MultiMix: Sparingly-Supervised, Extreme Multitask Learning From Medical Images

Ayaan Haque, Abdullah-Al-Zubaer Imran, Adam Wang, Demetri Terzopoulos

Saratoga High, Stanford University,
UCLA, VoxelCloud Inc.

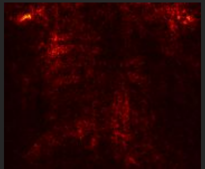

Clinical Challenges and Motivations

- Diagnosis
  - Diagnosis requires trained professionals
  - Diagnosis is expensive, subjective, time-consuming, and non-reproducible
- Segmentation
  - Annotating anatomical structures provides insights, is a challenge when done manually
  - Both of these can rely on chest x-rays
- Data in Medical Imaging
  - 90% of healthcare data is in medical images, but 97% are not analyzed
  - An automated, end-to-end systems
- Deep Learning comes into play!

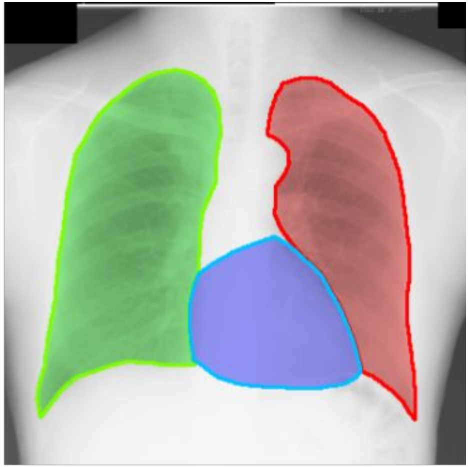

Stanford MEDICINE

Domain Generalization

- Datasets come from different domains
  - Techniques of taking medical images are different
  - Quality of images
  - Patient population
  - Statistical difference in data distribution
- Target Population Data
  - Training data not from target domain reduces effectiveness
  - Having strong performance on a totally unseen domain is of importance
    - Less models need to be created
    - Less data needs to be collected
- This has been a challenge in medical imaging, and there is still a lack of effective, domain-generalized methods

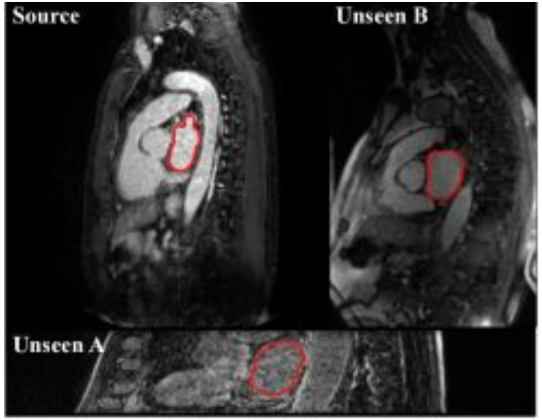

Stanford MEDICINE

Standard Approaches for DL in Medical Imaging

- CNNs for Classification
  - ResNets, SOTA Models
  - Feature Extraction
- U-Net for Segmentation
  - Encoder, Decoder
  - Downsampling feature extractor with Convolutions
  - Upsampling with skip-connections
    - Reuses features from previous layers

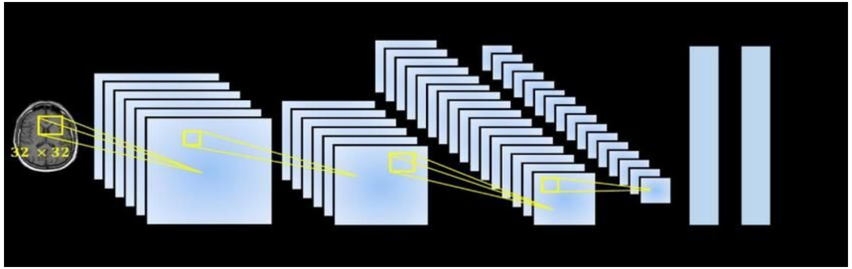

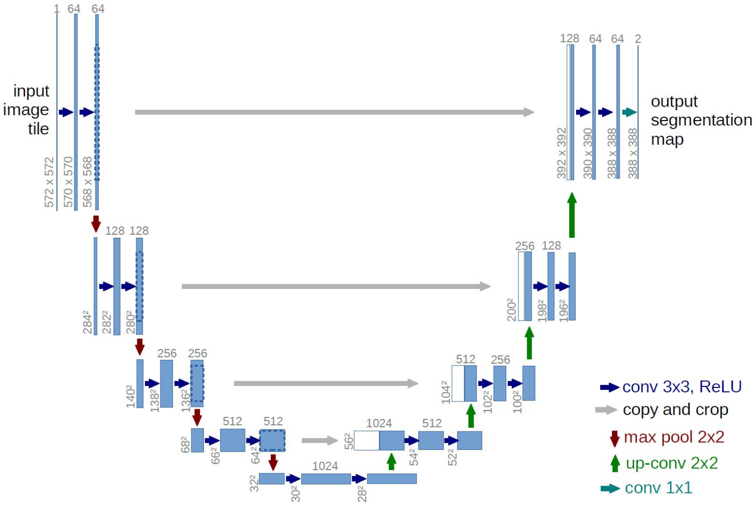

Stanford MEDICINE

Semi-Supervised Learning

- Partially labeled datasets
  - Many datasets have lots of unlabeled data, as obtaining labeled data is difficult
  - In the medical domain, this is a large challenge, labeling is expensive and time-consuming
- SSL maximizes knowledge gains by leveraging labeled data in tandem with unlabeled data
  - Lots of recent research on SSL methods
  - The main premise is that labeled data and unlabeled data are similar in their characteristics and distribution

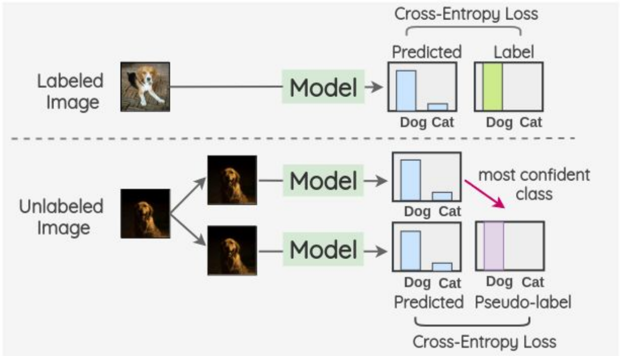

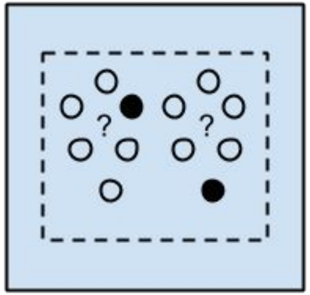

Popular SSL Methods

- Pseudo-Labeling (Lee 2013)
  - The model itself creates labels for the unlabeled data points
  - Method of Self-training
- Entropy Minimization (Grandvalet and Bengio 2005)
  - Trained to match the predictions of the unlabeled data to the labeled data
- Domain Adaptation (Beijbom 2012)
  - Pre-train a model on separate task in different domain, transfer learning
  - Similar to humans, learning one domain helps in learning others

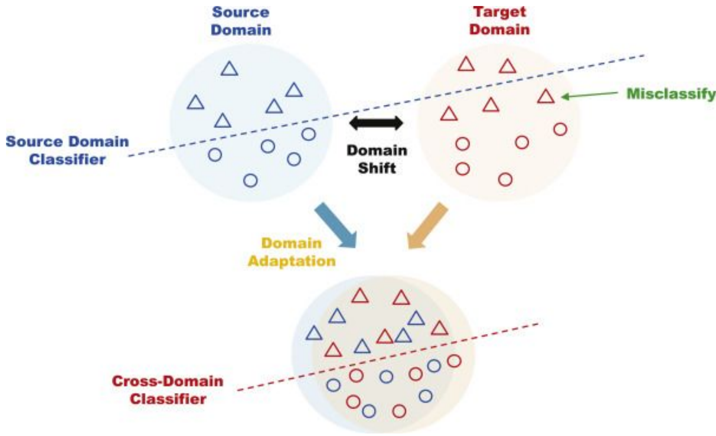

Stanford MEDICINE

Multi-Task Learning

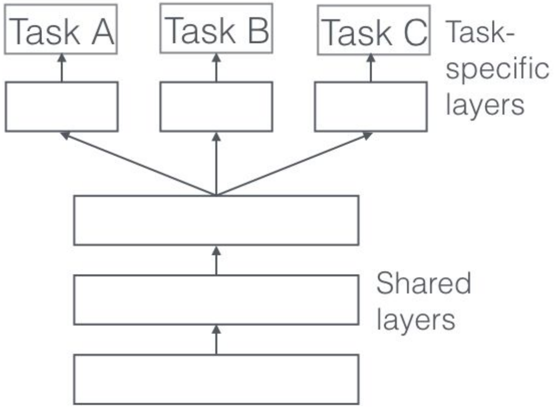

- Achieving two tasks within a single model (Ruder 2017)
  - Parameter Sharing (Soft and Hard)
  - Output a single vector with multiple dimensions
  - Multiple loss objectives are combined
- Improves generalization, accuracy, and optimizes both tasks
- Empirical evidence for its success
- Tasks must be similar or parameter sharing is inadequate
- Examples
  - Autonomous Driving

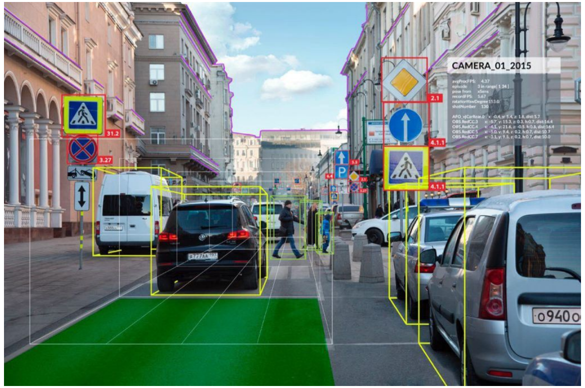

Related Work

- Single Task
  - U-Net (Ronneberger et al. 2015)
  - Myronenko 2018, VAE's for MRI segmentation
  - Many more
- Multi-Task
  - Y-Net (Mehta et al. 2018)
  - Joint Learning for Segmentation (Girard et al. 2019)
- Semi-Supervised, Multi-Task Learning
  - APPAU-Net (Imran and Terzopoulos 2019)
  - SSL-MTL Collaborative Learning (Zhou et al. 2019)
  - Very few, only in recent research in the last few years

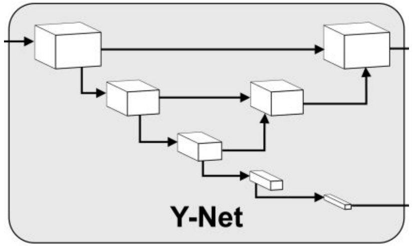

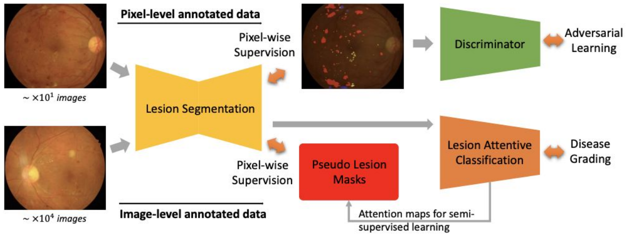

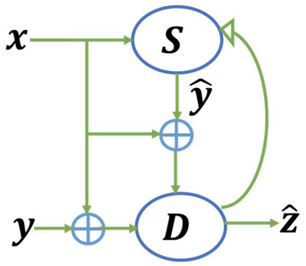

Stanford MEDICINE

Challenges

- Generalization between different domains
- Use of multiple datasets from different sources and domains
- Sparingly Supervised Learning
  - Extremely low amount of labels
- Multi-Task Learning
  - Specifically, bridging the two tasks into one

Key Contributions

A novel semi-supervised method exploiting consistency augmentation for jointly learning diagnostic classification and anatomical structure segmentation from multi-source, multi-domain medical image datasets.

Incorporation of an innovative saliency bridge module connecting the two related tasks leading to improved performances of either tasks, within the same model.

Extensive experimentation with varied quantities of labeled data and mixed sources for multiple tasks demonstrating better generalizability (both in- and cross-domain) of the proposed model.

Stanford MEDICINE

Proposed Model: MultiMix

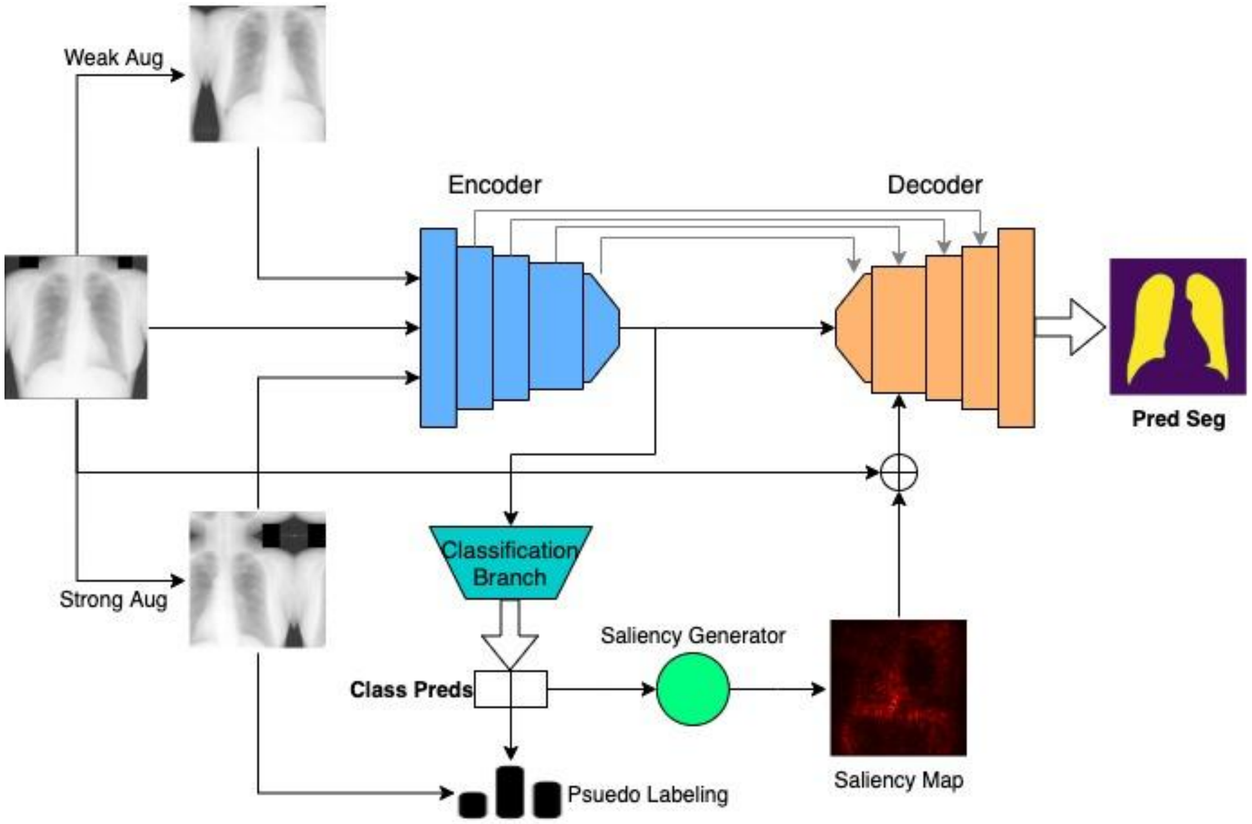

Consistency and Confidence Learning

- Augmentation based classification
  - Two augmentations on a single unlabeled image
    - Weak and Strong Augmentations
    - Random Augmentations
    - Labeled Images are weakly augmented
- Pseudo-Labeling
  - Predictions of the weakly augmented image are converted into pseudo-labels
  - Weakly augmented pseudo-label and strongly augmented image predictions are computed in the loss
  - Thresholding for Confidence

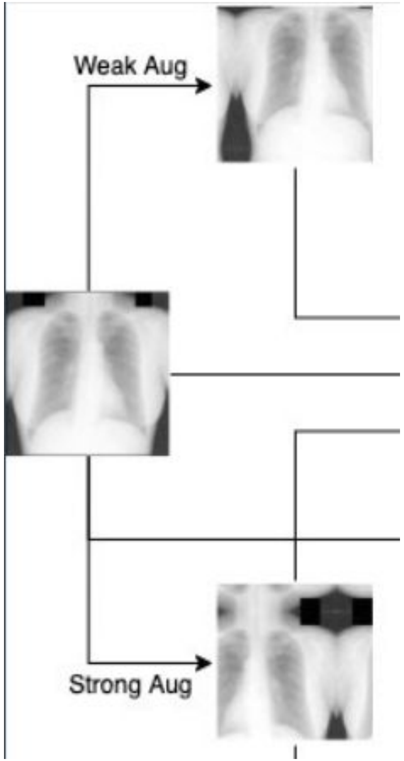

Stanford MEDICINE

Saliency Bridge Module

- Generation of Saliency Maps
  - Based on classification predictions, not segmentation
  - Computing gradients of predictions, creates a “heat map” or outline
  - Can be visualized
- Concatenation
  - Concatenated with input images
  - Input image provides additional context
- Bridges the two tasks, explainable and interconnected multi-tasking

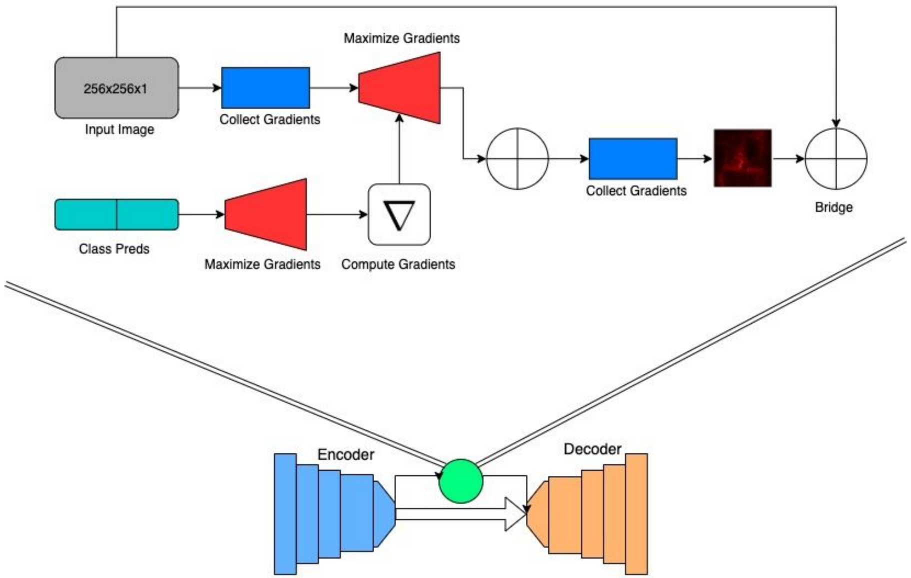

Loss Functions (Classification)

Supervised

$$L_l^c = \frac{1}{m}\sum_{i=1}^{m} L_c(\hat{c}_l^m, c_l^m)$$

Unsupervised

$$L_u^c = \lambda\frac{1}{m}\sum_{i=1}^{m} L_c(\hat{c}_s^m, arg\max(\hat{c}_w^m) \geq t)$$

- Supervised Loss
  - Standard classification loss
  - cross-entropy loss
- Unsupervised Loss
  - Weighted loss
  - Weak and Strong Predictions
  - Thresholding

Loss Functions (Segmentation)

- Supervised
  - Dice loss
  - Alpha is segmentation weight
- Unsupervised Loss
  - KL Divergence
    - based on distribution of the masks
  - Beta is the unsupervised weight

Supervised

$$L_l^s = \alpha \frac{1}{m} \sum_{i=1}^{m} L_s(\hat{y}_l^m, y_l^m)$$

Unsupervised

$$L_u^s = \beta \frac{1}{m} \sum_{i=1}^{m} L_{KL}(\hat{y}_l^m, \hat{y}_u^m)$$

Stanford MEDICINE

Algorithm

Algorithm 1 MultiMix Mini-Batch Training

Require:

Training set of labeled data $x_l^c, y_l^c \in \mathcal{D}_l^c$
Training set of labeled data $x_l^s, y_l^s \in \mathcal{D}_l^s$
Training set of unlabeled data $x_w^c \in \mathcal{D}_u^c$
Training set of unlabeled data $x_s^c \in \mathcal{D}_u^c$, where x_w^c and x_s^c augmented at different strengths
Training set of unlabeled inputs $x_u^s \in \mathcal{D}_u^s$
Network architecture $\mathcal{F}_\theta$ with learnable parameters θ

for each step **do**

Sample minibatch $x_l^c(i); x_l^c(1), \dots, x_l^c(m) \sim p_{\mathcal{D}^c(x)}$
Sample minibatch $x_l^s(i); x_l^s(1), \dots, x_l^s(m) \sim p_{\mathcal{D}^s(x)}$
Sample minibatch $x_w^c(i); x_w^c(1), \dots, x_w^c(m) \sim p_{\mathcal{D}_w^c(x)}$ $x_s^c(i); x_s^c(1), \dots, x_s^c(m) \sim p_{\mathcal{D}_s^c(x)}$
Sample minibatch $x_u^s(i); x_u^s(1), \dots, x_u^s(m) \sim p_{\mathcal{D}^s(x)}$

Compute model outputs for the labeled data: $\hat{c}_l, \hat{y}_l \leftarrow \mathcal{F}_\theta$
Compute model outputs for the unlabeled data: $\hat{c}_w, \hat{c}_s, \hat{y}_u \leftarrow \mathcal{F}_\theta$

Compute psuedo-label for weakly augmented classification predictions: $arg\max(\hat{c}_w) \geq t \leftarrow c_u$

Update $\mathcal{F}_\theta$ along its gradient:

$$\nabla_{\theta_{\mathcal{F}}} \frac{1}{|\mathcal{M}_l|} \sum_{i \in \mathcal{M}_\mathcal{L}} \left[L_{(\hat{c}_l, \hat{y}_l, y_l^c, y_l^s)} \right] + \alpha \frac{1}{|\mathcal{M}_u|} \sum_{i \in \mathcal{M}_u} \left[L_{(\hat{c}_s, \hat{y}_u, c_u, \hat{y}_l)} \right]$$

end for

Require

- Labeled subsets of classification and segmentation images
- Unlabeled subset of segmentation images
- Weakly and Strongly augmented versions of unlabeled images
- Model with parameters

During Each Step

- Compute outputs for all subsets
- Compute pseudo-label for weakly augmented preds
- Update the model along its gradient

Stanford MEDICINE

Datasets

Table 3: Details of the datasets used for training and testing.

mode	Dataset	Total	Normal	Abnormal	Train	Val	Test
in-domain	JSRT	247	–	–	111	13	123
	CheX	5,856	1583	4273	5216	16	624
cross-domain	MCU	138	–	–	93	10	35
	NIHX	4185	2754	1431	–	–	4185

In-Domain

- JSRT
  - Lung Segmentation
  - 247 images, 111 for training and 123 for testing
- CheX
  - Pneumonia Classification
  - 5800 images, 5126 for training
- Cross-Domain
  - MCU
    - Lung Segmentation
    - 138 images, 35 for testing
  - NIHX
    - Pneumonia classification
    - Subset of NIH Chest X-Ray Dataset

Stanford MEDICINE

Datasets

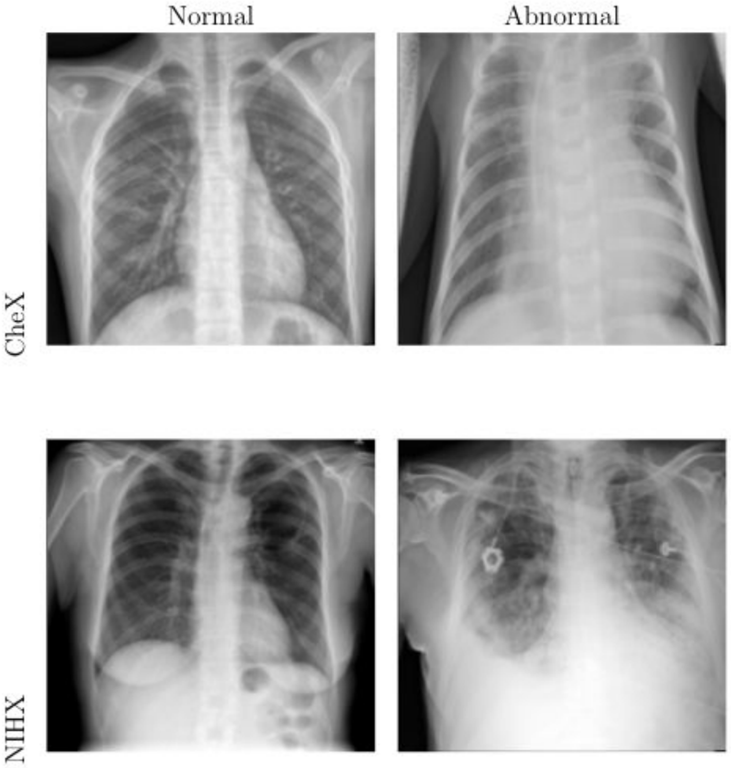

- Significant Difference in data distribution
- Visual Difference in Images
- Makes multi-source training difficult and high performance harder

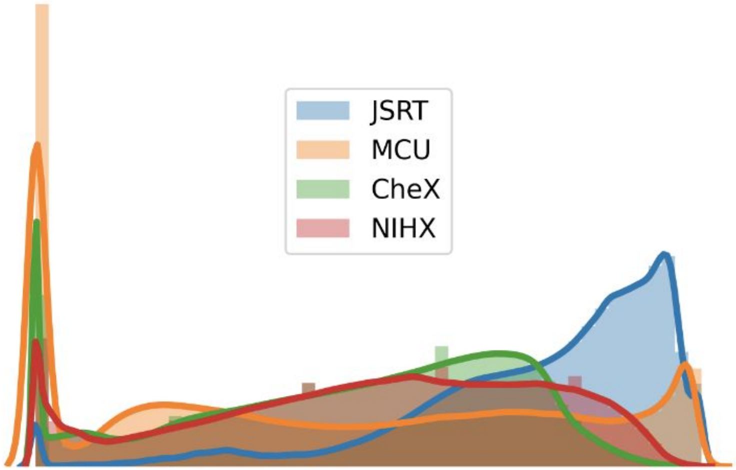

Stanford MEDICINE

Baseline Models

- U-Net
  - Baseline segmentation model, fully-supervised, single-task learning
- U-Net Encoder
  - Baseline for classification, fully supervised, single-task learning
  - Essentially a standard CNN, just the encoder portion of the U-Net
- U-MTL
  - Baseline multi-tasking model
  - U-Net encoder decoder structure with a classification branch

Implementation Details

- Inputs and Frameworks
  - Inputs normalized and resized 256x256x1
  - Coded in Pytorch (Python) and used a NVIDIA K80 GPU
- Training:
  - Varying training labels
    - seg (10, 50, full)
    - class (100, 1000, full)
  - Experiments repeated 5 times
- Hyperparameters:
  - Adam Optimizer, adaptive learning rates of 0.1 every 8 epochs, initial of 0.0001
  - LReLU negative slope = 0.2, dropout = 0.25
  - $t = 0.7$, $\lambda = 0.25$, $\alpha = 5.0$, $\beta = 0.01$, $M = 10$

Stanford MEDICINE

Results (Overview)

- Models
  - U-Net
  - U-Net Encoder
  - U-Net Encoder SSL
  - UMTL
  - UMTL-S
  - UMTL-SSL
  - UMTL-SSL-S
  - MultiMix
- In-Domain and Cross-Domain
  - CheX, JSRT
  - NIHX, MCU
- Metrics
  - Classification:
    - Acc, F1-Nor, F1-Abn
  - Segmentation:
    - DS, JS, SSIM, HD, P, R
- Qualitative Results
  - Box Plots
  - Boundary/Mask Visualizations
  - Classification and ROC Curves

Stanford MEDICINE

Results (In-Domain Evaluation)

Model	$\lvert\mathcal{D}_l^c\rvert$	Classification Acc	F1-Nor	F1-Abn	$\lvert\mathcal{D}_l^s\rvert$	Segmentation DS	JS	SSIM	HD	P	R
U-Net	—	—	—	—	10	0.634	0.695	0.810	2.899	0.779	0.865
	—	—	—	—	50	0.855	0.854	0.904	0.341	0.918	0.925
	—	—	—	—	Full	0.915	0.906	0.929	0.104	0.949	0.953
Enc	100	0.732	0.424	0.806	—	—	—	—	—	—	—
	1000	0.773	0.546	0.842	—	—	—	—	—	—	—
	Full	0.737	0.534	0.838	—	—	—	—	—	—	—
Enc-SSL	100	0.780	0.570	0.844	—	—	—	—	—	—	—
	1000	0.822	0.692	0.876	—	—	—	—	—	—	—
	Full	0.817	0.680	0.872	—	—	—	—	—	—	—
UMTL	100	0.707	0.443	0.797	10	0.626	0.871	0.908	4.323	0.900	0.964
	100	0.655	0.683	0.853	50	0.647	0.854	0.881	4.733	0.864	0.989
	100	0.706	0.416	0.804	Full	0.696	0.872	0.911	3.908	0.892	0.986
	1000	0.750	0.490	0.825	10	0.761	0.904	0.926	3.050	0.924	0.977
	1000	0.749	0.510	0.833	50	0.768	0.927	0.938	2.606	0.940	0.985
	1000	0.747	0.530	0.840	Full	0.759	0.928	0.930	2.955	0.924	0.981
	Full	0.744	0.515	0.828	10	0.909	0.919	0.521	0.903	0.912	0.962
	Full	0.738	0.438	0.820	50	0.930	0.948	0.954	0.444	0.969	0.977
	Full	0.731	0.447	0.822	Full	0.932	0.951	0.957	0.372	0.965	0.977
UMTL-S	100	0.704	0.358	0.806	10	0.922	0.848	0.891	4.005	0.871	0.966
	100	0.701	0.336	0.796	50	0.926	0.867	0.894	4.393	0.873	0.891
	100	0.713	0.442	0.794	Full	0.931	0.890	0.920	3.983	0.906	0.980
	1000	0.740	0.482	0.828	10	0.948	0.908	0.924	2.546	0.931	0.972
	1000	0.771	0.566	0.844	50	0.965	0.931	0.941	2.083	0.949	0.981
	1000	0.742	0.497	0.830	Full	0.962	0.925	0.935	1.758	0.958	0.985
	Full	0.747	0.500	0.830	10	0.955	0.914	0.936	0.568	0.954	0.956
	Full	0.737	0.433	0.820	50	0.972	0.944	0.953	0.560	0.966	0.977
	Full	0.723	0.413	0.817	Full	0.974	0.953	0.957	0.539	0.967	0.981
UMTL-SSL	100	0.790	0.618	0.856	10	0.906	0.925	0.940	0.626	0.954	0.953
	100	0.818	0.688	0.872	50	0.919	0.946	0.952	0.561	0.962	0.963
	100	0.852	0.670	0.868	Full	0.937	0.954	0.958	0.613	0.969	0.981
	1000	0.794	0.630	0.860	10	0.893	0.926	0.941	0.524	0.961	0.962
	1000	0.822	0.693	0.877	50	0.903	0.945	0.952	0.712	0.963	0.980
	1000	0.818	0.707	0.867	Full	0.899	0.953	0.958	0.724	0.968	0.982
	Full	0.812	0.688	0.870	10	0.905	0.921	0.935	0.627	0.946	0.973
	Full	0.813	0.683	0.873	50	0.927	0.947	0.954	**0.397**	0.968	0.977
	Full	0.816	0.678	0.873	Full	0.935	0.954	0.958	0.625	0.970	0.981
UMTL-SSL-S	100	0.798	0.628	0.860	10	0.951	0.911	0.935	0.792	0.940	0.963
	100	0.834	0.696	0.874	50	0.972	0.946	0.952	0.727	0.965	0.977
	100	0.817	0.688	0.860	Full	0.975	0.951	0.954	0.812	0.968	0.981
	1000	0.806	0.652	0.872	10	0.956	0.916	0.937	0.852	0.943	0.966
	1000	0.808	0.662	0.862	50	0.971	0.944	0.952	0.917	0.965	0.978
	1000	0.801	0.646	0.862	Full	0.975	0.952	0.954	0.753	0.969	0.981
	Full	0.796	0.632	0.864	10	0.960	0.923	0.940	0.782	0.954	0.967
	Full	0.808	0.662	0.868	50	0.972	0.945	0.953	0.645	0.966	0.978
	Full	0.800	0.632	0.628	Full	0.961	0.924	0.940	0.392	0.948	0.969
MultiMix	100	0.800	0.594	0.856	10	0.954	0.920	0.938	0.695	0.949	0.969
	100	0.824	0.613	0.854	50	0.971	0.943	0.951	0.681	0.964	0.976
	100	0.792	0.593	0.854	Full	0.973	0.948	0.954	0.636	0.966	0.981
	1000	0.817	0.647	0.865	10	0.954	0.910	0.932	0.902	0.942	0.968
	1000	0.825	0.650	0.860	50	0.970	0.941	0.950	0.811	0.964	0.977
	1000	0.830	0.586	0.856	Full	**0.974**	0.919	0.953	0.643	0.933	0.984
	Full	0.840	0.730	0.880	10	0.954	0.913	0.935	0.621	0.949	0.968
	Full	**0.854**	**0.760**	**0.890**	50	0.972	**0.950**	**0.956**	0.692	**0.970**	**0.980**
	Full	0.843	0.740	0.890	Full	0.975	0.952	0.960	0.528	0.970	0.982

Results (Cross-Domain Evaluation)

Model	$\lvert\mathcal{D}_l^c\rvert$	Classification Acc	Classification F1-Nor	Classification F1-Abn	$\lvert\mathcal{D}_l^s\rvert$	Segmentation DS	Segmentation JS	Segmentation SSIM	Segmentation HD	Segmentation P	Segmentation R
U-Net	—	—	—	—	10	0.555	0.480	0.680	8.691	0.553	0.866
	—	—	—	—	50	0.763	0.736	0.870	2.895	0.752	0.887
	—	—	—	—	Full	0.838	0.906	0.929	1.414	0.793	0.910
Enc	100	0.352	0.070	0.506	—	—	—	—	—	—	—
	1000	0.390	0.192	0.508	—	—	—	—	—	—	—
	Full	0.434	0.296	0.524	—	—	—	—	—	—	—
Enc-SSL	100	0.402	0.222	0.510	—	—	—	—	—	—	—
	1000	0.486	0.380	0.530	—	—	—	—	—	—	—
	Full	0.510	0.472	0.538	—	—	—	—	—	—	—
UMTL	100	0.350	0.045	0.510	10	0.586	0.708	0.836	7.156	0.731	0.950
	100	0.363	0.085	0.515	50	0.580	0.684	0.825	7.013	0.697	0.975
	100	0.342	0.015	0.508	Full	0.607	0.742	0.863	6.398	0.759	0.968
	1000	0.413	0.263	0.507	10	0.676	0.674	0.833	3.268	0.712	0.927
	1000	0.400	0.203	0.513	50	0.704	0.811	0.896	3.232	0.828	0.964
	1000	0.430	0.293	0.517	Full	0.638	0.795	0.890	3.893	0.810	0.966
	Full	0.455	0.365	0.525	10	0.737	0.765	0.879	0.917	0.801	0.930
	Full	0.444	0.332	0.522	50	0.868	0.793	0.894	0.742	0.898	0.946
	Full	0.443	0.328	0.520	Full	0.854	0.828	0.913	0.792	0.866	0.942
UMTL-S	100	0.344	0.006	0.510	10	0.797	0.670	0.807	5.754	0.698	0.938
	100	0.364	0.098	0.506	50	0.828	0.715	0.826	6.412	0.731	0.971
	100	0.342	0.008	0.510	Full	0.838	0.715	0.834	6.321	0.740	0.966
	1000	0.378	0.138	0.512	10	0.844	0.718	0.854	3.921	0.754	0.939
	1000	0.392	0.186	0.514	50	0.883	0.793	0.888	3.017	0.821	0.959
	1000	0.370	0.130	0.510	Full	0.898	0.831	0.905	4.150	0.845	0.970
	Full	0.470	0.398	0.524	10	0.881	0.785	0.888	0.862	0.830	0.939
	Full	0.413	0.270	0.510	50	0.917	0.848	0.919	0.658	0.966	0.888
	Full	0.433	0.315	0.513	Full	0.916	0.850	0.921	0.882	0.886	0.952
UMTL-SSL	100	0.442	0.316	0.524	10	0.833	0.778	0.884	0.895	0.810	0.948
	100	0.398	0.166	0.520	50	0.853	0.839	0.907	0.851	0.864	0.952
	100	0.385	0.165	0.515	Full	0.841	0.818	0.911	0.853	0.854	0.949
	1000	0.445	0.333	0.525	10	0.818	0.781	0.892	1.085	0.825	0.938
	1000	0.526	0.486	0.544	50	0.826	0.804	0.904	0.811	0.792	0.949
	1000	0.485	0.413	0.538	Full	0.843	0.837	0.924	0.983	0.882	0.953
	Full	0.526	0.504	0.546	10	0.824	0.765	0.873	0.994	0.790	0.943
	Full	0.530	0.514	0.542	50	0.867	0.839	0.917	0.566	0.881	0.945
	Full	<u>0.520</u>	<u>0.490</u>	0.542	Full	0.884	0.884	0.934	0.599	0.918	0.955
UMTL-S-SSL	100	0.370	0.114	0.510	10	0.853	0.747	0.866	1.048	0.782	0.944
	100	0.400	0.192	0.518	50	0.889	0.799	0.899	0.854	0.834	0.950
	100	0.370	0.114	0.514	Full	0.915	0.848	0.920	0.987	0.880	0.956
	1000	0.432	0.286	0.524	10	0.871	0.785	0.884	1.327	0.818	0.944
	1000	0.458	0.342	0.530	50	0.893	0.803	0.895	1.123	0.835	0.946
	1000	0.462	0.350	0.536	Full	0.930	0.860	0.925	1.042	0.912	0.955
	Full	0.482	0.412	0.536	10	0.880	0.765	0.885	0.745	0.818	0.941
	Full	0.490	0.426	0.540	50	0.912	0.845	0.909	0.956	0.881	0.952
	Full	0.510	0.474	0.540	Full	0.875	0.809	0.875	0.722	0.851	0.944
MultiMix	100	0.440	0.164	0.510	10	0.857	0.732	0.863	1.227	0.767	0.943
	100	0.370	0.036	0.510	50	0.889	0.790	0.890	1.061	0.866	0.947
	100	0.500	0.300	0.510	Full	0.899	0.825	0.906	0.647	0.852	0.952
	1000	0.520	0.386	0.530	10	0.862	0.775	0.878	1.307	0.816	0.939
	1000	0.540	0.500	0.536	50	0.912	0.831	0.907	1.293	0.865	0.955
	1000	**0.570**	**0.620**	0.510	Full	**0.936**	**0.880**	**0.932**	0.803	0.917	**0.979**
	Full	0.550	0.430	0.534	10	0.886	0.802	0.894	0.746	0.839	0.948
	Full	0.560	0.570	**0.550**	50	0.935	0.878	0.930	**0.515**	**0.928**	0.957
	Full	<u>0.520</u>	<u>0.490</u>	<u>0.550</u>	Full	<u>0.943</u>	<u>0.892</u>	<u>0.937</u>	<u>0.417</u>	<u>0.928</u>	<u>0.958</u>

Stanford MEDICINE

Results (Segmentation Consistency)

In-Domain

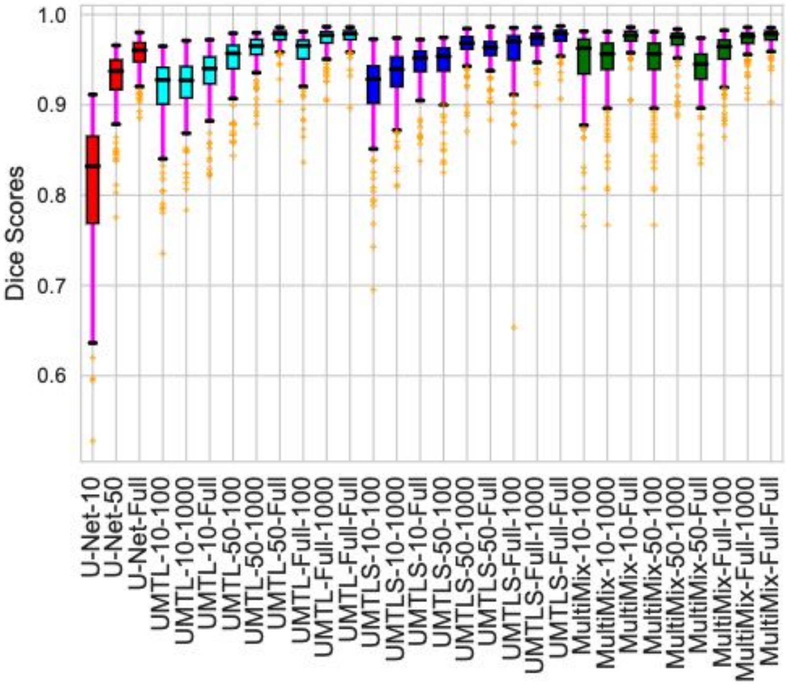

Cross-Domain

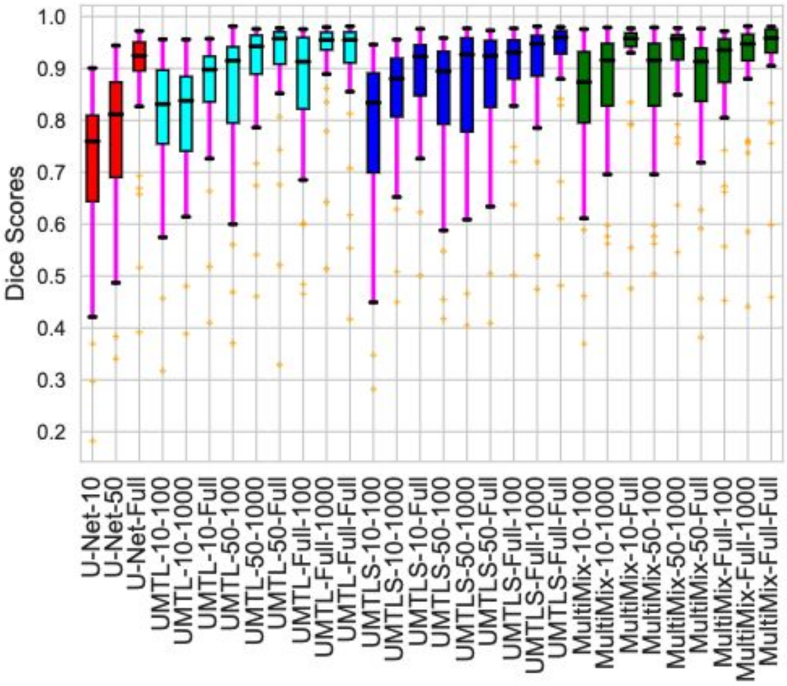

Stanford MEDICINE

Results (Segmentation Boundary)

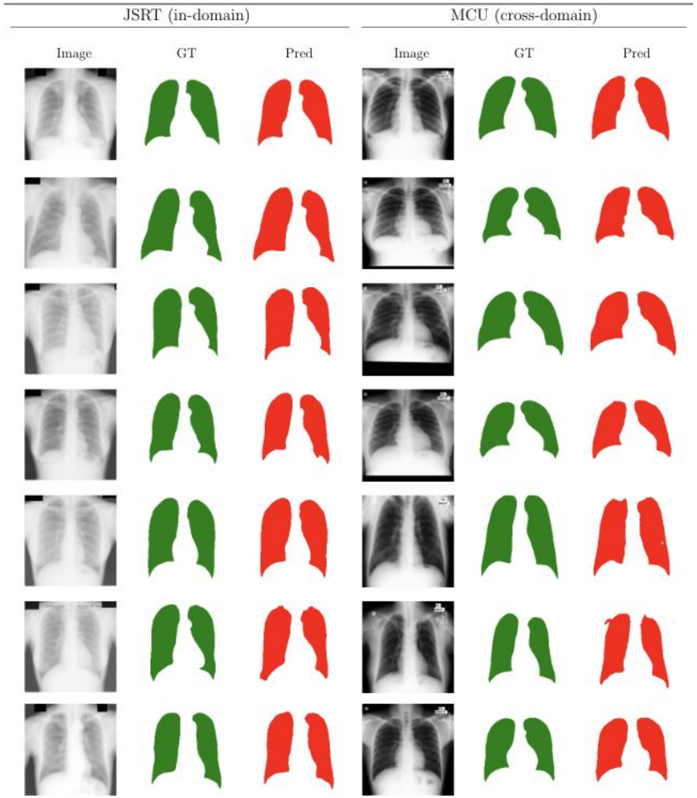

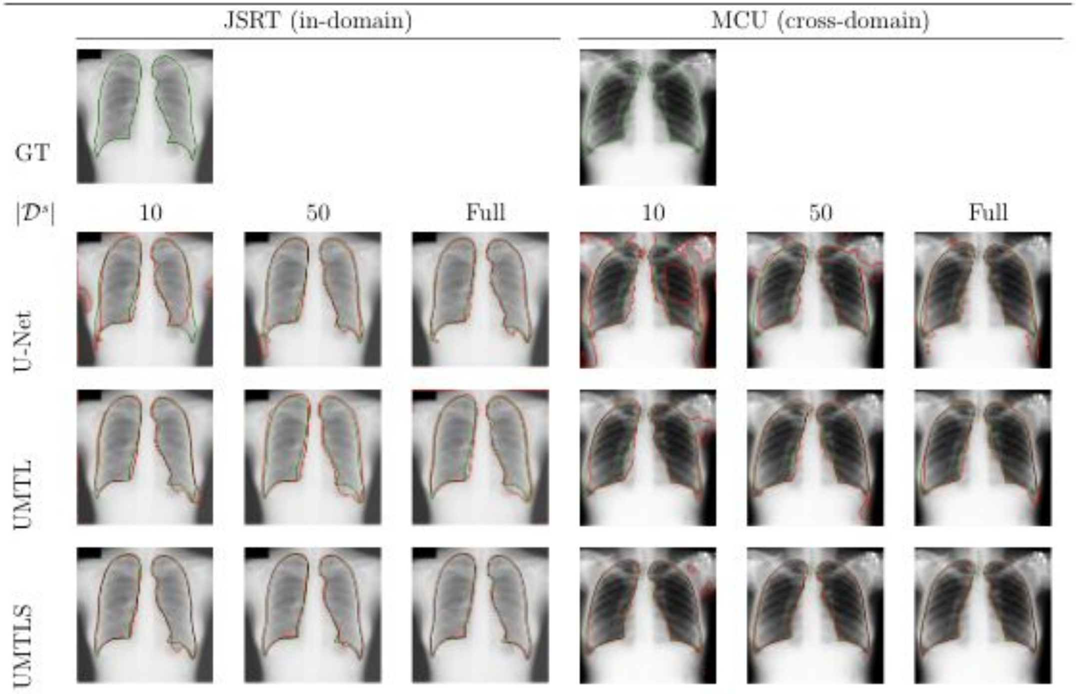

Stanford MEDICINE

Results (Bland Altman Plots)

In-Domain

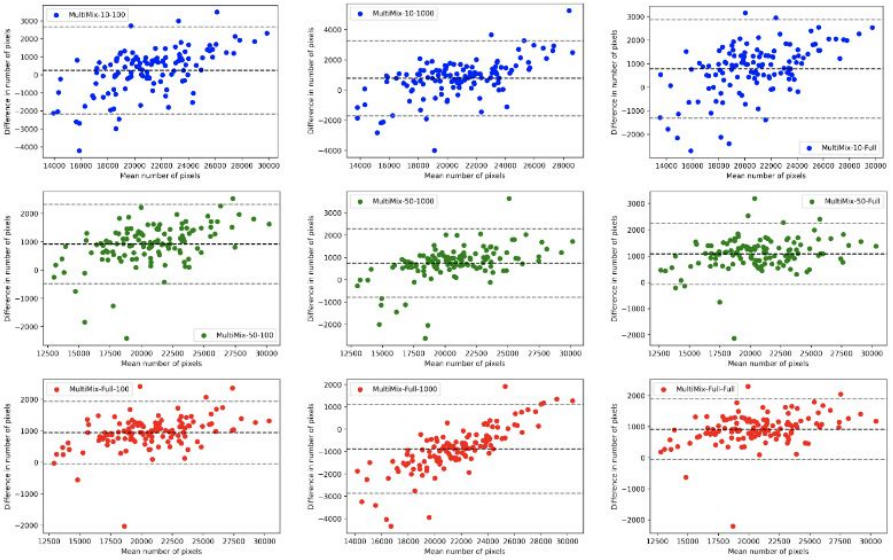

Cross-Domain

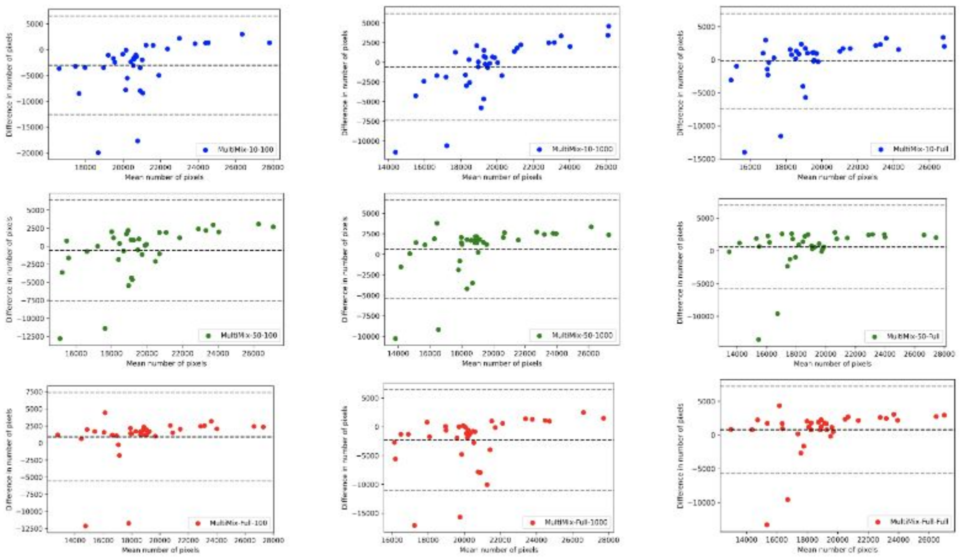

Stanford MEDICINE

Results (Consistency Curves)

In-Domain

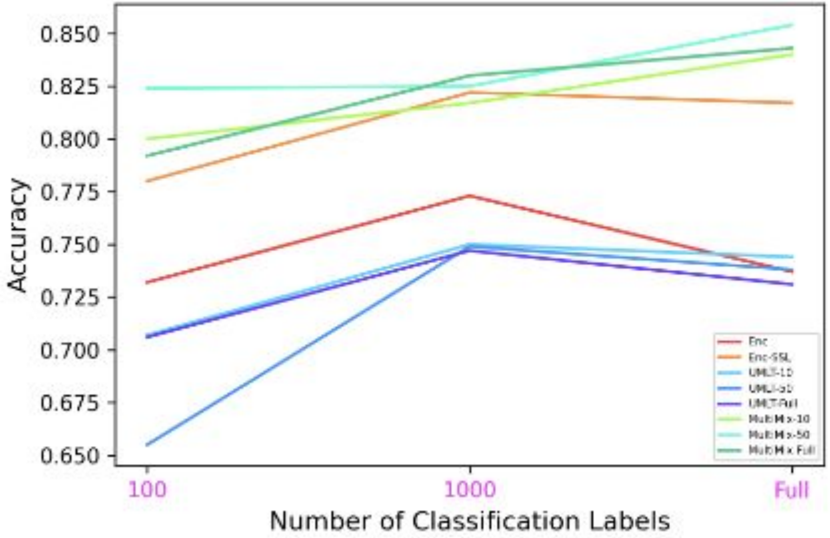

Cross-Domain

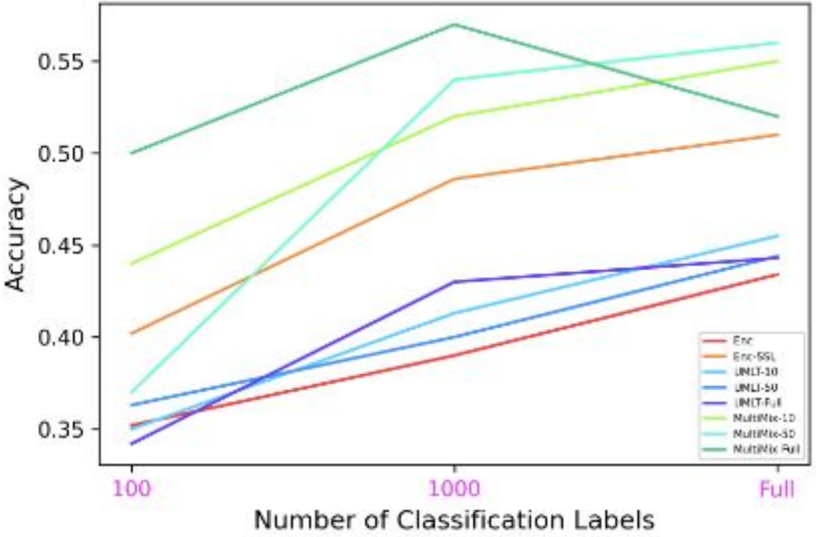

Stanford
MEDICINE

Results (ROC Curves)

In-Domain

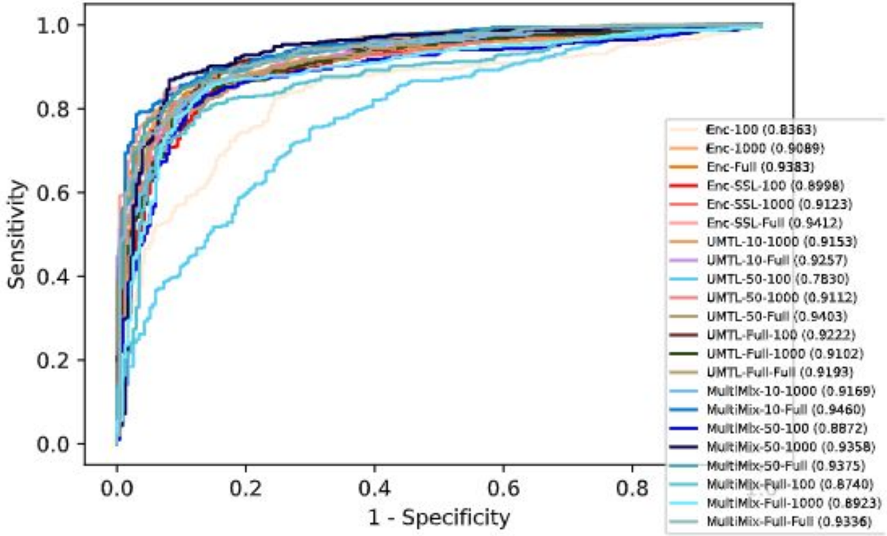

Cross-Domain

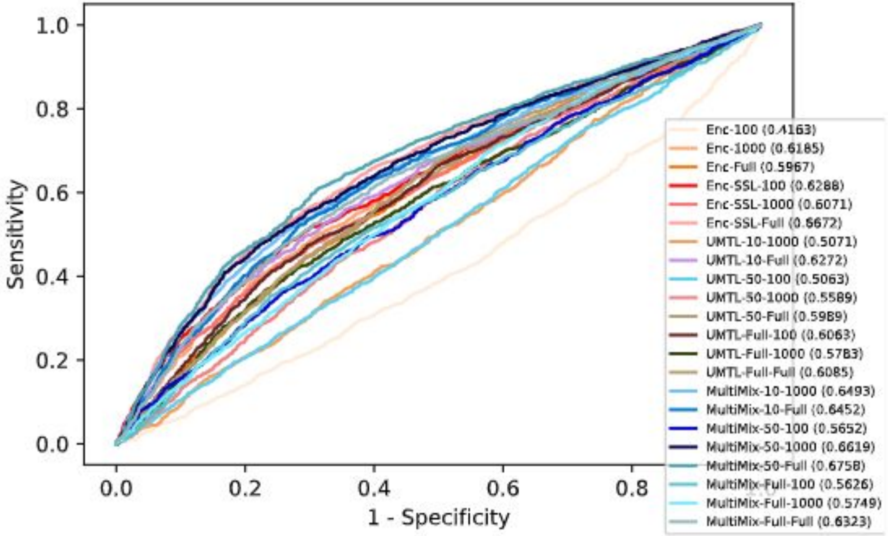

Strengths/Weaknesses

- Strengths
  - Strong Qualitative and Quantitative
  - performance
  - Extensive results supporting the effectiveness of the model
  - Outperformance of multiple baselines
  - Novel Methods
- Weaknesses
  - Cross-Domain classification can be improved
  - HD scores of MultiMix is lower than some baselines

Stanford MEDICINE

Acknowledgements

- Dr. Abdullah-Al-Zubaer Imran

- Prof. Adam Wang

- Prof. Demetri Terzopoulos

Stanford MEDICINE

Conclusions

- Novel, Sparingly Supervised, Multi-Tasking Model
  - Leveraging consistency augmentation and confidence based learning
  - Bridging two tasks through a saliency module
- Multi-Source, Multi-Domain Learning
  - Performs diagnostic classification (pneumonia) and segmentation (lungs)
- Extensive Experimentation and Results
  - Both qualitative and qualitative performance against multiple baselines
  - In-Domain and Cross-Domain performances demonstrate the generalization potential of the proposed MultiMix model

Thanks for Listening!

MultiMix: Sparingly-Supervised, Extreme Multitask Learning From Medical Images

Ayaan Haque, Abdullah-Al-Zubaer Imran, Adam Wang, Demetri Terzopoulos

Saratoga High, Stanford University, UCLA, VoxelCloud Inc.