

CREDANNO+: ANNOTATION EXPLOITATION IN SELF-EXPLANATORY LUNG NODULE DIAGNOSIS

Jiahao Lu^{1,2} Chong Yin³ Kenny Erleben¹ Michael Bachmann Nielsen² Sune Darkner¹

¹Department of Computer Science, University of Copenhagen, Denmark

²Department of Diagnostic Radiology, Rigshospitalet, Copenhagen University Hospital, Denmark

³Department of Computer Science, Hong Kong Baptist University, China

ABSTRACT

Recently, attempts have been made to reduce annotation requirements in feature-based self-explanatory models for lung nodule diagnosis. As a representative, cRedAnno achieves competitive performance with considerably reduced annotation needs by introducing self-supervised contrastive learning to do unsupervised feature extraction. However, it exhibits unstable performance under scarce annotation conditions. To improve the accuracy and robustness of cRedAnno, we propose an annotation exploitation mechanism by conducting semi-supervised active learning with sparse seeding and training quenching in the learned semantically meaningful reasoning space, to jointly utilise the extracted features, annotations, and unlabelled data. The proposed approach achieves comparable or even higher malignancy prediction accuracy with 10x fewer annotations, meanwhile showing better robustness and nodule attribute prediction accuracy under the condition of 1% annotations. Our complete code is open-source available: <https://github.com/diku-dk/credanno>.

Index Terms— Explainable AI, Lung nodule diagnosis, Self-explanatory model, Active learning, Semi-supervised learning

1. INTRODUCTION

Effective lung cancer screening requires accurate characterisation of pulmonary nodules in CT images [1]. Amongst recent efforts in explainable AI, post-hoc approaches that attempt to interpret well-performing “black boxes” [2] are not deemed trustworthy enough in clinical practice [3, 4]. In contrast, feature-based self-explanatory methods are trained to use a set of well-known human-understandable concepts to explain and derive their decisions [5, 6]. Although such semantic matching towards clinical knowledge is especially valuable in medical applications [7], the additional annotation requirements for features still limit the applicability of this approach.

The recently proposed cRedAnno [8] addresses this problem by introducing self-supervised contrastive learning [9] to alleviate the burden of learning most parameters from annotations. Albeit cRedAnno [8] achieves competitive performance

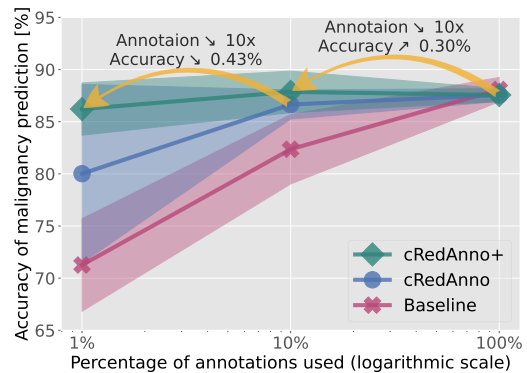


Fig. 1. Comparison between cRedAnno+ and cRedAnno, in terms of nodule malignancy prediction accuracy and annotation cost. cRedAnno+ achieves comparable or even higher accuracy with 10x fewer annotations, meanwhile being more robust.

using hundreds of nodule samples and considerably reduced annotations, its robustness awaits improvement when labelled samples are extremely scarce (Fig. 1), whereas the unlabelled data are not adequately utilised yet [10].

This paper aims to improve the prediction accuracy and robustness of cRedAnno [8] under scarce annotation conditions. We address this by conducting semi-supervised active learning [11] in the learned latent space that complies with radiologists’ reasoning for nodule malignancy. More specifically, we propose an efficient annotation exploitation mechanism, composing seeding by clustering [12], uncertainty sampling [13], and pseudo labelling [14] to jointly utilise the extracted features, annotations, and unlabelled data, facilitated by a quenching technique to update the pseudo labels and reinitialise the weights of predictors, as shown in Fig. 2.

Compared with cRedAnno [8], the proposed cRedAnno+ achieves comparable or even higher malignancy prediction accuracy with 10x fewer annotations while reaching simultaneously above 90% mean accuracy in predicting all nodule attributes, meanwhile being more robust under the condition of 1% annotations (Fig. 1).

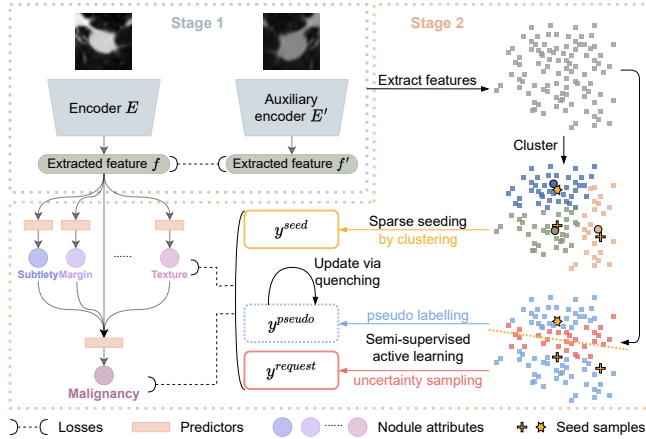


Fig. 2. Concept illustration. The proposed annotation exploitation mechanism conducts semi-supervised active learning with sparse seeding and training quenching in the semantically meaningful space learned in Stage 1, to jointly exploit the extracted features, annotations, and unlabelled data.

2. METHOD

As illustrated in Fig. 2, the proposed approach enriches cRedAnno [8] by replacing its second-stage supervised predictor training using random samples with an annotation exploitation mechanism. Therefore, here we only outline what is inherited from cRedAnno, whilst detailing the proposed approach and focusing on their differences.

2.1. Recapitulation of cRedAnno

The original cRedAnno [8] uses two-stage training, replacing the end-to-end training paradigm in previous works. In Stage 1, the majority of parameters are trained using self-supervised contrastive learning [9] as an encoder to map the input images to a latent space that complies with radiologists’ reasoning for nodule malignancy. In Stage 2, a small random portion of labelled samples is used to train a hierarchical set of predictors $\{G_{cls}, G_{exp}\}$, including a predictor $G_{exp}^{(i)}$ for each nodule attribute i , and a malignancy predictor G_{cls} whose input is the concatenation of extracted features f and the predicted human-understandable nodule attributes.

Limitations of cRedAnno [8]. The high prediction accuracy exhibited by cRedAnno [8] can attribute to the extracted features being highly separable in the learned space. Nevertheless, when only very few annotations are available, the random selection is likely biased and even risks not covering enough label space, which may lead to severe performance degradation. Furthermore, for training the predictors on the self-supervised model, randomly selected annotations are not necessarily informative enough [11]. In addition, the unlabelled data are not adequately used in training the predictors.

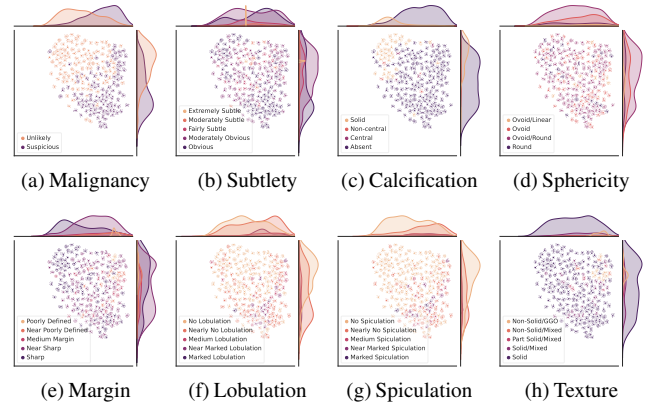


Fig. 3. t-SNE visualisation of features extracted from test images. Data points are coloured using ground truth annotations. Whilst malignancy shows highly separable and semantically correlates with the nodule attributes, a random selection of few samples may risk not covering enough label space.

2.2. Annotation exploitation mechanism

To jointly utilise the extracted features, annotations, and unlabelled data, we propose an annotation exploitation mechanism (Fig. 2) integrating the following components to address the aforementioned limitations.

Sparse seeding. To mitigate potential bias and randomness, we select seed samples by clustering [12] the extracted features in the learned space, which was underutilised in the previous cRedAnno [8]. The extracted features are clustered into n clusters, where n equals the number of seed samples to select. Then the sample closest to each cluster centroid (based on cosine similarity) is selected as f^{seed} , whose annotations $y^{seed} = \{y_{cls}^{seed}, y_{exp}^{seed}\}$ are used to train the predictors from the initial status st_0 to the seeded status st_1 :

$$G^{st_1} = \operatorname{argmin}_{\{G_{cls}, G_{exp}\}} \mathcal{L}(y^{seed}, \operatorname{softmax}(G^{st_0}(f^{seed}))), \quad (1)$$

where \mathcal{L} denotes the cross-entropy loss.

Semi-supervised active learning. Semi-supervised learning and active learning are conducted simultaneously [14] to exploit the available data. We adopt the classic yet effective uncertainty sampling by least confidence as acquisition strategy [13] to request annotations $y^{request}$ for the uncertain/informative samples $f^{request}$. Concurrently, other samples with relatively high confidence are assigned with pseudo annotations $y^{pseudo(st_t)}$ by the prediction of G^{st_t} at status st_t .

Quenching. To facilitate training under the restrictions of limited samples and complex annotation space, we propose “quenching” as a training technique. Similar to Curriculum Pseudo Labelling [17, 18], at a certain status st_t since st_1 , the predictor weights are reinitialised to G^{st_0} to avoid potential confirmation bias [19]. Meanwhile, the pseudo annotations

Table 1. Prediction accuracy (%) of nodule attributes and malignancy. The best in each column is **bolded** for full/partial annotation respectively. Dashes (-) denote values not reported by the compared methods. Results of our proposed **cRedAnno+** and the previous **cRedAnno** [8] are highlighted. Observe that with 1% annotations, cRedAnno+ reaches competitive accuracy in malignancy prediction and over 90% accuracy simultaneously in predicting all nodule attributes, meanwhile using the fewest nodules and no additional information.

	Nodule attributes							Malignancy	#nodules	No additional information
	Sub	Cal	Sph	Mar	Lob	Spi	Tex			
Full annotation										
HSCNN [5]	71.90	90.80	55.20	72.50	-	-	83.40	84.20	4252	χ^c
X-Caps [6]	90.39	-	85.44	84.14	70.69	75.23	93.10	86.39	1149	\checkmark
MSN-JCN [15]	70.77	94.07	68.63	78.88	94.75	93.75	89.00	87.07	2616	χ^d
MTMR [2]	-	-	-	-	-	-	-	93.50	1422	χ^e
cRedAnno+	96.32±0.61	95.88±0.15	97.23±0.20	96.23±0.23	93.93±0.87	94.06±0.60	97.01±0.26	87.56±0.61	730	\checkmark
Partial annotation										
WeakSup [16] (1:5) ^a	43.10	63.90	42.40	58.50	40.60	38.70	51.20	82.40	2558	χ^f
WeakSup [16] (1:3) ^a	66.80	91.50	66.40	79.60	74.30	81.40	82.20	89.10		
cRedAnno (10%^b)	96.06±2.02	93.76±0.85	95.97±0.69	94.37±0.79	93.06±0.27	93.15±0.33	95.49±0.85	86.65±1.39		
cRedAnno+ (10%^b)	96.23±0.45	92.72±1.66	95.71±0.47	90.03±3.68	93.89±1.41	93.67±0.64	92.41±1.05	87.86±1.99	730	\checkmark
cRedAnno (1%^b)	93.98±2.09	89.68±3.52	94.02±2.30	91.94±1.17	91.03±1.72	90.81±1.56	93.63±0.47	80.02±8.56		
cRedAnno+ (1%^b)	95.84±0.34	92.67±1.24	95.97±0.45	91.03±4.65	93.54±0.87	92.72±1.19	92.67±1.50	86.22±2.51		

^a 1 : N indicates that $\frac{1}{1+N}$ of training samples have annotations on nodule attributes. (All samples have malignancy annotations.)

^b The proportion of training samples that have annotations on nodule attributes and malignancy. ^c 3D volume data are used. ^d Segmentation masks and nodule diameter information are used. Two other traditional methods are used to assist training. ^e All 2D slices in 3D volumes are used. ^f Multi-scale 3D volume data are used.

are updated to the current prediction results:

$$G^{st_{t+1}} = \underset{\{G_{cls}, G_{exp}\}}{\operatorname{argmin}} \mathcal{L}(\{y^{request}, y^{pseudo}(st_t)\}, \operatorname{softmax}(G^{sto}(\{f^{request}, f^{pseudo}\}))), \quad (2)$$

to preserve the learned information and resume training.

3. EXPERIMENTAL RESULTS

Data pre-processing. The data pre-processing remains the same as in cRedAnno [8], following the common pre-processing procedure of the LIDC dataset [20] summarised in [21], resulting in 276/242 benign/malignant nodules for training and 108/104 benign/malignant nodules for testing.

Training settings. Our training settings in Stage 1 remain the same as in cRedAnno [8]. In Stage 2, K-means are used for clustering to select 1% annotation as seed samples. The predictors $G_{exp}^{(i)}$ and G_{cls} , each consisting of one linear layer, are first jointly trained using the seed samples for 100 epochs with SGD optimisers with momentum 0.9 and batch size 128. The learning rate follows a cosine scheduler with initial value 0.00025. After reaching the seeded status G^{st_1} , the predictors and optimisers are quenched for the first time. The training then resumes using the requested and dynamic pseudo annotations for 50 more epochs, where quenching happens every 10 epochs.

3.1. Analysis of extracted features in the learned space

The reduction of annotations relies heavily on the separability and semantic information of the learned feature space. We use t-SNE to visualise the learned feature as a qualitative evaluation. Feature f extracted from each testing image is mapped to

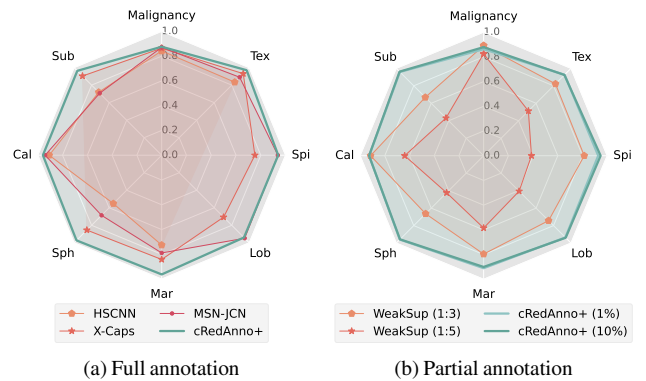


Fig. 4. Performance comparison, in terms of prediction accuracy (%) of nodule attributes and malignancy. Observe that **cRedAnno+** achieves simultaneously high accuracy in predicting malignancy and all nodule attributes, regardless of using either full or partial annotations.

a data point in 2D space. Fig. 3a to 3h correspond to these data points coloured by the ground truth annotations of malignancy to nodule attribute “texture”, respectively.

Fig. 3 intuitively demonstrates the underlying correlation between malignancy and nodule attributes. For instance, the cluster in Fig. 3c indicates that solid calcification negatively correlates with nodule malignancy. Similarly, Fig. 3f and Fig. 3g indicate that lobulation is associated with spiculation, both of which are positively correlated with malignancy. These semantic correlations coincide with the radiologists’ diagnostic process [1] and thereby further support the potential of the proposed approach as a trustworthy decision support system.

More importantly, Fig. 3a shows that even in this 2D space

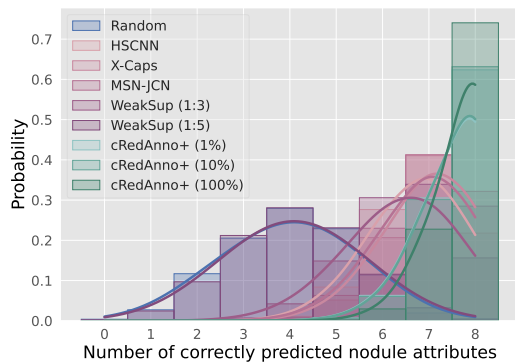


Fig. 5. Probabilities of correctly predicting a certain number of attributes for a given nodule sample. Observe that **cRedAnno+** shows a more prominent probability of simultaneously predicting all 8 nodule attributes correctly.

for visualisation, the samples show reasonably separable in both malignancy and nodule attributes. This provides the possibility to train the initial predictors using only a very small number of seed annotations, provided they are sufficiently dispersed and informative.

3.2. Prediction performance of nodule attributes and malignancy

The performance is evaluated quantitatively in terms of prediction accuracy of nodule attributes and malignancy. The evaluation procedure is the same as in the previous cRedAnno [8]: each annotation is considered independently [5]; the predictions of nodule attributes are considered correct if within ± 1 of aggregated radiologists’ annotation [6]; attribute “internal structure” is excluded from the results because its heavily imbalanced classes are not very informative [2, 5, 6, 15, 16].

Tab. 1 summarises the overall prediction performance and compares it with the state-of-the-art. The results show that when using only 518 among the 730 nodule samples and 1% of their annotations for training, cRedAnno+ reaches over 90% accuracy simultaneously in predicting all nodule attributes, which outperforms all previous works. Meanwhile, regarding nodule malignancy prediction accuracy, cRedAnno+ performs comparably with X-Caps [6] and already better than HSCNN [5], which uses 3D volume data. When using 10% annotations, our malignancy prediction accuracy surpasses all other explainable competitors using full annotations, among which MSN-JCN [15] is heavily supervised by additional information. Compared with the previous cRedAnno [8], cRedAnno+ achieves comparable or even higher malignancy prediction accuracy with 10x fewer annotations, as shown in Fig. 1.

The visualisation of the performance comparison is shown in Fig. 4. It can be observed that our approach demonstrates simultaneously high prediction accuracy in malignancy and all nodule attributes. This substantially increases the model’s

Table 2. Ablation study of proposed components, evaluated by the prediction accuracy of malignancy using 10% and 1% annotations. The best in each column is **bolded. Settings of our proposed **cRedAnno+** and the previous **cRedAnno** [8] are highlighted.**

Seed sample selection	Annotation acquisition strategy	Pseudo labelling	Quenching	Malignancy accuracy	
				(10%)	(1%)*
random	✗	✗	✗	86.65±1.39	80.02±8.56
random	malignancy confidence	dynamic	✓	82.71±7.47	79.50±11.10
sparse	integrated entropy	dynamic	✓	86.52±0.99	86.22±2.51
sparse	malignancy confidence	static	✗	85.91±1.66	85.35±1.93
sparse	malignancy confidence	dynamic	✓	87.86±1.99	86.22±2.51

* Does not contain requested annotations.

trustworthiness and has not been achieved in previous works.

In addition, we also calculate the probabilities of correctly predicting a certain number of attributes for a given nodule sample, as shown in Fig. 5. The probabilities are calculated from Tab. 1. To not underestimate the performance of other compared methods, their not reported values are all assumed to be 100% accuracy. It can be seen that cRedAnno+ demonstrates a more prominent probability of simultaneously predicting all 8 nodule attributes correctly. The probability of correctly predicting at least 7 attributes is higher than 90%, even under the extreme 1% annotation condition. In contrast, WeakSup(1:5), [16] despite achieving 82.4% accuracy in malignancy prediction, shows no significant difference in predicting nodule attributes compared to random guessing, which we consider to be the opposite of trustworthiness.

3.3. Ablation study

We validate the proposed annotation exploitation mechanism by ablating each component as a row shown in Tab. 2. The standard deviation when using 1% annotations shows that sparse seeding plays a crucial role in stabilising performance. The sum entropy [13] integrating malignancy and all nodule attributes was also experimented with as an alternative acquisition strategy, but exhibits impaired prediction accuracy. Quenching, which enables dynamic pseudo labelling, also proves necessary for the boosted performance.

4. CONCLUSION

In this paper, we propose cRedAnno+ to improve the prediction accuracy and robustness of the previous work in self-explanatory lung nodule diagnosis. Our experiments show that the proposed annotation exploitation mechanism enables comparable or even higher accuracy with 10x fewer annotations, meanwhile being more robust. Furthermore, cRedAnno+ is the first to reach over 90% accuracy simultaneously in predicting all nodule attributes with only hundreds of samples and 1% of their annotations, which adds significantly to the trustworthiness. The limitations of this work remain in its simple predictor architecture and not carefully-tuned hyperparameters, and its generalisability is yet to be verified on other suitable datasets.

5. ACKNOWLEDGEMENTS

No funding was received for conducting this study. The authors have no relevant financial or non-financial interests to disclose.

6. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted using the open-access LIDC dataset [20]. Ethical approval was not required as confirmed by the license attached with the open-access data.

7. REFERENCES

- [1] Ioannis Vlahos, Konstantinos Stefanidis, et al., “Lung cancer screening: Nodule identification and characterization,” *Translational Lung Cancer Research*, vol. 7, no. 3, pp. 288–303, June 2018.
- [2] Lihao Liu, Qi Dou, et al., “Multi-Task Deep Model With Margin Ranking Loss for Lung Nodule Analysis,” *IEEE Transactions on Medical Imaging*, vol. 39, no. 3, pp. 718–728, Mar. 2020.
- [3] Cynthia Rudin, “Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead,” *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, May 2019.
- [4] Bas H.M. van der Velden, Hugo J. Kuijf, et al., “Explainable artificial intelligence (XAI) in deep learning-based medical image analysis,” *Medical Image Analysis*, vol. 79, pp. 102470, July 2022.
- [5] Shiwen Shen, Simon X Han, et al., “An interpretable deep hierarchical semantic convolutional neural network for lung nodule malignancy classification,” *Expert Systems with Applications*, vol. 128, pp. 84–95, Aug. 2019.
- [6] Rodney LaLonde, Drew Torigian, and Ulas Bagci, “Encoding Visual Attributes in Capsules for Explainable Medical Diagnoses,” in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, Cham, 2020, Lecture Notes in Computer Science, pp. 294–304, Springer International Publishing.
- [7] Zohaib Salahuddin, Henry C. Woodruff, et al., “Transparency of deep neural networks for medical image analysis: A review of interpretability methods,” *Computers in Biology and Medicine*, vol. 140, pp. 105111, Jan. 2022.
- [8] Jiahao Lu, Chong Yin, et al., “Reducing Annotation Need in Self-explanatory Models for Lung Nodule Diagnosis,” in *Interpretability of Machine Intelligence in Medical Image Computing*, vol. 13611, pp. 33–43. Springer Nature Switzerland, Cham, 2022.
- [9] Mathilde Caron, Hugo Touvron, et al., “Emerging Properties in Self-Supervised Vision Transformers,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9650–9660.
- [10] Oriane Simeoni, Mateusz Budnik, et al., “Rethinking deep active learning: Using unlabeled data at model training,” in *2020 25th International Conference on Pattern Recognition (ICPR)*, Milan, Italy, Jan. 2021, pp. 1220–1227, IEEE.
- [11] Alex Tamkin, Dat Nguyen, et al., “Active Learning Helps Pretrained Models Learn the Intended Task,” Apr. 2022.
- [12] Rong Hu, Brian Mac Namee, and Sarah Jane Delany, “Off to a Good Start: Using Clustering to Select the Initial Training Set in Active Learning,” in *Twenty-Third International FLAIRS Conference*, May 2010.
- [13] Burr Settles, *Uncertainty Sampling*, pp. 11–20, Springer International Publishing, Cham, 2012.
- [14] Keze Wang, Dongyu Zhang, et al., “Cost-Effective Active Learning for Deep Image Classification,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 27, no. 12, pp. 2591–2600, Dec. 2017.
- [15] Wei Chen, Qiuli Wang, et al., “End-to-End Multi-Task Learning for Lung Nodule Segmentation and Diagnosis,” in *2020 25th International Conference on Pattern Recognition (ICPR)*, Milan, Italy, Jan. 2021, pp. 6710–6717, IEEE.
- [16] Aniket Joshi, Jayanthi Sivaswamy, and Gopal Datt Joshi, “Lung nodule malignancy classification with weakly supervised explanation generation,” *Journal of Medical Imaging*, vol. 8, no. 04, Aug. 2021.
- [17] Paola Cascante-Bonilla, Fuwen Tan, et al., “Curriculum Labeling: Revisiting Pseudo-Labeling for Semi-Supervised Learning,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 8, pp. 6912–6920, May 2021.
- [18] Bowen Zhang, Yidong Wang, et al., “FlexMatch: Boosting Semi-Supervised Learning with Curriculum Pseudo Labeling,” in *Advances in Neural Information Processing Systems*. 2021, vol. 34, pp. 18408–18419, Curran Associates, Inc.
- [19] Eric Arazo, Diego Ortego, et al., “Pseudo-Labeling and Confirmation Bias in Deep Semi-Supervised Learning,” in *2020 International Joint Conference on Neural Networks (IJCNN)*, Glasgow, United Kingdom, July 2020, pp. 1–8, IEEE.
- [20] Samuel G. Armato, Geoffrey McLennan, et al., “The Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI): A Completed Reference Database of Lung Nodules on CT Scans: The LIDC/IDRI thoracic CT database of lung nodules,” *Medical Physics*, vol. 38, no. 2, pp. 915–931, Jan. 2011.
- [21] Vasileios Baltatzis, Kyriaki-Margarita Bintsi, et al., “The Pitfalls of Sample Selection: A Case Study on Lung Nodule Classification,” in *Predictive Intelligence in Medicine*, vol. 12928, pp. 201–211. Springer International Publishing, Cham, 2021.