Learning Fixed Points in Generative Adversarial Networks: From Image-to-Image Translation to Disease Detection and Localization

Md Mahfuzur Rahman Siddiquee[1], Zongwei Zhou[2,4], Nima Tajbakhsh[2], Ruibin Feng[2], Michael B. Gotway[3], Yoshua Bengio[4], and Jianming Liang[2,4]

[1]School of Computing, Informatics, and Decision Systems Engineering; [2]Biomedical Informatics, College of Health Solutions, Arizona State University
[3]Department of Radiology, Mayo Clinic Arizona; [4]Mila - Quebec Artificial Intelligence Institute

ARIZONA STATE UNIVERSITY · MAYO CLINIC · GPSA Graduate & Professional Student Association · College of Health Solutions, Arizona State University · Mila · ICCV 2019 Seoul, Korea

The Question

Can GAN remove eyeglasses, if present, from an image **while otherwise** preserving the image?

Can we remove diseases as well?

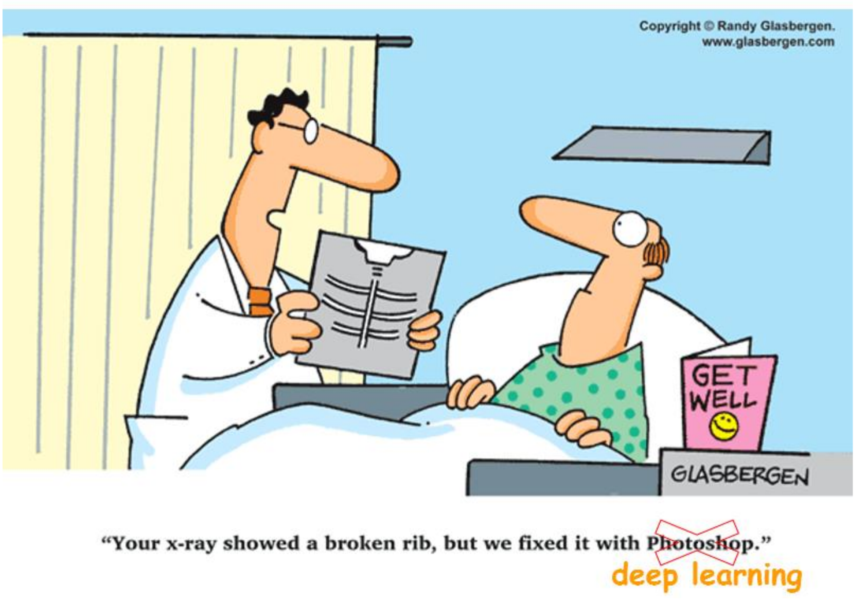

The Motivation

The eyeglasses, if present, can be localized by subtracting the given image from the generated image without the eyeglasses. Similarly, we can detect and localize the diseases using image-level annotations only.

Obtaining image-level annotation is more feasible and practical than lesion-level annotation, as it can be obtained automatically using natural language processing.

The Fixed-Point Translation

We formulate the problem of removing the object, if present, from an image while otherwise preserving the image as Fixed-Point Image-to-Image Translation, which has the following properties:

1. Handling unpaired images, as it would be too late to acquire a healthy image from a cancer-stricken patient.
2. Translating any image to any given target domain without the source domain label, as a GAN for virtual healing aims to turn any image, with an unknown health status, into a healthy one.
3. Conducting an identity transformation for same-domain translation, as a GAN for virtual healing must leave a healthy image intact, injecting neither artifacts nor new information into the image.
4. Transforming only cross-domain-relevant attributes, without affecting any unrelated attributes, as removing diseases from an image should have no impact on normal regions.

Overcoming the Limitations of the State of the Art in Image-to-Image Translation

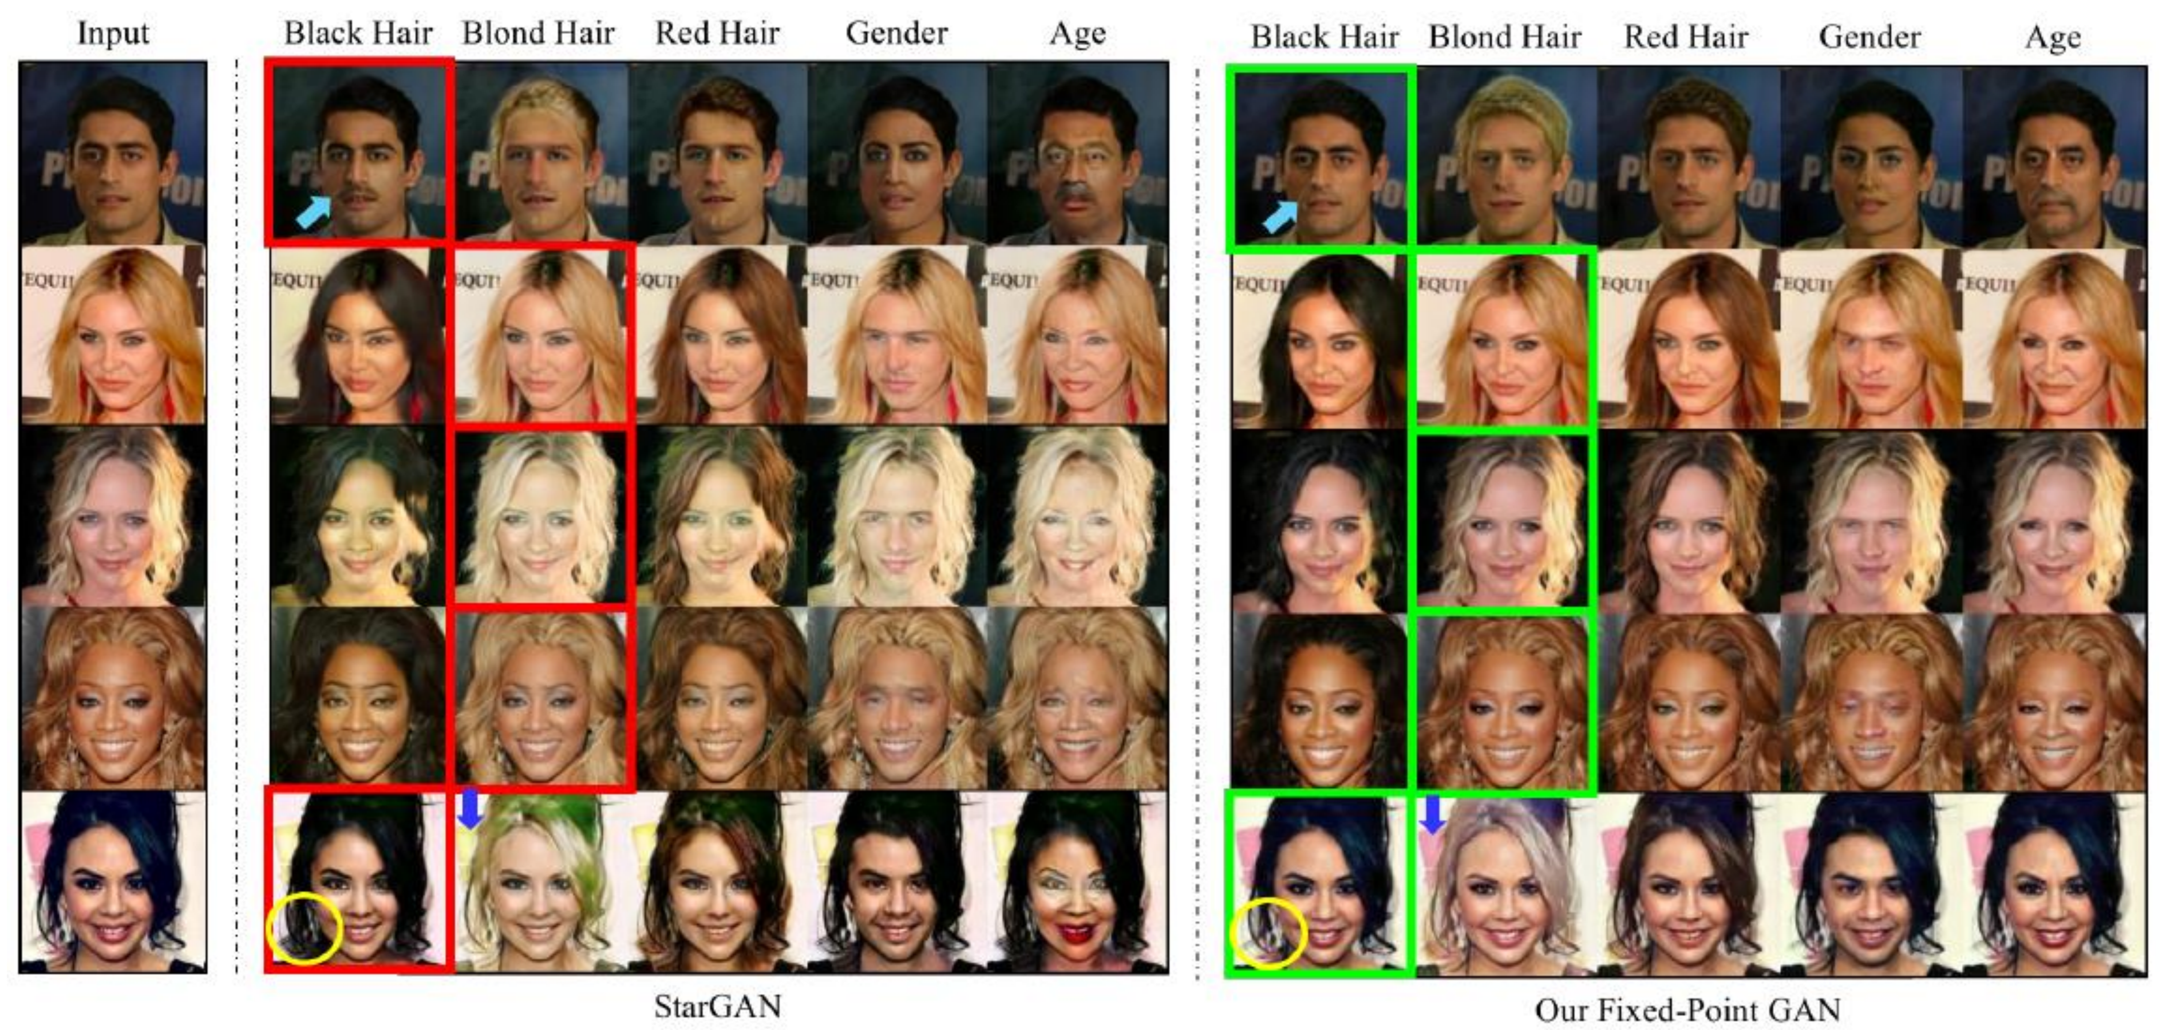

The current state-of-the-art multi-domain image-to-image translation method (StarGAN)[2] violates Requirements 3 and 4 and generates many artifacts (framed in red). Our Fixed-Point GAN overcomes these limitations (framed in green) via fixed-point translation learning.

Offering a Framework for Disease Detection & Localization with Image-Level Annotation

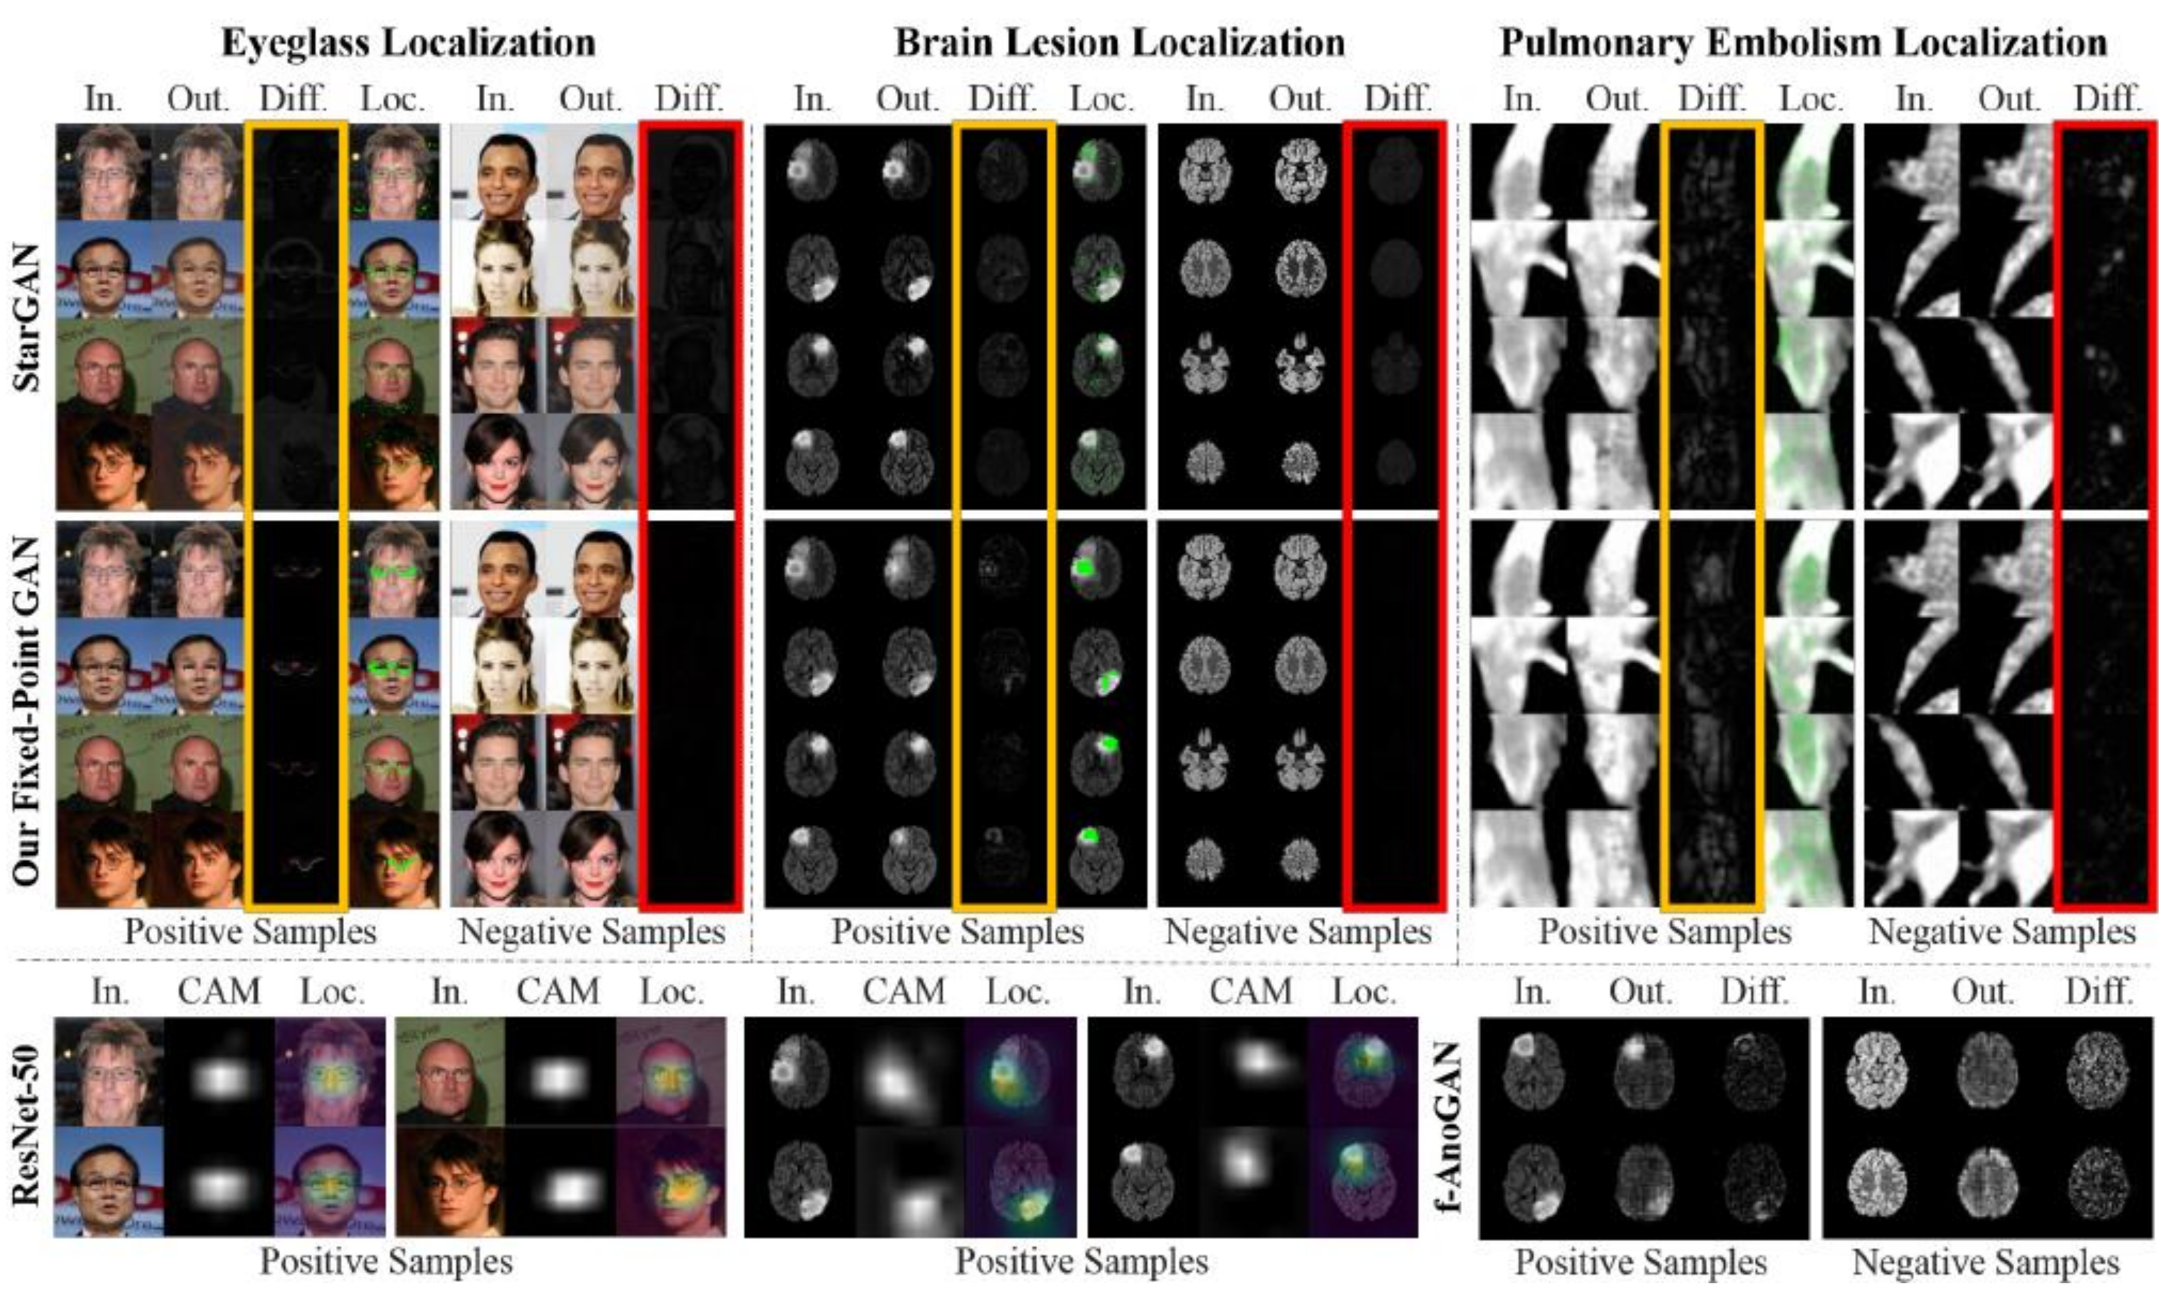

Through fixed-point translation learning, our Fixed-Point GAN aims to preserve healthy images during the translation, thereby few differences between the generated (healthy) images and the original (healthy) images are observed in the difference maps (framed in red). For diseased images, owing to the transformation learning from diseased images to healthy ones, disease locations are revealed in the difference maps (framed in yellow). Besides, Fixed-Point GAN is more precise than CAM and f-AnoGAN for localizing eyeglasses and diseases (bottom row).

Training our Fixed-Point GAN

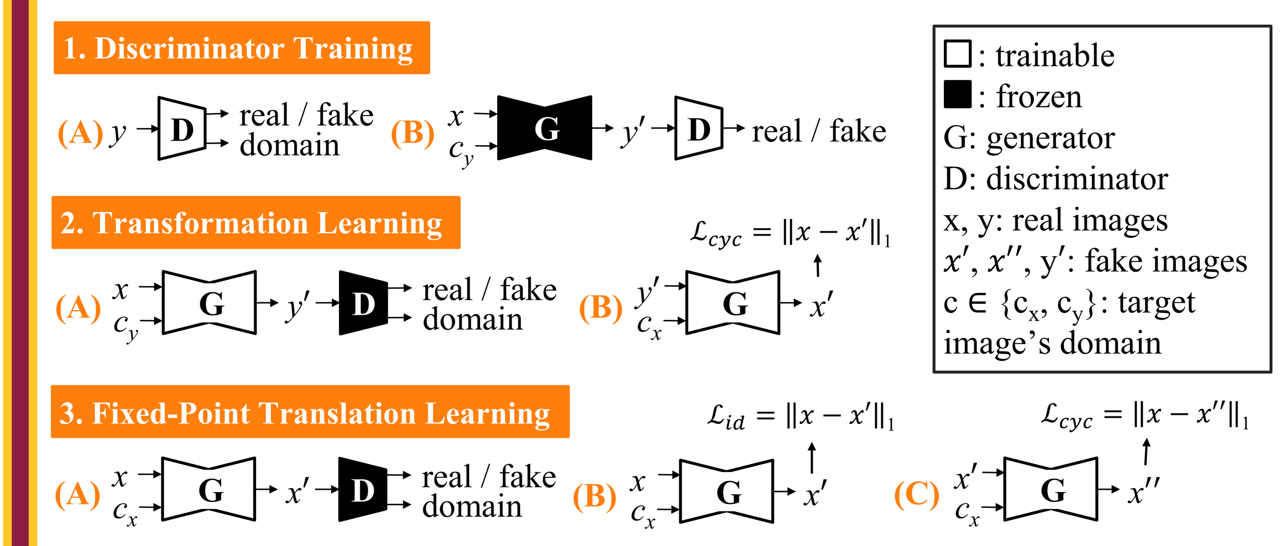

* This figure depicts the steps of a single iteration of the entire training process. After training, only the model **G** is utilized for inference.

Revised Objective Functions

	Loss	Definition
Eq. 1	$\mathcal{L}_{adv}$	$= \sum_{c\in\{c_x,c_y\}} \mathbb{E}_{x,c}\left[\log\left(1 - D_{r/f}(G(x,c))\right)\right] + \mathbb{E}_x\left[\log D_{r/f}(x)\right]$
Eq. 2	$\mathcal{L}^r_{domain}$	$= \mathbb{E}_{x,c_x}\left[\log D_{domain}(c_x\|x)\right]$
Eq. 3	$\mathcal{L}^f_{domain}$	$= \sum_{c\in\{c_x,c_y\}} \mathbb{E}_{x,c}[-\log D_{domain}(c\|G(x,c))]$
Eq. 4	$\mathcal{L}_{cyc}$	$= \sum_{c\in\{c_x,c_y\}} \mathbb{E}_{x,c_x,c}[\|G(G(x,c),c_x) - x\|_1]$
Eq. 5	$\mathcal{L}_{id}$	$= \mathbb{E}_{x,c}\,[\|G(x,c) - x\|_1]$ if $c = c_x$; 0 otherwise
Eq. 6	$\mathcal{L}_D$	$= -\mathcal{L}_{adv} + \lambda_{domain}\mathcal{L}^r_{domain}$
Eq. 7	$\mathcal{L}_G$	$= \mathcal{L}_{adv} + \lambda_{domain}\mathcal{L}^f_{domain} + \lambda_{cyc}\mathcal{L}_{cyc} + \lambda_{id}\mathcal{L}_{id}$

* Blue denotes our revisions to the loss functions of standard image-to-image translation training.

GitHub tinyurl.com/y62r5erw

Outperforming Existing Methods[1-5] for

Brain Lesion Detection and Localization (BRATS 2013)

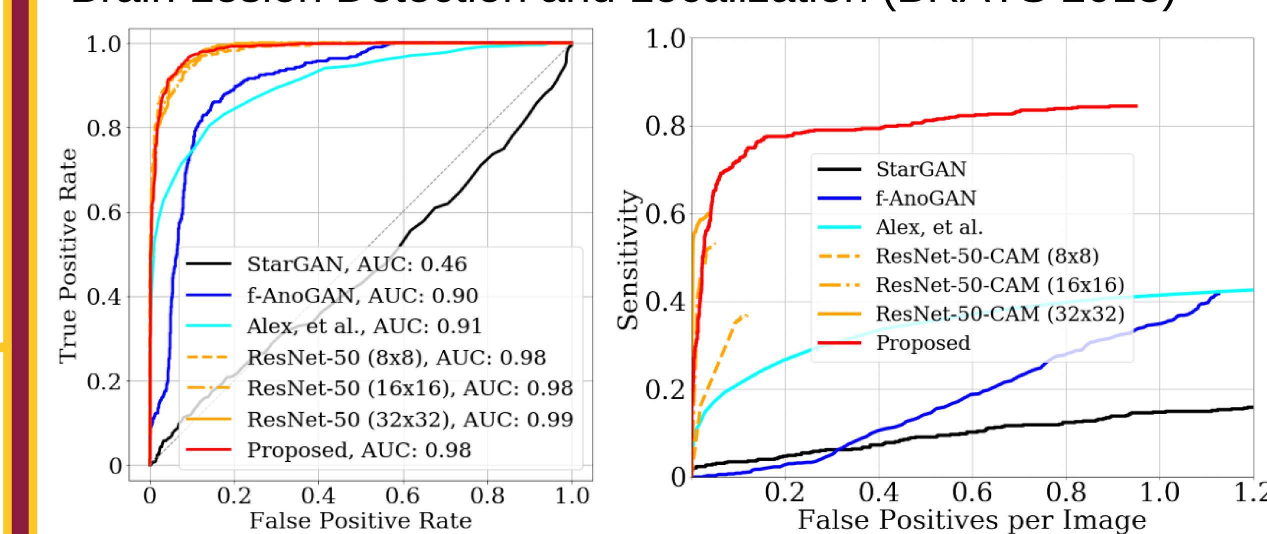

(a) Image-Level Detection on Synthetic Images

(b) Lesion-Level Localization on Synthetic Images

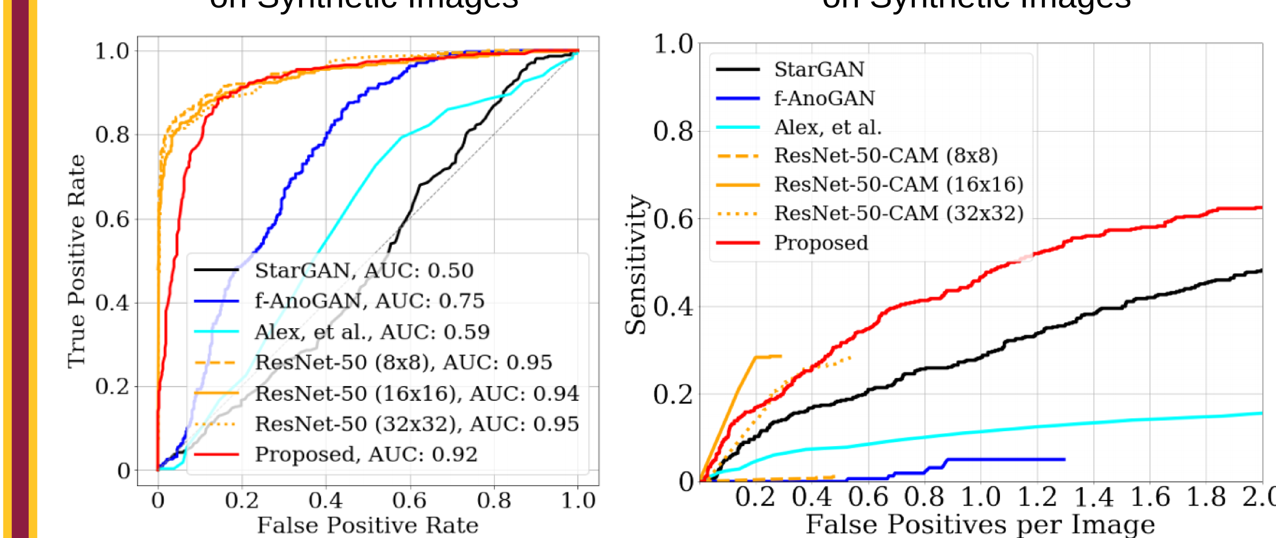

(a) Image-Level Detection on Real Images

(b) Lesion-Level Localization on Real Images

Outperforming Existing Methods[2,3] for

Pulmonary Embolism Detection and Localization

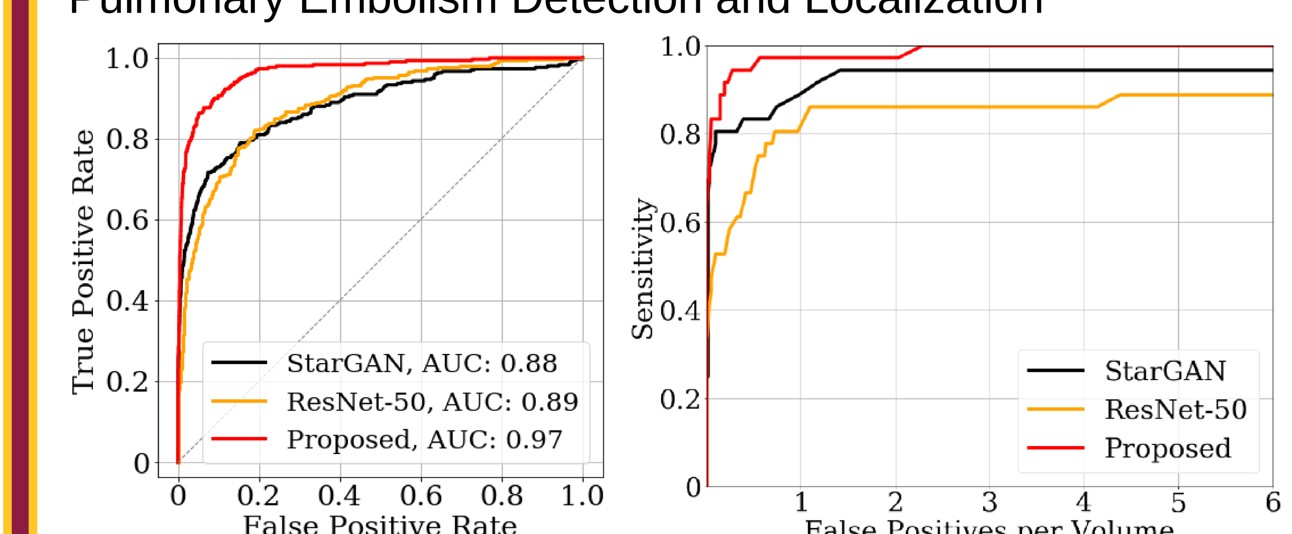

(a) Image-Level Detection

(b) Lesion-Level Localization

Improving Image-to-Image Translation

Real Images (Acc.)	Our Fixed-Point GAN	StarGAN
94.5%	92.31%	90.82%

(a) Cross-Domain Translation (Classification Accuracy)

Autoencoder	Our Fixed-Point GAN	StarGAN
0.11 ± 0.09	0.36 ± 0.35	2.40 ± 1.24

(b) Same-Domain Translation (L_1 Distance)

Ablation Study of Generator's Configuration

Dataset	StarGAN	w/ Delta	w/ Fixed-Point Translation	w/ Both
BRATS	0.4611	0.5246	0.9980	**0.9831**
PE	0.8832	0.8603	0.9216	**0.9668**

(a) Image-Level Detection (AUC)

Dataset	StarGAN	w/ Fixed-Point Translation	w/ Both
BRATS	13.6%	81.2%	**84.5%**
PE	88.9%	94.4%	**97.2%**

(b) Lesion-Level Loc. Sensitivity at 1 False Positive

Our Contributions

1. Our Fixed-Point GAN outperforms the state-of-the-art in image-to-image translation for both natural and medical images.
2. A novel method for disease detection and localization using image-level annotation, reducing annotation cost dramatically.

Conclusion

We introduced fixed-point translation and developed Fixed-Point GAN showing its effectiveness in image-to-image translation and disease detection and localization using image-level annotation.

References

1. Alex, V. et al. *Medical Imaging 2017: Image Processing, 10133*, 101330G (2017).
2. Choi, Y. et al. *CVPR*, 8789-8797 (2018).
3. He, K. et al. *CVPR*, 770-778 (2016).
4. Schlegl, T. et al. *Medical Image Analysis, 54*, 30-44 (2019).
5. Zhou, B. et al. *CVPR*, 2921-2929 (2016).

Acknowledgements

This research has been supported partially by ASU and Mayo Clinic through a Seed Grant and an Innovation Grant, and partially by NIH under Award Number R01HL128785.