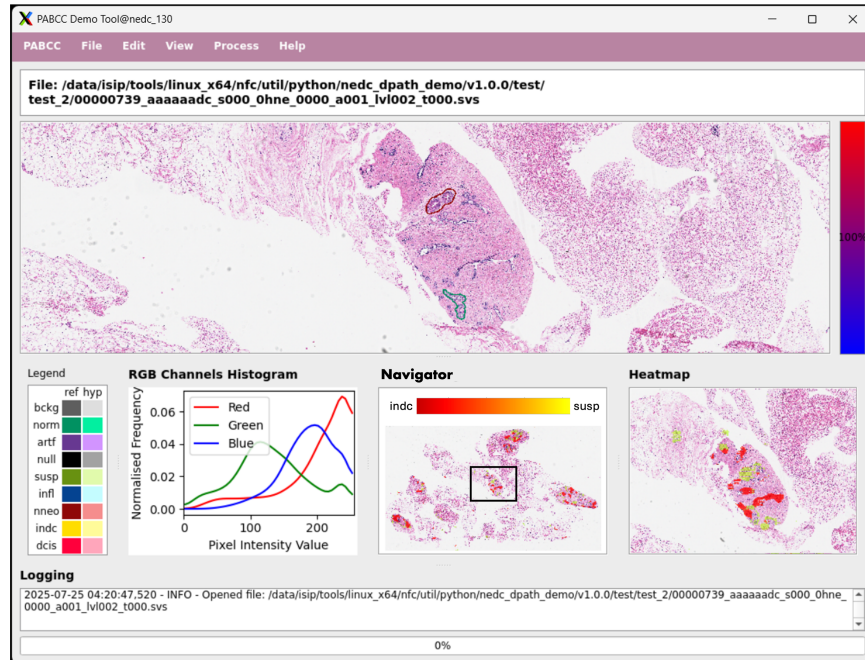


Rapid and Inexpensive Precision Breast Cancer Screening Using Machine Learning

August 1, 2025



Prepared By:


C. Dumitrescu, S. Purba, A. Al Mamun, D. Hackel, M. Bagritsevich,
P. Meng, D. Heathcote, I. Obeid and J. Picone

The Neural Engineering Data Consortium
College of Engineering, Temple University
1947 North 12th Street

Philadelphia, Pennsylvania 19122

Tel: 215-204-4841; Fax: 215-204-5960

Email: {claudia-anca.dumitrescu, sadia.afarin.purba, abdullah18, dmitry.hackel,
maria.bagritsevich, phuykong.meng, dylan.heathcote,
iobeid, picone}@temple.edu



**NEURAL ENGINEERING
DATA CONSORTIUM**

Part 1: Narrative Summary

The primary goal of the Pennsylvania Breast Cancer Consortium (PABCC) grant titled “Rapid and Inexpensive Precision Breast Cancer Screening Using Machine Learning” was to create software for automatically segmenting breast cancer pathology slides. Despite tremendous advances recently in artificial intelligence, high performance segmentation of high-resolution pathology images remains an elusive goal. In this one-year project, we have significantly advanced state of the art in segmentation, leading to an improved capability for triaging pathology slides and increasing the overall efficiency of manual review of these slides. In addition to advancing algorithm work, we have also developed a proof-of-concept demonstration of the technology, a screenshot of which is shown on the cover page of this document.

There were three components to our research work: data, algorithm research and demo. The tools and research developed have been placed in the public domain and are available from our industry-leading web site [1]. Acquisition of the data resources requires completion of an application form (a Temple University requirement). Downloads are performed several ways including via anonymous rsync. Software is freely available via downloadable archive files.

Data: The Temple University Hospital Breast Cancer Corpus (TUBR) is a subset of over 90,000 scanned slides containing a wide range of pathologies. The Fox Chase Cancer Center (FCCC) Breast Cancer Corpus (FCBR) is a subset of over 14,000 scanned slides from the FCCC Biosample Repository. FCBR is richer in terms of incidents of cancerous regions. TUBR provides a nice contrast since it was collected from a wider range of patients earlier in the diagnosis process. The differences between these two corpora is exemplified in the contrast between occurrences of the INDC and NNEO labels in Table 1.

Annotated training data plays a crucial role in the development of high performance machine learning systems. During this project, we have made significant enhancements to our industry-leading open source corpora. This is described in detail in [2]. We increased the number of annotated patches to an average of over 13 per slide. We also manually reviewed the annotations using our most experienced team of annotators and corrected several deficiencies (e.g., ensuring annotated regions form closed polygons). A summary of the statistics for the current versions of the corpora is given in Table 1.

Algorithm Research: Recognizing the critical nuances of histopathological microstructures, such as glandular architecture, nuclear morphology, and calcification, motivated us to evaluate a broad range of state-of-the-art deep learning architectures including a Residual Network (RsNet18) baseline architecture [3]-[5], EfficientNet (B0 and B7) [6], and Vision Transformer (ViT-16 and ViT-32) [7]. These systems were selected because they accurately represent a large range of popular deep learning architectures used in image processing. A summary of performance is given in Table 2. The computational complexity of the algorithms is shown in Table 3. A more extensive analysis of performance is given in [2].

The numbers in these tables represent an accuracy measure that ranges from 0% (poor) to 100% (good). Despite expectations that deeper or more complex models such as EfficientNet-B7 or ViT-32 might dominate, our findings revealed that EfficientNet-B0 consistently outperformed all other architectures across both accuracy and generalization metrics. Its superior performance was particularly notable given its smaller parameter count and computational efficiency, suggesting that deeper models may be prone to overfitting on histopathological textures, especially when working with high-resolution multi-scale images.

Table 1. A comparison of TUBR and FCBR

Attribute	TUBR (v5.0.0)	FCBR (v3.0.1)
No. Files:		
/train	1,652	765
/dev	932	373
/eval	921	325
TOTAL	3,505	1,463
Size (Gbytes)	1,257	3,357
No. Labels	46,666	20,074
Avg. No./Slide	13.31	13.72
Area/Slide (%)	2.26%	2.52%
No. Labels:		
norm	13,324 [28.55%]	316 [1.57%]
dcis	1,673 [3.59%]	1,209 [6.02%]
indc	2,552 [5.47%]	10,952 [54.56%]
nneo	14,293 [30.63%]	729 [3.63%]
infl	2,700 [5.79%]	1,214 [6.05%]
artf	2,185 [4.68%]	914 [4.55%]
null	2,910 [6.24%]	1,087 [5.41%]
susp	237 [0.51%]	24 [0.12%]
bckg	6,792 [14.55%]	3,629 [18.08%]

Building on this, we also developed a novel hybrid architecture combining a Swin Transformer backbone with a Unet decoder, augmented with both a classification head [8]-[10]. This dual-headed model sought to exploit both pixel-wise segmentation fidelity and region-level classification robustness, enabling it to deliver not just a label but spatial localization of pathological structures. Preliminary results demonstrate that this architecture substantially improves explainability and clinical trust, allowing regions of interest to be visualized with high alignment to annotated structures.

Demo: In addition to the scientific contributions, a major deliverable of this research cycle was the development of a real-time graphical user interface (GUI). A screenshot is shown on the cover page of this document. This open source GUI integrates the trained model and displays – in near real-time – regions predicted as malignant or pre-malignant. The assistive software was designed to map images to interactive heatmaps, providing actionable insights for pathologists, thereby improving diagnostic efficiency and reducing cognitive load. All software and models used in this demo are available from our web site.

Next Steps: Though we have made great strides in improving the performance of the classification system, there is still work to be done before the technology will be clinically acceptable. We plan to collect feedback from pathologists on the current demo system, and to explore ways to make the visualizations more effective and informative. The core microsegmentation problem remains an important research focus for our group. We are actively seeking funding to continue this work, as described below. We are also expanding this framework to encompass additional tissue types and cancer subtypes beyond breast histology. We are currently preparing a manuscript that describes an exploration of the role detection of calcification can play in improved cancer detection [11]. We are also preparing a manuscript that explores fine-tuning of a large language model based system to do automated segmentation [12]. A third study involving the interpretability of a hybrid Swin-Unet model is in its early stages [13].

Grant Applications: Two major proposals that are a direct outcome of this work are: (1) “Precision Breast Cancer Detection in High Resolution Digital Pathology Using Hierarchical Self-Attention and Lightweight Microsegmentation” submitted to the DoD Breast Cancer Research Program (\$1.7M for three years), and (2) “Prostate Cancer Screening Using Whole Slide Digital Pathology and Deep Learning” under development to be submitted to the DoD Congressionally Directed Medical Research Programs (CDMRP) (\$1.5M for three years). The former is a direct extension of this project. The latter extends the technology developed to a new application domain. We also used portions of this research for two Temple internal proposals: (1) “Robust Segmentation of Signal Data Using Foundation Models” (\$125K for two years) and (2) “Exploiting Quantum Entanglement in Machine Learning” (\$100K for one year). We included this work in a proposal titled “Generative Models for Automatic Classification of Biomedical Signals” to the Keck Foundation. Though we were selected to represent Temple in this highly competitive solicitation, and got past the first round of reviews, we were not selected as a finalist for a site visit. In addition to these proposals, we have applied much of the data and technology developed in this project to an on-going NSF grant involving the application of quantum computing to machine learning. Breast cancer classification is one of the two domains in which we are working for this grant. During Fall 2025, we expect to write two major NSF grants on machine learning applications in breast cancer classification. Finally, we continue to monitor additional DoD and foundation opportunities for relevant funding opportunities.

Table 2. A performance comparison

Arch	TUBR			FCBR		
	Train	Dev	Eval	Train	Dev	Eval
ResNet18	40.87	37.42	33.28	44.40	44.34	42.11
EB0	62.32	56.60	51.72	69.83	54.49	55.02
EB7	67.17	55.67	50.72	57.98	47.00	49.39
ViT-16	62.68	48.32	44.72	26.11	24.44	25.07
ViT-32	57.12	46.34	41.39	29.58	26.11	28.06

Table 3. Complexity analysis (TUBR)

Arch	No. Parameters	Training Time	Decoding Time
ResNet18	11.69M	2,356	1,597
EB0	5.30M	3,900	1,520
EB7	66.35M	4,883	1,735
ViT-16	86.60M	3,713	1,726
ViT-32	88.20M	3,869	1,623

Part 2: Publication Overview

We are now preparing to formally disseminate the outcomes of this NSF-funded work. Several manuscripts are currently in preparation. Our flagship publication [2] is a book chapter that represents an expanded version of [14]. This describes the corpora we have developed, and the experiments summarized in this report. A second manuscript is under development that will introduce three new classes to the FCBP Corpus that identify various levels of calcification. This crystallization work represents a strategic expansion of our project: by quantifying mineral composition and morphology [15][16] within annotated calcifications, we aim to elevate predictive accuracy and uncover novel biomarkers for in situ carcinoma. The maturity of our infrastructure – with over 100,000 WSIs digitized, over 66,000 labels manually annotated by a team of experts, and a GPU-accelerated AI pipeline – positions us to deliver robust, reproducible results. A third manuscript under development and expected to be submitted to IEEE SPMB 2025 (which we host) will discuss some very interesting experiments in which we used human-in-the-loop reinforcement learning to adapt a general segmentation foundational model [17] to breast cancer diagnosis.

These publications will be made available from our web site as soon as the drafts are complete. We will also provide PABCC with copies as soon as the papers have been accepted for publication.

Part 3: Impact

The data and software developed in this project will profoundly impact the international community. The corpora we have developed are unique in several ways: (1) the images coupled with the associated metadata will support many studies into the basic science, (2) the volume of data along with manual annotations will support training of the so-called large language models, and that has not been previously possible, and (3) manual annotations of high resolution images will motivate researchers to address the microsegmentation problem.

The algorithm baselines established in this work will serve as important reference points by which to measure progress. Though there is a widely held perception that artificial intelligence has solved these problems, our results indicate significant work needs to be done before the technology will be clinically relevant. Publicly available and reproducible benchmarks will accelerate progress because researchers will have very clear performance targets and a well-documented experimental infrastructure in which to conduct research. Our software has been designed to be easy to download, install and use. We provide online support as well through our well-known listserv *help@nedcdata.org*. Support requests are usually answered within a few hours.

The demonstration vehicle will allow us to communicate in meaningful ways with clinicians. It will help educate them on the technology and better inform us about their clinical needs.

Our experience has been that when these types of resources exist, the user base will expand significantly. As we increasingly publish our work, more people will become aware of these resources and motivated to improve technology. Commercial and non-profit entities will be engaged, and progress will accelerate. We expect within a couple of years we will have over 10,000 subscribers of these resources. Expanding the pool of researchers exploring these problems is an important first step to accelerating progress.

Acknowledgements

The development of data and resources of this scale has been a long-term effort that has been supported by several organizations over the years including the National Science Foundation (grants nos. CNS-1726188 and 1925494) and the Temple University Catalytic Collaborative Funding Initiative. The Pennsylvania Breast Cancer Coalition Breast and Cervical Cancer Research Initiative grant has been instrumental in our development of a wide range of machine learning solutions and our demonstration system. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of Temple University or the sponsors associated with this research.



References

- [1] Picone, J., & Obeid, I. (2025). NEDC Resources. The Neural Engineering Data Consortium. [Online]. url: <https://isip.piconepress.com/projects/nedc/html/resources.shtml>.
- [2] Hackel, D., Bagritsevich, M., Dumitrescu, C., Al Mamun, Md. A., Purba, S. A., Heathcote, D., Obeid, I., & Picone, J. (2026). Enabling Microsegmentation: Digital Pathology Corpora for Advanced Model Development. In *Signal Processing in Medicine and Biology: Applications of Artificial Intelligence in Medicine and Biology* (Vol. 1, p. 50). Springer. [Online]. url: https://isip.piconepress.com/publications/book_sections/2026/springer/dpath/.
- [3] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. doi: 10.1109/CVPR.2016.90.
- [4] Khalkhali, V., Shawki, N., Shah, V., Golmohammadi, M., Obeid, I., & Picone, J. (2021). Low Latency Real-Time Seizure Detection Using Transfer Deep Learning. In I. Obeid, I. Selesnick, & J. Picone (Eds.), *Proceedings of the IEEE Signal Processing in Medicine and Biology Symposium (SPMB)* (pp. 1–7). IEEE. doi: 10.1109/SPMB52430.2021.9672285.
- [5] Alexandrov, D., & Picone, J. (2024). The Impact of ECG Channel Reduction on Multi-Label Cardiac Diagnosis. *Proceedings of the IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*, 6. doi: 10.1109/SPMB62441.2024.10842233.
- [6] Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In K. Chaudhuri & R. Salakhutdinov (Eds.), *Proceedings of the International Conference on Machine Learning (ICML)* (Vol. 97, pp. 6105–6114). PMLR. url: <http://proceedings.mlr.press/v97/tan19a.html>.
- [7] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., & Houlsby, N. (2021). An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *Proceedings of the International Conference on Learning Representations (ICLR)*, 1–21. url: <https://iclr.cc/virtual/2021/oral/3458>.
- [8] Liu, Z., Hu, H., Lin, Y., Yao, Z., Xie, Z., Wei, Y., Ning, J., Cao, Y., Zhang, Z., Dong, L., Wei, F., & Guo, B. (2022). Swin Transformer V2: Scaling Up Capacity and Resolution. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11999–12009. doi: 10.1109/CVPR52688.2022.01170.
- [9] Atabansi, C. C., Nie, J., Liu, H., Song, Q., Yan, L., & Zhou, X. (2023). A survey of Transformer applications for histopathological image analysis: New developments and future directions. *BioMedical Engineering OnLine*, 22(1), 96. doi: 10.1186/s12938-023-01157-0.
- [10] Nguyen, C., Asad, Z., Deng, R., & Huo, Y. (2022). Evaluating transformer-based semantic segmentation networks for pathological image segmentation. *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 12032, 120323N. doi: 10.1117/12.2611177.
- [11] Dumitrescu, C., Hackel, D., Obeid, I., & Picone, J. (2025). Crystallization Signatures as Predictive Biomarkers in Cancer Pathology. *Proceedings of the IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*, 1, 1–5. url: https://isip.piconepress.com/publications/unpublished/conferences/2025/ieee_spmb/calcalcification/.

- [12] Purba, S. A., Melles, A.-M., Hackel, D., & Picone, J. (2025). Assessing the Visual Reasoning of Multimodal Language Models in Biomedical Applications. *Proceedings of the IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*, 1, 1–5.
url: https://isip.piconepress.com/publications/unpublished/conferences/2025/ieee_spmb/segmentation/.
- [13] Dumitrescu, C., Hackel, D., Obeid, I., & Picone, J. (2026). An Analysis of the Interpretability of a Hybrid Swin-Unet Model for Digital Pathology. *Proceedings of the IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*, 1, 1–5.
url: <https://isip.piconepress.com/publications/unpublished/conferences/2026/explainability>.
- [14] Bagritsevich, M., Hackel, D., Obeid, I., & Picone, J. (2024). Annotation of the Fox Chase Cancer Center Digital Pathology Corpus. *Proceedings of the IEEE Signal Processing in Medicine and Biology Symposium*, 1–4. doi: 10.1109/SPMB62441.2024.10842255.
- [15] Shin, K. S., Laohajatsang, M., Men, S., Figueroa, B., Dintzis, S. M., & Fu, D. (2020). Quantitative chemical imaging of breast calcifications in association with neoplastic processes. *Theranostics*, 10(13), 5865–5878. Doi: 10.7150/thno.43325.
- [16] Scott, R., Stone, N., Kendall, C., Geraki, K., & Rogers, K. (2016). Relationships between pathology and crystal structure in breast calcifications: An in situ X-ray diffraction study in histological sections. *Npj Breast Cancer*, 2(1), 16029. doi: 10.1038/npjbcancer.2016.29.
- [17] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W. -Y. Lo, P. Dollár, & R. Girshick. (2023). Segment Anything. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 3992–4003.
doi: 10.1109/ICCV51070.2023.00371.