

Longitudinal Multi-View Modeling for Breast Cancer Risk Prediction

Solveig Thrun¹[0009-0006-7349-9449], Zijun Sun¹[0009-0004-1848-6966], Suaiba A. Salahuddin¹[0000-0003-4407-0222], Kristoffer Wickstrøm¹[0000-0003-1395-7154], Elisabeth Wetzter¹[0000-0002-0544-8272], Stine Hansen²[0000-0002-0962-6160], Robert Jenssen^{1,3,4}[0000-0002-7496-8474], and Michael Kampffmeyer^{1,3}[0000-0002-7699-0405]

¹ Department of Physics and Technology, UiT The Arctic University of Norway, Tromsø, Norway solveig.thrun@uit.no

² SPKI The Norwegian Centre for Clinical Artificial Intelligence, University Hospital of North Norway, Tromsø, Norway

³ Norwegian Computing Center, Oslo, Norway

⁴ Pioneer Centre for AI, University of Copenhagen, Copenhagen, Denmark

Abstract. Accurate breast cancer risk prediction from screening mammography is critical for enabling personalized screening intervals and early detection. Recent deep learning methods have shown the value of longitudinal data and explicit temporal alignment. However, existing approaches either perform explicit alignment using a single mammographic view or model multiple views without explicit longitudinal alignment, limiting their ability to exploit the complementary spatial-temporal information used in clinical practice. To address this gap, we propose LMV-Net, a longitudinal multi-view breast cancer risk prediction model that jointly analyzes anatomically complementary CC and MLO views within an explicitly aligned longitudinal framework. We evaluate our approach on the public EMBED and CSAW-CC datasets, comparing it to state-of-the-art breast cancer risk prediction methods. Our model consistently outperforms existing approaches in overall risk prediction performance and across different breast density and cancer subgroups. Importantly, these improvements highlight the potential of longitudinal multi-view modeling to enhance risk stratification, paving the way for future work on personalized screening, earlier identification of high-risk patients, and more efficient screening resource allocation. The code is publicly available [here](#).

Keywords: Breast cancer risk prediction · Longitudinal mammogram · Multi-view analysis

1 Introduction

Breast cancer is the most common cancer among women worldwide, representing a major public health challenge [5]. Early detection through population-based screening has been critical in reducing mortality, with mammography, typically

acquired in craniocaudal (CC) and mediolateral oblique (MLO) views, remaining the gold standard [3,30]. Large-scale screening programs have therefore played a key role in identifying cancers at an earlier, better treatable stage [30]. However, effective risk stratification for personalized screening and early detection in high-risk populations remains limited [21], especially in women with dense breast tissue, where density both increases cancer risk and reduces mammographic sensitivity [4,27].

Recent advances in deep learning have improved breast cancer risk prediction from mammography, offering a data-driven path toward more individualized screening [13]. Early deep learning models analyzed current images in isolation (e.g., Mirai [29]), overlooking the valuable longitudinal information that radiologists exploit when comparing current and prior examinations [1,26]. To leverage this longitudinal information, recent models incorporate prior screenings, enabling them to capture gradual tissue changes over time and improve risk assessment [7,15,16,18,22,28]. These models consistently outperform single-time-point approaches, demonstrating the value of modeling temporal information.

Building on advances in mammography research, most longitudinal mammography models rely on implicit alignment (e.g., LoMaR [16], VMRA-MaR [22]), comparing prior and current images without enforcing spatial correspondence. Although non-rigid breast deformation makes explicit alignment challenging, recent evidence indicates that such strategies are better suited to capture subtle temporal variations, outperforming implicit approaches for modeling longitudinal change [23]. Nevertheless, existing explicit alignment methods (ImgFeatAlign [23,24], OA-BreaCR [28]) have been restricted to single-view formulations, without jointly modeling the complementary information available across CC and MLO views. As a result, current approaches do not fully exploit the combined temporal and cross-view cues that are critical for accurate breast cancer risk prediction.

To address this shortcoming, we introduce a breast-wise longitudinal multi-view risk prediction framework that explicitly aligns the feature maps of each CC and MLO view across multiple time points, capturing subtle longitudinal changes and more closely mirroring radiologists’ workflow. Our approach uses image-based deformation fields to align current and prior feature maps for each view, mitigating spatial misalignment while preserving discriminative information. By introducing a dual-stream attention mechanism, the model enables multi-view, breast-wise comparison over time. This design effectively captures longitudinal tissue changes and provides a foundation for more reliable risk prediction. Ultimately, it supports earlier identification of high-risk patients and facilitates personalized screening strategies.

Our **contributions** are:

1. LMV-Net, a breast-wise longitudinal multi-view model for breast cancer risk prediction, that facilitates multi-view, multi-time-point integration via explicit feature-space alignment and cross-view temporal fusion.
2. A Dual Stream Attention block that refines longitudinal features of each view (CC and MLO) using parallel self- and cross-attention.




Fig. 1: Schematic overview of the proposed LMV-Net.

- Extensive validation on two large public datasets, EMBED and CSAW-CC, demonstrating robust performance.

2 Method

We aim to predict an individual’s future breast cancer risk over a five-year horizon using mammograms from current and prior screening visits for a single breast, to support personalized, risk-based screening programs. Unlike prior approaches that rely on explicit alignment and use only a single view [23,28], our proposed LMV-Net (Fig. 1) incorporates both CC and MLO views, which provide complementary information. To achieve this, we introduce a dual-stream attention block for multi-view fusion that refines explicitly aligned longitudinal features through parallel self- and cross-attention. LMV-Net comprises six sequential stages.

Step 1: Feature extraction. A shared image encoder extracts feature maps from current and prior mammograms in the MLO and CC views, yielding four feature maps: current feature maps $\mathbf{f}_{\text{cur}}^{\text{MLO}}, \mathbf{f}_{\text{cur}}^{\text{CC}} \in \mathbb{R}^{C \times H \times W}$ and prior feature maps $\mathbf{f}_{\text{pri}}^{\text{MLO}}, \mathbf{f}_{\text{pri}}^{\text{CC}} \in \mathbb{R}^{C \times H \times W}$, where C is the number of channels, and H and W denote the height and width of the feature maps, respectively.

Step 2: Temporal feature alignment. Inspired by [23], we explicitly align longitudinal feature maps in feature space using image-based deformation fields. The deformation field is obtained using a pretrained and frozen MammoRegNet [24] by registering the prior mammogram to the current mammogram. A separate deformation field is estimated for each view (CC and MLO). Using a spatial transformer function [24,28], the prior feature maps ($\mathbf{f}_{\text{pri}}^{\text{MLO}}$ and $\mathbf{f}_{\text{pri}}^{\text{CC}}$) are aligned to the corresponding current feature maps ($\mathbf{f}_{\text{cur}}^{\text{MLO}}$ and $\mathbf{f}_{\text{cur}}^{\text{CC}}$), resulting in

the aligned prior feature maps $\mathbf{f}_{\text{pri-align}}^{\text{MLO}}$ and $\mathbf{f}_{\text{pri-align}}^{\text{CC}}$.

Step 3: Temporal subtraction. To capture temporal changes, we follow [23,28] and compute view-specific difference features via temporal subtraction: $\mathbf{f}_{\text{diff}}^{\text{CC}} = \mathbf{f}_{\text{cur}}^{\text{CC}} - \mathbf{f}_{\text{pri-align}}^{\text{CC}}$, $\mathbf{f}_{\text{diff}}^{\text{MLO}} = \mathbf{f}_{\text{cur}}^{\text{MLO}} - \mathbf{f}_{\text{pri-align}}^{\text{MLO}}$. These difference features encode longitudinal changes between aligned current and prior mammograms, and positional encoding [28] is applied to $\mathbf{f}_{\text{diff}}^{\text{CC}}$ and $\mathbf{f}_{\text{diff}}^{\text{MLO}}$ to account for variable time intervals between screenings.

Step 4: Temporal feature concatenation. Current, aligned prior, and difference features are concatenated along the channel dimension to form $\mathbf{f}_{\text{long}}^{\text{MLO}} \in \mathbb{R}^{3C \times H \times W}$ and $\mathbf{f}_{\text{long}}^{\text{CC}} \in \mathbb{R}^{3C \times H \times W}$, integrating current appearance, longitudinal context, and temporal changes.

Step 5: Dual Stream Attention. We propose a Dual Stream Attention block designed to enhance the modeling of longitudinal multi-view features. The module jointly refines the longitudinal representations $\mathbf{f}_{\text{long}}^{\text{MLO}}$ and $\mathbf{f}_{\text{long}}^{\text{CC}}$ by explicitly capturing both intra-view and inter-view dependencies. For each view, the block applies multi-head self-attention to model long-range intra-view interactions and improve spatial coherence within the same mammographic view. In parallel, multi-head cross-attention leverages features from the complementary view as keys and values, enabling each view to selectively attend to anatomically and semantically relevant regions in the other view. This mechanism facilitates effective information exchange between CC and MLO representations, enabling the model to leverage complementary perspectives of the same breast tissue. The outputs of self- and cross-attention are fused via element-wise summation, after which stochastic depth (DropPath [12]) is applied to improve generalization, and reduce overfitting. This attention output is added to the original residual features, followed by layer normalization [2]. Each view is subsequently processed through a feed-forward network (FFN) with a fourfold expansion of the channel dimension, increasing model expressiveness. A second residual connection and normalization are applied after the FFN. Through this dual-stream attention design, each view integrates both its own global spatial context and complementary cross-view information, resulting in enriched and more discriminative multi-view feature representations.

Step 6: Final fusion and risk prediction. The refined $\mathbf{f}_{\text{long}}^{\text{MLO}}$ and $\mathbf{f}_{\text{long}}^{\text{CC}}$ features are pooled and concatenated into a fused multi-view representation $\mathbf{f}$, which passes through a cumulative probability layer [18,29] for time-dependent risk estimation. This layer comprises two fully connected layers: one predicting hazards at each future time step (i.e. 1 year, 2 years, ...) and another computing a baseline hazard, with ReLU enforcing non-negativity. The cumulative probability at time t is then computed as

$$P(t | \mathbf{f}) = h_{\text{base}}(\mathbf{f}) + \sum_{j=1}^t h_j(\mathbf{f}), \quad (1)$$

where $h_{\text{base}}(\mathbf{f})$ denotes the baseline hazard and $h_j(\mathbf{f})$ the hazard at time step j , capturing the sequential accumulation of risk. Training uses a time-dependent binary cross-entropy (BCE) loss, with labels set to 1 if cancer occurs within t

Table 1: Breast-level CC+MLO exam splits and time-to-cancer labels.

EMBED	Split	Positive						Negative
		1Y	2Y	3Y	4Y	5Y	All	>5Y
	Train	370	96	83	46	34	629	12262
	Val	146	32	29	20	14	241	4907
	Test	220	76	54	37	29	416	7341

CSAW-CC	Split	Positive						Negative
		1Y	2Y	3Y	4Y	5Y	All	>5Y
	Train	352	350	94	224	88	1108	15686
	Val	140	142	48	92	28	450	6188
	Test	208	214	76	124	54	676	9402

years and 0 otherwise. For censored patients, a masking function $\delta(t)$ [28,29] is applied so that only valid time points contribute to the summed BCE loss. Following [6,28], we adopt multilevel joint learning to capture complementary dependencies between CC and MLO views. The network jointly learns a fused multi-view representation for risk estimation and view-specific features that serve as auxiliary supervision during training, encouraging discriminative per-view learning and enhancing the quality of the fused representation used at inference.

3 Experimental Setup and Results

Datasets and Pre-processing We evaluate LMV-Net on two large, publicly available mammography datasets: the Emory Breast Imaging Dataset (EMBED) [14] and the Cohort of Screen-Aged Women Case Control dataset (CSAW-CC) [8], ensuring reproducibility and an experimental setup mirroring prior work [16,22,23,24,28,29]. EMBED comprises mammograms from a diverse population acquired on Hologic, GE, and Fujifilm systems, while CSAW-CC contains cases collected at Karolinska University Hospital (Sweden) using Hologic systems. Following [28], we include only patients with at least five years of follow-up. All standard 2D mammograms are resized to 1664×2048 pixels while preserving aspect ratio. Data are split at the patient level into training, validation, and test sets (5:2:3). CSAW-CC provides LIBRA-based breast density estimates [17], grouped into three equally sized categories (low, medium, high), whereas EMBED uses BI-RADS density categories [9]. Time-to-cancer is defined as the interval between image acquisition and diagnosis. Table 1 summarizes dataset splits and time-to-cancer distributions at the breast-wise CC-MLO exam level. For patients with more than two time points, we construct all consecutive time-point pairs of images for training and evaluation. For training MammoRegNet, we exclusively used patients who were excluded from the risk prediction cohort, preventing any potential data leakage.

Evaluation Metrics Following [18,28,29], we evaluate performance using the C-index [25] and AUC [10] over 1–5 year risk horizons. 95% confidence intervals (CIs) are estimated via exam-level bootstrapping with 1,000 resamples.⁵ Statistical significant improvements over the second-best model is assessed using a paired non-parametric bootstrap test, with $p < 0.05$ considered significant.

Implementation Details All models were implemented in PyTorch 2.0.1 [20]

⁵ Alternatively, patient-level clustered bootstrapping yielded negligible differences to exam-level bootstrapping, and all significant results remained significant.

Table 2: 1–5 Year breast cancer risk prediction: C-index and AUC values with $\pm 95\%$ CI. ❄: frozen encoder, 🔥: fine-tuned encoder. Δ denotes % improvement of best LMV-Net over best baseline. * indicates statistically significant improvement of LMV-Net over the best baseline ($p < 0.05$). Training time in hours (h) and total number of model parameters (#, in millions).

		C-index (%) $\uparrow$	Follow-up year AUC (%) $\uparrow$					h	#
			1-Y	2-Y	3-Y	4-Y	5-Y		
EMBED	Mirai [29] (❄)	59.8 $\pm$ 8.7	56.7 $\pm$ 20.7	52.6 $\pm$ 15.3	54.0 $\pm$ 12.6	49.3 $\pm$ 12.5	49.2 $\pm$ 11.6	6	28
	Mirai [29] (🔥)	69.1 $\pm$ 10.6	71.0 $\pm$ 15.6	70.9 $\pm$ 15.5	70.6 $\pm$ 11.7	65.7 $\pm$ 11.3	65.9 $\pm$ 11.6	7	28
	VMRA-MaR [22] (🔥)	72.4 $\pm$ 12.5	77.9 $\pm$ 4.5	75.3 $\pm$ 9.7	73.1 $\pm$ 12.7	68.9 $\pm$ 14.3	67.1 $\pm$ 11.8	45	2802
	OA-BreaCR [28] (🔥)	72.3 $\pm$ 2.0	72.9 $\pm$ 2.8	72.3 $\pm$ 4.9	72.4 $\pm$ 2.5	72.3 $\pm$ 2.4	73.2 $\pm$ 2.4	9	15
	ImgFeatAlign [23] (❄)	74.7 $\pm$ 2.4	75.0 $\pm$ 2.9	75.5 $\pm$ 2.4	75.3 $\pm$ 2.1	75.9 $\pm$ 2.3	72.5 $\pm$ 3.2	12	52
	LMV-Net (❄)	78.3* $\pm$ 3.0	80.8* $\pm$ 3.7	78.3* $\pm$ 3.5	77.5* $\pm$ 3.3	76.6 $\pm$ 3.2	75.1* $\pm$ 4.9	15	65
	LMV-Net (🔥)	81.4*$\pm$2.8	83.9*$\pm$3.4	82.1*$\pm$3.0	80.8*$\pm$2.9	80.0*$\pm$2.8	77.2*$\pm$4.7	16	65
	Δ (%) $\uparrow$	+ 6.7	+ 6.0	+ 6.6	+ 5.5	+ 4.1	+ 4.0		
CSAW-CC	Mirai [29] (❄)	71.4 $\pm$ 2.7	56.8 $\pm$ 10.6	54.7 $\pm$ 5.8	56.4 $\pm$ 5.1	56.1 $\pm$ 4.2	55.4 $\pm$ 4.0	7	28
	Mirai [29] (🔥)	72.3 $\pm$ 2.9	69.9 $\pm$ 3.3	69.8 $\pm$ 3.9	69.2 $\pm$ 3.8	67.0 $\pm$ 3.4	66.9 $\pm$ 3.4	8	28
	VMRA-MaR [22] (🔥)	72.7 $\pm$ 2.6	70.9 $\pm$ 5.5	70.3 $\pm$ 5.6	69.5 $\pm$ 4.4	70.9 $\pm$ 4.3	64.5 $\pm$ 6.0	46	2802
	OA-BreaCR [28] (🔥)	61.6 $\pm$ 4.2	63.5 $\pm$ 4.4	61.7 $\pm$ 2.1	65.0 $\pm$ 2.0	67.0 $\pm$ 1.9	67.7 $\pm$ 2.2	10	15
	ImgFeatAlign [23] (❄)	70.4 $\pm$ 1.9	72.0 $\pm$ 2.4	71.5 $\pm$ 1.9	72.6 $\pm$ 1.9	72.0 $\pm$ 2.1	75.2 $\pm$ 2.3	13	52
	LMV-Net (❄)	73.4* $\pm$ 2.5	77.5* $\pm$ 3.3	73.5* $\pm$ 2.6	74.0 $\pm$ 2.5	73.1 $\pm$ 2.6	75.9$\pm$2.8	16	65
	LMV-Net (🔥)	74.4*$\pm$2.7	78.7*$\pm$3.6	74.1*$\pm$2.9	74.3*$\pm$2.7	73.5$\pm$2.8	75.0 $\pm$ 3.1	17	65
	Δ (%) $\uparrow$	+ 1.7	+ 6.7	+ 2.6	+ 1.7	+ 1.5	+ 0.7		

and on eight AMD Instinct MI210 GPUs (64 GB each) with a batch size of 4 per GPU (32 total). For fair comparison, we use the Mirai encoder [29], a ResNet-18 [11] with the final fully connected layer removed, following prior work [16,22,23,28]. Self- and cross-attention use four heads, and training uses AdamW [19] with a learning rate of 5×10^{-5} and weight decay of 1×10^{-4} . LMV-Net was trained for up to 30 epochs, with early stopping after 15 stagnant epochs based on the validation C-index and learning rate scheduling that halved the learning rate after 5 stagnant epochs. Data augmentation included random cropping, affine transformations, color jitter, and gamma correction. Baselines [22,24,28,29] and MammoRegNet [23,24] followed official implementations and hyperparameters, and LMV-Net was evaluated with both frozen and fine-tuned pretrained Mirai encoders.

3.1 Results

Comparison with State-of-the-Arts (SOTA) Table 2 compares the proposed LMV-Net with SOTA methods, including Mirai [29] (Sci. Transl. Med.’2021), OA-BreaCR [28] (MICCAI’24), VMRA-MaR [22] (MICCAI’25), and ImgFeatAlign [23] (MICCAI’25), on both datasets. LMV-Net consistently outperforms all baselines in C-index and AUC across all follow-up years on both datasets. Fine-tuning the encoder yields further improvements, achieving statistically significant gains ($p < 0.05$) over the second-best model. These results highlight the benefit of jointly modeling multi-view and longitudinal information for breast cancer risk prediction. Notably, the larger gains observed on EMBED are due to its greater heterogeneity across populations and imaging systems, whereas the single-scanner CSAW-CC cohort is more homogeneous, limiting the potential benefits of improved model generalization.




Fig. 2: C-index by density category (left; EMBED and CSAW-CC) and by cancer type (middle; EMBED) with 95% CI, and AUC over follow-up years for invasive cancer (right; EMBED). *: frozen encoder, 🔥: fine-tuned encoder.

C-index by density category Higher breast density is associated with increased cancer risk [4] and reduced mammographic sensitivity [27]. Fig 2 (left) shows C-index values with 95% CI across density categories for EMBED and CSAW-CC. Across both datasets, LMV-Net (🔥) achieves the highest C-index across nearly all density groups. Note, some model estimates are not shown for high-density EMBED cases (category D) due to insufficient comparable pairs for reliable estimation. Gains are most pronounced in medium- and high-density groups, highlighting LMV-Net’s robustness in challenging screening scenarios.

Performance across cancer types We also perform the first systematic comparison of recent state-of-the-art breast cancer risk prediction methods across different cancer types on the more diverse EMBED dataset. As shown in Fig. 2 (middle), LMV-Net outperforms SOTA methods in C-index for both types, with the fine-tuned encoder achieving the best results and competing methods exhibiting greater uncertainty. Fig. 2 (right) shows AUC over follow-up years for invasive cancer, the more aggressive and clinically relevant type. LMV-Net achieves higher and more stable AUCs, whereas other methods show lower performance and greater uncertainty. These findings demonstrate that integrating multiview and longitudinal information in alignment with clinical workflow yields more robust and reliable breast cancer risk prediction, particularly for invasive disease.

Attention map analysis Fig. 3 (left) visualizes the self- and cross-attention maps for a representative patient. LMV-Net assigns elevated attention to the tumor region (red bounding box) in both CC and MLO views. Self-attention primarily focuses on the lesion and its surrounding area within each view, capturing intra-view dependencies. Cross-attention, integrates information from the complementary view, where we see the model attending primarily to the tumor tissue and/or suspicious regions in the other view. We further compute the ratio of mean attention density within the tumor bounding box to that within the entire breast region, aggregated across the 12 test-set patients with tumor annotations. (Fig. 3, right). On average, attention density within the tumor region is 8–10 $\times$ higher than in the surrounding breast tissue. This supports that the




Fig. 3: Left: Attention maps showing LMV-Net focus on the tumor (red box). Right: Attention density in tumor bounding box relative to the whole breast.

proposed Dual Stream Attention block concentrates attention on anatomically corresponding tumor regions.

Ablation study We assess key architectural components through ablation studies: (i) *Repeating current features*: both temporal alignment (Step 2) and subtraction (Step 3) are removed, and Step 4 concatenates the current features three times; (ii) *Concatenating with zeros*: temporal alignment and subtraction are removed as above, but Step 4 concatenates the current features with two zero tensors; (iii) *Single-view modeling*: one branch of the model (CC or MLO) is removed, and the cross-attention block in the remaining branch is replaced with an additional self-attention block. Only a single view is processed, effectively disabling multi-view input while preserving the same number of attention layers. (iv) *Duplicated-view modeling*: the same view is fed into both branches; (v) *Dual-stream attention removed*: the dual-stream attention block is removed; (vi) *Multi-view learning*: using only $Y_{\text{multi-view}}$; and (vii) *Encoder fine-tuning*: freezing the pretrained Mirai backbone. Table 3 shows that removing longitudinal information, multi-view modeling, attention, or encoder fine-tuning consistently reduces C-index and AUC. The full model achieves the best performance on both EMBED and CSAW-CC, confirming the value of each component.

4 Conclusion and Outlook

In this study, we introduced LMV-Net, which mimics the radiologist’s workflow by jointly modeling longitudinal and multi-view information through explicit longitudinal alignment. Experiments on two large publicly available datasets demonstrate statistically significant improvements over SOTA methods, particularly in challenging scenarios such as patients with higher breast density. Clinically, this approach could improve early detection and personalized screening through more accurate risk assessment. While our results demonstrate the value of joint longitudinal multi-view modeling, further investigation is warranted. Incorporating multiple prior examinations may improve performance by providing richer temporal context. In addition, our framework relies on explicit feature-space alignment to enable accurate longitudinal modeling such that advances in robust explicit alignment methods could further enhance its performance. Future

Table 3: Ablation study for 1–5 year risk: mean C-index and AUC on EMBED and CSAW-CC. MV: Multi-view, DV: Different views in both branches, FT: Encoder fine-tuning, DSA: Dual Stream Attention.

	Variant	Prior	MV	DV	FT	DSA	$Y_{MLO}^{Y_{CC}}$	C-index (%) $\uparrow$	Follow-up year AUC (%) $\uparrow$				
									1-Y	2-Y	3-Y	4-Y	5-Y
EMBED	(i) No prior (current repeated)	✗	✓	✓	✓	✓	✓	77.8	80.3	78.1	76.9	76.3	73.6
	(ii) No prior (zero-padded)	✗	✓	✓	✓	✓	✓	77.5	80.2	78.3	76.4	76.2	73.8
	(iii) Only one view	✓	✗	✓	✓	✓	✓	70.8	72.2	70.8	68.9	65.9	61.9
	(iv) Same view both branches	✓	✗	✗	✓	✓	✓	77.4	78.8	78.1	77.7	76.2	74.7
	(v) Without DSA	✓	✓	✓	✓	✗	✓	72.8	73.9	73.7	72.8	70.3	65.8
	(vi) Use only $Y_{multi-view}$	✓	✓	✓	✓	✓	✗	80.1	83.0	80.4	79.4	78.8	76.5
	(vii) No finetuning	✓	✓	✓	✗	✓	✓	78.3	80.8	78.3	77.5	76.6	75.1
	Ours	✓	✓	✓	✓	✓	✓	81.4	83.9	82.1	80.8	80.0	77.2
CSAW-CC	(i) No prior (current repeated)	✗	✓	✓	✓	✓	✓	73.0	75.3	73.8	73.2	72.3	73.2
	(ii) No prior (zero-padded)	✗	✓	✓	✓	✓	✓	72.2	75.2	73.0	72.6	71.9	72.4
	(iii) Only one view	✓	✗	✓	✓	✓	✓	73.3	73.8	73.1	73.5	73.1	73.9
	(iv) Same view both branches	✓	✗	✗	✓	✓	✓	72.2	75.2	72.6	72.5	71.1	73.0
	(v) Without DSA	✓	✓	✓	✓	✗	✓	69.5	70.5	70.0	70.6	69.3	71.4
	(vi) Use only $Y_{multi-view}$	✓	✓	✓	✓	✓	✗	73.3	77.6	73.3	73.4	72.6	74.4
	(vii) No finetuning	✓	✓	✓	✗	✓	✓	73.4	77.5	73.5	74.0	73.1	75.9
	Ours	✓	✓	✓	✓	✓	✓	74.4	78.7	74.1	74.3	73.5	75.0

work will also focus on more comprehensive evaluation across additional clinically relevant metrics, including precision–recall AUC (PR-AUC), particularly given the class imbalance characteristic of breast cancer risk prediction.

Acknowledgments. This work was supported by The Research Council of Norway (Visual Intelligence, grant no. 309439, FRIPRO grants no. 315029 and no. 360068, and IKTPLUSS grant no. 303514).

Disclosure of Interests. The authors declare no competing interests.

References

1. Akwo, J., Trieu, P.Y., Barron, M., Reynolds, T., Lewis, S.: Access to prior screening mammograms affects the specificity but not sensitivity of radiologists’ performance. *Clinical Radiology* **79**(12) (2024)
2. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
3. Bhushan, A., Gonsalves, A., Menon, J.U.: Current state of breast cancer diagnosis, treatment, and theranostics. *Pharmaceutics* **13**(5) (2021)
4. Bodewes, F., van Asselt, A., Dorrius, M., Greuter, M., de Bock, G.: Mammographic breast density and the risk of breast cancer: A systematic review and meta-analysis. *The Breast* **66** (2022)
5. Bray, F., Laversanne, M., Sung, H., Ferlay, J., Siegel, R.L., Soerjomataram, I., Jemal, A.: Global cancer statistics 2022: Globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA: A Cancer Journal for Clinicians* **74**(3) (2024)
6. Chen, H., Martel, A.L.: Breast cancer detection from multi-view screening mammograms with visual prompt tuning. In: *Artificial Intelligence and Imaging for Diagnostic and Treatment Challenges in Breast Care. Deep-Breath.* (2025)

7. Dadsetan, S., Arefan, D., Berg, W.A., Zuley, M.L., Sumkin, J.H., Wu, S.: Deep learning of longitudinal mammogram examinations for breast cancer risk prediction. *Pattern Recognition* **132** (2022)
8. Dembrower, K., Lindholm, P., Strand, F.: A multi-million mammography image dataset and population-based screening cohort for the training and evaluation of deep neural networks—the cohort of screen-aged women (CSAW). *Journal of Digital Imaging* **33**(2) (2020)
9. D’Orsi, C., Sickles, E., Mendelson, E., Morris, E.: ACR BI-RADS Atlas: Breast Imaging Reporting and Data System. American College of Radiology (2014)
10. Hanley, J., McNeil, B.: The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology* **143** (1982)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR* (2016)
12. Huang, G., Sun, Y., Liu, Z., Sedra, D., Weinberger, K.Q.: Deep networks with stochastic depth. In: *ECCV* (2016)
13. Hussain, S., Ali, M., Naseem, U., Nezhadmoghadam, F., Jatoi, M., Gulliver, T., Tamez-Peña, J.: Breast cancer risk prediction using machine learning: a systematic review. *Frontiers in Oncology* **14** (2024)
14. Jeong, J.J., Vey, B.L., Bhimireddy, A., Kim, T., Santos, T., Correa, R., Dutt, R., Mosunjac, M., Oprea-Ilie, G., Smith, G., Woo, M., McAdams, C.R., Newell, M.S., Banerjee, I., Gichoya, J., Trivedi, H.: The EMory BrEast imaging Dataset (EMBED): A racially diverse, granular dataset of 3.4 million screening and diagnostic mammographic images. *Radiol Artif Intell* **5**(1) (2023)
15. Jiang, S., Bennett, D.L., Rosner, B.A., Colditz, G.A.: Longitudinal analysis of change in mammographic density in each breast and its association with breast cancer risk. *JAMA Oncology* **9**(6) (2023)
16. Karaman, B.K., Dodelzon, K., Akar, G.B., Sabuncu, M.R.: Longitudinal mammogram risk prediction. In: *MICCAI* (2024)
17. Keller, B.M., Nathan, D.L., Wang, Y., Zheng, Y., Gee, J.C., Conant, E.F., Kontos, D.: Estimation of breast percent density in raw and processed full field digital mammography images via adaptive fuzzy c-means clustering and support vector machine segmentation. *Medical Physics* **39**(8) (2012)
18. Lee, H., Kim, J., Park, E., Kim, M., Kim, T., Kooi, T.: Enhancing breast cancer risk prediction by incorporating prior images. In: *MICCAI* (2023)
19. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization (2019)
20. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library (2019)
21. Rubio, I., Drukker, C., Esgueva, A.: Risk-based breast cancer screening: what are the challenges? *Tumori J.* (2024)
22. Sun, Z., Thrun, S., Kampffmeyer, M.: VMRA-MaR: An Asymmetry-Aware Temporal Framework for Longitudinal Breast Cancer Risk Prediction. In: *MICCAI* (2025)
23. Thrun, S., Hansen, S., Sun, Z., Blum, N., Salahuddin, S.A., Wang, X., Wickstrøm, K., Wetzer, E., Jenssen, R., Stille, M., Kampffmeyer, M.: The impact of longitudinal mammogram alignment on breast cancer risk assessment. *arXiv preprint arXiv:2511.08328* (2025)
24. Thrun, S., Hansen, S., Sun, Z., Blum, N., Salahuddin, S.A., Wickstrøm, K., Wetzer, E., Jenssen, R., Stille, M., Kampffmeyer, M.: Reconsidering explicit longitudinal

- mammography alignment for enhanced breast cancer risk prediction. In: MICCAI (2025)
25. Uno, H., Cai, T., Pencina, M., D’Agostino, R., Wei, L.: On the c-statistics for evaluating overall adequacy of risk prediction procedures with censored survival data. *Statistics in Medicine* **30**(10) (2011)
 26. Varela, C., Karssemeijer, N., Hendriks, J.H., Holland, R.: Use of prior mammograms in the classification of benign and malignant masses. *European Journal of Radiology* **56**(2) (2005)
 27. Wang, A., Vachon, C., Brandt, K., Ghosh, K.: Breast density and breast cancer risk: a practical review. *Mayo Clin. Proc.* **89**(4) (2014)
 28. Wang, X., Tan, T., Gao, Y., Marcus, E., Han, L., Portaluri, A., Zhang, T., Lu, C., Liang, X., Beets-Tan, R., Teuwen, J., Mann, R.: Ordinal Learning: Longitudinal Attention Alignment Model for Predicting Time to Future Breast Cancer Events from Mammograms . In: MICCAI (2024)
 29. Yala, A., Mikhael, P.G., Strand, F., Lin, G., Smith, K., Wan, Y.L., Lamb, L., Hughes, K., Lehman, C., Barzilay, R.: Toward robust mammography-based models for breast cancer risk. *Sci. Transl. Med.* **13**(578) (2021)
 30. Zielonke, N., Gini, A., Jansen, E., Anttila, A., Segnan, N., Ponti, A., Veerus, P., de Koning, H., van Ravesteijn, N., Heijnsdijk, E., Heinävaara, S., Sarkeala, T., Cañada, M., Pitter, J., Széles, G., Voko, Z., Minozzi, S., Senore, C., van Balle-gooyen, M., Driesprong-de Kok, I., Lansdorp-Vogelaar, I., Ivanus, U., Jarm, K., Novak Mlakar, D., Primic-Žakelj, M., McKee, M., Priaulx, J.: Evidence for reducing cancer-specific mortality due to screening for breast cancer in europe: A systematic review. *Eur. J. Cancer* **127** (2020)