Data-Driven Optimization of Prostate Sectorization from MRI

Florencia Boccarato[1,*][0009−0002−6200−4536], Fahym Bounazou[1,3,4][0009−0003−1229−8449], Hye Lim Lee[1][0009−0009−3629−9626], Raphaele Renard Penna[2,4][0000−0001−8484−4590], and Hervé Delingette[1][0000−0001−6050−5949]

[1] Université Côte d'Azur, Inria, Epione Team, France
[2] Department of Radiology, Hôpital Pitié-Salpêtrière, Assistance Publique des Hôpitaux de Paris, France
[3] Unité de Recherche Clinique, Hôpital Pitié-Salpêtrière, Assistance Publique des Hôpitaux de Paris, France
[4] Sorbonne Université, GRC n5, Onco-Urologie Prédictive, France
[*] Corresponding Author: florencia.boccarato@inria.fr

Abstract. The PI-RADS sectorization scheme serves as the standard clinical language for reporting prostate lesion locations and guiding biopsies. However, the mapping of those discrete anatomical sectors into MR image coordinates has not been standardized and currently relies on rigid geometric scripts that are non-differentiable and often anatomically incomplete. In this paper, we propose a differentiable geometric framework that learns a global consensus sectorization template in a canonical coordinate space. By parameterizing sector boundaries as learnable primitives, our model distills an anatomical prior from expert annotations via gradient descent. During inference, this global template is projected onto patient-specific zonal masks, ensuring the sectors adapt to individual morphology while strictly enforcing PI-RADS topological constraints. Evaluation on internal (N=84) and external (ProstateX, N=50) cohorts demonstrates a significant improvement in Anterior-Posterior division over state-of-the-art baselines and provides a complete 24-sector mapping by incorporating previously omitted Lateral-Central partitions. By providing a bi-directional differentiable bridge between 3D MRI clinical nomenclature, this framework enables the spatial localization of sector-based biopsy results within the 3D volume. This unlocks massive clinical registries for weak-supervision tasks and establishes a foundational tool for standardized, spatially-indexed data curation. The source code and data are publicly available at https://gitlab.inria.fr/fboccara/data-driven-prostate-sectorization

Keywords: Prostate Sectorization · Differentiable Geometry

1 Introduction

Accurate localization of prostate tumors on MRI is essential to the diagnostic pipeline. In clinical practice, this localization is reported via sector-based

schemes that discretize the 3D prostate volume into predefined anatomical regions, referred to as *sectors*. While this facilitates communication between radiologists and urologists, they create a fundamental data-science bottleneck: clinical registries often contain millions of biopsy-derived ground truths (e.g., Gleason Scores) that are spatially indexed only by discrete sector labels. Because there is no native 3D coordinate for a biopsy core, this "weakly-labeled" data remains spatially locked and difficult to exploit for training high-fidelity AI models for cancer localization or grading. Recent work has shown that sector-based spatial reconstructions can partially address this limitation by converting sector annotations into volumetric supervision signals for weakly supervised learning [4, Chapter 5]. However, the broader challenge of defining a clinically consistent and anatomically faithful sectorization remains unresolved.

Studies have shown that even with structured guidelines, inter-observer variability remains high [3]. While previous work has demonstrated that automated partitioning can improve reporting consistency [8], these methods focus primarily on clinical assistance rather than the generation of 3D datasets. Furthermore, the evolution of the Prostate Imaging–Reporting and Data System (PI-RADS)—from the 16 regions of v1 [2] to the 41 sectors of v2.1 [9]—has created several de facto standards. Existing computational approaches generally fall into two categories: rigid rule-based scripts [5], which rely on hard-coded geometric assumptions and lack the flexibility to adapt to atypical prostate shapes; or flexible partitioning strategies [8], which prioritize fitting the anatomical boundaries at the expense of abandoning PI-RADS constraints. Because the former cannot be optimized and the latter lacks clinical consistency, neither provides the differentiable, image-to-sector mapping required to translate noisy, clinical records into high-fidelity image masks training data.

In this work, we propose the first data-driven, differentiable framework for prostate sectorization, bridging the gap between clinical standards and anatomical adaptability. We reformulate prostate sectorization as a constrained optimization problem, where a global anatomical template is learned in a canonical coordinate space from the knowledge of prostate lesion masks and their associated sectors. By parameterizing sector boundaries as differentiable geometric primitives, specifically planes and zonal masks, our model distills a latent anatomical consensus from noisy expert labels. We rely on a specialized objective function that maximizes tumor coverage within assigned sectors while explicitly modeling the spatial uncertainty inherent in human labeling and tumor "spillover". This ensures that the learned boundaries remain strictly compliant with PI-RADS topology while being optimized via gradient descent to match the "real-world" behavior of expert radiologists.

Our framework provides an interpretable and mathematically grounded model for mapping prostate sectors into 3D Regions of Interest (ROI). This is used to quantify the adequacy between lesion masks and assigned sectors and thus can serve as a quality control (QC) tool. Our contributions are as follows:

– **Differentiable Anatomical Mapping:** We introduce a 24-sector interpretable sectorization model of the prostate that provides a learnable map-

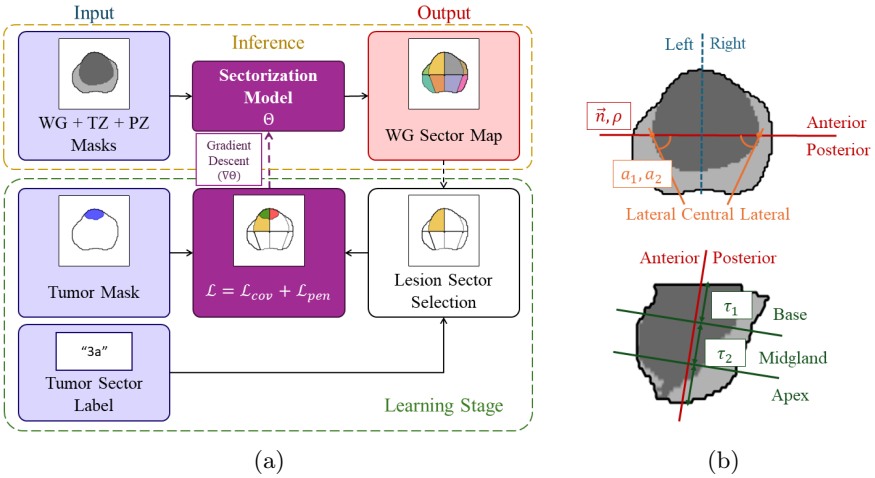

Fig. 1. (a) Learning and inference pipelines of the sectorization framework. (b) Geometric parameterization of prostate sector on axial (Top) and sagittal (Bottom) views.

ping between 3D MRI voxels and discrete prostate sectors "unlocking" sector-indexed biopsy data for high-resolution machine learning.

- **Latent Consensus Loss:** We propose a novel objective function specifically designed to optimize geometric primitives (planes and zonal masks) from weak, and noisy labels by maximizing tumor coverage while penalizing cross-sector spillover.
- **Quantitative Evaluation of Lesion Sectorization:** We demonstrate that our framework properly associates lesion masks with prostate sectors on both internal and external (ProstateX) datasets, showcasing better performances than prior work to separate anterior from posterior lesions.

2 Methodology

We propose a differentiable framework to learn a global anatomical consensus for prostate sectorization. Our model partitions the Transition Zone (TZ) and Peripheral Zone (PZ) into a 24-sector map, a clinically relevant subset of PI-RADS v1's 27-sectors proposed map, using a set of learnable geometric primitives. The parameter optimization visible in Fig. 1a is based on the comparison between the predicted sector ROI and actual prostate lesion masks.

2.1 Preprocessing and Canonical Mapping

For each patient, we assume the availability of a Whole Gland (WG) mask and its subdivision into TZ and PZ. Following clinical convention, the Central Zone is merged with the PZ and the Anterior Fibromuscular Stroma (AFS) with

the TZ. While various zonal segmentation backbones can be employed [1,6], we utilize the automated method by Hamzaoui et al. [5] on 3D T2-weighted MRI volumes. To standardize the input space, all volumes are resampled to a voxel spacing of $[1, 0.5, 0.5]$ mm and cropped to a $[96, 192, 192]$ grid centered on the prostate. To ensure parameter invariance across varying gland morphologies, each prostate's bounding box is then mapped to a canonical coordinate system $\mathcal{C} \in [-1, 1]^3$. All subsequent geometric operations are performed within $\mathcal{C}$, ensuring that our learned parameters Θ represent a generalized anatomical prior rather than patient-specific fits.

During the training phase, we also use lesion masks M_j, $j \in \{1 \dots J\}$ and their annotated main sector location following PI-RADS v1 notation. There can be several lesions associated with the same prostate mask but having different associated prostate sectors. Moreover, a lesion often encompasses several prostate sectors; however, only the main sector is considered here.

2.2 Differentiable Geometric Parameterization

We define the sectorization via 11 learnable parameters (Θ) governing five anatomical partition rules, as shown in Fig. 1b.

- **Anterior/Posterior (AP) Division:** Defined by a plane with normal $\mathbf{n} \in \mathbb{R}^3$ and offset $\rho \in \mathbb{R}$. To ensure a smooth optimization manifold, we utilize a 6D continuous rotation representation [10] for $\mathbf{n}$, resulting in seven learnable parameters.
- **Base/Midgland/Apex (BMA) Division:** We define a longitudinal centerline by projecting the superior-inferior axis onto the AP plane. The prostate and lesion volumes are projected onto this 1D axis, where two learnable scalar thresholds, τ_1 and τ_2, define the regional boundaries.
- **Posterior Lateral/Central (LC) Division:** Within the posterior region, lateral and central sectors are defined slice-wise by rays originating from the intersection of the TZ boundary and the AP plane. While this origin point is slice-specific, the angular direction of the rays is governed by a global learnable parameter ϕ. We parameterize the left ray using a 2D $(\cos\phi, \sin\phi)$ representation and derive the right ray by left-right symmetry. Therefore, this division results in two learnable parameters.
- **Left/Right and Zonal Divisions:** To preserve PI-RADS topology without increasing model complexity, we incorporate a fixed sagittal symmetry plane and the patient-specific TZ/PZ interface directly from the input masks.

To enable gradient-based optimization, we replace hard geometric thresholding with a sigmoid-based soft relaxation $\sigma(x)$. We define the probabilistic sector map $S_k(\mathbf{p}, \Theta)$ as the product of the binary zonal mask $Z_k(\mathbf{p})$ and a set of n_k differentiable indicator functions $f_{i,k}(\mathbf{p}, \Theta)$: $S_k(\mathbf{p}, \Theta) = Z_k(\mathbf{p}) \cdot \prod_{i=1}^{n_k} f_{i,k}(\mathbf{p}, \Theta)$.

The functions $f_{i,k}(\mathbf{p}, \Theta)$ are defined based on the boundary type. For planar divisions (AP and BMA): $f_{i,k}(\mathbf{p}) = \sigma\left(\alpha \cdot (\mathbf{n}_i \cdot \mathbf{p} + \rho_i)\right)$. For the LC division, applied to the posterior region: $f_{i,k}(\mathbf{p}) = \sigma(\alpha \cdot (\psi(\mathbf{p}) - \phi))$, where $\psi(\mathbf{p})$ is the axial angle of voxel $\mathbf{p}$ relative to the slice-specific origin.

The temperature parameter α controls the sharpness of the sector boundary thereby modulating the magnitude and locality of the gradients during optimization. Owing to the canonical coordinate system $\mathcal{C}$, this parameterization remains consistent across glands of varying size, enabling a unified representation for both small and large prostates. During inference, the final discrete sectorization is recovered via an *argmax* operation over $S_k(\mathbf{p}, \bar{\Theta})$.

2.3 Latent Consensus Objective Function

Standard voxel-wise losses (e.g., Dice or cross-entropy) are ill-suited for this task, as lesions are sub-regions of the sectors and do not define its total extent. Furthermore, since a lesion may physically span multiple anatomical regions, the loss must adjudicate boundaries based on the aggregate clinical assignments across the dataset. To address this particular setup, we design a custom objective function based on spatial containment.

We compare the soft sector masks $S_k(\mathbf{p}, \Theta)$ with the lesion masks $M_k(\mathbf{p})$ available for a given patient, where M_k denotes the union of all lesions clinically assigned to sector k. For a patient volume, let $\mathcal{K}$ be the set of sectors containing at least one lesion. We define the loss $L(\Theta)$ via two terms:

1. **Spatial Coverage (Sensitivity):** Maximizes the inclusion of annotated lesion voxels within their designated sector:

$$\mathcal{L}_{cov}(\Theta) = -\frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \frac{\sum_{\mathbf{p} \in \Omega} S_k(\mathbf{p}, \Theta) \cdot M_k(\mathbf{p})}{\sum_{\mathbf{p} \in \Omega} M_k(\mathbf{p}) + \epsilon}$$

2. **Cross-Sector Penalty (Specificity):** Penalizes the "spillover" of the sector S_k into the territory of lesions assigned to all other sectors $j \neq k$. This term prevents the sectors from over-expanding into neighboring anatomical regions, discouraging trivial degenerate solutions where one sector covers most of the gland:

$$\mathcal{L}_{pen}(\Theta) = \frac{1}{|\mathcal{K}|} \sum_{k=1}^{24} \frac{\sum_{\mathbf{p} \in \Omega} S_k(\mathbf{p}, \Theta) \cdot \sum_{j \neq k} M_j(\mathbf{p})}{\sum_{j \neq k} \sum_{\mathbf{p} \in \Omega} M_j(\mathbf{p}) + \epsilon}$$

The total objective $L = \mathcal{L}_{cov} + \lambda_{fp} \mathcal{L}_{pen}$ allows the framework to learn the optimal latent boundaries by minimizing labeling conflict across the entire population.

2.4 Optimization

Parameters Θ are initialized to estimated PI-RADS v1 geometric defaults. Optimization is performed globally using the Adam optimizer ($lr = 10^{-3}$) across a 5-fold cross-validation scheme. This global approach ensures that the learned $\bar{\Theta}$ represents a stable anatomical consensus across the population. We employ a learning-rate scheduler and early stopping based on validation loss over 100

epochs, with fixed hyperparameters $\alpha = 10$ and $\lambda_{fp} = 1$. The framework was implemented in PyTorch, and the final parameters $\bar{\Theta}$ correspond to the model achieving the highest aggregate coverage on the validation set.

3 Results

We evaluate our framework across three distinct datasets to assess robustness against temporal shifts and geographic variability. Our private cohort consists of 236 retrospective clinical prostate MRI studies collected at AP-HP centers, with institutional ethics committee approval (CESREES in 2020, file no. 2945786, and CNIL authorization in 2022, Decision DR-2022-044). All imaging data were fully anonymized prior to analysis. From this cohort, we define: (i) **Internal-Holdout** ($N = 37$); and (ii) **Internal-Temporal** ($N = 47$), collected two years prior to investigate institutional temporal shifts. For geographic generalizability, we additionally evaluate on the publicly available ProstateX dataset ($N = 50$) [7], following its research usage guidelines. Please note that the prostate masks (WG,TZ,PZ) and lesions masks of the ProstateX dataset have been created by different algorithms or raters than the internal dataset.

Since ProstateX lacks lesion sector labels, we established a "silver standard" via a multi-reader study. Two researchers independently assigned sector labels based on T2W MRI and lesion segmentations, with discrepancies resolved by a senior radiologist (>10 years experience). While we acknowledge that the annotators were not blinded to the study's geometric framework, the masks provide a necessary benchmark for external validation. Annotations are also publicly available at the corresponding repository.

3.1 Quantitative Analysis

Performance is measured by lesion coverage, i.e. the fraction of the annotated lesion volume contained within its clinically assigned sector. We report coverage per prostate division rule by re-aggregating lesions into simplified 2 (AP,PZ LC) or 3 (BMA) sub-regions. On the other hand, **Global** utilizes the full 24-sector map, a stringent metric where any lesion entirely displaced into a direct neighboring sector results in 0% coverage. As lesions are expected to span multiple sectors, 100% coverage is neither expected nor theoretically optimal for any single sector mask. We compare our learned model against two baselines: (1) **PI-Init**, our architecture initialized with standard PI-RADS v1 defaults without optimization; and (2) **Rule-Based**, the geometric method by Hamzaoui et al. [5].

Our model demonstrates high stability across all cohorts (Table 1). Notably, on the External-ProstateX set, our framework achieved a Global coverage of 0.56, outperforming PI-Init (0.46) and matching the rule-based baseline, despite differences in external annotation protocols. Statistical analysis (Fig. 2) confirms that for the Anterior/Posterior (AP) division, our model significantly outperforms both baselines (paired Wilcoxon $p<0.01$). While the BMA rule shows no

Table 1. Average lesion coverage. Bold indicates best average performance.

	Internal-Holdout			**Internal-Temporal**			**External-ProstateX**		
Rule	Ours	PI-Init	Rule-Based	Ours	PI-Init	Rule-Based	Ours	PI-Init	Rule-Based
AP	**0.95**	0.94	0.91	**0.93**	0.90	0.90	**0.93**	0.89	0.89
BMA	**0.69**	0.68	0.68	0.73	**0.74**	0.73	0.81	0.80	0.81
PZ LC	**0.67**	0.52	0.63	**0.73**	0.69	0.60	**0.67**	0.49	0.64
Global	0.42	0.37	0.42	**0.56**	0.53	0.49	0.56	0.46	0.56

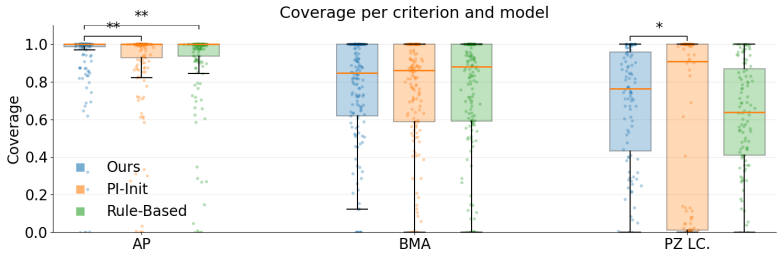

Fig. 2. Coverage distributions for each rule on the test set across our model, PI-Init and Rule-Based. Brackets indicate significant pairwise differences using a paired Wilcoxon test ($** : p < 0.01$, $* : p < 0.05$)

significant difference in mean coverage, our model produces significantly fewer outliers. For the PZ LC division, our optimization significantly improves upon PI-Init ($p<0.05$).

Failure Recovery Analysis. We evaluated "completely misaligned" cases (0% coverage), which primarily occurred in the PZ LC and AP divisions. Our model successfully recovered 20 out of 22 PZ LC failures exhibited by PI-Init. Even though Rule-Based already shows high robustness in the PZ LC division, our model achieved comparable failure rates (2 total vs. 2 total across all sets) while reducing the AP failure cases from 4 to 3. This suggests that our learned global parameters achieve the precision of per-case geometric methods while maintaining a differentiable, continuous representation.

3.2 Qualitative Assessment of Anatomical Fidelity

Quantitative metrics in sectorization are often constrained by the inherent subjectivity of clinical labels. Fig. 3a provides a visual benchmark of anatomical fidelity. While the method proposed by Hamzaoui et al. [5] frequently produces rigid, linear boundaries that bisect suspicious lesions (e.g., Lesion 4p and 6p), our differentiable model generates boundaries that respect the natural curvature of the prostate zones. Walter-Rittel et al. [8], on the other hand, proposes a sectorization that doesn't respect PI-RADS topology on the PZ.

Fig. 3b illustrates the sagittal alignment across different algorithms. In the approach by Hamzaoui et al. [5], the AP division is coupled to the TZ mask. While this reflects a common clinical heuristic, as posterior TZ lesions are rare,

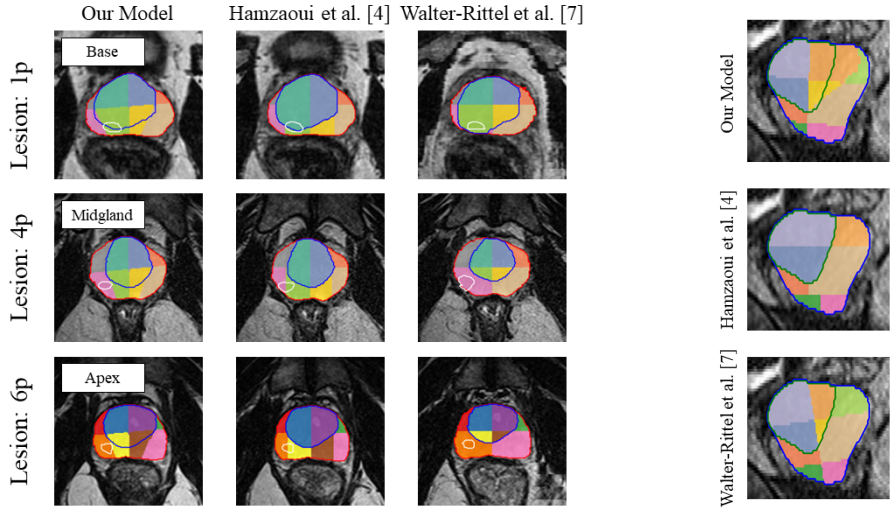

(a) (b)

Fig. 3. Visual comparison against prior methods [5,8]. (a) Axial slices for three subjects showcasing their lesion segmentations (white outlines). All methods misclassify the base as midgland. The axial slices are taken orthogonal to each model's centerline. (b) Sagittal view.

it deviates from the formal PI-RADS guidelines. Furthermore, outside of the TZ, their method employs a region-wise vertical AP division that introduces discontinuities across the longitudinal axis. In contrast, while methods like Walter-Rittel et al. [8] offer high precision via per-patient centerline regression, they are computationally intensive and highly sensitive to Whole Gland (WG) segmentation noise. Our framework provides a robust intermediate representation: a continuous, shared AP plane that preserves longitudinal consistency across the population. While current results are stable, future iterations will incorporate learnable parameters to account for the individual tilt of the prostate gland, further refining the alignment at the apex and base.

4 Conclusion and Future Work

In this work, we presented a novel data-driven framework for prostate MRI sectorization that bridges the gap between 3D imaging and standardized clinical nomenclature. By introducing a differentiable parameterization of PI-RADS boundaries, we move beyond the limitations of rigid geometric scripts. Our model provides a continuous, anatomically loyal representation of the prostate sectors that remains strictly consistent with clinical guidelines while adapting to population-wide anatomical distributions.

Several avenues for future work remain. First, extending our training to multi-center datasets would allow the model to better quantify and adapt to institutional variability in reporting criteria. Second, while the current implementation utilizes T2W-weighted images for geometric consistency, the framework is natively compatible with multi-parametric inputs. Adding supplementary modalities would better reflect the radiologist's holistic diagnostic process, potentially mitigating discrepancies caused by inter-modality registration. Finally, we aim to incorporate learnable parameters to compensate for prostate tilt and eventually map all cases to a standardized population atlas.

Acknowledgments. This research was made possible thanks to Bpifrance (AICOO project), and by the French government through the National Research Agency (ANR), under the IA Cluster projects (ANR-23-IACL-0001).

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

References

1. Adams, L.C., Makowski, M.R., Engel, G., Rattunde, M., Busch, F., Asbach, P., Niehues, S.M., Vinayahalingam, S., van Ginneken, B., Litjens, G., et al.: Prostate158-an expert-annotated 3t mri dataset and algorithm for prostate cancer detection. Computers in biology and medicine **148**, 105817 (2022)
2. Dickinson, L., Ahmed, H.U., Allen, C., Barentsz, J.O., Carey, B., Futterer, J.J., Heijmink, S.W., Hoskin, P.J., Kirkham, A., Padhani, A.R., et al.: Magnetic resonance imaging for the detection, localisation, and characterisation of prostate cancer: recommendations from a european consensus meeting. European urology **59**(4), 477–494 (2011)
3. Greer, M.D., Shih, J.H., Barrett, T., Bednarova, S., Kabakus, I., Law, Y.M., Shebel, H., Merino, M.J., Wood, B.J., Pinto, P.A., et al.: All over the map: An interobserver agreement study of tumor location based on the pi-radsv2 sector map. Journal of Magnetic Resonance Imaging **48**(2), 482–490 (2018)
4. Hamzaoui, D.: AI-based diagnosis of prostate cancer from multiparametric MRI. Theses, Université Côte d'Azur (Jun 2023). https://doi.org/10.70675/8312130fzd494z4cb5z9ef4zf0b398d0145c
5. Hamzaoui, D., Montagne, S., Renard-Penna, R., Ayache, N., Delingette, H.: Automatic zonal segmentation of the prostate from 2d and 3d t2-weighted mri and evaluation for clinical use. Journal of Medical Imaging **9**(2), 024001–024001 (2022)
6. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. Nature methods **18**(2), 203–211 (2021)
7. Litjens, G., Debats, O., Barentsz, J., Karssemeijer, N., Huisman, H.: Prostatex challenge data. The cancer imaging archive **10**, K9TCIA (2017)
8. Walter-Rittel, T.C., Frisch, A., Hamm, C.A., Baumgärtner, G.L., Hartenstein, A., Dräger, F., Haas, M., Cash, H., Hofbauer, S., Hamm, B., et al.: Automated 24-sector grid-map algorithm for prostate mpmri improves precision and efficacy of prostate lesion location reporting. European Journal of Radiology **183**, 111897 (2025)

9. Weinreb, J., Barentsz, J., Choyke, P., Cornud, F., Haider, M., Macura, K., Margolis, D., Schnall, M., Shtern, F., Tempany, C., et al.: Pi-rads prostate imaging-reporting and data system, v2. 1. American College of Radiology (2019)
10. Zhou, Y., Barnes, C., Lu, J., Yang, J., Li, H.: On the continuity of rotation representations in neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5745–5753 (2019)