

High-Quality Medical Image Generation from Free-hand Sketch

Quan Huu Cap, Atsushi Fukuda

AI Development Department, Aillis, Inc., Tokyo, Japan

Email: {quan.cap, atsushi.fukuda}@aillis.jp

Abstract—Generating medical images from human-drawn free-hand sketches holds promise for various important medical imaging applications. Due to the extreme difficulty in collecting free-hand sketch data in the medical domain, most deep learning-based methods have been proposed to generate medical images from the synthesized sketches (e.g., edge maps or contours of segmentation masks from real images). However, these models often fail to generalize on the free-hand sketches, leading to unsatisfactory results. In this paper, we propose a practical free-hand sketch-to-image generation model called Sketch2MedI that learns to represent sketches in StyleGAN’s latent space and generate medical images from it. Thanks to the ability to encode sketches into this meaningful representation space, Sketch2MedI only requires synthesized sketches for training, enabling a cost-effective learning process. Our Sketch2MedI demonstrates a robust generalization to free-hand sketches, resulting in high-quality and realistic medical image generations. Comparative evaluations of Sketch2MedI against the pix2pix, CycleGAN, UNIT, and U-GAT-IT models show superior performance in generating pharyngeal images, both quantitative and qualitative across various metrics.

Index Terms—medical image generation, free-hand sketch-to-image, pharyngeal images, generative adversarial networks.

I. INTRODUCTION

Human-drawn free-hand sketches are an effective way of visualizing target objects or concepts. In the domain of medical imaging, the synthesis of medical images from sketches offers potential for various important applications. One such application is for patient communication, where physicians can interactively explain conditions or medical concepts to patients through the generated images from free-hand sketches. Furthermore, sketch-to-image generation enables the application of sketch-based medical image retrieval. This approach allows medical professionals to sketch and retrieve images based on their partial memory, enhancing the efficiency of image retrieval processes.

In recent years, many deep learning-based methods have been proposed to generate images from free-hand sketches and achieved remarkable results [1–4]. However, unlike general computer vision tasks, collecting medical images and the corresponding sketches to build such models is very difficult and time-consuming since it requires a lot of expert knowledge as well as drawing effort. To mitigate this issue, some studies proposed using synthesized sketches such as edge maps [5–7] or semantic contours (i.e., boundaries of ground-truth semantic masks) [8] from the target photos as pseudo sketch-

replacement for model training. However, human-drawn free-hand sketches of objects are usually in different levels of abstractions and deformations, whereas synthesized sketches ideally align with target image boundaries. Because of this large domain difference, images generated from free-hand sketches by those models easily follow the abstract boundaries, yielding unrealistic results. Moreover, some methods are data-dependent and thus, cannot be applied to other modalities such as medical data [7, 9]. We hypothesize that if there is a way to represent the synthesized sketches in a more abstract manner, the generated results from free-hand sketches could be significantly improved.

It has been well-known that the latent space in StyleGAN [10] expresses semantically rich representations, allowing different levels of abstractions to be manipulated. In recent years, many studies have proposed encoding real images into this latent space to facilitate diverse and expressive image manipulations [11–14]. Inspired by these works, in this paper, we propose an encoder-decoder-based model called Sketch2MedI to encode sketches into StyleGAN’s meaningful latent space and subsequently, generate medical images from this manifold to ensure photorealistic quality. To enable cost-effective learning, we propose using a sketch synthesizer which is a trained segmentation model to synthesize sketch data (i.e., predicting semantic contours) from real medical images to train our Sketch2MedI model. Despite training entirely on the synthesized sketch data, we demonstrate that our proposal generalizes very well in generating pharyngeal images from human-drawn free-hand sketches. In comparison to the pix2pix [5], CycleGAN [15], UNIT [16], and U-GAT-IT [17] models, the generated results by our Sketch2MedI showed substantial improvements, both quantitatively and qualitatively.

II. RELATED WORK

In medical imaging, several studies have been proposed utilizing sketch to image generation [18–20]. Zhang *et al.* [18] proposed generating fake edge maps from noises and using them to generate various medical image modalities. However, using edge maps differs largely from free-hand sketches. Moreover, this system lacks user control as the generated results are from random noises. Liang *et al.* [19] designed a sketch-guided model for realistic and editable ultrasound image synthesis. Nevertheless, edge maps from real images must be extracted beforehand to synthesize the results. Furthermore, training their model requires segmenta-

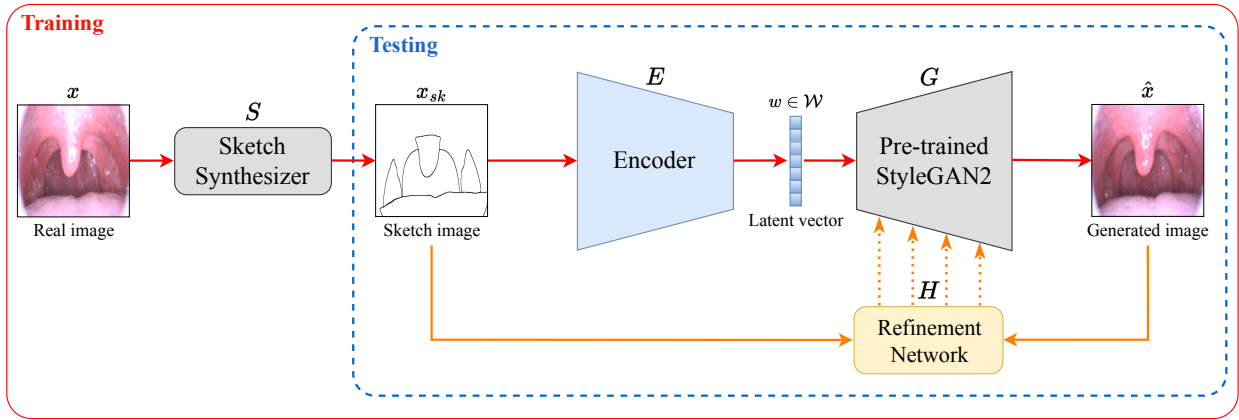


Fig. 1: The overview of our proposed Sketch2MedI. At test time, we only input the human-drawn free-hand sketches into the model (ignore using S).

tion annotations for all images, resulting in substantial costs. In contrast, we train our Sketch2MedI using sketch data synthesized through a segmentation model that was trained on a subset of the dataset, providing a more cost-effective option. Toda *et al.* [20] trained a model to generate lung cancer computed tomography (CT) images from edge maps for data augmentation. Despite showing promising results on free-hand sketches, the practicality of their model remains uncertain since at test time, each sketch needs to be carefully adjusted with visual feedback from physicians to avoid missing relevant information in the drawing, thus introducing complexity to the process. Moreover, their model was evaluated on a notably limited dataset with only 20 images. Therefore, the model’s generalizability on diverse styles of free-hand sketches remains unknown.

III. PROPOSED METHOD

A. Overview

Our work is built on the idea of StyleGAN inversion [11] where real images are encoded into the expressive and meaningful latent space \mathcal{W} for image manipulation. In StyleGAN, the latent vector $w \in \mathcal{W}$ is mapped from a random noise vector $z \in Z = \mathcal{N}(\mu, \sigma^2)$ via a non-linear mapping function. The main objective of this study is to encode sketch images into this \mathcal{W} space to enable realistic medical image generation.

We designed our Sketch2MedI model as an encoder-decoder network to generate pharyngeal images from human-drawn free-hand sketches. Fig. 1 illustrates the overview of our proposal. Given the real pharyngeal image x , the sketch synthesizer S , and the pre-trained StyleGAN2 [21] generator G , we first apply S to synthesize the grayscale sketch image $x_{sk} = S(x)$. Then, the encoder E encodes x_{sk} into the \mathcal{W} space to obtain the latent vector $w = E(x_{sk}) \in \mathcal{W}$. From this, the generator G is used to generate the output image $\hat{x} = G(w)$.

In general, using the encoder E alone (i.e., without H employed) is sufficient to generate realistic images. However, those generated results occasionally are not well-aligned with

pharyngeal image object boundaries. To further improve the generated quality, the refinement network H takes the x_{sk} and \hat{x} images as inputs, and then gradually adjusts the weights of the generator G in an *iterative* manner, refining the quality of the output \hat{x} .

We train our Sketch2MedI model in two phases. In the first phase, the encoder E is trained without employing the refinement network H . In the second phase, the refinement network H is trained with the encoder E frozen. Note that the S and G networks remain frozen during both phases. At test time (Fig. 1 (dashed blue box)), we instead use the human-drawn free-hand sketches as inputs to our Sketch2MedI model to generate pharyngeal images (ignore using S).

B. The sketch synthesizer

Since collecting medical images and their corresponding free-hand sketches is a challenging task, we propose the sketch synthesizer S , which is a trained instance segmentation model to instead synthesize sketch data (i.e., semantic contours) from real images for training the Sketch2MedI. We chose to utilize semantic contours since they are less noisy than edge maps and have a closer appearance to free-hand sketches.

Specifically, S predicts the segmentation masks from a real pharyngeal image x , the synthesized grayscale sketch x_{sk} is then formed by extracting the object contours from all predicted instance masks. Our sketch synthesizer S is based on the SOLOv2 model [22] and was trained on a subset of the dataset. Compared to collecting medical free-hand sketches, using S to synthesize training data enables much more cost-effective learning.

C. The encoder

The encoder E is designed to encode the sketch x_{sk} into StyleGAN’s latent space \mathcal{W} to enable realistic medical image generation. We adopt the ResNet50-based [23] encoder from [12] which takes x_{sk} as input and outputs a latent vector $w \in \mathbb{R}^{512}$. A similarity loss is employed to train E , which ensures low distortion between real image x and generated image \hat{x} . Following [12], the similarity loss \mathcal{L}_{sim} includes a

cosine similarity loss \mathcal{L}_{cos} , a L_2 loss, and an LPIPS perceptual loss $\mathcal{L}_{\text{LPIPS}}$ [24]. The \mathcal{L}_{sim} is as follows:

$$\mathcal{L}_{\text{sim}} = \lambda_{l_2} L_2(x, \hat{x}) + \lambda_{l_{\text{pips}}} \mathcal{L}_{\text{LPIPS}}(x, \hat{x}) + \lambda_{\text{cos}} \mathcal{L}_{\text{cos}}(x, \hat{x}), \quad (1)$$

where λ_{l_2} , $\lambda_{l_{\text{pips}}}$, λ_{cos} are the loss coefficients. The objective of \mathcal{L}_{cos} is to minimize the cosine similarity between the feature embeddings of x and \hat{x} . The \mathcal{L}_{cos} is defined as:

$$\mathcal{L}_{\text{cos}} = 1 - \cos(C(x), C(\hat{x})), \quad (2)$$

where C is a pre-trained self-supervised MOCOv2 model [25].

In addition, the encoder E is trained with an adversarial loss [26] that encourages the encoded latent vector w to lie close to the latent space \mathcal{W} of StyleGAN [12]. A 4-layer fully connected network from [27] is applied as the latent discriminator D to discriminate between real samples of the \mathcal{W} space (generated from noise vector $z \in Z$) and the encoder's outputs. The adversarial loss is the GAN loss [26] with R_1 regularization [28]:

$$\mathcal{L}_{\text{adv}}^D = -\mathbb{E}_{w \sim \mathcal{W}}[\log D(w)] - \mathbb{E}_{x \sim p(x)}[\log(1 - D(E(x)))] + \frac{\gamma}{2} \mathbb{E}_{w \sim \mathcal{W}}[\|\nabla_w D(w)\|_2^2], \quad (3)$$

$$\mathcal{L}_{\text{adv}}^E = -\mathbb{E}_{x \sim p(x)}[\log D(E(x))]. \quad (4)$$

Finally, the total objective loss \mathcal{L}_E for training the encoder E is:

$$\mathcal{L}_E = \mathcal{L}_{\text{sim}} + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}}^E. \quad (5)$$

Here, γ and λ_{adv} are the loss coefficients. For more details on the architecture and implementation of the encoder E , please refer to the literature [12].

D. The refinement network

We build our refinement network H based on the hyper-network from [13]. The task of H is to further improve the quality of the generated image \hat{x} in an *iterative* manner. The network H receives a 4-channel image (concatenated x_{sk} and \hat{x}) and predicts a set of weight offsets with respect to the weights of the pre-trained StyleGAN2 [21] generator G . We perform several steps through H for a single sketch x_{sk} (e.g., T steps). Each refinement step allows H to gradually refine its predicted weight offsets, helping the generator G to produce a more precise result.

Let $G(w; \theta)$ be the output image of G parameterized by weights θ . For each sketch input x_{sk} , we first obtain the generated image $\hat{x}_0 = G(w; \theta)$. At each refinement step $1 \leq t \leq T$, the network H predicts a set of weight offsets $\Delta_t = H(x_{sk}, \hat{x}_{t-1})$. The new weights $\hat{\theta}_t$ of G at step t are the accumulated modulation across all previous steps [13]:

$$\hat{\theta}_t := \theta \cdot \left(1 + \sum_{i=1}^t \Delta_i \right). \quad (6)$$

The new generated image is updated by $\hat{x}_t = G(w; \hat{\theta}_t)$ and is used to predict weight offsets at the next step. We apply the same \mathcal{L}_{sim} loss in Eq. (1) to train H , and the loss

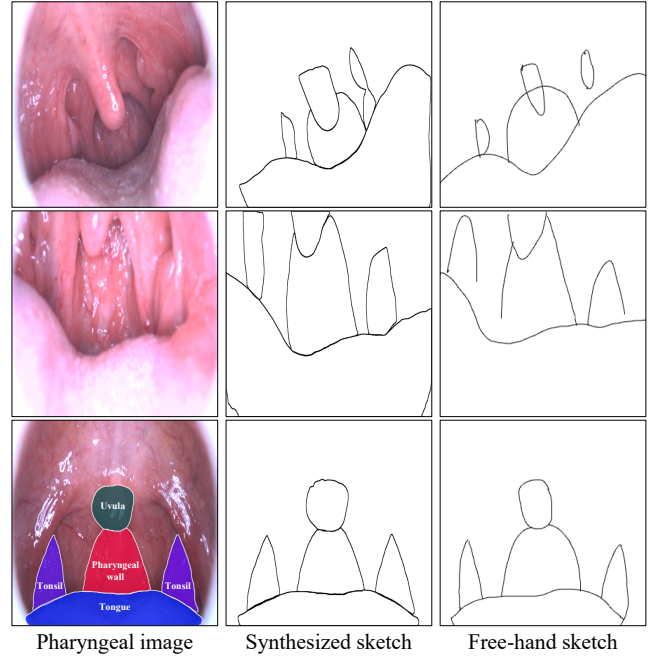


Fig. 2: Samples from the pharyngeal, synthesized sketch, and free-hand sketch datasets. An example of the predicted masks of four pharyngeal objects (i.e., uvula, pharyngeal wall, tonsil, and tongue) is illustrated in the bottom-left image. The sketch images depict these objects.

is calculated at each refinement step. Again, note that the S , E , and G networks are frozen during training. The total refinement step for each sketch input is set to $T = 3$ in both training and testing. At test time, the final generated image is obtained at the last step. For more details on the architecture and implementation of the refinement network H , please refer to the literature [13].

IV. EXPERIMENTAL SETTINGS

A. Datasets

In this study, we analyzed pharyngeal images and created various datasets. Fig. 2 illustrates some samples from these datasets. The specific of each dataset is described as follows:

1) *Pharyngeal dataset*: We collected a large number of high-quality 60,000 pharyngeal RGB images from over 5,000 patients. These images were acquired using a camera specifically designed for taking pharyngeal images. Eventually, all images were resized to 256×256 pixels.

2) *Segmentation dataset*: Another 2,000 pharyngeal images were annotated with segmentation annotations by experienced experts. The labels include four classes namely: *uvula*, *pharyngeal wall*, *tonsil*, and *tongue*. From this data, 1,600 and 400 images were used for training and validating the sketch synthesizer S , respectively.

3) *Synthesized sketch dataset*: After the training, S is applied to synthesize 60,000 corresponding sketch images from the above pharyngeal dataset. Given the predicted segmentation masks from a pharyngeal image with confidence

TABLE I: Evaluation results of quality and originality aspects on generated images from the free-hand sketch dataset

Method	Quality aspect				Originality aspect			
	FID (\downarrow)	MOS* (\uparrow)	TReS (\uparrow)	CLIP-IQA (\uparrow)	DICE (\uparrow)	LPIPS (\downarrow)	DISTS (\downarrow)	STLPIPS (\downarrow)
pix2pix [5]	105.02	2.77 \pm 1.10	65.28 \pm 6.07	0.52 \pm 0.06	46.58	0.42 \pm 0.06	0.24 \pm 0.03	0.36 \pm 0.07
CycleGAN [15]	87.06	1.99 \pm 0.96	59.78 \pm 15.44	0.46 \pm 0.12	38.19	0.48 \pm 0.11	0.28 \pm 0.05	0.41 \pm 0.10
UNIT [16]	117.80	2.18 \pm 0.92	63.27 \pm 7.08	0.39 \pm 0.09	48.39	0.50 \pm 0.09	0.27 \pm 0.03	0.40 \pm 0.09
U-GAT-IT [17]	118.83	1.42 \pm 0.70	50.53 \pm 9.38	0.36 \pm 0.09	19.03	0.58 \pm 0.12	0.33 \pm 0.06	0.47 \pm 0.10
Sketch2MedI w/o H (proposed)	38.39	N/A	76.91 \pm 6.47	0.62 \pm 0.07	56.53	0.36 \pm 0.05	0.20\pm0.03	0.30 \pm 0.07
Sketch2MedI (proposed)	40.83	3.99\pm0.98	78.04\pm6.28	0.63\pm0.07	58.28	0.35\pm0.05	0.20\pm0.03	0.29\pm0.07

* The MOS of the ‘‘Sketch2MedI w/o H ’’ model is unavailable since it is part of the proposal.

† **Bold** text indicates the best performance.

scores $\geq \delta$, a synthesized sketch is formed by extracting the object contours from all instance masks. We set $\delta = 0.35$ in our work.

4) *Free-hand sketch dataset*: The free-hand sketches are drawn on blank backgrounds by observing the reference pharyngeal images. Here, 12 users (exclusive from the segmentation annotators) were asked to draw the boundaries of four pharyngeal class objects as in the segmentation dataset (i.e., uvula, pharyngeal wall, tonsil, and tongue), resulting in 495 sketch images. Note that the reference pharyngeal images used to create this dataset are exclusive from the pharyngeal dataset.

Typically, free-hand sketches contain different levels of abstractions and drawing styles from multiple users. Thus, these sketches are usually misaligned with the objects of the referenced pharyngeal images (see Fig. 2).

B. Evaluation methods

We evaluate the generated results from the free-hand sketch dataset in two aspects: *quality* (e.g., clearness, naturalness of outputs) and *originality* (e.g., structure similarity between outputs and reference pharyngeal images), similar to [29]. To measure the quality aspect, we used the subjective mean opinion score (MOS) and non-reference image quality assessment (NR-IQA) criteria, including FID [30], TReS [31], and CLIP-IQA [32].

For the originality aspect, we measured the mean DICE score between predicted segmentation masks of the reference image x and the generated image \hat{x} by applying the trained sketch synthesizer S . Along with that, we also used the full-reference image quality assessment (FR-IQA) criteria, including LPIPS [24], DISTS [33], and STLPIPS [34]. Note that the *originality* aspect in this case is not ideal since free-hand sketches are not precisely aligned with the reference pharyngeal object boundaries, however, we still consider this a practical scenario.

C. Training details

1) *The sketch synthesizer*: The synthesizer S is based on the SOLOv2 instance segmentation model [22]. We trained S on the training set of the segmentation dataset (i.e., 1,600 images) using the default hyperparameters of SOLOv2 with ResNet101 [23] backbone. After the training, we achieved the mean DICE score of 86.49 on the validation set.

2) *The generator*: The pharyngeal dataset is used to train our generator G . The output size of G is set to 256×256 pixels. We applied training with the default configurations from the StyleGAN2 literature [21] and finished the training after 100 epochs.

3) *The encoder and refinement networks*: The pharyngeal and synthesized sketch datasets are used to train both E and H . We trained E together with the latent discriminator D in an adversarial manner using the losses in Eq. (3) and (5), while H is trained with the loss from Eq. (1). We set the coefficients of those losses to $\lambda_{l_2} = 1.0$, $\lambda_{l_{lips}} = 0.8$, $\lambda_{cos} = 0.5$, $\gamma = 10$, and $\lambda_{adv} = 0.1$. Note here again that H is trained after we have done training E and both models were trained for 100 epochs.

4) *Comparison models*: For comparison purposes, we also trained four other sketch-to-image generation models using pix2pix [5], CycleGAN [15], UNIT [16], and U-GAT-IT [17]. We trained these models on the same pharyngeal and synthesized sketch datasets with their default configurations. All models were trained for 100 epochs.

V. RESULTS AND DISCUSSION

For the subjective MOS evaluation, 20 users participated in scoring the quality (e.g., clearness, naturalness) of the generated images from pix2pix, CycleGAN, UNIT, U-GAT-IT, and our Sketch2MedI models on a scale of 1 to 5 (bad to excellent). Because of the time-consuming process, we randomly selected 80 images from the free-hand sketch dataset for the MOS test. Subsequently, each user evaluated a total of 400 images, and the final score for each model was averaged across all participants. Table I shows the numerical evaluations of the quality and originality aspects. To confirm the effect of the refinement network H , we also included the results of our proposal without using H , noted as ‘‘Sketch2MedI w/o H ’’ in Table I. Here, for all cases, we refer to ‘‘Sketch2MedI’’ as our complete model with both E and H employed. Overall, the use of network H is effective in improving the quality and originality aspects. Visual comparisons of the generated images from pix2pix, CycleGAN, UNIT, U-GAT-IT, and our Sketch2MedI are shown in Fig. 3.

Due to the large distribution shift between the synthesized and free-hand sketches, generated images by pix2pix, CycleGAN, UNIT, and U-GAT-IT are often deformed and

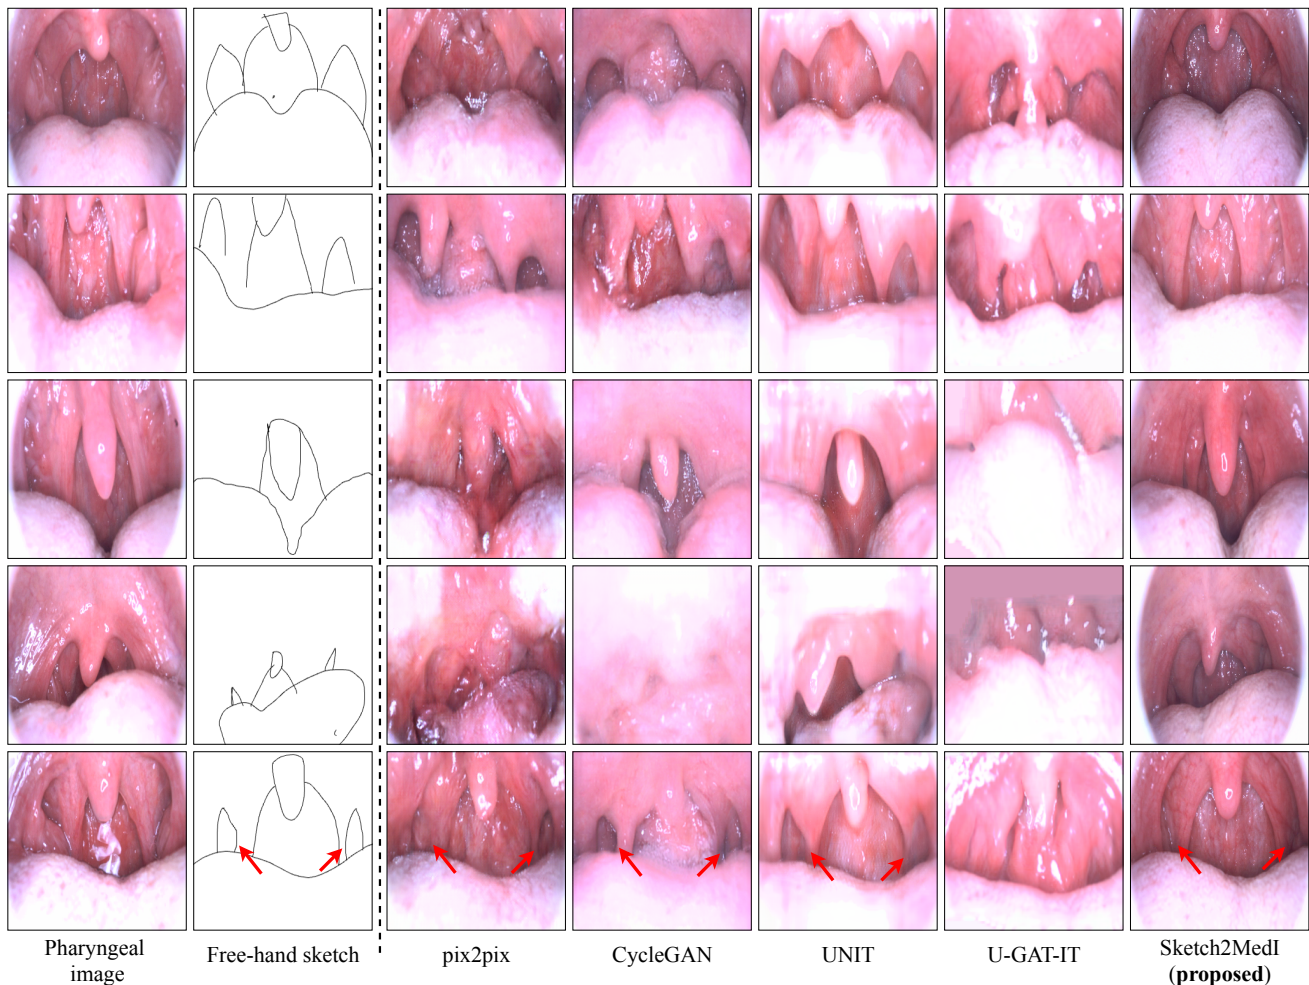


Fig. 3: Visual comparison of the generated images from the free-hand sketch dataset by pix2pix [5], CycleGAN [15], UNIT [16], U-GAT-IT [17], and our proposed Sketch2MedI models.

unrealistic. From the experimental results, the outputs of the supervised pix2pix model somewhat align with the sketch edges. However, its overall quality remains unsatisfactory. Concurrently, the unpaired image translation methods (i.e., CycleGAN, UNIT, and U-GAT-IT) severely failed to generalize on the free-hand sketches. These methods were not able to learn meaningful representations from sketch data, thus yielding inferior results (Table I, Fig. 3).

On the contrary, thanks to the ability to encode the input sketches into StyleGAN’s latent space which has expressive and semantically rich representations, the Sketch2MedI model is robust against the distribution shift, consistently ensuring the generation of visually pleasing and high-quality results. Our method performed efficiently under various user drawing styles and significantly surpassed all other methods in terms of the quality aspect (FID, MOS, and NR-IQA). We believe the high-quality medical image generation of the Sketch2MedI method could be a practical application in patient communication, as previously discussed.

For the originality aspect, our method also achieved superior

results to other methods (in DICE and FR-IQA). The utilization of H has been proven to improve originality preservation compared to when using the encoder E alone (Sketch2MedI w/o H in Table I). In addition, the same performance trend has been observed where H helped to improve the mean DICE score by 5.03 points (from 71.59 to 76.62) on 6,000 test synthesized sketches. Although this originality evaluation is considerably not ideal as the free-hand sketches are not well-aligned with target image objects (see Fig. 2), these results suggest the possibility of utilizing our Sketch2MedI for downstream tasks such as sketch-based medical image retrieval.

Despite achieving practical results, our Sketch2MedI occasionally generated images with loosened sketch information. For instance, in Fig. 3 (last row, red arrows), the results from Sketch2MedI do not precisely express the intended tonsil shapes (i.e., generated smaller shapes) compared to other models. This drawback might be because the sketch data is encoded into the abstract latent space before the generation, causing the inadequacy in capturing sketch information. We believe

introducing some constraints during training can mitigate this issue and intend to improve our model in future works.

VI. CONCLUSION

In this study, we proposed the practical Sketch2MedI model for realistic free-hand sketch to medical image generation. By encoding sketch inputs into StyleGAN's latent space which has semantically rich representations and then generating medical images from it, the Sketch2MedI proved to be effective in generalizing to the free-hand sketches even though it was trained entirely on the different synthesized sketch dataset. Our proposal not only generates high-quality and realistic images but also exhibits its applicability to useful utilization in the field of medical imaging.

ACKNOWLEDGMENT

We would like to express our sincere thanks to all members at Aillis Inc. who contributed to this project, especially Hiroshi Yoshihara (M.Sc.) and Wataru Takahashi (M.Eng.) for their valuable comments and feedback.

REFERENCES

- [1] W. Chen and J. Hays, "Sketchygan: Towards diverse and realistic sketch to image synthesis," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9416–9425.
- [2] C. Gao, Q. Liu, Q. Xu, L. Wang, J. Liu, and C. Zou, "Sketchycoco: Image generation from freehand scene sketches," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 5174–5183.
- [3] C. Ham, G. C. Tarres, T. Bui, J. Hays, Z. Lin, and J. Collomosse, "Cogs: Controllable generation and search from sketch and style," in *European Conference on Computer Vision*, 2022, pp. 632–650.
- [4] S. Koley, A. K. Bhunia, A. Sain, P. N. Chowdhury, T. Xiang, and Y.-Z. Song, "Picture that sketch: Photorealistic image generation from abstract sketches," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2023, pp. 6850–6861.
- [5] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1125–1134.
- [6] Y. Li, X. Chen, F. Wu, and Z.-J. Zha, "Linestofacephoto: Face photo generation from lines with conditional self-attention generative adversarial networks," in *ACM International Conference on Multimedia*, 2019, pp. 2323–2331.
- [7] S.-Y. Chen, W. Su, L. Gao, S. Xia, and H. Fu, "Deepfacedrawing: Deep generation of face images from sketches," *ACM Transactions on Graphics*, vol. 39, no. 4, pp. 72–1, 2020.
- [8] Y. Li, X. Chen, B. Yang, Z. Chen, Z. Cheng, and Z.-J. Zha, "Deepfacepencil: Creating face images from freehand sketches," in *ACM International Conference on Multimedia*, 2020, pp. 991–999.
- [9] S.-Y. Chen, F.-L. Liu, Y.-K. Lai, P. L. Rosin, C. Li, H. Fu, and L. Gao, "Deepfaceediting: Deep face generation and editing with disentangled geometry and appearance control," *ACM Transactions on Graphics*, vol. 40, no. 4, pp. 1–15, 2021.
- [10] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4401–4410.
- [11] R. Abdal, Y. Qin, and P. Wonka, "Image2stylegan: How to embed images into the stylegan latent space?" in *IEEE International Conference on Computer Vision*, 2019, pp. 4432–4441.
- [12] O. Tov, Y. Alaluf, Y. Nitzan, O. Patashnik, and D. Cohen-Or, "Designing an encoder for stylegan image manipulation," *ACM Transactions on Graphics*, vol. 40, no. 4, pp. 1–14, 2021.
- [13] Y. Alaluf, O. Tov, R. Mokady, R. Gal, and A. Bermano, "Hyperstyle: Stylegan inversion with hypernetworks for real image editing," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022, pp. 18 511–18 521.
- [14] H. Liu, Y. Song, and Q. Chen, "Delving stylegan inversion for image editing: A foundation latent space viewpoint," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2023, pp. 10 072–10 082.
- [15] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *IEEE International Conference on Computer Vision*, 2017, pp. 2223–2232.
- [16] M.-Y. Liu, T. Breuel, and J. Kautz, "Unsupervised image-to-image translation networks," in *Advances in Neural Information Processing Systems*, 2017, pp. 1–9.
- [17] J. Kim, M. Kim, H. Kang, and K. Lee, "U-gat-it: Unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation," in *International Conference on Learning Representations*, 2020, pp. 1–19.
- [18] T. Zhang, H. Fu, Y. Zhao, J. Cheng, M. Guo, Z. Gu, B. Yang, Y. Xiao, S. Gao, and J. Liu, "Skrgan: Sketching-rendering unconditional generative adversarial networks for medical image synthesis," in *Medical Image Computing and Computer Assisted Intervention*, 2019, pp. 777–785.
- [19] J. Liang, X. Yang, Y. Huang, H. Li, S. He, X. Hu, Z. Chen, W. Xue, J. Cheng, and D. Ni, "Sketch guided and progressive growing gan for realistic and editable ultrasound image synthesis," *Medical Image Analysis*, vol. 79, p. 102461, 2022.
- [20] R. Toda, A. Teramoto, M. Kondo, K. Imaizumi, K. Saito, and H. Fujita, "Lung cancer ct image generation from a free-form sketch using style-based pix2pix for data augmentation," *Scientific Reports*, vol. 12, no. 1, p. 12867, 2022.
- [21] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of stylegan," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 8110–8119.
- [22] X. Wang, R. Zhang, T. Kong, L. Li, and C. Shen, "Solov2: Dynamic and fast instance segmentation," in *Advances in Neural Information Processing Systems*, 2020, pp. 17 721–17 732.
- [23] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [24] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 586–595.
- [25] X. Chen, H. Fan, R. Girshick, and K. He, "Improved baselines with momentum contrastive learning," *arXiv:2003.04297*, 2020. [Online]. Available: <https://arxiv.org/abs/2003.04297>
- [26] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, 2014, pp. 2672–2680.
- [27] Y. Nitzan, A. Bermano, Y. Li, and D. Cohen-Or, "Face identity disentanglement via latent space mapping," *ACM Transactions on Graphics*, vol. 39, pp. 1–14, 2020.
- [28] L. Mescheder, A. Geiger, and S. Nowozin, "Which training methods for gans do actually converge?" in *International Conference on Machine Learning*, 2018, pp. 3481–3490.
- [29] Q. H. Cap, A. Fukuda, and H. Iyatomi, "A practical framework for unsupervised structure preservation medical image enhancement," *arXiv:2304.01864*, 2023. [Online]. Available: <https://arxiv.org/abs/2304.01864>
- [30] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," in *Advances in Neural Information Processing Systems*, 2017, pp. 1–12.
- [31] S. A. Golestaneh, S. Dadsetan, and K. M. Kitani, "No-reference image quality assessment via transformers, relative ranking, and self-consistency," in *IEEE Winter Conference on Applications of Computer Vision*, 2022, pp. 1220–1230.
- [32] J. Wang, K. C. Chan, and C. C. Loy, "Exploring clip for assessing the look and feel of images," in *AAAI Conference on Artificial Intelligence*, 2023, pp. 2555–2563.
- [33] K. Ding, K. Ma, S. Wang, and E. P. Simoncelli, "Image quality assessment: Unifying structure and texture similarity," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 5, pp. 2567–2581, 2020.
- [34] A. Ghildyal and F. Liu, "Shift-tolerant perceptual similarity metric," in *European Conference on Computer Vision*, 2022, pp. 91–107.