

Learning Where to Look: A Reinforcement Learning Framework for Robust Micro-Ultrasound Prostate Cancer Detection

Mohammad Mahdi Abootorabi^{1,3}, Sina Namazi¹, Armin Saadat¹, Lyuyang Wang¹, Obed Dzikunu^{1,3}, Paul F. R. Wilson^{2,3}, Zhuoxin Guo^{2,3}, Brian Wodlinger⁴, Parvin Mousavi^{2,3}, and Purang Abolmaesumi¹

¹ The University of British Columbia, Vancouver, BC, Canada

{mahdi.abootorabi, purang}@ece.ubc.ca

² Queen's University, Kingston, ON, Canada

³ Vector Institute, Toronto, ON, Canada

⁴ Exact Imaging, Markham, ON, Canada

Abstract. Micro-ultrasound (μ US) is a new, emerging, and promising imaging modality for prostate cancer (PCa) detection, but accurate identification of suspicious tissue remains highly dependent on clinical experience, leading to substantial inter-observer variability. Machine-learning assistance can reduce this variability; however, training reliable deep models is challenging because supervision is sparse and noisy—typically limited to core-level histopathology outcomes (e.g., cancer grade and its percentage in a biopsy core) without pixel-level lesion annotations and under severe class imbalance. We introduce Prost-RL, which re-frames μ US PCa detection as a spatially aware, policy-driven inference problem by learning where to look before decoding. Prost-RL integrates a lightweight reinforcement-learning policy into a foundation-model encoder-decoder to generate interpretable spatial attention maps that act as soft prompts for both cancer-likelihood heatmap prediction and image-level classification. We further propose Adaptive Policy Optimization (APO) to stabilize hybrid supervised-RL training and a noise-robust objective combining symmetric cross-entropy with negative-entropy regularization to mitigate weak-label noise and encourage sharp localization. On a cohort of 6,607 biopsy cores from 693 patients across five clinical sites, Prost-RL achieves 79.0 ± 3.5 AUROC with $64.6 \pm 6.3\%$ sensitivity at 80% specificity for core-level detection (+2.1 AUROC and +4.5 sensitivity points over the strongest baseline), and 79.3 ± 5.8 AUROC for clinically significant cancer classification. The learned policy highlights biopsy-aligned regions, providing transparent, spatially grounded evidence alongside quantitative risk predictions. Code is available at: <https://github.com/DeepRCL/Prost-RL>.

Keywords: Prostate Cancer · Ultrasound · Reinforcement Learning.

1 Introduction

Prostate cancer (PCa) remains a leading cause of cancer mortality in men [2]. While multi-parametric MRI (mpMRI) improves risk stratification [15, 6], its cost and infrastructure requirements limit broad access [11], motivating biopsy-targeting workflows based on ultrasound alone. High-resolution micro-ultrasound (μ US) is an emerging and promising modality for PCa detection, operating at frequencies up to 29 MHz to resolve prostate micro-architecture with MRI-comparable accuracy [4, 13]. However, interpretation of μ US is highly dependent on clinical experience: speckle noise, artifacts, and heterogeneous tumor appearance make suspicious regions subtle and variable [10, 7], leading to substantial inter-observer variability. Automating μ US interpretation with machine learning can help standardize targeting and reduce observer variability.

A central barrier is supervision. In routine biopsy, pathology is available only at the *core level*: each acquired core provides a global cancer label, grade, and involvement percentage, but no pixel-level lesion delineation. Supervision is sparse (restricted to the needle region), noisy (small tumors may be missed), and strongly imbalanced (benign tissue dominates), making it difficult to train deep models that localize cancer and predict clinically significant disease reliably.

Recent work has progressed from patch-based CNNs to foundation-model (FM) adaptations that mitigate weak labels and domain shift via self-supervision, multiple-instance learning (MIL), temporal/radio-frequency modeling, prompting with clinical priors, and knowledge distillation [5, 9, 10, 20–22]. Yet, most approaches remain *passive*: they process the full image holistically and rely on implicit attention to discover lesions. This becomes brittle under weak and noisy labels, where fine-tuning large decoders can amplify spurious regions and degrade localization [8, 18]. These observations motivate an explicit mechanism that *learns where to look* before decoding.

In this paper, we introduce **Prost-RL**, a policy-driven framework that integrates a lightweight reinforcement-learning (RL) policy into an FM encoder-decoder to produce interpretable spatial attention maps serving as soft prompts for heatmap localization and image-level classification. Prost-RL is trained end-to-end with **Adaptive Policy Optimization (APO)**, a hybrid supervised-RL objective inspired by alignment methods, and with noise-robust loss (symmetric cross-entropy and negative-entropy regularization) to stabilize learning from sparse, weak core-level labels. On a multi-center retrospective μ US cohort of 6,607 biopsy cores from 693 patients across five sites, Prost-RL improves core-level detection to 79.0 ± 3.5 AUROC with $64.6 \pm 6.3\%$ sensitivity at 80% specificity (+2.1 AUROC and +4.5 sensitivity points over the strongest baseline) and achieves 79.3 ± 5.8 AUROC for clinically significant cancer classification, while providing biopsy-aligned attention maps for transparent decision support.

2 Materials and Method

In this section, first we discuss the dataset used in 2.1, then we present our proposed method in 2.2, and finally we describe our experiments in 2.3.




Fig. 1. Prost-RL overview. Image and clinical prompt encoders produce features F and embedding c . A spatial policy π_θ generates an attention map α that modulates F into shared embeddings E , decoded by a heatmap head and a csPCa classifier.

2.1 Data

We utilized a private multi-center μ US dataset of 6,607 biopsy cores from 693 patients, collected prospectively across five clinical sites (2013–2016) under an approved trial (ClinicalTrials.gov: NCT02079025 [16]) with informed consent. Biopsies were performed using the ExactVu μ US system (Exact Imaging, Markham, Canada), acquiring B-mode sagittal-plane images (depth 28 mm, width 46.06 mm) at 10–12 cores per patient; the frame immediately preceding needle firing was extracted per core. Clinical metadata included patient age, PSA, PSA density (PSAD), and anatomical core location (POS) as 2D base-apex and medial-lateral coordinates. Core-level histopathology provided ground-truth ISUP Grade Group (GG) and proportional cancer involvement, assigned to annotated needle trace regions as weak spatial supervision. Positive cores ($GG \geq 1$) comprised clinically insignificant (isPCa, $GG=1-2$; $n = 480$) and clinically significant (csPCa, $GG \geq 3$; $n = 400$) cases, yielding 880 positives against 5,727 benign cores (86% benign). Biopsy needle trace and whole-gland binary masks were provided alongside images. Evaluation followed patient-level five-fold cross-validation with center-stratified folds [22]; benign cores from patients with any malignant core were excluded from training to reduce label ambiguity, with no involvement-threshold filtering or benign undersampling applied. Images were resized to 256×256 and normalized to $[0, 1]$; masks to 64×64 to match attention map resolution.

2.2 Method

Fig. 1 provides an overview of *Prost-RL*. Our design is encoder-decoder agnostic; we instantiate it on top of ProstNFound+ [21], currently the strongest baseline for μ US PCa detection, because it couples a pretrained medical foundation

model (MedSAM [14]) with lightweight adapter layers and a clinical prompt encoder. *Prost-RL* contributes three additions to this backbone: (i) an *attention policy network* that learns *where to look* before features reach the decoder, (ii) a noise-robust training objective to handle weak, proportional biopsy labels, and (iii) reinforcement-learning-tuning that optimizes a ranking-aware reward. Training proceeds with a supervised warm-up, followed by RL refinement.

Given an input image $\mathbf{x}$, the MedSAM encoder produces spatial features $\mathbf{F} = \text{Enc}(\mathbf{x}) \in \mathbb{R}^{C \times H \times W}$. Rather than passing these raw features directly for prediction, they are first explicitly modulated by the attention policy. The resulting modulated dense embeddings, together with shared sparse prompt embeddings (encoding clinical metadata), are passed simultaneously to two distinct heads: (i) a heatmap decoder for pixel-level cancer likelihood, and (ii) an image-level classifier specifically for clinically significant prostate cancer (csPCa) detection. Joint training on these shared inputs ensures that the explicitly learned spatial focus actively supports both precise tumor localization and global risk grading.

Spatial Attention Policy and Feature Modulation. Under weak proportional supervision, models tend to produce spatially diffuse predictions. We introduce an explicit attention policy π_θ that learns a spatial distribution over encoder locations. The policy processes $\mathbf{F}$ and clinical metadata $\mathbf{c}$. Clinical features are embedded into a channel-wise gate $\mathbf{g} \in \mathbb{R}^C$, which modulates processed features via element-wise multiplication. A convolutional head produces spatial logits that are masked outside the prostate region and normalized to obtain attention weights α . Attention-weighted features are computed and injected through residual modulation: $\tilde{\mathbf{F}} = \mathbf{F}_{\text{proc}} \odot \alpha$, $\mathbf{E} = \mathbf{F} + \phi(\tilde{\mathbf{F}})$,

where $\mathbf{F}_{\text{proc}}$ denotes the policy-processed features (after convolution and clinical gating), and $\phi(\cdot)$ is a bias-free 1×1 projection with GELU nonlinearity. This additive modulation ensures zero attention yields zero contribution while preserving base features. Modulated embeddings $\mathbf{E}$ is shared by both decoder heads, and α provides an interpretable spatial map of the model’s attention.

Noise-Robust Weakly Supervised Training. Ground-truth is available only as a core-level involvement proportion $q \in [0, 1]$ per needle trace. Pixel-level BCE forces the model to predict cancer everywhere in a positive core, producing diffuse heatmaps. ProstNFound+ [21] relaxes this by applying BCE against q over the needle-prostate intersection, but standard CE fully trusts noisy labels: the gradient $\partial \mathcal{L} / \partial \hat{p} = -q / \hat{p} + (1 - q) / (1 - \hat{p})$ diverges when the model confidently disagrees, causing memorization of erroneous proportions.

We introduce a Symmetric Cross-Entropy (SCE) [19] objective combined with entropy minimization as our heatmap loss. SCE augments CE with a Reverse Cross-Entropy (RCE) term, $\mathcal{L}_{\text{SCE}} = \alpha \mathcal{L}_{\text{CE}}(\hat{p}, q) + \beta \mathcal{L}_{\text{RCE}}(\hat{p}, q)$, where $\mathcal{L}_{\text{RCE}}(\hat{p}, q) = -\hat{p} \log q_\epsilon - (1 - \hat{p}) \log(1 - q_\epsilon)$ and $q_\epsilon = \text{clip}(q, \epsilon, 1 - \epsilon)$. Because the RCE gradient with respect to $\hat{p}$ is $-\log q_\epsilon + \log(1 - q_\epsilon)$, a constant independent of the prediction, the corrective signal remains bounded under strong model-label disagreement, providing implicit noise tolerance without discarding supervision. However, proportion-based losses only constrain the average prediction: a model outputting q uniformly across all pixels satisfies $\mathcal{L}_{\text{SCE}}$ identically to one that lo-

calizes precisely. To resolve this degeneracy, we add a pixel-level entropy regularizer, $\mathcal{L}_{\text{ent}} = \lambda_H |\mathcal{R}|^{-1} \sum_{(i,j) \in \mathcal{R}} \mathcal{H}(\hat{y}_{i,j})$, where $\mathcal{R}$ is the valid spatial region defined by the intersection of the prostate and needle masks, $\hat{y}_{i,j}$ is the predicted pixel probability, and $\mathcal{H}$ is the binary entropy. By minimizing $\mathcal{L}_{\text{ent}}$, we penalize uncertain, mid-range per-pixel predictions and drive the model toward spatially sharp, decisive heatmaps. The total supervised objective is $\mathcal{L}_{\text{sup}} = \mathcal{L}_{\text{SCE}} + \mathcal{L}_{\text{ent}} + \mathcal{L}_{\text{clf}}$, where $\mathcal{L}_{\text{clf}}$ is a balanced CE on the classification head.

Adaptive Policy Optimization. While our noise-robust supervised objective stabilizes learning from weak proportional labels, it inherently optimizes absolute prediction confidence on a strictly per-sample basis. However, clinical diagnosis heavily relies on relative risk stratification and is effectively a ranking problem. Furthermore, supervision is available only at the core level, providing no direct guidance on which spatial regions truly correspond to malignant tissue. Consequently, this supervision can lead to suboptimal spatial focus.

To refine the spatial policy and optimize this ranking capability, we introduce APO as a post-warm-up training stage. The goal is not merely to increase confidence, but to bias the spatial policy toward attention patterns that produce preferable predictions under weak supervision. Rather than relying solely on proxy classification losses, we design a pairwise ranking reward that serves as an in-batch proxy for ranking. For each sample, the reward reflects the fraction of correctly ordered positive-negative pairs, with a multiplicative bonus ($\gamma = 2$) applied to clinically significant (csPCa, $\text{GG} \geq 3$) cases to explicitly prioritize high-risk tumor detection. For a mini-batch with positive set $\mathcal{P}$ and negative set $\mathcal{N}$, we compute a pairwise reward based on the ordering of classification scores:

$$r_i = \begin{cases} \frac{2}{|\mathcal{N}|} \sum_{j \in \mathcal{N}} \mathbf{1}[p_i > p_j] - 1, & y_i = 1, \\ \frac{2}{|\mathcal{P}|} \sum_{j \in \mathcal{P}} \mathbf{1}[p_j > p_i] - 1, & y_i = 0, \end{cases} \quad (1)$$

where p_i is the predicted cancer probability. This reward promotes attention configurations that improve separation between malignant and benign cores. Clinically significant cancers receive amplified reward to prioritize high-grade disease.

To optimize this reward, we adopt a group-relative policy optimization strategy inspired by Group Relative Policy Optimization (GRPO) [17], which computes advantages by comparing multiple rollouts of the same image. However, a challenge arises: our continuous spatial attention is mostly deterministic. Passing the same image through the policy network multiple times would yield identical attention maps, resulting in zero within-image reward variance and collapsing the policy gradients. To break this degeneracy and enable exploration, we inject continuous Gaussian noise directly into the attention logits prior to the softmax activation during the rollout phase: $\tilde{\mathbf{a}} = \mathbf{a} + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma^2)$, where σ regulates the exploration scale. This stochastic injection generates K different spatial configurations per forward pass, allowing the model to observe how subtle shifts in focal attention influence downstream classification rankings.

Finally, applying standard GRPO is suboptimal in our setting due to dataset class imbalance (86% benign) and severe difficulty heterogeneity (e.g., small 10%

vs. large 80% tumor involvements). Standard rollout normalization treats all samples equally, allowing trivial benign cases to overwhelm the gradient signal. We address this by adopting Domain-aware Group Relative Policy Optimization (DRPO) [3].

DRPO applies hierarchical scaling on top of standard GRPO normalization, clustering samples by domain and intra-domain difficulty (via K-Means on the rollout reward vectors). In our setting, we instantiate the domain axis as cancer *vs.* benign, so that hierarchical temperature scaling on the computed advantages dynamically upweights rare cancer cases and hard-to-classify borderline lesions. A differentiable ranking surrogate would only refine the current attention map along its local gradient. RL instead generates K stochastic rollouts per image (exploration) and uses their per-sample reward vectors for the cluster-based scaling, enabling difficulty-aware reweighting that has no per-sample-loss analog. During RL fine-tuning, we jointly optimize the DRPO policy loss alongside the supervised objectives, ensuring the model maintains core detection stability while explicitly nudging the policy toward optimally ranked predictions.

2.3 Experiments

We compare against MedSAM-UNETR [1], MedSAM-FT [14], MicroSegNet [12], ProstNFound [22], and ProstNFound+ [21], and conduct ablation studies on each component. Evaluation follows patient-level five-fold cross-validation with center-stratified splits; core-level scores are computed as mean heatmap activation within the needle–prostate intersection, and image-level csPCa scores from the classification head. We report AUROC and sensitivity at fixed specificities. Supervised training runs for 35 epochs (batch size 8, AdamW, $\text{lr} = 2 \times 10^{-5}$, encoder $\text{lr} = 1 \times 10^{-5}$, cosine annealing, weight decay 10^{-3} , $\alpha = \beta = 1$). RL fine-tuning runs for 35 additional epochs (batch size 16, $K=4$ rollouts, $\sigma=0.15$). Model selection uses validation AUC on high-involvement ($\geq 40\%$) cores.

3 Results and Discussion

Results are summarized in Table 1. Prost-RL consistently outperforms all state-of-the-art baselines across both evaluation heads. For core-level detection, it achieves 79.0 ± 3.5 AUROC, surpassing ProstNFound+ by +2.1%, with sensitivity at 80% specificity improving from 60.1% to 64.6%. Gains extend to high-involvement cores (84.9 *vs.* 83.6) and csPCa detection (79.9 *vs.* 79.6). MedSAM-based approaches without explicit spatial guidance lag considerably behind, underscoring the benefit of learning where to look before decoding. The classification head mirrors this trend: Prost-RL achieves 79.3 ± 5.8 AUROC for csPCa *vs.* non-csPCa (78.5 ± 5.3 for ProstNFound+), with sensitivity at 80% specificity improving from 58.2% to 62.8%, confirming that the shared attention policy consistently benefits both localization and global risk grading.

Fig. 2 part (a) plots mean core-level AUROC as a function of minimum cancer involvement threshold, isolating performance on progressively higher-confidence

Table 1. Prostate cancer detection performance in micro-ultrasound. Results are Mean $\pm$ Standard Deviation across five folds. Statistical significance of Prost-RL over ProstNFound+ was assessed using a two-sided paired t -test across the five folds: $*p < 0.05$, $***p < 0.001$.

Auxiliary Image-Level Classification Head					
Method	All Cores		csPCa <i>vs.</i> non-csPCa		
	AUROC	Sens@80	AUROC	Sens@60	Sens@80
ProstNFound+ [21]	75.4 $\pm$ 4.1	56.3 $\pm$ 7.3	78.5 $\pm$ 5.3	81.8 $\pm$ 8.3	58.2 $\pm$ 10.6
Prost-RL (Ours)	76.6 $\pm$ 3.4	58.0 $\pm$ 4.6	79.3 $\pm$ 5.8	83.1 $\pm$ 8.4	62.8 $\pm$ 12.6
Heatmap Decoder (Core-Level Detection)					
Method	All Cores		High Involvement		csPCa
	AUROC	Sens@80	AUROC	Sens@80	AUROC
MedSAM-UNETR [1]	71.4 $\pm$ 2.7	50.7 $\pm$ 5.0	79.4 $\pm$ 1.6	63.9 $\pm$ 5.7	71.5 $\pm$ 4.6
MedSAM-FT [14]	71.6 $\pm$ 2.0	50.9 $\pm$ 4.2	78.4 $\pm$ 0.9	62.7 $\pm$ 1.9	71.3 $\pm$ 3.8
MicroSegNet [12]	73.5 $\pm$ 2.2	52.7 $\pm$ 3.2	80.0 $\pm$ 0.6	62.6 $\pm$ 1.7	74.3 $\pm$ 5.2
ProstNFound [22]	76.8 $\pm$ 3.8	61.2 $\pm$ 6.9	82.0 $\pm$ 3.2	70.7 $\pm$ 7.8	78.0 $\pm$ 6.5
ProstNFound+ [21]	76.9 $\pm$ 3.5	60.1 $\pm$ 5.6	83.6 $\pm$ 2.4	72.5 $\pm$ 4.9	79.6 $\pm$ 4.1
Prost-RL (Ours)	79.0 $\pm$ 3.5***	64.6 $\pm$ 6.3*	84.9 $\pm$ 2.5	76.3 $\pm$ 5.9	79.9 $\pm$ 5.5

positive cores. All methods improve with higher involvement, consistent with larger tumors being more visually conspicuous. Parts (a) and (b), illustrate the contribution of our two main additions: the supervised stage alone (Prost-RL Supervised, incorporating the attention policy and noise-robust losses) surpasses both baselines, ProstNFound and ProstNFound+, at every threshold, demonstrating that explicit spatial attention modulation and SCE meaningfully improve localization. The full Prost-RL model widens this gap further across all thresholds, with the margin over the supervised variant growing at lower involvement levels, precisely the regime where weak, noisy labels are most problematic, suggesting that APO does not merely boost overall confidence but specifically improves the model’s ability to rank and localize subtle, low-involvement lesions that supervised training alone struggles to distinguish.

Ablation Studies and Qualitative Evaluation. Fig. 2 (part c–e) ablates the key design choices of Prost-RL. **Prompt ablation (c):** even without any clinical prompts, Prost-RL surpasses ProstNFound+ by a clear margin, confirming that the attention policy and noise-robust training provide independent gains. Performance improves monotonically as prompts are added, with the full set yielding the best results, consistent with each feature providing complementary clinical context. **Policy architecture (d):** all tested architectures (transformer and convolutional variants) perform comparably, indicating that the RL training strategy drives the gains rather than a specific network design. **Loss function (e):** our proposed loss outperforms the needle-region cross-entropy of ProstNFound and the MIL proportion BCE of ProstNFound+, demonstrating that




Fig. 2. Top: comparison of Prost-RL (Full) vs. supervised-only variant and other base-lines. Bottom: ablation study results for different setting choices.

explicit noise robustness via SCE and negative-entropy regularization is beneficial in this weakly supervised, label-noisy setting. Fig. 3 compares outputs qualitatively. For malignant cases, Prost-RL produces sharper, more spatially focused heatmaps confined to biopsy-confirmed regions, while ProstNFound+ yields diffuse activations. The policy attention map provides an interpretable spatial prior correlated with the model’s decision, supporting clinician trust. For the benign case, Prost-RL correctly suppresses needle-region activation where ProstNFound+ does not.

Prost-RL’s gains arise from explicit spatial attention guiding decoding under weak supervision, SCE with entropy regularization enforcing sharp heatmaps, and APO refining the policy toward better malignant-benign separation under class imbalance. These components are encoder-decoder agnostic and can apply broadly to weakly supervised tasks with sparse, noisy, proportion-based labels. Since prostatectomy outcomes are not systematically available for multi-center μ US cohorts, core-level labels are intrinsically noisy, and needle biopsy can disagree with radical-prostatectomy grading. This is precisely the regime our noise-robust design targets. Validation on a surgical sub-cohort and extension to 3D volumes are important future directions.

Acknowledgments. This work was supported in part by the Canadian Institutes of Health Research (CIHR), the Natural Sciences and Engineering Research Council of




Fig. 3. Example heatmaps generated by ProstRL and ProstNFound+, shown alongside ground-truth labels and radiologist annotations.

Canada (NSERC), Vector Institute, and through computational resources and services provided by Advanced Research Computing at the University of British Columbia. P. Mousavi is supported in part by Canada CIFAR AI Chair and Canada Research Chair.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Alzate-Grisales, J.A., Mora-Rubio, A., García-García, F., Tabares-Soto, R., De La Iglesia-Vayá, M.: Sam-unetr: Clinically significant prostate cancer segmentation using transfer learning from large model. *IEEE Access* **11**, 118217–118228 (2023)
2. 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* pp. 229–263 (2024)
3. Dai, W., Chen, P., Ekbote, C., Liang, P.P.: Qoq-med: Building multimodal clinical foundation models with domain-aware grpo training. *arXiv preprint arXiv:2506.00711* (2025)
4. DuBois, J., Smani, S., Golos, A., Rivera Lopez, C., Lokeshwar, S.D.: Micro-ultrasound in the detection of clinically significant prostate cancer: A comprehensive review and comparison with multiparametric mri. *Tomography* **11**(7), 80 (2025)
5. 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)
6. Gavade, A.B., Nerli, R., Kanwal, N., Gavade, P.A., Pol, S.S., Rizvi, S.T.H.: Automated diagnosis of prostate cancer using mpMRI images: A deep learning approach for clinical decision support. *Computers* **12**(8), 152 (2023)
 7. Gilany, M., Wilson, P., Jamzad, A., Fooladgar, F., To, M.N.N., Wodlinger, B., Abolmaesumi, P., Mousavi, P.: Towards confident detection of prostate cancer using high resolution micro-ultrasound. In: International conference on medical image computing and computer-assisted intervention. pp. 411–420. Springer (2022)
 8. Han, K., Wang, S., Chen, J., Qian, C., Lyu, C., Ma, S., Qiu, C., Sheng, V.S., Huang, Q., Liu, Z.: Region uncertainty estimation for medical image segmentation with noisy labels. *IEEE Transactions on Medical Imaging* (2025)
 9. 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 cine-loops. In: 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2025)
 10. Harmanani, M., Wilson, P.F., To, M.N.N., Gilany, M., Jamzad, A., Fooladgar, F., Wodlinger, B., Abolmaesumi, P., Mousavi, P.: Trusworthy: toward clinically applicable deep learning for confident detection of prostate cancer in micro-ultrasound. *International Journal of Computer Assisted Radiology and Surgery* **20**(5), 981–989 (2025)
 11. Hutchinson, R., Lotan, Y.: Cost consideration in utilization of multiparametric magnetic resonance imaging in prostate cancer. *Translational andrology and urology* **6**(3), 345 (2017)
 12. 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)
 13. Klotz, L., Lughezzani, G., Maffei, D., Sánchez, A., Pereira, J.G., Staerman, F., Cash, H., Luger, F., Lopez, L., Sanchez-Salas, R., et al.: Comparison of micro-ultrasound and multiparametric magnetic resonance imaging for prostate cancer: a multicenter, prospective analysis. *Canadian Urological Association Journal* **15**(1), E11 (2020)
 14. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**, 654 (2024)
 15. Reisæter, L.A., Fütterer, J.J., Losnegård, A., Nygård, Y., Monssen, J., Gravdal, K., Halvorsen, O.J., Akslen, L.A., Biermann, M., Haukaas, S., et al.: Optimising preoperative risk stratification tools for prostate cancer using mpMRI. *European Radiology* **28**(3), 1016–1026 (2018)
 16. Rohrbach, D., Wodlinger, B., Wen, J., Mamou, J., Feleppa, E.: High-frequency quantitative ultrasound for imaging prostate cancer using a novel micro-ultrasound scanner. *Ultrasound in medicine & biology* **44**(7), 1341–1354 (2018)
 17. Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang, H., Zhang, M., Li, Y., Wu, Y., et al.: Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300> **2**(3), 5 (2024)
 18. Wang, H., Huai, L., Li, W., Qi, L., Jiang, X., Shi, Y.: Weakmedsam: Weakly-supervised medical image segmentation via sam with sub-class exploration and prompt affinity mining. *IEEE Transactions on Medical Imaging* (2025)

19. Wang, Y., Ma, X., Chen, Z., Luo, Y., Yi, J., Bailey, J.: Symmetric cross entropy for robust learning with noisy labels. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 322–330 (2019)
20. Willis, E., Elghareb, T., Wilson, P.F.R., To, M.N.N., Abootorabi, M.M., Jamzad, A., Wodlinger, B., Mousavi, P., Abolmaesumi, P.: Guide-us: Grade-informed unpaired distillation of encoder knowledge from histopathology to micro-ultrasound (2026), <https://arxiv.org/abs/2602.19005>
21. Wilson, P.F., 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* pp. 1–10 (2026)
22. 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)