

# Lung Cancer Risk Estimation with Incomplete Data: A Joint Missing Imputation Perspective

Riqiang Gao<sup>1</sup>, Yucheng Tang<sup>1</sup>, Kaiwen Xu<sup>1</sup>, Ho Hin Lee<sup>1</sup>, Steve Deppen<sup>2</sup>, Kim Sandler<sup>2</sup>, Pierre Massion<sup>2</sup>, Thomas A. Lasko<sup>1,2</sup>, Yuankai Huo<sup>1</sup>, Bennett A. Landman<sup>1</sup>

<sup>1</sup> EECS, Vanderbilt University, Nashville, TN, USA 37235

<sup>2</sup> Vanderbilt University Medical Center, Nashville, TN, USA 37235


riqiang.gao@vanderbilt.edu

**Abstract.** Data from multi-modality provide complementary information in clinical prediction, but missing data in clinical cohorts limits the number of subjects in multi-modal learning context. Multi-modal missing imputation is challenging with existing methods when 1) the missing data span across heterogeneous modalities (e.g., image vs. non-image); or 2) one modality is largely missing. In this paper, we address imputation of missing data by modeling the joint distribution of multi-modal data. Motivated by partial bidirectional generative adversarial net (PBiGAN), we propose a new Conditional PBiGAN (C-PBiGAN) method that imputes one modality combining the conditional knowledge from another modality. Specifically, C-PBiGAN introduces a conditional latent space in a missing imputation framework that jointly encodes the available multi-modal data, along with a class regularization loss on imputed data to recover discriminative information. To our knowledge, it is the first generative adversarial model that addresses multi-modal missing imputation by modeling the joint distribution of image and non-image data. We validate our model with both the national lung screening trial (NLST) dataset and an external clinical validation cohort. The proposed C-PBiGAN achieves significant improvements in lung cancer risk estimation compared with representative imputation methods (e.g., AUC values increase in both NLST (+2.9%) and in-house dataset (+4.3%) compared with PBiGAN,  $p < 0.05$ ).

**Keywords:** Missing data · Multi-modal · GAN · Lung Cancer

## 1 Introduction

Lung cancer has the highest cancer death rate [1] and early diagnosis with low-dose computed tomography (CT) can reduce the risk of dying from lung cancer by 20% [2,3]. Risk factors (e.g., age and nodule size) are widely used in machine learning and established prediction models [4,5,6,7]. With deep learning techniques, CT image features can be automatically extracted at the nodule-level [8], scan-level [9], or patient-level with longitudinal scans [10]. Previous studies



**Fig. 1.** Missing data in multiple modalities. The upper panel shows a general screening process. In practice, missing data can happen at different phases (as red text). Lower panel shows that patient may miss risk factors or/and follow-up CT scans.

demonstrated that CT image features and risk factors provide complementary information, which is combined to improve lung cancer risk estimation [11].

In the clinical screening process (Fig. 1), patients’ demographic information (e.g., age and gender) is captured in electronic medical records (EMR). In the shared decision-making (SDM) visit, lung cancer risk factors (e.g., smoke status) are collected to determine if a chest CT is necessary. For each performed CT scan, a radiology report is created. Then, such a process might recur according to clinical guidelines. Extensive efforts have been made to collect comprehensive information for patients. However, data can be missing due to multiple issues from data entry, data exchange, data description, et cetera.

Missing data mechanisms were categorized into three types [12]: 1) missing completely at random (MCAR): the missing has no dependency on data, 2) missing at random (MAR): the missing only depends on observed variables, 3) missing not at random (MNAR): the missing may be affected by unobserved variables. To address missing data problems, various imputation approaches were proposed to “make-up” missing data for downstream analyses [13,14,16,15,17,18]. Mean imputation is widely used to fill missing data with population averages. Last observation carried forward (LOCF) [13] takes the last observation as a replacement for missing data, which has been used in clinical serial trials. Soft-imputer [14] provides a convex algorithm for minimizing the reconstruction error corresponding to a bound on the nuclear norm. Recently, deep learning based imputation methods have been developed using generative models [17,18] (e.g., variants of variational auto-encoder (VAE) [19] and generative adversarial net (GAN) [20]). The partial bi-directional GAN (PBiGAN) [18], an encoder-decoder imputation framework, has been validated as a state-of-the-art performance of imputations. However, majority methods have limited imputation within a single modality, which can lead to two challenges in multi-modal context: 1) it is hard to integrate data spanning across heterogeneous modalities (e.g., image vs. non-image) into a single-modal imputation frame-

work, 2) recovering discriminative information is unattainable when data are largely missing in target modality (limiting case: data are completely missing).

We posit that essential information missed in one modality can be maintained in another. In this paper, we propose the Conditional PBiGAN (C-PBiGAN) to model the joint distribution across modalities by introducing 1) a conditional latent space in multi-modal missing imputation context; 2) a class regularization loss to capture discriminative information during imputation. Herein, we focus on lung cancer risk estimation, where risk factors and serial CT scans are two essential modalities for rendering clinical decisions. C-PBiGAN achieves superior predicting performance of downstream multi-modal learning tasks in three broad settings: 1) missing data in image modality, 2) missing data in non-image modality, and 3) both modalities have missing data. With C-PBiGAN, we validate that 1) CT images are conducive to impute missed factors for better risk estimation, and 2) lung nodules with malignancy phenotype can be imputed conditioned on risk factors.


Our contributions are three folds: **(1)** To our knowledge, we are the first to impute missing data by modeling joint distribution of image and non-image data with adversarial training; **(2)** Our model can impute visually realistic data and recover discriminative information, even when the target data in target modality are completely missing; **(3)** Our model achieves superior downstream predicting performance compared with benchmarks with simulated missing (MCAR) and missing in practice (MNAR).

## 2 Theory

**Encoder-Decoder and PBiGAN framework.** PBiGAN [18] is a recently proposed imputation method with encoder-decoder framework based on bidirectional GAN (BiGAN) [21]. Our conditional PBiGAN (C-PBiGAN) is shown in Fig. 2, where the PBiGAN [18] is consist of “black text” components. Note that PBiGAN only deals with a single modality (i.e., modality  $A$  in Fig. 2).

The generator of PBiGAN includes a separate encoder and decoder. The decoder  $g^A$  transforms a latent code  $z$  into a complete data space  $X^A$ , where  $z$  is a feature space (e.g.,  $z_o^A$ ) or sampled from a simple distribution (e.g., Gaussian). The encoder  $q^A(z_o^A|x^A, m)$ , denoted as  $q^A$  for simplification, maps the missing distribution  $p_m$  of an incomplete data  $(x^A, m)$  into a latent vector  $z_o^A$ , where  $x^A \in \mathbb{R}^n$  denotes complete data, and  $m \in \{0, 1\}^n$  is a missing indicator with same dimension of  $x^A$  that determines which entries in  $x^A$  are missing (i.e., 1 for observed, 0 for missing).

The discriminator  $D$  of PBiGAN takes the observed data  $[x^A, m]$  and its corresponding latent code  $z_o^A$  as the “real” tuple in adversarial training. The “fake” tuple  $(\hat{x}^A, \hat{m}, \hat{z})$  is comprised of 1) a random latent code  $\hat{z}$  sampled from a simple distribution  $p_{\hat{z}}$  (e.g., Gaussian), 2) missing indices  $\hat{m}$  from a missing distribution  $p_{\hat{m}}$ , and 3) the generated data  $\hat{x}^A$  based on random latent code  $\hat{z}$ . The loss function of PBiGAN is defined as follows, which is minimax optimized:



**Fig. 2.** Structure of the proposed C-PBiGAN. The orange and green characters highlight our contributions compared with PBiGAN [18].  $m$  is the missing index of target modality  $A$  and  $z$  is the corresponding latent space.  $\tilde{x}^B$  is the complete data of conditional modality  $B$ , which can be fully observed or imputed.  $\tilde{x}^A$  is the imputed data of  $A$  based on observed data  $[x^A, m]$  and  $\tilde{x}^B$ .  $\hat{x}^A$  is the generated data of  $A$  based on  $\tilde{x}^B$  and noise distributions of  $p_{\tilde{z}}$  and  $p_{\hat{m}}$ .  $C$  is a classifying module along with cross-entropy loss regularizing the generator for keeping the identities of imputed data.

$$L(D, g^A, q^A) = \mathbb{E}_{(x^A, m) \sim p_m} \mathbb{E}_{z_o^A \sim q^A(z_o^A | x^A, m)} [\log D(x^A, m, z_o^A)] + \mathbb{E}_{(\cdot, \hat{m}) \sim p_{\hat{m}}} \mathbb{E}_{\hat{z} \sim p_{\tilde{z}}} [\log(1 - D(g^A(\hat{z}, \hat{m}), \hat{m}, \hat{z}))] \quad (1)$$

**The Proposed Conditional PBiGAN.** The original PBiGAN [18] imputes data within a single modality, which does not utilize complementary information from multiple modalities. Herein, we propose C-PBiGAN to impute one modality conditioned on another, and a cross-entropy loss is optimized during generator training to effectively preserve discrimination for imputed data.

As Fig. 2, when imputing  $A$  (target modality), the conditional data  $\tilde{x}^B$  is complete, either fully observed or imputed. Two encoders  $q^A$  and  $q^B$  are used to map data space to latent space for modality  $A$  and  $B$ , respectively. The GAN loss of our method  $L_G(D, g^A, q^A, q^B)$ , also denoted as  $L_G$ , is written as follows:

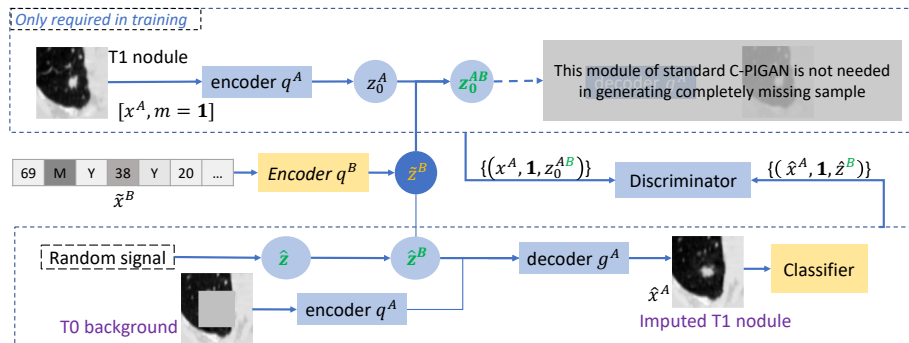
$$L_G = \mathbb{E}_{(x^A, m) \sim p_m} \mathbb{E}_{z_o^{AB} \sim [q^A(z_o^A | x^A, m), q^B(\tilde{z}^B | \tilde{x}^B)]} [\log D(x^A, m, z_o^{AB})] + \mathbb{E}_{(\cdot, \hat{m}) \sim p_{\hat{m}}} \mathbb{E}_{\hat{z}^B \sim [p_{\tilde{z}}, q^B(\tilde{z}^B | \tilde{x}^B)]} [\log(1 - D(g^A(\hat{z}^B, \hat{m}), \hat{m}, \hat{z}^B))] \quad (2)$$

Different from Eq. (1) of PBiGAN focusing on single modality  $A$ , the latent space  $z_o^{AB}$  in Eq.(2) includes the knowledge from two modalities.

To enforce the imputed  $\tilde{x}^A$  or generated  $\hat{x}^A$  having the same identity with  $x^A$  even when data are largely missing, we further introduce a feature extraction net  $C$  along with cross-entropy loss (the second term in Eq. 3) when training the generator. Specifically, C-PBiGAN is optimized with:

$$\min_{g^A, q^A} (\max_D (L_G(D, g^A, q^A, q^B)) - \mathbb{E}_{\tilde{x}^A \sim g^A(\cdot)} [\log p(y|C(\tilde{x}^A))]) \quad (3)$$

where  $y$  is class label and  $p(y|C(\tilde{x}^A))$  is the prediction from  $C$ . Modules  $q^B$ ,  $C$  can be pretrained or trained with  $g^A, q^A$  simultaneously.



**Fig. 3.** An instantiation of limiting C-PBiGAN: imputing TP1 nodule in longitudinal context.  $\tilde{x}^B$  is the imputed risk factor of TP1.  $[x_A, m = 1]$  is complete TP1 data only used in training, as the upper dashed box. “TP0 background” is the observed TP0 (or TP1 in training phase) image with center masked, which is fed to  $q^A$  to make the imputed TP1 with a similar background as TP0. A comparable setting C-PBiGAN<sup>#</sup> is fed with TP0 without masking center.

Different from conditional GAN [22], 1) our model can utilize the partially observed data in the imputation context, and 2) a module  $C$  along with cross-entropy loss is introduced to highlight identity preservation of imputed data.

A limiting case of C-PBiGAN is to impute data that is completely missing (i.e.,  $m = 0$ ). In this case, complete data for training (i.e.,  $m = 1$ ) are needed, and it is the generated  $\hat{x}^A$ , rather than  $\tilde{x}^A$  as in Fig. 2, that used for downstream task. In Eq. (3), the  $\tilde{x}^A$  is replaced with  $\hat{x}^A$ . One of our tasks imputing nodules belongs to this limiting case, as Fig. 3 (details in Section 3).

### 3 Experiment Designs and Results

**Datasets.** We consider two longitudinal CTs (TP0 for previous, TP1 for current) as the complete data for image modality. The non-image modality includes the following 14 risk factors: *age, sex, education, body mass index, race, quit smoke time, smoke status, pack-year, chronic obstructive pulmonary disease, personal cancer history, family lung cancer history, nodule size, spiculation, upper lobe of nodule*. The first two, the middle nine, and the last three factors come from EMR, SDM visit, and radiology report (Fig. 1), respectively.

Two datasets are studied, 1) the national lung screening trail (NLST) [3] and 2) an in-house screening dataset from Vanderbilt Lung Screening Program (VLSF, <https://www.vumc.org/radiology/lung>). Patients in NLST are selected if 1) they have 14 selected risk factors available, 2) have a tissue-based

diagnosis, and 3) the diagnosis happened within 2 years of the last scan if it is a cancer case. *Note that selected subjects are all high-risk patients (all received biopsies), the distinction between cancer / non-cancer in our cohort is hard than in the whole NLST population.* In total, we have 3889 subjects from NLST in which 601 were diagnosed with cancer. 404 subjects from the in-house dataset are evaluated, in which 45 were diagnosed with lung cancer. Due to issues as Fig. 1, the available factors have an average of 32% missing rate, and 60% of patients do not have complete longitudinal scans.

**Method Implementations.** C-PBiGAN has been instantiated to impute risk factors and longitudinal images. Risk factor imputation follows the general C-PBiGAN (Fig. 2), as the factors can be partially observed even when some data are missing. In this case, we only replace modality  $A$  with partially observed risk factors and modality  $B$  with CT in Fig. 2. Image imputation is under the limiting case of C-PBiGAN as described in Section 2 (as Fig. 3), since the “nodule” of interest cannot be partially observed. We follow the C-PBiGAN theory in Section 2 for image imputation, and we also utilize information from longitudinal context in practice. We assume the background of a nodule would not substantially change between TP0 and TP1. Thus, motivated by masking strategies of [24,25], nodule background is borrowed from observed CT (i.e., TP0 image) of the same patient by masking its center when generating the target time point (i.e., TP1 image), see “TP0 background” in Fig. 3. In brief, we target at the problem of missing whole image, while the implementation is kind of central inpainting based on our assumption. We have reconstruction regularization motivated by PBiGAN and UNet [23] skip connections in image-modality implementation.

Given a CT scan, we follow Liao’s pipeline [9] to preprocess the data and detect the top five confidence nodule regions for downstream work. Rather than imputing a whole 3D CT scan, we focus on imputing the nodule areas of interest in 2D context with axial/coronal/sagittal directions as 3 channels (i.e.,  $3 \times 128 \times 128$ ). Considering 1) radiographic reports regarding TP0 are rarely available, and 2) TP1 plays a more important role in lung cancer risk estimation [10], we focus on the imputation on TP1 of image modality in this study. The TP0 image is copied with the TP1 image when TP1 is observed and TP0 is missing.

**Networks.** The structures of encoder, decoder, and discriminator are 1) adapted from face example in PBiGAN [18] for image modality, and 2) separately comprised of four dense layers for non-image modality. A unified multi-modal longitudinal model (MLM), including an image path and a non-image path, is used for lung cancer risk estimation to evaluate the effectiveness of imputations. The image path includes a backbone of ResNeTP18 [26] to extract image features and a LSTM [27] to integrate longitudinal images (from TP0 and TP1). Risk factor features are extracted by a module with four dense layers. The image path and non-image path in the MLM are validated to be effective by comparing with representative prediction models (i.e., AUC in NLST: image-path model (0.875) vs. Liao et al. [9] (0.872) with image data only, non-image path model (0.883) vs.

**Table 1.** AUC results (%) of the test set of NLST (upper, a case of MCAR mechanism) and external in-house set (lower, a case of MNAR mechanism). Generally, each row or each column represents an imputation option for image-missing or risk-factor-missing, respectively. “Image-only” or “factor-only” represents predicting only use imputed longitudinal-images or factors, respectively.

Method	image-only	Mean-imput	Soft-imputer	PBiGAN	<b>C-PBiGAN</b>	fully-observed
test set of longitudinal NLST (30% factors, 50% TP1 image are missing, MCAR)						
factor-only	N/A	79.73	79.46	79.14	83.04	86.24
LOCF	73.45	83.76	83.80	83.79	84.00	86.21
PBiGAN	76.54	83.02	83.82	83.29	83.51	85.90
C-PBiGAN#	82.70	85.00	85.62	85.17	85.87	86.72
<b>C-PBiGAN</b>	84.15	85.72	85.90	85.91	<b>86.20</b>	88.27
fully-observed	87.48	88.23	88.40	88.44	88.46	89.57
external test of in-house dataset (MNAR)						
factor-only	N/A	75.17	83.46	84.40	86.56	N/A
LOCF	75.52	82.83	87.11	86.99	87.63	N/A
PBiGAN	73.44	80.85	84.43	84.88	85.86	N/A
C-PBiGAN#	80.59	83.87	86.57	87.19	87.69	N/A
<b>C-PBiGAN</b>	82.61	85.29	88.11	88.49	<b>89.19</b>	N/A

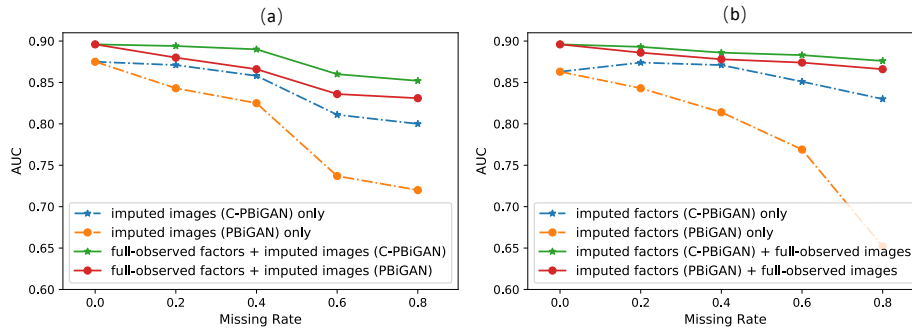
Mayo clinical model [7] (0.829)). The image and non-image features are combined for the final prediction.

**Settings and Evaluations.** The NLST is randomly split into train / validation / test sets with 2340 / 758 / 791 subjects. The in-house dataset of 404 subjects is externally tested when training is finished in NLST. We follow the experimental setup of PBiGAN opensource code [18] when training C-PBiGAN, e.g., use Adam optimizer with a learning rate of  $1e-4$ . The max number of training epochs is set to 200. Our experiments are performed with Python 3.7 and PyTorch 1.5 on GTX Titan X. The mask size of “TP0 background” is  $64 \times 64$ . The area under the receiver operating characteristic (AUC) [28] for lung cancer risk estimation is used to quantitatively evaluate the effectiveness of imputations.


**Imputation Baselines.** Representative imputations (introduced in Sec. 1) of image (i.e., LOCF [13] and PBiGAN [18]) and non-image (i.e., mean imputation, soft-imputer [14] and PBiGAN [18]) are combined for comparison as in Table 1. As a comparable setting of ours, C-PBiGAN# denotes feeding TP0 module without masking the center, rather than “TP0 background” in Fig. 3.

**Results and Discussion.** Table 1 shows 1) tests of NLST (upper) with 30% of missing in risk factors and 50% of missing in longitudinal TP1 and 2) external tests of in-house data with missing in practice. The C-PBiGAN combination (**bold** in Table 1) significantly improves all imputation combinations without C-PBiGAN across the image and non-image modalities ( $p < 0.05$ , bootstrapped two-tailed test,  $n=2000$  [29]) in both NLST and external clinical dataset (e.g., C-PBiGAN increases 4.3% AUC on PBiGAN in the external cohort). Those indicate our model effectively imputes data when missing in both modalities for cancer risk estimation.

Fig. 4 compares proposed C-PBiGAN with PBiGAN in terms of the lung cancer predicting performance in NLST under (a) various TP1 missing rates when factors are fully observed, (b) various factor missing rates when longitudinal images are fully observed. Our model outperforms PBiGAN in the image-missing



**Fig. 4.** (a) AUCs of various TP1-image missing rates when factors are fully observed in NLST, and (b) AUCs of various factor missing rates when images are fully observed in NLST. The left start point is under condition that data are not missing.



**Fig. 5.** Qualitative results of imputed TP1 nodules (upper: malignant, bottom: benign). Malignant/benign cases from C-PBiGAN are most distinguishable.

and factor-missing contexts of different rates. A more obvious superiority can be found when only using the imputed modality for prediction (e.g., C-PBiGAN: 0.830 vs. PBiGAN: 0.652 when risk factors have missing rate of 80%), and the imputed factors conditioned on images can even achieve higher AUC than the fully observed factors at some missing rates. Those indicate the information from conditional modality in C-PBiGAN does help the imputation.

Fig. 5 shows malignant and benign cases from NLST and in-house dataset. Both PBiGAN and proposed C-PBiGAN can reconstruct visually realistic images, while malignant and benign cases from PBiGAN are harder to distinguish.

As a comparable setting, C-PBiGAN# is less effective than C-PBiGAN (Table 1, Fig. 5) given the current setting and network structure. It is probably because when feeding TP0 without masking center to provide nodule background (i.e., C-PBiGAN#), the central nodule region of imputed TP1 can be fit to the center of TP0, just like the nodule background of imputed TP1 is designed to fit TP0 nodule background. This limits the discrimination of imputed TP1, as the examples in Fig. 5. Thus, it is essential to separate “background” and “nodule” during learning, since we want the “background” of imputed TP1 to be close to observed TP0 while the “nodule” of imputed TP1 should mainly be conditioned



on risk factors. Motivated by strategies in [24,25], our C-PBiGAN is fed with TP0 background masking the center when imputing the TP1 (in Fig. 3).

## 4 Conclusion

We propose a novel deep learning based missing imputation model for multi-modal data. By modeling the joint distribution of multiple modalities, the proposed C-PBiGAN can effectively impute the missing data across image and non-image modalities. We validate our method on a large-scale NLST dataset (MCAR) and an external clinical cohort (MNAR). Given no restriction on data type, our model can be readily extended to other multi-modal missing contexts.

**Acknowledgements.** This research was supported by NSF CAREER 1452485, R01 EB017230 and R01 CA253923. This study was supported in part by U01 CA196405 to Massion. This project was supported in part by the National Center for Research Resources, Grant UL1 RR024975-01, and is now at the National Center for Advancing Translational Sciences, Grant 2 UL1 TR000445-06. This study was funded in part by the Martineau Innovation Fund Grant through the Vanderbilt-Ingram Cancer Center Thoracic Working Group and NCI Early Detection Research Network 2U01CA152662 to PPM.

## References

1. Siegel, R.L., Miller, K.D., Jemal, A.: Cancer statistics, 2019. *CA. Cancer J. Clin.* 69, 7–34 (2019).
2. Aberle, D.R., Adams, A.M., Berg, C.D., Black, W.C., Clapp, J.D., Fagerstrom, R.M., Gareen, I.F., Gatsonis, C., Marcus, P.M., Sicks, J.R.D.: Reduced lung-cancer mortality with low-dose computed tomographic screening. *N. Engl. J. Med.* 365, 395–409 (2011).
3. N.L.S.T.R.T.J.: The national lung screening trial: Overview and study design. *Radiology.* 258, 243–253 (2011).
4. Huang, P. et al.: Prediction of lung cancer risk at follow-up screening with low-dose CT: a training and validation study of a deep learning method. *Lancet Digit. Heal.* 1, e353–e362 (2019).
5. Tammemägi, M.C., Katki, H.A., Hocking, W.G., Church, T.R., Caporaso, N., Kvale, P.A., Chaturvedi, A.K., Silvestri, G.A., Riley, T.L., Commins, J., Berg, C.D.: Selection criteria for lung-cancer screening. *N. Engl. J. Med.* 368, 728–736 (2013).
6. Swensen, S.J.: The Probability of Malignancy in Solitary Pulmonary Nodules. *Arch. Intern. Med.* 157, 849 (1997).
7. McWilliams, A. et al.: Probability of cancer in pulmonary nodules detected on first screening CT. *N. Engl. J. Med.* 369, 910–919 (2013).
8. Liu, L., Dou, Q., Chen, H., Qin, J., Heng, P.A.: Multi-Task Deep Model with Margin Ranking Loss for Lung Nodule Analysis. *IEEE Trans. Med. Imaging.* 39, 718–728 (2020).
9. Liao, F., Liang, M., Li, Z., Hu, X., Song, S.: Evaluate the Malignancy of Pulmonary Nodules Using the 3-D Deep Leaky Noisy-or Network. *IEEE Trans. Neural Networks Learn. Syst.* 1–12 (2019).

10. Gao, R., Tang, Y., Xu, K., Huo, Y., Bao, S., Antic, S.L., Epstein, E.S., Deppen, S., Paulson, A.B., Sandler, K.L., Massion, P.P., Landman, B.A.: Time-Distanced Gates in Long Short-Term Memory Networks. *Med. Image Anal.* 65, 101785 (2020).
11. Gao, R. et al.: Deep Multi-path Network Integrating Incomplete Biomarker and Chest CT Data for Evaluating Lung Cancer Risk. arxiv 2010.09524. (2021)
12. Rubin, D.B.: Inference and missing data. *Biometrika.* 63, 581–592 (1976). <https://doi.org/10.1093/biomet/63.3.581>.
13. Van Buuren, S.: Flexible imputation of missing data. CRC Press. (2018).
14. Mazumder, R., Hastie, T., Edu, H., Tibshirani, R., Edu, T., Jaakkola, T.: Spectral Regularization Algorithms for Learning Large Incomplete Matrices. *J. Mach. Learn. Res.* 11, 2287–2322 (2010).
15. Yoon, J., Jordon, J., Van Der Schaar, M.: GAIN: Missing data imputation using generative adversarial nets. In: International Conference on Machine Learning. pp. 9042–9051. International Machine Learning Society (IMLS) (2018).
16. Stekhoven, D.J., Bühlmann, P.: Missforest-Non-parametric missing value imputation for mixed-type data. *Bioinformatics.* 28, 112–118 (2012).
17. Mattei, P.A., Freiißen, J.: Miwae: Deep generative modelling and imputation of incomplete data sets. In: 36th International Conference on Machine Learning, ICML 2019. pp. 7762–7772 (2019).
18. Cheng, S., Li, -Xian, Marlin, B.M.: Learning from Irregularly-Sampled Time Series: A Missing Data Perspective. *Int. Conf. Mach. Learn.* (2020).
19. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: International Conference on Learning Representations. International Conference on Learning Representations, ICLR (2014).
20. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in Neural Information Processing Systems. pp. 2672–2680 (2014).
21. Donahue, J., Krähenbühl, P., Darrell, T.: Adversarial Feature Learning. (2016).
22. Mirza, M., Osindero, S.: Conditional Generative Adversarial Nets. arXiv Prepr. arXiv1411.1784. (2014).
23. Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. "U-net: Convolutional networks for biomedical image segmentation." International Conference on Medical image computing and computer-assisted intervention. Springer, Cham, 2015.
24. Jin, D., Xu, Z., Tang, Y., Harrison, A.P., Mollura, D.J.: CT-Realistic Lung Nodule Simulation from 3D Conditional Generative Adversarial Networks for Robust Lung Segmentation. *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics).* 11071 LNCS, 732–740 (2018).
25. Mirsky, Y., Mahler, T., Shelef, I., Elovici, Y.: CT-GAN: Malicious Tampering of 3D Medical Imagery using Deep Learning. *Proc. 28th USENIX Secur. Symp.* 461–478 (2019).
26. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016).
27. Hochreiter, S., Schmidhuber, J.: Long Short-Term Memory. *Neural Comput.* 9, 1735–1780 (1997). <https://doi.org/10.1162/neco.1997.9.8.1735>.
28. Fawcett, T.: An introduction to ROC analysis. *Pattern Recognit. Lett.* 27, 861–874 (2006). <https://doi.org/10.1016/j.patrec.2005.10.010>.
29. Mateusz buda: Statistical functions based on bootstrapping for computing confidence intervals and p-values comparing machine learning models and human readers, <https://github.com/mateusz buda/ml-stat-util>, last accessed 2021/02/27.