



Mining Negative Temporal Contexts for False Positive Suppression in Real-Time Ultrasound Lesion Detection

Haojun Yu¹, Youcheng Li¹, QuanLin Wu^{2,5}, Ziwei Zhao², Dengbo Chen⁴,
Dong Wang¹, and Liwei Wang^{1,3}(✉)

¹ National Key Laboratory of General Artificial Intelligence, School of Intelligence
Science and Technology, Peking University, Beijing, China

{haojunyu,wanglw}@pku.edu.cn

² Center of Data Science, Peking University, Beijing, China

³ Center for Machine Learning Research, Peking University, Beijing, China

⁴ Yizhun Medical AI Co., Ltd., Beijing, China

⁵ Pazhou Laboratory, Huangpu, Guangdong, China

Abstract. During ultrasonic scanning processes, real-time lesion detection can assist radiologists in accurate cancer diagnosis. However, this essential task remains challenging and underexplored. General-purpose real-time object detection models can mistakenly report obvious false positives (FPs) when applied to ultrasound videos, potentially misleading junior radiologists. One key issue is their failure to utilize negative symptoms in previous frames, denoted as *negative temporal contexts* (NTC) [15]. To address this issue, we propose to extract contexts from previous frames, including NTC, with the guidance of inverse optical flow. By aggregating extracted contexts, we endow the model with the ability to suppress FPs by leveraging NTC. We call the resulting model *UltraDet*. The proposed UltraDet demonstrates significant improvement over previous state-of-the-arts and achieves real-time inference speed. We release the code, checkpoints, and high-quality labels of the CVA-BUS dataset [9] in <https://github.com/HaojunYu1998/UltraDet>.

Keywords: Ultrasound Video · Real-time Lesion Detection · Negative Temporal Context · False Positive Suppression

1 Introduction

Ultrasound is a widely-used imaging modality for clinical cancer screening. Deep Learning has recently emerged as a promising approach for ultrasound lesion detection. While previous works focused on lesion detection in still images [25] and offline videos [9, 11, 22], this paper explores real-time ultrasound video lesion detection. Real-time lesion prompts can assist radiologists during scanning, thus

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-43987-2_1.

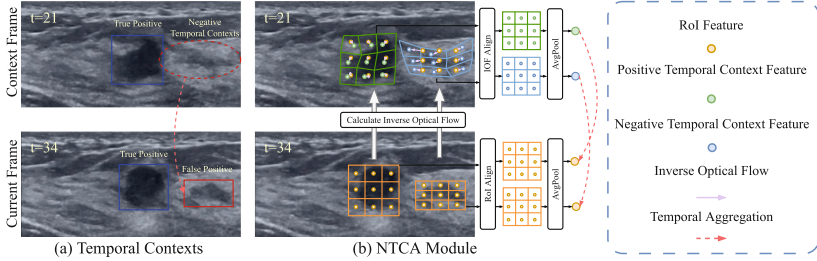


Fig. 1. Illustration of Negative Temporal Context Aggregation (NTCA) module. (a) Our motivation: mining negative temporal contexts for FP suppression. (b) The NTCA module leverages temporal contexts to suppress the FP. (Color figure online)

being more helpful to improve the accuracy of diagnosis. This task requires the model to infer faster than 30 frames per second (FPS) [19] and only previous frames are available for current frame processing.

Previous general-purpose detectors [1, 2] report simple and obvious FPs when applied to ultrasound videos, *e.g.* the red box in Fig. 1(a). These FPs, attributable to non-lesion anatomies, can mislead junior readers. These anatomies appear like lesions in certain frames, but typically show negative symptoms in adjacent frames when scanned from different positions. So experienced radiologists will refer to corresponding regions in previous frames, denoted as *temporal contexts* (TC), to help restrain FPs. If TC of a lesion-like region exhibit negative symptoms, denoted as *negative temporal contexts* (NTC), radiologists are less likely to report it as a lesion [15]. Although important, the utilization of NTC remains unexplored. In natural videos, as transitions from non-objects to objects are implausible, previous works [1, 2, 20] only consider inter-object relationships. As shown in Sect. 4.4, the inability to utilize NTC is a key issue leading to the FPs reported by general-purpose detectors.

To address this issue, we propose a novel *UltraDet* model to leverage NTC. For each Region of Interest (RoI) \mathcal{R} proposed by a basic detector, we extract temporal contexts from previous frames. To compensate for inter-frame motion, we generate deformed grids by applying inverse optical flow to the original regular RoI grids, illustrated in Fig. 1. Then we extract the RoI features from the deformed grids in previous frames and aggregate them into \mathcal{R} . We call the overall process *Negative Temporal Context Aggregation* (NTCA). The NTCA module leverages RoI-level NTC which are crucial for radiologists but ignored in previous works, thereby effectively improving the detection performance in a reliable and interpretable way. We plug the NTCA module into a basic real-time detector to form *UltraDet*. Experiments on CVA-BUS dataset [9] demonstrate that *UltraDet*, with real-time inference speed, significantly outperforms previous works, reducing about 50% FPs at a recall rate of 0.90.

Our contributions are four-fold. (1) We identify that the failure of general-purpose detectors on ultrasound videos derives from their incapability of utilizing negative temporal contexts. (2) We propose a novel *UltraDet* model, incorpo-

rating an NTCA module that effectively leverages NTC for FP suppression. (3) We conduct extensive experiments to demonstrate the proposed UltraDet significantly outperforms the previous state-of-the-arts. (4) We release high-quality labels of the CVA-BUS dataset [9] to facilitate future research.

2 Related Works

Real-Time Video Object Detection is typically achieved by single-frame detectors, often with temporal information aggregation modules. One-stage detectors [5, 8, 16, 21] use only intra-frame information, DETR-based detectors [20, 26] and Faster R-CNN-based detectors [1, 2, 7, 14, 23, 28] are also widely utilized in video object detection. They aggregate temporal information by mining inter-object relationships without considering NTC.

Ultrasound Lesion Detection [10] can assist radiologists in clinical practice. Previous works have explored lesion detection in still images [25] and offline videos [9, 11, 22]. Real-time video lesion detection is underexplored. In previous works, YOLO series [17, 24] and knowledge distillation [19] are used to speed up inference. However, these works use single-frame detectors or post-process methods while learnable inter-frame aggregation modules are not adopted. Thus their performances are far from satisfactory.

Optical Flow [3] is used to guide ultrasound segmentation [12], motion estimation [4] and elastography [13]. For the first time, we use inverse optical flow to guide temporal context information extraction.

3 Method

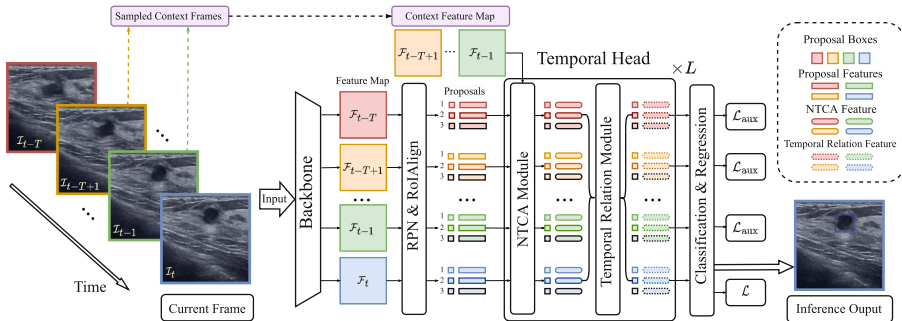


Fig. 2. Illustration of UltraDet model. The yellow and green frames are sampled as context frames, and their feature maps are inputs of the NTCA module. (Color figure online)

In real-time video lesion detection, given the current frame \mathcal{I}_t and a sequence of T previous frames as $\{\mathcal{I}_\tau\}_{\tau=t-T}^{t-1}$, the goal is to detect lesions in \mathcal{I}_t by exploiting the temporal information in previous frames as illustrated in Fig. 2.

3.1 Basic Real-Time Detector

The basic real-time detector comprises three main components: a lightweight backbone (*e.g.* ResNet34 [6]), a Region Proposal Network (RPN) [14], and a Temporal Relation head [2]. The backbone is responsible for extracting feature map \mathcal{F}_τ of frame \mathcal{I}_τ . The RPN generates proposals consisting of boxes \mathcal{B}_τ and proposal features \mathcal{Q}_τ using RoI Align and average pooling:

$$\mathcal{Q}_\tau = \text{AvgPool}(\text{RoIAlign}(\mathcal{F}_\tau, \mathcal{B}_\tau)) \quad (1)$$

where $\tau = t - T, \dots, t - 1, t$. To aggregate temporal information, proposals from all $T + 1$ frames are fed into the Temporal Relation head and updated with inter-lesion information extracted via a relation operation [7]:

$$\mathcal{Q}^l = \mathcal{Q}^{l-1} + \text{Relation}(\mathcal{Q}^{l-1}, \mathcal{B}) \quad (2)$$

where $l = 1, \dots, L$ represent layer indices, \mathcal{B} and \mathcal{Q} are the concatenation of all \mathcal{B}_τ and \mathcal{Q}_τ , and $\mathcal{Q}^0 = \mathcal{Q}$. We call this basic real-time detector *BasicDet*. The BasicDet is conceptually similar to RDN [2] but does not incorporate relation distillation since the number of lesions and proposals in this study is much smaller than in natural videos.

3.2 Negative Temporal Context Aggregation

In this section, we present the Negative Temporal Context Aggregation (NTCA) module. We sample T_{ctxt} context frames from T previous frames, then extract temporal contexts (TC) from context frames and aggregate them into proposals. We illustrate the NTCA module in Fig. 3 and elaborate on details as follows.




Fig. 3. Illustration of the Negative Temporal Context Aggregation module.

Inverse Optical Flow Align. We propose the Inverse Optical Flow Align (IOF Align) to extract TC features. For the current frame \mathcal{I}_t and a sampled context frame \mathcal{I}_τ with $\tau < t$, we extract TC features from the context feature map \mathcal{F}_τ with the corresponding regions. We use inverse optical flow $\mathcal{O}_{t \rightarrow \tau} \in \mathbb{R}^{H \times W \times 2}$

to transform the RoIs from frame t to τ : $\mathcal{O}_{t \rightarrow \tau} = \text{FlowNet}(\mathcal{I}_t, \mathcal{I}_\tau)$ where H, W represent height and width of feature maps. The $\text{FlowNet}(\mathcal{I}_t, \mathcal{I}_\tau)$ is a fixed network [3] to predict optical flow from \mathcal{I}_t to \mathcal{I}_τ . We refer to $\mathcal{O}_{t \rightarrow \tau}$ as *inverse optical flow* because it represents the optical flow in inverse chronological order from t to τ . We conduct IOF Align and average pooling to extract $\mathcal{C}_{t,\tau}$:

$$\mathcal{C}_{t,\tau} = \text{AvgPool}(\text{IOFAlign}(\mathcal{F}_\tau, \mathcal{B}_t, \mathcal{O}_{t \rightarrow \tau})) \quad (3)$$

where $\text{IOFAlign}(\mathcal{F}_\tau, \mathcal{B}_t, \mathcal{O}_{t \rightarrow \tau})$ extracts context features in \mathcal{F}_τ from deformed grids generated by applying offsets $\mathcal{O}_{t \rightarrow \tau}$ to the original regular grids in \mathcal{B}_t , which is illustrated in the Fig. 1(b).

Temporal Aggregation. We concatenate $\mathcal{C}_{t,\tau}$ in all T_{ctx} context frames to form \mathcal{C}_t and enhance proposal features by fusing \mathcal{C}_t into \mathcal{Q}_t :

$$\mathcal{Q}_{\text{ctx},t}^l = \mathcal{Q}_{\text{ctx},t}^{l-1} + \text{Attention}(\mathcal{Q}_{\text{ctx},t}^{l-1}, \mathcal{C}_t, \mathcal{C}_t) \quad (4)$$

where $l = 1, \dots, L$ represent layer indices, $\mathcal{Q}_{\text{ctx},t}^0 = \mathcal{Q}_t$, and $\text{Attention}(Q, K, V)$ is Multi-head Attention [18]. We refer to the concatenation of all TC-enhanced proposal features in $T+1$ frames as \mathcal{Q}_{ctx} . To extract consistent TC, the context frames of T previous frames are shared with the current frame.

3.3 UltraDet for Real-Time Lesion Detection

We integrate the NTCA module into the BasicDet introduced in Sect. 3.1 to form the UltraDet model, which is illustrated in Fig. 2. The head of UltraDet consists of stacked NTCA and relation modules:

$$\mathcal{Q}^l = \mathcal{Q}_{\text{ctx}}^l + \text{Relation}(\mathcal{Q}_{\text{ctx}}^l, \mathcal{B}). \quad (5)$$

During training, we apply regression and classification losses $\mathcal{L} = \mathcal{L}_{\text{reg}} + \mathcal{L}_{\text{cls}}$ to the current frame. To improve training efficiency, we apply auxiliary losses $\mathcal{L}_{\text{aux}} = \mathcal{L}$ to all previous T frames. During inference, the UltraDet model uses the current frame and T previous frames as inputs and generates predictions only for the current frame. This design endows the UltraDet with the ability to perform real-time lesion detection.

4 Experiments

4.1 Datasets

CVA-BUS Datasets. We use the open source CVA-BUS dataset that consists of 186 valid videos, which is proposed in CVA-Net [9]. We split the dataset into train-val (154 videos) and test (32 videos) sets. In the train-val split, there are 21423 frames with 170 lesions. In the test split, there are 3849 frames with 32 lesions. We focus on the lesion detection task and do not utilize the benign/malignant classification labels provided in the original dataset.

High-Quality Labels. The bounding box labels provided in the original CVA-BUS dataset are unsteady and sometimes inaccurate, leading to jiggling and inaccurate model predictions. We provide a new version of high-quality labels that are re-annotated by experienced radiologists. We reproduce all baselines using our high-quality labels to ensure a fair comparison. Visual comparisons of two versions of labels are available in supplementary materials. To facilitate future research, we will release these high-quality labels.

Table 1. Quantitative results of real-time lesion detection on CVA-BUS [9].

Model	Type	Pr80	Pr90	FP80	FP90	AP50	R@16	FPS
One-Stage Detectors								
YOLOX [5]	Image	69.7 _{3.7}	43.4 _{7.7}	23.8 _{4.8}	87.6 _{24.5}	80.4 _{1.6}	97.5 _{0.5}	59.8
RetinaNet [8]	Image	75.7 _{2.5}	57.2 _{2.9}	9.3 _{2.0}	32.8 _{6.5}	84.5 _{1.0}	95.1 _{0.6}	53.6
FCOS [16]	Image	87.2 _{2.2}	72.2 _{5.1}	11.0 _{2.4}	23.0 _{3.7}	89.5 _{1.4}	98.8 _{0.3}	56.1
DeFCN [21]	Image	81.5 _{1.8}	67.5 _{2.3}	21.1 _{3.2}	33.4 _{4.3}	86.4 _{1.3}	99.3 _{0.3}	51.2
Track-YOLO [24]	Video	75.1 _{2.7}	47.0 _{3.1}	18.1 _{1.9}	74.2 _{14.7}	80.1 _{1.0}	94.7 _{0.9}	46.0
DETR-Based Detectors								
DeformDETR [27]	Image	90.1 _{3.2}	72.7 _{10.6}	5.6 _{2.2}	37.8 _{20.9}	90.5 _{2.0}	98.7 _{0.3}	33.8
TransVOD [26]	Video	92.5 _{2.2}	77.5 _{7.2}	3.1 _{1.3}	23.7 _{11.5}	90.1 _{1.8}	98.4 _{0.4}	24.2
CVA-Net [9]	Video	92.3 _{2.6}	80.2 _{6.1}	4.7 _{2.6}	19.6 _{5.6}	91.6 _{1.9}	98.6 _{0.8}	23.1
PTSEFormer [20]	Video	93.3 _{1.9}	85.4 _{6.0}	2.8 _{1.1}	12.5 _{9.8}	91.5 _{1.6}	97.9 _{1.2}	9.1
FasterRCNN-Based Detectors								
FasterRCNN [14]	Image	91.3 _{0.9}	75.2 _{3.6}	6.9 _{1.4}	34.4 _{6.7}	88.0 _{1.4}	92.4 _{1.0}	49.2
RelationNet [7]	Image	91.4 _{1.3}	79.2 _{2.9}	6.2 _{2.0}	24.4 _{5.6}	87.6 _{1.7}	92.4 _{0.9}	42.7
FGFA [28]	Video	92.9 _{1.5}	82.2 _{4.1}	4.4 _{1.6}	13.3 _{3.7}	90.5 _{1.1}	93.6 _{0.9}	33.8
SELSA [23]	Video	91.6 _{1.7}	80.2 _{2.5}	7.5 _{1.5}	23.3 _{5.5}	89.2 _{1.1}	92.6 _{0.8}	43.8
MEGA [1]	Video	93.9 _{1.5}	86.9 _{2.3}	3.1 _{1.7}	11.7 _{3.0}	90.9 _{1.0}	93.6 _{0.7}	40.2
BasicDet (RDN) [2]	Video	92.4 _{1.0}	83.6 _{2.2}	3.8 _{1.2}	13.4 _{3.2}	88.7 _{1.4}	92.7 _{0.6}	42.2
UltraDet (Ours)	Video	95.7 _{1.2}	90.8 _{1.4}	1.9 _{0.4}	5.7 _{1.6}	91.6 _{1.6}	93.8 _{1.3}	30.4

4.2 Evaluation Metrics

Pr80, Pr90. In clinical applications, it is important for detection models to be sensitive. So we provide frame-level precision values with high recall rates of 0.80 and 0.90, which we denote as Pr80 and Pr90, respectively.

FP80, FP90. We further report lesion-level FP rates as critical metrics. Frame-level FPs are linked by IoU scores to form FP sequences [24]. The number of FP sequences per minute at recall rates of 0.80 and 0.90 are reported as FP80 and FP90, respectively. The unit of lesion-level FP rates is seq/min.

AP50. We provide AP50 instead of mAP or AP75 because the IoU threshold of 0.50 is sufficient for lesion localization in clinical practice. Higher thresholds like 0.75 or 0.90 are impractical due to the presence of blurred lesion edges.

R@16. To evaluate the highest achievable sensitivity, we report the frame-level average recall rates of Top-16 proposals, denoted as R@16.


4.3 Implementation Details

UltraDet Settings. We use FlowNetS [3] as the fixed FlowNet in IOF Align and share the same finding with previous works [4, 12, 13] that the FlowNet trained on natural datasets generalizes well on ultrasound datasets. We set the pooling stride in the FlowNet to 4, the number of UltraDet head layers $L = 2$, the number of previous frames $T = 15$ and $T_{\text{ctxt}} = 2$, and the number of proposals is 16. We cached intermediate results of previous frames and reuse them to speed up inference. Other hyper-parameters are listed in supplementary materials.

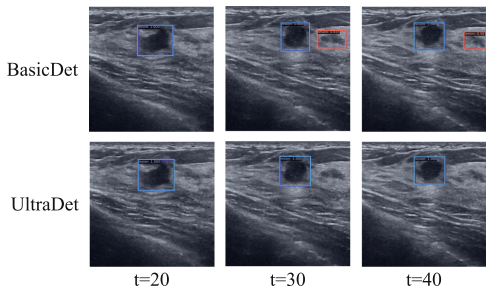
Shared Settings. All models are built in PyTorch framework and trained using eight NVIDIA GeForce RTX 3090 GPUs. We use ResNet34 [6] as backbones and set the number of training iterations to 10,000. We set the feature dimensions of detection heads to 256 and baselines are re-implemented to utilize only previous frames. We refer to our code for more details.

4.4 Main Results

Quantitative Results. We compare performances of real-time detectors with the UltraDet in Table 1. We perform 4-fold cross-validation and report the mean values and standard errors on the test set to mitigate fluctuations. The UltraDet outperforms all previous state-of-the-art in terms of precision and FP rates. Especially, the Pr90 of UltraDet achieves 90.8%, representing a 5.4% absolute improvement over the best competitor, PTSEFormer [20]. Moreover, the FP90 of UltraDet is 5.7 seq/min, reducing about 50% FPs of the best competitor, PTSEFormer. Although CVA-Net [9] achieve comparable AP50 with our method, we significantly improve precision and FP rates over the CVA-Net [9].



(a) Ratios of suppressible FPs



(b) Visual comparisons of BasicDet and UltraDet

Fig. 4. (a) Ratios of FPs that are suppressible by leveraging NTC. (b) Visual comparisons of BasicDet and UltraDet prediction results at recall 0.90. Blue boxes are true positives and red boxes are FPs. (Color figure online)

Importance of NTC. In Fig. 4(a), we illustrate the FP ratios that can be suppressed by using NTC. The determination of whether FPs can be inhibited

by NTC is based on manual judgments of experienced radiologists. We find that about 50%–70% FPs of previous methods are suppressible. However, by utilizing NTC in our UltraDet, we are able to effectively prevent this type of FPs.

Inference Speed. We run inference using one NVIDIA GeForce RTX 3090 GPU and report the inference speed in Table 1. The UltraDet achieves an inference speed of 30.4 FPS and already meets the 30 FPS requirement. Using TensorRT, we further optimize the speed to 35.2 FPS, which is sufficient for clinical applications [19].

Qualitative Results. Figure 4(b) visually compares BasicDet and UltraDet. The BasicDet reports FPs at $t = 30$ and 40 as it fails to leverage NTC when $t = 20$, while the UltraDet successfully suppresses FPs with the NTCA module.

4.5 Ablation Study

Table 2. Ablation study of each NTCA sub-module.

IOFAlign	TempAgg	Pr80	Pr90	FP80	FP90	AP50	R@16	FPS
-	-	92.4 _{1.0}	83.6 _{2.2}	3.8 _{1.2}	13.4 _{3.2}	88.7 _{1.4}	92.7 _{0.6}	42.2
-	✓	93.7 _{1.8}	84.3 _{1.4}	3.4 _{1.0}	12.5 _{0.8}	90.0 _{1.9}	93.0 _{1.3}	37.2
✓	-	94.5 _{2.3}	88.7 _{2.2}	2.6 _{0.6}	9.0 _{1.5}	90.5 _{1.9}	92.9 _{1.4}	32.3
✓	✓	95.7_{1.2}	90.8_{1.4}	1.9_{0.4}	5.7_{1.6}	91.6_{1.6}	93.8_{1.3}	30.4

Effectiveness of Each Sub-module. We ablate the effectiveness of each sub-module of the NTCA module in Table 2. Specifically, we replace the IOF Align with an RoI Align and the Temporal Aggregation with a simple average pooling in the temporal dimension. The results demonstrate that both IOF Align and Temporal Aggregation are crucial, as removing either of them leads to a noticeable drop in performance.

Table 3. Design of the NTCA Module.

Num	Pr80	Pr90	FP80	FP90	AP50	R@16	FPS
Feature-level	94.0 _{0.9}	84.6 _{2.9}	2.9 _{0.8}	11.7 _{3.0}	90.8 _{0.7}	93.3 _{0.6}	30.6
RoI-level	95.7_{1.2}	90.8_{1.4}	1.9_{0.4}	5.7_{1.6}	91.6_{1.6}	93.8_{1.3}	30.4
Both-level	94.6 _{1.0}	88.7 _{1.8}	2.5 _{0.9}	7.9 _{2.4}	90.8 _{1.5}	93.8_{0.9}	26.9

Design of the NTCA Module. Besides RoI-level TC aggregation in UltraDet, feature-level aggregation is also feasible. We plug the optical flow feature warping proposed in FGFA [28] into the BasicDet and report the results in Table 3. We find RoI-level aggregation is more effective than feature-level, and both-level aggregation provides no performance gains. This conclusion agrees with radiologists’ skills to focus more on local regions instead of global information.

5 Conclusion

In this paper, we address the clinical challenge of real-time ultrasound lesion detection. We propose a novel Negative Temporal Context Aggregation (NTCA) module, imitating radiologists' diagnosis processes to suppress FPs. The NTCA module leverages negative temporal contexts that are essential for FP suppression but ignored in previous works, thereby being more effective in suppressing FPs. We plug the NTCA module into a BasicDet to form the UltraDet model, which significantly improves the precision and FP rates over previous state-of-the-arts while achieving real-time inference speed. The UltraDet has the potential to become a real-time lesion detection application and assist radiologists in more accurate cancer diagnosis in clinical practice.

Acknowledgements. This work is supported by National Key R&D Program of China (2022ZD0114900) and National Science Foundation of China (NSFC62276005).

References

1. Chen, Y., Cao, Y., Hu, H., Wang, L.: Memory enhanced global-local aggregation for video object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 10337–10346 (2020)
2. Deng, J., Pan, Y., Yao, T., Zhou, W., Li, H., Mei, T.: Relation distillation networks for video object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 7023–7032 (2019)
3. Dosovitskiy, A., et al.: FlowNet: learning optical flow with convolutional networks. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2758–2766 (2015)
4. Evain, E., Faraz, K., Grenier, T., Garcia, D., De Craene, M., Bernard, O.: A pilot study on convolutional neural networks for motion estimation from ultrasound images. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **67**(12), 2565–2573 (2020)
5. Ge, Z., Liu, S., Wang, F., Li, Z., Sun, J.: YOLOX: exceeding yolo series in 2021. arXiv preprint [arXiv:2107.08430](https://arxiv.org/abs/2107.08430) (2021)
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
7. Hu, H., Gu, J., Zhang, Z., Dai, J., Wei, Y.: Relation networks for object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3588–3597 (2018)
8. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2980–2988 (2017)
9. Lin, Z., Lin, J., Zhu, L., Fu, H., Qin, J., Wang, L.: A new dataset and a baseline model for breast lesion detection in ultrasound videos. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) MICCAI 2022, Part III. LNCS, vol. 13433, pp. 614–623. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-16437-8_59

10. Liu, S., et al.: Deep learning in medical ultrasound analysis: a review. *Engineering* **5**(2), 261–275 (2019)
11. Movahedi, M.M., Zamani, A., Parsaei, H., Tavakoli Golpaygani, A., Haghghi Poya, M.R.: Automated analysis of ultrasound videos for detection of breast lesions. *Middle East J. Cancer* **11**(1), 80–90 (2020)
12. Nguyen, A., et al.: End-to-end real-time catheter segmentation with optical flow-guided warping during endovascular intervention. In: 2020 IEEE International Conference on Robotics and Automation (ICRA), pp. 9967–9973. IEEE (2020)
13. Peng, B., Xian, Y., Jiang, J.: A convolution neural network-based speckle tracking method for ultrasound elastography. In: 2018 IEEE International Ultrasonics Symposium (IUS), pp. 206–212. IEEE (2018)
14. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. In: *Advances in Neural Information Processing Systems*, vol. 28 (2015)
15. Spak, D.A., Plaxco, J., Santiago, L., Dryden, M., Dogan, B.: BI-RADS® fifth edition: a summary of changes. *Diagn. Interv. Imaging* **98**(3), 179–190 (2017)
16. Tian, Z., Shen, C., Chen, H., He, T.: FCOS: fully convolutional one-stage object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9627–9636 (2019)
17. Tiyyattachai, T., et al.: The feasibility to use artificial intelligence to aid detecting focal liver lesions in real-time ultrasound: a preliminary study based on videos. *Sci. Rep.* **12**(1), 7749 (2022)
18. Vaswani, A., et al.: Attention is all you need. In: *Advances in Neural Information Processing Systems*, vol. 30 (2017)
19. Vaze, S., Xie, W., Namburete, A.I.: Low-memory CNNs enabling real-time ultrasound segmentation towards mobile deployment. *IEEE J. Biomed. Health Inform.* **24**(4), 1059–1069 (2020)
20. Wang, H., Tang, J., Liu, X., Guan, S., Xie, R., Song, L.: PTSEFormer: progressive temporal-spatial enhanced transformer towards video object detection. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *ECCV 2022, Part VIII*. LNCS, vol. 13668, pp. 732–747. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-20074-8_42
21. Wang, J., Song, L., Li, Z., Sun, H., Sun, J., Zheng, N.: End-to-end object detection with fully convolutional network. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15849–15858 (2021)
22. Wang, Y., et al.: Key-frame guided network for thyroid nodule recognition using ultrasound videos. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) *MICCAI 2022, Part IV*. LNCS, vol. 13434, pp. 238–247. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-16440-8_23
23. Wu, H., Chen, Y., Wang, N., Zhang, Z.: Sequence level semantics aggregation for video object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9217–9225 (2019)
24. Wu, X., et al.: CacheTrack-YOLO: real-time detection and tracking for thyroid nodules and surrounding tissues in ultrasound videos. *IEEE J. Biomed. Health Inform.* **25**(10), 3812–3823 (2021)
25. Yap, M.H., et al.: Automated breast ultrasound lesions detection using convolutional neural networks. *IEEE J. Biomed. Health Inform.* **22**(4), 1218–1226 (2017)
26. Zhou, Q., et al.: TransVOD: end-to-end video object detection with spatial-temporal transformers. *IEEE Trans. Pattern Anal. Mach. Intell.* **45**(6), 7853–7869 (2023)

27. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable DETR: deformable transformers for end-to-end object detection. arXiv preprint [arXiv:2010.04159](https://arxiv.org/abs/2010.04159) (2020)
28. Zhu, X., Wang, Y., Dai, J., Yuan, L., Wei, Y.: Flow-guided feature aggregation for video object detection. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 408–417 (2017)