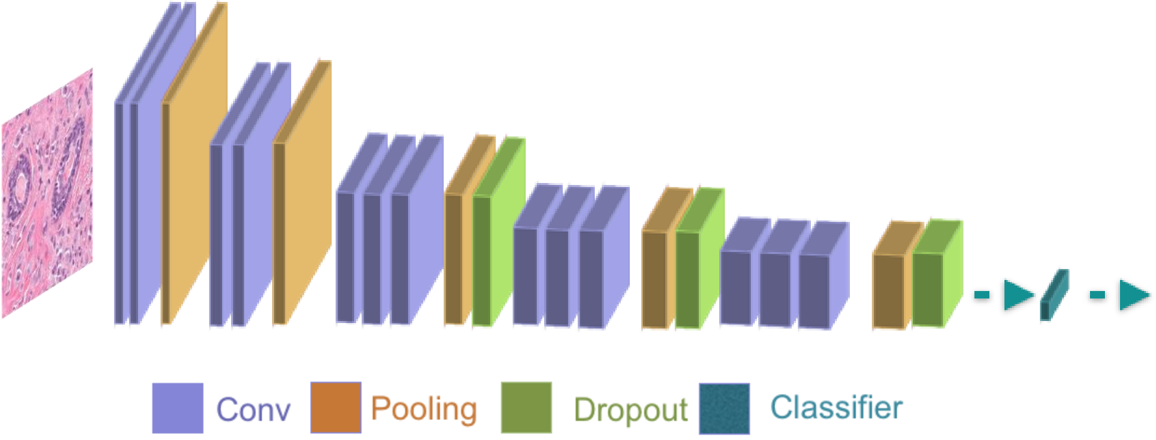

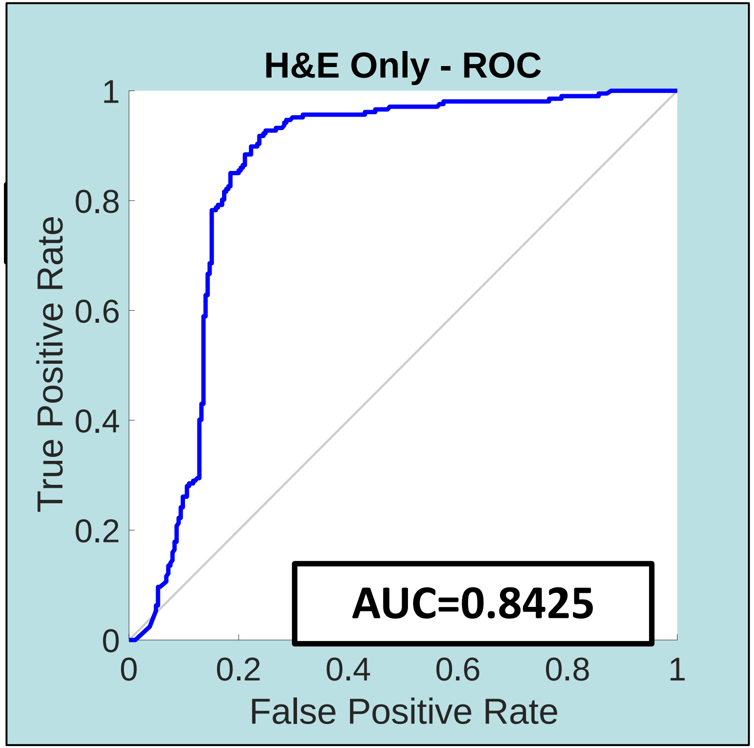

Data: 2,210 H&E Slide ROIs (TCGA).
Labeling: Labeled by 2 pathologists.

Predicted Feature Fusion

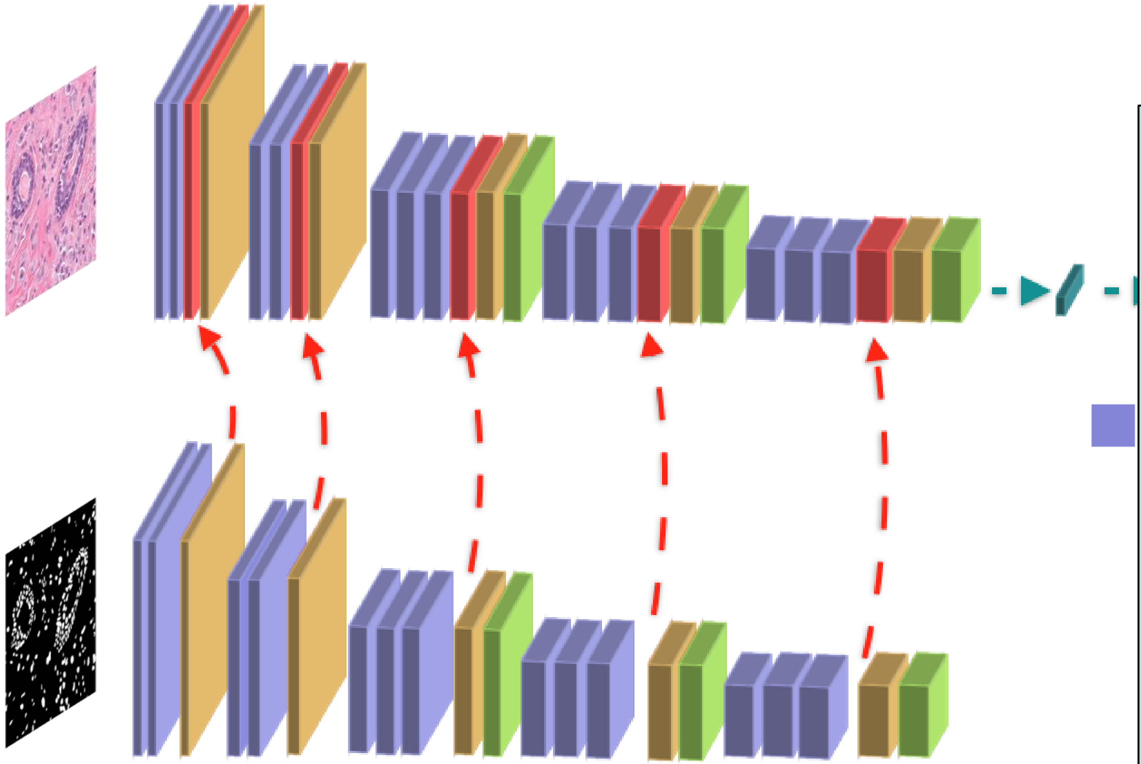

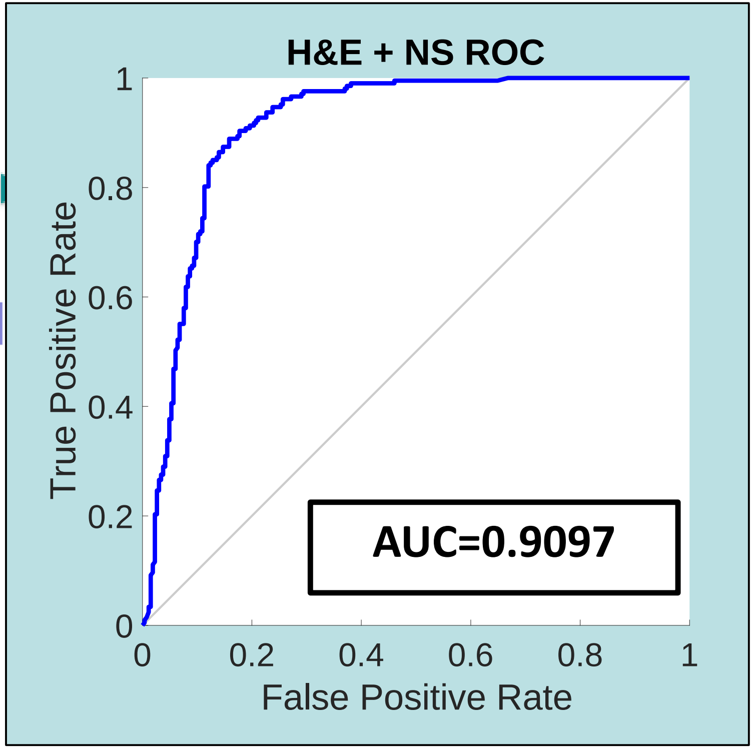

AUC Increases by 7.98% - Same Data.

Predicted Feature Fusion

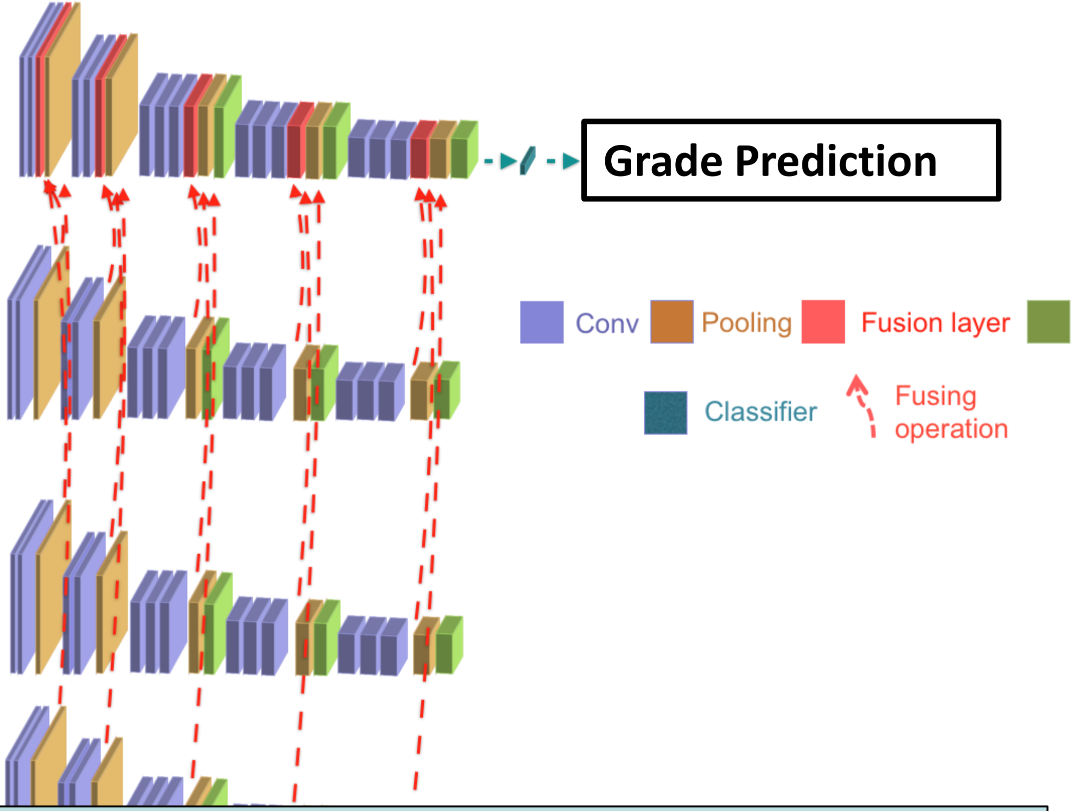

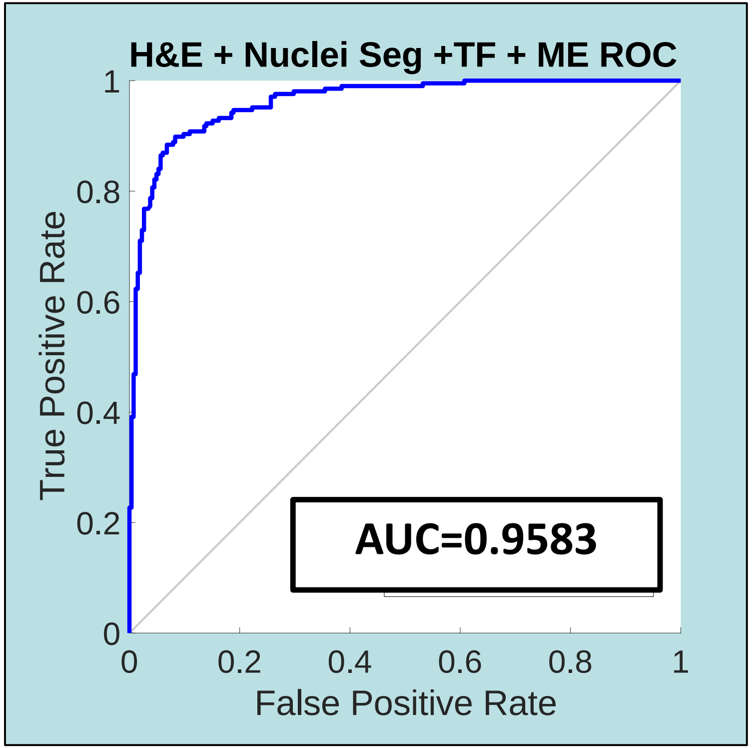

AUC Increases by 13.74% - Same Data.

Multimodal Densenet: General Framework for Multimodal Data Fusion

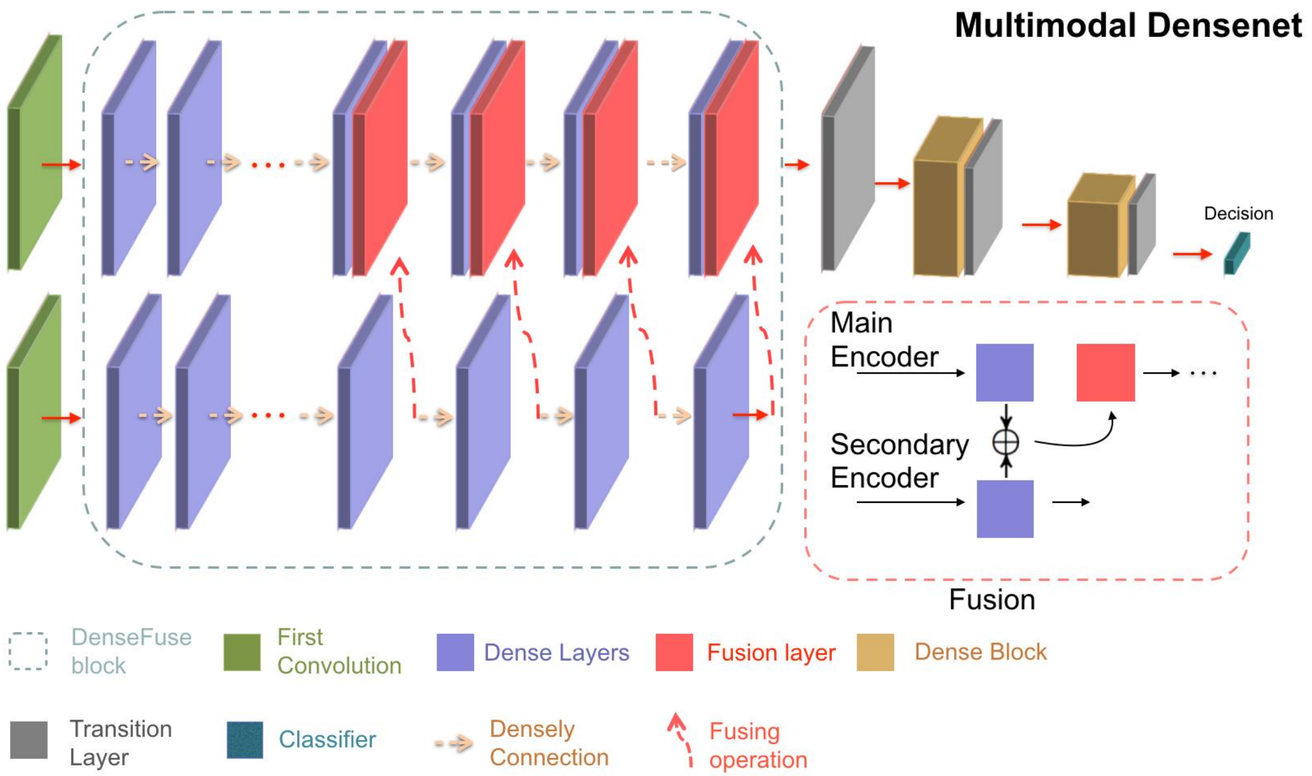

(F. Mahmood et al., 2018)

Multimodal Densenet: General Framework for Multimodal Data Fusion

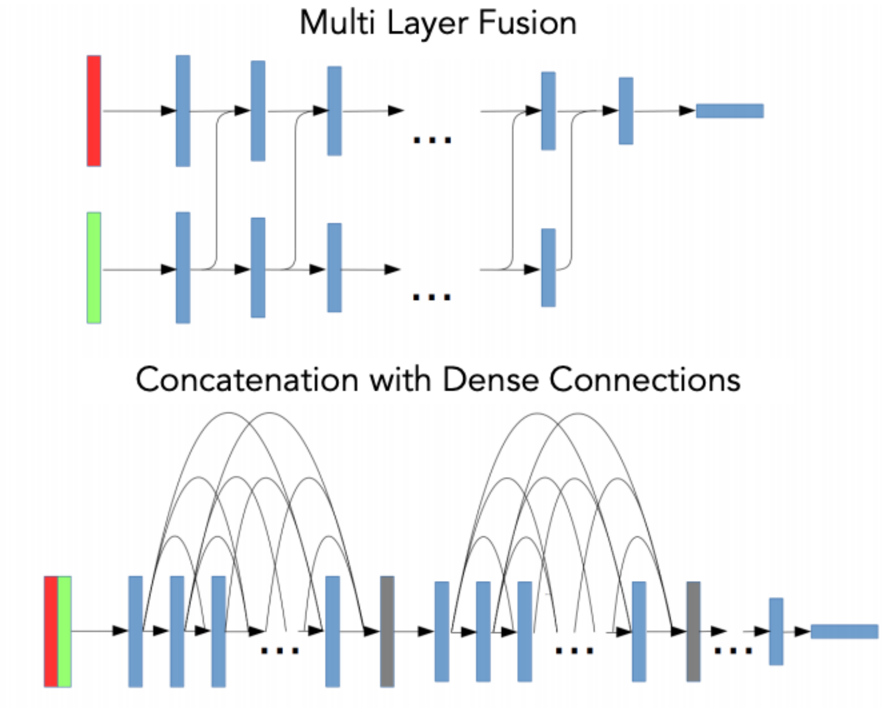

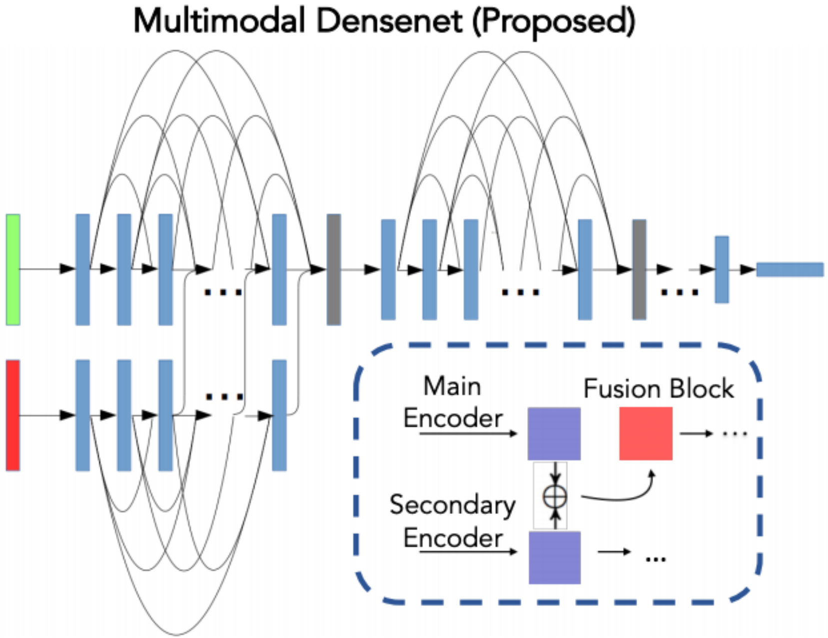

Multimodal Fusion Library

(F. Mahmood et al., 2018)

Deep Learning for Medical Imaging – Major Challenges

1. **Limited Annotated Data**
   - Under representation of rare conditions.
   - Limited experts available for annotation.
   - Privacy Issues
   - Cost
2. **Domain Adaptation**
   - Diversity in data, different sensors, cites and patients.
   - Patient specific texture and color information.
3. **Structured Prediction**
   - Global vs Local features.
4. **Incorporating Multimodal Information**
   - Subjective diagnosis is multimodal.

Future Work - Clinical

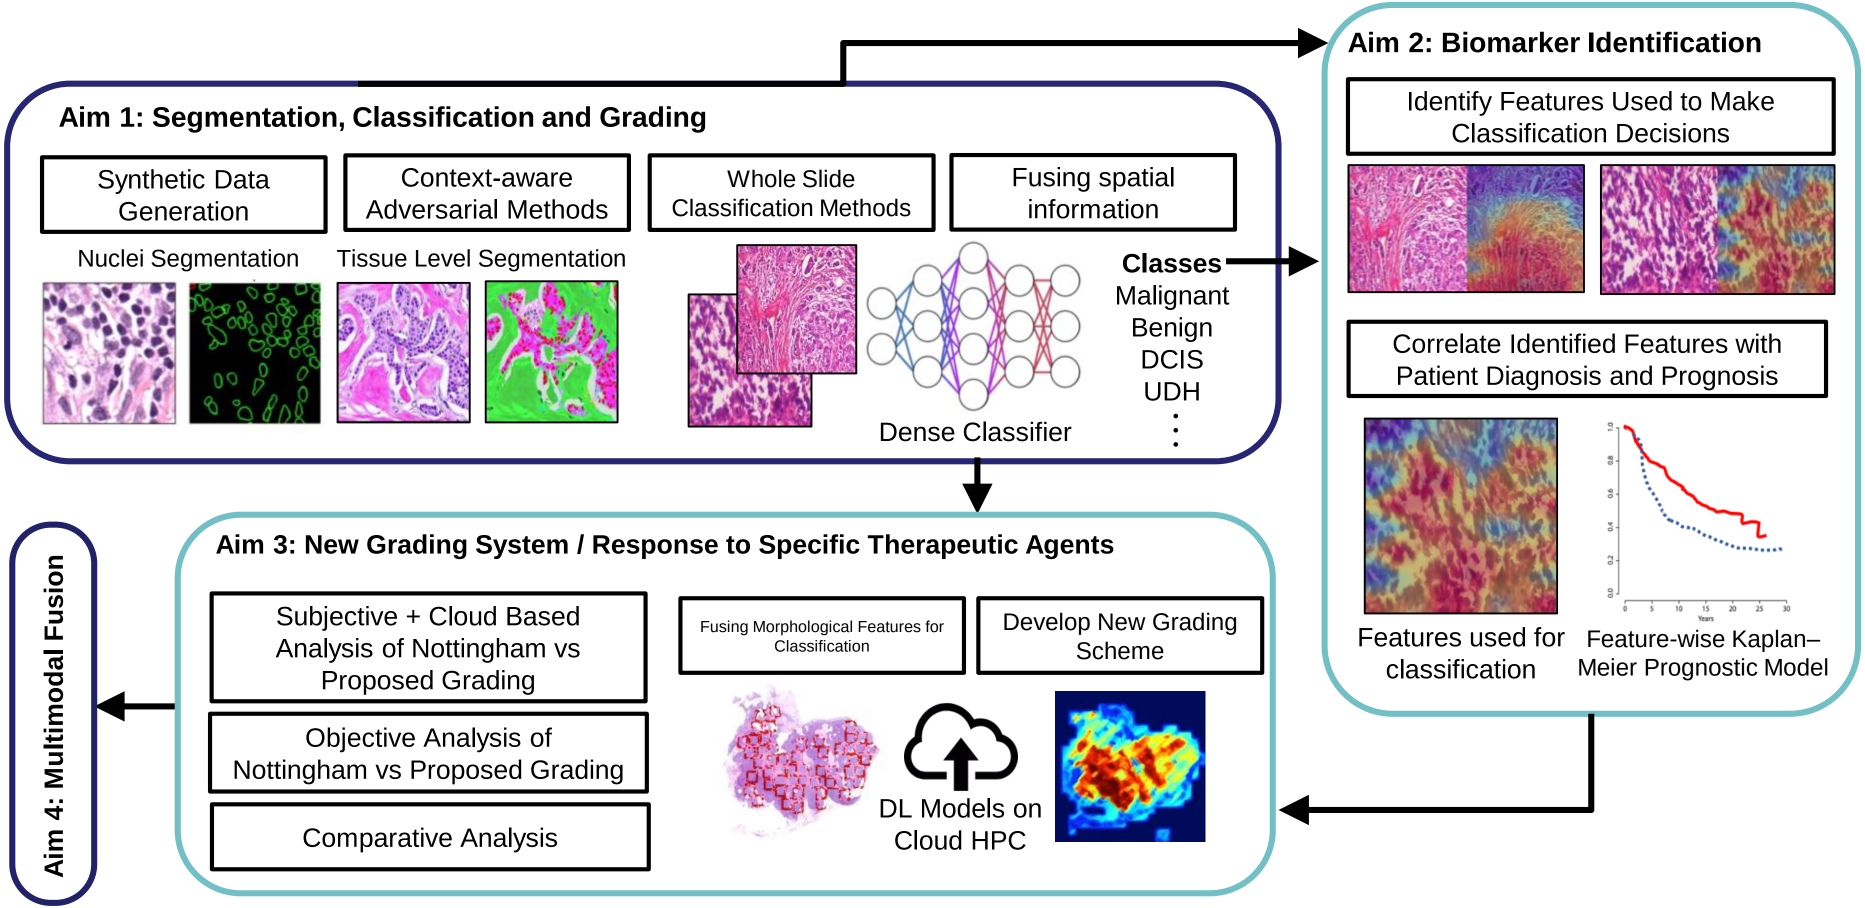

Acknowledgements

Computational Biophotonics Lab

Nicholas J. Durr
Richard Chen
Greg N. McKay
Jordan Sweer
Taylor Bobrow
Mason Chen
Bailey Surtees
Eric Chiang

Siemens

Sandra Sudarsky
Daphne Yu

JHU School of Medicine

Alexander S. Baras
Susan Harvey
Saowanee Ngamruengphong

Department of Computer Science

Alan Yuille
Greg Haiger

Olympus

Lee Zhen

Google

Jesus Trujillo Gomez

JOHNS HOPKINS BIOMEDICAL ENGINEERING
JOHNS HOPKINS MEDICINE
SIEMENS
OLYMPUS
Google
NVIDIA

Thank You.

faisalm@jhu.edu

Code / Data Available at: http://faisal.ai