J.P.Morgan

ST-FL: Style Transfer Preprocessing in Federated Learning

Antonios Georgiadis[1,*], Varun Babbar[1,2,*], Fran Silavong[1], Sean Moran[1], Rob Otter[1]

CTO Applied Research, JP Morgan Chase.[1] University of Cambridge[2],

[*]These authors contributed equally to this work.

Introduction

Federated Learning

- Due to stringent privacy laws, the sharing of confidential data (such as CT scans that can potentially be used to identify patients) between institutions and countries is fraught with difficulties, and is generally considered impossible.
- Federated learning solves the problem of how to learn a single model based on data that is locked away in data silos without revealing per-client private data to other clients or the central server.
- In FL, only the model weights are shared between clients and a central server (where they are averaged) during training and not the actual training data, which is considered private and highly confidential.

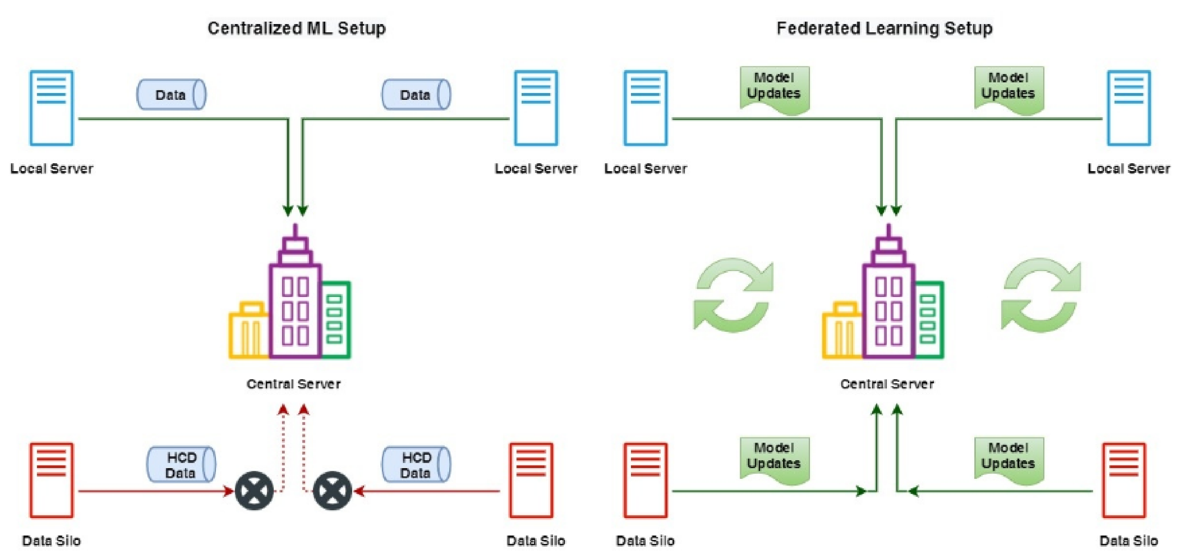

Figure: Centralized vs Federated Training

Challenges Associated with Federated Learning

- Medical CT scans are often sourced from different institutions, which may use different imaging techniques, instruments, and post-processing algorithms. This inevitably leads to variability in datasets sourced from different clients.
- In this scenario, averaging client model weights in a federated setting during training may lead to a one-size-fits-none situation.
- Prior research has explored the benefits of federated learning for leveraging disparate datasets for the purpose of COVID-19 chest CT scan segmentation, eg [2], [4], and [6].
- However, there is no previous research that accounts for the differing factors of variation of CT images that are distributed across client nodes.

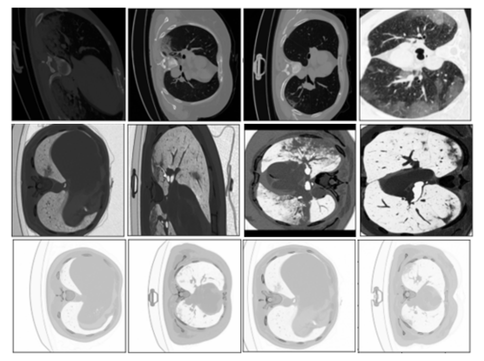

Figure: CT Scans of Variable Noise Patterns Commonly Seen in Practice

Contributions

- **Mixed CT image data quality & the effect on FL**: We highlight the negative effect of differing quality images on client nodes on the accuracy of a federated U-Net [5] for CT image segmentation.
- **Noise agnostic FL framework for different types of noise**: We present *ST-FL*, a federated learning framework that incorporates a denoising CycleGAN[7] at each client node, standardising image quality per client and increasing the robustness of federated learning to mixed data quality observed in practice.

Proposed Scheme

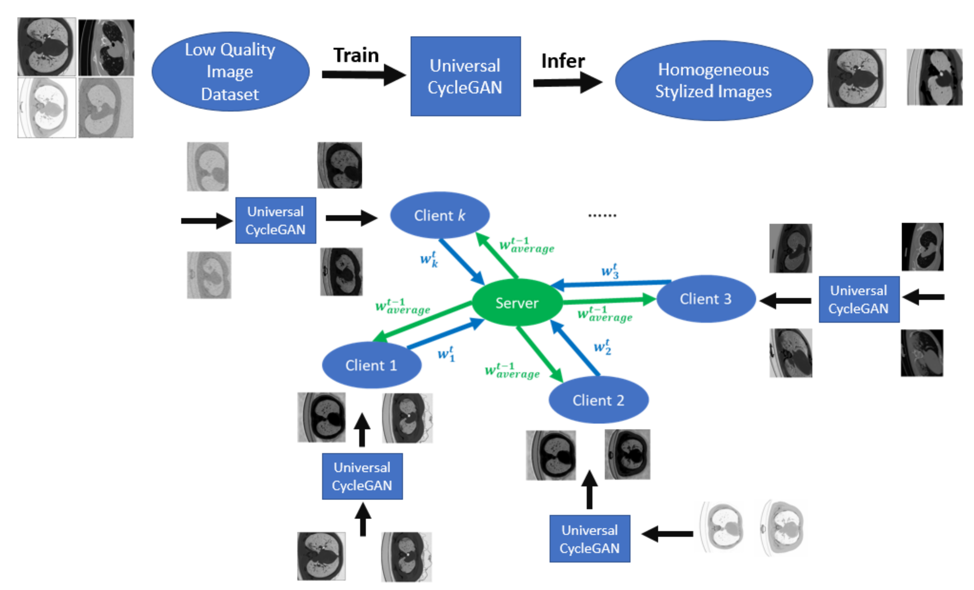

Figure: Federated Learning with Universal CycleGAN

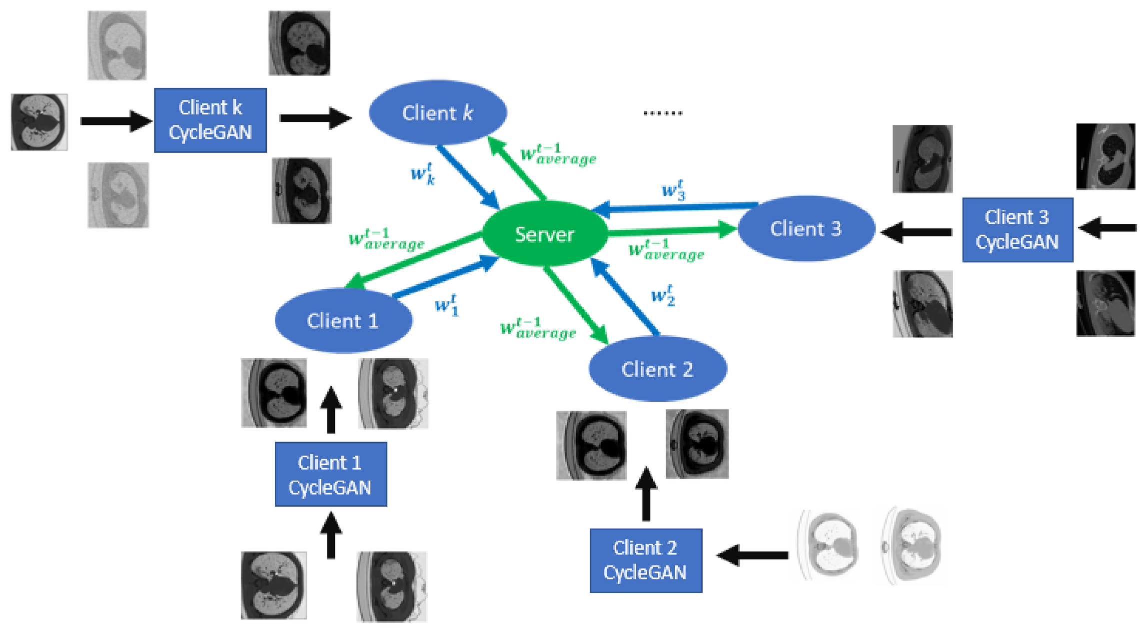

Figure: Federated Learning with Client Specific CycleGAN

- Given a low quality domain LQ represented by $D_{agg} \subseteq \cup_{k=1}^{N} \mathcal{D}^k$ for client datasets D_k and the target domain T (**where T is the same for all clients**), we learn generators - $\mathcal{G}_{LQT}(x_{LQ}) : LQ \to T$ and $\mathcal{G}_{TLQ}(x_T) : T \to LQ$ - and discriminators $D_{LQ}(x_{LQ}||\mathcal{G}_{TLQ}(x_T))$ and $D_T(x_T||\mathcal{G}_{\mathcal{LQ}T}(x_{LQ}))$.
- The resulting style transferred images are concatenated channelwise with the original images and the input fed into the **client specific segmentation network**, which is trained in the normal federated setting.
- In this approach, we train an individual CycleGAN for each client. Specifically, for a client style domain C_k represented by dataset D_k and target domain T (**where T is the same for all clients**), we learn generators - $\mathcal{G}_{C_kT}(x_{C_k}) : C_k \to T$ and $\mathcal{G}_{TC_k}(x_T) : T \to C_k$ - and discriminators $D_{C_k}(x_{C_k}||\mathcal{G}_{TC_k}(x_Y))$ and $D_T(x_T||\mathcal{G}_{C_{k}T}(x_{C_k}))$.
- The resulting style transferred images are concatenated channelwise with the original images as before and normal federated training then takes place.

Experimental Methodology

In order to test the efficacy of our schemes, we perform experiments with 2 different types of client datasets: Synthetic and Semi-Synthetic. We also compare our schemes to the vanilla FedAvg model and a centralised model that is trained on Client Specific CycleGAN style transfered versions of all client datasets. We perform experiments with 3,4, and 5 clients and report all results averaged over 10 trials (with 95% CI).

Synthetic Dataset

- We use the **Coronacases** [3] dataset of COVID-19 patient chest scans, both in its vanilla form and in an augmented form wherein each client dataset represents different noise patterns added to the dataset.
- This scenario models a situation where different client institutions may have chest scans with **similar structural characteristics but differing style characteristics** (eg. differing imaging modalities and noise levels).

Figure: Sample Images from Synthetic Datasets: (Left to Right: Vanilla Coronacases, Mixed Noise Coronacases, Noisy Coronacases, Inversion Coronacases, Contrast Enhanced Coronacases - Style Target)

Semi-Synthetic Dataset

- We use the **Coronacases** and **MedSeg** [1] dataset as client datasets and augment them with similar noise patterns as above to create additional client datasets.
- Compared to the Coronacases dataset, the MedSeg dataset was seen to have **noisy labels and some structural differences** that can potentially hinder effective training.
- These data-centric properties model the scenario where client institutions may have scans with **differing structural and stylistic characteristics** and some improperly labelled examples.

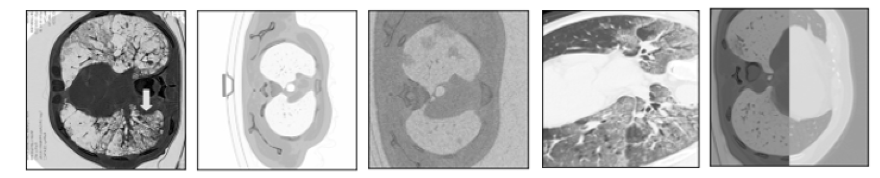

Figure: Sample Images from Semi-Synthetic Datasets: (Left to Right: Vanilla Medseg - Style Target, Vanilla Coronacases, Noisy Coronacases, Inversion MedSeg, Mixed Noise Coronacases)

Validation Curves Averaged All Experiments

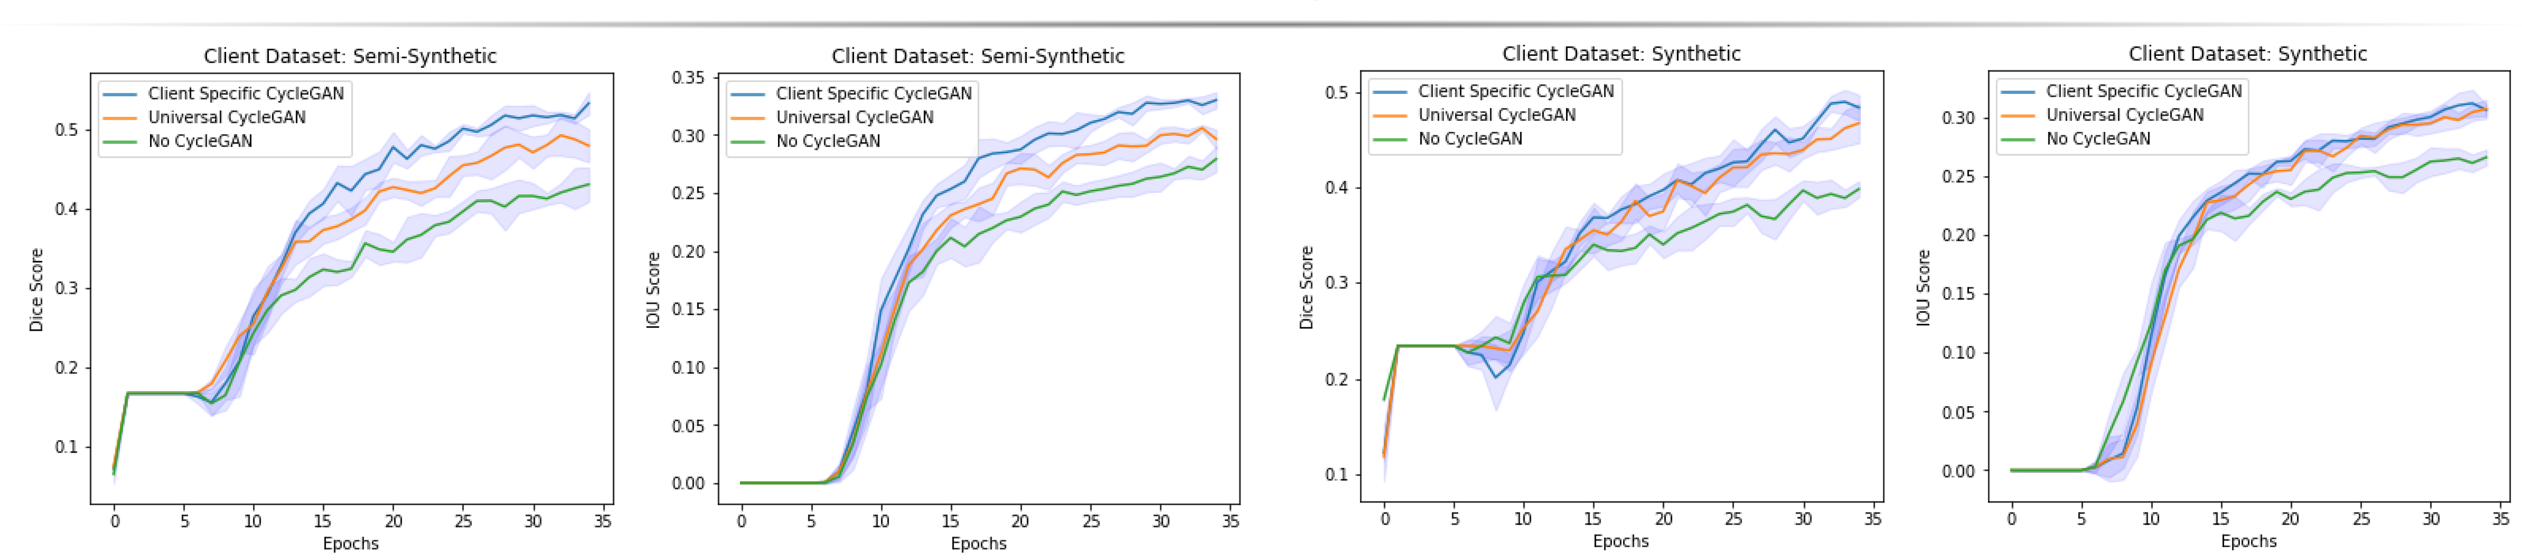

Quantitative results

Number of Clients	Metric	Dataset Type	Federated Training: Vanilla FedAvg	Federated Training: Universal CycleGAN	Federated Training: Client Specific CycleGAN	Centralised Training
3	Dice	Synthetic	0.414 ±0.012	0.505 ±0.027	**0.533 ±0.041**	0.560 ±0.023
		Semi-Synthetic	0.497 ±0.009	0.520 ±0.012	**0.539 ±0.012**	0.494 ±0.009
	IOU	Synthetic	0.274 ±0.011	0.329 ±0.017	**0.336 ±0.019**	0.347 ±0.016
		Semi-Synthetic	0.337 ±0.014	0.355 ±0.016	**0.366 ±0.006**	0.338 ±0.010
4	Dice	Synthetic	0.430 ±0.008	0.493 ±0.026	**0.514 ±0.037**	0.528 ±0.017
		Semi-Synthetic	0.462 ±0.014	0.483 ±0.021	**0.489 ±0.014**	0.450 ±0.013
	IOU	Synthetic	0.287 ±0.007	0.296 ±0.013	**0.332 ±0.017**	0.329 ±0.007
		Semi-Synthetic	0.306 ±0.007	0.319 ±0.006	**0.321 ±0.004**	0.305 ±0.008
5	Dice	Synthetic	0.378 ±0.012	0.460 ±0.025	**0.465 ±0.006**	0.490 ±0.006
		Semi-Synthetic	0.390 ±0.018	0.420 ±0.017	**0.462 ±0.017**	0.392 ±0.018
	IOU	Synthetic	0.255 ±0.008	0.240 ±0.009	**0.282 ±0.013**	0.310 ±0.012
		Semi-Synthetic	0.265 ±0.007	0.281 ±0.012	**0.301 ±0.010**	0.261 ±0.010

Table: Federated and Centralised Performance Metrics for Differing Numbers of Clients (Best Epoch)

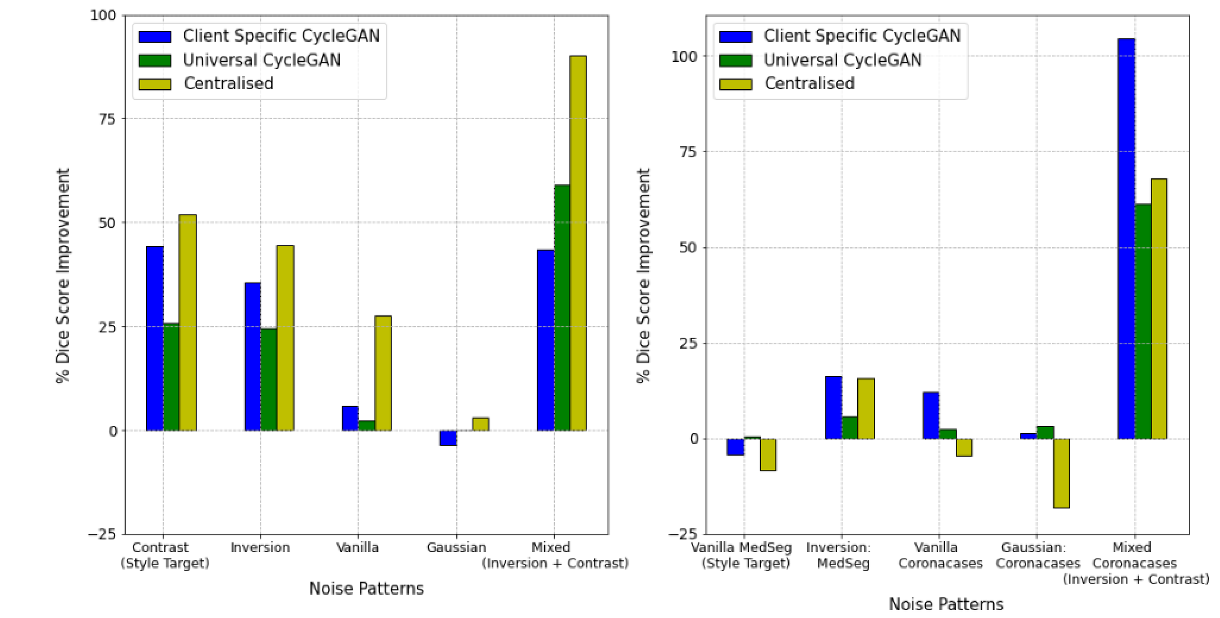

Figure: % Noise-Specific Performance Improvement of the Different Techniques Tested Relative to Vanilla FedAvg (Best Epoch, Left: Synthetic Datasets, Right: Semi-Synthetic Datasets)

Qualitative results

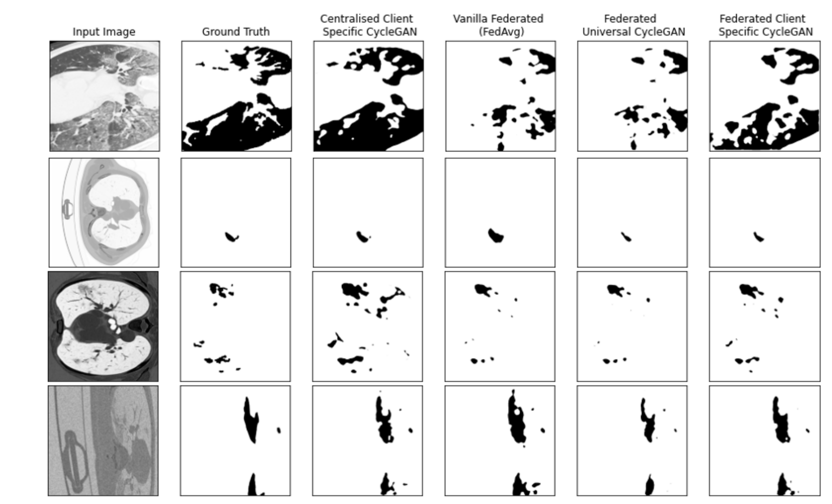

Figure: Thresholded Segmentations Produced Using the Different Schemes Tested

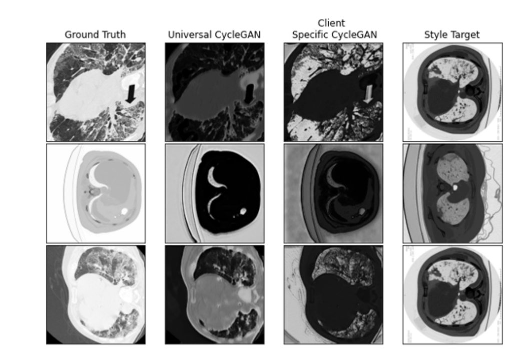

Figure: A Comparison Between Style Transferred Images Produced by the Universal and Client Specific CycleGAN. Images produced by the Client Specific CycleGAN bear greater resemblance to the style target than those produced by the Universal CycleGAN.

References

[1] Covid-19 CT scan image data and segmentation dataset. Free to download. `http://medicalsegmentation.com/covid19/`.

[2] Wonyong Jeong, Jaehong Yoon, Eunho Yang, and Sung Ju Hwang. Federated semi-supervised learning with inter-client consistency, 2020.

[3] Ma Jun, Ge Cheng, Wang Yixin, An Xingle, Gao Jiantao, Yu Ziqi, Zhang Minqing, Liu Xin, Deng Xueyuan, Cao Shucheng, et al. Covid-19 ct lung and infection segmentation dataset. *Zenodo, Apr*, 20, 2020.

[4] Rajesh Kumar, Abdullah Aman Khan, Sinmin Zhang, Wenyong Wang, Yousif Abuidris, Waqas Amin, and Jay Kumar. Blockchain-federated-learning and deep learning models for COVID-19 detection using CT imaging. *CoRR*, abs/2007.06537, 2020.

[5] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention 2015*, pages 234–241, Cham, 2015. Springer International Publishing.

[6] Dong Yang, Ziyue Xu, Wenqi Li, Andriy Myronenko, Holger R Roth, Stephanie Harmon, Sheng Xu, Baris Turkbey, Evrim Turkbey, Xiaosong Wang, et al. Federated semi-supervised learning for covid region segmentation in chest ct using multi-national data from china, italy, japan. *Medical image analysis*, 70:101992, 2021.

[7] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In