

Multiplexed Expansion Microscopy for Drug Response Prediction in MIBC

Nouha Tiyal

CMU-CS-25-127

August 2025

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:

Russell S. Schwartz, Chair
Min Xu

*Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Computer Science.*

Copyright © 2025 **Nouha Tiyal**

Keywords: Computational biology, medical imaging, deep learning, multiplexed fluorescence microscopy, expansion microscopy, ExPATH, whole-slide imaging, nuclear morphology, DAPI, TelC, CenpB, WGA, ablation, muscle-invasive bladder cancer.

To Allah, the Most Merciful, for His guidance and countless blessings throughout this journey.

“My Lord, increase me in knowledge.” (Qur’an 20:114)

To my beloved parents, whose love, sacrifices, and prayers made this work possible.

Abstract

Expansion microscopy (ExPath) enables nanoscale resolution of tissue architecture using conventional microscopes, offering a powerful alternative to traditional histopathology. In this thesis, we present a deep learning pipeline that leverages ExPath imaging combined with a biologically informed, four-channel multiplexed staining panel: DAPI, TelC, CENPB, and WGA to classify tissue types and predict chemotherapy response in muscle-invasive bladder cancer (MIBC). We propose that nuclear morphology, when captured at high resolution and enriched by chromatin and membrane-specific markers, contains sufficient information to compete with H&E and generalize across diagnostic and prognostic tasks. To test this hypothesis, we construct a preprocessing pipeline that transforms 16-bit 4-channel TIFF WSIs into normalized, pseudo-RGB 1024×1024 patches compatible with ImageNet-pretrained models. We evaluate multiple architectures (ResNet34, ResNet50, ViT-tiny, EfficientNet) and demonstrate that ResNet-based models trained on ExPath outperform simulated non-ExPath baselines and DAPI-only variants by a significant margin. Through controlled ablation experiments, we quantify the contribution of each channel and find that multiplexing substantially boosts classification accuracy. Our models achieve 89.52% tissue classification accuracy and 0.9 ROC-AUC for drug response prediction. Furthermore, we observe cross-cancer generalizability when applying MIBC-trained models to lung carcinoma ExPath images. This work establishes the feasibility of compact, multiplexed, ExPath-driven classification pipelines as a viable alternative to costly multi-modal diagnostics. It offers an early step toward a DAPI-first foundation model for computational pathology, with potential to scale across cancer types and tissue conditions using minimal staining and high-content imaging.

Acknowledgments

First and foremost, all praise and thanks are due to Allah, the Most Gracious, the Most Merciful, for His blessings during my studies and in completing this thesis; without His grace, none of this would have been possible. I ask Allah to make this knowledge beneficial to me and a means by which I benefit others.

I owe a profound debt to my beloved parents. Their unwavering love, sacrifices, and prayers are the foundation of my life and learning. I can never repay what they have given me; I pray that Allah rewards them abundantly and grants them the highest ranks in Jannah.

My heartfelt thanks go to my advisor, Professor Russell S. Schwartz, for steadfast mentorship, patient guidance, and exacting feedback throughout this journey. I also thank Professor Min Xu, committee member and former instructor, for his incisive insights and encouragement. I am grateful to the Schwartz Lab for their collaboration and daily support, and to Professor Leon Zhao and the Zhao Lab for their collaboration and for providing the data and documentation that enabled this work.

I gratefully acknowledge the dedicated efforts of peers Jason Nguyen, Hanbing Zou, Nishat Anjum Bristy, Thomas Rachman, Arpitha Rao, Alex Guo, and others whose help strengthened this project.

I am thankful to Carnegie Mellon University for its scholarly community and resources, and to Professor Peter Steenkiste, Head of the Fifth-Year Master's Program, for his guidance and assistance from the very beginning. I also acknowledge with gratitude the support of my master's studies by the "Zakat & Charity" Fund under the Hadia Abdul Latif Jameel Group.

Last, but certainly not least, I extend my sincere thanks to my family members, my friends in Pittsburgh and beyond, and my professors at the Qatar campus for their constant encouragement and kindness. Their presence, even from afar—has been a steady source of strength and motivation.

Contents

1	Background	1
1.1	Digital Pathology and Computational Approaches	1
1.2	Expansion Pathology	2
1.3	Muscle-Invasive Bladder Cancer (MIBC)	2
2	Related Work	5
3	Dataset	7
3.1	Preprocessing and Image Normalization	8
3.2	Data Composition and Annotation	8
3.3	From Dataset to Proof-of-Concept	8
4	Methods	9
4.1	Data Processing	9
4.2	Model Architectures and Training	11
4.2.1	Training Details	11
5	Experiments Setup	13
5.1	Experimental Tasks	13
5.1.1	ExPath vs. Simulated Non-ExPath Imaging	13
5.1.2	Marker Contribution via Ablation Studies	13
5.1.3	Classification Generalizability	14
5.2	Data Splits and Sampling	14
5.3	Evaluation Metrics	14
6	Results	15
6.1	ExPath Outperforms Simulated Non-ExPath Imaging	15
6.2	Marker Contribution Analysis via Ablation Studies	15
6.2.1	Post-hoc Evaluation Ablations	15
6.2.2	Training-Time Ablations	17
6.3	Classification Tasks	19
6.3.1	Tissue Classification	19
6.3.2	Drug Response Prediction	19
6.3.3	Cross-Cancer Generalization (MIBC → Lung)	20

6.4	UMAP Visualization of Patch Embeddings	21
7	Discussion and Limitations	23
7.1	Biological and Clinical Implications	23
7.2	Dataset Limitations	23
7.3	Training and Evaluation Constraints	24
7.4	Pipeline-Level Limitations	24
7.5	Interpretability and Trustworthiness	24
8	Future Work	27
9	Conclusion	29
	Bibliography	31

List of Figures

4.1	Data Processing Pipeline	9
4.2	Images after processing steps	10
6.1	Accuracy Confusion Matrix for Model trained with ExPath	16
6.2	Overall Post-hoc Ablation Results	17
6.3	Per-class Post-hoc Ablation Results	18
6.4	Training Ablation Results	19
6.5	Accuracy Confusion Matrix for Tissue Classification (3-way and 4-way)	20
6.6	Accuracy Confusion Matrix for Drug Response Prediction	20
6.7	UMAP Visualization of Patch Embeddings	22

List of Tables

Chapter 1

Background

1.1 Digital Pathology and Computational Approaches

Digital pathology is the process of scanning tissue slides to create high-resolution images that can be viewed, shared, and analyzed on a computer. This transformation has laid the foundation for modern computational pathology, enabling the application of machine learning to large-scale image datasets for diagnosis, prognosis, and biomarker discovery. By shifting pathology into a digital format, it becomes possible to develop algorithms that systematically capture phenotypic patterns across thousands of patients, moving toward more data-driven healthcare. Despite this progress, most computational pathology pipelines today rely on hematoxylin and eosin (H&E) staining as the dominant modality. H&E provides rich morphological context, but it is limited in two ways: (1) it does not capture molecularly specific signals in the nucleus such as chromatin organization or telomere integrity, and (2) it is fundamentally constrained by the diffraction limit of light microscopy. DAPI-based models, which directly capture the nucleus, are rare and typically used only in combination with other biomarkers, since they allow multiplexing. Meanwhile, multimodal pipelines that fuse histology with genomic or transcriptomic data, for example, prognostic modeling frameworks developed by the Mahmood Lab [Lu et al., 2022], offer molecular depth but require costly sequencing workflows and complex integration strategies. This creates a gap between scalable, image-only methods and high-cost, resource-intensive multimodal approaches. Our work addresses this gap by using Expansion Pathology (ExPath) on DAPI stained tissue as an alternative to conventional H&E-driven computational pathology. ExPath offers physical expansion of tissue and DAPI staining allows for co-staining with targeted fluorescent labeling, providing nanoscale resolution and DNA-FISH like molecular readouts (via DAPI, TelC, CENPB, and WGA). We hypothesize that this combination of higher resolution and molecularly informative fluorescent labeling can generate image-derived features that are both scalable and prognostically meaningful. The clinical problem we focus on is Muscle-Invasive Bladder Cancer (MIBC), a highly aggressive form of bladder cancer for which predictive biomarkers of treatment response remain lacking. In the following sections, we first provide background on ExPath as a methodological innovation, and then on MIBC as the motivating clinical context for our study.

1.2 Expansion Pathology

Expansion Pathology (ExPath) is a technique that leverages the principles of expansion microscopy (ExM) to physically magnify biological specimens, enabling nanoscale imaging of tissues using conventional microscopes. Originally developed to visualize fine neuronal structures, ExM has since been adapted to pathological contexts to enable detailed spatial mapping of cells and subcellular components in disease tissues [Chen et al., 2015; Zhao et al., 2017].

The core idea behind ExPath is to embed tissue sections in a swellable polymer matrix and then isotropically expand them several-fold, usually by 4× or more. This physical expansion circumvents the diffraction limit of light microscopy, allowing researchers to observe structures at an effective resolution of 70 nm with standard confocal or widefield microscopes. Notably, this allows subcellular features like chromatin structure, protein complexes, and cell–cell junctions to be visualized with high fidelity [Tillberg et al., 2016].

ExPath by itself is a physical technique: tissue is embedded in a swellable hydrogel and expanded isotropically, enabling high-resolution imaging with conventional microscopes. When combined with fluorescent labeling, this approach becomes particularly powerful for computational pathology. In our study, we use ExPath together with a four-channel panel designed to capture both structural coverage and biologically relevant nuclear features: DAPI (4',6-diamidino-2-phenylindole) for nuclear shape and chromatin texture, TelC (a telomeric repeat probe) as a telomeric probe reflecting replicative potential and genome stability [Shay, 2016], CENPB (centromere protein B antibody) to highlight centromeric heterogeneity implicated in oncogenesis [McGovern et al., 2012; Briasoulis et al., 2008; Wallander et al., 2021], and WGA (wheat germ agglutinin) to outline cell membrane boundaries through glycoprotein binding. Together, these markers provide a comprehensive view of nuclear and cellular organization, complementing the resolution gains offered by expansion.

Compared to traditional hematoxylin and eosin (H&E) staining, which offers a more generalized and morphologically rich overview of tissue context, ExPath emphasizes fine-grained nuclear and subnuclear detail. While H&E is commonly used for clinical histopathology, it does not support multiplexing or nanoscale imaging. In contrast, ExPath and DAPI enable both, making it especially powerful for high-content, quantitative analysis of tissue microenvironments and cellular phenotypes [Zhao et al., 2017].

This flexibility makes our approach particularly promising for computational pathology tasks that rely on detailed nuclear morphology, such as segmentation, classification, and prognosis prediction in cancer.

1.3 Muscle-Invasive Bladder Cancer (MIBC)

Bladder cancer is the tenth most common cancer worldwide and exhibits a dichotomy in clinical behavior based on the depth of tumor invasion. Muscle-Invasive Bladder Cancer (MIBC) refers to tumors that have invaded the detrusor muscle (muscularis propria) of the bladder wall, which marks a shift toward more aggressive disease with a significantly worse prognosis than non-muscle invasive forms [Siegel et al., 2024; Babjuk et al., 2022].

MIBC accounts for approximately 25-30% of newly diagnosed bladder cancer cases and re-

quires more intensive treatment due to its propensity for metastasis. The current standard of care involves neoadjuvant chemotherapy (typically cisplatin-based), followed by radical cystectomy and lymph node dissection [Grossman et al., 2003]. Common first-line regimens include MVAC (methotrexate, vinblastine, doxorubicin, and cisplatin) and dose-dense GC (gemcitabine and cisplatin). However, therapeutic response varies widely, and there is currently no robust imaging-based biomarker to predict patient response pre-operatively [Stein et al., 2001].

There is increasing interest in integrating histopathological and molecular data to stratify MIBC patients and guide treatment decisions. Although molecular subtyping (e.g., luminal vs. basal) has yielded predictive insights in some studies [Robertson et al., 2017], these approaches often require costly and labor-intensive sequencing workflows. A more scalable alternative would be to use computational analysis of pathology images, especially if such images could be collected with minimal stain or modality burden and yet still capture prognostically relevant phenotypes.

This sets the stage for imaging-based classifiers that can leverage the nuclear, chromosomal, and glycoprotein features revealed by multiplexed expansion microscopy to predict treatment response in MIBC. Unlike traditional approaches that depend on coarse morphological patterns or require expensive multi-modal imaging, we hypothesize that the micro- and nano-scale features captured in DAPI-based multiplexed ExPath images can serve as a rich basis for tissue- and drug-response classification in bladder cancer and beyond.

Chapter 2

Related Work

Hematoxylin and eosin (H&E) staining remains the cornerstone of clinical histopathology, enabling visualization of tissue architecture, cellular morphology, and pathological abnormalities at scale. Hematoxylin stains nuclei dark blue-purple by binding to DNA and nuclear proteins, while eosin stains extracellular matrix and cytoplasmic components pink [Fischer et al., 2008]. This combination provides rich contextual information about both cellular and stromal components, making H&E suitable for diagnosis, subtyping, and prognosis across many cancer types.

Deep learning (DL) models trained on H&E-stained whole-slide images (WSIs) have demonstrated high accuracy in classification, segmentation, and survival prediction tasks [Campanella et al., 2019; Coudray et al., 2018; Lu et al., 2021]. These models benefit from the ubiquity and generalizability of H&E data, which is standardized across most clinical labs. However, despite its broad adoption, H&E has intrinsic limitations: it lacks molecular specificity and cannot be easily multiplexed to track biomarkers beyond basic morphology. Moreover, H&E does not capture subnuclear chromatin patterns or precise molecular states relevant to targeted treatment strategies.

In response, many recent studies have turned to expensive multimodal workflows, combining H&E with immunohistochemistry (IHC), RNA sequencing, or spatial transcriptomics, to extract richer, more predictive features [He et al., 2021; Zhao et al., 2017]. While powerful, these approaches are resource-intensive, require tissue consumption across modalities, and do not scale well in clinical settings.

On the other hand, DAPI (4',6-diamidino-2-phenylindole) is a fluorescent stain that binds strongly to A-T-rich regions of DNA, providing a high-resolution view of nuclear morphology. It is commonly used in fluorescence microscopy for identifying nuclei, segmenting cells, and analyzing chromatin states [Kapuscinski, 1995]. In contrast to H&E, DAPI staining enables multiplexing with other fluorescence-based markers and is compatible with advanced imaging modalities such as confocal microscopy, super-resolution microscopy, and expansion microscopy.

In digital pathology, DAPI has been leveraged for segmentation tasks, cell counting, and morphometric analysis. For example, Cellpose [Stringer et al., 2021] and StarDist [Schmidt et al., 2018] are two widely-used nucleus segmentation models trained on DAPI-stained images. The simplicity and consistency of DAPI staining make it an ideal anchor channel for high-throughput nucleus-based analysis. Moreover, because DAPI highlights nuclear shape, size, and chromatin texture, it can serve as a basis for downstream classification tasks, particularly in cancers where

nuclear morphology correlates with malignancy grade and prognosis [Zink et al., 2004].

However, DAPI alone lacks the tissue context and cytoplasmic or stromal information necessary for generalized diagnosis or classification across organs. It is not sensitive to extracellular matrix composition, immune infiltration, or spatial tissue architecture, features that H&E naturally encodes. As such, DAPI-based classification models have often been limited to highly controlled settings or narrowly-defined tasks.

Several deep learning models have been proposed for cell or tissue classification using only morphological features, particularly from nuclear stains. These approaches typically rely on instance-level segmentations followed by handcrafted or learned features (e.g., shape descriptors, intensity texture) to distinguish between healthy and diseased cells or tissues [Sirinukunwattana et al., 2016; Khoshdeli et al., 2017].

Biologically, the rationale behind morphology-only classification is rooted in the observation that malignant transformation often manifests through nuclear atypia, such as irregular nuclear contours, hyperchromasia, and increased nucleolar size. These traits have been used for decades in manual cytopathological scoring systems. However, relying solely on nuclear features may not distinguish between phenotypes that require stromal, vascular, or immune context, limiting the accuracy and robustness of morphology-only models for prognosis or subtype prediction [Aeffner et al., 2019].

We aim to address the limitations of both H&E-based and DAPI-only approaches by leveraging Expansion Pathology (ExPath) on DAPI-stained tissue combined with DNA-FISH probes. This integration unites nanoscale resolution with multiplexed fluorescent labeling to capture structural and molecular features relevant across cancer types. DAPI serves as the multiplexing anchor, consistently highlighting nuclear morphology while enabling co-staining. ExPath expands tissues to achieve 70 nm effective resolution with standard microscopes, overcoming diffraction limits. To enrich nuclear features, we incorporate biologically motivated markers: TelC for telomere length (replicative potential and genome stability), CENPB for centromere organization (chromosomal integrity and oncogenesis), and WGA for glycoprotein-rich cellular boundaries. Together, these markers provide a richer molecular and structural context than H&E alone and avoid the cost and complexity of multimodal sequencing workflows.

Our goals are fourfold: first, to build classification models for tissue and chemotherapy response prediction in MIBC that perform competitively with state-of-the-art image-based methods; second, to demonstrate that ExPath confers a measurable performance edge compared to standard-resolution pipelines; third, to quantify the contribution of each channel through systematic ablation experiments; and finally, to test the cross-cancer generalizability of MIBC-trained models, using lung carcinoma as an initial case study. These efforts collectively lay the groundwork for a DAPI-first foundation model in pathology, where ExPath-enabled nuclear detail and multiplexed probes offer a scalable, generalizable alternative to H&E or costly multimodal diagnostics. The following sections describe the dataset, preprocessing pipeline, experimental methods, results, and implications of this approach.

Chapter 3

Dataset

A central question in computational pathology is whether morphological information alone is sufficient for accurate diagnosis, prognosis, or treatment prediction. Traditional H&E imaging offers rich morphological and contextual features but is inherently limited in molecular specificity and multiplexing flexibility. In contrast, fluorescence imaging based on DAPI highlights nuclear morphology at high resolution and allows multiplexing with additional markers, but lacks cytoplasmic and stromal context. This creates a fundamental tradeoff between generalizability and biological specificity.

To investigate whether nuclear morphology can be enriched and made generalizable through selective multiplexing, we use expansion microscopy (ExPath) images generated by the Zhao Lab at Carnegie Mellon University. These images are stained with a four-channel panel comprising:

- DAPI: labels DNA and highlights nuclear shape and chromatin architecture
- TelC: a telomeric repeat probe indicating telomere integrity and nuclear aging patterns
- CENPB: a centromere protein that reveals chromosomal structure and nuclear organization
- WGA: wheat germ agglutinin, a membrane/glycoprotein stain that outlines cell boundaries and extracellular structures

This panel is designed to overcome the limitations of DAPI-only imaging by capturing biologically informative markers relevant to cancer progression and cellular heterogeneity. Importantly, all four markers are compatible with expansion microscopy and were selected for their ability to (1) retain structural integrity post-expansion, and (2) encode features that contribute to classification and generalization.

The ExPath imaging protocol expands physical tissue sections by approximately 4 $\times$ linearly, resulting in 70 nm resolution with standard confocal microscopy. This level of detail enables precise visualization of subnuclear and membrane-level features that are difficult to resolve in standard 40 $\times$ H&E slides. As a result, a single 1024 $\times$ 1024 ExPath patch contains cellular features analogous to a 256 $\times$ 256 H&E patch in terms of content density but at higher resolution and specificity. To accommodate this, our pipeline operates on 1024 $\times$ 1024 patches instead of the standard 224 $\times$ 224 or 256 $\times$ 256 used in typical WSIs.

3.1 Preprocessing and Image Normalization

Raw images are stored as 16-bit grayscale TIFF files across four fluorescence channels. However, directly training on this format is computationally expensive and incompatible with most pretrained image models (e.g., ImageNet-based ResNets). To address this, we implement a pre-processing pipeline that includes: Bit Depth Reduction: From 16-bit to 8-bit per channel. Despite potential overflows, this simplification dramatically reduces memory usage and speeds up downstream training without substantial performance loss. Previous work (e.g., [Siddiqui et al., 2021]) suggests bit depth reduction can be effective when paired with normalization. Channel Normalization: Each channel is normalized both across the entire slide and individually to control for staining variance and ensure numerical stability during training. RGB Conversion: Channels are recombined into pseudo-RGB format for compatibility with pretrained CNN architectures. While not physically meaningful in terms of color, this enables transfer learning from conventional image models (e.g., ResNet34, ViT-tiny). Slide Tiling: WSIs are tiled into non-overlapping 1024×1024 patches. Given the expanded size of ExPath cells, this patch size captures sufficient tissue context for classification while avoiding excessive downsampling.

This pipeline reduces image storage size by over 150% and improves training speed by nearly 5×, making it feasible to scale learning experiments across full WSIs.

3.2 Data Composition and Annotation

The dataset includes multiple WSIs from patients with muscle-invasive bladder cancer, labeled for two major tasks:

- Tissue-Type Classification: Each patch is assigned one of 3 classes based on expert annotation and visual inspection: Cancerous, Lymphocyte, Smooth muscle.
- Drug Response Prediction: Patches are aggregated at the slide level to infer response to neoadjuvant chemotherapy (progressive, partial, or complete), serving as a supervised signal for treatment outcome classification.

Annotation was performed using expert review.

3.3 From Dataset to Proof-of-Concept

While our long-term goal is to develop a general-purpose DAPI-based foundation model, the current dataset is not yet sufficient in scale or diversity to support self-supervised pretraining. Instead, we propose a proof-of-concept model that tests the core hypothesis: that a carefully selected DAPI + multiplex panel imaged under ExPath can enable accurate tissue classification and drug response prediction, potentially replacing more expensive multi-modal systems.

To this end, we train both patch-wise and slide-wise classifiers to evaluate the impact of individual markers, the effectiveness of ExPath imaging, and the generalizability of our approach. These experiments are designed to validate the value of our dataset and motivate further data collection for full foundation model training.

Chapter 4

Methods

Our pipeline is designed to transform high-resolution, multiplexed ExPath whole-slide images (WSIs) into patch-based inputs suitable for classification tasks such as tissue-type labeling and chemotherapy response prediction. We evaluate several deep learning architectures on this task and compare their performance to identify models best suited for downstream experiments.

The core of our computational pipeline follows an end-to-end workflow comprising:

1. WSI Tiling: Full slides are divided into non-overlapping patches of size 1024×1024 pixels.
2. Preprocessing and Normalization: Images are converted to pseudo-RGB format, reduced from 16-bit to 8-bit, and channel-normalized.
3. Model Inference: Each patch is passed through a deep learning model (e.g., ResNet34, ResNet50, ViT-tiny).
4. Prediction Aggregation: Patch-wise predictions are aggregated for downstream slide-level tasks (e.g., drug response).

This design allows for scalable inference across full slides and compatibility with ImageNet-pretrained models while preserving subcellular resolution.

4.1 Data Processing




Figure 4.1: Data Processing Pipeline

Raw ExPath images are stored as 16-bit grayscale TIFFs across four fluorescence channels:

DAPI, TelC, CENPB, and WGA. These images are first preprocessed as follows:

- **Bit Depth Reduction:** We convert from 16-bit to 8-bit per channel, which reduces memory usage and file size by 50%. While overflows may occur during reduction, the effect is mitigated by robust downstream normalization.
- **Channel Normalization:** For each patch, we normalize pixel intensities per channel using min-max scaling across the slide. We also explore dataset-wide channel normalization to standardize global intensity statistics.
- **RGB Conversion:** To enable compatibility with pretrained RGB-based models, we convert the four grayscale channels into a three-channel pseudo-RGB format. Unlike simple one-to-one channel mappings, we overlay the DAPI signal uniformly across all three RGB channels to preserve nuclear morphology as a consistent background feature. The remaining markers (e.g., TelC, WGA, CENPB) are then selectively added to individual channels (R, G, or B) depending on the permutation. Various combinations are explored in our ablation studies to assess the impact of marker-channel assignments.

Most histopathology models operate on 224×224 or 256×256 patches. However, due to the $4\times$ linear expansion in ExPath, cells appear significantly larger than in conventional H&E slides at the same magnification. A 224×224 patch in ExPath typically contains only a few nuclei and lacks sufficient context for reliable classification.

To mitigate this, we tile the WSIs into 1024×1024 pixel patches, which provide roughly equivalent cellular coverage to 256×256 patches in H&E. This patch size strikes a balance between resolution, field-of-view, and GPU memory constraints. Patches are extracted without overlap to maximize coverage and computational efficiency.

This transformation allows us to directly use ImageNet-pretrained weights for initialization, improving convergence and reducing training time. It also allows us to streamline our ablation experiments.




Figure 4.2: Images after processing steps

4.2 Model Architectures and Training

We evaluate several commonly used architectures to identify the most effective model for classifying ExPath patches:

- ResNet34: A moderately deep convolutional neural network (CNN) with skip connections. Well-suited for medium-sized datasets and provides strong baselines for tissue classification tasks.
- ResNet50: A deeper variant that offers improved representation capacity. It is slower to train but often achieves better accuracy.
- EfficientNet-B0: A parameter-efficient CNN with compound scaling. However, in our setting, it consistently underperforms ResNet models likely due to the difference in image statistics between natural images and fluorescence WSIs.
- ViT-tiny: A transformer-based architecture with global self-attention. We use the ViT-tiny variant for its computational feasibility on 1024×1024 inputs. While it captures long-range dependencies, it does not outperform CNNs in this task, likely due to limited data size.

Model Initialization and Input Processing: All models are initialized with ImageNet-pretrained weights. For ViT, we use HuggingFace’s `vit_tiny_patch16_224` checkpoint, initially upsampling input patches to 224×224 to match the expected input shape. This upsampling is used only during the model selection phase to identify the most promising architecture. Once selected, the original patch resolution is retained, and the model is fine-tuned to adapt to the native input size. For ResNets and EfficientNet, patches are resized and normalized using the standard ImageNet statistics (mean and standard deviation).

Although pretrained on natural RGB images, ImageNet models have proven useful in medical imaging tasks when input channels are mapped appropriately. In our pipeline, we exploit this by converting ExPath grayscale channels into pseudo-RGB inputs and adapting intensity distributions via normalization. Empirically, pretrained ResNet models converge faster and outperform models trained from scratch on our ExPath dataset.

However, we note that this benefit is diminished when training directly on raw TIFFs or high-bit-depth formats. TIFF loading overhead, memory usage, and the lack of pretrained 4-channel models contribute to this performance gap. Our preprocessing pipeline thus serves as both a computational optimization and a domain adaptation step for model compatibility.

4.2.1 Training Details

- Loss Function: We use cross-entropy loss for both multi-class classification and binary response prediction tasks.
- Optimizer: Adam optimizer is used with an initial learning rate of 1e-4 and weight decay of 1e-5.
- Batch Size: Batch size is set to 16 for ResNet models and 8 for ViT due to GPU memory limits.
- Early Stopping: Training is monitored using validation accuracy with early stopping after 10 epochs of no improvement.

Chapter 5

Experiments Setup

To evaluate the utility of multiplexed DAPI-based ExPath images for classification and treatment response prediction, we design a sequence of experiments grounded in three core hypotheses:

- ExPath imaging provides superior morphological information compared to non-ExPath fluorescence microscopy.
- The combination of DAPI with TelC, CENPB, and WGA offers a biologically meaningful and discriminative feature space for classification.
- Our data composition enables generalization across tasks such as tissue classification, drug response prediction, and cross-cancer inference.

This section outlines how we validate these claims using controlled experiments and ablation studies across multiple models, tissue types, and evaluation metrics.

5.1 Experimental Tasks

We define three primary classification tasks:

5.1.1 ExPath vs. Simulated Non-ExPath Imaging

We test whether the physical resolution gains from expansion microscopy result in better model performance. To simulate non-ExPath images, we artificially downsample and blur the original images to match conventional microscopy resolution and quality. We train identical models on both versions and compare classification accuracy and visual separability.

5.1.2 Marker Contribution via Ablation Studies

To assess the contribution of individual markers (DAPI, TelC, CENPB, WGA), we conduct both channel ablation and combination experiments:

- Ablation Study: We train and evaluate models after removing one channel at a time from the 4-channel input. For instance, comparing DAPI+TelC+CENPB+WGA (full input) to DAPI+TelC+CENPB (without WGA) allows us to quantify WGA’s marginal contribution.

- **Subset Combinations:** We test all relevant 1-, 2-, and 3-channel combinations (e.g., DAPI only, DAPI+CENPB, TelC+WGA) to explore performance trade-offs and guide future panel design.

This allows us to measure not only the overall importance of multiplexing, but also the biological informativeness of each channel in isolation.

5.1.3 Classification Generalizability

We evaluate whether our composed dataset supports generalization across tasks:

- **Tissue Classification:** We train models to classify each patch into one of seven tissue classes: progressive disease, partial response, complete response, lymphocyte, squamous, smooth muscle, and empty.
- **Drug Response Classification:** We train models to predict the drug response label of a slide (complete, partial, or progressive) based on its aggregated patch predictions.
- **Cross-Cancer Generalization:** To test generalizability, we apply our trained MIBC models to lung carcinoma ExPath images. The goal is to determine whether features learned from the bladder tissue context retain predictive utility in a different cancer type.

5.2 Data Splits and Sampling

To ensure robust evaluation, we employ a strict patient-level split between training, validation, and test sets. No patches from the same WSI or patient are shared across sets. We use an 80/10/10 split unless otherwise noted.

To address class imbalance in tissue-type labels (e.g., fewer healthy or lymphocyte-rich patches), we implement:

- Downsampling of majority classes (e.g., cancer)
- Oversampling of minority classes using weighted sampling during training
- Patch Aggregation for drug response to preserve clinical labels at the slide level

5.3 Evaluation Metrics

We report both patch-wise and patient-wise metrics: **Patch-Level Metrics:**

- Accuracy, AUC
- Confusion matrices to visualize misclassifications

Patient-Level Metrics (for drug response):

- Majority vote and softmax averaging across patches
- ROC-AUC and balanced accuracy for each drug response class

Visual Metrics:

- UMAP plots of learned representations to assess clustering and separability

Chapter 6

Results

This section presents the outcomes of the experiments defined in the previous section. We evaluate our pipeline’s ability to (1) leverage the enhanced resolution of ExPath, (2) utilize the discriminative power of multiplexed nuclear and membrane markers, and (3) generalize across tissue classification and treatment prediction tasks. We present both quantitative metrics and qualitative visualizations.

6.1 ExPath Outperforms Simulated Non-ExPath Imaging

In our first experiment, we evaluated the effect of physical expansion on classification performance by comparing models trained on simulated non-ExPath images downsampled to 224×224 against those trained on full-resolution ExPath patches at 1024×1024 . While the downsampled non-ExPath models converged more rapidly, their performance plateaued at approximately 83% accuracy, suggesting a limitation in the morphological detail available at this resolution. In contrast, ExPath models trained on larger patches exhibited slower convergence but ultimately achieved close to 93.33% accuracy, demonstrating that the additional subcellular resolution afforded by expansion microscopy provides a measurable performance gain despite increased training complexity.

6.2 Marker Contribution Analysis via Ablation Studies

To quantify the contribution of individual channels in our multiplexed panel, we performed two sets of ablation experiments: (1) post-hoc evaluation ablations using previously trained models, and (2) training-time ablations in which models were retrained with systematically reduced channel inputs.

6.2.1 Post-hoc Evaluation Ablations

In the first experiment, we ablated channels during evaluation of previously trained models and summarized the effects using heatmaps across multiple aggregation baselines (mean, max, min, median). Among these, the median baseline produced the most reliable and stable results, so we




Figure 6.1: Accuracy Confusion Matrix for Model trained with ExPath

focus on it here. The x-axis of the heatmaps represents the channels ablated, while the y-axis shows the corresponding change in accuracy or loss. Across all settings, removal of the DAPI channel caused the most pronounced accuracy drop, with the entire left panel of the heatmap collapsing highlighting its essential role as the anchor channel.

When examining class-specific sensitivity, we observed heterogeneous dependencies on the markers. For complete response, performance was particularly sensitive to ablation of combinations of markers, suggesting that reliable identification of this class requires integrated nuclear and chromatin context. For progressive disease, accuracy dropped most sharply when DAPI was ablated, followed by combined removal of TelC and CENPB, with accuracy falling to 62%. This likely reflects the biological relevance of these markers: DAPI encodes general nuclear morphology, while TelC and CENPB provide information on telomeric and centromeric organization, both of which are tightly coupled to malignant transformation and genomic instability in cancer. In contrast, the marker with the least perturbation effect for complete response was WGA, where class performance dropped to 34% accuracy upon ablation. This is consistent with WGA primarily capturing membrane glycoproteins, which are less central to nuclear-level processes driving




Figure 6.2: Overall Post-hoc Ablation Results

class separation. Interestingly, for lymphocytes, accuracy was more scattered: ablating TelC, CENPB, or WGA reduced performance substantially (to 35%), while ablating DAPI alone still maintained 63% accuracy. Moreover, ablating CENPB alone had minimal effect, with lymphocyte classification retaining 93% accuracy. One potential explanation is that lymphocyte identity is strongly determined by cell shape and boundary delineation, features captured by WGA and to some extent TelC, rather than centromeric structure, which is less variable in non-malignant immune cells.

6.2.2 Training-Time Ablations

In the second experiment, we trained separate models on reduced channel sets to directly evaluate the effect of marker exclusion during learning. As expected, the all-channel model achieved the highest accuracy (93%), outperforming all other channel subsets. A DAPI-only model reached 85% accuracy, confirming DAPI’s role as a strong baseline but highlighting the value of multiplexing. Notably, even the weakest multi-channel combination excluding DAPI still outperformed the DAPI-only baseline, achieving 90% accuracy. This result underscores that while DAPI is indispensable, complementary markers (TelC, CENPB, WGA) encode sufficient discriminative information to elevate performance beyond what nuclear morphology alone provides.

Taken together, these experiments demonstrate that (1) DAPI remains the single most critical channel, anchoring both nuclear detail and overall performance, (2) TelC and CENPB provide




Figure 6.3: Per-class Post-hoc Ablation Results

strong complementary signals for cancer-related phenotypes, and (3) WGA contributes less to classification accuracy overall, but plays a subtle role in delineating non-cancerous classes such as lymphocytes.




Figure 6.4: Training Ablation Results

6.3 Classification Tasks

To evaluate the broader utility of multiplexed ExPath imaging, we assessed model performance across three levels of generalization: tissue classification, drug response prediction, and cross-cancer transfer. All experiments were conducted under strict patient-level splits to prevent data leakage, with models trained on slide-level labels and final predictions aggregated to the patient level.

6.3.1 Tissue Classification

We first trained models to distinguish among three major tissue classes: cancerous tissue, lymphocytes, and smooth muscle. Patient-level aggregation yielded an accuracy of 89.52% and an ROC-AUC of 0.95, underscoring the discriminative power of ExPath-enabled nuclear and sub-nuclear features. These results highlight that even with minimal staining and no stromal context, nuclear morphology enriched by TelC, CENPB, and WGA suffices for accurate tissue stratification.

6.3.2 Drug Response Prediction

We next examined the ability of the model to predict neoadjuvant chemotherapy response (complete, partial, or progressive disease) using aggregated slide-level predictions. Despite the challenge of weak supervision and heterogeneous patient responses, the model achieved 85.66%




Figure 6.5: Accuracy Confusion Matrix for Tissue Classification (3-way and 4-way)

accuracy and an ROC-AUC of 0.90 at the patient level. This suggests that multiplexed ExPath captures treatment-relevant morphological signals, offering a scalable imaging-based alternative to costly molecular profiling.




Figure 6.6: Accuracy Confusion Matrix for Drug Response Prediction

6.3.3 Cross-Cancer Generalization (MIBC → Lung)

Finally, we tested the generalizability of MIBC-trained models by applying them to lung carcinoma ExPath images. Here, the task was simplified to binary classification (cancerous vs. non-cancerous). Remarkably, the transferred model achieved 72% accuracy, despite no retraining on lung data. While performance was lower than in bladder-specific tasks, this result demonstrates that nuclear and chromatin features learned from MIBC are at least partially conserved across cancer types, pointing to the feasibility of training broader, cross-organ foundation models. To-

gether, these findings establish that multiplexed ExPath models not only achieve high accuracy in MIBC-specific classification but also extend to treatment prediction and cross-cancer inference, highlighting their potential as a generalizable platform for computational pathology.

6.4 UMAP Visualization of Patch Embeddings

To complement quantitative performance metrics, we examined the latent feature space of our trained model using UMAP dimensionality reduction. The resulting 2D embedding (Figure above) illustrates how patches from different tissue classes cluster in the learned representation space. We observe that certain classes, such as those corresponding to cancerous regions, form relatively dense and well-separated clusters, while others, particularly those involving lymphocyte-rich or smooth muscle regions, appear more intermixed with neighboring classes. This suggests that the model captures discriminative features that align with biological distinctions, but also highlights areas of overlap where morphology alone may not suffice for clean separation. Notably, the large contiguous cluster of yellow points (label 6 – progressive disease) indicates a dominant tissue class that the model consistently embeds in a compact region of feature space, reflecting strong intra-class consistency. Conversely, smaller peripheral clusters reflect minority classes and suggest heterogeneity in their representation. Overall, the UMAP plot supports the model’s ability to learn biologically meaningful separations while also revealing the boundaries where multiplexed information may need to be further leveraged to improve class distinction.

Summary of Key Findings

1. ExPath imaging enhances classification performance due to higher effective resolution.
2. Multiplexing with TelC, CENPB, and WGA improves both accuracy and interpretability.
3. Patient-level models trained on bladder tissue generalize to lung, suggesting potential for broader foundation modeling.




Figure 6.7: UMAP Visualization of Patch Embeddings

Chapter 7

Discussion and Limitations

Our experiments demonstrate that multiplexed DAPI-based imaging using expansion microscopy (ExPath) provides a powerful alternative to traditional histopathological approaches for tissue classification and drug response prediction. The combination of nuclear and sub-nuclear markers (DAPI, TelC, CENPB) with a membrane stain (WGA) enables models to learn discriminative and generalizable features from a minimal set of channels, offering strong performance even without cytoplasmic or stromal context. In this section, we discuss the broader implications of our findings, the challenges we encountered, and key limitations of our current dataset and pipeline.

7.1 Biological and Clinical Implications

The success of our pipeline in distinguishing fine-grained tissue classes and predicting chemotherapy response suggests that nuclear morphology and chromatin organization when captured at high resolution encode information that is both biologically meaningful and clinically actionable. Traditional H&E-based workflows, while more general-purpose, do not afford the subnuclear or telomeric resolution necessary to extract such features.

Our results also reinforce the idea that certain molecular signatures (e.g., telomere organization, centromeric alignment) are tightly coupled with treatment response phenotypes in muscle-invasive bladder cancer (MIBC). Moreover, we observe that models trained exclusively on ExPath patches from MIBC generalize well to lung carcinoma tissue, implying that these nuclear features may be conserved across cancer types, an encouraging sign for building broader foundation models.

7.2 Dataset Limitations

The dataset exhibits significant imbalance in tissue-type labels: cancer patches are overrepresented, while lymphocytes and smooth muscle are underrepresented. We addressed this via downsampling and oversampling during training; however, imbalance may still influence feature learning and model calibration. This effect is particularly evident in misclassifications between biologically similar or transition-state classes (e.g., partial vs. complete response).

The dataset was also only weakly-supervised, whole slides were attributed a single label (responsive, partially responsive, on-responsive).

Additionally, our current dataset size is insufficient to train a full DAPI-based foundation model in a self-supervised manner. All models in this study were trained in a supervised setting with explicit labels, limiting our ability to explore transfer learning or contrastive pretraining paradigms. More slides from diverse cancer types, multiple organs, and varying marker panels will be necessary to move toward a general-purpose feature extractor.

7.3 Training and Evaluation Constraints

While slide-level labels were used for drug response prediction, some models were trained patch-wise without enforcing patient- or slide-level disjointness. This could lead to leakage of patient-specific features and artificially inflate accuracy. However, we mitigate this by reporting only slide-level metrics for treatment prediction tasks and by using patient-level splitting for evaluation.

To enable faster training and reduce memory usage, we convert 16-bit TIFFs to 8-bit images. This can cause information loss, especially in dimmer channels or when signal intensities are concentrated in narrow dynamic ranges. Although normalization partially mitigates this, overflow and saturation effects remain a concern. In the future, we propose exploring hybrid 12-bit representations or intensity-aware compression to preserve information while maintaining compatibility with RGB-based pretrained models.

7.4 Pipeline-Level Limitations

The TelC, CENPB, and WGA markers used in this study were chosen based on literature knowledge and availability, but may not be optimal for all cancer types or tissue states. It is possible that other combinations (e.g., epigenetic markers, cell-cycle indicators, immune-specific stains) could outperform our current set. Our ablation results are a step toward understanding marker utility, but more systematic marker selection strategies (e.g., marker attribution, information gain) are warranted.

While effective, our reliance on pseudo-RGB conversion and ImageNet-pretrained backbones is a compromise. These models are not optimized for bioimaging and may fail to fully exploit fluorescence intensity distributions, spatial correlation across channels, or nuclear substructures. Future work should explore bio-specific pretraining or architectures designed for multi-channel imaging.

7.5 Interpretability and Trustworthiness

Although saliency maps and Grad-CAM overlays offer some interpretability, deeper analysis is needed to ensure trust in model decisions especially in clinical contexts. For example, attention-based models could be paired with domain-specific priors (e.g., nuclear morphology scoring systems) to ensure that models attend to biologically meaningful regions.

In summary, our findings provide a compelling proof-of-concept that multiplexed DAPI-based imaging under expansion microscopy can support accurate and generalizable computational pathology. Nonetheless, this work is an early step in a broader vision toward lightweight, interpretable, and biologically grounded imaging models especially in contexts where tissue is limited and multiplexing is preferred over multi-modal pipelines.

Chapter 8

Future Work

While this work provides a strong proof-of-concept demonstrating the potential of multiplexed DAPI-based ExPath imaging for computational pathology, several avenues remain for further refinement, expansion, and clinical translation. A first priority concerns improving image fidelity and preserving dynamic range. Our current preprocessing pipeline reduces all 16-bit TIFF channels to 8-bit images to ensure compatibility with pretrained RGB-based backbones. Although this simplification facilitates rapid prototyping, it inevitably discards information through dynamic range compression and channel overflow. Future iterations of the pipeline should therefore explore hybrid representations such as 12-bit or adaptive bit-depth encodings that strike a balance between information preservation and computational feasibility. Complementary normalization schemes, including histogram equalization or log-transformed intensity scaling, may further enhance weak signals in low-intensity channels such as WGA or TelC, thereby strengthening the biological informativeness of the data. A second direction is the development of a DAPI-based foundation model. This study has shown that a carefully selected DAPI + multiplex panel, when paired with ExPath, can serve as a robust supervised testbed. However, the long-term vision is to extend beyond task-specific supervised models and instead establish a general-purpose feature extractor trained through large-scale self-supervised learning (SSL). Such a model would learn representations of nuclear morphology, chromatin texture, and membrane organization across diverse cancer types, offering a modular foundation that can later be fine-tuned for downstream tasks such as tissue classification, segmentation, and drug response prediction. The advantage of this approach is its scalability: unlike H&E, which is difficult to multiplex, DAPI-based imaging can flexibly integrate additional fluorescence markers depending on the biological question, making it well suited for foundation-level modeling. The robustness and scalability of the computational pipeline also require attention. Current preprocessing assumes standardized, high-quality ExPath TIFF images, yet real-world deployments must accommodate variability in staining quality, illumination, imaging resolution, and sample preparation. Future work should therefore include the development of adaptive normalization strategies capable of automatically detecting and correcting for artifacts such as uneven illumination, staining variability, or blur. Automated quality control pipelines could further ensure that problematic slides or patches are flagged prior to model training, improving reliability in both research and clinical workflows. In addition, it will be important to explore alternative computational approaches beyond standard CNNs and transformer baselines. Architectures tailored for multi-channel biomedical images

such as multi-branch CNNs, channel-attention networks, or hybrid vision transformer models may prove more effective at exploiting correlations across nuclear and membrane markers than ImageNet-pretrained RGB backbones. Self-distillation frameworks and graph-based methods that encode cell-cell interactions also hold promise, particularly in capturing spatial relationships between nuclei in tissue microenvironments. Future computational directions could thus combine morphology-aware architectures with biologically informed priors to maximize both accuracy and interpretability.

A key limitation of this study is dataset scale and diversity. The current dataset, while sufficient for supervised proof-of-concept modeling, lacks the size, heterogeneity, and annotation depth needed to support large-scale pretraining or robust evaluation across cancer subtypes. An ideal dataset for this line of work would comprise thousands of ExPath-prepared whole-slide images spanning multiple cancer types and tissue conditions, annotated at both the patch and slide level for diagnosis, treatment response, and molecular subtyping. Such a dataset would ideally include balanced representation across tissue classes (tumor, immune cells, smooth muscle, etc.), curated clinical metadata (treatment regimens, survival outcomes), and longitudinal samples capturing treatment progression. Crucially, it would also incorporate diverse staining panels, enabling systematic evaluation of which multiplex markers generalize best across cancers. The assembly of such a dataset would allow for the training of a true DAPI-first foundation model capable of generalization across both biological and clinical domains.

In summary, advancing this work will require parallel progress on imaging fidelity, computational methods, and data scale. Improvements in preprocessing and representation will preserve critical biological signals, while the development of a DAPI-based foundation model will enable general-purpose representation learning. Enhanced robustness and quality control will make the pipeline clinically viable, and the creation of a large-scale, diverse ExPath dataset will unlock the potential for scalable, generalizable models. Together, these steps chart a path from the current proof-of-concept toward a clinically deployable ExPath-driven computational pathology framework.

Chapter 9

Conclusion

In this work, we demonstrated that multiplexed Expansion Microscopy (ExPath) combined with a biologically informed four-channel panel (DAPI, TelC, CENPB, WGA) provides a powerful and scalable alternative to traditional H&E-based computational pathology. By physically expanding tissue to nanoscale resolution and anchoring analysis on DAPI, we showed that it is possible to extract subnuclear and chromatin-level features that are both discriminative and clinically meaningful. Our preprocessing pipeline transformed high-bit-depth, multi-channel whole-slide images into pseudo-RGB patches suitable for ImageNet-pretrained architectures, enabling efficient training and evaluation across multiple classification tasks.

Through systematic experiments, we established three core findings. First, ExPath imaging delivers a measurable performance edge over simulated non-ExPath baselines, with models achieving nearly 90% accuracy in tissue classification. Second, ablation studies revealed that while DAPI is indispensable, complementary markers such as TelC and CENPB significantly enhance performance, particularly for cancer-related phenotypes, whereas WGA plays a more limited but non-negligible role in delineating non-cancerous cell types. Third, our models generalized across tasks and domains: achieving 89.5% accuracy and 0.95 AUC in three-way tissue classification, 85.7% accuracy and 0.90 AUC in drug response prediction, and 72% accuracy when transferred to lung carcinoma data, demonstrating cross-cancer portability of learned features.

Together, these results support the central hypothesis that nuclear morphology, when captured at high resolution and enriched with targeted fluorescent markers, contains sufficient information to match and even surpass traditional H&E workflows for diagnostic and prognostic modeling. Moreover, our findings establish a proof-of-concept for the long-term vision of a DAPI-first foundation model in pathology: a general-purpose, scalable representation trained on ExPath-enabled multiplexed imaging, capable of adapting across tissue types, cancer contexts, and downstream clinical tasks.

While this study is constrained by dataset size, class imbalance, and reliance on RGB-adapted pretrained backbones, it provides a clear framework for scaling. Future work should focus on assembling larger and more diverse datasets, improving image fidelity through adaptive bit-depth and normalization strategies, and exploring architectures optimized for multi-channel biomedical imaging. By advancing along these directions, multiplexed ExPath has the potential to transform computational pathology into a high-resolution, biologically grounded, and clinically impactful

discipline.

Bibliography

- [1] Florian Aeffner, Mark D. Zarella, Andrew M. Buchan, et al. Introduction to digital image analysis in whole slide imaging: A white paper from the digital pathology association. *Journal of Pathology Informatics*, 10:9, 2019. doi: 10.4103/jpi.jpi_82_18.
- [2] Marek Babjuk, Eva Comperat, Richard Zigeuner, et al. Eau guidelines on non-muscle-invasive bladder cancer (tat1 and cis) – 2022 update. *European Urology*, 81(1):75–94, 2022. doi: 10.1016/j.eururo.2021.08.010.
- [3] Evangelos Briasoulis, Sofia Rontogianni, Dimitrios Dafni, Nikolaos Baziotis, Michael M. Konstadoulakis, Anastasia Kondi-Pafitis, Konstantinos G. Pavlakis, Rena Trigidou, and Georgios Peros. Anti-centromere antibodies heralding small cell lung carcinoma: a case report. *Lung Cancer*, 60(2):302–306, 2008. doi: 10.1016/j.lungcan.2007.09.014.
- [4] Gabriele Campanella, Matthew G. Hanna, Luke Geneslaw, Allen Miraflor, Vitor Werneck Krauss Silva, Klaus J. Busam, Edi Brogi, Victor E. Reuter, David S. Klimstra, and Thomas J. Fuchs. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature Medicine*, 25(8):1301–1309, 2019. doi: 10.1038/s41591-019-0508-1.
- [5] Fei Chen, Paul W. Tillberg, and Edward S. Boyden. Expansion microscopy. *Science*, 347(6221):543–548, 2015. doi: 10.1126/science.1260088.
- [6] Nicolas Coudray, Paolo S. Ocampo, Theodore Sakellaropoulos, et al. Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nature Medicine*, 24(10):1559–1567, 2018. doi: 10.1038/s41591-018-0177-5.
- [7] Andrew H. Fischer, Kenneth A. Jacobson, Jack Rose, and Rolf Zeller. Hematoxylin and eosin staining of tissue and cell sections. *Cold Spring Harbor Protocols*, 2008: pdb.prot4986, 2008. doi: 10.1101/pdb.prot4986.
- [8] H. Barton Grossman, Ralph B. Natale, Cheryl M. Tangen, et al. Neoadjuvant chemotherapy plus cystectomy compared with cystectomy alone for locally advanced bladder cancer. *New England Journal of Medicine*, 349(9):859–866, 2003. doi: 10.1056/NEJMoa022148.
- [9] Bing He, Ludvig Bergenstråhle, Linnea Stenbeck, Abubakar Abid, Alma Andersson, Åke Borg, Jonas Maaskola, Joakim Lundeberg, James Zou, and Carolina Wählby. Integrating spatial gene expression and breast tumor morphology via deep learning. *Nature Biomedical Engineering*, 4:827–834, 2020. doi: 10.1038/s41551-020-0578-x.
- [10] Jan Kapuscinski. Dapi: A dna-specific fluorescent probe. *Biotechnic & Histochemistry*, 70(5):220–233, 1995. doi: 10.3109/10520299509108199.

- [11] Ming Y. Lu, Drew F. K. Williamson, Tiffany Y. Chen, Richard J. Chen, Matteo Barbieri, and Faisal Mahmood. Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature Biomedical Engineering*, 5:555–570, 2021. doi: 10.1038/s41551-020-00682-w.
- [12] Amirhossein Mahbod, Ilkay Öksüz, Chiheb Trabelsi, Olivier Bousselham, Chandramouli S. Ávila, Christos Kampel, Andreas Langs, Georg Dorffner, Isabella Ellinger, Michal Kozubek, Robert Sablatnig, and Pedro H. G. Ávila. Impact of bit-depth reduction on fluorescence microscopy nuclei segmentation using deep learning. *Diagnostics*, 11(6):967, 2021. doi: 10.3390/diagnostics11060967.
- [13] Stephen L. McGovern, Yuan Qi, Lajos Pusztai, W. Fraser Symmans, and Thomas A. Buchholz. Centromere protein-a, an essential centromere protein, is a prognostic marker for relapse in estrogen receptor-positive breast cancer. *Breast Cancer Research*, 14(2):R72, 2012. doi: 10.1186/bcr3161.
- [14] A. Gordon Robertson and The Cancer Genome Atlas Research Network. Comprehensive molecular characterization of muscle-invasive bladder cancer. *Cell*, 171(3):540–556.e25, 2017. doi: 10.1016/j.cell.2017.09.007.
- [15] Uwe Schmidt, Martin Weigert, Coleman Broaddus, and Gene Myers. Cell detection with star-convex polygons. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 11071 of *Lecture Notes in Computer Science*, pages 265–273. Springer, 2018. doi: 10.1007/978-3-030-00934-2_30.
- [16] Jerry W. Shay and Woodring E. Wright. Telomeres and telomerase: three decades of progress. *Nature Reviews Genetics*, 20:299–309, 2019. doi: 10.1038/s41576-019-0099-1.
- [17] Rebecca L. Siegel, Angela N. Giaquinto, and Ahmedin Jemal. Cancer statistics, 2024. *CA: A Cancer Journal for Clinicians*, 74(1):12–49, 2024. doi: 10.3322/caac.21820.
- [18] K. Sirinukunwattana, S. E. A. Raza, Y. W. Tsang, D. R. J. Snead, I. A. Cree, and N. M. Rajpoot. Locality sensitive deep learning for detection of nuclei in routine colon cancer histology images. *IEEE Transactions on Medical Imaging*, 35(5):1196–1206, 2016. doi: 10.1109/TMI.2016.2527906.
- [19] James P. Stein, Gary Lieskovsky, Robert Cote, et al. Radical cystectomy in the treatment of invasive bladder cancer: Long-term results in 1,054 patients. *Journal of Clinical Oncology*, 19(3):666–675, 2001. doi: 10.1200/JCO.2001.19.3.666.
- [20] Carsen Stringer, Tim Wang, Michalis Michaelos, and Marius Pachitariu. Cellpose: a generalist algorithm for cellular segmentation. *Nature Methods*, 18:100–106, 2021. doi: 10.1038/s41592-020-01018-x.
- [21] Paul W. Tillberg, Fei Chen, Kiryl D. Piatkevich, Yongxin Zhao, Chih-Chieh Jay Yu, Brian P. English, Linyi Gao, Anthony Martorell, Ho-Jun Suk, Fumiaki Yoshida, Ellen M. DeGennaro, Douglas H. Roossien, Guanyu Gong, Uthpala Seneviratne, Steven R. Tannenbaum, Robert Desimone, Dawen Cai, and Edward S. Boyden. Protein-retention expansion microscopy of cells and tissues labeled using standard fluorescent proteins and antibodies. *Nature Biotechnology*, 34(9):987–992, 2016. doi: 10.1038/nbt.3625.

- [22] Karin Wallander, Jessada Thutkawkorapin, Ellika Sahlin, Annika Lindblom, and Kristina Lagerstedt-Robinson. Massive parallel sequencing in a family with rectal cancer. *Hereditary Cancer in Clinical Practice*, 19:23, 2021. doi: 10.1186/s13053-021-00181-2. Reports CENP-B p.(Glu438Lys) as a candidate variant.
- [23] Yongxin Zhao, Octavian Bucur, Humayun Irshad, Fei Chen, Astrid Weins, Andreea L. Stancu, Eun-Young Oh, Marcello DiStasio, Vanda Torous, Benjamin Glass, Isaac E. Stillman, Stuart J. Schnitt, Andrew H. Beck, and Edward S. Boyden. Nanoscale imaging of clinical specimens using pathology-optimized expansion microscopy. *Nature Biotechnology*, 35(8):757–764, 2017. doi: 10.1038/nbt.3892.
- [24] Daniele Zink, Andrew H. Fischer, and Jeffrey A. Nickerson. Nuclear structure in cancer cells. *Nature Reviews Cancer*, 4(9):677–687, 2004. doi: 10.1038/nrc1430.