

FoldIt: Haustral Folds Detection and Segmentation in Colonoscopy Videos

Shawn Mathew^{*1}, Saad Nadeem^{*2}, and Arie Kaufman¹

¹ Department of Computer Science, Stony Brook University

² Department of Medical Physics, Memorial Sloan Kettering Cancer Center

Abstract. Haustral folds are colon wall protrusions implicated for high polyp miss rate during optical colonoscopy procedures. If segmented accurately, haustral folds can allow for better estimation of missed surface and can also serve as valuable landmarks for registering pre-treatment virtual (CT) and optical colonoscopies, to guide navigation towards the anomalies found in pre-treatment scans. We present a novel generative adversarial network, FoldIt, for feature-consistent image translation of optical colonoscopy videos to virtual colonoscopy renderings with haustral fold overlays. A new transitive loss is introduced in order to leverage ground truth information between haustral fold annotations and virtual colonoscopy renderings. We demonstrate the effectiveness of our model on real challenging optical colonoscopy videos as well as on textured virtual colonoscopy videos with clinician-verified haustral fold annotations. All code and scripts to reproduce the experiments of this paper will be made available via our Computational Endoscopy Platform at <https://github.com/nadeemlab/CEP>.

Keywords: Colonoscopy · Haustral Folds Segmentation.

1 Introduction

High polyp miss rates during colonoscopy procedure are mainly attributed to colon wall protrusions, formed by circumferential contraction of the colon inner muscular layer, also known as haustral folds. These folds are extremely difficult to segment in optical colonoscopy (OC) videos due to texture and lighting variations along with specular reflections, fluid motion, and organ movements. If segmented, however, these folds can guide endoscope navigation towards high-occlusion areas, potentially reducing the polyp miss rate. Moreover, even though the colon can stretch and distort shape considerably in OC versus the pre-treatment CT virtual colonoscopy (VC), the haustral folds remain intact across the two modalities and hence can serve as useful landmarks for registration.

Mathew et al. [15] recently introduced a new unsupervised model, XDCycleGAN, for inferring scale-consistent depth maps from OC video frames using geometric information from 3D VC data, extracted from abdominal CT scans.

^{*}Equal contribution. Email: nadeems@mskcc.org

The unsupervised model was shown to handle variations in texture, lighting and specular reflections much more effectively than previous supervised approaches that were trained completely on OC or on VC datasets but not both simultaneously. Xu et al. [19] also showed superior performance by using cycle-consistency and optical flow for spatially- and temporally-consistent translation of simulated VC flythroughs to real OC videos.

In this work, we present FoldIt, a new generative adversarial network that can accurately detect and segment haustral folds in OC videos using unpaired image-to-image translation of OC frames to VC haustral fold renderings. *We show that the haustral fold segmentation via our model leads to feature-consistent domain translation for OC video sequences; the feature-consistency refers to consistency of haustral fold annotations between consecutive frames.* FoldIt is available on GitHub via our [Computational Endoscopy Platform](#). The contributions of this work are as follows:

1. A method for haustral fold detection and segmentation in real OC images.
2. A semi-supervised approach to perform image-to-image domain translation via a common domain.
3. A transitive loss to drive the domain translation while retaining/preserving haustral fold features in the common VC domain.

2 Related Works

Deep learning approaches have recently shown promising results in various endoscopic intervention tasks such as depth estimation, 3D reconstruction, and surface coverage [2,6,11,14]. Deep learning models are data driven and the supervised category requires ground truth information. The issue with the supervised approaches, specifically for colonoscopy, is the need for realistic training data that models the specular reflections, fluid motion, and erratic camera and spasm movements. The ground truth data creation is extremely time consuming and even then the network can easily fail since the input data is not representative of the real domain. To overcome this, the supervised methods require additional measures to handle real data. Rau et al. [18] trained a pix2pix variant on textured VC data which needed extra training on OC images to handle real OC data. Mahmood et al. [13] required an additional network to transform their data into a synthetic-like (VC) domain before being able to estimate depth maps. Chen et al. [3] trained on VC input and tested on porcine phantom models, so it is unclear how well their approach performs on real OC data.

In contrast, unsupervised approaches already work on the input domain, so no additional modules are needed to handle real OC images. They, however, require more complicated network design and careful assessment to avoid overfitting. In FoldIt, we present a semi-supervised approach which leverages real OC images. Ground truth haustral fold annotations on 3D VC triangular mesh models were computed using Fiedler vector representation [17] and verified by a clinician. We capture the best aspects of both the unsupervised and supervised approaches in our FoldIt model.

Generative Adversarial Networks (GAN) [7] have shown promising results for image-to-image translation. GANs help the generator output match the training data distribution. Approaches such as pix2pix [10], CycleGAN [20], and StarGAN [4] all utilize adversarial learning in their approaches. Pix2pix does paired image-to-image translation while CycleGAN can handle unpaired data. Both pix2pix and CycleGAN assume a one-to-one mapping between domains. Mathew et al. [15] introduced a one-to-many mapping to handle OC images since there are large variations in texture, lighting, and specular reflections.

Haustral fold annotation requires many-to-many mapping since the depth of the folds can be unclear from a single frame. Travel-GAN [1] introduced a many-to-many approach, however, their translation is not constrained, making it less suitable for our task. To address this, we present a model that performs many-to-many translation via a common domain. This common domain has a one-to-many mapping with the other two domains to help constrain the output. Fang et al. [5] have recently proposed a triple translation loss, similar to our transitive loss, but they focus on face generation based on age in a single domain which does not require translation between multiple domains.

3 Data

The OC and VC data were obtained from 10 patients who underwent a VC procedure followed by an OC procedure. Data from 7 patients were used for training and 3 for testing. The OC frames are cropped and rescaled to 256x256 to remove the borders in the frame. VC meshes are synthesized from abdominal CT scans using a pipeline described in [16]. These meshes do not align with the OC videos, as the shape of the colon changes between the two procedures. Flythroughs are created in the VC mesh along the centerline with random camera rotations, and 2 light sources on left and right side of the camera to replicate the OC procedure. These lights follow the inverse square fall-off property to create realistic lighting [13]. The haustral folds on VC meshes were segmented using Fiedler vector and corresponding levelset representations, as outlined in [17], and verified by a clinician. The segmented folds were overlaid on the VC mesh rendering and used as haustral fold annotation frames. Each video was split into 300 frames for a total of 3000 frames. Example frames from each domain are shown in Fig. 1b.

4 Methods

FoldIt uses 4 generators and discriminators. G_{AB} translates images from domain A to domain B , and G_{BA} acts in the opposite direction. G_{AC} and G_{BC} translate from input domains, A and B , to the common domain, C . Our approach can be broken down into four losses: adversarial losses, transitive losses, a ground truth loss, and identity losses. The complete objective function is as follows:

$$\mathcal{L}_{obj} = \lambda_{adv} \mathcal{L}_{adv} + \lambda_T \mathcal{L}_T + \lambda_{GT} \mathcal{L}_{GT} + \lambda_{idt} \mathcal{L}_{idt} \quad (1)$$

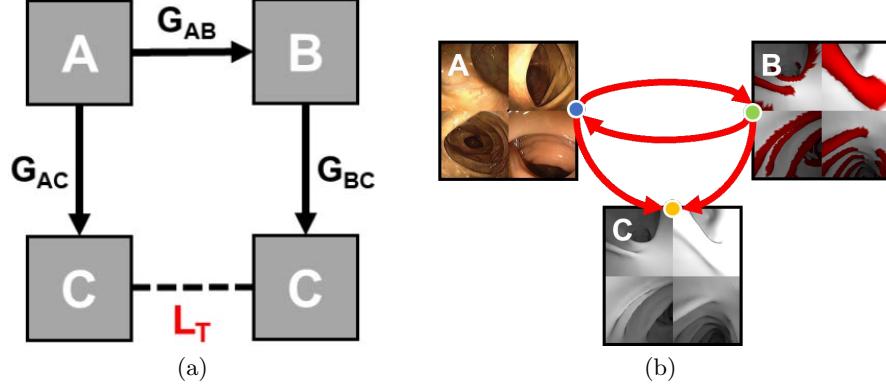


Fig. 1. (a) An image from domain A is converted to domain B and then C through generators G_{AB} and G_{BC} . This resulting image is compared against the image directly translated to domain C via G_{AC} to compute the transitive loss. (b) The relationship between the domains A , B , and C . We can translate between domains A and B , however, the result in domain C must remain the same. For FoldIt, A is the OC domains, B is the Haustral Fold annotation (red overlay) domains, and C is the VC domain.

Each of the 4 generators have an associated discriminator, D , which applies an adversarial loss to ensure the output image encompasses the features of the output domain. The adversarial loss for each of these GANs is described as:

$$\mathcal{L}_{GAN}(G, D, A, B) = \mathbb{E}_{y \sim p(B)} [\log(D(y))] + \mathbb{E}_{x \sim p(A)} [\log(1 - D(G(x)))] \quad (2)$$

where $y \sim p(B)$ represents the output domain B data distribution and $x \sim p(A)$ represents the input domain A data distribution. An adversarial loss is applied to each generator to create the adversarial component of the objective loss:

$$\begin{aligned} \mathcal{L}_{adv} = & \mathcal{L}_{GAN}(G_{AB}, D_{AB}, A, B) + \mathcal{L}_{GAN}(G_{BA}, D_{BA}, B, A) \\ & + \mathcal{L}_{GAN}(G_{AC}, D_{AC}, A, C) + \mathcal{L}_{GAN}(G_{BC}, D_{BC}, B, C) \end{aligned} \quad (3)$$

Cycle consistency losses have shown great results for the task of image-to-image translation. Our approach uses a modified cycle consistency loss, which resembles the transitive property, to learn the domain translation via a common domain. The main insight here is that the translation to the common domain, C , should be consistent between our domains A and B . Here, C is the VC domain. When we do the translation, $A \rightarrow B \rightarrow C$, the result should be the same as $A \rightarrow C$ (see Fig. 1a). We can express this transitive loss as ($\|\cdot\|_1$ is ℓ_1 norm):

$$\mathcal{L}_T(G_{AB}, G_{BC}, G_{AC}, A) = \mathbb{E}_{x \sim p(A)} \|G_{BC}(G_{AB}(x) - G_{AC}(x))\|_1 \quad (4)$$

This loss is applied on both directions in the total objective loss.

$$\mathcal{L}_T = \mathcal{L}_T(G_{AB}, G_{BC}, G_{AC}, A) + \mathcal{L}_T(G_{BA}, G_{AC}, G_{BC}, B) \quad (5)$$

\mathcal{L}_{GT} is our ground truth loss, meant to utilize the ground truth pairing between the common domain and one of the domains A and B . It simply applies the ℓ_1 norm between the paired domains. Here, we have ground truth correspondence between haustral fold annotations and VC. The ground truth loss is:

$$\mathcal{L}_{GT}(G_{BC}, B, C) = \mathbb{E}_{x, z \sim p(B, C)} \|(G_{BC}(x) - z)\|_1 \quad (6)$$

where $x, z \sim p(A, C)$ represents the paired data distribution (A, C) .

Lastly, we have the identity loss, which is meant for additional stability during training as described in [15,20]. An image from the output domain is expected to be unchanged when passed as input through the network under an identity loss. This is applied only to G_{AC} and G_{BC} as alterations in the output for G_{AB} and G_{BA} may still be valid. The identity loss is as follows:

$$\mathcal{L}_{idt}(G_{AC}, C) = \mathbb{E}_{z \sim p(C)} \|(G_{AC}(z) - z)\|_1 \quad (7)$$

The complete identity loss is:

$$\mathcal{L}_{idt} = \mathcal{L}_{idt}(G_{AC}, C) + \mathcal{L}_{idt}(G_{BC}, C) \quad (8)$$

Each generator used up 18 MB for its ResNet architecture [8] with 9 blocks. We use a PatchGAN discriminator [10] similar to [20], which used 3MB each. The network was trained for 100 epochs and used the following weights: $\lambda_{adv} = 1, \lambda_T = 10, \lambda_{GT} = 1, \lambda_{idt} = 1$. The inference time for a single image is 0.04 seconds on an Nvidia RTX 6000 GPU.

5 Results

There is no ground truth correspondence between OC and VC due to the colon shape changes between the two procedures making quantitative analysis difficult. We use textured VC mesh with clinician-verified haustral fold annotations to evaluate our results. We render the VC mesh with two different textures to test the sensitivity of our model to texture and lighting variations. The haustral folds were segmented and overlaid on the VC rendering. Fig. 2 shows results from our FoldIt versus XDCycleGAN [15] on these textured VC meshes.

FoldIt results are closer to the ground truth for both textures, while the haustral fold features remain preserved/consistent across frames. The green bounding boxes indicate locations where XDCycleGAN is not feature consistent and removes/adds folds between neighboring frames. As seen in Table 1, FoldIt achieves higher Dice and IoU scores when compared with XDCycleGAN on both textures. Since the textured VC colons have the same underlying geometry, the network should have similar predictions for both videos. The Dice and IoU scores are calculated between Texture 1 and Texture 2 and shown in the third row of Table 1. Again, FoldIt is more consistent in its predictions with varying textures. The complete video sequences are provided in the **supplementary video**³.

³ Supplementary Video: https://youtu.be/_iWBJnDMXjo

Table 1. Dice and IoU scores on ground truth VC model with the two textures shown in Fig. 2. Dice and IoU scores between Texture 1 and 2 is shown on the third row.

	Dice		IoU	
	FoldIt	XDCycleGAN	FoldIt	XDCycleGAN
Texture 1	0.47 ± 0.11	0.25 ± 0.11	0.31 ± 0.09	0.15 ± 0.07
Texture 2	0.50 ± 0.10	0.21 ± 0.10	0.33 ± 0.09	0.12 ± 0.06
Consistency	0.77 ± 0.10	0.64 ± 0.16	0.64 ± 0.12	0.49 ± 0.16

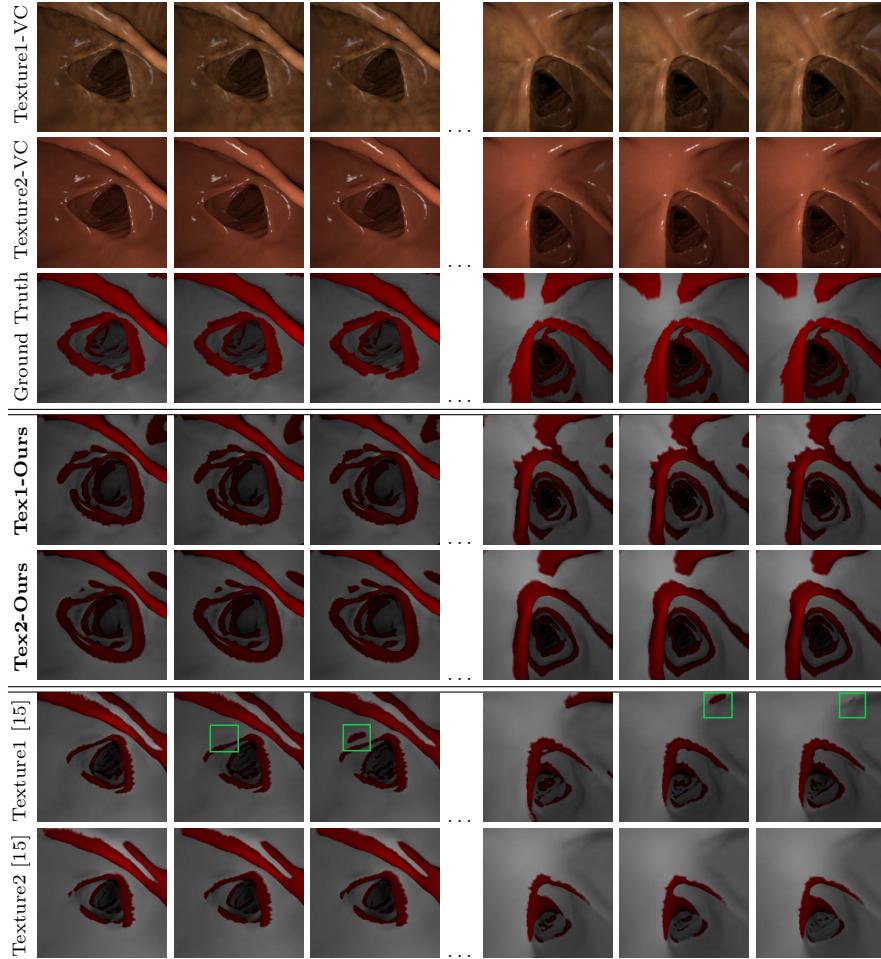


Fig. 2. Evaluation of FoldIt (ours) and XDCycleGAN model [15] on VC with two different textures and ground truth/clinician-verified haustral fold (red) overlays. Green bounding boxes indicate locations where XDCycleGAN is not feature consistent and drops folds between neighboring frames. Complete video sequences are provided in the **supplementary video**.

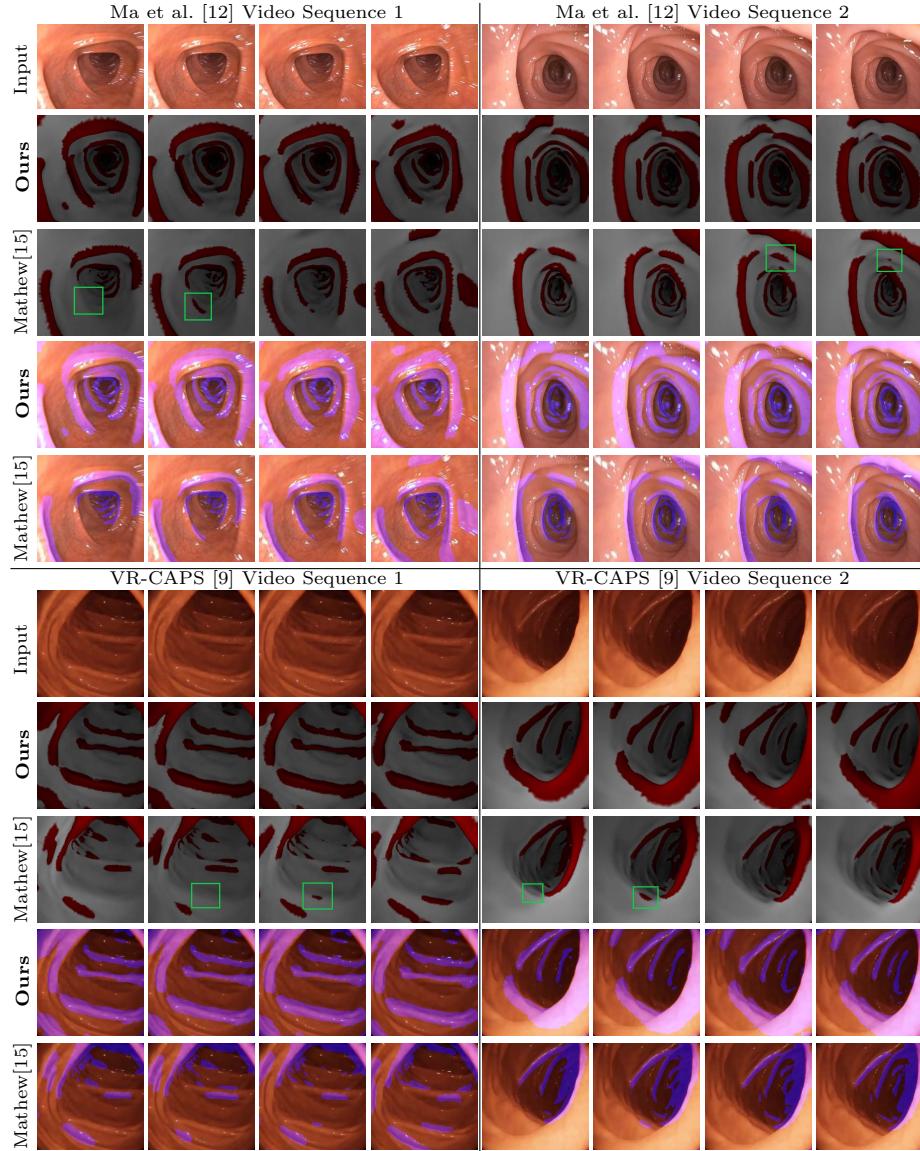


Fig. 3. Results of FoldIt (ours) and XDCycleGAN model [15] are shown on video sequences from Ma et al. [12] and recently released VR-CAPS colon simulator [9]. Complete video sequences are provided in the **supplementary video**. The OC frames are followed by FoldIt and XDCycleGAN output and the fold segmentations overlaid on the OC images. Green bounding boxes indicate examples of XDCycleGAN losing haustral folds information across frames whereas our FoldIt preserves these features.

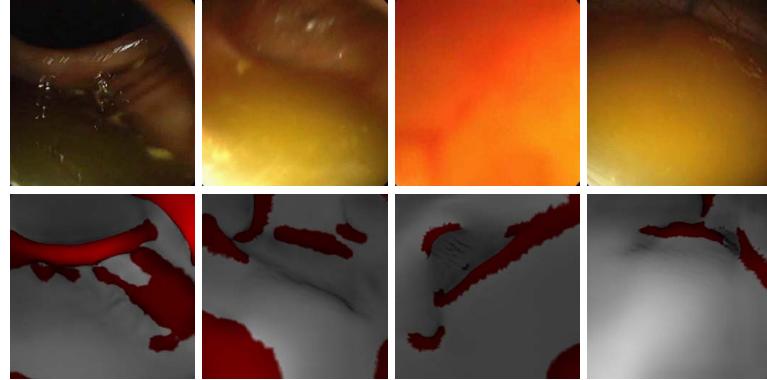


Fig. 4. Examples of our network failure in blurry and fluid-motion frames.

We also show results in Fig. 3 for 11 real OC video sequences from Ma et al. [12] and 2 from recently-released VR-CAPS simulator [9]. For both, FoldIt and XDCycleGAN, we show an overlay view, where the fold segmentation is extracted from the network output and superimposed in blue on the OC input. While XDCycleGAN is capable of doing scale-consistent depth maps for these video sequences, it struggles to retain haustral fold features across frames, as shown by the green bounding boxes and results in flickering in the final output (as shown in the **supplementary video**⁴). FoldIt delineates the folds more accurately and is more consistent in preserving fold features across frames (and does not result in the flickering effect seen with XDCycleGAN).

6 Limitations and Future Work

In this work, we introduced FoldIt, a new model to detect and segment haustral folds in colonoscopy videos. Our model has some limitations. It fails to handle cases where there are large amounts of fluid or blurriness, as shown in Fig. 4. This is due to the fact that our network is trained on clean VC images where there is no fluid or blur. In effect, these types of OC frames do not have valid corresponding frames in our VC training data. In the future, fluid and blur can be introduced in our VC renderings to add a valid correspondence for these failed OC frames. Moreover, even though our current model gives consistent results across video frames, we have not incorporated any explicit temporal constraint. In the future, we will incorporate temporal consistency in our model to track haustral folds across frames; the tracked folds will be used to register OC and VC videos. Finally, we will use additional anatomical features such as taenia coli (three longitudinal muscles that run across colon) to further enrich our model.

⁴ Supplementary Video: https://youtu.be/_iWBJnDMXjo

Acknowledgements

This project was supported by MSK Cancer Center Support Grant/Core Grant (P30 CA008748), and NSF grants CNS1650499, OAC1919752, and ICER1940302.

References

1. Amodio, M., Krishnaswamy, S.: Travelgan: Image-to-image translation by transformation vector learning. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8983–8992 (2019)
2. Bae, G., Budvytis, I., Yeung, C.K., Cipolla, R.: Deep multi-view stereo for dense 3D reconstruction from monocular endoscopic video. International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 774–783 (2020)
3. Chen, R.J., Bobrow, T.L., Athey, T., Mahmood, F., Durr, N.J.: Slam endoscopy enhanced by adversarial depth prediction. arXiv preprint arXiv:1907.00283 (2019)
4. Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8789–8797 (2018)
5. Fang, H., Deng, W., Zhong, Y., Hu, J.: Triple-GAN: Progressive face aging with triple translation loss. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 804–805 (2020)
6. Freedman, D., Blau, Y., Katzir, L., Aides, A., Shimshoni, I., Veikherman, D., Golany, T., Gordon, A., Corrado, G., Matias, Y., Rivlin, E.: Detecting deficient coverage in colonoscopies. arXiv preprint arXiv:2001.08589 (2020)
7. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in Neural Information Processing Systems. pp. 2672–2680 (2014)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
9. İncetan, K., Celik, I.O., Obeid, A., Gokceler, G.I., Ozyoruk, K.B., Almalioglu, Y., Chen, R.J., Mahmood, F., Gilbert, H., Durr, N.J., et al.: VR-Caps: A virtual environment for capsule endoscopy. Medical Image Analysis. p. 101990 (2021)
10. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition pp. 1125–1134 (2017)
11. Liu, X., Stiber, M., Huang, J., Ishii, M., Hager, G.D., Taylor, R.H., Unberath, M.: Reconstructing sinus anatomy from endoscopic video—towards a radiation-free approach for quantitative longitudinal assessment. International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 3–13 (2020)
12. Ma, R., Wang, R., Pizer, S., Rosenman, J., McGill, S.K., Frahm, J.M.: Real-time 3D reconstruction of colonoscopic surfaces for determining missing regions. International Conference on Medical Image Computing and Computer-Assisted Intervention pp. 573–582 (2019)
13. Mahmood, F., Chen, R., Durr, N.J.: Unsupervised reverse domain adaptation for synthetic medical images via adversarial training. IEEE Transactions on Medical Imaging **37**(12), 2572–2581 (2018)

14. Mathew, S., Nadeem, S., Kaufman, A.: Visualizing missing surfaces in colonoscopy videos using shared latent space representations. IEEE 18th International Symposium on Biomedical Imaging (ISBI). pp. 329–333 (2021)
15. Mathew, S., Nadeem, S., Kumari, S., Kaufman, A.: Augmenting colonoscopy using extended and directional cyclegan for lossy image translation. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4696–4705 (2020)
16. Nadeem, S., Kaufman, A.: Computer-aided detection of polyps in optical colonoscopy images. SPIE Medical Imaging **9785**, 978525 (2016)
17. Nadeem, S., Marino, J., Gu, X., Kaufman, A.: Corresponding supine and prone colon visualization using eigenfunction analysis and fold modeling. IEEE Transactions on Visualization and Computer Graphics **23**(1), 751–760 (2016)
18. Rau, A., Edwards, P.E., Ahmad, O.F., Riordan, P., Janatka, M., Lovat, L.B., Stoyanov, D.: Implicit domain adaptation with conditional generative adversarial networks for depth prediction in endoscopy. International Journal of Computer Assisted Radiology and Surgery **14**(7), 1167–1176 (2019)
19. Xu, J., Anwar, S., Barnes, N., Grimpenn, F., Salvado, O., Anderson, S., Armin, M.A.: Ofgan: Realistic rendition of synthetic colonoscopy videos. International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 732–741 (2020)
20. Zhu, J.Y., Park, T., Isola, P., Efros, A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. IEEE Int Conf Computer Vision. pp. 2223–2232 (2017)