

Compass: Prostate Cancer Detection Needs Multi-View Context

Paul F.R. Wilson^{1,6*}, Mohamed Harmanani^{1,6*(✉)}, Zhuoxin Guo^{1,6},
Obed K. Dzikunu^{2,6}, Hannes Cash³, Adam Kinnaird⁴, Brian Wodlinger⁵,
Purang Abolmaesumi^{2†}, and Parvin Mousavi^{1,6†}

¹ Queen's University, Kingston, Canada

² University of British Columbia, Vancouver, Canada

³ Otto von Guericke University Magdeburg, Germany

⁴ University of Alberta, Edmonton, Canada

⁵ Exact Imaging, Markham, Canada

⁶ Vector Institute, Toronto, Canada

Abstract. Artificial intelligence (AI) analysis of micro-ultrasound (μ US) has shown promise for prostate cancer (PCa) detection. However, most existing AI methods focus on the analysis of single μ US images in isolation. By contrast, expert μ US readers typically assess a full recorded video study, which provides three-dimensional context, to improve PCa detection compared to single-frame analysis. Inspired by this clinical workflow, we propose Compass, a novel AI methodology which models a μ US study as a stream of 2D images. Compass jointly integrates rotational sweep videos of the prostate with μ US frames acquired at the moment of biopsy, and performs evidence aggregation across the study using a transformer conditioned on the probe's rotational angle. Finally, a decoder head predicts frame-level and study-level risk scores for the patient. The model is trained and evaluated using a multi-center clinical trial dataset of μ US studies, including continuous rotational scans of the prostate and videos captured during biopsy acquisition. We compare the proposed method to baseline AI methods from the literature and to risk scores provided by clinical experts. Our framework shows strong performance, highlighting the value of multi-view context for μ US PCa detection, and providing a potentially powerful tool to complement human expertise in μ US-based PCa diagnosis. Our code is available at: <https://github.com/mharmanani/Compass>.

Keywords: Multi-view context · Foundation models · Prostate cancer · Ultrasound

1 Introduction

Prostate cancer (PCa) is the second most commonly diagnosed cancer in men worldwide and one of the leading causes of cancer-related deaths [9]. Accurate

* co-first authors (interchangeable order)

† co-senior authors (✉) 22mh5@queensu.ca

early detection and risk stratification of PCa are essential to guide treatment planning and ensuring favorable outcomes. Medical imaging plays a crucial role in the clinical workflow by localizing suspicious regions to guide biopsies, after which histopathology provides the definitive diagnosis and grade that drive treatment decisions.

High-resolution micro-ultrasound (μ US) is a new and promising modality for PCa detection, offering real-time visualization and lesion localization at substantially higher spatial resolution than conventional ultrasound, with a lower cost than multi-parametric MRI (mpMRI). The PRI-MUS [3] protocol and scoring system assigns risk scores from μ US by interpreting prostate echotexture, and has shown non-inferior performance to mpMRI PI-RADS [7]. Despite these advances, μ US interpretation of PRI-MUS score is subject to operator expertise and can suffer from significant inter-observer variability. To facilitate reliable lesion targeting across diverse clinical settings and workflows, there is a need for an objective, user-agnostic tool for μ US analysis and PCa detection.

Modern artificial intelligence (AI) methods have demonstrated strong diagnostic accuracy across medical imaging, fueled in part by the recent rise of foundation models (FMs). In PCa detection, FM-based approaches have led progress in mpMRI [10, 8], as well as conventional US [4], temporal-enhanced US [2], and μ US [12, 11, 5]. Most FM-based methods in this space operate primarily on single-plane 2D B-mode images to produce e.g., per-image heatmaps and risk scores [11]. However, μ US interpretation is inherently three-dimensional (3D) and *spatially continuous*: clinicians collect cine-loops of systematic prostate sweeps to localize suspicious patterns across multiple views, and determine biopsy targets. While current literature in US-based PCa detection have explored the integration of information from cine-loops [4], temporal dynamics [2], and heuristic slice-level aggregation into AI models [5], the analysis of full prostate sweeps remains largely untapped. There is a clear gap between how clinicians collect and interpret ultrasound data versus the state-of-the-art in AI-based PCa detection.

To bridge this gap, we propose **Compass**, a multi-view framework that leverages the spatio-temporal structure of complete prostate studies for accurate cancer detection. Compass aggregates evidence from two sources: (i) cine-loops of μ US sweeps of the prostate, recorded while the probe rotates around its main axis, providing rich multi-view context, and (ii) μ US 2D biopsy frames. We evaluate our framework on a multi-center clinical trial dataset, comparing against prior baselines and expert risk assessments. Our main contributions are:

1. A multi-view reframing of μ US PCa detection that leverages the rich clinical context from 3D cine-loops, enabling robust patient-level decision-making from spatially continuous evidence rather than isolated frames.
2. A novel architecture for mixed-evidence reasoning that jointly integrates 3D sweeps with 2D biopsy frames, conditioned on the ultrasound probe’s rotational angle, effectively aligning *global* context with *local* evidence.
3. Validation on a new multi-center dataset to compare with human experts and state-of-the-art (SOTA) AI methods, establishing the practical value of AI-based multi-view decision support in μ US PCa detection.

2 Materials

Data. We used a subset of data from the OPTIMUM [7] clinical trial, comprising proprietary μ US-guided prostate biopsy studies from two clinical centers: University of Alberta, Edmonton, Canada (UA) and Professionelle Urologie, Berlin, Germany (PU). The cohort includes 118 patients (76 UA, 42 PU). The ExactVu micro-ultrasound platform (Exact Imaging, Markham, Canada), which uses a high-frequency 29 MHz side-firing linear-array transrectal probe (512 elements, 90 μ m pitch), was used. During each study, at least one initial transverse scan was performed by rotating the side-firing probe about its longitudinal axis, sweeping the imaging plane from the left lateral aspect of the prostate through midline to the right lateral aspect. Subsequently, a median of 15 biopsy samples per patient were acquired from systematic sextant biopsy locations and from any suspicious lesions identified by the performing urologist using the PRI-MUS [3] protocol. UA used transrectal while PU used transperineal biopsies. For each biopsy core, multi-frame B-mode cine sequences were recorded. For each core, the needle-trace region was annotated, and we use the frame immediately preceding needle firing as the canonical biopsy image for core-level modeling. Throughout the procedure, an inertial measurement unit continuously records the roll angle about the probe’s longitudinal axis, providing an approximation of the left-right position of the imaging plane.

Labels and clinical biomarkers. Histopathology provides per-core diagnosis and ISUP Grade Group (GG). For a biopsy sample we define clinically significant prostate cancer (csPCa) as $GG \geq 2$ for a biopsy core. For a patient we define csPCa as having at least one biopsy sample containing csPCa, following [7]. The dataset also provides routinely available clinical biomarkers such as the patient’s age and prostate specific antigen (PSA).

3 Methods

An overview of our method is shown in Figure 1.

Sweep-level feature extraction. Each μ US rotational sweep is treated as a set of frames with known transducer roll angles. Let $\{(F_t, \theta_t)\}_{t=1}^T$ denote the selected sweep frames and their roll angles. A frozen image encoder produces per-frame visual tokens $U_t = \Phi(F_t) \in \mathbb{R}^{M \times d}$, where M is the number of image tokens and $d = 256$. We adopt ProstNFound⁺ [11] as the backbone, due to its strong performance on and robust prostate US representations.

To encode acquisition geometry, we embed the continuous roll angle with a sinusoidal encoder $r_t = \text{PE}(\theta_t) \in \mathbb{R}^d$. We then summarize each frame using a lightweight attention module $\text{Dec}_{\text{sw}}(\cdot)$, and add the roll embedding and shared sweep token embedding τ^{sw} to obtain a sweep token z_t^{sw} per timestep:

$$z_t^{\text{sw}} = \text{Dec}_{\text{sw}}(U_t) + r_t + \tau^{\text{sw}} \in \mathbb{R}^d.$$

Biopsy-level frame encoding. We denote the biopsy frames and their roll angles by $\{(I_i, \theta_i)\}_{i=1}^N$. Following ProstNFound⁺, we embed the biomarker vector




Fig. 1. Compass: Dual-branch patient-level framework that integrates two complementary sources of evidence: (i) rotational μ US sweeps that provide dense multi-view coverage of the gland, and (ii) biopsy core frames paired with clinical metadata.

$m = (\text{Age}, \text{PSA})$ using a prompt encoder Π that maps normalized scalars to a small set of K learnable prompt tokens $P = \Pi(m) \in \mathbb{R}^{K \times d}$.

We encode each biopsy frame and roll angle with the same frozen ProstateNet encoder and positional embedding module to obtain visual tokens and roll embeddings $V_i = \Phi(I_i) \in \mathbb{R}^{M \times d}$ and $r_i = \text{PE}(\theta_i)$.

We then fuse clinical prompts with image evidence using a lightweight decoder $\text{Dec}_{\text{bx}}(\cdot)$ that performs self-attention and cross-attention between the visual tokens and the prompt set. The roll angle embedding and shared biopsy token embedding τ_{bx} is added, resulting in the per-biopsy tokens $z_i^{\text{bx}} \in \mathbb{R}^d$:

$$z_i^{\text{bx}} = \text{Dec}_{\text{bx}}(V_i, P) + r_i + \tau_{\text{bx}} \in \mathbb{R}^d.$$

Mixed-token transformer reasoning. Given biopsy tokens $\{z_i^{\text{bx}}\}_{i=1}^N$ and sweep tokens $\{z_t^{\text{sw}}\}_{t=1}^T$, we form a single mixed-token sequence by concatenation and prepend a learnable [CLS] token:

$$X = [\text{CLS}, z_1^{\text{bx}}, \dots, z_N^{\text{bx}}, z_1^{\text{sw}}, \dots, z_T^{\text{sw}}] \in \mathbb{R}^{(1+N+T) \times d}.$$

We pass the sequence through a transformer encoder $\text{Trans}(\cdot)$ that applies global multi-head self-attention over *all* tokens:

$$H = \text{Trans}(X), \quad H \in \mathbb{R}^{(1+N+T) \times d}.$$

This design enables cross-branch evidence exchange: biopsy tokens can contextualize sparse biopsy cues with broader gland coverage, while sweep tokens can

emphasize regions that are most consistent with biopsy-confirmed malignancy. The patient-level score is predicted from the [CLS] state using a class decoder head, and biopsy-level scores are predicted from the biopsy token states using a frame/biopsy decoder head.

$$\hat{y}_{\text{pat}} = h_{\text{pat}}(H_{\text{CLS}}), \quad \hat{y}_{\text{bio}}^{(i)} = h_{\text{bio}}(H_i), \quad i = 1, \dots, N.$$

Loss function. We train Compass with a joint objective combining patient-level supervision and biopsy-level auxiliary supervision:

$$\mathcal{L} = \mathcal{L}_{\text{Pat}} + \lambda \mathcal{L}_{\text{Bio}} = \text{CE}(\hat{y}_{\text{pat}}, y_{\text{pat}}) + \frac{\lambda}{N} \sum_{i=1}^N \text{CE}(\hat{y}_{\text{bio}}^{(i)}, y_{\text{bio}}^{(i)}),$$

where λ is a constant weighting hyperparameter.

4 Experiments

Baselines & Evaluation. We compare Compass to three baseline families: (i) frame-level classifiers (CLIP ViT-L/14, MedSAM, Cinepro, ProstNFound⁺ [11]); (ii) multi-instance learning (MIL) variants using transformer aggregation over frame/core embeddings (Cinepro-MIL, MedSAM-MIL, ProstNFound⁺-MIL); and (iii) video models, including ViViT [1], 3D-ResNet18, and hybrid encoder–transformer models (e.g., MicroSegnet [6], ProstNFound⁺ + ViViT). All experiments used 5-fold patient-level cross-validation. We report patient-level AUROC, balanced accuracy, and sensitivity at 60% specificity (Sen@60), as well as biopsy-level AUROC and balanced accuracy at the epoch with the best mean patient AUROC across folds.

Implementation Details. All experiments were conducted in PyTorch using a single NVIDIA GPU with mixed-precision training. We optimized all models with AdamW (weight decay 10^{-4}). Compass was trained for 100 epochs with a constant learning rate of 3×10^{-5} and a 5-epoch linear warmup, using a batch size of 2. The backbone was frozen during training, gradient clipping was set to 1.0, and the token-level loss weight λ was 0.05.

5 Results and Discussion

Diagnostic Performance. Table 1 compares our performance against several prominent baselines from the literature. Among baselines that use 2D biopsy planes only, CLIP is the weakest, likely because it has been trained mainly with non-ultrasound data. Cinepro exhibits moderate performance at the core-level due to its robust training strategy, but falls behind MedSAM at the patient-level. Interestingly, MedSAM obtains very low sensitivity, highlighting the weaknesses of naive FM fine-tuning and the importance of domain-specific training strategies. Finally, ProstNFound⁺ achieves the highest performance among baselines (+4.1% Patient AUROC, +12.9% Sen@60) since it integrates relevant domain

Table 1. Patient and biopsy-level classification performance across relevant baselines. We report the mean and standard deviation over 5-fold cross-validation.

Method	Patient-level			Biopsy-level	
	AUROC	B.Acc.	Sen@60	AUROC	B.Acc.
CLIP ViT-L/14	67.9±7.9	71.3±5.3	57.9±18.0	62.6±11.4	62.9±8.0
Cinepro	72.6±6.5	76.9±8.3	68.2±13.1	68.6±7.7	67.6±4.2
MedSAM	74.6±5.8	76.8±5.3	43.1±24.4	61.7±4.3	61.6±3.1
ProstNFound ⁺	78.7±15.4	82.9±12.5	76.7±17.8	68.8±10.3	65.9±6.6
Cinepro-MIL	76.3±13.9	76.1±11.2	73.3±19.2	–	–
MedSAM-MIL	75.7±13.7	76.4±10.3	73.3±19.2	–	–
ProstNFound ⁺ -MIL	77.3±14.5	78.1±9.4	82.5±15.3	–	–
ViViT	61.9±13.9	66.0±10.4	61.7±23.1	–	–
3D-ResNet18	71.1±6.0	72.3±4.6	73.2±11.9	–	–
MicroSegNet+ViViT	67.7±7.9	69.0±5.2	66.8±11.8	–	–
ProstNFound ⁺ +ViViT	66.4±15.3	68.8±9.5	66.6±17.6	–	–
PRI-MUS (Experts)	78.5±11.1	78.3±6.9	83.8±15.8	73.1±8.9	71.9±7.5
Compass (ours)	87.2±8.7	85.0±8.2	89.9±10.0	<u>70.5±8.6</u>	<u>68.3±6.9</u>


knowledge via prompts [11, 12].

Compared to frame-level baselines, MIL primarily improves sensitivity by aggregating evidence over multiple instances, reducing per-frame noise and stabilizing threshold behavior. For example, MedSAM-MIL yields a large gain in sensitivity (+30%) over MedSAM, and ProstNFound⁺-MIL substantially improves sensitivity (+5.8%) over its single-frame variant. This highlights the benefit of set-based aggregation in biopsy workflows: by pooling evidence across multiple frames, MIL reduces noise and sampling variance, and better reflects the clinical reality that csPCa cues are sparse and heterogeneous across the prostate gland. The video baselines underperform compared to MIL ($\approx 7\%$ Patient AUROC, Bal.Acc. for ProstNFound⁺+ViViT), which is expected in low-data regimes, as spatio-temporal models are data hungry and struggle under weak supervision.

Finally, Compass achieves the highest *patient-level* performance compared to other baselines in this cohort, due to several key reasons: (i) it leverages a strong clinical backbone to produce stable per-image evidence tokens, and (ii) incorporates sweep information rather than treating frames independently. Crucially, (iii) Compass performs *cross-branch reasoning* by fusing biopsy and sweep tokens with roll-aware conditioning in a single transformer. This lets Compass align localized biopsy evidence with rotation-consistent sweep patterns, yielding a large improvement over ProstNFound⁺, MIL baselines, and generic video models. With respect to *core-level* performance, Compass exhibits comparable but lower performance than PRI-MUS (-2.6% AUROC). This suggests that Compass’s value lies not in replacing PRI-MUS, but in complementing it through comprehensive multi-view context that core-level methods do not capture.

Table 2. Ablation experiments at the patient-level show each component’s importance.

Method	Components			AUROC	Bal.Acc.
	Sweeps	Angle PE	Transformer		
No angle embedding	✓	×	✓	82.7	75.1
No transformer	✓	✓	×	74.3	66.2
No sweeps	×	✓	✓	78.8	70.6
Compass (ours)	✓	✓	✓	87.2	85.0

**Fig. 2.** Outcome distributions across score bins for PRI-MUS and Compass at the patient (top) and core (bottom) levels. Bars show percentages of csPCa, insignificant PCa, and benign cases; n indicates bin counts.

Ablation Study. Table 2 isolates the contribution of key components in the Compass framework. Removing the roll-angle positional encoding reduces performance (-4.8% AUROC, -1.4% Bal.Acc.), indicating that encoding probe rotation allows the model to exploit sweep structure beyond treating frames as an unordered set. Removing the transformer causes the largest degradation (-13.2% AUROC, -10.3% Bal.Acc.), underscoring that *cross-branch reasoning* is critical: integrating sweep and biopsy evidence across branches is key to Compass’ success. Finally, the biopsy-only ProstNFound⁺ variant underperforms Compass by 8.7% in AUROC and 5.9% in Bal.Acc., confirming that sweep information provides substantial complementary signal at the patient level. Overall, these ablations suggest Compass’s gains arise from the combination of (i) sweep evidence, (ii) rotation-aware structure, and (iii) transformer-based fusion that enables biopsies and sweep frames to condition one another.

Clinical evaluation. Figure 2 summarizes how csPCa prevalence varies across score strata for PRI-MUS and Compass at both the patient and core levels, complementing the scalar metrics in Table 1. At the patient level, Compass produces a clear monotonic risk gradient: high-score bins concentrate a larger fraction of csPCa, while low-score bins are dominated by benign and clinically




Fig. 3. (a) Model and PRI-MUS score distribution compared to grade group ground truth across biopsies (left to right) for patients with no cancer (top 2 row groups) and csPCa (bottom two row groups). (b) Image-level cancer detection via post-hoc heatmaps conditioned on Compass embeddings.

insignificant disease, indicating more effective study-level separation under the same clinical burden. This is reflected by Compass’s higher patient-level AUROC and improved Sen@60 score in Table 1, suggesting it can prioritize more csPCa-positive patients while maintaining a similar false-positive burden. In contrast, PRI-MUS exhibits greater overlap between csPCa and non-csPCa categories across intermediate bins, consistent with weaker patient-level prioritization.

At the biopsy core level, higher scores are associated with larger percentages of csPCa for both Compass and PRI-MUS. PRI-MUS exhibits a sharper concentration of csPCa at higher scores, and a clearer benign separation at lower scores, corresponding to its higher core-level AUC. This reinforces Compass’s complementary role in the clinical workflow: Compass is best suited for sweep-informed patient triage, while PRI-MUS remains more informative for localized core-level prediction.

Qualitative evaluation. To highlight the efficacy of Compass for patient-level analysis, we show the distribution of its risk scores across biopsies *within* patients (Figure 3a). Compass’s biopsy risk scores typically align with PRI-MUS risk scores and with ground truth pathology findings. We also generate post-hoc heatmaps by training an off-the-shelf mask decoder on top of our model’s frame tokens. Figure 3b shows two representative biopsy cores, where the benign core (upper) has lower activations than the csPCa core (lower).

Limitations and Future Work. This study involved retrospective data from 118 patients at two clinical sites. Future work should validate the method prospectively on larger cohorts from more clinical sites to ensure generalization across potential deployment scenarios. Compass was trained using histopathology findings of biopsy specimens as ground truth, which cannot localize cancer

tissue labels at the pixel level. To improve cancer localization and core-level diagnosis, future work should investigate the use of pixel-level labels, acquired from PRI-MUS annotations or whole-mount histopathology, in model training.

6 Conclusion

We introduced Compass, a study-level framework for μ US prostate cancer detection that integrates two complementary evidence streams at biopsy time: rotational sweeps and biopsy cores. Compass substantially outperforms frame-level, MIL, and end-to-end video baselines. By explicitly modeling sweep context, Compass shifts AI-based μ US analysis from isolated-frame scoring toward study-level risk stratification, aligning more closely with how clinicians interpret full examinations. This opens the door to practical decision support for patient triage and targeted biopsy planning.

Acknowledgments. This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Canadian Institutes of Health Research (CIHR). Parvin Mousavi is supported by the Canada CIFAR AI Chair program, a Canada Research Chair, and the Vector Institute.

Disclosure of Interests. Brian Wodlinger is VP Clinical and Engineering at Exact Imaging. Other authors have no potential conflicts of interest to disclose.

References

1. Arnab, A., Dehghani, M., Heigold, G., Sun, C., Lučić, M., Schmid, C.: Vivit: A video vision transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6836–6846 (2021)
2. Elghareb, T., Harmanani, M., To, M.N.N., Wilson, P., Jamzad, A., Fooladgar, F., Abdelsamad, B., Dzikunu, O., Sojoudi, S., Reznik, G., et al.: Proteus: A spatio-temporal enhanced ultrasound-based framework for prostate cancer detection. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 405–414. Springer (2025)
3. Ghai, S., Eure, G., Fradet, V., Hyndman, M.E., McGrath, T., Wodlinger, B., Pavlovich, C.P.: Assessing Cancer Risk on Novel 29 MHz Micro-Ultrasound Images of the Prostate: Creation of the Micro-Ultrasound Protocol for Prostate Risk Identification. *The Journal of Urology* **196**(2), 562–569 (Aug 2016). <https://doi.org/10.1016/j.juro.2015.12.093>
4. Harmanani, M., Jamzad, A., To, M.N.N., Wilson, P.F., Guo, Z., Fooladgar, F., Sojoudi, S., Gilany, M., Chang, S., Black, P., et al.: Cinepro: Robust training of foundation models for cancer detection in prostate ultrasound cineloops. In: 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2025)
5. Imran, M., Brisbane, W.G., Su, L.M., Joseph, J.P., Shao, W.: Ai-enhanced micro-ultrasound improves detection of clinically significant prostate cancer at biopsy. *BJUI Compass* **7**(2), e70133 (2026)

6. Jiang, H., Imran, M., Muralidharan, P., Patel, A., Pensa, J., Liang, M., Benidir, T., Grajo, J.R., Joseph, J.P., Terry, R., et al.: Microsegnet: A deep learning approach for prostate segmentation on micro-ultrasound images. *Computerized Medical Imaging and Graphics* **112**, 102326 (2024)
7. Kinnaird, A., Luger, F., Cash, H., Ghai, S., Urdaneta-Salegui, L.F., Pavlovich, C.P., Brito, J., Shore, N.D., Struck, J.P., Schostak, M., others, OPTIMUM Investigators: Microultrasonography-Guided vs MRI-Guided Biopsy for Prostate Cancer Diagnosis: The OPTIMUM Randomized Clinical Trial. *JAMA* **333**(19), 1679–1687 (May 2025). <https://doi.org/10.1001/jama.2025.3579>
8. Lee, J.H., Li, C.X., Jahanandish, H., Bhattacharya, I., Vesal, S., Zhang, L., Sang, S., Choi, M.H., Soerensen, S.J.C., Zhou, S.R., Sommer, E.R., Fan, R., Ghanouni, P., Song, Y., Seibert, T.M., Sonn, G.A., Rusu, M.: Prostate-specific foundation models for enhanced detection of clinically significant cancer (2025), <https://arxiv.org/abs/2502.00366>
9. Schafer, E.J., Laversanne, M., Sung, H., Soerjomataram, I., Briganti, A., Dahut, W., Bray, F., Jemal, A.: Recent patterns and trends in global prostate cancer incidence and mortality: an update. *European Urology* **87**(3), 302–313 (2025)
10. Shao, L., Liang, C., Yan, Y., Zhu, H., Jiang, X., Bao, M., Zang, P., Huang, X., Zhou, H., Nie, P., Wang, L., Li, J., Zhang, S., Ren, S.: An mri–pathology foundation model for noninvasive diagnosis and grading of prostate cancer. *Nature Cancer* **6**(10), 1621–1637 (Oct 2025). <https://doi.org/10.1038/s43018-025-01041-x>, epub 2025-09-02
11. Wilson, P.F.R., Harmanani, M., To, M.N.N., Jamzad, A., Elghareb, T., Guo, Z., Kinnaird, A., Wodlinger, B., Abolmaesumi, P., Mousavi, P.: ProstNFound+: A prospective study using medical foundation models for prostate cancer detection. *International Journal of Computer Assisted Radiology and Surgery* (Feb 2026). <https://doi.org/10.1007/s11548-025-03561-4>
12. Wilson, P.F., To, M.N.N., Jamzad, A., Gilany, M., Harmanani, M., Elghareb, T., Fooladgar, F., Wodlinger, B., Abolmaesumi, P., Mousavi, P.: Prostnfound: integrating foundation models with ultrasound domain knowledge and clinical context for robust prostate cancer detection. In: *International conference on medical image computing and computer-assisted intervention*. pp. 499–509. Springer (2024)