

SecondOpinion: Anatomy-Aware Gated Reasoning for Efficient Medical Image Analysis

Siam Tahsin Bhuiyan^{1,2}[0009-0000-1298-1991], Rashedur Rahman^{1,2*}[0000-0003-0267-2612], Sefatul Wasi²[0009-0004-6949-0607], Riyadhul Islam¹[0009-0007-8702-9330], Syoji Kobashi³[0000-0003-3659-4114], Ashraful Islam^{1,2}[0000-0003-2367-2013], Saadia Binte Alam^{1,2}[0009-0007-0358-7635]

¹ Center for Computational & Data Sciences, Independent University, Bangladesh, Dhaka, Bangladesh

² Department of Computer Science and Engineering, Independent University, Bangladesh, Dhaka, Bangladesh

³ Graduate School of Engineering, University of Hyogo, Kobe, Japan
rashed@iub.edu.bd

Abstract. Deep learning models for medical image analysis typically apply a fixed amount of computation to every input, regardless of case difficulty. Anatomy-guided dual-stream architectures have been shown to improve diagnostic performance, but they evaluate both streams unconditionally, even on cases a single stream could already resolve confidently. We propose SecondOpinion, a framework in which a fast primary stream processes every case, while a second, anatomy-guided stream is invoked only when GateKeeper, a gating mechanism trained explicitly as a binary correctness classifier, judges that the primary stream’s prediction needs additional scrutiny, much as a clinician might seek a second opinion on a difficult case. When activated, the two streams are combined through a lightweight cross-attention fusion module. We evaluate SecondOpinion on a unified five-class chest X-ray dataset and a pelvic fracture dataset, the latter including a held-out, harder subset of fractures that are invisible on X-ray but confirmed via CT. SecondOpinion matches or exceeds prior state-of-the-art performance on both tasks, while activating its anatomy-guided stream on only 9.23% of chest X-ray cases, rising to 24.12% on visible fractures and 45.71% on invisible fractures, an activation rate that tracks task difficulty directly. These results suggest that supervising a gating signal toward correctness, rather than relying on unsupervised confidence, allows a model to allocate anatomical reasoning where it is actually needed.

Keywords: Medical Image Analysis · Conditional Computation · Anatomy-guided Learning · Dual-stream Networks · Chest X-ray Classification · Fracture Detection.

* Corresponding author

1 Introduction

Computer-aided diagnosis from medical images has advanced rapidly, with deep learning models achieving expert-comparable performance across diverse imaging modalities and pathologies [12]. Most diagnostic models, however, apply a fixed amount of computation to every input regardless of case difficulty: an obvious finding is processed at the same cost as an ambiguous one. In clinical practice, diagnostic effort is naturally adaptive, with straightforward cases resolved quickly while ambiguous ones prompt a clinician to seek a second opinion. This motivates our central question: can a diagnostic model learn to allocate additional anatomical reasoning only when it is actually needed?

Anatomy-guided deep learning has been shown to improve recognition performance, particularly for spatially sparse or low-contrast structures [6], by infusing anatomic knowledge into the network design itself rather than leaving it to be discovered unsupervised [2]. PelFANet [2] follows this direction for pelvic fracture detection, using anatomy-aware feature guidance to improve sensitivity to subtle fractures. Dual-stream architectures are a common mechanism for incorporating such guidance, processing anatomical and pathological information through separate representations before fusion [19, 4]. A consistent limitation across this line of work is that both streams are evaluated on every input unconditionally, doubling inference cost even for the large fraction of cases a single stream could already resolve confidently.

A separate line of work addresses efficiency directly through conditional computation, adapting computational effort to each input rather than reducing cost uniformly for all inputs [8]. Early-exit networks are the most established instance of this idea, using a gating mechanism to decide whether further computation can be skipped once a confident prediction is reached [18, 7]. Recent work shows that aligning training dynamics with the inference-time gating policy improves early-exit reliability [7], motivating our use of explicit, ground-truth-supervised confidence signals rather than relying on implicit task-loss gradients alone.

While confidence-gated dual-path designs exist in retrieval-augmented classification [16], correctness-supervised gating over anatomy-guided streams has not been explored. We bridge this gap with SecondOpinion: a fast primary stream handles every case, and an anatomy-guided second stream is invoked only when GateKeeper, a learned gating mechanism, flags low confidence in the primary stream’s prediction, much as a clinician might seek a second opinion on a difficult case.

Our contributions are as follows:

1. We propose SecondOpinion, an adaptive dual-stream framework in which an anatomy-guided "second opinion" stream is invoked conditionally, rather than on every input as in prior dual-stream designs.
2. We introduce GateKeeper, a lightweight gating mechanism trained explicitly as a binary correctness classifier, predicting the probability that the primary stream’s prediction is correct, rather than relying on an unsupervised or heuristically thresholded confidence score.

3. We design a lightweight cross-attention fusion module, invoked only when GateKeeper activates the second stream, that allows the primary stream’s representation to selectively attend to anatomical context.
4. We validate SecondOpinion on two datasets of increasing diagnostic difficulty: multi-class chest disease detection (COVID-19, viral pneumonia, lung opacity, tuberculosis, normal) and pelvic fracture detection, generalizing from visible to invisible fractures. GateKeeper’s activation rate tracks this difficulty gradient (9.23% to 24.12% to 45.71%), while matching or exceeding prior state-of-the-art at a fraction of the always-on computational cost.

2 Methodology




Fig. 1. SecondOpinion Full Pipeline.

SecondOpinion processes every input through a fast primary stream, and invokes a second, anatomy-guided stream only when a learned gating mechanism, GateKeeper, flags the primary stream’s prediction as uncertain. Figure 1 illustrates the full pipeline: a primary stream (Stream A) on the raw image, an anatomy-guided stream (Stream B) on a segmented representation, GateKeeper, and a cross-attention fusion module invoked only when Stream B is activated.

2.1 Stream A: Primary Pathway

Stream A processes the raw input $x_a \in \mathbb{R}^{3 \times 224 \times 224}$ through an ImageNet-pretrained EfficientNet-B0 backbone [15], producing a pooled feature vector

$h_a \in \mathbb{R}^{1280}$ and logits $z_a \in \mathbb{R}^C$ via a linear head, where C is the number of classes. Stream A is always evaluated and constitutes the primary diagnostic pathway.

2.2 Stream B: Anatomy-Guided Pathway

Stream B processes an anatomy-guided representation x_b , obtained from a U-Net (MiT-B0 backbone) segmentation model trained on ground-truth masks for all classes except tuberculosis, where masks were instead produced via inference; in all cases, Stream B’s input is the predicted segmentation, not ground truth. This representation is passed through a second, independently weighted EfficientNet-B0 backbone, producing a pooled feature vector $h_b \in \mathbb{R}^{1280}$ and spatial map $S_b \in \mathbb{R}^{1280 \times 7 \times 7}$. Stream B shares no intermediate features with Stream A, reflecting the clinical analogy that a second opinion should be formed independently of the first.

2.3 GateKeeper

GateKeeper estimates $P(\text{Stream A is correct})$ from Stream A’s pooled features h_a , softmax probabilities $p_a = \text{softmax}(z_a)$, and the Shannon entropy [14] of p_a ,

$$H(p_a) = - \sum_{i=1}^C p_{a,i} \log(p_{a,i} + \epsilon) \quad (1)$$

concatenated and passed through a small MLP $\phi(\cdot)$ to produce

$$\alpha = \sigma(\phi([h_a; p_a; H(p_a)])) \in (0, 1) \quad (2)$$

GateKeeper activates Stream B whenever $\alpha < 0.5$, the natural decision boundary of the binary cross-entropy objective used to train it, rather than a tuned threshold. This objective targets correctness prediction, not efficiency directly; efficiency follows as a consequence, since a well-calibrated gate activates Stream B only where Stream A is unreliable.

2.4 Cross-Attention Fusion

When activated, Stream B’s spatial features are fused with Stream A’s via single-head cross-attention [17], with Stream A as query and Stream B as key/value, reflecting Stream A’s role as the primary reasoner attending selectively to anatomical context. Given flattened, projected features $Q = S_a W_Q$, $K = S_b W_K$, $V = S_b W_V \in \mathbb{R}^{N \times d}$ ($N=49$ tokens), the fused output is

$$\tilde{F} = \text{softmax}\left(\frac{QK^\top}{\sqrt{d}}\right) V \quad (3)$$

$\tilde{F}$ is mean-pooled and layer-normalized to $\tilde{f} \in \mathbb{R}^d$, concatenated with h_a, h_b , and passed through a classification head to produce $z_{\text{fused}} \in \mathbb{R}^C$.

2.5 Training Procedure

Training proceeds in two phases. In **Phase 1**, both streams run on every sample with GateKeeper frozen (no reliable correctness signal exists yet), optimizing

$$\mathcal{L}_{\text{phase1}} = \mathcal{L}_{\text{CE}}(z_a, y) + \mathcal{L}_{\text{CE}}(z_{\text{fused}}, y) \quad (4)$$

In **Phase 2**, GateKeeper is unfrozen and trained jointly. A soft-gated prediction $z_{\text{final}} = \alpha z_a + (1 - \alpha) z_{\text{fused}}$ preserves gradient flow, and GateKeeper is supervised with binary cross-entropy against $\hat{g} = \mathbf{1}[\arg \max(z_a) = y]$, giving the unweighted four-term objective

$$\mathcal{L}_{\text{phase2}} = \mathcal{L}_{\text{CE}}(z_{\text{final}}, y) + \mathcal{L}_{\text{CE}}(z_a, y) + \mathcal{L}_{\text{CE}}(z_{\text{fused}}, y) + \text{BCE}(\alpha, \hat{g}) \quad (5)$$

The two auxiliary terms keep both streams learning even when α saturates toward 0 or 1; no efficiency-penalty term is needed, since low activation on correctly-handled cases follows directly from the gate’s supervision objective. All terms are unweighted to avoid additional tunable hyperparameters. At inference, hard gating replaces Equation 5’s soft combination: Stream B and fusion run only when $\alpha < 0.5$.

3 Experiments

3.1 Datasets

Chest X-ray. We combine two public sources, the COVID-19 Radiography Database [5, 11] and the Tuberculosis Chest X-ray Database [10], into a unified five-class dataset of 21,865 images (10,192 normal, 6,012 lung opacity, 1,345 viral pneumonia, 3,616 COVID-19, 700 tuberculosis), resized to 224×224 . Tuberculosis lacked segmentation ground truth; its anatomy-guided input was inferred from the model trained on the other four classes (Section 2.2). We use stratified 5-fold cross-validation with no augmentation.

Pelvic Fracture. We use the private AMERI PXR dataset (228 images: 168 fracture, 60 normal) from the Steel Memorial Hirohata Hospital in Japan, evaluated under stratified 5-fold cross-validation with expert-annotated bone segmentation. A held-out Invisible Fracture (INVIS) subset of 35 cases, confirmed via 3D-CT but not visible on X-ray, evaluates generalization only. Class imbalance was addressed via augmentation ($2 \times$ fracture, $6 \times$ normal) using random rotation, shear, flip, and translation.

3.2 Implementation Details

All models were trained on a single NVIDIA A100 GPU (Google Colab) with AdamW and CosineAnnealing scheduling, batch size 16, learning rate 1×10^{-4} for both phases. Chest X-ray training used 25/35 epochs (Phase 1/2); pelvic fracture used 15/25 epochs, reflecting dataset scale. All other settings were held constant across tasks.

Table 1. Chest X-ray results. Mean over 5-fold CV; 95% CI shown in the row below each of our models’ results.

Model	Acc.	Prec.	Rec.	Spec.	F1	AUROC
EfficientNetB3 [1]	93.35%	95.99%	93.81%	95.11%	0.949	0.987
CheXNet [1]	94.12%	94.76%	93.05%	95.81%	0.939	0.988
CheXNet-CBAM [3]	96.21%	96.14%	94.67%	97.55%	0.954	0.989
Stream A only	93.25%	95.02%	93.56%	95.01%	0.943	0.981
	$\pm 0.32\%$	$\pm 0.40\%$	$\pm 0.53\%$	$\pm 0.52\%$	± 0.0044	± 0.0061
Stream B only	94.06%	95.53%	94.28%	95.56%	0.949	0.983
	$\pm 0.19\%$	$\pm 0.22\%$	$\pm 0.61\%$	$\pm 0.58\%$	± 0.0040	± 0.0042
Always-On	99.05%	98.56%	98.32%	99.42%	0.984	0.993
	$\pm 0.06\%$	$\pm 0.16\%$	$\pm 0.24\%$	$\pm 0.02\%$	± 0.0018	± 0.0054
SecondOpinion	98.41%	97.41%	97.13%	98.91%	0.973	0.990
	$\pm 0.07\%$	$\pm 0.26\%$	$\pm 0.08\%$	$\pm 0.02\%$	± 0.0013	± 0.0046

3.3 Evaluation Metrics

We report Accuracy, Precision, Recall, Specificity, F1, and AUROC with 95% CIs (z -approximation) over 5-fold CV. For pelvic fracture, F1 is the harmonic mean of Recall and Specificity, matching the convention used by PelFANet [2] and the ImageNet/DRR20 baselines [9] for direct comparability; chest X-ray F1 uses the standard Precision-Recall form. We also report Params and FLOPs as static per-path values and activation-weighted averages.

4 Results & Discussion

4.1 Main Results

Stream A only and Stream B only share identical architectures, differing only in input (raw image versus segmentation-guided representation), and Stream B only consistently outperforms Stream A only across all tasks, confirming that anatomy-guided inputs provide independent diagnostic value.

Chest X-ray. SecondOpinion reaches 98.41% accuracy and 0.990 AUROC (Table 1), exceeding CheXNet-CBAM [3] while activating Stream B on only 9.23% of samples, trailing the always-on configuration by just 0.6 accuracy points. We view this gap as the explicit cost of efficiency: SecondOpinion does not aim to outperform an unconditional dual-stream model, but to approximate it closely while avoiding unnecessary computation on cases that do not require it.

Pelvic Fracture. SecondOpinion outperforms PelFANet [2] on both subsets in AUROC and F1 (Tables 2, 3), trailing only on recall. This is a trade-off: the recall gap is offset by a larger gain in specificity, yielding higher F1 overall. Stream B activates on 24.12% of visible and 45.71% of invisible cases, the latter reflecting generalization to unseen, harder cases, with the always-on gap widening to 0.8 accuracy points at this higher activation rate.

Table 2. Visible fracture (VIS) results. Mean over 5-fold CV; 95% CI shown in the row below each of our models’ results.

Method	Acc.	Prec.	Rec.	Spec.	F1	AUROC
ImageNet+DRR20 [9]	—	—	—	—	0.8390	0.9280
DRR20 [9]	—	—	—	—	0.8520	0.9290
PelFANet [2]	88.68%	92.49%	92.21%	78.33%	0.8471	0.9334
Stream A only	86.40%	85.90%	84.00%	73.30%	0.763	0.867
	$\pm 3.8\%$	$\pm 4.3\%$	$\pm 9.0\%$	$\pm 20.3\%$	± 0.121	± 0.020
Stream B only	86.60%	86.20%	85.80%	76.60%	0.806	0.871
	$\pm 3.6\%$	$\pm 4.2\%$	$\pm 7.8\%$	$\pm 9.6\%$	± 0.071	± 0.008
Always-On	92.60%	93.90%	89.80%	90.70%	0.901	0.951
	$\pm 2.7\%$	$\pm 0.9\%$	$\pm 1.9\%$	$\pm 6.4\%$	± 0.027	± 0.006
SecondOpinion	91.40%	93.00%	87.90%	88.30%	0.879	0.949
	$\pm 3.5\%$	$\pm 1.5\%$	$\pm 2.9\%$	$\pm 6.5\%$	± 0.022	± 0.011

Table 3. Invisible fracture (INVIS) results, evaluated as a held-out generalization set. Mean over 5-fold CV; 95% CI shown in the row below each of our models’ results.

Method	Acc.	Prec.	Rec.	Spec.	F1	AUROC
ImageNet+DRR20 [9]	—	—	—	—	0.7210	0.7140
DRR20 [9]	—	—	—	—	0.7860	0.8002
PelFANet [2]	82.29%	88.36%	84.35%	78.33%	0.8123	0.8688
Stream A only	76.60%	77.90%	82.10%	71.60%	0.756	0.744
	$\pm 4.1\%$	$\pm 6.8\%$	$\pm 3.8\%$	$\pm 13.2\%$	± 0.063	± 0.050
Stream B only	77.00%	80.10%	82.10%	72.60%	0.763	0.752
	$\pm 3.7\%$	$\pm 6.8\%$	$\pm 3.8\%$	$\pm 12.3\%$	± 0.057	± 0.036
Always-On	85.40%	89.90%	83.50%	85.30%	0.842	0.901
	$\pm 0.8\%$	$\pm 1.1\%$	$\pm 5.1\%$	$\pm 3.0\%$	± 0.016	± 0.020
SecondOpinion	84.60%	89.20%	83.20%	83.30%	0.830	0.898
	$\pm 0.9\%$	$\pm 1.7\%$	$\pm 5.1\%$	$\pm 5.1\%$	± 0.026	± 0.020

4.2 Efficiency and Activation Analysis

Table 4 reports per-path Params/FLOPs and activation-weighted averages. SecondOpinion’s cost scales with task difficulty, from 4.78M params/0.442 GFLOPs on chest X-ray (9.23% activation) to 7.15M params/0.608 GFLOPs on invisible fracture (45.71% activation), 41–71% of the always-on cost while retaining performance within 0.6–2.2 points of the ceiling. The corresponding AUROC gap between Stream A and the always-on model increased from 1.18 to 8.40 and 15.78 points across the three tasks, indicating that GateKeeper activates Stream B primarily when the additional anatomical information yields the greatest performance gain, particularly under the distribution shift to unseen invisible fractures. Note that reported costs cover the classification pipeline only and exclude the upstream segmentation model, which runs offline prior to inference.

4.3 Qualitative Analysis

Figure 2 presents GradCAM visualizations [13] for two correctly classified cases. For a COVID-19 chest X-ray case (Stream B inactive), Stream A attends cor-

Table 4. Parameters, FLOPs, and activation rate per configuration.

Configuration	Params (M)	FLOPs (G)	Activation
Single Stream Only	4.18	0.400	–
Always-On	10.68	0.855	100%
SecondOpinion – CXR	4.78	0.442	9.23%
SecondOpinion – VIS	5.75	0.510	24.12%
SecondOpinion – INVIS	7.15	0.608	45.71%




Fig. 2. GradCAM examples. (a) Chest X-ray (Stream B inactive): Stream A attends correctly to both lungs. (b) Invisible fracture (Stream B active): Stream A attends to irrelevant background corners. (c) Same case post-fusion: attention refocuses onto the bone region.

rectly to both lung fields. For an invisible fracture case (Stream B active), Stream A instead attends to irrelevant background corners; after fusion, attention shifts to the central bone region, and the case is correctly classified. This suggests the anatomy-guided second opinion corrects an otherwise poorly localized prediction.

Limitations. The invisible fracture subset contains only 35 cases, so despite reporting 95% confidence intervals, its point estimates should be interpreted with caution. The pelvic fracture dataset is private, our evaluation uses a single backbone and gating architecture, and our qualitative analysis is limited to two representative cases intended as a sanity check rather than a systematic interpretability evaluation.

5 Conclusion

We presented SecondOpinion, a framework that processes every case through a fast primary stream and consults a second, anatomy-guided stream only when GateKeeper judges the primary prediction unreliable. Across two datasets spanning a gradient of diagnostic difficulty, SecondOpinion matches or exceeds prior state-of-the-art performance while activating its anatomy-guided stream on only a fraction of cases. Future work includes extending GateKeeper to additional

backbones and modalities, and exploring learned, sample-specific anatomy representations beyond fixed segmentation pipelines.

Data Use Declaration

The public datasets were obtained from their official repositories and used in accordance with their respective licenses; all required dataset citations are included in the References. The private AMERI PXR dataset was used under IRB approval (No. 2019-1-52) from the Steel Memorial Hirohata Hospital, Japan, and is available from the corresponding author upon reasonable request and with permission from the Steel Memorial Hirohata Hospital, Japan.

Disclosure of Interests

The authors have no competing interests to declare.

References

1. Bhuiyan, S.T., Rahman, R., Islam, R., Haque, M., Islam, M., Kobashi, S., Alam, S.B.: Lung disease diagnosis from cxr using convolutional neural network (cnn): A comparative study. *Soft Computing* **29**(23), 6333–6345 (2025)
2. Bhuiyan, S.T., Rahman, R., Wasi, S., Yagi, N., Kobashi, S., Islam, A., Alam, S.B.: Invisible yet detected: Pelfanet with attention-guided anatomical fusion for pelvic fracture diagnosis. In: *Workshop on Empowering Medical Image Computing and Research through Early-Career Expertise*. pp. 103–113. Springer (2025)
3. Bhuiyan, S.T., Zaman, Z., Rahman, R., Choudhury, M.S., Wasi, S., Islam, A., Alam, S.B., Kobashi, S.: Chexnet-cbam: Enhancing large-scale medical pre-training with attention for lung disease diagnosis. In: *2026 IEEE 56th International Symposium on Multiple-Valued Logic (ISMVL)*. pp. 105–110. IEEE (2026)
4. Bruno, P., Macrì, M., Dodaro, C.: A dual-stage deep learning framework for breast ultrasound image segmentation and classification. *Journal of Medical Systems* **49**(1), 162 (2025)
5. Chowdhury, M.E., Rahman, T., Khandakar, A., Mazhar, R., Kadir, M.A., Mahbub, Z.B., Islam, K.R., Khan, M.S., Iqbal, A., Al Emadi, N., et al.: Can ai help in screening viral and covid-19 pneumonia? *Ieee Access* **8**, 132665–132676 (2020)
6. Jin, C., Udupa, J., Zhao, L., Tong, Y., Odhner, D., Pednekar, G., Nag, S., Lewis, S., Poole, N., Mannikeri, S., et al.: Anatomy-guided deep learning for object localization in medical images. In: *Medical Imaging 2022: Image Processing*. vol. 12032, pp. 591–597. SPIE (2022)
7. Mokssit, S., Karrakchou, O., Mousist, A., Ghogho, M.: Confidence-gated training for efficient early-exit neural networks. *arXiv preprint arXiv:2509.17885* (2025)
8. Passalis, N., Raitoharju, J., Tefas, A., Gabbouj, M.: Efficient adaptive inference for deep convolutional neural networks using hierarchical early exits. *Pattern Recognition* **105**, 107346 (2020)
9. Rahman, R., Yagi, N., Hayashi, K., Maruo, A., Muratsu, H., Kobashi, S.: Enhancing fracture diagnosis in pelvic x-rays by deep convolutional neural network with synthesized images from 3d-ct. *Scientific Reports* **14**(1), 8004 (2024)

10. Rahman, T., Khandakar, A., Kadir, M.A., Islam, K.R., Islam, K.F., Mazhar, R., Hamid, T., Islam, M.T., Kashem, S., Mahbub, Z.B., et al.: Reliable tuberculosis detection using chest x-ray with deep learning, segmentation and visualization. *Ieee Access* **8**, 191586–191601 (2020)
11. Rahman, T., Khandakar, A., Qiblawey, Y., Tahir, A., Kiranyaz, S., Kashem, S.B.A., Islam, M.T., Al Maadeed, S., Zughaier, S.M., Khan, M.S., et al.: Exploring the effect of image enhancement techniques on covid-19 detection using chest x-ray images. *Computers in biology and medicine* **132**, 104319 (2021)
12. Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D., Bagul, A., Langlotz, C., Shpanskaya, K., et al.: Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225* (2017)
13. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE international conference on computer vision*. pp. 618–626 (2017)
14. Shannon, C.E.: A mathematical theory of communication. *The Bell system technical journal* **27**(3), 379–423 (1948)
15. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: *International conference on machine learning*. pp. 6105–6114. PMLR (2019)
16. Tang, Z., Zhao, S.: T-dumpra: Teacher-guided dual-path multi-prototype retrieval augmented framework for fine-grained medical image classification. *arXiv preprint arXiv:2604.17360* (2026)
17. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
18. Wołczyk, M., Wójcik, B., Bałazy, K., Podolak, I.T., Tabor, J., Śmieja, M., Trzcinski, T.: Zero time waste: Recycling predictions in early exit neural networks. *Advances in Neural Information Processing Systems* **34**, 2516–2528 (2021)
19. Zafar, J., Koc, V., Zafar, H.: Dual-stream contrastive latent learning generative adversarial network for brain image synthesis and tumor classification. *Journal of Imaging* **11**(4), 101 (2025)