

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/355585385>

Anatomical and Diagnostic Bayesian Segmentation in Prostate MRI — Should Different Clinical Objectives Mandate Different Loss Functions?

Conference Paper · October 2021

CITATIONS

0

READS

13

5 authors, including:



Anindo Saha

Radboud University Medical Centre (Radboudumc)

8 PUBLICATIONS 16 CITATIONS

SEE PROFILE



Matin Hosseinzadeh

Sharif University of Technology

16 PUBLICATIONS 143 CITATIONS

SEE PROFILE

Anatomical and Diagnostic Bayesian Segmentation in Prostate MRI —Should Different Clinical Objectives Mandate Different Loss Functions?

Anindo Saha¹, Joeran Bosma¹, Jasper Linmans^{1,2}, Matin Hosseinzadeh¹, Henkjan Huisman¹

¹Diagnostic Image Analysis Group, Radboud University Medical Center, The Netherlands

²Computational Pathology Group, Radboud University Medical Center, The Netherlands

{anindya.shaha,joeran.bosma,jasper.linmans,
matin.hosseinzadeh,henkjan.huisman}@radboudumc.nl

Abstract

We hypothesize that probabilistic voxel-level classification of anatomy and malignancy in prostate MRI, although typically posed as near-identical segmentation tasks via U-Nets, require different loss functions for optimal performance due to inherent differences in their clinical objectives. We investigate distribution, region and boundary-based loss functions for both tasks across 200 patient exams from the publicly-available ProstateX dataset. For evaluation, we conduct a thorough comparative analysis of model predictions and calibration, measured with respect to multi-class volume segmentation of the prostate anatomy (whole-gland, transitional zone, peripheral zone), as well as, patient-level diagnosis and lesion-level detection of clinically significant prostate cancer. Notably, we find that distribution-based loss functions (in particular, focal loss) are well-suited for diagnostic or panoptic segmentation tasks such as lesion detection, primarily due to their implicit property of inducing better calibration. Meanwhile, (with the exception of focal loss) both distribution and region/boundary-based loss functions perform equally well for anatomical or semantic segmentation tasks, such as quantification of organ shape, size and boundaries.

1 Introduction

Anatomical versus Diagnostic Anatomical (organ-level) and diagnostic (pathology-level) segmentation are fundamentally different, with respect to their clinical objectives. For instance, within the clinical workflow of prostate MRI, segmentation of the prostate anatomy {whole-gland (WG), transitional zone (TZ), peripheral zone (PZ)} is used to estimate its volume, boundaries and geometry—enabling the calculation of prostate-specific antigen (PSA) density, guiding treatment planning and future interventions [1–3]. Meanwhile, segmentation of clinically significant prostate cancer (csPCa) is primarily used to detect the number of malignant lesions present in the prostate gland (if any), and characterize each instance for diagnosis, risk stratification and/or targeted biopsies (similar to PI-RADS guidelines [4]). In other words, while the former anatomical task leans towards *semantic* single-object quantification, the latter diagnostic task can be framed as *panoptic* multi-object detection [5, 6]. Despite these differences, in medical image computing, these objectives are often treated and trained as near-identical segmentation tasks using U-Nets, and evaluated in a similar manner using the Dice Similarity Coefficient (DSC) [7–10]. Recent studies confirm that while DSC can be appropriate for single-object quantification, its inability to measure multi-object detection makes it unsuitable for lesion localization [5, 11–13]. We hypothesize that such considerations should not only be made during evaluation, but should also be incorporated at train-time by using task-specific loss function(s).

Segmentation Loss Functions We can distinguish between loss functions that minimize mismatch in underlying distributions and those that minimize mismatch in segmentation regions/boundaries [14]. Distribution-based loss functions, such as the widely adopted cross-entropy (CE) loss, minimize dissimilarity between the target distribution (ground-truth) and its approximation (model predictions), over the course of training. Balanced cross-entropy (BCE) loss is an extension of CE, introducing class weights α to address imbalanced class frequencies in the dataset. Focal loss (FL) represents a further extension of BCE, introducing an additional hyperparameter γ to differentiate between easy and difficult examples [15]. When $\gamma = 0$, FL reduces down to CE or BCE depending on the value of α . On the other hand, region or boundary-based loss functions are derived from integrals over segmentation regions (e.g. soft Dice loss [16]) or distance metrics in the space of contours (e.g. boundary loss [17]). Unlike distribution-based losses, which can benefit from class re-weighting, the sum of Dice and boundary loss (DB) is implicitly impervious to class imbalance, and in turn, recommended for highly imbalanced datasets [17].

2 Experiments and Analysis

Materials We used 200 prostate bpMRI (T2W, high b-value DWI, ADC) exams from the publicly-available ProstateX dataset [18], paired with voxel-level delineations of WG, TZ, PZ and csPCa [19]. All images were resampled to $0.5 \times 0.5 \times 3.0 \text{ mm}^3$ resolution, center-cropped to $160 \times 160 \times 20$ voxels and intensity-normalized (T2W, DWI: z-score; ADC: linear) [20], prior to usage.

Bayesian Segmentation Model We used a probabilistic adaptation [21] of the deep attentive 3D U-Net, developed and validated specifically for prostate bpMRI in our previous work [20]. Monte-Carlo dropout nodes were added to capture both *epistemic* and *aleatoric* uncertainty during inference (as recommended by Hu et al. [22]). Cosine annealing learning rate [23] (decaying from 10^{-4} to 10^{-7}) and *AMSGrad* optimizer [24] were used to train the model. Data augmentations comprised of additive Gaussian noise (standard deviation 0-0.5), horizontal flip, rotation ($\pm 7.5^\circ$), translation (0-15% horizontal and/or vertical shifts) and scaling (0-15%) centered along the axial plane. Fig. 1 illustrates the train-time schematic of the model, and its source code is made publicly-available at https://github.com/DIAGNijmegen/prostateMR_3D-CAD-csPCa.

Experiments We investigated four different loss policies (CE, BCE, FL, DB) as the segmentation loss function (L_S) in our model for anatomical ($\{WG, TZ, PZ\}$ in T2W MRI) and diagnostic (csPCa in bpMRI) segmentation. For the former task, α is set as $[0.05, 0.30, 0.65]$ as per the empirical class distribution of the dataset, and for the latter task, α is set as $[0.75, 0.25]$ as per the findings of previous studies [20, 25]. In both cases, γ is set to its default value of 2.00 as used in [15, 20, 25]. All metrics were computed in 3D over 3 independent runs \times 5-fold cross-validation \times mean of 100 executions of probabilistic inference per image. Identical data splits were maintained for all configurations across both tasks.

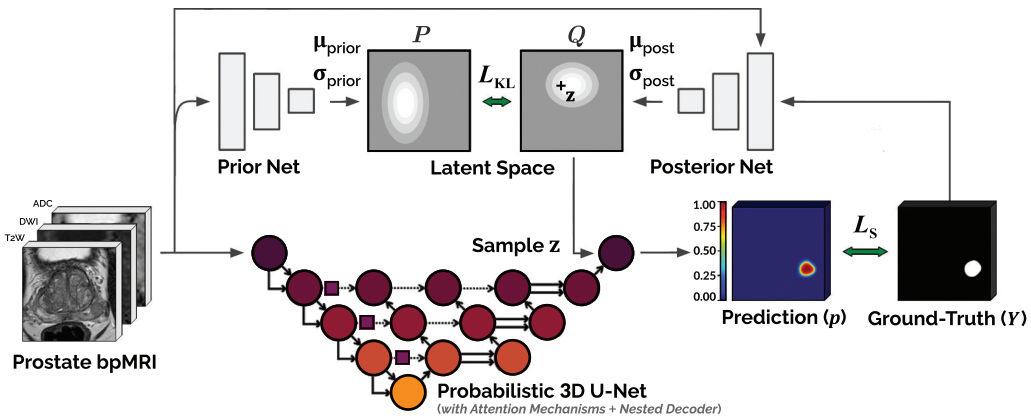


Figure 1: Train-time schematic for the Bayesian segmentation model. L_{KL} denotes the Kullback-Leibler divergence loss between prior distribution P and posterior distribution Q . L_S denotes the segmentation loss between prediction p and ground-truth Y . For each execution of the model, one sample $z \in Q$ (train-time) or $z \in P$ (test-time) is drawn to predict one segmentation mask p [21].

3 Results and Discussion

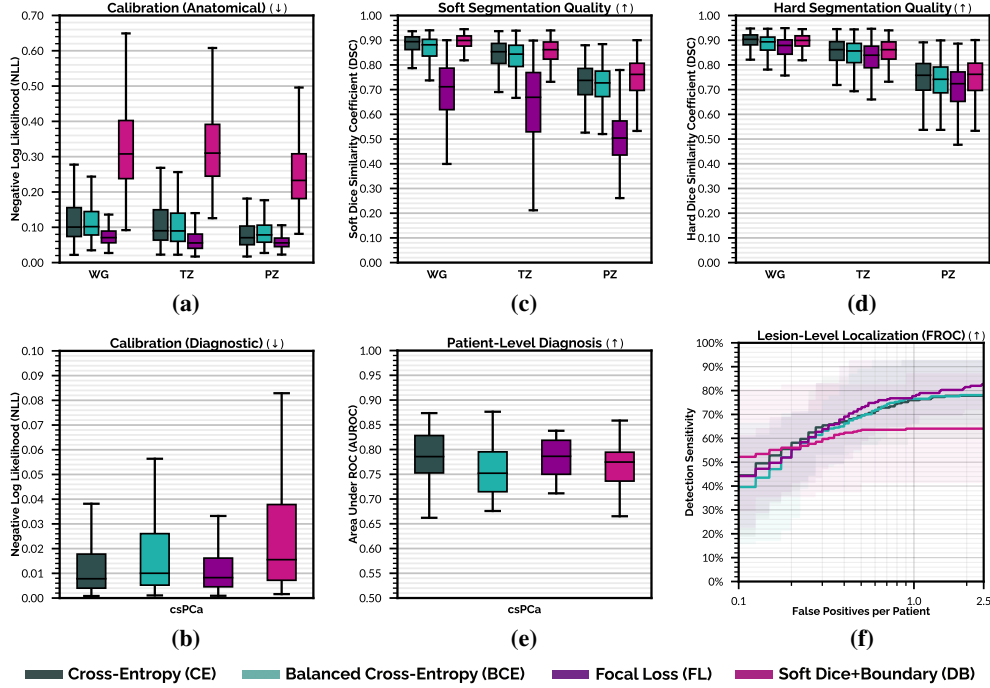


Figure 2: Anatomical (**top-row**) and diagnostic (**bottom-row**) segmentation of {WG, TZ, PZ} in T2W MRI and csPca in bpMRI, respectively. For each L_S configuration, we evaluated (**a, b**) model calibration with negative log-likelihood (NLL), (**c, d**) segmentation quality with mean soft/hard Dice Similarity Coefficient (DSC) across all 600 observations (3 runs \times 5 folds \times 40 validation samples), (**e**) patient-level diagnosis with Area Under Receiver Operating Characteristic (AUROC), (**f**) lesion-level localization with Free-Response Receiver Operating Characteristic (FROC). Transparent areas in (**f**) FROC indicate 95% confidence intervals. Arrows (\uparrow or \downarrow) across (**a-f**) indicate whether the corresponding metric should be maximized/minimized for ideal performance.

Calibration We find that distribution-based loss functions (CE, BCE, FL) produce implicitly well calibrated predictions, in comparison to DB [26] (as shown in Fig. 2 (a, b) and *Appendix A*). Among them, FL exhibits notably better calibration. Our findings concur with that of Mukhoti et al. [27], where the authors demonstrate that FL is formulated as such, that it minimizes dissimilarity between ground-truth and predicted distributions, while *increasing entropy* of the latter –thereby limiting NLL overfitting and inducing both calibration and regularization. Such a property can also lead to more useful uncertainty estimates [27, 28].

Anatomical Segmentation From Fig. 2 (c), it is clear that while CE, BCE and DB achieve similar segmentation quality with respect to raw *softmax* predictions ({WG: 81-89; TZ: 82-85; PZ: 71-74} mean DSC), FL performs worse ({WG: 70; TZ: 64; PZ: 50} mean DSC). When considering binarized *argmax(softmax)* predictions, as seen in Fig. 2 (d), overall FL performance is only marginally lower than that of CE, BCE and DB ({WG: 88-90; TZ: 84-86; PZ: 72-76} mean DSC). We attribute this decline in FL performance to its high predictive uncertainty near class boundaries (as seen in *Appendix A*), which limits its ability to produce precise contour definitions.

Diagnostic Segmentation From Fig. 2 (e, f), we observe that while all loss functions are similar in patient-level diagnosis (76-78 mean AUROC), results indicate that distribution-based loss functions, especially FL, perform substantially better than region/boundary-based losses (DB) at lesion detection (with an average increase of 15.7% in maximum detection sensitivity and 0.25 in partial area under FROC between 0.1-2.5 false positives per patient). We attribute this improvement to the implicit train-time calibration induced by FL, which can facilitate better risk stratification and higher detection sensitivities across difficult diagnostic tasks (e.g. via retaining low-confidence lesion predictions, as opposed to muting them out –refer to *Appendix A*). We believe that the growing success of csPca detection models trained using FL derivatives (such as the *FocalNet* [25], among others

[20, 29, 30]) can, in part, be attributed to this phenomenon rather than architectural enhancements or class weighting alone (note BCE performs near-identical to CE, if not worse, across both tasks). On the other hand, while DB loss is a natural fit for maximizing DSC, it is ill-posed to optimize panoptic segmentation objectives [5, 11]. Subsequently, its miscalibration (refer to Fig. 2 (a, b)) or polarized predictions (refer to *Appendix A*), translate to a relatively flat FROC curve with lower maximum detection sensitivity (refer to Fig. 2 (f)) –presumably detection of low-confidence, difficult or small lesions are skipped in favour of maximizing confidence and overlap of clear lesions [11].

In conclusion, we recommend distribution-based loss functions (in particular, FL) for diagnostic or panoptic segmentation tasks. For anatomical or semantic segmentation tasks, we observe that although distribution-based losses improve calibration, (with the exception of FL) they are equivalent to region/boundary-based losses in terms of standard/hard DSC. Hence, CE, BCE, DB or presumably their composites (as applied by the *nnU-Net* [31] and investigated by J. Ma et al. [14]), can be recommended. Further analyses are required, using larger datasets and multiple medical imaging modalities, to draw out definitive conclusions.

Broader Impact

Prostate cancer is one of the most prevalent cancers in men worldwide [32]. In the absence of experienced radiologists, its multifocality, morphological heterogeneity and strong resemblance to numerous non-malignant conditions in MR imaging, can lead to low inter-reader agreement (< 50%) and sub-optimal interpretation [33–35]. The development of automated, reliable detection algorithms has therefore become an important research focus in medical image computing —offering the potential to support radiologists with consistent quantitative analysis, improve diagnostic accuracy, and in turn, minimize unnecessary biopsies in patients [36, 37].

To the best of our knowledge, this study has no foreseeable negative societal effects.

Acknowledgments and Disclosure of Funding

This research is supported, in parts, by the European Union H2020: ProCancer-I project (EU grant 952159), Health~Holland (LSHM20103), ContextVision AB (Linköping, Sweden) and Siemens Healthineers (CID: C00225450).

References

- [1] M. Eklund, F. Jäderling, A. Discacciati, and M. Bergman et al. MRI-Targeted or Standard Biopsy in Prostate Cancer Screening. *New England Journal of Medicine*, 2021. doi: 10.1056/NEJMoa2100852.
- [2] U. G. Falagario, I. Jambor, A. Lantz, and O. Ettala et al. Combined Use of Prostate-specific Antigen Density and Magnetic Resonance Imaging for Prostate Biopsy Decision Planning: A Retrospective Multi-institutional Study Using the Prostate Magnetic Resonance Imaging Outcome Database (PROMOD). *European Urology Oncology*, 2020. ISSN 2588-9311. doi: 10.1016/j.euo.2020.08.014.
- [3] M. Valerio and Y. Cerantola and S. E. Eggener and H. Lepor et al. New and Established Technology in Focal Ablation of the Prostate: A Systematic Review. *European Urology*, 71(1):17–34, 2017. ISSN 0302-2838. doi: 10.1016/j.eururo.2016.08.044.
- [4] J.C. Weinreb, J.O. Barentsz, P.L. Choyke, and F. Cornud. PI-RADS Prostate Imaging – Reporting and Data System: 2015, Version 2. *European Urology*, 69(1):16 – 40, 2016.
- [5] A. Kirillov, K. He, R. Girshick, and C. Rother et al. Panoptic Segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [6] X. Yu, B. Lou, D. Zhang, D. Winkel, and L. Joskowicz. Deep Attentive Panoptic Model for Prostate Cancer Detection Using Biparametric MRI Scans. In *Medical Image Computing and Computer Assisted Intervention (MICCAI 2020)*, pages 594–604, 2020. ISBN 978-3-030-59719-1.
- [7] Y. Artan and I. S. Yetik. Prostate Cancer Localization Using Multiparametric MRI based on Semisupervised Techniques With Automated Seed Initialization. *IEEE Transactions on Information Technology in Biomedicine*, 16(6):1313–1323, 2012. doi: 10.1109/TITB.2012.2201731.
- [8] A. G. Chung, F. Khalvati, M. J. Shafiee, M. A. Haider, and A. Wong. Prostate Cancer Detection via a Quantitative Radiomics-Driven Conditional Random Field Framework. *IEEE Access*, 3:2531–2541, 2015. doi: 10.1109/ACCESS.2015.2502220.
- [9] S. Kohl, D. Bonekamp, H. P. Schlemmer, and K. Maier-Hein et al. Adversarial Networks for the Detection of Aggressive Prostate Cancer, 2017.
- [10] P. F. Jäger. Challenges and Opportunities of End-to-End Learning in Medical Image Classification, 2020.
- [11] A. Reinke and M. D. Tizabi and M. Eisenmann and L. Maier-Hein. Common Pitfalls and Recommendations for Grand Challenges in Medical Artificial Intelligence. *European Urology Focus*, 7(4):710–712, 2021. ISSN 2405-4569. doi: 10.1016/j.euf.2021.05.008.

- [12] A. Carass, S. Roy, A. Gherman, and I. Oguz et al. Evaluating White Matter Lesion Segmentations with Refined Sørensen-Dice Analysis. *Scientific Reports*, 10(1):8242, 2020. ISSN 2045-2322.
- [13] V. Yeghiazaryan and I. D. Voiculescu. Family of Boundary Overlap Metrics for the Evaluation of Medical Image Segmentation. *Journal of Medical Imaging*, 5(1):1 – 19, 2018. doi: 10.1117/1.JMI.5.1.015006.
- [14] J. Ma and J. Chen and M. Ng and A. L. Martel et al. Loss Odyssey in Medical Image Segmentation. *Medical Image Analysis*, 71:102035, 2021. ISSN 1361-8415. doi: 10.1016/j.media.2021.102035.
- [15] T. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal Loss for Dense Object Detection. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2999–3007, 2017.
- [16] F. Milletari, N. Navab, S. Ahmadi. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. In *4th International Conference on 3D Vision (3DV)*, 2016.
- [17] H. Kervadec and J. Bouchtiba and C. Desrosiers and E. Granger et al. Boundary Loss for Highly Unbalanced Segmentation. *Medical Image Analysis*, 67:101851, 2021. ISSN 1361-8415.
- [18] S. G. Armato and H. Huisman and K. Drukker and L. Hadjiiski et al. PROSTATEx Challenges for Computerized Classification of Prostate Lesions from Multiparametric Magnetic Resonance Images. *Journal of Medical Imaging*, 5(4):1 – 9, 2018. doi: 10.1117/1.JMI.5.4.044501.
- [19] R. Cuocolo, A. Stanzione, A. Castaldo, and D. R. De Lucia et al. Quality Control and Whole-Gland, Zonal and Lesion Annotations for the PROSTATEx Challenge Public Dataset. *European Journal of Radiology*, 138:109647, 2021. doi: <https://doi.org/10.1016/j.ejrad.2021.109647>.
- [20] A. Saha, M. Hosseinzadeh, and H. Huisman. End-to-end Prostate Cancer Detection in bpMRI via 3D CNNs: Effects of Attention Mechanisms, Clinical Priori and Decoupled False Positive Reduction. *Medical Image Analysis*, 73:102155, 2021. ISSN 1361-8415. doi: 10.1016/j.media.2021.102155.
- [21] S. A. A. Kohl, B. Romera-Paredes, C. Meyer, and O. Ronneberger et al. A Probabilistic U-Net for Segmentation of Ambiguous Images. In *Medical Imaging Meets NeurIPS Workshop – 32nd Conference on Neural Information Processing Systems (NeurIPS)*, 2018.
- [22] S. Hu, D. Worrall, S. Knecht, and M. Welling et al. Supervised uncertainty quantification for segmentation with multiple annotations. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, pages 137–145, 2019. ISBN 978-3-030-32245-8.
- [23] I. Loshchilov and F. Hutter. SGDR: Stochastic Gradient Descent with Warm Restarts. In *International Conference on Learning Representations (ICLR)*, 2016.
- [24] S. J. Reddi, S. Kale, and S. Kumar. On the Convergence of Adam and Beyond. In *International Conference on Learning Representations (ICLR)*, 2018.
- [25] R. Cao, A. Mohammadian Bajgiran, S. Afshari Mirak, S. Shakeri, X. Zhong, D. Enzmann, S. Raman, and K. Sung. Joint Prostate Cancer Detection and Gleason Score Prediction in mp-MRI via FocalNet. *IEEE Transactions on Medical Imaging*, 38(11):2496–2506, 2019.
- [26] A. Mehrtash, W. M. Wells, C. M. Tempny, and T. Kapur et al. Confidence calibration and predictive uncertainty estimation for deep medical image segmentation. *IEEE Transactions on Medical Imaging*, 39(12):3868–3878, 2020. doi: 10.1109/TMI.2020.3006437.
- [27] J. Mukhoti, V. Kulharia, A. Sanyal, and P. K. Dokania et al. Calibrating deep neural networks using focal loss. In *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*, 2020.
- [28] V. Kuleshov, N. Fenner, and S. Ermon. Accurate Uncertainties for Deep Learning Using Calibrated Regression. In *35th International Conference on Machine Learning (ICML 2018)*, 2018.
- [29] R. Cao, X. Zhong, S. Shakeri, and K. Sung et al. Prostate Cancer Detection and Segmentation in Multiparametric MRI via CNN and Conditional Random Field. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 1900–1904. IEEE, 2019.
- [30] Y. Qian, Z. Zhang, and B. Wang. ProCDet: A New Method for Prostate Cancer Detection based on MR Images. *IEEE Access*, 2021. doi: 10.1109/ACCESS.2021.3114733.
- [31] F. Isensee, S. A. A. Kohl, J. Petersen, P. F. Jaeger, and K. H. Maier-Hein. nnU-Net: A Self-Configuring Method for Deep Learning-based Biomedical Image Segmentation. *Nature Methods*, 18(2):203–211, Feb 2021. ISSN 1548-7105. doi: 10.1038/s41592-020-01008-z.
- [32] K.D. Miller, L. Nogueira, A.B. Mariotto, and R.L. et al. Siegel. Cancer Treatment and Survivorship Statistics, 2019. *CA: A Cancer Journal for Clinicians*, 69(5):363–385, 2019.
- [33] B. Israël, M. van der Leest, M. Sedelaar, A.R. Padhani, P. Zámecnik, and J.O. Barentsz. Multiparametric Magnetic Resonance Imaging for the Detection of Clinically Significant Prostate Cancer: What Urologists Need to Know. Part 2: Interpretation. *European Urology*, 77(4):469–480, 2020.
- [34] C.P. Smith, S.A. Harmon, T. Barrett, and L.K. Bittencourt. Intra- and Interreader Reproducibility of PI-RADS v2: A Multireader Study. *Journal of Magnetic Resonance Imaging*, 49(6):1694–1703, 2019.
- [35] A.B. Rosenkrantz, L.A. Ginocchio, D. Cornfeld, and A.T. Froemming. Interobserver Reproducibility of the PI-RADS Version 2 Lexicon. *Radiology*, 280(3):793–804, 2016.
- [36] M.M.C. Elwenspoek, A.L. Sheppard, M.D.F. McInnes, and P. Whiting. Comparison of Multiparametric Magnetic Resonance Imaging and Targeted Biopsy With Systematic Biopsy Alone for the Diagnosis of Prostate Cancer: A Systematic Review and Meta-Analysis. *JAMA Network Open*, 2(8):e198427, 2019.
- [37] P. Schelb, J.P. Kohl, S. and Radtke, and D. Bonekamp. Classification of Cancer at Prostate MRI: Deep Learning versus Clinical PI-RADS Assessment. *Radiology*, 293(3):607–617, 2019.

A Appendix: Model Predictions

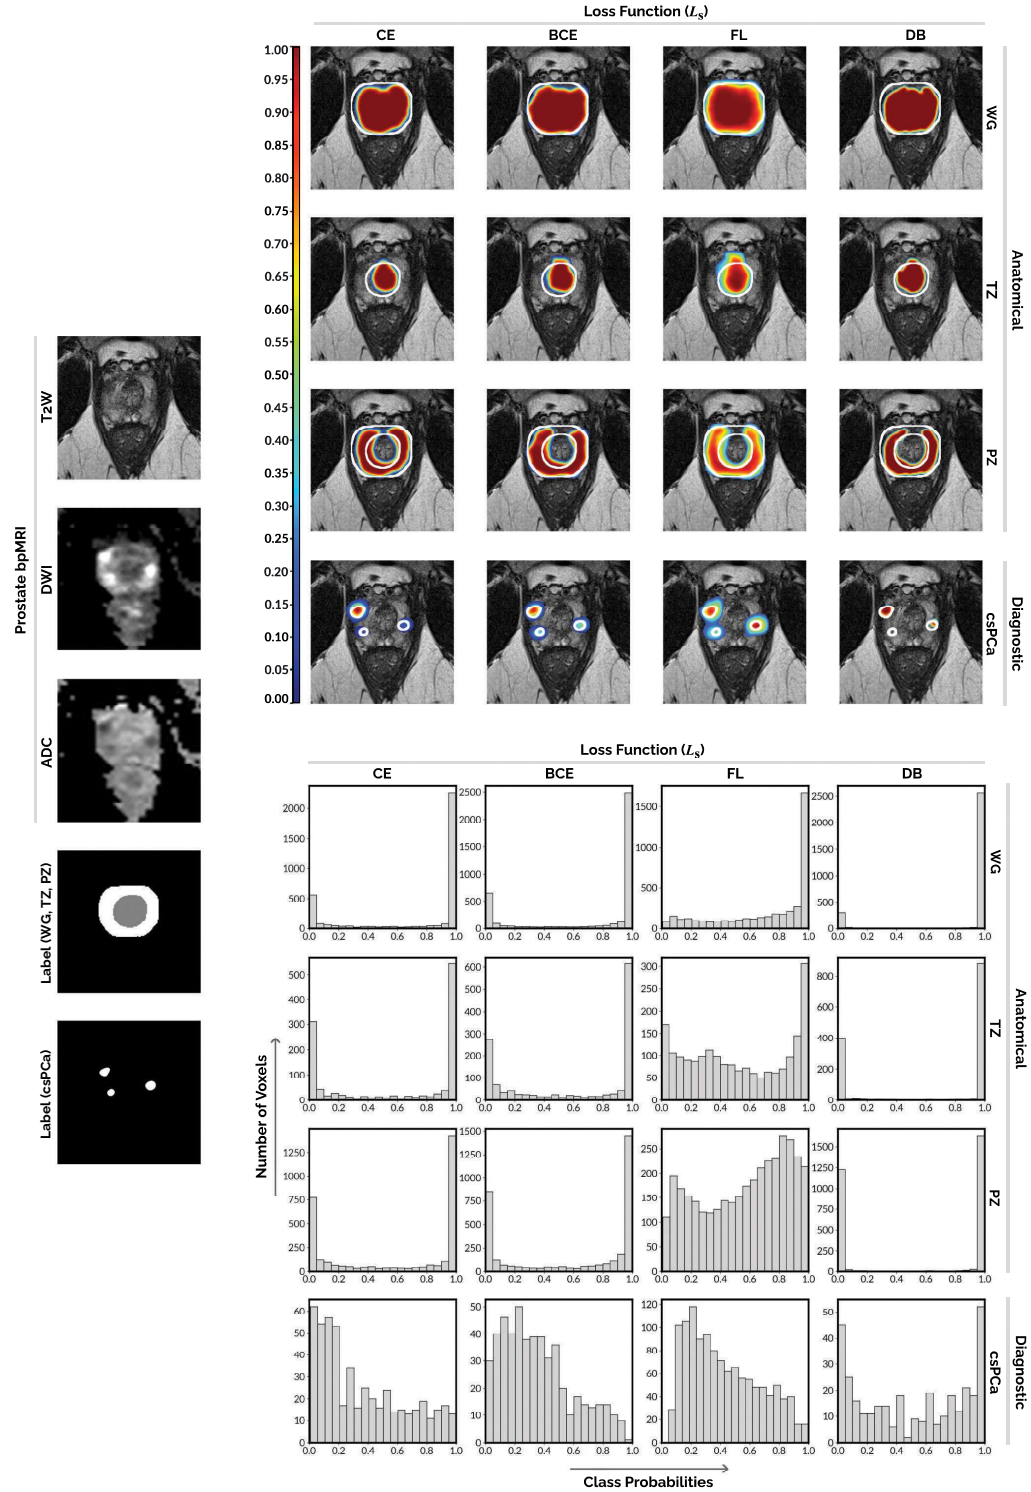


Figure 3: Anatomical (WG, TZ, PZ) and diagnostic (csPca) segmentations for a single validation scan, across all four loss functions (CE, BCE, FL, DB), are shown above. Predictions are overlaid on the T2W scan as reference, where white lines indicate the corresponding ground-truth annotation for the given task. For each case, a histogram of class probabilities over the predicted segment(s) has been illustrated below, indicating the degree of implicit calibration induced by the loss function.