

E-BayesSAM: Efficient Bayesian Adaptation of SAM with Self-Optimizing KAN-Based Interpretation for Uncertainty-Aware Ultrasonic Segmentation

Bin Huang^{1,3}, Zhong Liu¹, Huiying Wen², Bingsheng Huang^{3(✉)}, Xin Chen^{1(✉)}, and Shuo Li⁴

¹ the National-Regional Key Technology Engineering Laboratory for Medical Ultrasound, Guangdong Key Laboratory of Biomedical Measurements and Ultrasound Imaging, School of Biomedical Engineering, Shenzhen University Medical School, Shenzhen University, Shenzhen, China
chenxin@szu.edu.cn

² the Institute of Maternal and Child Medicine, Shenzhen Maternity and Child Healthcare Hospital, Southern Medical University, Shenzhen, China

³ Medical AI Lab, School of Biomedical Engineering, Shenzhen University Medical School, Shenzhen University, Shenzhen 518000, China
huangb@szu.edu.cn

⁴ the Department of Biomedical Engineering and the Department of Computer and Data Science, Case Western Reserve University, Cleveland, OH, USA

Abstract. Although the Segment Anything Model (SAM) has advanced medical image segmentation, its Bayesian adaptation for uncertainty-aware segmentation remains hindered by three key issues: (1) instability in Bayesian fine-tuning of large pre-trained SAMs; (2) high computation cost due to SAM’s massive parameters; (3) SAM’s black-box design limits interpretability. To overcome these, we propose E-BayesSAM, an efficient framework combining Token-wise Variational Bayesian Inference (T-VBI) for efficiently Bayesian adaptation and Self-Optimizing Kolmogorov-Arnold Network (SO-KAN) for improving interpretability. T-VBI innovatively reinterprets SAM’s output tokens as dynamic probabilistic weights and reparameterizes them as latent variables without auxiliary training, enabling training-free VBI for uncertainty estimation. SO-KAN improves token prediction with learnable spline activations via self-supervised learning, providing insight to prune redundant tokens to boost efficiency and accuracy. Experiments on five ultrasound datasets demonstrated that E-BayesSAM achieves: (i) real-time inference (0.03s/image), (ii) superior segmentation accuracy (average DSC: Pruned E-BayesSAM’s 89.0% vs. E-BayesSAM’s 88.0% vs. MedSAM’s 88.3%), and (iii) identification of four critical tokens governing SAM’s decisions. By unifying efficiency, reliability, and interpretability, E-BayesSAM bridges SAM’s versatility with clinical needs, advancing deployment in safety-critical medical applications. The source code is available at GitHub.

Keywords: Medical Imaging · Segmentation · Uncertainty · Segment Anything Model.

1 Introduction

Developing a Bayesian adaptation of the Segment Anything Model (SAM) for uncertainty estimation faces critical computational and architectural challenges: its massive pre-trained parameters destabilize variational training dynamics, while its deterministic, black-box design inherently obscures uncertainty modeling—essential for clinical reliability (Fig. 1A). Although SAM and its variants [10,14,2,24,8,28,11,15] have demonstrated outstanding performance, their lack of probabilistic uncertainty estimates limits their implementation in medical imaging, where ambiguous boundaries (e.g., tumor margins) and noisy inputs are commonplace, especially in ultrasound imaging [16]. This gap restricts deployment in safety-critical tasks where reliable uncertainty estimation is essential to avoid misdiagnosis or overtreatment.

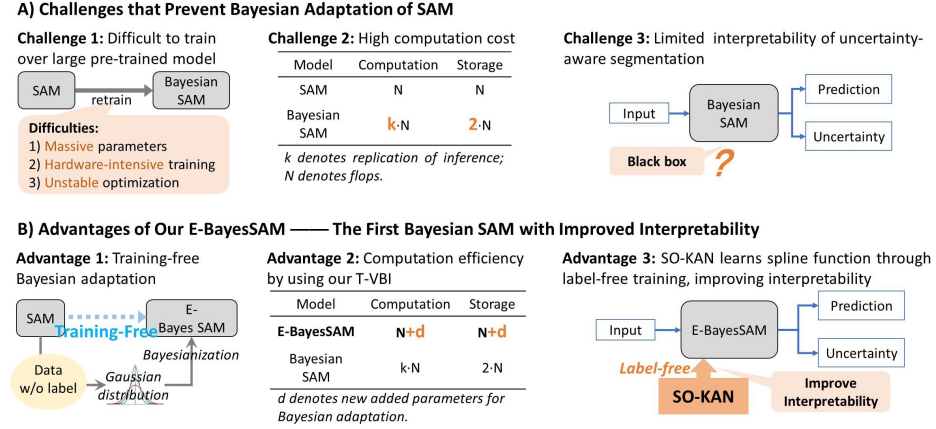


Fig. 1. E-BayesSAM represents the first Bayesian adaptation of the Segment Anything Model (SAM) for uncertainty-aware segmentation, achieved through two proposed modules: Token-wise Variational Bayesian Inference (T-VBI) for efficient Bayesian adaptation and Self-Optimizing Kolmogorov-Arnold Network (SO-KAN) for improving interpretability.

Current methods [27,3,25,29] often estimate the uncertainty of SAM using test-time augmentation, which incurs high computational costs and remains a black-box approach. Meanwhile, conventional Bayesian methods [6,9,22] face practical challenges: full-network Bayesian adaptation may disrupt SAM’s pre-trained knowledge, while reliance on labeled data for distribution modeling limits scalability in large-scale models [5,17].

To overcome the aforementioned challenges, we leverage the strength of Variational Bayesian Inference (VBI) and Kolmogorov-Arnold Network (KAN) [12]. However, they cannot be directly applied to SAM, as both require supervised retraining—a process that risks disrupting SAM’s zero-shot segmentation abil-

ity due to unstable optimization. To address this, we propose Token-wise VBI (T-VBI) and Self-Optimizing KAN (SO-KAN) to develop an Efficient Bayesian SAM (E-BayesSAM) with three key advantages (Fig. 1B): 1) We innovatively reformulate SAM’s output tokens as dynamic Bayesian weights by reparameterizing them as Gaussian distributions: the original tokens serve as means (μ), while standard deviations (σ) are inferred from unlabeled dataset token statistics, enabling training-free Bayesian adaptation (Advantage 1 in Fig. 1B). 2) Built upon Advantage 1, the proposed T-VBI introduces only minimal auxiliary parameters to SAM (Advantage 2 in Fig. 1B) while avoiding the computational overhead of full-network Bayesian methods. 3) We improve the multilayer perceptron (MLP) for token prediction with SO-KAN, using its transparent spline activations to explicitly interpret tokens and prune redundant tokens (Advantage 3 in Fig. 1B). Critically, spline parameters are trained via self-supervised learning on unlabeled medical images, eliminating the dependency on manual annotations.

To the best of our knowledge, E-BayesSAM is the first Bayesian SAM framework, achieving efficient uncertainty-aware segmentation and improved interpretability. Our contributions are threefold: 1) We creatively reformulate SAM’s output tokens as dynamic Gaussian-distributed weights, introducing fewer auxiliary parameters and accelerating VBI compared to conventional methods. 2) Our designed T-VBI preserves SAM’s zero-shot capability while bypassing unstable retraining and annotation cost for uncertainty-aware segmentation. 3) SO-KAN learns spline-based activations trained via self-supervised learning to provide interpretable activation, which guides pruning of redundant tokens for enhanced segmentation accuracy.

2 Methods

Our E-BayesSAM (Fig. 2) achieves Bayesian adaptation by reformulating SAM’s output tokens as dynamic Bayesian weights, forming the theoretical foundation of our framework. This framework integrates two innovations—T-VBI and SO-KAN—to leverage the strengths of KAN and VBI into SAM, enabling efficient uncertainty-aware segmentation and interent interpretability without requiring manual annotations.

2.1 Theoretical Foundation: Output Tokens are Dynamic Weights

We creatively reinterpret the output tokens of SAM as the dynamic weights, forming the theoretical foundation of the E-BayesSAM. In conventional SAM, the segmentation mask is obtained via the point-wise product of output tokens $\mathbf{T}$ and enhanced image features $\mathbf{F}$, where $\mathbf{T}$ encodes segmentation-critical attributes (e.g., anatomical boundaries, texture patterns). Let $\mathcal{F}_{\text{SAM}}$ be the function class of SAM’s decoder with output tokens $T \in \mathbb{R}^{k \times d}$, and $\mathcal{F}_{\text{NN}}$ the class of conventional neural networks with static weights $W \in \mathbb{R}^{k \times d}$. For any $f \in \mathcal{F}_{\text{NN}}$,




Fig. 2. E-ByteSAM employs two innovations—Token-wise Variational Bayesian Inference (T-VBI) and Self-Optimizing Kolmogorov-Arnold Network (SO-KAN)—to leverage the strengths of KAN and VBI into SAM, enabling efficient uncertainty-aware and interpretable segmentation.

there exists $f_{\text{SAM}} \in \mathcal{F}_{\text{SAM}}$ such that:

$$\sup_{x \in \mathcal{X}} \|f_{\text{SAM}}(x) - f(x)\| < \epsilon \quad (1)$$

where x denotes the input of model, and $\epsilon > 0$ depends on SAM’s token generation network capacity.

SAM’s tokens $T = g(x; \theta)$ are produced by a neural network g . By the universal approximation theorem, g can approximate any continuous mapping $x \rightarrow W$ with error $\|g(x; \theta) - W\|_F < \epsilon_1$.

SAM computes the predicted masks via $M = \phi(TQ^\top)$, where $Q = \psi(F_{\text{image}})$. For a conventional network $M_{\text{NN}} = \phi(WQ^\top)$, the error is bounded as:

$$\|M - M_{\text{NN}}\|_F \leq L_\phi \|TQ^\top - WQ^\top\| \leq L_\phi \|T - W\|_F \|Q\|_F \quad (2)$$

Since $\|Q\|_F$ is bounded and $\|T - W\|_F < \epsilon_1$, the total error satisfies $\|M - M_{\text{NN}}\|_F < \epsilon$ with $\epsilon = L_\phi \cdot \epsilon_1 \cdot \|Q\|_F$.

In conclusion, SAM’s tokens T emulate static weights W by dynamically adapting to inputs, preserving approximation power under the assumptions of the universal approximation theorem.

2.2 Token-wise VBI (T-VBI) for Efficient Uncertainty Estimation

Our T-VBI reparameterizes SAM’s output tokens $\mathbf{T}$ as dynamic Gaussian distributions to enable Monte Carlo-based uncertainty estimation while minimizing computational overhead. To bypass the computational inefficiency associated with Bayesian adaptation of the SAM decoder, we reformulate the output tokens $\mathbf{T}$ as dynamic probabilistic weights—the minimal subset requiring Gaussian modeling for uncertainty-aware segmentation.

The framework operates in two phases: 1) **Standard Deviation Statistics**: Compute token-wise standard deviations σ from SAM’s output token distribution across unlabeled training data. 2) **Reparameterized Sampling**: For new inputs x , derive means μ from predicted output tokens $\mathbf{T}$ while employing σ pre-computed in Phase 1. The two phases are formulized as:

$$\mu = \mathbf{T} = g(x; \theta), \quad \sigma = \sqrt{\frac{1}{N} \sum_{i=1}^N (t_{ij} - \mu_j)^2}, \quad (3)$$

where $\mu_j = \frac{1}{N} \sum_{i=1}^N t_{ij}$ (t_{ij} : j -th token dimension).

The latent segmentation variables (stochastic output tokens $\mathbf{T}'$) are sampled via the reparameterization trick:

$$\mathbf{T}' = \mu + \sigma \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, 1), \quad (4)$$

where $\odot$ denotes element-wise multiplication, and $\mathcal{N}(0, 1)$ denotes the standard normal distribution. The segmentation mask $\hat{y}$ and uncertainty map u are decoded as:

$$\hat{y} = \frac{1}{K} \sum_{i=1}^K \text{Sigmoid}(\mathbf{T}'_i \mathbf{Q}^\top), \quad u = \sqrt{\frac{1}{K} \sum_{i=1}^K (y_i - \mu_t)^2}, \quad (5)$$

where $\text{Sigmoid}(\cdot)$ is the sigmoid activation function, K is the number of variational inference samples, and y_i is the i -th prediction, and $\mu_t = \frac{1}{K} \sum_{i=1}^K y_i$. By confining Bayesian computation to the token space, T-VBI reduces memory overhead by 99% compared to full-network variational inference, aligning with Fig. 1B.

2.3 Self-Optimizing KAN (SO-KAN) for Interpretability and Pruning

SO-KAN learns spline parameters via self-supervised learning while preserving SAM’s zero-shot segmentation capability, providing insights to prune tokens for lightweight, uncertainty-aware segmentation. To retain SAM’s pre-trained knowledge, we initialize KAN’s base weight with SAM’s MLP weights, freeze these weight, and train only residual spline via self-supervised learning.

To interpret and prune the tokens, spline activations are visualized to compute the positive ratio—critical since SAM’s sigmoid output implies predictions rely on positive activation values. Tokens with high positive ratio (e.g., Top-4) are identified as high-contribution tokens governing boundary/texture attention for accurate segmentation; the remaining tokens are pruned to derive Pruned E-BayesSAM.

2.4 Model Development

We initialize E-BayesSAM using MedSAM [14] through three core steps: (1) computing output token statistics from the SAM-MED2D [2] training set; (2) adapting SO-KAN via self-supervised learning; (3) pruning E-BayesSAM using the spline activations of KAN. For token variance estimation, per-dimension standard deviations are calculated as Equation (3). For SO-KAN training, two losses are designed: uncertainty-weighted loss $\mathcal{L}_{\text{UnW}}$ to utilize the pseudo labels and a feature-consistency loss $\mathcal{L}_{\text{feat}}$ to maintain alignment with SAM’s original features. The composite loss $\mathcal{L} = \mathcal{L}_{\text{UnW}} + \mathcal{L}_{\text{feat}}$ is defined as:

$$\mathcal{L}_{\text{UnW}} = \frac{1}{N} \sum_{i=1}^N \frac{1 - \mu_i}{\epsilon + \mu_i} \left\| \hat{y}_i - y_i^{\text{pseudo}} \right\|_2^2, \quad \mathcal{L}_{\text{feat}} = \frac{1}{N} \sum_{i=1}^N \left\| \mathbf{T}_i^{\text{original}} - \mathbf{T}_i^{\text{KAN}} \right\|^2, \quad (6)$$

where $\mu_i = \mathbb{E}[\mathbf{z}_i]$, $\epsilon = 1 \times 10^{-6}$, and y_i^{pseudo} denotes conventional SAM’s prediction. Data augmentation enhance data diversity via scaling (0.75x–1.25x), rotation ($\pm 45^\circ$), gamma (0.5–1.5), brightness (± 0.1), and Gaussian noise ($\sigma \in [0.01, 0.2]$).

Implementation Details We implemented our framework in PyTorch on an Nvidia RTX 3070 GPU (8GB VRAM) with a batch size of 1, initial learning rate of 0.0001, and maximum iteration to 1,000,000, using AdamW optimizer [13].

3 Experimental Results and Discussion

3.1 Datasets and Evaluation Metrics

We evaluate E-BayesSAM on five publicly available ultrasound datasets: two thyroid nodule datasets (DDTI [18] and TN3K [7]) and three breast nodule datasets (UDIAT [26], BUSI [1], and OASBUD [19]). Segmentation accuracy is quantified using the Dice Similarity Coefficient (DSC) [4].

3.2 Compared Methods

We benchmark E-BayesSAM against medical SAM baseline and uncertainty SAM variants under identical settings. LiteMedSAM (lightweight MedSAM) serves as the medical SAM baseline. For Pruned E-BayesSAM, we retain high-contribution tokens identified via KAN’s spline-based saliency analysis. Inference time is compared against Prompt-based augmentation [3], AutoSAM-U [21], and BayesFormer [20]. All methods use box prompts generated by expanding ground-truth masks by 10 pixels to simulate annotation variability.

Table 1. Performance comparison of our E-BayesSAM with other methods in terms of DSC (%). BayesMedSAM denotes the Bayesian MedSAM by individually using our T-VBI (Section 2.2).

Method	Thyroid Dataset		Breast Dataset			DSC _{Avg}
	DDTI	TN3K	UDIAT	BUSI	OASBUD	
MedSAM	90.0	88.9	91.0	90.6	80.9	88.3
BayesMedSAM	89.8	88.8	90.8	90.5	80.8	88.1
E-BayesSAM	89.4	88.7	90.6	90.5	80.8	88.0
Pruned E-BayesSAM	90.3	88.2	91.8	90.2	84.7	89.0




Fig. 3. E-BayesSAM achieves uncertainty-aware segmentation without training, and its pruned version even shows the outstanding performance.

3.3 Results

Performance of E-BayesSAM As demonstrated in Fig. 3, our E-BayesSAM successfully estimates segmentation uncertainty via T-VBI while preserving SAM’s zero-shot segmentation capability (Table 1). Quantitative results show a marginal decrease in average DSC (E-BayesSAM: 88.0% vs. MedSAM: 88.3%) but crucially enhance prediction reliability for clinical deployment. Qualitative analysis in Fig. 3 reveals E-BayesSAM’s dual behavior: improved segmentation of ambiguous boundaries (row 1 and 3 in Fig. 3) at the cost of slightly reducing accuracy on well-defined structures (row 2 and 4 in Fig. 3). We hypothesize this stems from KAN layers’ targeted tokens modifications—prioritizing uncertainty-sensitive features over anatomically stable regions, aligning with clinical prioritization of challenging cases.




Fig. 4. Four high-contribution tokens identified by KAN encode critical segmentation cues—boundary localization, foreground-background differentiation, and anatomical context—achieving accurate segmentation.

SO-KAN for Improving Interpretability and Accuracy Contribution analysis of KAN activations identifies four high-contribution tokens governing SAM’s segmentation decisions (Fig. 4). Since retaining five tokens decreases DSC by 5%, we select four tokens to optimize performance. These tokens correlate with anatomical boundary detection and background suppression—features essential for robust segmentation. Pruning the 28 lowest-contribution tokens reduces computation by 80% while improving baseline DSC by 0.7%. A similar observation was reported in DToP [23], which adaptively prunes tokens in SegViT and achieves a 0.5% improvement in mIoU.

Efficiency of E-BayesSAM for Uncertainty Estimation E-BayesSAM achieves efficient uncertainty estimation by confining Bayesian computation to SAM’s output tokens, avoiding full-network stochasticity. It processes 256×256 inputs in 0.03s/sample— $23\times$ faster than BayesFormer (0.70s) and $18\times$ faster than prompt augmentation (0.55s). This efficiency arises from token-wise reparameterization of T-VBI, adding 256 new parameters to SAM (9.8M), resulting in a model lighter than conventional Bayesian adaptations (>13.9 M). This design enables deployment on mobile ultrasound systems with <4 GB VRAM.

4 Conclusion

We present E-BayesSAM, a novel framework that addresses the critical challenges of Bayesian adaptation of SAM—computational inefficiency and inter-

pretability limitations—for ultrasonic image segmentation. By integrating our designed T-VBI and SO-KAN, E-BayesSAM leverages the strengths of VBI and KAN to achieve three key advances: (1) real-time inference (0.03s/sample), (2) improved segmentation accuracy (DSC: Pruned E-BayesSAM 89.0% vs. SAM’s 88.3%), and (3) intrinsic interpretability via token-wise saliency mapping. Crucially, KAN’s spline visualizations reveal that four tokens govern boundary detection and background suppression, enabling targeted token pruning without performance degradation. E-BayesSAM’s lightweight design and efficient inference bridge SAM’s versatility with the reliability required in clinical practice. It also inherits MedSAM’s multi-modality generalizability, enabling direct extension to other modalities, and provides a deployable solution for safety-critical medical applications.

Acknowledgments. This study was funded by Shenzhen Basic Science Research (Key Program) (grant number JCYJ20220818095612027).

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

References

1. Al-Dhabyani, W., Gomaa, M., Khaled, H., Fahmy, A.: Dataset of breast ultrasound images. *Data in brief* **28**, 104863 (2020)
2. Cheng, J., Wang, S., Liu, Y., Chen, Y., Chen, H., Li, Y.: Sam-med2d. *arXiv preprint arXiv:2308.16184* (2023)
3. Deng, G.D., Zou, K., Ren, K., Wang, M., Yuan, X., Ying, S., Fu, H.: Sam-u: Multi-box prompts triggered uncertainty estimation for reliable sam in medical image. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023 Workshops*. vol. 14394 (2023)
4. Dice, L.R.: Measures of the amount of ecologic association between species. *Ecology* **26**, 297–302 (1945)
5. Dusenberry, M., Jerfel, G., Wen, Y., Ma, Y., Snoek, J., Heller, K., Lakshminarayanan, B., Tran, D.: Efficient and scalable bayesian neural nets with rank-1 factors. In: *Proceedings of Machine Learning Research (PMLR)*. pp. 2782–2792 (2020)
6. Gawlikowski, J., Tassi, C.R., Ali, M., Lee, J.K., Humt, M., Feng, K., Kruspe, A., Triebel, R., Jung, P., Roscher, R., Shahzad, M., Yang, L., Bamler, R., Ng, A.Y., Kim, D., Meiner, J., Mikut, R., Heinrich, S., Weber, J., Brendel, M., Vogt, J.E.: A survey of uncertainty in deep neural networks. *Artificial Intelligence Review* **56**(1), 1513–1589 (2023)
7. Gong, H., Chen, J., Chen, G., Li, H., Li, G., Chen, F.: Thyroid region prior guided attention for ultrasound segmentation of thyroid nodules. *Computers in biology and medicine* **155**, 106389 (2023)
8. Huang, Y., Yang, X., Liu, L., Zhou, H., Chang, A., Zhou, X., Chen, R., Yu, J., Chen, J., Chen, C., Liu, S., Chi, H., Hu, X., Yue, K., Li, L., Grau, V., Fan, D.P., Dong, F., Ni, D.: Segment anything model for medical images? *Medical Image Analysis* **92**, 103061 (2024)

9. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? In: *Advances in Neural Information Processing Systems*. vol. 30 (2017)
10. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollar, P., Girshick, R.: Segment anything. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 4015–4026 (2023)
11. Lee, H.H., Gu, Y., Zhao, T., Xu, Y., Yang, J., Usuyama, N., Wong, C., Wei, M., Landman, B.A., Huo, Y., Santamaria-Pang, A.: Foundation models for biomedical image segmentation: A survey. *arXiv preprint arXiv:2401.07654* (2024)
12. Liu, Z., Wang, Y., Vaidya, S., Ruehle, F., Halverson, J., Soljačić, M., Hou, T.Y., Tegmark, M.: Kan: Kolmogorov-arnold networks. *arXiv preprint arXiv:2404.19756* (2024)
13. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: *International Conference on Learning Representations* (2019)
14. Ma, J., Wang, R., Gao, J., Song, W., Zhang, H., Ma, Q., Liu, Y., Chen, H., Chen, Y., Li, Y., Wang, S., Zhang, Y.: Segment anything in medical images. *Nature Communications* **15**(1), 654 (2024)
15. Mazurowski, M.A., Pardo, D., Chen, H., Li, Y., Liu, Y., Wang, S., Zhang, H., Ma, J.: Segment anything model for medical image analysis: An experimental study. *Medical Image Analysis* **89**, 102918 (2023)
16. Ning, G., Liang, H., Jiang, Z., Zhang, H., Liao, H.: The potential of segment anything (sam) for universal intelligent ultrasound image guidance. *BioScience Trends* **17**(3), 230–233 (2023)
17. Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Nowozin, S., Dillon, J., Lakshminarayanan, B., Snoek, J.: Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. In: *Advances in neural information processing systems*. vol. 32 (2019)
18. Pedraza, L., Vargas, C., Narváez, F., Durán, O., Muñoz, E., Romero, E.: An open access thyroid ultrasound image database. In: *In 10th International symposium on medical information processing and analysis*. vol. 9287, pp. 188–193 (2015)
19. Piotrkowska-Wróblewska, H., Dobruch-Sobczak, K., Byra, M., Nowicki, A.: Open access database of raw ultrasonic signals acquired from malignant and benign breast lesions. *Medical physics* **44**(11), 6105–6109 (2017)
20. Sankararaman, K.A., Wang, S., Fang, H.: Bayesformer: Transformer with uncertainty estimation. *arXiv preprint arXiv:2206.00826* (2022)
21. Shaharabany, T., Dahan, A., Giryas, R., Wolf, L.: Autosam: Adapting sam to medical images by overloading the prompt encoder. *arXiv preprint arXiv:2306.06370* (2023)
22. Shridhar, K., Laumann, F., Liwicki, M.: A comprehensive guide to bayesian convolutional neural network with variational inference. *arXiv preprint arXiv:1901.02731* (2019)
23. Tang, Q., Zhang, B., Liu, J., Liu, F., Liu, Y.: Dynamic token pruning in plain vision transformers for semantic segmentation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 777–786 (October 2023)
24. Wang, H., Cheng, J., Liu, Y., Chen, Y., Li, Y., Chen, H.: Sam-med3d. *arXiv preprint arXiv:2310.15161* (2023)
25. Yao, X., Liu, H., Hu, D., Lu, D., Lou, A., Li, H., Deng, R., Arenas, G., Oguz, B., Schwartz, N., Byram, B., Oguz, I.: Fnpc-sam: Uncertainty-guided false negative/positive control for sam on noisy medical images. In: *Proc SPIE Int Soc Opt Eng*. vol. 12926 (2024)

26. Yap, M.H., Pons, G., Martí, J., Ganau, S., Sentís, M., Zwigelaar, R., Davison, A.K., Martí, R.: Automated breast ultrasound lesions detection using convolutional neural networks. *IEEE Journal of Biomedical and Health Informatics* **22**(4), 1218–1226 (2018)
27. Zhang, Y., Hu, S., Jiang, C., Cheng, Y., Qi, Y.: Segment anything model with uncertainty rectification for auto-prompting medical image segmentation. *CoRR abs/2311.10529* (2023)
28. Zhang, Y., Shen, Z., Jiao, R.: Segment anything model for medical image segmentation: Current applications and future directions. *Computers in Biology and Medicine* **171**, 108238 (2024)
29. Zhou, N., Zou, K., Ren, K., Luo, M., He, L., Wang, M., Chen, Y., Zhang, Y., Chen, H., Fu, H.: Medsam-u: Uncertainty-guided auto multi-prompt adaptation for reliable medsam. *arXiv preprint arXiv:2409.00924* (2024)