



Uncertainty-Guided Progressive GANs for Medical Image Translation

Uddeshya Upadhyay¹(✉), Yanbei Chen¹, Tobias Hepp^{1,2}, Sergios Gatidis^{1,2},
and Zeynep Akata^{1,2}

¹ University of Tübingen, Tübingen, Germany

² Max Planck Institute for Intelligent Systems, Tübingen, Germany

Abstract. Image-to-image translation plays a vital role in tackling various medical imaging tasks such as attenuation correction, motion correction, undersampled reconstruction, and denoising. Generative adversarial networks have been shown to achieve the state-of-the-art in generating high fidelity images for these tasks. However, the state-of-the-art GAN-based frameworks do not estimate the uncertainty in the predictions made by the network that is essential for making informed medical decisions and subsequent revision by medical experts and has recently been shown to improve the performance and interpretability of the model. In this work, we propose an uncertainty-guided progressive learning scheme for image-to-image translation. By incorporating aleatoric uncertainty as attention maps for GANs trained in a progressive manner, we generate images of increasing fidelity progressively. We demonstrate the efficacy of our model on three challenging medical image translation tasks, including PET to CT translation, undersampled MRI reconstruction, and MRI motion artefact correction. Our model generalizes well in three different tasks and improves performance over state of the art under full-supervision and weak-supervision with limited data. Code is released here: <https://github.com/ExplainableML/UncerGuidedI2I>.

Keywords: Image-to-image translation · Uncertainty estimation · Progressive GANs · PET · CT · MRI · Artefact correction

1 Introduction

In the medical domain, each imaging modality reflects particular physical properties of the tissue under examination. This results in images with different dimensionality, spatial resolution, and contrast. Various imaging modalities provide a complimentary stream of information for clinical diagnostics or technical pre and post-processing steps. Moreover, acquiring medical images is susceptible to various kinds of noise and modality-specific artefacts. To remedy these issues, translating images between different domains is of great importance.

Inter-modal image-to-image translation can potentially replace additional acquisition procedures, reducing examination costs and time. Besides, intra-modality image-to-image translation enables complex artefact and noise correction. For example, attenuation correction of positron emission tomography

(PET) data is challenging in situations where no density distribution is available from computed tomography (CT) data, as in the case for stand-alone PET scanners or combined PET/magnetic resonance imaging (MRI). In these situations, the generation of pseudo-CTs from PET data can be helpful. Further examples are related to image reconstruction and/or correction in MRI: Reconstruction of undisturbed artifact-free images is hard to achieve with traditional methods; deep-learning-based image-to-image translation can solve this challenge. In particular, generative adversarial networks (GAN) based on convolutional neural networks (CNN) have proven to provide a high visual quality of the generated synthetic images. However, predictions of GANs can be unreliable, and particularly in medical applications, the quantification of uncertainty is of high importance for the interpretation of the results. In this work, we propose a generic end-to-end model that introduces high-capacity conditional progressive GANs to synthesize high-quality images, using aleatoric uncertainty estimates as the guide to focus on improving image quality in regions where the network is highly uncertain about the prediction. We perform experiments on three challenging and vital medical imaging tasks: PET to CT translation, undersampled MRI reconstruction, and motion correction in MRI. Moreover, we empirically demonstrate the efficacy of our model under weak supervision with limited data.

2 Related Works

Traditional machine learning techniques for medical image translation rely on explicit feature representations [6, 12, 17, 35]. More recently, convolutional neural networks have been proposed for various image translation tasks [3–5, 8, 13, 19] and state-of-the-art performance is achieved by generative adversarial networks [1, 2, 7, 9, 15, 21, 22, 29–32]. The existing methods propose conditional GAN architectures with deterministic outputs that typically uses $\mathcal{L}_1/\mathcal{L}_2$ -based fidelity loss for the generator assumes a pixel-wise *homoscedasticity* and also assumes the pixel-wise error (i.e., residual) to be *independent and identically distributed* (i.i.d) following a Laplace or Gaussian distribution. This is a limiting assumption as explained in [10, 23, 25]. While these methods can provide synthetic images of high visual quality, the image content may still deviate significantly from the corresponding ground-truth. This results in overconfidence or misinterpretation with negative consequences, particularly in the medical domain. There have been recent works on quantifying aleatoric and epistemic uncertainty in task-specific medical imaging algorithms like classification, segmentation, super-resolution etc. [14, 20, 25–27] quantifying it for general image-to-image translation problem largely remains unexplored. Thus, the central motivation of our work is to provide measures of uncertainty for image-to-image translation tasks that can contribute to safe applications of results.

Moreover, recent work has shown that high-capacity generators that are progressive in nature lead to high-quality results as described in [1, 2, 9]. However, the progressive generation of high-quality images remains unguided without specifically attending to poorly translated regions. Prior works indicate a correlation between estimated uncertainty and prediction error [20, 23, 33]. We exploit

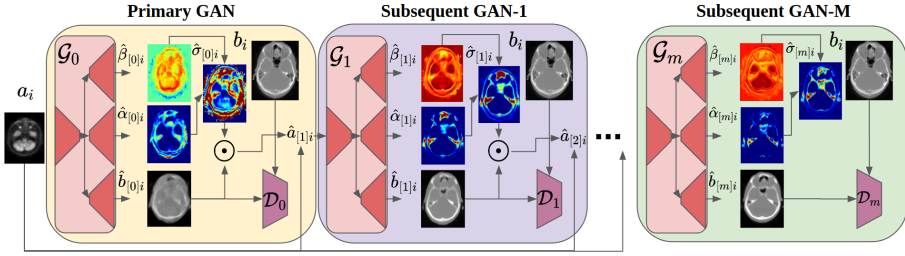


Fig. 1. Uncertainty-guided Progressive GANs (UP-GAN): The primary GAN takes the input image from domain A , while subsequent GANs absorb outputs from the preceding GAN (see Eq. 3 and 4). Explicitly guided by the attention maps, the uncertainty maps are estimated from the preceding GAN.

this relationship for the progressive enhancement of synthetic images, which has not been investigated by prior work before.

3 Uncertainty-Guided Progressive GAN (UP-GAN)

Let A and B be two image domains with a set of images $S_A := \{a_1, a_2 \dots a_n\}$ and $S_B := \{b_1, b_2 \dots b_m\}$ where a_i and b_i represent the i^{th} image from domain A and B respectively. Let each image drawn from an underlying *unknown* probability distribution \mathcal{P}_{AB} , i.e., $(a_i, b_i) \sim \mathcal{P}_{AB} \forall i$ have K pixels, and u_{ik} represent the k^{th} pixel of a particular image u_i . Our goal is to learn a mapping from domain A to B ($A \rightarrow B$) in a paired manner, i.e., learning the underlying conditional distribution $\mathcal{P}_{B|A}$ from the set of given samples $\{(a_i, b_i)\}$, following the distribution \mathcal{P}_{AB} . For a given image a_i in domain A , the estimated image in domain B is called \hat{b}_i . The pixel wise error is defined as $\epsilon_{ij} = \hat{b}_{ij} - b_{ij}$. While the existing framework models the residual as the i.i.d as described above, we relax that assumption by modelling the residual as non i.i.d variables and learning the optimal distribution from the dataset, as described in the following.

Figure 1 shows our model that consists of cascaded GANs, where each generator is capable of estimating the aleatoric uncertainty, along with generating images. Our solution alleviates the aforementioned limitations of recent methods by modelling the underlying per-pixel residual distribution as *independent* but *non-identically* distributed *zero-mean generalized Gaussian distribution* (GGD) as in [23], where the network learns to predict the optimal *scale* (α) and *shape* (β) of the GGD for every pixel. Therefore, $\hat{b}_{ij} = b_{ij} + \epsilon_{ij}$ with, $\epsilon_{ij} \sim GGD(\epsilon; 0, \alpha_{ij}, \beta_{ij}) \equiv \beta_{ij} (2\alpha_{ij} \Gamma(\beta_{ij}^{-1}))^{-1} \exp(-\alpha_{ij}^{-1} |\epsilon|^{\beta_{ij}})$. We generate images in multiple phases, with each phase generating output images along with the aleatoric uncertainty estimates. The outputs from one phase serve as the input to the subsequent GAN in the next phase, explicitly guided by the attention map derived from uncertainty estimates. Importantly, this uncertainty-based guidance enforces the model to focus on refining the uncertain regions that are likely to be poorly synthesized, resulting in progressively improving quality.

Our framework is composed of a sequence of M GANs, where the m^{th} GAN is represented by a pair of networks, generator and discriminator, given by, $(\mathcal{G}_m(\cdot; \theta_m), \mathcal{D}_m(\cdot; \phi_m))$. Both the generator and discriminator can have arbitrary network architecture as long as generator can estimate *aleatoric uncertainty* as described in [23]. We choose all the discriminators to be the patch discriminators from [7] and generators to be modified U-Net [16], where the head is split into three to estimate the parameters of the GGD as shown in Fig. 1 and in [23].

Primary GAN. We train the first GAN (\mathcal{G}_0) using the dataset S_A and S_B . The predictions of the generator are given by $(\hat{\alpha}_{[0]i}, \hat{\beta}_{[0]i}, \hat{b}_{[0]i})$. The network is trained with an adaptive fidelity loss function $\mathcal{L}_{\alpha\beta}^G$ [23] and an adversarial loss $\mathcal{L}_{\text{adv}}^G$ [36], combined as $\mathcal{L}_{\text{tot}}^G$ for the generator ($\mathcal{G}_0(\cdot; \theta_0) : A \rightarrow B$):

$$\mathcal{L}_{\alpha\beta}^G(\hat{b}_{[0]i}, \hat{\alpha}_{[0]i}, \hat{\beta}_{[0]i}, b_i) = \frac{1}{K} \sum_j \left(\frac{|\hat{b}_{[0]ij} - b_{ij}|}{\hat{\alpha}_{[0]ij}} \right)^{\hat{\beta}_{[0]ij}} - \log \frac{\hat{\beta}_{[0]ij}}{\hat{\alpha}_{[0]ij}} + \log \Gamma(\hat{\beta}_{[0]ij}^{-1}) \quad (1)$$

$$\mathcal{L}_{\text{adv}}^G = \mathcal{L}_2(\mathcal{D}_1(\hat{b}_{[0]i}), 1) \text{ and } \mathcal{L}_{\text{tot}}^G = \lambda_1 \mathcal{L}_{\alpha\beta}^G + \lambda_2 \mathcal{L}_{\text{adv}}^G. \quad (2)$$

The patch discriminator (\mathcal{D}_1) is trained using the adversarial loss from [36] given by $\mathcal{L}_{\text{adv}}^D = \mathcal{L}_2(\mathcal{D}^A(b_i), 1) + \mathcal{L}_2(\mathcal{D}^A(\hat{b}_{[0]i}), 0)$.

Subsequent GANs. The m^{th} GAN (where $m > 0$) takes the output produced by the $(m-1)^{\text{th}}$ GAN, i.e. $(\hat{\alpha}_{[m-1]i}, \hat{\beta}_{[m-1]i}, \hat{b}_{[m-1]i})$, along with the original sample a_i from domain A as its input and generates a refined output. The image estimated by the $(m-1)^{\text{th}}$ GAN along with its uncertainty map learns to create the input feature $f_{[m]i}$ for the m^{th} GAN, where the uncertainty map serves as an attention mechanism to highlight the uncertain regions in the image. The input $a_{[m]i}$ for the m^{th} generator is given by concatenating a_i and $f_{[m]i}$, i.e.,

$$\hat{\sigma}_{[m-1]i} = \hat{\alpha}_{[m-1]i} \sqrt{\frac{\Gamma(3/\hat{\beta}_{[m-1]i})}{\Gamma(1/\hat{\beta}_{[m-1]i})}}, \text{ and } f_{[m]i} = \hat{b}_{[m-1]i} \odot \frac{\hat{\sigma}_{[m-1]i}}{\sum_j \hat{\sigma}_{[m-1]ij}} \quad (3)$$

$$a_{[m]i} = \text{concat}(f_{[m]i}, a_i) \quad (4)$$

The input $a_{[m]i}$ for the m^{th} GAN encourages the generator to further refine the highly uncertain regions in the image given the original input context. The generator and the discriminator are trained using $\mathcal{L}_{\text{tot}}^G$ and $\mathcal{L}_{\text{adv}}^D$, respectively.

Progressive Training Scheme. We initialize the parameters $\theta \cup \phi$ sequentially. First, we initialize $\theta_1 \cup \phi_1$ using the training set (S_A, S_B) to minimize the loss function given by $\mathcal{L}_{\text{tot}}^G$ and $\mathcal{L}_{\text{adv}}^D$. Then, for the subsequent GANs, we initialize the $\theta_m \cup \phi_m$ ($m > 1$) by fixing the weights of all the previous generators and training the m^{th} GAN alone (see Eq. 3 and 4 with losses $\mathcal{L}_{\text{tot}}^G$ and $\mathcal{L}_{\text{adv}}^D$). Once all the parameters have been initialized (i.e., $\theta_m \cup \phi_m \forall m$), we do further fine tuning by training all the networks end-to-end by combining the loss functions of all the intermediate phases and a significantly smaller learning-rate.

4 Experiments

In this section, we first detail the experimental setup and comparative methods in Sect. 4.1, and present the corresponding results in Sect. 4.2.

4.1 Experimental Setup

Tasks and Datasets. We evaluate our method on the following three tasks.

- (i) PET to CT translation: We synthesize CT images from PET scans to be used for the attenuation correction, e.g. for PET-only scanners or PET/MRI. We use paired data sets of non-attenuation-corrected PET and the corresponding CT of the head region of 49 patients acquired on a state-of-the-art PET/CT scanner (Siemens Biograph mCT), approved by ethics committee of the Medical Faculty of the University of Tübingen. Data is split into 29/5/15 for training/val/test sets. Figure 2 shows exemplary slices for co-registered PET and CT.
- (ii) Undersampled MRI reconstruction: We translate undersampled MRI images to fully-sampled MRI images. We use MRI scans from the open-sourced IXI¹ dataset that consists of T1-weighted (T1w) MRI scans. We use a cohort of 500 patients split into 200/100/200 for training/val/test, and retrospectively create the undersampled MRI with an acceleration factor of 12.5 \times , i.e., we preserve only 8% of the fully-sampled k-space measurement (from the central region) to obtain the undersampled image.
- (iii) MRI Motion correction: We generate sharp images from motion corrupted images. We retrospectively create the motion artefacts in the T1w MRI from IXI following the transformations in the *k-space* as described in [18]. Figure 3-(ii) shows the input MRI scan with artefacts and ground-truth.

Training Details and Evaluation Metrics. All GANs are first initialized using the aforementioned progressive learning scheme with (λ_1, λ_2) in Eq. 2 set to $(1, 0.001)$. We use Adam [11], with the hyper-parameters $\beta_1 := 0.9$, $\beta_2 := 0.999$, an initial learning rate of 0.002 for initialization and 0.0005 post-initialization that decays based on cosine annealing over 1000 epochs, using a batch size of 8. We use three widely adopted metrics to evaluate image generation quality: PSNR measures $20 \log \text{MAX}_I / \sqrt{\text{MSE}}$, where MAX_I is the highest possible intensity value in the image and MSE is the mean-squared-error between two images. SSIM computes the structural similarity between two images [28]. MAE computes the mean absolute error between two images. Higher PSNR, SSIM, and lower MAE indicate a higher quality of the generated images (wrt ground-truth).

Compared Methods. We compare our model to representative state-of-the-art methods for medical image translation, including Pix2pix [7], a baseline conditional adversarial networks for image-to-image translation tasks using GANs,

¹ From <https://brain-development.org/ixi-dataset/>.

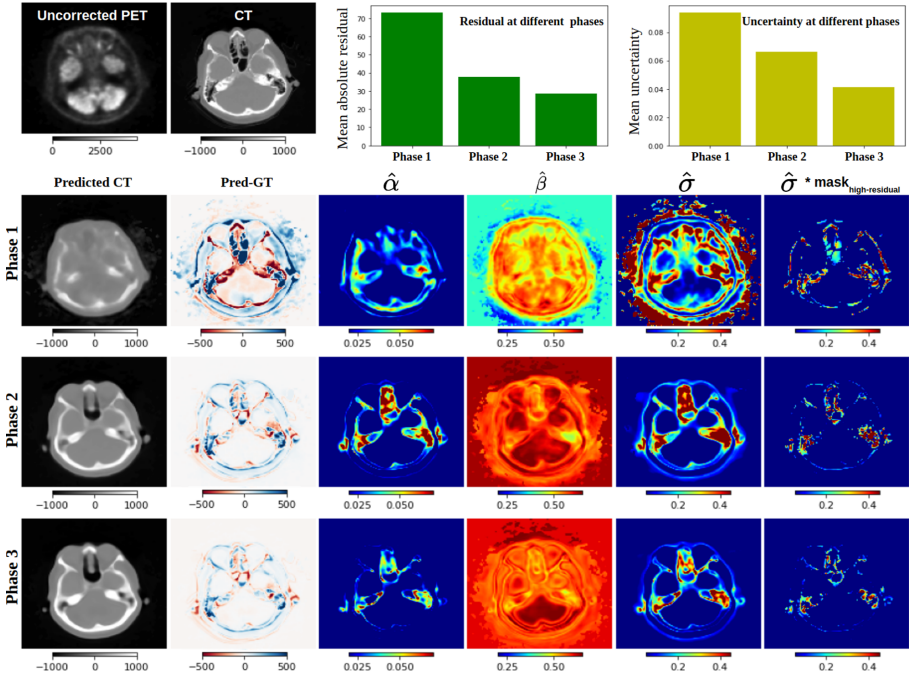


Fig. 2. Outputs from different phases of UP-GAN (with $M = 3$). (Top) The input (uncorrected PET), the corresponding ground-truth CT, mean residual values over different phases, mean uncertainty values over different phases. (Bottom) Each row shows the predicted output, the residual between the prediction and the ground-truth, the predicted scale ($\hat{\alpha}$) map, the predicted shape ($\hat{\beta}$) map, the uncertainty map, and the uncertainty in high residual regions.

PAN [24], and MedGAN [2], a GAN-based method that relies on *external-pre-trained feature extractors*, with a generator that refines the generated images progressively. MedGAN is shown to perform superior to methods like, FilsGAN [34], ID-cGAN [32], and achieve state-of-the-art performance for several medical image-to-image translation problems.

4.2 Results and Analysis

Qualitative Results. Figure 2 visualizes the (intermediate) outputs of the generators at different phases of the framework. The visual quality of the generated image content increasingly improves along the network phases (as shown in the first column, second row onward). At the same time, prediction error and uncertainty decrease continuously (second column and fifth column, second row onward, respectively). High uncertainty values are found in anatomical regions with fine osseous structures, such as the nasal cavity and the inner ear in the

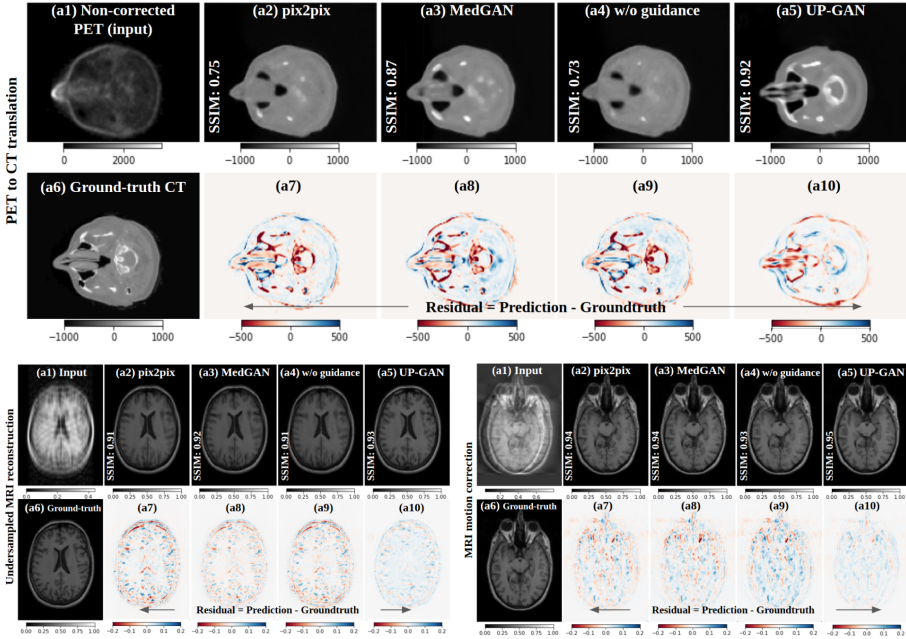


Fig. 3. Qualitative results. (Top) PET to CT translation. (Bottom) Undersampled MRI reconstruction (left), and MRI motion correction (right).

petrous portion of the temporal bone. Particularly in such regions of high uncertainty, we achieve a progressive improvement in the level of detail.

Figure 3-(Top) visualizes the generated CT images from the PET for all the compared methods along with our methods. We observe that more high-frequency features are present in our prediction compared to the previous state-of-the-art model (MedGAN). We also observe that the overall residual is significantly lower for our method compared to the other baselines. MedGAN performs better than pix2pix in synthesizing high-frequency features and sharper images. Figure 3-(Bottom) shows similar results for the undersampled MRI reconstruction task and MRI motion correction task. In both cases, our model yields superior images, as can be seen via relatively neutral residual maps.

Quantitative Results. Table 1 shows the quantitative performance of all the methods on the three tasks; for all the tasks, our method outperforms the recent models. In particular, for the most challenging task, PET to CT translation, our method with uncertainty-based guide outperforms the previous state-of-the-art method, MedGAN (that relies on task-specific external feature extractor), *without using any external feature extractor*. Therefore, the uncertainty guidance reduces the burden of having an externally trained task-specific feature extractor to achieve high fidelity images. The same trend holds for undersampled MRI reconstruction and motion correction in MRI. The statistical tests on




Fig. 4. Quantitative results in the presence of limited labeled training data.

Table 1. Evaluation of various methods on three medical image translation tasks.

Methods	PET to CT			Undersampled MRI Recon			MRI Motion Correction		
	SSIM	PSNR	MAE	SSIM	PSNR	MAE	SSIM	PSNR	MAE
pix2pix [7]	0.89±0.04	26.0±2.0	38.5±10.7	0.92±0.03	28.5±0.9	27.6±9.3	0.94±0.06	29.6±1.4	26.3±8.2
PAN [24]	0.90±0.08	26.5±4.5	37.2±15.6	0.93±0.05	28.8±0.7	26.2±10.4	0.95±0.10	30.1±2.8	24.9±9.7
MedGAN [2]	0.90±0.04	27.1±2.5	35.4±11.8	0.94±0.02	29.7±1.9	24.2±8.7	0.95±0.04	30.8±1.8	23.6±9.1
UP-GAN	0.95±0.05	28.9±0.4	24.7±12.9	0.97±0.07	29.4±2.1	24.1±7.5	0.96±0.03	32.1±0.3	22.8±11.1

SSIM values of MedGAN and our UP-GAN gives us a p -value of 0.016 for PET-to-CT translation, 0.021 for undersampled MRI reconstruction, and 0.036 for MRI motion correction. As all the p -values are < 0.05 , results are statistically significant.

Ablation Study. We study the model that does not utilize the estimated uncertainty maps as attention maps and observe that the model without the uncertainty as the guide performs inferior to the UP-GAN with a performance (SSIM/PSNR/MAE) of (0.87/25.4/40.7), (0.93/27.3/38.7), and (0.92/26.2/35.1) for PET to CT translation, undersampled MRI reconstruction, and MRI motion correction, respectively. UP-GAN model leverages the uncertainty map to refine the predictions where the model is uncertain, which is also correlated to the regions where the translation is poor. The model without uncertainty-based guidance does not focus on the regions mentioned above in the prediction and is unable to perform as well as UP-GAN.

Evaluating Models with Weak Supervision. We evaluate all the models for PET to CT synthesis by limiting the number of paired image samples used for training. We define *five* supervision levels corresponding to different amounts of cross-domain pairwise training sample slices. For this experiment, we train the recent state-of-the-art models with a varying number of patients in the training stage, i.e., we use 5, 10, 15, 20, and 29 patients, respectively. Figure 4 shows the performance of all the models at varying supervision levels. We observe that our model with uncertainty guidance outperforms all the baselines at full supervision (with 29 patients). Moreover, our model sharply outperforms the baselines with limited training data (with < 29 patients). UP-GAN produces intermediate uncertainty maps that have higher values under weak supervision (compared to the full supervision case), but this still allows UP-GAN to focus

on highly uncertain regions, that the current state-of-the-art models do not have access to, hence are not able to leverage that to refine the predicted images.

5 Conclusion

In this work, we propose a new generic model for medical image translation using uncertainty-guided progressive GANs. We demonstrate how uncertainty can serve as an attention map in progressive learning schemes. We demonstrate the efficacy of our method on three challenging medical image translation tasks, including PET to CT translation, undersampled MRI reconstruction, and motion correction in MRI. Our method achieves state-of-the-art in various tasks. Moreover, it allows the quantification of uncertainty and shows better generalizability with smaller sample sizes than recent approaches.

Acknowledgements. This work has been partially funded by the ERC (853489 - DEXIM) and by the DFG (2064/1 - Project number 390727645). The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support.

References

1. Armanious, K., Jiang, C., Abdulatif, S., Küstner, T., Gatidis, S., Yang, B.: Unsupervised medical image translation using cycle-MedGAN. In: European Signal Processing Conference (EUSIPCO) (2019)
2. Armanious, K., et al.: MedGAN: medical image translation using GANs. *Comput. Med. Imaging Graph.* **79**, 101684 (2020)
3. Chen, H., et al.: Low-dose CT with a residual encoder-decoder convolutional neural network. *IEEE TMI* **36**, 2524–2535 (2017)
4. Dou, Q., Chen, H., Yu, L., Qin, J., Heng, P.A.: Multilevel contextual 3-D CNNs for false positive reduction in pulmonary nodule detection. *IEEE Trans. Biomed. Eng.* **64**, 1558–1567 (2016)
5. Havaei, M., et al.: Brain tumor segmentation with deep neural networks. *Med. Image Anal.* **35**, 18–31 (2017)
6. Huynh, T., et al.: Estimating CT image from MRI data using structured random forest and auto-context model. *IEEE TMI* **35**, 174–183 (2015)
7. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *IEEE CVPR* (2017)
8. Kamnitsas, K., et al.: Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Med. Image Anal.* **36**, 61–78 (2017)
9. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: *ICLR* (2018)
10. Kendall, A., Gal, Y.: What uncertainties do we need in Bayesian deep learning for computer vision? In: *NIPS* (2017)
11. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
12. Küstner, T., et al.: MR-based respiratory and cardiac motion correction for pet imaging. *Med. Image Anal.* **42**, 129–144 (2017)

13. Litjens, G., et al.: A survey on deep learning in medical image analysis. *Med. Image Anal.* **42**, 60–88 (2017)
14. Nair, T., Precup, D., Arnold, D.L., Arbel, T.: Exploring uncertainty measures in deep networks for multiple sclerosis lesion detection and segmentation. *Med. Image Anal.* **59**, 101557 (2020)
15. Nie, D., et al.: Medical image synthesis with deep convolutional adversarial networks. *IEEE Trans. Biomed. Eng.* **65**, 2720–2730 (2018)
16. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015*. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
17. Rueda, A., Malpica, N., Romero, E.: Single-image super-resolution of brain MR images using overcomplete dictionaries. *Med. Image Anal.* **17**, 113–132 (2013)
18. Shaw, R., Sudre, C., Ourselin, S., Cardoso, M.J.: MRI k-space motion artefact augmentation: model robustness and task-specific uncertainty. In: *MIDL (2019)*
19. Shin, H.C., et al.: Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. *IEEE TMI* **35**, 1285–1298 (2016)
20. Tanno, R., et al.: Bayesian image quality transfer with CNNs: exploring uncertainty in dMRI super-resolution. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017*. LNCS, vol. 10433, pp. 611–619. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66182-7_70
21. Upadhyay, U., Awate, S.P.: A mixed-supervision multilevel GAN framework for image quality enhancement. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11768, pp. 556–564. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32254-0_62
22. Upadhyay, U., Awate, S.P.: Robust super-resolution GAN, with manifold-based and perception loss. In: *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pp. 1372–1376. IEEE (2019)
23. Upadhyay, U., Chen, Y., Akata, Z.: Uncertainty-aware generalized adaptive CycleGAN. preprint [arXiv:2102.11747](https://arxiv.org/abs/2102.11747) (2021)
24. Wang, C., Xu, C., Wang, C., Tao, D.: Perceptual adversarial networks for image-to-image transformation. *IEEE TIP* **27**, 4066–4079 (2018)
25. Wang, G., Li, W., Aertsen, M., Deprest, J., Ourselin, S., Vercauteren, T.: Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. *Neurocomputing* **338**, 34–45 (2019)
26. Wang, G., Li, W., Ourselin, S., Vercauteren, T.: Automatic brain tumor segmentation using convolutional neural networks with test-time augmentation. In: Crimi, A., Bakas, S., Kuijff, H., Keyvan, F., Reyes, M., van Walsum, T. (eds.) *BrainLes 2018*. LNCS, vol. 11384, pp. 61–72. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-11726-9_6
27. Wang, G., Li, W., Ourselin, S., Vercauteren, T.: Automatic brain tumor segmentation based on cascaded convolutional neural networks with uncertainty estimation. *Front. Comput. Neurosci.* **13**, 56 (2019)
28. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE TIP* **13**, 600–612 (2004)
29. Wolterink, J.M., Dinkla, A.M., Savenije, M.H.F., Seevinck, P.R., van den Berg, C.A.T., Išgum, I.: Deep MR to CT synthesis using unpaired data. In: Tsaftaris, S.A., Gooya, A., Frangi, A.F., Prince, J.L. (eds.) *SASHIMI 2017*. LNCS, vol. 10557, pp. 14–23. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68127-6_2

30. Yang, Q., et al.: Low-dose CT image denoising using a generative adversarial network with Wasserstein distance and perceptual loss. *IEEE TMI* **37**, 1348–1357 (2018)
31. Yi, X., Walia, E., Babyn, P.: Generative adversarial network in medical imaging: a review. *Med. Image Anal.* **58**, 101552 (2019)
32. Zhang, H., Sindagi, V., Patel, V.M.: Image de-raining using a conditional generative adversarial network. *IEEE Trans. Circuits Syst. Video Technol.* **30**, 3943–3956 (2019)
33. Zhang, Z., Romero, A., Muckley, M.J., Vincent, P., Yang, L., Drozdal, M.: Reducing uncertainty in undersampled MRI reconstruction with active acquisition. In: *IEEE CVPR* (2019)
34. Zhao, H., Li, H., Cheng, L.: Synthesizing filamentary structured images with GANs. preprint [arXiv:1706.02185](https://arxiv.org/abs/1706.02185) (2017)
35. Zhong, L., et al.: Predict CT image from MRI data using KNN-regression with learned local descriptors. In: *IEEE ISBI* (2016)
36. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *IEEE ICCV* (2017)