# Probabilistic Recovery of Missing Images in Contrast-Enhanced CT

Dhruv V. Patel Chiao-Chih Hsu Bino Varghese Steven Cen

**Darryl Hwang** 

Inderbir Gill

Vinay Duddalwar

Assad A. Oberai

University of Southern California dhruvvpa, hsuchiao, bino.Varghese, cen, darrylhw, igill, Vinay.Duddalwar, aoberai@usc.edu

#### Abstract

Contrast-Enhanced CT (CECT) imaging is used in the diagnosis of renal cancer and planning of surgery. A sequence of CECT phase images which captures the movement of contrast agent inside an organ often completely misses some of the phase images or has corrupted phase images making the entire sequence useless. We propose a probabilistic deep generative model for imputing such missing phase images in CECT imaging. Our proposed model recovers the missing phase images with quantified uncertainty estimates enabling medical decision-makers make better-informed decisions. Furthermore, we propose a novel style-based adversarial loss to learn very fine-scale features unique to CECT imaging resulting in better recovery. We demonstrate the efficacy of this algorithm using a patient dataset collected in an IRB-approved retrospective study.

# 1 Introduction

Medical image data acquired from ultrasound, X-rays (CT), MR, and other modalities are used routinely in detecting, diagnosing, and planning treatment for myriad diseases. The problem of missing data is ubiquitous in medical imaging. Missing image data can be in the form of missing images in a sequence of images, or missing regions within a single image, or artifacts like blurring. In all these cases, missing data leads to the loss in the utility of the images, and an accompanying loss in the accuracy of detection, diagnosis, and treatment planning for a disease.

There are many reasons for missing or lost data. In some cases, patients may be initially scanned under one protocol, while the final management of disease might require additional, or more thorough, scans. However, this may not be feasible due to the patient's inability to tolerate additional scans, logistical issues, and restrictions imposed by the insurance provider. In addition to this, in [1], the authors refer to missing image data as the "leaky" radiological pipeline. They point to several causes for missing image data that include incompatibility between different vendors of medical imaging equipment, a saturation of the bandwidth of a device, and collateral damage during events like server errors. In all these cases, a portion of image data is missing and because of this the portion that is collected is rendered useless, or of little value.

While all imaging modalities have their strengths and limitations, Contrast-Enhanced CT (CECT) imaging is remarkably effective in detection of renal cancers; CECT images are generated by injecting an intravenous contrast agent into the subject and then imaging the organ during four

distinct time points. The images corresponding to these different time points are called pre-contrast, corticomedullary, nephrographic and excretory phase images. We note that the term "phase images" in this article refers to these images obtained during different time instances capturing the dynamics of contrast agent in an organ and not the images corresponding to the phase component of measured signal/image as is typically used in traditional medical imaging.

Conventional diagnosis whether a tumor is benign or malignant is based on the qualitative visual inspection of the four CECT phase images. Further, once a decision has been made to treat or resect the renal mass, these images are used by the surgeon to plan the surgery. The loss of any one or more CECT phase images due to any of the reasons discussed above negatively impacts the management of renal masses. It leads to less accurate diagnosis of malignant masses [2] and adversely effects surgical planning in cases where surgical intervention is necessary. In this manuscript we present a statistical deep-learning based technique for imputing missing phase images in a sequence of renal CECT image. This technique provides the best guess for the missing phase and also quantifies the confidence in this guess. This additional information can allow the clinician to make an informed decision about how much to trust the imputed data when delivering their final decision.

**Related work and our contribution** Image imputation refers to the task of recovering the missing/corrupted part of an image from the part that is available/not corrupted [3, 4]. It is an ill-posed inverse problem and Bayesian inversion provides a principled approach of solving it with quantified uncertainty estimates. However, when dealing with high-dimensional data involving complex prior information, which is typically the case in medical imaging, it faces significant challenges. Recently deep generative priors have shown considerable promise in learning complex probability distributions and are successfully being used to solve stochastic inverse problems in physical science, computer vision, and medical imaging applications [5, 6, 7, 8, 9]. Motivated by the success of these deep generative models in such diverse applications, in this work we use deep Genearative Adversarial Network (GAN) as a prior in Bayesian update for Contrast-enhancecd CT (CECT) image imputation task. Our main contributions are: (1) The development of a new stylized loss function for a Wasserstein GAN (WGAN) that produces more realistic CECT images than the vanilla WGAN.(2) The use of the WGAN as a prior in a Bayesian inference method to determine the most likely missing phase image, given an incomplete sequence of images. (3) Quantification of uncertainty in each pixel of the recovered image.

#### 2 **Problem formulation**

Let  $x \in \mathbb{R}^{n_x \times n_y \times 4}$  denote the true  $(n_x \times n_y)$  CECT image with all four phases, and let  $y \in \mathbb{R}^{n_x \times n_y \times n_p}$  denote the observed CECT image which contains  $n_p$  phases, where  $n_p = \{1, 2, 3\}$ . The observation y is related to true data x by, y = Mx, where M is the boolean masking operator. The goal of probabilistic image imputation is to infer the conditional probability distribution p(x|y). That is given an observed CECT image, y, with missing phases, recover the probability distribution of true underlying image containing all four phases x. Using Bayes' rule we represent this conditional density as  $p^{post}(x|y) = p^{like}(y|x)p^{prior}(x)/p(y)$ , where  $p^{prior}(x)$  is the prior density representing prior belief about the inferred signal x, and  $p^{like}(y|x)$  is the likelihood distribution.

In [9] an algorithm which uses a GAN to learn the prior distribution from multiple samples of x was proposed. In this work we use these GAN priors to infer the missing CECT phase image. Specifically, we first train a GAN using a sample set S containing 4-phased CECT images to learn the prior distribution directly from data. Thereafter, we reformulate the inference problem in the low-dimensional latent space of the GAN for efficient posterior sampling. Since the dimension of the latent space is much smaller than that of image space, one can use sampling-based techniques such as Markov Chain Monte Carlo (MCMC) to compute the statistic efficiently.

**Style-based loss** One of the unique characteristic of CECT images is the presence of fine-scale features. These features cannot be captured simply by training WGAN model using standard adversarial loss because the adversarial loss of a WGAN only encourages the model to minimize the Wasserstein distance between the training data density and the learned data density. Since this fine-scale structure is often times crucial in making important diagnostic decisions, it is desirable to have a model which can learn these features as well.

In this work we propose a novel style-based loss in addition to standard adversarial loss for learning the true prior density. Similar to previous works on deep style transfer and texture synthesis [?



Figure 1: Left and right panel corresponds to two different patients. Row 1: True phase images with missing Phase 1. Row 2: MAP estimate with vanilla GAN (MAP-WGAN), true image (True) and MAP and standard deviation estimates for a GAN with Style loss (MAP: WGAN + Style and SD: WGAN + Style ). Row 3: True phase images with missing Phase 2. Row 4: same as Row 2.

10, 11] we propose to use Gramm matrix based style loss. However, unlike any of the previous work we do not rely on the pre-trained classification network (VGG-16) to build the Gramm matrix. Instead we rely on the features extracted from certain layers of the discriminator to build the Gramm matrix. In other words, discriminator serves the dual purpose of a critic for real versus fake image classification and a feature extractor for style transfer. Specifically, we define the Gramm matrix as  $\mathcal{G}_{ij}{}^{l} = \sum_{k} \mathcal{F}_{ik}{}^{l}\mathcal{F}_{jk}{}^{l}$ , where  $\mathcal{F}_{ik}{}^{l}$  is the activation of  $i^{th}$  filter in layer l of the discriminator at location j. We then define the style loss by minimizing the Gramm matrices of a batch of real and fake samples. Specifically,  $\mathcal{L}_{style} = \sum_{n=1}^{b} ||\mathcal{G}_{ij}{}^{l}(real) - \mathcal{G}_{ij}{}^{l}(fake)||^{2}$ . In our study we use first 3 layers to compute style loss. The total loss is now given as  $\mathcal{L} = \mathcal{L}_{style} + \mathcal{L}_{adv}$ . Further, in order to introduce very fine-scale speckle pattern seen in the training data, we adapt the recent ideas from the state-of-the-art StyleGAN architectures [12] and inject a fixed amount of random noise in the final layer of the generator.

## **3** Results

This IRB-approved retrospective study included patients with renal lesions who had preoperative multiphase contrast enhanced CT. All scans were obtained on the same scanner (Brilliance 64, Philips Healthcare) during patient breath-holding with the following parameters: 120 kVp, variable tube current, slice thickness of 0.5 mm with reconstruction interval of 2 mm. An unenhanced CT scan of the of the abdomen was obtained first, followed by three contrast-enhanced scans in the corticomedullary (30 seconds), nephrographic (90 seconds), and excretory (5–7 minutes) phases. Approximately 100–150 mL of nonionic water-soluble IV iodinated contrast medium (iopamidol, Isovue 350, Bracco Imaging) dosed to weight was administered with a power injector at a rate of 5 mL/s. Tumor segmentation and phase co-registration Using Synapse 3D software (FujiFilm, Stamford CT), an experienced radiologist manually segmented the renal tumors as 3D ROIs. The nephrographic phase was used as the reference template for subsequent co-registering in other phases. Two-dimensional images of all tumors capturing the largest tumor diameters in each phase in the axial projection were selected and used as inputs to the WGAN model.

The GAN was trained using 375 unique CECT sequences (each with 4 phase images) which were augmented by a factor of 8 by translation and rotation resulting in 3,000 training sequences. Two different GAN models were considered: (a) WGAN model trained only using adversarial loss, and (b) WGAN model trained using adversarial and style loss. Once the GAN prior was learned, the posterior sampling was performed using Hamiltonian Monte Carlo (HMC). Figure 1 shows the representative

results for two patients from the test set for both GAN models. In each case we pretend that one phase (Phase 1 and then Phase 2) are missing and use the HMC sampler and the GAN models to recover them, and the compare the results with the true phase image. We observe that both GAN models are able to recover the overall shape and intensity of the missing phase. However, the GAN with the style loss is able to capture these fine-scale features in addition to the overall shape and average intensity. Furthermore, we also show the uncertainty in the recovered phase image for each case in the form of estimated standard deviation and we observe that it is highest at the boundary of the lesion. It is also large in regions where the intesity of the recovered image is large.

## 4 Conclusions

We present a novel probabilistic strategy for recovering missing phase images in CECT imaging of renal tumors. This strategy relies on a WGAN with a novel style loss to learn the prior distribution from a collection of complete CECT image sequences, and then implements a Bayesian update to recover the missing phase conditioned on the knowledge of the phases that are known. We present initial results on patient data where we recover the missing image and the pixel-wise uncertainty. We note that this idea can be extended to other applications in medical imaging as well.

## **Broader impact statement**

The work proposed in this manuscript aims at tackling the problem of recovery of missing phase images in contrast-enhanced CT imaging with quantified uncertainty estimates. This work could have a significant impact on diagnosis of renal cancer and planning of surgery, as it provides medical decision-makers the missing data in the medical images. It also estimates the error in this data allowing them to make more informed decisions, resulting in better patient care and improvement of public health at large. The proposed algorithm could also be extended to many different medical imaging modalities such as MRI, X-ray, and ultrasound.

# Acknowledgements

The support from ARO grant W911NF2010050 and the Ming-Hsieh Institute at USC is acknowledged.

## References

- [1] Catherine Oglevee and Oleg Pianykh. Losing Images in Digital Radiology: More than You Think. *Journal of Digital Imaging*, 28(3):264–271, jun 2015.
- [2] Marta E Heilbrun, Erick M Remer, David D Casalino, Michael D Beland, Jay T Bishoff, M Donald Blaufox, Courtney A Coursey, Stanley Goldfarb, Howard J Harvin, Paul Nikolaidis, et al. Acr appropriateness criteria indeterminate renal mass. *Journal of the American College of Radiology*, 12(4):333–341, 2015.
- [3] A. Dalca, K. Bouman, W. Freeman, N. Rost, M. Sabuncu, and P. Golland. Medical image imputation from image collections. *IEEE Transactions on Medical Imaging*, 38:504–514, 2019.
- [4] Dongwook Lee, Junyoung Kim, Won-Jin Moon, and Jong Chul Ye. CollaGAN : Collaborative GAN for Missing Image Data Imputation. jan 2019.
- [5] Erik M. Lindgren, Jay Whang, and Alexandros G. Dimakis. Conditional Sampling from Invertible Generative Models with Applications to Inverse Problems. feb 2020.
- [6] Francesco Tonolini, Jack Radford, Alex Turpin, Daniele Faccio, and Roderick Murray-Smith. Variational Inference for Computational Imaging Inverse Problems. apr 2019.
- [7] Chen Zhang and Bangti Jin. Probabilistic Residual Learning for Aleatoric Uncertainty in Image Restoration. aug 2019.
- [8] Jonas Adler and Ozan Öktem. Deep bayesian inversion. *arXiv preprint arXiv:1811.05910*, 2018.
- [9] Dhruv V. Patel and Assad A. Oberai. GAN-based Priors for Quantifying Uncertainty. mar 2020.
- [10] Wenqi Xian, Patsorn Sangkloy, Varun Agrawal, Amit Raj, Jingwan Lu, Chen Fang, Fisher Yu, and James Hays. Texturegan: Controlling deep image synthesis with texture patches. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8456–8465, 2018.
- [11] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016.
- [12] Tero Karras NVIDIA and Samuli Laine NVIDIA. A Style-Based Generator Architecture for Generative Adversarial Networks Timo Aila NVIDIA. Technical report.