

Learning cell representations in temporal and 3D contexts

Kristofer delas Peñas

Wolfson College
University of Oxford

Doctor of Philosophy

January 2023

Abstract

Cell morphology and its changes under different circumstances is one of the primary ways by which we can understand biology. Computational tools for characterization and analysis, therefore, play a critical role in advancing studies involving cell morphology.

In this thesis, I explored the use of representation learning and self-supervised methods to analyze nuclear texture in fluorescence imaging across different contexts and scales. To analyze the cell cycle using 2D temporal imaging data, as well as DNA damage in 3D imaging data, I employed a simple model based on the VAE-GAN architecture. Through the VAE-GAN model, I constructed manifolds in which the latent representations of the data can be grouped and clustered based on textural similarities without the need for exhaustive training annotations. I used these representations, as well as manually engineered features, to perform various analyses both at the single cell and tissue levels.

The application on the cell cycle data revealed that common tasks such as cell cycle staging and cell cycle time estimation can be done even with minimal fluorescence information and user annotation. On the other hand, the texture classes derived to characterize DNA damage in 3D histology images unveiled differences between control and treated tissue regions. Lastly, by aggregating cell-level information to characterize local cell neighborhoods, interactions between DNA-damaged cells and immune cells can be quantified and some tissue microstructures can be identified.

The results presented in this thesis demonstrated the utility of the representations learned through my approach in supporting biological inquiries involving temporal and 3D spatial data. The quantitative measurements computed using the presented methods have the potential to aid not only similar experiments on the cell cycle and DNA damage but also in exploratory studies in 3D histology.

Learning cell representations in temporal and 3D contexts



Kristofer delas Peñas

Wolfson College

University of Oxford

Doctor of Philosophy

January 2023

To my family and Daniel
for picking me up when I was at my lowest.

And to Rigoberto Tiglao
for calling me an idiot.
I hope this proves you wrong.

Acknowledgements

I want to share the joy of completing this DPhil to the people who supported me in this journey.

Big thanks to my supervisors for the guidance throughout this project. Same goes with my collaborators who had generously shared not only datasets but also insights that ultimately led to completion of this DPhil. Thanks as well to all the PhD students and postdocs in the lab who have offered helpful critiques and suggestions.

I also thank my friends and family in the Philippines for entertaining my video calls even in the oddest of hours because of the time difference. Your stories and you listening to mine had been the cornerstone of my sanity during the pandemic.

Thank you to the members of the Oxford Philippines Society who had become my family away from home.

Thanks to all the friends I made along the way during my stay in Oxford.

Lastly, I thank Daniel for the love and support.

Abstract

Cell morphology and its changes under different circumstances is one of the primary ways by which we can understand biology. Computational tools for characterization and analysis, therefore, play a critical role in advancing studies involving cell morphology.

In this thesis, I explored the use of representation learning and self-supervised methods to analyze nuclear texture in fluorescence imaging across different contexts and scales. To analyze the cell cycle using 2D temporal imaging data, as well as DNA damage in 3D imaging data, I employed a simple model based on the VAE-GAN architecture. Through the VAE-GAN model, I constructed manifolds in which the latent representations of the data can be grouped and clustered based on textural similarities without the need for exhaustive training annotations. I used these representations, as well as manually engineered features, to perform various analyses both at the single cell and tissue levels.

The application on the cell cycle data revealed that common tasks such as cell cycle staging and cell cycle time estimation can be done even with minimal fluorescence information and user annotation. On the other hand, the texture classes derived to characterize DNA damage in 3D histology images unveiled differences between control and treated tissue regions. Lastly, by aggregating cell-level information to characterize local cell neighborhoods, interactions between DNA-damaged cells and immune cells can be quantified and some tissue microstructures can be identified.

The results presented in this thesis demonstrated the utility of the representations learned through my approach in supporting biological inquiries involving temporal and 3D spatial data. The quantitative measurements computed using the presented

methods have the potential to aid not only similar experiments on the cell cycle and DNA damage but also in exploratory studies in 3D histology.

Statement of Originality

This thesis is submitted to the University of Oxford in partial fulfilment of the requirements of the degree of Doctor of Philosophy. This thesis contains no material which has been previously submitted for a degree or diploma at this University or any other institution. This research represents my own original work towards this research degree.

Contents

List of Figures	x
List of Tables	xii
Glossary	xiii
1 Introduction	17
1.1 Questions, aims, and outline	21
2 Methods	24
2.1 Convolutional Neural Networks	26
2.2 Representation Learning	27
2.3 Generative Modelling	28
2.3.1 Variational autoencoders	29
2.3.2 Generative adversarial networks	32
2.4 Segmentation	33
2.4.1 Semantic Segmentation	33
2.4.2 Instance Segmentation	37
2.5 Projections and Clustering	40
2.5.1 Uniform Manifold Approximation and Projection	41
2.5.2 Gaussian Mixture Models	42
2.6 Shape Analysis	43
2.6.1 Landmarks	44
2.6.2 Parametric descriptors	44
2.6.3 Deformation-based representations	47
2.7 Texture Analysis	47

2.7.1	Statistical methods	47
2.7.2	Model-based methods	48
2.7.3	Transformations	49
2.8	Summary	50
3	Learning a Morphological Manifold of the Cell Life Cycle	51
3.1	Cell Cycle	52
3.2	Methods	59
3.2.1	VAE-GAN	59
3.2.2	Temporal Cycle Consistency	61
3.3	Experiment Setup	64
3.3.1	Dataset	64
3.3.2	Training and Evaluation	65
3.4	Analyses and Discussion	66
3.4.1	Classifying visually distinct clusters	66
3.4.2	Identifying Key Cell Cycle Stages from Single Cell Trajectories	76
3.4.3	Quantifying Cycle Times and Extracting Population Statistics	78
3.5	Summary	80
4	Phenotyping Cells in 3D 4T1 Histology	82
4.1	Object Representations in 3D	84
4.1.1	Voxels	84
4.1.2	Surfaces and meshes	85
4.2	4T1 Data	85
4.3	3D Nuclei Segmentation	86
4.4	Profiling DNA Damage in Tumor Tissues	89
4.4.1	Analysis Pipeline	92
4.4.2	Building cell profiles of DNA damage	95
4.4.3	Pseudo-class labels	98
4.4.4	Comparing cell profiles in control and treated tissues	98
4.5	Revisiting the Cell Cycle	102

4.6	Summary	106
5	Characterizing tissue regions and microstructures	108
5.1	Marker readout of γ H2AX and CD8 markers	110
5.2	Enumerating cells in thick 4T1 tissues	112
5.3	γ H2AX-CD8 proximity analysis	115
5.4	Comparing 2D and 3D tissue analysis	117
5.4.1	2D Histology	118
5.4.2	3D Histology	119
5.4.3	Extracting thin slices from 3D 4T1 tissue data	123
5.4.4	Proximity analysis in 2D and 3D	124
5.5	Segmenting microstructures	125
5.5.1	Using orientation of cells in local neighborhoods to infer structure	126
5.5.2	Recognizing vascular regions and tears in tissues	127
5.6	Summary	130
6	Discussion	131
	References	137

List of Figures

2.1	An autoencoder	29
2.2	U-Net architecture	36
2.3	Stardist	39
2.4	Cellpose segmentation	41
2.5	Spherical Harmonics	46
3.1	The cell cycle	53
3.2	PCNA fluorescence during the cell cycle stages	54
3.3	Reconstruction using the encoder-decoder setup	62
3.4	VAE-GAN model for cell morphology manifold	63
3.5	Temporal Cycle Consistency	64
3.6	Clustering on encodings from Resnet, VGG	67
3.7	Clustering on encodings from DeepCycle	68
3.8	Clustering on encodings from the autoencoder	69
3.9	VAE-GAN reconstructions	70
3.10	Clustering on encodings from VAE-GAN	71
3.11	Trajectory plots of individual cells as they undergo the cell life cycle	72
3.12	The VAE-GAN-TCC manifold.	73
3.13	Initial clustering using K-Means	73
3.14	Clusters formed by GMM	75
3.15	UMAP1 vs time plot reveals some key cell cycle stages	77
3.16	UMAP2 vs time plot reveals oscillating patterns in the cell cycle	79
3.17	Estimated cycle time distribution	80
4.1	3D image volume of a 4T1 tissue section	87

4.2	3D nuclei segmentation	94
4.3	VAE-GAN model for DNA damage morphology manifold	97
4.4	DNA damage clusters from GLCM and VAE-GAN	100
4.5	γ H2AX profiles	101
4.6	Representative cell images for each of the VAE-GAN DNA damage clusters.	101
4.7	Representative cell images for each of the “cell cycle-inspired” texture classes derived from DAPI fluorescence.	104
4.8	Cell population in our 4T1 data grouped by the derived “cell cycle-inspired” DAPI classes	104
4.9	γ H2AX+ cell population in our 4T1 data grouped by the derived “cell cycle-inspired” DAPI classes	105
4.10	Joint profiling of cells in 4T1 tissue using texture classes derived from γ H2AX (DDR) and DAPI (cell cycle).	106
5.1	Depth-wise image contrast stretching	112
5.2	Readout of CD8 and γ H2AX markers in control and treated tissue regions	113
5.3	Enumerating γ H2AX+ and CD8+ cells	114
5.4	Proximity analysis of γ H2AX+ and CD8+ cells	116
5.5	Conventional pathology versus non-destructive 3D pathology	119
5.6	3D pathology can outperform 2D pathology	120
5.7	Elongated nuclei	121
5.8	Virtual sectioning of 4T1 image volume into thin 2D sections	124
5.9	Comparing proximity analysis on 2D (thin) slices with 3D image volumes	125
5.10	Orientation map of cells in a 3D 4T1 histology volume	128
5.11	Similarity scores for patches based on nuclei orientation	129
5.12	Grouping of nuclei based on orientation and locality segments holes in tissues	129

List of Tables

3.1	Confusion matrix for cell cycle stage classification	74
3.2	Confusion matrix for cell cycle stage classification with G1 and G2 constraint	78
4.1	Precision and recall of 3D nuclei segmentation	89
4.2	Segmentation performance on selected tiles of tissue	90

Glossary

2D, 3D	Two- or three-dimensional, referring in this thesis to spatial dimensions in an image.
DNA	Deoxyribonucleic acid
VAE	Variational autoencoder
GAN	Generative adversarial networkl
V1	Primary visual cortex
MR, MRI . . .	Magnetic resonance, magnetic resonance imaging
4T1	A breast cancer cell line
CD8, CD4 . .	Cluster of differentiation 8, cluster of differentiation 4
γH2AX	Gamma H2AX, phosphorylated histone H2AX protein
CNN	Convolutional neural network
BSD	Berkeley Segmentation Dataset
U-Net	An encoder-decoder segmentation network
KLD	KullbackLeibler divergence
T1-WI	T1-weighted magnetic resonance image
FLAIR	Fluid-attenuated inversion recovery
ReLU	Rectified linear unit
CT	Computed tomography
StarDist	A segmentation network based on star-convex polygons
Mask R-CNN	A segmentation network based on region proposals (object bounding boxes)

NMS	Non-maximum suppression
ResNet	Residual network, a CNN architecture
Stardist3D	A variant of Stardist for 3D image volumes
U-Net3D	A variant of U-Net for 3D image volumes
IFT	Inverse Fourier transform
SARS-CoV-2	Severe acute respiratory syndrome coronavirus 2
Cellpose	A spatial gradient-based segmentation network
UMAP	Uniform Manifold Approximation and Projection
t-SNE	t-distributed stochastic neighbor embedding
GMM	Gaussian mixture model
EM	Expectation-Maximization
LR	Left-right
IVD	Intervertebral discs
GLCM	Gray level co-occurrence matrix
RNA, mRNA	Ribonucleic acid, messenger RNA
G,G0,G1,G2	Growth stage , growth stage 0, growth stage 1, growth stage 2
S	DNA replication stage
M	Mitosis stage
PCNA	Proliferating cell nuclear antigen
H2B	Histone H2B, a nuclear protein
mCherry	A red fluorescent protein
EGFP	Enhanced green fluorescent protein
SVM	Support vector machines
HMM	Hidden Markov model
FUCCI	Fluorescent Ubiquitination-based Cell Cycle Indicator

Cdt1	Chromatin licensing and DNA replication factor 1
mKO2	An orange fluorescent protein
mAG1	DNA-3-methyladenine glycosylase 1
MSE	Mean squared error
MNIST	A dataset of images of handwritten single digits 0-9, Modified National Institute of Standards and Technology
TCC	Temporal cycle consistency
ADAM	Adaptive moment estimation, an optimization method
VGG	A convolutional neural network architecture
RMSE	Root mean squared error
DAPI	4',6-diamidino-2-phenylindole, stain for DNA-rich cell regions
NIH	National Institutes of Health
MMTV	Murine mammary tumor virus
BALB/C	An albino, laboratory-bred strain of the house mouse
B6C3H,C3H	A general purpose mouse strain used in a wide variety of research areas including cancer studies
L-NAME	N(ω)-nitro-L-arginine methyl ester
NOS	Nitric oxide synthase
LaminB	protein anchored in the inner surface of the nuclear membrane
DDR	DNA damage response
PARP	Poly-ADP ribose polymerase
DSB	Double strand break
CTL	Cytotoxic T cell
FAS	A death receptor on the surface of cells that leads to apoptosis when bound to a ligand

PD-1		Programmed death 1
H&E		Hematoxylin and eosin
EMT		Epithelial-mesenchymal transition
HPC		High performance computing
EM		Electron microscopy
RI		Refractive index

1

Introduction

Contents

1.1 Questions, aims, and outline	21
--	----

Recognizing patterns is how we navigate the world. The primary function of all animal brains, including the human brain, is to perceive, encode, and integrate information from sensory stimuli to consequently elicit a response [1]. Through this process, a human who first experiences encountering an object, say a dog or their parent's face, encodes the visual stimuli. Later on, when faced with another object, not necessarily the exact same objects (perhaps a different dog, a photo of the parent, or a histology sample), the human brain can pull up the encoded information, make comparisons, and draw conclusions (*recognition*).

Central to recognizing patterns is the encoding of stimuli into meaningful

representations. While our understanding of such representations in the brain—how they are formed and stored—is still lacking, several theories exist attempting to explain the recognition process [1, 2]. Through machine learning algorithms, we attempt to create a parallel to human representation learning and pattern recognition. Several methods exist that are loosely based on theories on human perception and cognition. One such method is the Gabor filter. Gabor filters, named after Dennis Gabor, are linear filters used in texture analysis. Their applications are widely used in modelling the behavior of primary visual cortex (V1) simple cells as the filters exhibit similar responses to stimuli [3, 4]. Most commonly, a multi-scale pyramid approach using Gabor filters is used to mimic the invariance to orientation and scale of V1 simple cells [5]. Gabor filters were used for different purposes in past works. In [6] and [7], representations using Gabor filters were used to recognize face and characters from images. Detection of cars in roads is another application of Gabor filters [8]. In the biomedical domain, they were applied to segment intervertebral discs in MR images [9] and detect blood vessels in retinal fundus images [10]. However, with lack of other solid cognitive science theories to build upon, learning representations and recognizing patterns through machine learning over the years favored solid foundations in mathematics and statistics.

Prior to the widespread application of deep learning, features were manually engineered and extracted from data. Engineering such features is key to meaningful representations of data points and machine learning algorithms are ultimately limited by the quality of these features. In image analysis, features such as shape are quantified through the extraction of edges and boundaries, and some geometrical and/or statistical measurements. For texture, signal decomposition techniques and comparisons relative to a finite texture filter bank are common ways for manual characterization. While manually engineered features have proven to be pivotal in studies in different domains, higher performance scores in tasks were demonstrated using deep learning approaches that take in raw input data and do not require

the process of manual feature engineering.

Representation learning is an important part of machine learning. Its crucial role has become more prominent with the wide usage of deep learning architectures. Moreover, in conjunction with self-supervised approaches, learning representations from data has been shown to unravel similarities and dissimilarities in data points with very little to no injection of prior knowledge. This plays well with scientific experiments that are data-driven as biases can be minimized and the representations learned will not be limited by set assumptions from prior knowledge and can be tools for novel discoveries.

Constructing representation and finding distinctive patterns are also indispensable to many biomedical studies. The analysis of cell morphology and its dynamics is an important task in biology and is currently an active area of research developing techniques for different imaging modalities and problem domains [11]. The characterization of morphology can give clues on cell functions and responses [12–14]. In cancer, for example, alterations in the morphology of the nucleus were established as characteristic of some tumor types [15–17]. Researchers in this area look at the morphological changes in the cells induced by drugs, searching for ones that may revert the alterations in morphology due to cancer. On the other hand, analysis of the morphological variations during the cell cycle lead to our renewed understanding of cell death [18]. In pathology, diagnosis of certain diseases such as cancer and fibrosis is driven by morphological analysis of regions and landmarks such as nuclei and vascular and scar tissues.

In the automation of the analysis of cell morphology, measurements on nuclear shape, patterns in the cytoplasm, localization of the organelles, and cell wall structure are extracted [11]. Traditionally, this is done using image processing pipelines for microscopy in order to support and to provide evidence for cell research such as in discriminating between phenotypes. More recently, as I have previously alluded to, the application of deep learning methods to this task proved to be

successful [19, 20] and freed us from manual engineering of features. In a previous work [21], I developed approaches to quantify the axial depth and trajectory trend of vesicles as part of computational tools for the analysis of protein movement in cells. These approaches involved a convolutional network to denoise the microscopy data and another network to predict the axial motion of vesicles. Further to the wider usage of deep learning, many attempts have been done to extend the process, requiring little to no *a priori* knowledge, to look at data distributions, patterns, and clusters in a learned latent space [22].

As cells constantly change and adapt to the environment, the analysis of cell morphology must also consider the changes and transitions in morphology. More than the spatial distribution by which cell morphology can be represented, the temporal information on morphological transitions may describe some dynamic cellular processes, such as growth and differentiation. Building a model to look into these dynamic cellular processes can push forward our understanding of the mechanisms behind these processes.

In recent years, we have also seen the emergence of 3D imaging techniques for microscopy and histology. While the usual analyses on 2D images of cells and tissues have given evidence to many biological hypotheses, 2D imaging cannot capture the entirety of the cell, resulting in loss of information that might be useful for cell characterizations such as volume, orientation, and 3D spatial distribution of subcellular structures. Extending models and analyses from 2D and 3D, however, is not trivial as the additional dimension requires additional computational power, space requirements, and data registration.

In these various contexts I place the development of the analyses pipelines in this thesis. I will explore in this thesis several applications of learning representations of single cells and tissues. By building tools for these applications, I address gaps that I have identified. In a well-studied phenomenon like the cell cycle, I addressed a methodological gap posed by the need to analyze data with less user

intervention and to provide meaningful results and visualizations. On the other hand, another methodological gap is addressed in providing analysis for a recent development in histology and in laying out foundations and tools for subsequent studies involving 3D histology in experimental medicine settings. The novelty of the work described in this thesis lies on the development of image analysis and self-supervised machine learning pipelines to address these gaps. I argue and substantiate the utility of the pipelines by comparing them with existing works and additional annotations where they exist.

1.1 Questions, aims, and outline

In this section, I introduce the specific research questions that surround learning and utilizing cell representations in spatio-temporal and 3D contexts. Simultaneously, I outline the chapters of this thesis and discuss how they address the specific research questions.

Chapter 2. Much of the cell and tissue analyses in the past decade use machine learning algorithms to extract information from data that are necessary to give evidence to biological hypotheses. To build the pipelines to address the aims of Chapters 3-5, I adapted existing methods previously proven successful in a similar domain (or even in a completely different one). In this chapter, I provide background to the algorithms that would make up crucial parts of our pipelines. In discussing these algorithms, I cite specific examples of their successful demonstration that ultimately influenced the decision to utilize the algorithms.

Chapter 3. The cell cycle is a well-studied biological phenomenon and holds an important place in studies such as in cancer and aging. In this chapter, I interrogate the possibility of extracting information for cell cycle analysis with fewer input and information compared to previous works in this area. I argue that by limiting the

requirements for analysis, the setup of experiments can be made simpler and more open for additional markers. More importantly, with the absence of typical labels used to train models for cell cycle analysis, the models are freed up from possible biases that may arise from using prior knowledge of the phenomenon. Specifically, I attempt to answer the following question: *How can we build a model with fewer input and annotations and integrate temporal information with spatial features to visualize and to analyze the cell cycle?* I detail how the self-supervised model performs at common cell cycle analysis tasks, such as cell cycle staging of individual cells and cell cycle time estimation. The methods and results presented in this chapter were submitted to Biological Imaging journal and are currently under review.

Chapter 4. Three-dimensional imaging enables us to interrogate structure and spatial relationships that cannot be done using conventional 2D imaging. In this chapter, I extend the use of the built pipeline from Chapter 2 to accommodate 3D data. Within the context of building a pipeline for the analysis of 3D histology data, I investigate building profiles for DNA damage in cells, aiming to establish a methodology to verify and quantify response to treatments in tumor tissues. Moreover, I utilize a similar architecture to stage cells similar to the cell cycle analysis done in Chapter 3 to further subdivide the cell population and provide additional points of comparison between control and treated tissues. Specifically, I aim to answer the following question: *How can we extend the representation learning pipeline for phenotyping cells in 3D histology?* I present multiple ways utilizing the different derived textured classes to phenotype the cells in 3D 4T1 histology data to distinguish treated from control tissue regions. The heatmaps representing DNA damage profiles in tissues were presented at and published as part of the proceedings of the 1st International Workshop on Medical Optical Imaging and Virtual Microscopy Image Analysis (MOVI) at the 25th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI).

Additionally, I showed a joint analysis of texture classes from two fluorescent channels to derive more complex cell phenotypes linking DNA damage with the cell cycle.

Chapter 5. Single cell analysis, although useful, cannot capture certain biological processes and structures that involve multiple cells and cell types, such as in the case of response to treatments and DNA damage. Several tissue-level analysis pipelines exist but training and building a model for tissue-level analysis when single-cell analysis models already exist is not only wasteful of resources but also of processed information. In this chapter, I explore an extension of the work described in Chapter 4 on a different biological scale—from single cells to tissues. Specifically, I aim to answer the following question: *How can we aggregate and utilize cell-level information to characterize and to differentiate higher-level structures in tissues?* I detail how the developed method quantified the interaction of CD8+ and γ H2AX+ cells and identify holes (tears and vasculatures) in 3D histology. The CD8+ and γ H2AX+ proximity analysis described in this chapter was also published in the workshop paper presented at MOVI-MICCAI.

Chapter 6. In this final chapter, I summarize the main results from this thesis, highlight the contribution of the works from the individual chapters towards supporting biological inquiries as analytical tools, and consider potential applications and avenues for extension in future research.

2

Methods

Contents

2.1	Convolutional Neural Networks	26
2.2	Representation Learning	27
2.3	Generative Modelling	28
2.3.1	Variational autoencoders	29
2.3.2	Generative adversarial networks	32
2.4	Segmentation	33
2.4.1	Semantic Segmentation	33
2.4.1.1	Thresholding	33
2.4.1.2	Deep Learning-based Semantic Segmentation .	35
2.4.1.3	U-Net	35
2.4.2	Instance Segmentation	37
2.4.2.1	Stardist	37

2.4.2.2	Cellpose	39
2.5	Projections and Clustering	40
2.5.1	Uniform Manifold Approximation and Projection	41
2.5.2	Gaussian Mixture Models	42
2.6	Shape Analysis	43
2.6.1	Landmarks	44
2.6.2	Parametric descriptors	44
2.6.2.1	Fourier decomposition	45
2.6.2.2	Spherical harmonic decomposition	45
2.6.3	Deformation-based representations	47
2.7	Texture Analysis	47
2.7.1	Statistical methods	47
2.7.1.1	Histogram	47
2.7.1.2	Gray Level Co-occurrence Matrices	48
2.7.2	Model-based methods	48
2.7.3	Transformations	49
2.7.3.1	Texture Filters	49
2.7.3.2	Deep Filters and CNN Features	49
2.8	Summary	50

In this chapter, we look at an overview of the algorithms that served as building blocks for the methods in the succeeding chapters. I present convolutional networks, generative modelling, and representation learning to give background to the analysis pipelines used in this thesis. We also discuss segmentation, dimensionality reduction, and clustering algorithms. This is followed by an overview of shape and texture analysis. In the discussion, I cite some specific examples from literature that successfully demonstrate the advantages of these methods both in general computer vision and biomedical domains. Ultimately, I ground the discussion in relation to the problems to be presented in the succeeding chapters.

2.1 Convolutional Neural Networks

Convolutional neural networks (CNNs) are a class of artificial neural networks primarily applied on image and video data. CNNs are one of the most important advancements in machine learning with pioneers being awarded with the Turing Award in 2018.

Convolutional neural networks have been applied in various tasks such as activity recognition [23], natural language processing and text recognition [24, 25], and object detection and localization [26, 27]. CNNs consist of neurons that have learnable weights and biases. They contain an input layer, an output layer and multiple hidden layers, where a hidden layer consists of a convolutional, pooling, and normalization layers, and fully connected layers. A convolutional layer performs a convolution operation, patterned after feedback from visual stimuli. Pooling reduces dimensionality of feature maps across the network.

Over the years, many architectures of convolutional neural networks were developed since AlexNet [28] in 2012. Influx of newer architectures with various tweaks and novel modifications continues, with new architectures outperforming older ones in classification or segmentation using standard image dataset benchmarks like ImageNet [29] and Berkeley Segmentation Dataset (BSD) [30].

In biomedical image analysis, convolutional neural networks are central to recent findings and advancements in different areas over the last decade. Image registration is important in many biomedical studies. In MR images, correction of motion artefacts prior to quantitation and other downstream tasks can improve the performance of analysis pipelines. In [31], 3D abdominal MR images are registered using a CNN to obtain motion-free images throughout the duration of a respiratory cycle. For pathology, CNN is applied to identify and classify melanoma [32, 33], to diagnose and grade prostate cancer [34], and to detect fibrosis [35, 36]. Applications of CNN in microscopy include super-resolution techniques [37, 38], denoising [39–41],

and particle depth and trajectory estimation [42, 43].

Convolutional neural networks will play a central role in the pipelines to be presented in Chapters 3 and 4. As textural information from microscopy images are needed for my analyses, convolutional neural networks are advantageous as they have been demonstrated to capture texture in meaningful ways. For this reason, I utilized CNNs in this thesis whenever applicable.

2.2 Representation Learning

The success of convolutional neural networks and machine learning models, in general, is heavily influenced by the quality of data representations. Prior to CNNs, representations or features were manually crafted using prior knowledge and mathematical principles. Engineering these features, however, is labor-intensive and highly dependent on prior knowledge. It is desirable to abstract the process of feature engineering and to disentangle learning these features from models for downstream tasks to facilitate faster and more robust development of models and applications [44].

A core principle in representation learning is that high-dimensional data, especially in the real-world, tend to lie within the vicinity of a manifold with much lower dimensions. This *manifold hypothesis* serves as a primary motivation for machine learning and to some extent explains how machine learning techniques perform well on tasks involving high dimensional data. Much of the recent efforts in the area of representation learning employ techniques such as convolutional neural networks to construct a low-dimensional manifold from data with high dimensions such as images [44].

In the case of biological data, particularly in close relation to experimental medicine, representation learning can be beneficial. As I will discuss in Chapters 3, 4, and 5, reducing large amounts of information from pixels and voxels in images into

low dimensional representation vectors is paramount to a multitude of downstream tasks. In Chapter 3, I will detail loss functions that I have employed to influence these data representations such that dissimilarities and similarities between relevant data points are preserved. Using the learned representations from cell and histology imaging data, I will present analyses and visualizations in the following chapters.

2.3 Generative Modelling

We often see two categorizations of machine learning algorithms, dependent on the assumptions adopted during the construction of each algorithm, and their perceived goal. These categorizations are the *discriminative* models and the *generative* models. Statistically, discriminative models directly learn the conditional posterior distribution $p(y|x)$ of the labels (or targets) y for input x . This formulation is a direct way of solving classification and regression tasks where we ask for labels or values for some unseen input x . For most tasks, discriminative models are sufficient as they lend easily to intuition and logical reasoning—based on our observation x (and our past experiences), we decide y without the need to know or to approximate the underlying relations between the variables. However, this limits the extent by which we can model under the discriminative approach. To effectively model the conditional posterior, each data point used in training needs a corresponding label. This makes the approach highly supervised and unsuitable for problems where labels are scarce or unavailable. Generative models approach these tasks in a different manner. Instead of learning the posterior, generative models learn a probability distribution of the input variables $p(x)$, and for labelled data the joint probability with label (target) variable $p(y, x)$. To solve the same classification or regression problem, Bayes' rule can be applied to the generative model to infer the posterior.

While the generative capability of generative models can be attractive, even trendy recently, I limited the scope of this work to using these generative models




Figure 2.1: An autoencoder. An input image is encoded in a compressed representation which is then used to approximate the input. Source: Adapted from [45].

to learn representations from cell and histology images. Much of the data to be presented in this work involve little to no annotations; generative models and self-supervised learning helped in this aspect.

2.3.1 Variational autoencoders

The variational autoencoder is one of the most extensively used architectures for generative modelling. Variational autoencoders share the same structure with other architectures in the autoencoder family. Autoencoders are a type of artificial neural network comprised of two components: an encoder and a decoder. Given an input X , the encoder maps it to a hidden (latent) representation h . The decoder then maps h to an approximation of a target output. The two components are trained together to realize a good mapping between input and output pairs by minimizing a loss metric such as cross-entropy or mean squared error. By forcing the dimensionality of h to be smaller than that of X , the model is also forced to compress information and to filter out noise from the signal. In Figure 2.1, a diagram of an autoencoder is shown. In this example, we can see how images can be compressed while maintaining a level of integrity in the stored information. The autoencoder architecture has been extensively used for segmentation tasks such as applications using the U-Net model to be discussed in Section 2.4.1.3.

Despite having the same structure as other autoencoder architectures, the mathematical foundation of variational autoencoders separates it from the rest. When autoencoders are trained, the encoder and decoder are trained jointly to generate good reconstructions from the learned latent representation of the data points. These mapped latent representations are free to move around the $\mathbb{R}^n$ space, where n is the dimensionality of the latent space, as the model is not regularized to structure the latent space. Given a powerful enough decoder and high degree of freedom from the network parameters, the model can easily overfit and maximize reconstruction fidelity, but with points in the latent space without semantic content. This then affects the generative ability of the model and the disentanglement of the learned representation as data points are not mapped to the latent space semantically. To prevent this and to allow for a more structured latent space, VAEs use an explicit regularization to constrain the latent code produced by the network.

Consider a data point $x^{(i)}$ from a dataset $X = \{x^{(i)}\}_{i=1}^N$, where N is the number of data points. Using a VAE, $x^{(i)}$ is mapped not to a single point in the latent space but to a continuous random variable z . The assumption is that $x^{(i)}$ is generated through a random process involving z : (1) a value $z^{(i)}$ is drawn from a prior distribution $p(z)$, and then (2) $x^{(i)}$ is generated from a conditional distribution $p(x|z)$. Inference on z can then be done by approximating the posterior $q(z|x)$. We assume that $p(x|z)$ and $q(z|x)$ are parametric distributions and can be approximated by neural networks. In this perspective, $q(z|x)$ can be thought of as an encoder and $p(x|z)$ the decoder for the latent code z .

The loss function of the VAE, which involves reconstruction and Kullback-Leibler divergence (KLD), is given by the following equation:

$$\mathcal{L} = \mathbb{E}_{q(z|x)}[\log p(x|z)] - D_{KL}[q(z|x)||p(z)], \quad (2.1)$$

where $\mathbb{E}_{q(z|x)}[\log p(x|z)]$ is the reconstruction loss, $D_{KL}[q(z|x)||p(z)]$ is the KLD of the

posterior $q(z|x)$ from the prior $p(z)$. While VAEs can be used with any continuous and differentiable distribution, most applications of VAE assume multivariate Gaussian with diagonal covariance for simplicity (closed form for KLD). This assumption proves to hold in several previous works. There are of course other works that demonstrate the use of different distributions but they require additional tricks for sampling and computation of the KLD [46, 47]. In the succeeding chapters, the same assumption on Gaussian priors and posteriors is held in favor of simplicity and to be able to focus on addressing some pitfalls in training VAEs.

Using VAEs in generative modelling is not without disadvantage. Due to favoring simple posteriors, optimizing the learning objective sacrifices the generative quality. In contrast with other generative techniques, VAEs tend to output over-smoothed reconstructions. This is addressed in variants and improvements over the vanilla VAE where stricter bounds [48] or looser bounds [49] for the learning objective are derived; a weighted KLD term using hyper-parameter is used [50], a hierarchical latent space is assumed [51–53], representations are discretized [54, 55], and additional networks and loss terms are used [56]. In the experiments in Chapter 3, I explored the last technique to improve the quality of the reconstructions; however, I deemed the manifold structure more important than the generative quality of the model so I focused more on the interpretation and manipulation of the latent space rather than reconstruction fidelity.

VAEs have also been successfully applied in other previous biomedical research. In [20], conditional VAEs were used to encode variations in cell and nuclear shape and the localization of some subcellular structures. This work uses two VAEs: one to encode cell and nuclear shape, and the other to learn the relationship between subcellular structures and the cell-nuclear-shape encoding. This design allows the model to predict the probable localization of unobserved structures in fluorescence microscopy images of cells based on the detected cell and nuclear shapes and other fluorescence. Hira et al [57] used a modified VAE to reduce the

dimensionality of multi-omics data and performed analyses on the latent code to classify transcriptional subtypes of ovarian cancer with high accuracy.

In light of the mentioned successful demonstrations of the VAE as a tool for representation learning, I chose the VAE as the foundation of the models for manifold construction in this thesis. I started with the standard VAE and build on top of it as I impose further constraints during training. In most experiments, I encountered similar pitfalls as previous works when using the VAE. I drew inspiration from these works and applied some regularizations which will be described in Chapter 3.

2.3.2 Generative adversarial networks

Generative adversarial networks (GANs) [58] are another approach to generative modelling. Like a VAE, a GAN is also composed of two networks but instead of an encoder-decoder setup, we have a generator-discriminator setup that is trained in an adversarial manner. The generator network G tries to capture the data distribution and samples this estimation to produce artificial data. On the other hand, the discriminator D estimates how likely is the sample coming from the training data and not from the distribution in G . G is then trained to maximize the probability of D making a mistake. This lends to a minimax game between G and D :

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{data}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))]. \quad (2.2)$$

, where $\mathbb{E}_{x \sim p_{data}(x)} [\log D(x)]$ is the reconstruction loss and $\mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))]$ is the discriminative loss.

This adversarial approach in training generative models changed the landscape of machine learning and pioneered an influx of clever mechanisms to improve the generative and reconstruction quality of the models. This is one key aspect where GAN outperforms VAE. Where VAE would produce blurry images, GANs would generate crisp and highly-detailed ones. However, as there is no explicit

structuring and encoding in the way standard GANs are trained, exploring the latent space for analysis can be challenging.

In Chapter 3, I will explore the joint training of a VAE and a GAN to learn representations from cell cycle data.

2.4 Segmentation

Image segmentation is a key problem in computer vision and essential to other downstream tasks. In the pipelines I will present in the succeeding sections of Chapter 4, segmentation is a primary step prior to extraction of the relevant measurements and representations. Segmentation assigns individual pixels (or in some cases, groups of pixels) into disjoint and discrete classes. In medical imaging, the need and use for segmentation ranges from delineating cell nuclei in microscopy to looking at brain structure and stimuli in magnetic resonance imaging (MRI). In Chapter 4, individual cell nuclei are segmented from the 3D 4T1 histology data.

2.4.1 Semantic Segmentation

Semantic segmentation describes the process of assigning each pixel of an image to a class label, such as dog, sky, or car. In the case of cell microscopy, pixels are assigned to either cell (or a specific cell type) or background. Through semantic segmentation, applications requiring scene understanding can be built, like self-driving vehicles and virtual reality. In biomedical imaging, semantic segmentation enables us to focus our analysis on target areas in biomedical images.

2.4.1.1 Thresholding

One of the simplest ways to segment an image is through thresholding. In a binary case, we assign a pixel to the foreground if the intensity exceeds a set threshold, otherwise the pixel would belong to the background. Formally, given an image I

and a threshold value v , the binary segmentation is computed by

$$I(i, j) = \begin{cases} 0, & < v, \\ 1, & \geq v. \end{cases},$$

where i and j are row and column indices of a pixel location in the image. For multiple segmentation classes, instead of using a single threshold value, multiple intensity ranges can be defined and a pixel is assigned to the corresponding segmentation class if the intensity falls within the range.

Segmentation by thresholding depends heavily on the cutoff intensity values that define the segmentation classes. Several methods exist to automatically compute the values to use. The basic idea is to find the value that minimizes the variances in the segmentation classes. One such method is Otsu's [59]. Algorithm 1 describes Otsu's method for image thresholding.

Algorithm 1 Otsu's method

Let I be an $M \times N$ image with H intensity levels.

Step 1: Compute histogram for a 2D image.

Step 2: Calculate foreground and background variances (measure of spread) for a single threshold t .

i) Calculate weight of background pixels and foreground pixels:

$$W = \frac{\sum w}{M \times N}$$

ii) Calculate mean of background pixels and foreground pixels:

$$\mu = \frac{\sum h \times w}{W}$$

iii) Calculate variance of background pixels and foreground pixels:

$$\sigma^2 = \frac{\sum (h - \mu)^2 w}{W}$$

Step 3: Calculate within class variance:

$$\text{within class variance} = W_{\text{background}} \sigma_{\text{background}}^2 + W_{\text{foreground}} \sigma_{\text{foreground}}^2$$

Step 4: Choose threshold t that maximizes within class variance.

While simple, Otsu's method has been used in many applications. In [60], background area from an X-ray dental image was removed using Otsu. It was also applied on T1-weighted (T1-WI) brain MR images registered with their

corresponding fluid attenuated inversion recovery (FLAIR) sequence to automate the segmentation and detection of glioblastoma. Since the registered images would contain tumor regions with high intensity values, a pipeline that includes a simple method like Otsu were able to segment tumors with high Dice score [61]. However in recent years, we have seen wider acceptance and usage of more complex algorithms that can achieve even higher segmentation accuracy. Still, in very simple and controlled datasets, thresholding offers good and fast segmentation of image data. In our case, I found that thresholding is sufficient to segment some channels in our microscopy data to identify specific cells. This will be presented in Chapter 4.

2.4.1.2 Deep Learning-based Semantic Segmentation

In recent years, as a natural progression of the popularity of applications of CNNs, several CNN-based architectures were developed for semantic segmentation. A step further from the classification problems addressed by CNN architectures like AlexNet[28], VGG[62], Inception[63] and ResNet[64], architectures coupling classification with localization, detection, and segmentation were built. Most of these segmentation architectures, like U-Net [65] described in Section 2.4.1.3, were designed as encoder-decoder models. The encoder is a classification network, oftentimes utilizing the aforementioned popular CNN architectures. This network learns discriminative features for segmentation in a lower-dimensional space. To achieve a full-resolution segmentation, the output of the encoder is passed to the decoder that would project it back to the original resolution space.

2.4.1.3 U-Net

U-Net [65] is a fully convolutional neural network architecture that aims to semantically label individual pixels in the image. The network consists of two paths: a contracting path (encoder) and an expansive path (decoder). Figure 2.2 illustrates the U-Net architecture. The contracting path is the typical CNN with a series




Figure 2.2: U-Net architecture (example for 32×32 pixels in the lowest resolution). Each blue box corresponds to a multichannel feature map. The number of channels is denoted on top of the box. The x-y-size is provided at the lower left edge of the box. White boxes represent copied feature maps. The arrows denote the different operations. Image and caption adapted from [65].

of convolutional layers with rectified linear unit (ReLU) activations and pooling layers to downsample. The expansive path is an upsampling route, taking the output of the contracting path. This path performs a series of up-convolutions, a combination of upsampling and a 2×2 convolution.

The key feature in U-Net is the shortcut connections between the two paths for each resolution scale as shown in Figure 2. At each scale, concatenating the output from the contracting path with the upsampled output from the previous scale in the expansive path ensures that the finer details lost through downsampling can be recovered and be used to fine-tune the segmentation.

U-Net was first applied on microscopy images but since then has been applied to different imaging modalities like MRIs and computed tomography (CT) scans [66, 67]. U-Net serves as backbone for newer segmentation networks, such as

Cellpose [68] and Stardist [69].

2.4.2 Instance Segmentation

Instance segmentation builds on top of semantic segmentation. Not only are class labels assigned to each pixel, instances of the same class are also discriminated. This addresses problems beyond simple scene understanding, where quantification and isolation of objects are needed. In the case of cell microscopy images, we want instance segmentation to be able to count individual cells, and to classify and to characterize the morphology of each. One key issue that algorithms for instance segmentation are still trying to improve on is the resolution of overlapping instances. The following subsections will discuss two CNN-based instance segmentation algorithms that differ in how they resolve these overlapping cell instances. In Chapters 4 and 5, the identification of individual nuclei will enable the enumeration and proximity analysis of different cell types.

2.4.2.1 Stardist

Several prior attempts on instance segmentation [70, 71] proposed instance segmentation as a refinement of semantic segmentation. In such methods, once an image is semantically segmented, pixels are then grouped together to potentially identify different instances. This grouping can be done by connected components analysis or applying domain-specific constraints (expected shape/size), among others. Stardist [69] follows an alternative approach, similar to Mask R-CNN [72], wherein instead of doing segmentation as the initial step, objects are identified first with a rough delineation of their boundaries. This is then followed by a refinement of the object shape.

The output of object detection networks are predominantly axis-aligned bounding boxes; and hence instance segmentation networks such as Mask R-CNN that depend

on these detection boxes use non-maximum suppression (NMS) to prevent multiple detection of the same object. In NMS, low confidence detections are suppressed by overlapping detections with higher confidence. The problem in NMS arises due to the fact that the boxes are axis-aligned. Figure 2.3a illustrates this problem. Because boxes are not good representatives of cell shapes, NMS may suppress boxes containing actual cell instances. Stardist addresses this by representing cell instances as star-convex polygons.

Using a U-Net backbone, Stardist estimates from cell images object probabilities, together with radial distances as shown in Figure 2.3b. Like the original U-Net where the output is the probability of each pixel belonging to the foreground, Stardist outputs a similar probability for pixels. However, Stardist also includes a configuration of a star-convex polygon for each pixel. For each pixel with index i, j , radial distances $r_{i,j}$ for n predetermined equidistant angles are computed. Evaluating on three datasets, Stardist outperformed U-Net and Mask R-CNN.

In [73], Stardist was extended for 3D cell imaging data. Instead of using U-Net as backbone, a modified version of ResNet [64] served as the backbone architecture to predict object probabilities and radial distances. n rays were approximated from equidistant points on a spherical Fibonacci lattice defined by

$$\begin{aligned} z_k &= -1 + \frac{2k}{n-1} \\ y_k &= \sqrt{1 - z_k^2} \sin[2\pi(1 - \phi^{-1})k] \\ x_k &= \sqrt{1 - z_k^2} \cos[2\pi(1 - \phi^{-1})k], \end{aligned}$$

where z_k , y_k , and x_k are the coordinates of the lattice point k , and ϕ is the golden ratio. Similar to the original Stardist, experiments on Stardist3D yielded better scores compared to U-Net3D [74] and IFT-Watershed [75].

Stardist has been applied to study tissue-resident macrophages and their possible contribution to early lung cancer [76], the role of cilia damage in SARS-CoV-2




Figure 2.3: (a) Potential segmentation errors for images with crowded nuclei: Merging of touching cells (upper right) or suppression of valid cell instances due to large overlap of bounding box localization (lower right). (b) The proposed StarDist method predicts object probabilities and star-convex polygons parameterized by the radial distances. (c) We densely predict using a simple U-Net architecture and then select the final instances via non-maximum suppression. Image and caption adapted from [69].

spread to the deeper lung parenchyma [77], and tissue fluidity in carcinoma [78].

2.4.2.2 Cellpose

Cellpose [68] is another segmentation algorithm based on convolutional neural networks. Similar to Stardist, Cellpose is also based on U-Net, having upsampling and downsampling modules. The difference lies on the information from the images it is trained on.

For a given image, maps representing vector flow are computed. Cellpose borrows from tracking studies that use vector fields to relate object parts in object tracking and coarse segmentation. Both horizontal and vertical gradients to create the vector field, as well as a semantic segmentation, are needed for training. In cell images, computed vector fields relate all pixels in a cell image to the center of the cell. This seems very similar to the ray tracing in Stardist, but the advantage lies on the non-reliance on pre-defined set of angles and direct relation to the cell center. In Cellpose, the gradients in the vector field does not necessarily need to point towards the cell center. This becomes useful for non-convex shapes such as

protrusions in cells. In such cases, the gradient translates to another pixel in the cell whose gradient points to the cell center, and these gradient traversals can be as long as necessary to capture complex cell shapes. Using a U-Net-based architecture, Cellpose predicts the gradients and semantic segmentation from input cell images. Inference can then be done by tracking the gradients over the generated semantic segmentation. Cellpose can also be extended to analyze 3D images by computing not only the horizontal (x) and vertical (y) gradients but also the gradient in z . Figure 2.4 visualizes the details behind Cellpose.

Cellpose has been applied in segmenting myofibers in skeletal muscles [79], studying epithelial morphogenesis in relation to cell matrix and cell-cell adhesion [80], and in revealing human-specific cell-cell proximity patterns and enriched interactions between neurons and non-neuronal cells in the human cortex [81]. In Chapter 4, I used Cellpose to segment nuclei from histology data.

2.5 Projections and Clustering

Often, the chosen dimensionality of the manifold is greater than what we can visualize (2- or 3- dimensional). To make sense of such manifold, we apply projection to the data points from this higher dimension to 2 or 3 dimensions, while still preserving structure and relative distances. On the other hand, we perform clustering on the latent representations to group similar data points together and extract semantic meaning from each cluster. Many of the analyses in this thesis required the projection of points in 2D for visualization. Clustering is also important in Chapters 3 and 4 as part of the annotation-free routine to discover texture similarities in cell cycle and histology imaging data.




Figure 2.4: Cellpose segmentation. (a-b) procedure and examples to compute gradients, (c-d) model architecture, (e-f) inference. Image adapted from [68].

2.5.1 Uniform Manifold Approximation and Projection

Uniform Manifold Approximation and Projection (UMAP) [82] is a technique for dimensionality reduction. The algorithm assumes that the data is distributed normally in a Riemannian manifold that is locally connected and with a locally constant metric. The embedding is computed by finding the lower dimensional projection of the data that best matches the fuzzy topological structure in higher dimension.

Prior to the development of UMAP, t-distributed stochastic neighbor embedding

(t-SNE) [83] is the dimensionality reduction and visualization technique used in most biological studies [84–86]. t-SNE suffers from getting stuck at local minima, slow computation, and poor scalability [82]. UMAP offers tricks to circumvent these issues and claims to preserve global structure better than t-SNE. For this reason, we used the UMAP technique instead of t-SNE in our experiments in the succeeding chapters. Full derivation and mathematical foundations of UMAP can be found in [82].

As the points in the latent space are encoded as multivariate Gaussians in the experiments in the next chapters, the Wasserstein distance was used as custom metric in the computation of the UMAP projection. The Wasserstein metric is a distance function defined between probability distributions on a given metric space. For two normal distributions, μ_1 and μ_2 , with means m_1 and m_2 and symmetric positive semi-definite covariance matrices C_1 and C_2 , the 2nd Wasserstein distance is given by

$$W_2(\mu_1, \mu_2)^2 = \|m_1 - m_2\|_2^2 + \text{trace}(C_1 + C_2 - 2(C_2^{1/2}C_1C_2^{1/2})^{1/2}) \quad (2.3)$$

For the special case where the covariances are zero, as in point masses in space, the 2nd Wasserstein distance becomes equivalent to the Euclidean distance between the two points. In the case of the 16-dimensional latent space of the VAE-GAN model, as the data points are represented by a diagonal covariance, σ^2 , the 2nd Wasserstein distance is simplified as

$$W_2(\mu_1, \mu_2)^2 = \|m_1 - m_2\|_2^2 + \Sigma(\sigma_1^2 + \sigma_2^2 - 2\sqrt{\sigma_1^2\sigma_2^2}) \quad (2.4)$$

2.5.2 Gaussian Mixture Models

In many experiments in this thesis, Gaussian mixture models (GMM) were used for clustering. GMM is a clustering technique that assumes data points are generated from a finite set of Gaussian distributions in an n -dimensional space. The technique is particularly useful in modelling subpopulations that are assumed to be normally

distributed and possibly overlapping.

GMMs can be considered as a relaxation of the K-means clustering. In K-means, we start with an initial set of cluster centers and assign each data point to the nearest cluster. We then recompute new cluster centers based on this assignment and repeat the assignment-recomputation process until convergence to a local minimum. This process results in hard boundaries between clusters and hard label assignments for data points. In GMM, points are modelled as a linear combination of n component Gaussians and solved using the Expectation-Maximization (EM) algorithm, and the weights associated with each component can be used to describe cluster labelling probabilistically and to define soft boundaries between clusters.

2.6 Shape Analysis

A key step in many computer vision problems is the analysis of object shapes. In biomedical imaging, the shape of particular regions of interest can support and provide evidence to biological and pathological inquiries. For example, the characterization of cell shapes gives rise to phenotyping for various studies on gene expression and signalling. Mechanotransduction is a fundamental biological process in which cells sense mechanical stimuli and induce response. The cytoskeleton, microtubules, and actin stress fibers can change and deform in response to the stimuli - changing the shape of the cell. The mechanotransduction process and the factors that can alter it such as diseases and chemical signals can therefore be studied by observing the cell shape [87].

In many studies, shape analysis is done by extracting descriptors from segmented object instances. These various descriptors for shape can be categorized into three types: landmark-based, parametric, deformation-based [88].

2.6.1 Landmarks

Landmarks are points in objects that can be accurately identified. From these points, several geometrical properties can be manually computed. These include the following:

1. Area
2. Perimeter
3. Elongation
4. Eccentricity
5. Orientation
6. Convexity
7. Curvature
8. Central Moments

Orientation will be used in Chapter 5 to analyze tissue regions and to recognize vasculature and tears in histology.

A key drawback from using landmark-based descriptors is the heavy assumption on the correctness of the landmark identification and placement. For analysis of segmented regions in images, the quality of the landmark-based descriptors is ultimately defined by the accuracy of the segmentation methods upstream.

2.6.2 Parametric descriptors

Parametric shapes are generalized models that return different shapes for different parameter settings. Reversing the process, complex shapes can be estimated and represented by a unique realization of parameter settings. Function bases are defined and discrete coefficients for each base are recorded and serve as shape

descriptor. In the following subsections, commonly used basis functions for shape descriptors will be discussed.

2.6.2.1 Fourier decomposition

The idea behind Fourier descriptor is to use the Fourier transform given below to estimate the boundary of an object:

$$\begin{aligned} f(x, y) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F(u, v) e^{j2\pi(ux+vy)} \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} [\cos 2\pi(ux + vy) + j \sin 2\pi(ux + vy)], \end{aligned}$$

where $f(x, y)$ is a shape function in 2D space and $F(u, v)$ is a function in the frequency domain.

For an image $f(x, y)$, the discrete analog below can be used.

$$\begin{aligned} f(x, y) &= \sum_{u=0}^M \sum_{v=0}^N F(u, v) e^{j2\pi(ux+vy)} \\ &= \sum_{u=0}^M \sum_{v=0}^N [\cos 2\pi(ux + vy) + j \sin 2\pi(ux + vy)], \end{aligned}$$

where $f(x, y)$ is the image intensity at pixel location (x, y) and F is a function in the frequency domain composed of sines and cosines, and (u, v) are coordinates in the frequency domain.

The image $f(x, y)$ can then be decomposed into a weighted sum of 2D orthogonal basis functions (sine and cosine). These weights we keep as representation for the image. For an extracted boundary of an object, the weights (or coefficients) form the Fourier shape descriptor [89, 90].

2.6.2.2 Spherical harmonic decomposition

Spherical harmonics are another set of orthogonal basis functions like the Fourier series. They are special functions defined on a sphere. Like the Fourier shape




Figure 2.5: Examples of spherical harmonics bases up to the third degree. Image adapted from [92].

descriptor, the idea is that any spherical function (region boundary in the case of describing shapes) can be represented as the sum of its harmonics [91], given below

$$f(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \hat{f}_{l,m} Y_l^m(\theta, \phi)$$

where $\hat{f}_{l,m}$ is the Fourier coefficient and $Y_l^m(\theta, \phi)$ is the spherical harmonics base computed from Legendre functions and complex exponentials at level l and degree m for angular coordinates θ and ϕ . Figure 2.5 visualizes examples of spherical harmonics bases up to the third degree.

In initial work on the histology data used in Chapter 4, I utilized the spherical harmonic coefficients in an attempt to differentiate elongated from non-elongated nuclei in tissues but I found that the representation is highly sensitive and reliant on the segmentation. While it was able to profile individual nuclei, distinguishing treated from control was not successful. The project has then pivoted to looking at texture, which I discuss in this thesis, more than the shape of the nuclei in histology. In future work, revisiting spherical harmonic coefficients representation can be done in attempting to bridge shape and texture analysis of our data.

2.6.3 Deformation-based representations

Both parametric and landmark-based descriptors capture object shapes from binary segmentation and object boundaries. Deformation-based descriptors characterize shape using the entire image. The approach assumes that shape can be inferred from deformation information that can be extracted from raw images. It works by matching a known template to an image and applying constraints to the deformation field. Applications include assessing hippocampal differences in semantic variant of primary progressive aphasia and Alzheimer’s disease in brain MR images [93], and revealing that initial heart looping is achieved by left-right (LR) asymmetry in the direction of deformation within the myocardial tube [94].

2.7 Texture Analysis

Another way to characterize objects is through texture. Texture in images refers to the spatial arrangement of different colors or intensities. There are many ways to analyze image texture. Approaches include statistical, model-based, and transformation-based methods.

2.7.1 Statistical methods

Statistical methods to capture texture aim to compute a measure or a set of measures from the distribution of pixel intensities in a region.

2.7.1.1 Histogram

The most straightforward way to look at intensity distribution is to calculate a histogram. Histograms are simple visualizations where the frequency of each intensity level can be shown and compared. In [95], the analysis of histograms in MR images of the spine was able to show differences among degeneration grades

of intervertebral discs (IVD).

Although simple, histograms have several disadvantages as a texture descriptor. First, local information is not captured as only the global pixel intensity distribution is accounted for. Second, the spatial relationships of intensities in pixel neighborhoods are discarded. Lastly, completely unrelated scenes in images may result in similar histograms [96].

2.7.1.2 Gray Level Co-occurrence Matrices

A gray level co-occurrence matrix (GLCM) [97] captures the spatial relationship of pixels in an image. It is as a statistical function used as a quantitative measurement for image texture. As a second-order statistic, the frequency of occurrence of pairs of image pixels are computed. A GLCM for an image is an $h \times h \times n_d \times n_\theta$ matrix, where h is the number of intensity levels (bins), n_d is the number of pixel distances to consider in neighborhoods, and n_θ is the number of angles to consider for pixel neighbors. From the computed matrix, some properties can be extracted [98]. GLCM properties have been used in ecology [99], remote sensing [100], manufacturing [101], and biomedical image analysis [102, 103]. In Chapter 4, we utilized some of the GLCM properties to characterize the distribution of fluorescence intensities in nuclei in 4T1 tissues.

2.7.2 Model-based methods

In model-based methods, a texture model is defined using a set of parameters that is set using domain knowledge or learned from the variation of pixels in images. Texture analysis methods under this family include autoregressive model, fractals, Gibbs random fields, Markov random fields, and Hidden Markov model [96, 104].

2.7.3 Transformations

Texture can be represented as (and computed from) transformations from the original space to another [96]. In the new transformed space, texture can be defined as corresponding response from filter banks or transform coefficients similar to parametric shape analysis as discussed in Section 2.6.2.

2.7.3.1 Texture Filters

Texture filters aim to provide local representations invariant to contrast, rotation, and scale. In order to do this, filters are generated from the same function but using different parameter settings.

Gabor filters are one of the most popular texture descriptors primarily due to their relation to models of early mammalian visual systems. The Gabor wavelets are generated by appropriate rotations and dilations of the following function:

$$\phi(x, y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp\left(-\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2}\right) \exp(j2\pi\omega),$$

where σ_x and σ_y are standard deviations for dilation, ω is the angle of rotation, and x and y are Cartesian coordinates. A filter bank defined by different rotations and dilations for this function can be applied to images to extract responses that would comprise a “bag-of-words” descriptor.

2.7.3.2 Deep Filters and CNN Features

With the increasing popularity of deep learning models and CNNs, the number of studies that use texture representations computed by the previously discussed approaches is declining. Several comparative studies demonstrated that features from CNN and filter responses from the convolutional layers of CNNs outperform the older approaches in texture analysis [105–107]. In the same light, I will compare texture features from a deep learning model with hand-crafted features in Chapter 4.

2.8 Summary

In this chapter I presented the collection of methods that will be relevant to the analyses in the succeeding chapters of this thesis. For each method, I discussed motivations and assumptions, with some mathematical detail, and then grounded them on specific applications in general computer vision problem domains and in areas related to biomedical imaging. I also briefly introduced some of the specific applications of these methods in this thesis such as the segmentation of the cell nuclei in 4T1 tissues and constructing representations for cell nuclei.

In the next chapter, I present a machine learning pipeline that uses the generative modelling, representation learning, clustering, and projection methods I introduced previously to analyze cell cycle imaging data.

3

Learning a Morphological Manifold of the Cell Life Cycle

Contents

3.1	Cell Cycle	52
3.2	Methods	59
3.2.1	VAE-GAN	59
3.2.2	Temporal Cycle Consistency	61
3.3	Experiment Setup	64
3.3.1	Dataset	64
3.3.2	Training and Evaluation	65
3.4	Analyses and Discussion	66
3.4.1	Classifying visually distinct clusters	66

3.4.2	Identifying Key Cell Cycle Stages from Single Cell Trajectories	76
3.4.3	Quantifying Cycle Times and Extracting Population Statistics	78
3.5	Summary	80

Cell division is key to life. Without it, no species would be able to reproduce, and growth and healing would be impossible. On the other hand, repeated cycles of cell division can result in tumors and cancer if left unregulated. Strict timing of cell division is therefore crucial to preventing life-threatening consequences. This timing of growth, replication, division, and death in cells is what we call the *cell life cycle*. Its critical role in biology is exhibited by the number of previous studies devoted to identifying checkpoints, pathways, and necessary conditions, within this cycle [108–111]. In this chapter, I investigate the application of a deep learning model to learn representations from a cell cycle dataset in a self-supervised manner through which common tasks in cell cycle analysis can be done and relevant information such as staging and cycle times can be visualized. Through the construction of a cell cycle morphological manifold, the representations we can extract incorporate spatio-temporal information from single channel data that can be visualized and analyzed for common tasks in cell cycle analysis, such as staging and cycle times.

3.1 Cell Cycle

The cell cycle has four known stages: G1, S, G2, and M. In G1 (gap 1 or growth 1), the cell grows in size and synthesizes mRNAs and proteins needed for DNA replication. It is in G1 when the decision to divide or not occurs. When the cell receives signal to remain undivided, the cell goes dormant (G0 phase). Otherwise, the cell continues protein synthesis in preparation for the S stage. During the S




Figure 3.1: The cell cycle. Image adapted from [112].

stage, DNA is copied. After checkpoints to check for DNA damage, the cell enters the G2 phase where further growth and protein synthesis occur. The phase serves as the final stage for checks and repairs on DNA strand breaks before division. Finally, after passing all checkpoints, the cell enters the M (mitosis) stage and divides into daughter cells. Figure 3.1 illustrates the cell cycle.

During the different stages within the cell cycle, cells undergo molecular changes which would then manifest as morphological changes. During the G1 and G2 stages, the volume of the cell increases. During mitosis, a rapid sequence of changes occurs: chromosomes condense, spindle fibers emerge, chromosomes move toward poles, and ultimately the cell divides. In the S stage, as the replicated DNA needs to be packaged into nucleosomes, histone production is increased. Moreover, proliferating cell nuclear antigen, or PCNA, is also maximally expressed during the S stage as it associates with the DNA polymerase during DNA synthesis. These two, histone and PCNA, when labelled serve as perfect targets to analyze the cell cycle progression. Labelled histone and PCNA manifest in microscopy images as high intensity and scattered localized regions in the nuclei. Figure 3.2 shows the PCNA fluorescence




Figure 3.2: PCNA fluorescence during the cell cycle stages. Image contrast was adjusted for visualization.

and variations in cells during the different stages. These characteristic variations across different stages enable the automated analysis for cell cycle research. While we know such variations exist during the cell cycle from the collective results of multiple cell cycle studies, cell features may present differently from one cell to another, or from one study to another. Furthermore, our perspective on these variations is limited by what is currently known about the previously identified discrete cell stages. We may also observe some variations that are subtle and whose differences may be hard to quantify without using automated analysis pipelines and annotation of individual cells.

In recent years, we have seen advancements in microscopy imaging enabling us to acquire data at a larger scale. With this comes the greater need to interpret and to process the large amounts of data in an automated manner. Machine learning-based approaches have been demonstrated to leverage this large amount of data used in training to build more robust models for tasks such as cell phenotyping [113, 114]. We have also seen several successes in machine learning-based cell models to decrease our reliance on annotations and to simplify biological experiment setups while retaining the same level of information that can be extracted [20]. With all these methods, we are now able to fill more gaps in our knowledge of biological processes and more easily revisit, challenge, and update our understanding of known phenomena such as the cell life cycle in light of new discoveries.

A handful of previous studies revisited the analysis of the cell life cycle [22, 115–117]. Gut *et al* [115] and Eulenberg *et al* [116] created independent models that can generate cell cycle trajectories and identify cell cycle stages. CellCognition [117] is an open-source framework that can classify live-cell images into life cycle stages based on their morphology quantified using extracted statistical texture features. Using chromatin (H2B-mCherry) and DNA replication factory (EGFP-PCNA) markers, a support vector machine (SVM) was trained in a supervised way to automatically label cell trajectories over time. This classification is further refined using a hidden Markov model (HMM). This approach demonstrates a way to include temporal context to resolve ambiguous classifications that arise due to model confusion and the inherent low signal-to-noise ratio of live imaging. Key limitation of this model, however, is the expert annotation requirement to train the SVM and HMM. Aside from being time-consuming, expert annotation requires the explicit discretization of cellular states which may not be suitable in recent efforts to understand the cell life cycle better, beyond what is currently known. Therefore, it is ideal to disentangle cell and cell morphology models from biases that could arise from annotations and our prior knowledge of the cell life cycle. Newer

techniques in machine learning, under the family of self-supervised techniques, that uses label-free modelling is best suited in this regard.

DeepCycle [22] is one example of such self-supervised methods. The method uses classification into four virtual classes based on intensity as an auxiliary task to train a convolutional neural network to encode cell images into a four-dimensional latent space. Projecting the points on this latent space into a two-dimensional plane reveals cyclic trajectories as cells undergo the life cycle. While this approach presents well-formed trajectories in the manifold, it still requires a form of discretization during training due to the arbitrary number of virtual classes. Moreover, the setup requires four channels: brightfield and Hoechst to train their model, and FUCCI channels to create the pseudo-labels. DeepCycle exploits the wide usage of the FUCCI setup in cell cycle studies. FUCCI uses the phase-dependent nature of replication licensing factors Cdt1 and Geminin. Cdt1 (amino acids 30-120) is fused with the fluorescent protein monomeric Kusabira-Orange 2 (mKO2) to serve as marker for G1 and Geminin (amino acids 1-110 or 1-60) is fused with the fluorescent protein monomeric Azami-Green 1 (mAG1) to observe the S, G2, and M phases. The combination of these two probes, their replication and degradation over time, reliably visualizes the cell cycle stages. The disadvantage in using the FUCCI setup is that two markers are already allotted to characterize the cell cycle. This means two fewer channels available for use to look at specific cell regions and processes tied to the cell cycle. DeepCycle attempts to free up this need for two dedicated channels by estimating cell cycle stage from only brightfield and Hoechst. Similarly, in this work, we further this reduction by looking at only one channel to estimate cell cycle stage.

Apart from the morphology itself, morphological changes over time is a critical information in the analysis of the cell cycle. CellCycleGAN [118] is one work in this area. From carefully annotated cell sequences using landmarks, statistical models for each of the cell cycle stages were computed. To construct a morphology at time t , a weighted sum of the models is calculated and passed to a GAN to generate a

corresponding image. Their results indicated high similarity between generated and real microscopy data. However, their approach focuses heavily on the generation of synthetic sequences. In contrast, the work in this chapter puts greater importance on using the latent space for morphological analyses. Moreover, supervision and reliance on landmarks to describe cell shape limit the CellCycleGAN approach.

Motivated by these gaps in addressing critical parts in cell morphology modelling, I re-explored the analysis on cell life cycle and propose a different approach based on self-supervised techniques for manifold learning. To construct this manifold, I exploit the power of generative techniques to compress high dimensional data to a lower dimensional space in semantically meaningful ways where spatio-temporal patterns in morphological changes during the cell life cycle can be better analyzed and generalized. Generative adversarial networks (GANs) [58] and variational autoencoders (VAEs) [119], were successfully demonstrated to be applicable to this task [19, 20]. Modelling through these techniques makes it possible to approximate the underlying distributions in the data and group data points into meaningful clusters. In the case of the cell life cycle data, these clusters can demonstrate and explain the different stages in the cycle and the variance in the visual appearance in each stage across cell samples, without the injection of any prior knowledge of known stages and their underlying molecular states. With the use of generative models in a self-supervised manner, I deviated from previous works. Instead of relying on pixel-wise cell segmentation, manual extraction of features, and discrete labels like in CellCognition [117], I trained an encoder that can automatically extract meaningful representations of live cell images. In contrast with the DeepCycle [22] method, I used the reconstruction of cell images as an auxiliary task to learn the morphological manifold. In addition to this, I further simplified the model to only utilize the DNA replication factory channel, compared to the combined brightfield and Hoechst channels in their setup.

Furthermore, I adopted concepts from manifold manipulation and video align-

ment to utilize the temporal information, together with the spatial features used in the generative model, to further structure the morphological manifold for the cell life cycle. The manipulation of trajectories in manifolds is explored in [120]. With a model-agnostic approach, the study presented a general way to force models to generate trajectories in the manifold for two related input sequences similarly. With the goal of representing the cell life cycle data as trajectories in the manifold, this temporal approach is a good way to structure the manifold such that embeddings for events such as mitosis and DNA replication will be constrained together and transitions between these events can be smoothened. I demonstrate that utilizing the temporal context of the cell morphology data to fine-tune the constructed manifold from extracted spatial features can be an alternative to the pseudo-labels and the auxiliary classification task in [22] and the explicit graphical model of discrete cell cycle stages in [118]. Moreover, in the preparation of cell samples for experiments, such as in the analysis of the cell life cycle, synchronizing multiple cells to start with and to exhibit desired visual characteristics at the same time is difficult and hence, this results in unaligned events across single-cell profiles and trajectories. Ideally, we want a cell event from different single cell trajectories to be as close as possible in the morphological manifold. For this purpose, I utilize the temporal loss from [120] to constrain the projection using not only the spatial features from the cell imaging data but also the temporal relationship among frames in a single-cell trajectory. In summary, the novelty of the approach lies on disentanglement of manifold construction from expert annotation and our current knowledge of the known cell cycle stages, and using temporal ordering of data to structure the manifold to extract information from cell trajectories.

In the succeeding sections, I discuss a model that uses only single- channel microscopy images and is based on a VAE variant. I also detail the dataset used, additional regularizations on the model, and choices in training parameters and their justification that ultimately led to the final manifold model for cell

cycle morphology. Lastly, I present analyses and experimental results that can be obtained with the trained model, compare with CellCognition [117], and argue the advantages of the annotation-free model.

3.2 Methods

Building on algorithms discussed in Chapter 2, I present in this section the algorithms and additional objectives to construct the manifold of the cell cycle.

3.2.1 VAE-GAN

As mentioned in Chapter 2.3.1, standard variational autoencoders suffer from blurry generated images because of the assumption on the Gaussian posterior. However, there are approaches that demonstrated how to increase the level of detail in image reconstructions. In this chapter, I employed a technique with additional regularization and a surrogate loss function for reconstruction in a GAN-like module.

A generative adversarial network (GAN), like VAE [58] is an approach to generative modelling. Discussed in Chapter 2.3.2, a GAN is composed of a generator-discriminator setup that is trained in an adversarial manner. GAN outperforms VAE in generating crisp and highly-detailed images but lacks structure in the latent space.

In an attempt to bridge the two techniques, VAE-GAN is proposed in [56]. Here, the choice of loss function for the reconstruction term in the VAE objective is highlighted. The authors argue that while pixel-wise similarity functions like mean squared error (MSE) is a common choice for the reconstruction loss, they do not reflect human perception of similarities. In the case of translation, for example, a human might not be able to discern small amounts of translation even though this can translate to a high MSE score. As an alternative, they proposed an approach to learn a higher-level metric that is sufficiently invariant and demonstrated that switching MSE with this metric in training VAEs yields better results. The approach

constitutes combining a VAE and a GAN with a shared decoder/generator module. As the discriminator is being trained to distinguish between true and fake data points, the authors note that the features the discriminator extracts can be used as a substitute to the MSE score for the VAE objective. Training routine for the VAE-GAN is shown in Algorithm 2.

Algorithm 2 Training a VAE-GAN model

```

 $\theta_{Enc}, \theta_{Dec}, \theta_{Dis}$  initialize network parameters
repeat
   $X \leftarrow$  random mini-batch from dataset
   $Z \leftarrow \text{Enc}(X)$ 
   $\mathcal{L}_{prior} \leftarrow \mathbb{D}_{KL}[q(Z|X), p(Z)]$ 
   $\tilde{X} \leftarrow \text{Dec}(Z)$ 
   $\mathcal{L}_{llike}^{Disl} \leftarrow -\mathbb{E}_{q(Z|X)}[p(\text{Dis}_l(X)|Z)]$ 
   $Z_p \leftarrow$  samples from prior  $N(0, I)$ 
   $X_p \leftarrow \text{Dec}(Z_p)$ 
   $\mathcal{L}_{GAN} \leftarrow \log(\text{Dis}(X)) + \log(1 - \text{Dis}(\tilde{X})) + \log(1 - \text{Dis}(X_p))$ 

  // Update parameters according to gradients
   $\theta_{Enc} \stackrel{+}{\leftarrow} -\nabla_{\theta_{Enc}} (\mathcal{L}_{prior} + \mathcal{L}_{llike}^{Disl})$ 
   $\theta_{Dec} \stackrel{+}{\leftarrow} -\nabla_{\theta_{Dec}} (\gamma \mathcal{L}_{llike}^{Disl} - \mathcal{L}_{GAN})$ 
   $\theta_{Dis} \stackrel{+}{\leftarrow} -\nabla_{\theta_{Dis}} \mathcal{L}_{GAN}$ 
until convergence

```

For my experiments, I used VAE-GAN instead of the original VAE architecture to construct the manifold of the cell cycle. As several losses and networks are used in the model, I opted to train the VAE-GAN in four steps, with each step optimizing one loss before proceeding with the addition of other loss terms. My early experiments showed that training with all losses at the same time results in instability and in loss terms overpowering each other, so I decided to construct and manipulate the manifold one step at a time. First, I trained the network to minimize the reconstruction loss (MSE) and let the network output rough reconstructions. This is equivalent to training a simple autoencoder to compress information in a latent space. Second, given this initial latent space, I introduced regularization by using the KLD loss. Third, I added the GAN loss and swapped the learned

similarity metric with MSE. Lastly, I further regularized by using the temporal cycle consistency loss discussed in the next section.

The VAE-GAN model used to construct the manifold is composed of simple components. The encoder contains three convolutional layers with batch normalization, followed by a series of dense feed-forward layers. For all the experiments, the dimension of the latent space is set to 16. The dimensionality of the latent space is chosen as a compromise between the usual dimensionality in literature for experiments on general image classification challenges (100) and MNIST (10). The value 16 is chosen for convenience: a power of 2 and a perfect square but still higher than the number of known cell cycle stages. The decoder mirrors the encoder with three transposed convolutional layers. To ensure that this simple encoder-decoder setup can effectively compress the cell data in the 16-dimensional latent space, an initial experiment without sampling and reparametrization trick in VAE is done. Shown in Figure 3.3 are some input cell images and their reconstructions from the encoder-decoder setup.

For the discriminator, a model with five convolutional layers with batch normalization was constructed. Feature maps after the fourth convolutional layer are extracted and used for the learned similarity metric described in [56]. Figure 4.3 shows a diagram of the layers of the VAE-GAN model used in the construction of the cell morphology manifold.

3.2.2 Temporal Cycle Consistency

Temporal cycle consistency (TCC) [120] is a self-supervised representation learning approach introduced to temporally align similar videos in an embedding space. The approach details a temporal loss based on cycles formed from different corresponding points in a set of trajectories. For embedding sequences $U = \{u_1, u_2, \dots, u_N\}$ and $V = \{v_1, v_2, \dots, v_M\}$, a reference point $u_i \in U$ is picked and the nearest neighbor




Figure 3.3: Input images (odd columns) and their reconstruction (even columns) from the VAE model without reparametrization. The reconstruction shows that the 16-dimensional latent representation is able to preserve some crucial characteristics of some cell life cycle stages such as the multiple foci and high intensity areas during DNA synthesis.

$v_j \in V$ is determined. Next, the opposite is also attempted to be established, i.e. u_i is the nearest neighbor of v_j . If both ways hold true, then the pair of points are cycle consistent. This concept is visualised in Figure 3.5. The goal is to maximize the number of such point pairs that are cycle consistent. To influence structure on the manifold, however, the process needs to be differentiable. The approach offers a cycle back classification where each frame in U is considered a separate class and reformulate the problem as finding the soft nearest neighbor. In building the manifold for the cell life cycle data, the TCC approach was used to encode the temporal context and to improve the alignment of specific stages in the life cycle in the manifold.




Figure 3.4: The VAE-GAN architecture (A) used to construct the manifold of cell morphology. The network is composed of three components: encoder (B), generator (C), and discriminator (D). The encoder maps input images to a 16-dimensional Gaussian with diagonal covariance. The generator produces a reconstruction from sampled points in the latent space. The discriminator forces the generator to output images as similar to the input as possible.




Figure 3.5: Temporal cycle consistency. Image adapted from [120].

3.3 Experiment Setup

This section describes the training and experimental setup, as well as the data used to construct the model for the cell cycle.

3.3.1 Dataset

To demonstrate the utility of the VAE-GAN model in the construction of a manifold of the cell cycle, I worked on a dataset from a previous cell cycle study. This dataset is from CellCognition [117]. In this dataset, EGFP-PCNA was used to mark the DNA replication factories and H2B-mCherry for the chromatin of Human Hela Kyoto cells imaged using widefield epifluorescence, 10x dry objective (0.645 $\mu\text{m}/\text{pixel}$), 1392x1040 pixel, 482 frames imaged at 5.9 min intervals. In this work, I focused only the PCNA channel from this data to build a model that uses only minimal information, further extending the goal of DeepCycle [22] in reducing the marker dependencies in models. This dataset has 408 labelled single cell trajectories with 255 frames. I used 60% of these trajectories for training and the remaining for testing.

3.3.2 Training and Evaluation

All single cell images are resized to and centered on 64×64 patches, and scaled to the $[0,1]$ range. Images are augmented via rotations of 90° increments. All network modules are trained with ADAM optimizers with learning rate of 0.0001. The loss function used for training, $\mathcal{L}_{training}$, which incorporates losses for the VAE, GAN, and TCC modules is described below:

$$\mathcal{L}_{training} = \alpha\mathcal{L}_{recon} + \beta\mathcal{L}_{KLD} + \gamma\mathcal{L}_{GAN} + \kappa\mathcal{L}_{TCC} \quad (3.1)$$

where $\mathcal{L}_{recon}$ is the MSE (first step in scheme in Section 3.2.1) or learned similarity metric loss (from second step in scheme in Section 3.2.1), $\mathcal{L}_{KLD}$ is the KLD loss, $\mathcal{L}_{GAN}$ is the discriminative loss, $\mathcal{L}_{TCC}$ is the temporal cycle consistency loss, and α , β , γ and κ are adjustable hyperparameter coefficients. α , β , γ and κ were initially set to 0 and scheduled values to implement the training scheme in Section 3.2.1.

For each step in the four-step training scheme described in Section 3.2.1, all modules are trained for 1000 epochs, with checkpoints and early stopping criterion every 10 epochs. In my experiments, I found that the model converges after 500 epochs in the first two steps of the training process, and after 800 epochs in the last step. I used the reconstructed images like in Figures 3.3 and 3.9 to qualitatively assess model convergence and reconstruction fidelity. Together with a biologist, these images were inspected to make sure that they still present noticeable visual cues with respect to the cell cycle stages. In particular, images in S stage were checked for isolated bright foci characteristic of this stage.

Training and evaluation were done 5 times to capture the variability of the model performance. To quantitatively evaluate my approach, I tested it against common tasks in the analysis of the cell cycle detailed in next section. Cell representations constructed from the approach were clustered and these clusters are compared to the ground truth labels. For this classification task, accuracy values were reported.

Cell cycle times were also computed and compared against the ground truth and evaluated with root mean squared error (RMSE).

3.4 Analyses and Discussion

In this section, I detail how the trained model can be used for common tasks in cell cycle analysis.

3.4.1 Classifying visually distinct clusters

I projected the images in the dataset to the learned manifolds and analyzed the clusters formed. I start with the standard VAE model and demonstrate qualitative improvements on this manifold as I apply the different constraints and additional modules. I also present projections using pre-trained CNN architectures to serve as baseline for comparison. Qualitative assessment of the manifolds was done by visually inspecting the clusters formed using ground truth labels and their separation. While qualitative assessment worked and visually showed the incremental improvements from the addition of constraints, losses and modules, a quantitative assessment in a future work, such as cluster cohesion and separation scores, inter- and intra-class variances, may unravel patterns in model construction and provide quantitative scoring for different approaches.

The baseline manifolds are constructed using the VGG [62] and ResNet-50 [64] pre-trained models. After using these models as feature extractor to output 25088 and 512 representative vectors for each of the images and using the CellCognition labels to cluster the vectors with respect to cell cycle stages, I visualized the data using UMAP with default parameters (2 components, 15 neighbors, 0.1 distance, euclidean metric). I plotted the resulting points in 2D as shown in Figure 3.6. We see that in both models, the constructed manifold does not produce visually distinct clusters of the cell cycle stages. As the VGG and ResNet models are not




Figure 3.6: Clustering on encodings from Resnet, VGG. Colors correspond to CellCognition labels.

fine-tuned for cell images and were trained to classify natural image classes, the poor clustering is to be expected and further demonstrates the need for specifically-trained model for the cell cycle manifold.

I next looked at an attempt to utilize the technique described in DeepCycle [22]. For this experiment, since the dataset does not have FUCCI channels from which to derive the pseudo-labels for each data point, I tried deriving 4 intensity classes from




Figure 3.7: Clustering on encodings from DeepCycle. Colors correspond to CellCognition labels.

the H2b channel. I trained the DeepCycle network to output these pseudo-labels in a similar fashion as the original study. However, as can be seen in Figure 3.7, the network did not produce good clusters. In the absence of FUCCI markers for the cell cycle, deriving usable alternative pseudo-labels is non-trivial. This opens up the door for another approach that does not depend on these markers.

I next looked at the latent space formed by an autoencoder. I used the encoder and decoder layers of our VAE model and removed the sampling mechanism, effectively turning the network to an autoencoder. I trained this autoencoder to minimize the reconstruction error between the input image and the image generated by the decoder. For this, I used the MSE loss function. Aside from analyzing the manifold to be constructed, I also looked at how good the reconstructed images are. The quality of the reconstructions of the autoencoder will give us assurance that the encoder and decoder are powerful enough to capture the details in the input images and reconstruct from the latent points with an acceptable level of fidelity. In Figure 3.3 I show select input images and their reconstruction from the autoencoder. While some fine-grained details are not present in the reconstructed




Figure 3.8: Clustering on encodings from the autoencoder. Colors correspond to CellCognition labels.

images, key features in the nuclei in different cell cycle stages were preserved. These include the shape of the nuclei, multiple foci, and the distribution of fluorescence during the different stages of the cell cycle. Examining the manifold in Figure 3.8 we see that, similar with the pre-trained CNN models, no distinct cluster groups correspond to the cell cycle stages.

With just the reconstruction loss, the custom network was not able to produce a meaningful manifold. To improve on this, I introduced the variational mechanism of the VAE and treated the latent points as Gaussian distributions in the manifold. For reconstruction, I sampled from these distributions, and used these sampled points as input to the decoder. I further constrained the network to produce distributions as close as possible to the unit Gaussian. This constraint structures the manifold to be smoother as it prevents the network from placing points freely in space.

As discussed in Section 2.3.1, one challenge that arises during training is the balance between the two losses - KLD and reconstruction. In my experiments, while decent reconstructions can be generated using the autoencoder version of the network, turning the sampling on and optimizing the KLD loss results in over-blurred images.




Figure 3.9: VAE-GAN reconstructions. Odd columns show input images. Even columns show the reconstructions. In this image set, while the reconstructions are far from perfect, the characteristic localized high intensity regions during S is captured (top 2 rows).

To address this, I added a learned similarity metric as defined in [56]. For this, a simple discriminator module was added to the trained VAE network. I retrained the network with the additional GAN loss and substituted the MSE loss with the similarity metric from the discriminator for the reconstruction loss. I show in Figure 3.9 the reconstructions from the VAE-GAN network. While the reconstructions still suffer from the blurry trait of VAE-generated images, the model was able to *recover* more fine-grained details when compared with the VAE output. I further revisited this reconstruction improvements in later experiments but for the purpose of constructing the manifold for the cell cycle, I observed that the introduction of the GAN module resulted in a better structure. I attribute this to the learned similarity metric from the discriminator. As discussed in [56], instead of forcing the network to focus on pixel-level details, the substitute metric enables the network to learn semantic similarity in the images. There are still large overlaps between pairs of clusters in the learned manifold but looking closely at individual cell trajectories in this manifold, the clustering on the cell appearance during the different stages becomes more apparent (see Figure 3.11a). From this point onwards, as the latent points in VAEs are Gaussians, UMAP was used with the Wasserstein metric.




Figure 3.10: Clustering on encodings from VAE-GAN. Colors correspond to CellCognition labels.

While single cell analysis can already reveal good clusters representative of the cell cycle stages, looking at the cell population level show that corresponding clusters (and hence cell cycle stages) in pairs of trajectories do not have clear overlaps and instead occupy different subspaces in the manifold. This means that at the population level, the overlap of cell cycle stage clusters are still big despite the clear and distinct clusters at the single-cell level. To remedy this, I fine-tuned the VAE-GAN network and added another loss. The intuition is that single cells should behave and appear similarly when compared with respect to time. Some stages of the cycle may occur longer for some cells but this should be reflected as staying in the same subspace for longer instead of jumping to another. Borrowing from video alignment studies, I used the temporal cycle consistency [120] to affect this constraint. In Figure 3.11, we can see the result of the addition of the TCC loss to the manifold and the individual cell trajectories. While there are still rough jumps from one point to another that we can attribute to the noise in the image data, we can see that corresponding clusters are now located in the same subspaces in the manifold. This improves the semantic detail in the manifold that we can




Figure 3.11: Trajectory plots of individual cells as they undergo the cell life cycle. (a) Despite distinct visual clusters in individual trajectories, there is no spatial cluster correspondences without temporal regularisation. (b) Using TCC resulted in more consistent visual clusters across individual cell tracks. Colors correspond to CellCognition labels.

infer, both at the population and single-cell level.

Aiming to arrive at roughly the same clustering without using the provided labels in the data, I performed GMM clustering on the 16-dimension latent representations of the data points obtained from the trained VAE-GAN model. As discussed in Section 2.5.2, GMM results to soft and probabilistic clusters instead of the hard-clusters in K-Means (shown in Figure 3.13). Since I want to model the continuous progression during the cell cycle, I decided that soft boundaries in clusters is better suited for the task. To classify, I looked at the probability values of each




Figure 3.12: The VAE-GAN-TCC manifold. Colors correspond to CellCognition labels.




Figure 3.13: Initial clustering using K-Means resulted in groupings that do not capture distinct cell cycle stages well.

cluster for a data point. In training the GMM, 8 components were used. This number of components is greater than the number of labels in the dataset (6). The extra components gives flexibility in clustering; as I have not optimised for (and discretised) clusters during training the VAE-GAN model, the model may identify subpopulations of cell images within a cell cycle stage. Instead of forcing the GMM to output 6 component Gaussians to correspond to the class labels, I let

	M	G1	G2	late S	S
M	357	212	43	3	46
G1	13	3269	46	7	566
G2	348	1015	962	30	756
lateS	6	12	3	369	59
S	7	562	103	217	2974

Table 3.1: Confusion matrix for cell cycle stage classification. Rows and columns correspond to different cell cycle states. Values shown are representative examples from the 5 independent training and evaluation runs. Due to the visual similarity of the G1 and G2 classes, the model misclassified many G2 data points as G1.

the GMM form additional clusters that can be manually merged once cluster-class correspondence is established. Figure 3.14 shows the clusters formed by the trained GMM. Of the clusters formed, I observed strongly distinct groupings for the late S (blue), G2 (red), and M (magenta) stages. The yellow and green clusters we observed to contain both early S and mid S images. G1 appeared to be spread across 3 clusters (cyan, light green and purple).

To validate this clustering, I compared the cluster labels for individual points with the ground truth labels provided in the CellCognition [117] study. The ground truth separates the data into 6 classes: early S, mid S, late S, M, G1, and G2. I established correspondence of clusters with class labels by looking at the majority of labels within each cluster. Doing so resulted in 3 clusters matching the G1 class. I merged these clusters to form one G1 cluster and combined two clusters for the early S - mid S (S) cluster, reducing the number of components from 8 to 5. With each component corresponding to a cell cycle class, I evaluated the classification accuracy of the VAE-GAN network using the test images. I extracted the latent representations of these images and used the GMM to label each. With the established cluster-class correspondence, the clustering on the learned manifold resulted in 67.03% accuracy (SD: $\pm 2.12\%$, $n=5$). Table 3.1 shows the confusion matrix for classification of cell cycle stages.

The accuracy rate in identifying the cell cycle stage is below the reported




Figure 3.14: Clusters formed by GMM. Shown on the left are representative points from each cluster.

performance of the SVM model in CellCognition. From Table 3.1, we can see how the trained model confuses G2 as G1. Moreover, the model finds it hard to delineate between two successive stages (M-G1, G1-S) in the cycle, with confused and spread out decision between such stages. The model performed better in identifying S and late S with 76% and 82% accuracy. For studies that only focus on DNA replication, the network’s accuracy may be enough to stage single cell time-lapse videos.

While the accuracy suffers in comparison with CellCognition, the advantage lies on the self-supervision of the model. Even without the labels, the VAE-GAN network was able to capture to some extent the variations in the M, S, and late S stages. Furthermore, additional analyses can be done using the same model to extract additional information to augment the results of this staging. I discuss these in the following subsections.

3.4.2 Identifying Key Cell Cycle Stages from Single Cell Trajectories

An integral part in the analysis of cell cycle experiments is the staging of individual cell trajectories. In my experiments, I used the constructed manifold to identify key stages in the cycle from the trajectories of the individual cells. While the GMM approach described in Section 3.4.1 was able to capture known cell stages as clusters in the manifold, the approach relies on the whole population of cell samples and the number of components to use for clustering. Generally, deciding on the best number of components for GMM is based on some form of prior knowledge. In our case, it was based on the classes available from the previous study and the addition of extra components for flexibility. In this section, I discuss another way to look at the manifold from a single cell perspective.

To do this, I drew trajectories in 2D plots using the first and second UMAP dimensions versus time. While projections in 3D (2 UMAP + time) may be more




Figure 3.15: Plotting UMAP1 versus time reveals some key cell cycle stages.

visually appealing, following single cell trajectories in 3D would suffer from rough curvatures and oscillations as the manifold was learned without any smoothing constraint for single cell trajectories. For this reason, I looked at projections of trajectories on the first and second UMAP dimensions (UMAP1 and UMAP2) with respect to time. I observed that plotting the trajectories of cells with respect to the UMAP dimensions is enough to extract some information that can be useful for staging.

To suppress noise that would otherwise cause rough oscillations, I used B-spline interpolation to smooth the trajectory curve of individual cells. Utilizing UMAP1, I identified the peaks and troughs of each trajectory. Analogous to the Waddington epigenetic landscape, these peaks and troughs correspond to known cell states and the ascent or descent would correspond to transitional states in between. Figure 3.15 shows a UMAP1 plot for a single cell. In the figure, mitosis stages marking a full cycle appeared as the two lowest troughs in the curve. From the first trough, the curve ascends to the first peak which corresponds to the G1 stage. The curve continues towards another trough that shows early S. It proceeds towards the highest peak (late S) and descends to another deep trough (M).

	M	G1	G2	lateS	S
M	357	212	43	3	46
G1	13	3290	25	7	566
G2	348	94	1883	30	756
lateS	6	12	3	369	59
S	7	562	103	217	2974

Table 3.2: Confusion matrix for cell cycle stage classification with G1 and G2 constraint

In Table 3.1, we saw how the VAE-GAN confuses G1 and G2 stages. As G1 and G2 are separated by the S stages, we can enforce a constraint to distinguish between the two growth stages: growth stage that follows S is G2. As we can infer roughly when late S occurs in single cell trajectories using UMAP1, I improve on the classification accuracy from 67.03 (SD: $\pm 2.12\%$, $n=5$) to 73.98 (SD: $\pm 1.60\%$, $n=5$), or around 75% of the supervised model’s accuracy. I do this by correcting the model’s prediction for G1/G2: any predicted G1 occurring after the inferred time for the late S will be reclassified as G2, and any predicted G2 before the late S will be reclassified as G1. Table 3.2 shows the confusion matrix after applying the constraint. Initially, 1015 from the G2 class were misclassified by our model as G1. Similarly, 46 G1 data were misclassified as G2. Applying the G1-G2 constraint brings down the number of misclassifications to 94 and 25 for G2 and G1, respectively. The addition of the constraint brings the accuracy closer to that of the fully-supervised two-channel CellCognition model.

3.4.3 Quantifying Cycle Times and Extracting Population Statistics

In the previous section we have seen that by using the 2D plots involving the UMAP1, the time when mitosis occurred in a single cell trajectory can be inferred. This is done by looking at the deepest troughs in the smoothed curve. In this section, I show that UMAP2 can infer the same information and moreover, present the cell cycle in an oscillating pattern in the UMAP2 dimension.




Figure 3.16: Plotting UMAP2 versus time reveals the oscillating nature of the cell cycle. Blue line is the 10-frame running average. Cycles are marked end to end by sudden drops in the running average curve.

Similar to Section 3.4.2, I plotted the UMAP2 value of each data point with respect to time and extracted the troughs in the curve. Shown in Figure 3.16 is a UMAP2 plot for a single cell. In the figure we can see troughs that correspond to the G1 stage and as the cell continues to grow and undergoes S, and G2, it traverses the UMAP2 axis until peaking during mitosis. This is followed by a drastic drop to another trough, representing the start of another cycle.

Using the smoothed curves for each cell trajectory, we can then quantify the cycle time by computing the period of the cycle. As discussed in Section 3.3.1, each trajectory is composed of 255 images which totals about 25 hours of imaging - more than enough for one complete cycle. I identified for each curve two troughs and computed the time difference between the two. This time difference between the two points can serve as an estimate for cycle time. Comparing with the cycle time computed in CellCognition SVM model, I estimate the cycle time using my approach with RMSE of 18.80 minutes or about 2% difference. I plotted the distribution of the computed cycle times in Figure 3.17 where we can see a unimodal distribution with peak at around 17 hours.




Figure 3.17: The estimated cycle times of the cell trajectories exhibit a unimodal distribution.

3.5 Summary

In this chapter, an approach to construct a manifold of the cell cycle is presented. The approach consists of a VAE-GAN model trained with an additional temporal loss to encapsulate not only the spatial information in image data but also the temporal context available in live imaging as cells progress through their life cycle. The approach requires only the DNA replication factory marker (PCNA) and does not use exhaustive form of annotation during training. This disentanglement from complex experiment setups and expert annotation of data can aid discovery and confirmatory studies for the cell cycle as analysis can be free from annotation bias and prior knowledge of underlying molecular mechanisms.

I highlight the contribution of this work as follows:

1. The learning of the cellular manifold for morphology was disentangled from expert annotation of cell images. This approach in the construction of the manifold based on VAE-GAN readily renders the cell model to discovery and confirmatory studies in a data-driven manner, free from annotation bias and prior knowledge of underlying molecular mechanisms.
2. Even with using only single-channel image data, information relevant to the analysis of the cell cycle can be extracted.
3. A temporal constraint was applied to better structure the manifold for the cell life cycle data.
4. The time-series cell morphology data is represented in meaningful ways that can be utilized in a broad range of analyses, without the need for supervision and for complex imaging and staining.

4

Phenotyping Cells in 3D 4T1 Histology

Contents

4.1	Object Representations in 3D	84
4.1.1	Voxels	84
4.1.2	Surfaces and meshes	85
4.2	4T1 Data	85
4.3	3D Nuclei Segmentation	86
4.4	Profiling DNA Damage in Tumor Tissues	89
4.4.1	Analysis Pipeline	92
4.4.1.1	Regions of Interest	92
4.4.1.2	Isolating γ H2AX+ and CD8+ cells	92
4.4.2	Building cell profiles of DNA damage	95
4.4.2.1	Statistical Texture Features	95

4.4.2.2 Latent Features	96
4.4.3 Pseudo-class labels	98
4.4.4 Comparing cell profiles in control and treated tissues	98
4.5 Revisiting the Cell Cycle	102
4.6 Summary	106

The previous chapter is a good example of analyzing 2D images and extracting enough information from the limited 2D view. Even with one single thin optical section, information relevant to the cell cycle can be extracted. In this chapter, we look at characterizing the morphology of cells in a different type of data. Wherein the last chapter, the focus is on 2D data with temporal element, here we look at a dataset that extends beyond two dimensions.

Analyzing 3D data offers many advantages over 2D. In histology, a 2D image of a thin tissue slice of about 5 micrometers gives us a very shallow view of tissues. Features of cells (such as elongation) and cell groups (such as margins) may not be accurately measured as information captured in 2D may not adequately represent the cells and cell groups that may occupy depths outside of the selected image or tissue sections. This can lead to a failure to examine the correct regions of interest in biopsies and false negatives in diagnoses. With 3D imaging, we can capture information from thicker tissue sections, enabling us to model cells and cell neighborhoods in three dimensions more accurately. In experimental medicine and drug discovery, the effects of drug candidates [121–123] exhibited as patterns in cell and tissue images are profiled. This image-based profiling and phenotyping has been shown to have the potential to expedite the selection of target molecules in drug discovery pipelines [124, 125]. In the same light as the motivation for 3D histology, the capability to observe the effect of treatments in thick tissue samples in experimental medicine can offer a clearer picture of the disease, microenvironments, and drug penetration.

Here, I detail and assess the development of an image processing pipeline for 3D histology. Working within the context of assimilating 3D imaging into clinical pathology workflow in collaboration with the Frederick National Laboratory for Cancer Research, National Institute of Health, this pipeline was used to analyze tissue sections of 4T1 mouse models for DNA damage and immune response. In this chapter, I derived morphological and texture classes for cell phenotypes in our 3D data from which treated and control samples can be differentiated—a primary task that can have great impact on experimental medicine studies.

This chapter is divided into five sections. The first section gives an overview of typical approaches in representing objects in 3D data. This section gives background to how cells in 3D histology can be represented for subsequent analysis. The second section describes the 3D data used in our analyses. The third section describes the assessment of two state-of-the-art 3D segmentation algorithms. The fourth section combines the extracted information from methods described in the preceding sections to phenotype cells in the 3D context. The last section adapts the methods described in the previous chapter for use in the 4T1 data to extract “cell cycle-like” texture classes from the DAPI signal to further enumerate different cell phenotypes.

4.1 Object Representations in 3D

Several approaches to represent objects in 3D exist. Depending on the task at hand, these approaches offer advantages, as well as challenges, in using them in analysis pipelines.

4.1.1 Voxels

Voxels are the most common representation of objects. As a direct extension of 2D pixels in images, voxels are simple representations of 3D objects that can capture complex volumes and shapes. It also renders very easily to manipulations

and transforms. The drawback, however, is the complexity of the representation. Commonly, the vast majority of the voxels correspond to areas not relevant to viewpoints and analysis—in shape analysis, for example. Moreover, with $\Omega(n^3)$ complexity, it scales poorly [126].

In cells analyses, representing cells in 3D are most commonly done as voxels. Segmentation methods such as U-Net [127] directly estimate voxel representations of cells from 3D images.

4.1.2 Surfaces and meshes

An alternative is to represent surfaces as meshes. In this approach, only select points in the object boundaries are kept, together with the adjacencies of these points. This alternative representation with only a subset of the object points addresses the high complexity of voxel representations. However, mesh representations also come with their own drawbacks, owing to their graph-like structure. Generation of meshes is more difficult and transformations are not straight-forward [126]. Because of this, for common tasks, surface and mesh representations are rasterized as voxels, such as in the case of Cellpose and Stardist, which I have introduced in Chapter 2.4.2.2. In our collaboration with the NIH, surface representations as meshes were derived from the segmentation using voxel2mesh [128] for visualization purposes; however, for the analyses to be presented in the next sections, I found that voxel representations provide a more straightforward way to profile nuclei.

4.2 4T1 Data

4T1 is originally isolated by Fred Miller et al [129] as one of the sublines derived from the 410.4 tumor from a single spontaneously arising tumor of an MMTV+ BALB/C mouse foster nursed on a C3H mother (BALB/BfC3H) [130]. It is transplantable, highly tumorigenic, and invasive. It can spontaneously metastasize to several organs

including lungs, liver, brain, and lymph nodes [131]. Due to this nature, it greatly resembles advanced human breast cancer and is therefore used as a model for stage IV breast cancer studies [131, 132]. Using 4T1 models, anti-breast cancer agents such as eupatorin [132] and resveratrol [133, 134] were investigated for their cytotoxic effects. Several studies have also used the 4T1 model to interrogate strategies and mechanisms of activations of CD4+ and CD8+ T cells [135–137].

In the experiments detailed in this thesis, I analyzed 4T1 cells grafted in the hind legs of mice. Tissues were treated with irradiation, indomethacin and L-NAME. Indomethacin is an anti-inflammatory drug. It is commonly used as a prescription medication to reduce fever, pain, stiffness, and swelling from inflammation. Indomethacin works by inhibiting the production of prostaglandin, a vasodilator. On the other hand, L-NAME or N(ω)-nitro-L-arginine methyl ester is an L-arginine analogue used to inhibit nitric oxide synthase (NOS) activity [138]. Control samples were also taken to provide baseline numbers for comparison. Tissues were cleared and then marked for CD8, γ H2AX, pan protein, and DAPI. Figure 4.1 shows a 3D image volume of a 4T1 tissue section.

Image volumes were acquired using a spinning disk confocal microscope at 63X magnification and binning every 3 pixels in the lateral axes. A total of 25 volumes for the control tissues and 49 volumes for the treated tissues were utilized in this work. For most of the experiments, instead of using the original resolution of 605×605 pixels with 391 z -slices, smaller regions of interest were extracted. This was done to focus the analysis on certain regions with the desired level of fluorescence, indicating the presence or occurrence of the biological activities in question.

4.3 3D Nuclei Segmentation

As is the case for a typical image analysis pipeline, the first step in the analysis of the 4T1 data is to isolate individual nuclei. In Chapter 2, I discussed two




Figure 4.1: 3D image volume of a 4T1 tissue section. Shown in red are local (CD8+) immune cells in the region in close proximity to DNA-damaged cells (γ H2AX+ shown in yellow).

state-of-the-art algorithms for 3D segmentation: Cellpose and Stardist. To decide which segmentation algorithm to use, we evaluated and compared the segmentation output of the two methods. This task was done initially by Dr. Ralf Haeusler as part of our NIH collaboration. Building on top of his work, further evaluation and qualitative assessments were done and downstream analyses were integrated with the segmentation pipeline.

As detailed in Chapter 2, Cellpose and Stardist differ in problem formulation and their approach to the solution and determining the best suited method is task- and data-dependent. We found that Stardist segmentation heavily relies on the number of angles and increasing this number resulted in slow impractical processing times. For the succeeding steps in the pipeline, I utilized the Cellpose segmentation

as we deemed it more suitable qualitatively and in terms of speed.

Initial data from our NIH collaborators included an additional marker to delineate nuclei. While DAPI expresses well within the 3D volume of the nuclei, resolving adjacent and overlapping nuclei poses a challenge. In order to measure how well the segmentation algorithms locate, count, and segment individual nuclei, a marker for LaminB was introduced. LaminB serves as a scaffolding for the nuclear envelope. With an Otsu-thresholded LaminB channel as a surrogate for ground truth masks, nuclei segmentation masks were evaluated. In Table 4.1, the precision and recall of the segmentation are reported across the different image tiles used for this assessment. Segmentation precision range from 0.70 to 0.84 while recall scores range from 0.69 to 0.97. While higher scores would be ideal, the segmentation performance was deemed reasonable for use for downstream tasks. Apart from that, visual inspection of results reveal that despite the varying nuclear sizes, the segmentation was able to capture this variation well. This is important particularly for the succeeding analysis on DNA damage across different nuclear patterns and sizes.

Another important measurement that we looked into is the accuracy in the number of segmented cells. Part of the analysis I will discuss in this chapter and the next is the measurement of proximity and identification and enumeration of some cell types. For the analysis to report reliable scores, the number of segmented cell nuclei should be as close as possible to the ground truth. In Table 4.2, the number of segmented cell nuclei for selected tiles is reported with respect to manually counts to evaluate the segmentation. For all tiles, the segmentation correctly accounts for about 70% of the nuclei. Through consultations with our collaborators, this percentage was deemed enough to characterize our tissue data.

Table 4.1: The precision and recall of the segmentation are reported across the different image tiles. The Cellpose segmentation was validated using thresholded LaminB fluorescence images that represent cell boundaries. With precision ranging from 0.70 to 0.84 and recall ranging from 0.69 to 0.97, the segmentation output was deemed adequate for succeeding analysis steps.

Horizontal tile index	Vertical tile index	Precision	Recall
0	0	0.77	0.81
0	1	0.75	0.91
0	2	0.76	0.9
0	3	0.78	0.83
0	4	0.75	0.82
1	0	0.84	0.91
1	1	0.76	0.95
1	2	0.73	0.92
1	3	0.77	0.88
1	4	0.75	0.96
2	0	0.83	0.97
2	1	0.77	0.88
2	2	0.79	0.95
2	3	0.83	0.9
2	4	0.77	0.87
3	0	0.77	0.77
3	1	0.76	0.71
3	2	0.8	0.88
3	3	0.75	0.73
3	4	0.73	0.71
4	0	0.70	0.88
4	1	0.73	0.75
4	2	0.72	0.69
4	3	0.71	0.72
4	4	0.71	0.79

4.4 Profiling DNA Damage in Tumor Tissues

In this section, I will describe a downstream task—cell phenotyping to assess DNA damage—that utilizes the previously described segmentation of nuclei from our 3D 4T1 histology data.

Similar to the cell cycle described in Chapter 3, another cellular process that manifests differently in terms of morphology due to many underlying factors is DNA damage and repair. Normally, a DNA-damaged cell will initiate repair once it

Table 4.2: Segmentation performance on selected tiles of tissue. The number of segmented cell nuclei for selected tiles is validated with respect to manually counts to evaluate the segmentation. The segmentation output using Cellpose correctly accounts for about 70% of the nuclei.

Tile	Confocal layer (z)	Score
	z - coordinate	number segmented cells / all cells
F10	53	15/20 (75%)
F10	130	11/16 (67.85%)
F13	70	17/23 (73.91%)
F15	112	9/10 (90%)
F15	147	18/20 (90%)
F19	145	15/17 (88.23%)

detects alterations such as breaks, fragmentation, translocation and deletions in the DNA. This process can occur across different cell cycle checkpoints and arrests. If excessive damage has occurred, programmed cell death can be triggered. In tumor contexts, studying DNA repair mechanisms is specially important as evidences of impaired repairing capability in tumor cells have been presented [139]. Disregulation of DNA damage response (DDR) promotes genomic instability, increased mutation rate, and intra-tumor heterogeneity [140, 141].

Our understanding of DDR defects in tumors is not yet complete but several anti-cancer therapies already exploit this mutation. Ionizing radiation and chemotherapy aim to induce cell death by damaging DNA. This process however is not guaranteed as tumor cells can still initiate DNA repair pathways and become resistant to these therapies. Inhibiting such pathways and specific targeting of tumor cells can therefore enhance the anticancer effect of DNA damage-based therapy [142]. Targeting specific repair enzymes, e.g. topoisomerases and poly-ADP ribose polymerase (PARP), has proven to be a useful strategy for treating patients with solid tumors, and a number of drugs are widely used in clinical practice. DNA topoisomerase enzymes deal with topological issues linked to DNA transcription and replication. Several mechanisms to inhibit topoisomerase enzymes are explored, including stabilization of the cleavage complex and build up of drugs (anthracycline,

camptothecin) in this complex that ultimately result in DNA strand break and cellular death [143, 144].

The quantification of DNA damage is of particular importance to many previous biological and clinical studies. DNA damage in cells is frequently investigated through imaging with a marker for γ H2AX [145–149]. Phosphorylated histone H2AX or γ H2AX serves as a biomarker for DNA double strand breaks (DSB). Traditionally, the quantification of DNA damage is done by counting the γ H2AX foci in cells either manually or automatically as each DNA DSB has been found to correspond to a γ H2AX focus [150]. Another approach is counting the number of cells in an image field that are positive for γ H2AX signal above some cutoff value assigned on a per cell basis [151, 152].

Alternatively, in the next sections I investigated the use of representation learning and texture features computed from fluorescence intensity statistics (gray-level co-occurrence matrices, GLCM) and latent encoding by adapting the VAE-GAN model discussed previously in Chapter 3 towards building profiles of cells in 3D histology. My approach aims to use representation learning to reduce the need for manual measurements and foci counting to capture similarities and dissimilarities in cells in histology images. Using these discovered features from representation learning in an unsupervised manner would enable building phenotypes without the need for annotation or additional fluorescent markers. I deviated from previous works by profiling DNA damage using visual texture, instead of directly computing for foci density. I analyzed the 3D volumes of 125 μ m thick 4T1 tissue sections provided by collaborators from the NIH-NCI, identified cells and cell clusters with DNA damage exhibited by γ H2AX+ cells in both control and treated samples.

4.4.1 Analysis Pipeline

In order to profile DNA damage in tissues, a number of steps, detailed in the next sections, were performed. First, regions of interest were selected to focus the analysis on regions with high fluorescence intensity. Next, target cell types were isolated from the 3D nuclei segmentation. Cells positive for DNA damaged are then profiled initially using γ H2AX signal, and then jointly with DAPI signal to derive further phenotypes.

4.4.1.1 Regions of Interest

From the acquired 3D confocal images, regions of size $302\text{px} \times 302\text{px} \times 130\text{px}$ ($W \times H \times D$) with high level of fluorescence for γ H2AX and CD8 were manually curated and extracted. Twelve sub-volumes were selected from the control and 14 sub-volumes were selected from the treated samples. The number of cells in each selected region ranged from 500 to 900.

4.4.1.2 Isolating γ H2AX+ and CD8+ cells

As introduced earlier, γ H2AX is a commonly used marker for DNA double strand breaks in cells. H2AX is a histone H2A variant that rapidly becomes phosphorylated at serine 139. This phosphorylation event can be detected using antibodies to γ H2AX and manifests as discrete nuclear foci that can be used to enumerate DSBs in cells [153]. The high sensitivity of γ H2AX detection of DNA DSBs has been successfully demonstrated in primary, cancer cultured cells, and tissues [154–156].

CD8, on the other hand, is a glycoprotein in the surface of cells. It is found on T cells, natural killer cells, cortical thymocytes, and dendritic cells. Serving as a co-receptor to the T-cell receptor, CD8 plays a critical role in T cell signalling and T cell-antigen interactions. Cells positive for CD8 (CD8+), particularly cytotoxic T cells (CTL), are important for inducing death on cancerous and virally infected

cells. Cytotoxicity is mediated by T cells in two ways: perforingranzyme-dependent granule exocytosis mechanism and the FAS-mediated mechanism [157, 158]. These mechanisms are explored in cancer studies to develop anti-cancer therapies. Many of these studies had shown evidence that high local density of CD8+ cells in tumors indicates successful effector function of these cells [159, 160]. Quantifying CD8 signal and enumerating CD8+ cells in images and volumes of tumor tissue are therefore important tasks in this area. In Chapter 5, I will revisit this cell enumeration task as part of our tissue-level analysis.

Figure 4.2E shows a slice of a 3D Cellpose nuclei segmentation of a histology volume using its DAPI channel (shown in Figure 4.2B). From the segmented nuclei using Cellpose, we wanted to focus on cells positive for γ H2AX (see Figure 4.2C). I observed that simple thresholding on the γ H2AX channel is sufficient to segment γ H2AX+ nuclei and using the Cellpose model on the γ H2AX channel actually resulted in “hallucinated” nuclei, false segmentations in areas without nuclei. One drawback of the thresholding, however, is that cell outlines are rougher compared to the Cellpose output. To solve this, I superimposed the thresholded γ H2AX volumes with the previous nuclei segmentation. Cells with a significant overlap in this superimposition are retained in the final γ H2AX segmentation mask. Figure 4.2F shows segmented γ H2AX+ nuclei.

For each of the identified γ H2AX+ nuclei, I extracted bounding boxes to be used for the computation of texture feature representation. A total of 372 γ H2AX+ nuclei were extracted with high confidence. Augmentation by rotation was done to increase this number for training our model.

I also applied a similar segmentation approach to the CD8 channel (see Figures 4.2D and 4.2G) to isolate CD8+ cells. As the immunomarker for CD8 binds on the cell membrane, the nuclei segmentation was dilated first before superimposition to ensure overlap.




Figure 4.2: Using the DAPI channel (a), Cellpose was used to generate segmentation of all nuclei (b). Using this segmentation and threshold masks on the γ H2AX channel (c) and the CD8 channel (e), γ H2AX+ cells and CD8+ are identified (d,f).

4.4.2 Building cell profiles of DNA damage

Techniques that look solely at shape exist and were demonstrated to capture the heterogeneity and variations of cells in tissues. In Chapter 2, I introduced common feature sets that are manually extracted to characterize shape. Wavelet-based features were extracted from breast cancer histopathology images to distinguish between mitotic and non-mitotic cells [161, 162]. Protrusions, concavity, aspect ratio, roughness, and area variance, as well as intensity, density, and orientation of actin cytoskeleton were extracted in [163] to predict cell lines and distinguish between cancer and non-cancer types. However, their approaches rely heavily on accurate segmentation of the nuclear envelope. In our case, as γ H2AX+ is primarily analyzed for its distribution and density across nuclei, texture more than shape would capture this information.

4.4.2.1 Statistical Texture Features

A popular way to quantify texture in images is to compute and analyze GLCM. In Chapter 2, I introduced GLCM as a second-order statistic that captures pairwise relationship of intensity levels within a specific neighborhood size in an image. Its use in medical image analysis is widespread, ranging from brain magnetic resonance images [102] to liver ultrasound [103]. From the computed matrices, features encapsulating the intensity co-occurrences in different ways can be measured [97]. In this work, I used the following GLCM features: energy, contrast, prominence, and correlation.

Energy. Energy measure the uniformity in an image. This is computed by summing the squared elements of the GLCM.

$$Energy = \sum_{i,j} p(i,j)^2, \quad (4.1)$$

where $p(i, j)$ is the frequency of co-occurrence of intensity bins i and j .

Contrast. Contrast looks at variance in intensities between neighboring pixels, given by

$$Contrast = \sum_{i,j} |i - j|^2 p(i, j), \quad (4.2)$$

where $p(i, j)$ is the frequency of co-occurrence of intensity bins i and j .

Prominence. Cluster prominence is a measure of asymmetry or skewness in the GLCM.

$$Prominence = \sum_{i,j} (i + j - \mu_i - \mu_j)^4 p(i, j), \quad (4.3)$$

where $p(i, j)$ is the frequency of co-occurrence of intensity bins i and j , and μ_i and μ_j are mean intensities of intensity bins i and j .

Correlation. Correlation measures how correlated a pixel is to its neighbor over the whole image.

$$Correlation = \sum_{i,j} \frac{(i - \mu_i)(j - \mu_j)p(i, j)}{\sigma_i \sigma_j}, \quad (4.4)$$

where $p(i, j)$ is the frequency of co-occurrence of intensity bins i and j , and μ_i , μ_j , σ_i , σ_j are mean intensities and standard deviations of intensity bins i and j .

4.4.2.2 Latent Features

In a similar fashion as for the cell cycle data in the previous chapter, I explored the use of latent features to characterize the morphology of DNA-damaged cells. The aim is to map the segmented nuclei into a latent space with representations capturing well the density, intensity, and localization of DNA DSB foci and group together nuclei exhibiting similar features. I adapted the VAE-GAN model and replaced the




Figure 4.3: The VAE-GAN architecture (A) used to construct the manifold of DNA damage in cells. The network is composed of three components: encoder, generator, and discriminator. The encoder (B) maps input images to 16-dimensional Gaussian distributions with diagonal covariance. The generator (C) produces a reconstruction from sampled points in the latent space. The discriminator (D) forces the generator to output images as similar to the input as possible. Latent encoding of the image volumes are then clustered to form pseudo-classes of DNA damage subtypes (E).

2D convolutional layers with their 3D convolutional counterpart to accommodate the 3D nature of our histology data. Figure 4.3 shows the structure of the model. The model’s encoder contains three convolutional layers with batch normalization, followed by a series of dense feed-forward layers. For all the experiments, the dimension of the latent space is set to 16. The generator/decoder mirrors the encoder with three transposed convolutional layers. For the discriminator, a network with five convolutional layers with batch normalization was constructed. Feature maps after the fourth convolutional layer are extracted and used for the learned similarity metric. For training, the images of segmented γ H2AX nuclei were resized to and centered on 64×64 patches, and scaled to the $[0,1]$ range. All network modules are trained with ADAM optimizers with learning rate of 0.0001 for 500 epochs.

4.4.3 Pseudo-class labels

Extracted representations for individual cells were then clustered into 5 pseudo texture classes using KMeans. These clusters are then projected and visualized into 2D using principal components analysis (PCA), shown in Figure 4.4. PCA was used in this work, in place of UMAP, as our collaborators routinely use PCA and adopting PCA in this work provides for a better integration with their setup. Since the data lacks expert annotation for repair and apoptotic classes, I aimed to use the formed clusters to find texture similarities in foci across individual cells, acting as class label surrogates.

4.4.4 Comparing cell profiles in control and treated tissues

The last step in my DNA damage analysis is to construct profiles of local neighborhoods of cells to distinguish between regions from treated and control samples. Constructing spherical neighborhoods of various radii lengths, centered on identified γ H2AX+ cells, I looked at the proximity of other γ H2AX+ from

different texture classes.

Using the texture subtypes identified by clustering on GLCM-based features, I constructed heatmaps for control and treated samples. Shown in Figure 4.5A are co-occurrence heatmaps of texture subtypes. Visually, there seems to be a clear difference between control and treated tissues. Clusters of γ H2AX+ cells with labels 0 and 3 are more prominent in treated tissues. On the other hand, control tissues exhibit clusters of γ H2AX+ cells with labels 0 and 2.

I also observed differences in the heatmaps using pseudo-class labels generated from clustering VAE-GAN encodings as shown in Figure 4.5B. The number of neighboring γ H2AX+ cells with labels 2 and 4 are elevated in control tissues while prominence in treated is in γ H2AX+ cells with labels 1 and 2.

In Figure 4.6, representative nuclei from each cluster are shown. If we compare these nuclei images, particularly clusters 1 and 2, we can see that while both seem to represent punctate nuclei, fluorescence signal is dimmer in class 1. Meanwhile, clusters 3 and 4 seem to represent larger regions of bright spots instead of isolated dots. Cluster 0 seems to represent small or fragmented cells that might be from inaccurate segmentation or apoptotic cells. In the heatmaps in Figure 4.5, the texture class represented by cluster 2 seems to be more prevalent in control. Additional markers and further tests can interrogate why this particular texture class appear more as local clusters in control tissues than in treated.

Both representations resulted in distinct profiles for control and treated regions. This demonstrates that the identified texture subtypes capture information that can be utilized for a more refined analysis of tissue sections. It is worth noting a key difference between the two feature sets that may be relevant in the interpretation of results and extending in a succeeding work. While both approaches construct representations from pixel intensities, GLCM features were computed on masked cell images. This was done because the raw (box-cropped) images contain pixels from the background and without masking, the background pixels comprising about




Figure 4.4: Clustering on GLCM representations (A) and VAE-GAN encodings (B) of DNA damage. Based from these clustering, individual cells are profiled into γ H2AX+ texture classes.




Figure 4.5: Heatmaps, using pseudo-classes from GLCM (C) and VAE-GAN(D), show differences in γ H2AX profiles in neighborhoods in control and treated samples.




Figure 4.6: Representative cell images for each of the VAE-GAN DNA damage clusters. Small nuclei with bright spots are represented in cluster 0. Punctate nuclei are represented in clusters 1 and 2, with cluster 2 of higher intensity values. Cluster 3 and 4 show nuclei with solid regions of high intensity.

30%-40% of the images may skew and overpower the pixel information within the nuclei. In contrast, input for the VAE-GAN model did not require masking as doing so may introduce artifacts from the sudden drop in intensities in the nuclei margins. The comparison between the two feature sets should therefore consider this difference in input context. Lacking annotation for the data used in this work, however, explaining the profile differences will be difficult. It will be worth exploring how these texture subtypes correspond to DNA damage response types. Our current collaboration with the NIH aim to explore this in the future but this would entail modifying the preparation, acquisition, and analysis setup to allow for multiplexing of fluorescence markers.

4.5 Revisiting the Cell Cycle

In Chapter 3, we looked at a manifold approach to classify cells into cell cycle stages. In this section, I will describe how the same approach can be applied to generate “cell cycle-inspired” texture classes from the DAPI channel of the 4T1 data. As our data lack specific markers for cell cycle, such as Geminin and Cdt1, alternative approaches to estimate cell cycle stage can be advantageous.

In the context of 4T1 models and our analysis of DNA damage, the capability of classifying cells into approximate cell cycle stages can help us characterize response levels to treatment and possibly interrogate DDR mechanisms across different stages. This has potential impact in experimental medicine as this kind of analysis not only informs of the response to drug candidates but also subdivides this information and visualizes the effect against the different cell cycle states. For example, most tissue images will show cells in the growth stages. If we quantify without accounting for this majority of cells in G1 or G2, we might miss out on opportunities to target and study the effects on other cell cycle stages.

In the preceding section, texture classes from γ H2AX are derived to build profiles

for cells and cell neighborhoods. While this result already demonstrates potential benefit for 3D histology analysis through the visualization of differences in local neighborhoods in control and treated tissues and which texture classes seem to be distinctive and discriminative, further enumeration of cell phenotypes that involve surrogate classes for the cell cycle stages can enrich even more the analysis of 3D tumor tissues and can open up many avenues for biological inquiry. For example, we know that certain mechanisms are preferred and differing conditions for checkpoint activation are necessary for the different cell cycle stages—homologous recombination pathway is activated only in late S/G2 and nonhomologous end-joining in other stages when chromatids are not available for use as templates.

Many previous studies aim to infer cell cycle stage from DAPI signal in 2D images. In [164], cell cycle stage is determined by quantifying the DNA content in segmented nuclei. Two features, nuclear area and total intensity, were computed to estimate DNA content. The method relies on the biological fact that the DNA content within nuclei varies as cells enter the different stages in their life cycle. From the computed features, a GMM with three components, one component each for G1, G2/M, and S stages, is estimated. An assumption that the DNA content in G2 nuclei is twice the DNA content in G1 nuclei was introduced to guide the initialization of the GMM components.

In the attempt to derive ‘cell cycle-inspired’ DAPI classes, I used a VAE-GAN model similar to Chapter 3. As a surrogate for true labels, I used the output of the approach in [164] as pseudo-labels to evaluate this model. Clusters are derived using a GMM with 3 components. Figure 4.7 shows representative nuclei images for each cluster. The cluster labels match 86% of the pseudo-labels.

I next looked at the distribution of cells in our 3D 4T1 data with respect to the derived DAPI texture classes. Histograms for both treated and control tissues show similar profiles, with G1 cells comprising the majority of the analyzed tissues. However, the dominance of G1 cells is much more pronounced in treated tissues.




Figure 4.7: Representative cell images for each of the “cell cycle-inspired” texture classes derived from DAPI fluorescence. Nuclei in the G1-inspired class are small nuclei with few punctations. The G2/M-inspired class shows larger nuclei with more punctations. The S-inspired class represent nuclei with characteristics in between the other two classes.




Figure 4.8: Cell population in our 4T1 data grouped by the derived “cell cycle-inspired” DAPI classes. In both treated and control tissue regions, the majority of cells are in the G1-inspired class and the minority in G2/M-inspired. The dominance of G1-inspired class is more pronounced in treated.

Figure 4.8 shows these similar histograms for the “cell cycle-inspired” texture classes.

Focusing on only the identified γ H2AX+ nuclei tells a different story. In Figure 4.9, we can see that in control tissue regions, half of the γ H2AX+ cells are in the G2/M-inspired texture class, while in treated tissues γ H2AX+ cells are almost evenly split across the three texture classes.

Using the derived cluster labels, I incorporated this “cell cycle-inspired” texture information with the γ H2AX texture classes to further subdivide the cell population




Figure 4.9: γ H2AX+ cell population in our 4T1 data grouped by the derived “cell cycle-inspired” DAPI classes. Half of the γ H2AX+ cells in control tissue regions are in the G2/M-inspired texture class, while in treated tissues γ H2AX+ cells are almost evenly split across the three texture classes

in our 3D 4T1 histology data. Shown in Figure 4.10 are histograms of cells with joint texture information from DAPI and γ H2AX. Profiles for the G2/M-inspired class for control and treated appear very similar, with prevalence on DDR class 1. As for the G1-inspired class, there seems to be a shift of prevalence from DDR class 3 to DR class 1 in treated tissues compared to control. We can observe a similar shift in the S-inspired class from DDR class 0 to DR class 1.

The further subdivision of the cell population with joint analysis results from the DAPI and γ H2AX channels reveal additional insight that has potential benefit for the characterization of tumor tissues and treatment response. However, the constructed phenotypes heavily rely on the textural and visual differences captured by the VAE-GAN model. Grounding these classes to actual biological differences using additional markers ideally will strengthen the insights that the histograms and profiling can offer. Nevertheless, the distinctive profiles from the derived phenotypes can serve as initial indicators of treatment response. In the context of our ongoing collaboration with the NIH, the phenotyping and profiling has pushed forward the direction of our experiments and served as a validation tool for our 3D histology protocol and analysis pipeline.




Figure 4.10: Joint profiling of cells in 4T1 tissue using texture classes derived from γ H2AX (DDR) and DAPI (cell cycle). Profiles for the G2/M-inspired class for control and treated appear very similar, with prevalence on DDR class 1. In the G1-inspired class, a shift of prevalence from DDR class 3 to DDR class 1 in treated tissues compared to control can be seen. A similar shift can also be observed in the S-inspired class from DDR class 0 to DR class 1.

4.6 Summary

In this chapter I presented a different context in which the analysis of morphology poses several challenges. 4T1 is a model cell line widely used in cancer studies and the demonstration of working protocols and analysis pipelines specific to this model has the potential to further enrich its sea of discoveries. I demonstrated using a 4T1 dataset that texture representations using GLCM and VAE-GAN have a potential

impact on DNA repair analysis on tissue microenvironments. Using my approach, I identified γ H2AX+ and learned texture subtypes that show notable differences between control and treated tissues. This can enable a more precise quantification of experiments, particularly on response to anti-cancer drugs and therapies.

I highlight the contribution of this work as follows:

1. State-of-the-art artificial intelligence methods are integrated with thick 3D histology analysis by developing a seamless pipeline, from segmentation of 3D 4T1 histology volumes to profiling DNA damage in local neighborhoods of cells.
2. 3D segmentation methods were tested and demonstrated to achieve high accuracy and sensitivity in localizing and segmenting nuclei in 3D 4T1 histology volumes.
3. Representations of morphological variations of DNA damage in cells are constructed using a VAE-GAN model and GLCM features.
4. Clustering on the extracted features using GLCM and VAE-GAN provided pseudo-labels for texture classes that can be used to characterize tissues and to differentiate between control and treated tissue regions.
5. By constructing surrogate classes for cell cycle stages from DAPI, finer-grained cell phenotypes can be used to further discriminate between control and treated tissues, and visualize the differences in their phenotypic profiles.

5

Characterizing tissue regions and microstructures

Contents

5.1	Marker readout of γH2AX and CD8 markers	110
5.2	Enumerating cells in thick 4T1 tissues	112
5.3	γH2AX-CD8 proximity analysis	115
5.4	Comparing 2D and 3D tissue analysis	117
5.4.1	2D Histology	118
5.4.2	3D Histology	119
5.4.3	Extracting thin slices from 3D 4T1 tissue data	123
5.4.4	Proximity analysis in 2D and 3D	124
5.5	Segmenting microstructures	125

5.5.1	Using orientation of cells in local neighborhoods to infer structure	126
5.5.2	Recognizing vascular regions and tears in tissues	127
5.6	Summary	130

In the previous chapter, I demonstrated several ways by which individual cells can be analyzed through the extraction of distinctive texture representations for phenotypes, particularly in control and treated tissues. While these approaches shed some light on the states of individual cells and their differences when compared, particularly in experimental medicine settings, they cannot represent structures, interactions and processes that extend beyond the single cell. For example, we know that certain interactions between different cell types occur in key biological processes such as DNA damage and immune response. In the last chapter, two cell types involved in these processes, γ H2AX+ and CD8+ cells, were discussed.

In this chapter, I consolidate the extracted information from individual cells in local neighborhoods in tissues and aggregate them towards meaningful neighborhood representations that can be used for the analysis of tissue regions. In particular, I investigated the use of such representations to interrogate interactions between DNA damaged cells and immune cells and to identify key vascular structures in tumor tissues. I also briefly looked into the assessment of analyzing 3D histology data in place of 2D histology in traditional histology analysis pipelines. I argue and provide quantitative evidence of the advantage of tissue representations using 3D data over 2D. I divide the chapter into the following sections: (i) background on transitioning from single cell analysis to tissue-level analysis; (ii) comparison of the readouts of fluorescence markers between control and treated tissues; (iii) enumeration of γ H2AX+ and CD8+ cells and comparison of the frequencies in control and treated tissues; (iv) an approach to aggregate single cell-level information in tissue regions to identify holes in tissues.

5.1 Marker readout of γ H2AX and CD8 markers

In the previous chapter, a 3D 4T1 histology data with control and treated samples was introduced. The tissues were marked with antibodies for CD8 and γ H2AX to help us identify which nuclei have damaged DNA and which ones are immune cells.

In this section I discuss a common but simple task necessary in establishing experimental setups in fluorescence microscopy. While the choice of fluorophores and all experiment details were left to the expertise of our collaborators, I assessed quantitatively the marker intensity and penetration to ensure that downstream tasks such as the analyses presented in the previous chapter, as well as in this chapter, can be done accurately.

Marker readout, or fluorescence readout, is an important technique to detect molecules in a microscopy experimental setup. In the context of our work, the readout of our selected markers can give indication and feedback on study parameters: antibody concentration, tissue preparation, image acquisition. In a much broader sense, in drug screening in experimental medicine, a quick marker readout can demonstrate drug activity and response. In this work, the inclusion of marker readout in the pipeline for 3D histology can perform this exact function.

Our 4T1 data presents a challenge in marker readout. Although a careful process of making the tissue samples as transparent as possible was done, the thick nature of the tissue and light scattering from other factors resulted in uneven intensity throughout the z -stack. Without correction, deeper z -slices appear dim while the first few slices appear brighter. This difference in contrast would be problematic in our analysis tasks, particularly in segmentation. In Chapter 4.4.1.2 I presented the thresholding approach to isolating CD8+ and γ H2AX+ cells. Although Otsu's method is adaptive to image contrast, I found that similar levels of contrast resulted in better masks. Moreover, without contrast correcting based on depth, the thresholding needs to be done on each plane separately.

To address this contrast gradient, I performed a depth-wise contrast stretching on the individual channels. For each of the z -slices, I computed a histogram and selected the intensities of the 5th and 95th percentiles. I also selected a reference z -slice. All other slices will be stretched to the contrast of this reference slice. For most of the image volumes in the dataset, the reference slice is the 9th or 10th slice. I did not use the earlier slices as they contain empty regions due to the uneven tissue surface. The equation below details the formula used to correct the contrast across the image volumes:

$$P_{out} = (P_{in} - c) \left(\frac{b - a}{d - c} \right) + a, \quad (5.1)$$

where P_{out} is the stretched intensity value, P_{in} is the old intensity value, a and b are the intensities of the 5th and 95th percentile of the reference slice and c and d are the intensities of the 5th and 95th percentile of the slice for correction.

Figure 5.1 shows a slice from different imaging depths before and after correction, as well as a YZ plane before and after the same correction.

Using the contrast-corrected images volumes, I then looked at the intensity distribution in control and treated tissues. In our experiments, we expect higher intensity levels of γ H2AX in treated regions. This follows from the established connection between irradiation and DNA DSBs, which γ H2AX measures. Similarly, we expect higher intensity levels of CD8 in treated tissues as immune cells are activated to mount a response to the DNA damage from irradiation. Figure 5.2 compares the intensity values for γ H2AX and CD8 in treated and control regions. Control tissues have mean intensity values of 248 and 249 for CD8 and γ H2AX. Meanwhile, treated tissues express higher fluorescence levels, with 389 and 361 mean intensity values for CD8 and γ H2AX.

The heightened fluorescence intensity levels in treated tissue regions indicate that the experimental parameters induce the desired response and activity in treated




Figure 5.1: Slices from an image volume before and after image stretching. With thick tissue specimen, the acquired image volumes have varying fluorescence intensity levels. To correct for subsequent tasks, the contrast for each slice was corrected with respect to a chosen reference plane. After correction, we can see in (b) and (d) similar contrast levels between different imaging depths.

tissues with a solid baseline control values. This result then justifies the subsequent analysis on our 4T1 histology data detailed in this thesis.

5.2 Enumerating cells in thick 4T1 tissues

Counting cells is another important task in histology analysis. More importantly, counting cells of certain types can characterize tissues and activities within tissues. In this section, I describe how the enumeration of cells in our 4T1 histology data can be




Figure 5.2: Marker readout reveals higher intensities of γ H2AX and CD8 in treated tissue regions compared to control. This supports expectations as irradiation has been linked to DNA double strand breaks and immune response.

another descriptor to characterize and differentiate treated and control tissue regions.

Enumerating cell types has clinical significance. More commonly, this is routinely done in hematological analysis to obtain diagnostic information ranging from anemia and infection to cancer. In our case, and in experimental medicine in general, we look at specific cell types that are exhibiting desired patterns and effects in order to study diseases and treatments.

In Chapter 4.4.1.2, I presented an approach for the segmentation of individual nuclei in the 4T1 image volumes. From this segmentation, we can establish an estimate on the number of cells in an image volume. In Table 4.2, select tiles were manually counted to validate the number of segmented cells.

In some experiments, the number of cells within a region of interest or field of view must exceed a set threshold before performing certain analyses. In experimental medicine, for example, as assays and WSIs can be time-consuming to examine exhaustively, smaller regions are processed instead. In such cases, we focus on regions with high levels but ensure that the scale at which we look still represent a meaningful view of the environment, represented by the constraint on the number of cells and consequently the region volume. This constraint is particularly important in the proximity analysis in Chapter 5.3 as we investigate not only the cell types




Figure 5.3: Enumerating γ H2AX+ and CD8+ cells in control and treated tissues. The increased number of γ H2AX+ and CD8+ supports our expectation of observing increased damage and immune response in treated tissues due to irradiation.

but their interaction at local scales.

As discussed in Chapter 4, γ H2AX is a sensitive marker for DNA DSBs. Looking at the number of cells exhibiting high intensity values for γ H2AX can give us an indication of the degree of damage due to irradiation, especially when we compare it with control baselines. Similarly, the concentration of CD8+ cells in regions can help us understand the mounted immune response to the induced DNA damage.

Using the isolated and identified γ H2AX+ and CD8+ nuclei from the segmentation in Chapter 4, I enumerated each cell type for each of the selected regions of interest in control and treated tissues. In Figure 5.3, we can see how the cell enumeration for γ H2AX and CD8 compare between control and treated tissues. In control, an average 0.5% of the cells in image volumes are CD8+ and about 1.5% are γ H2AX+. On the other hand, 1.25% of cells are CD8+ and 6% are γ H2AX+ in treated regions. This increased ratios of cells positive for CD8 and γ H2AX in treated tissues support the expectation of observing higher degree of damage and heightened immune response due to irradiation.

In the same light as how the heightened levels of intensity in treated tissue regions presented in the previous section demonstrate and substantiate the increased DNA damage and mounted immune response after irradiation, the enumeration

of the γ H2AX+ and CD8+ cell types in regions of interest defines a profile for treated regions that is distinguishable from control.

5.3 γ H2AX-CD8 proximity analysis

In the previous section I looked at what cell types and how many from each cell type can be observed within a region. In this section, I extend this analysis to also incorporate the spatial relationship between these types to indirectly quantify the interactions between cell types. The results of this section was presented as part of our workshop paper for MICCAI-MOVI [165].

While cells are considered the basic blocks of life, they do not exist and function on their own, particularly in complex organisms. Cells communicate with each other through physical and biochemical means to start or react to processes and conditions. In tumor microenvironments, for example, our current understanding has been driven by observing how tumor cells interact with normal cells and the extracellular matrix [166]. Previously in Chapter 4, I have introduced DNA damage response and the involvement of immune cells in this complex biological process. In the context of cancer and experimental medicine, studying the relationship and interaction between DNA damage and the mounted immune response can offer several opportunities. One of the primary aims of our collaborative project with the NIH is to examine in thick 3D histology the role of immune cells and their activation during DNA damage response. With our data labelled for γ H2AX and CD8, initial work was done to establish the relationship between these two cell types.

As interactions can either be direct or mediated by biochemical molecules such as ketones and cytokines, the spatial analysis between CD8+ and γ H2AX+ cells should not be limited to small scales that only account for adjacent cells. Moreover, the occurrence of adjacent CD8+ and γ H2AX+ cells in the 4T1 data is very low. Although the cytotoxic attacks of T cells have been shown in a direct contact




Figure 5.4: Analysis of the proximity of γ H2AX+ and CD8+ reveal a higher concentration of γ H2AX+ cells near activated immune (CD8+) cells. This increase becomes more pronounced in larger neighborhood scales (128px or about $80\mu\text{m}$), suggesting range of interaction possibly longer than previously thought that can be studied in a future work.

mechanism, this has not been exclusively established to be the only mechanism involved as activated T cells also release signalling molecules to recruit other T cells and other immune cell types in the area of detection. In my analysis, I looked at three scales: 32px, 64px, and 128px. This translates to roughly $20\mu\text{m}$, $40\mu\text{m}$, and $80\mu\text{m}$. At these scales, direct interactions, as well as indirect, can be interrogated and quantified.

Figure 5.4 shows the proximity analysis across different scales. At a small scale of 32px or about $20\mu\text{m}$, we can observe an increase in γ H2AX+ cells clustered near CD8+ cells. For most of the CD8+ cells in control tissue regions, there are no adjacent γ H2AX+ cell. On the other hand, an average of one γ H2AX+ cell is adjacent to CD8+ cells in treated tissue regions. At double the scale (64px or about $40\mu\text{m}$), the increase in treated tissue regions becomes more significant. For each of the CD8+ cells in treated tissue regions, an average of 7 γ H2AX+ cells are within its vicinity. At 128px or $80\mu\text{m}$, this trend continues—the distribution of γ H2AX+ near CD8+ in treated is much higher than in control ($p = 6.35 \times 10^{-24}$).

In the previous section the increase in the number of γ H2AX+ and CD8+ in treated tissue regions compared to control was established. The proximity analysis I presented in this section can supplement this information. In addition to the

increase in number in the two cell types, the spatial clustering of these cells implies the existence of an underlying mechanism that involves the two cell types.

To fully capture the interaction between DNA damage and immune cells, we would need to look at apoptosis induced by cytotoxic T cells. Apoptotic cells can also be marked with γ H2AX, similar to DNA damage. They appear as punctated and fragmented nuclei in fluorescence images. In Chapter 4, we looked at DNA damage phenotypes. Grounding and validating one of the texture classes to roughly correspond to apoptotic cells can further enrich the proximity analysis presented in this section as we can also visualize the percentage of apoptotic cells within the neighborhood from the enumerated DNA-damaged cells. In the context of tumor treatments, this additional information can help us quantify even more the effectiveness of drugs and irradiation as we can report not only the induction of DNA damage but also the reduced repair capability in cells. This validation, however, would require an additional marker for apoptosis, such as PD-1 (programmed cell death protein 1). Plans for our future experiments with a multiplexed setup include the use of PD-1.

5.4 Comparing 2D and 3D tissue analysis

The histological analyses presented in this chapter and the previous one work directly on the acquired 3D histology data. However, much of the analyses in histology are still done in 2D in clinical and experimental medicine settings. The long established 2D histology workflow still remains to be the gold standard for diagnostic and experimental purposes. In recent years, we have seen newer technologies in microscopy and tissue preparation that can enable the interrogation of thicker tissues and thus allow the analysis to take place in 3D. As with any type of new technology, its adoption would rely on the adequate demonstration of advantages and improvements over the existing ones. In this section, I argue and provide a

quantification of the advantage of analyzing thick 3D sections over the conventional (thin) 2D sections in the context of the 4T1 histology data I have used in this thesis.

In the following subsections, I will briefly introduce a typical 2D histology workflow, a 3D workflow that improves on the conventional 2D, how I extracted pseudo-thin slices from our 3D 4T1 data, and the comparison of the proximity analysis I have previously presented in both 2D and 3D.

5.4.1 2D Histology

A typical 2D histology workflow involves fixation, dehydration, and embedding of tissues. This is followed by sectioning and staining of the mounted sections [167].

Fixation. Fixation refers to the process of using chemicals such as formalin to preserve the structure of the tissue. This step prevents tissue degradation and hardens the sample, making it easier to be cut into thin sections.

Dehydration. Tissues are then dehydrated to further harden the sample and to ensure proper embedding as paraffin is not miscible in water. This process is typically done by the addition of ethanol, followed by xylene.

Embedding. After the tissue is dehydrated, it is embedded in paraffin or resin to enhance the extraction of cellular structures.

Sectioning. Embedded tissues are then mounted on a microtome and cut into sections, typically 4-5 μ m.

Staining. Thin tissue sections are stained depending on the specific structures and diagnosis wished to be observed. Common stains are hematoxylin and eosin (H&E), immunohistochemical stains, Giemsa, and Gram.




Figure 5.5: Conventional pathology versus non-destructive 3D pathology. (a) The conventional histology workflow was developed over a century ago and involves the use of harsh fixatives and dehydration agents (such as xylene) followed by wax embedding, destructive sectioning and the staining of slide-mounted sections with chromogens, such as H&E, or for immunohistochemistry. In addition to being time consuming and destructive, conventional 2D imaging allows the imaging of only a small fraction of the clinical specimen. (b) Advances in optical clearing and fluorescent labelling, along with high-throughput volumetric microscopy, have enabled entire specimens (for instance, core-needle biopsies) to be imaged in three dimensions, with minimal requirements for tissue processing and mounting. 3D pathology provides rich structural and molecular information from large intact specimens, and preserves valuable clinical specimens for downstream assays (such as DNA and RNA sequencing). Image and caption adapted from [168].

5.4.2 3D Histology

Two-dimensional histology, despite its limitations, remains the standard in clinical practice. In recent years, however, studies demonstrating the potential of 3D histology workflows have been increasing in number [168–170].

The rationale for 3D histology or histopathology over the conventional 2D workflow boils down to three arguments which is best illustrated in Figure 5.6, adapted from [168]. First, as 3D histology involves imaging thick tissue specimens in multiple optical or physical tissue sections, convoluted structures can be more accurately studied. These convoluted structures include vascular and glandular networks, as well as complex-shaped cell surfaces. In viewing these structures in 2D, artefacts, ambiguities, and misrepresentations are to be expected. In Figure 5.6a, a complex vasculature or tubule may appear disjointed or present regions that are completely misleading. The same is the case for elongated nuclei, an example




Figure 5.6: 3D pathology can outperform 2D pathology. (a) For convoluted 3D structures, such as prostate glands, tumor vasculature and kidney tubules, 2D cross-sectional views can be misleading. (b) For distributions of cells and other structures, 2D cross-sectional views might preclude the accurate quantification of complex spatial relationships, such as those found within tumor immune microenvironment and the cancer stem cell niche. (c) For rare objects, such as tumor-initiating cells, lymphovascular invasion and minimal residual disease, 2D sections might not provide adequate sampling to identify and quantify such targets. Image and caption adapted from [168]

of which is shown in Figure 5.7. Enumerating elongated nuclei and quantifying their elongation are often done to study nuclear stability, epithelial-mesenchymal transition (EMT), and protein regulations and expressions [171]. As the elongation can be at an angle with respect to the imaging plane, an elongated nuclei may appear rounder and the degree of elongation be perceived lesser in 2D sections. Second, the spatial distribution of different cell phenotypes can be interrogated better in 3D. With larger volumes of tissues to analyze, we may capture better the heterogeneity in tissues. Lastly, sparse objects or cell types may be sampled better in 3D compared to 2D histology. In our 4T1 data, for example, there are z -slices not containing any CD8+ cell particularly in control tissues. In a 2D histology workflow, if we unluckily sample only in these sections, the enumeration of CD8+ cells and comparison between control and treated tissues may be severely skewed.

3D reconstructions from images of histology sections had been attempted to extend tissue analysis beyond 2D [172–174]. Using machine learning methods, the 2D sections can be registered to make 3D image volumes of tissues. In [175], a high performance computing (HPC) workflow was proposed to give the analysis of 2D




Figure 5.7: An elongated nuclei well captured in 3D but not in 2D XY optical plane. With limited view on 2D (thin) sections, nuclei elongated along the z -axis will appear normal and round. This negatively impacts the accurate characterization of elongation which is important in tumor studies.

histology sections the 3D context of the imaged tissue samples. With a procedure to detect macrostructures in brain tissues and to extract features from these macrostructures, the registration of pairs of histology sections were done. After this, further refinement of the registration is performed to ‘repair’ regions with artefacts (distortion, missing cell information) by using another model trained on intact tissues. Autocuts-LM [176] is a similar system for end-to-end 3D reconstruction and analysis

of tissues. Originally designed for scanning electron microscopy (EM) tomography [177], tissue sections, embedded in resin, were serially sliced with thickness of around 300-400 nm, collected and imaged using lightsheet microscopy. These sections are then aligned by a slice-to-slice image-based reconstruction approach. Image pairs are matched by translation and rotation, followed by a finer-grained registration. Both of these systems demonstrate high performance measures for their registration output. In both approaches however, the reconstruction is ultimately limited by the destructive nature of the sectioning process. Moreover, because the sectioning is irreversible, the tissue becomes unavailable for future and further analyses.

With newer tissue preparation and imaging protocols and more powerful computational tools, we are now able to develop pipelines that aim to address the destructive nature of serial sectioning and provide an alternative to analyze 3D tissues intact. The transparency of biological tissues is paramount to acquiring high contrast and high-resolution images from thick 3D tissues sections [178]. While tissues are composed mainly of water, free and unwanted (for the purposes of imaging) substrates, such as lipids and colored pigments, increase the scattering (and/or absorption) of light that contributes to the reduction of the signal-to-noise ratio in images. But more importantly, with water comprising a big part of the tissues, there is a wide mismatch in refractive indices (RI) of target proteins (1.44) and lipids (1.45) and the surrounding water medium (1.33). This mismatch can further worsen the scattering of light [179–181]. Other factors that facilitated the possibility of 3D histology include more sensitive and deeper penetrating labelling protocols and the increased imaging capability using multiphoton and lightsheet microscopy.

With these advancements, we now have the capability to interrogate thick tissue samples without the need for destructive workflows, opening up avenues to improve existing quantitative analyses on tissues and to investigate biological and medical phenomena which were previously difficult or impossible to examine in 2D

[168, 182]. While 3D histology workflows seem promising, they come with a few challenges. Trade-offs on resolution, cost, speed, and photo-bleaching need to be carefully considered when choosing which 3D imaging technology to use and the choice varies across different tissues and purposes. Confocal imaging offers high contrast and spatial resolution but is constrained by speed and photobleaching. Lightsheet microscopy, on the other hand, acquires images faster, deeper, and with lower photobleaching but often with lower resolution compared to confocal. The amount and size of acquired data also pose problems in storage, quality control, and processing. Lastly, sufficient motivation and evidence towards the adoption of 3D histology is yet to be seen in clinical settings. Analyses and computational tools that highlight the potential of 3D histology in specific biological and clinical contexts are therefore crucial to forward the impact and acceptance of 3D histology.

5.4.3 Extracting thin slices from 3D 4T1 tissue data

In order to compare analysis on thin sections with analysis on thick sections, I simulated the sectioning process in a typical 2D histology workflow. This was done by extracting $5\mu\text{m}$ virtual sections from the 3D 4T1 image volumes and discarding the z -slices at $20\mu\text{m}$ intervals. This is visualized in Figure 5.8. In this sectioning, we can see how many of the cells do not become part of the data for analysis and structures such as the vascular hole in the figure are not well represented by the sampled slices.

I performed the virtual sectioning on each of the image tile, with each tile having 5 thin sections each. I used these extracted thin sections to also analyze the proximity of $\gamma\text{H2AX+}$ and CD8+ cells and compared the results with results using 3D volumes in Chapter 5.3.




Figure 5.8: To enable the comparative analysis between 2D (thin) sections with 3D thick sections, virtual sections of $5\mu\text{m}$ at $20\mu\text{m}$ intervals were extracted. This sectioning simulates a typical 2D histology workflow that samples limited sections and discard significant amount of data. Many cells are ignored in analysis and tissue structures are not represented well.

5.4.4 Proximity analysis in 2D and 3D

The proximity analysis done on 3D image volumes of 4T1 tissues revealed an increase and possible clustering of $\gamma\text{H2AX}+$ cells near $\text{CD8}+$ cells. With the three specified scales; ($20\mu\text{m}$, $40\mu\text{m}$, $80\mu\text{m}$) well within the thickness of the tissue volume, I was able to demonstrate that larger neighborhood regions show more pronounced increase in clustered $\gamma\text{H2AX}+$ cells. In my analysis on the extracted thin sections, the spherical neighborhood regions become almost cylindrical as the thickness of the thin virtual section ($5\mu\text{m}$) is way lower than the set neighborhood scale. Similar to the 3D proximity analysis, I counted the number of $\gamma\text{H2AX}+$ within the vicinity of $\text{CD8}+$ cells. Using 3D, even in the smallest neighborhood scale ($20\mu\text{m}$), there is a significant increase of nearby $\gamma\text{H2AX}+$ cells in treated tissue regions. It would therefore be adequate to show the 2D analysis on this scale as going bigger would only mean larger tissue regions not captured in the virtual sections.

In Figure 5.9, we can see that while there is still an observable increase in the number of clustered $\gamma\text{H2AX}+$ cells near $\text{CD8}+$ cells in $20\mu\text{m}$ neighborhoods, this is not as well pronounced compared to the results using the full 3D image volumes of thick tissue regions. This is further quantified with a p value of 1.81×10^{-1} , which is orders higher than the 3D counterpart shown in Figure 5.4.




Figure 5.9: Comparing proximity analysis on 2D (thin) slices with 3D image volumes. Using extracted thin virtual sections, the same analysis on 3D 4T1 image volume was done on 2D data to demonstrate the advantage of 3D histology in capturing spatial relationships of target cell types. In this work, I quantified the advantage of 3D by showing the more significant increase of neighboring γ H2AX+ cells near CD8+ cells in treated observed using 3D data compared to 2D virtual sections.

This quantification of γ H2AX+ cells proximal to CD8+ cells in both 3D and 2D (thin) histology sections is a demonstration of the advantage of 3D analysis in capturing spatial information that cannot be adequately modelled in 2D. Using the typical thickness of $5\mu\text{m}$ in 2D histology, much of the structural information in tissues is lost. As discussed in Chapter 5.4.2, methods that demonstrate the advantage of 3D histology is crucial in the progress and acceptance of such workflows in clinical and experimental medicine contexts. The comparative results presented in this section serve this specific purpose.

5.5 Segmenting microstructures

Tissue microstructures, as the name implies, are small structures or landmarks formed from organized set of cells and holes in tissues. In this section, I detail a simple way to detect microstructures, specifically vascular regions and tears, in tissues.

Vascular networks in tissues provide nutrients necessary to sustain life in cells and to remove waste products of cellular metabolism. These blood vessels are

comprised of complex tubular networks that can vary in size and composition depending on the anatomical location and function [183]. In tumors, the vascular network is part of the microenvironment that supports and promote their growth. The ability to map and to model these networks, possibly relative to target cell types (tumor, hypoxic, immune, apoptotic, and damaged), can impact how we understand the disease and how we can deliver and administer drugs and treatments [184, 185].

Another feature in tissue that can be observed in histology are tears. Tears and holes are artifacts that can result from mishandling during tissue preparation. During the dehydration step, when not fully dehydrated, the tissues become more prone to tearing as the paraffin wax was not able to infiltrate properly. Excessive exposure to embedding media can also result to excessive tissue hardening. When too hard, cracks during sectioning become more frequent. The microtome can also puncture tissue during sectioning [186].

The identification of both vascular networks and tears in tissues can have great impact in histological analysis. While the identification of these microstructures can easily be done with patch-based classification and segmentation with a supervised model, I will demonstrate that the readily available cell-level information from the segmented nuclei from Chapter 4 has the potential to achieve the same goal.

5.5.1 Using orientation of cells in local neighborhoods to infer structure

From the 3D nuclei segmentation done in Chapter 4, the 3D orientation of each nuclei can be estimated by computing the principal axes of an ellipse that encloses the segmentation. In Figure 5.10, I visualized the segmented nuclei together with the orientation information. From the same volume, we can see holes that are possible vasculature or tears. Using the orientation information I attempt to highlight these holes from the image volumes. I examined the similarity of an individual

cell's orientation with respect to its neighbors. To do this, cosine similarity was computed and averaged to represent the “structure” of the neighborhood of this cell. The assumption is that nuclei near boundaries of vascular holes, tears, and other microstructures would exhibit varying orientations. In the case of vasculature, surrounding cells would orient following the curvature of the hole. In tears, the disruption of tissue continuity would result in abrupt rotations.

In Figure 5.11, two patches are the corresponding similarity score is shown. The first patch shows a region around a cell near a possible tear or vasculature. As depicted by the lines for the principal axes of fitted ellipsoids, the patch region is “disorganized” compared to the second patch with very similar orientations of neighboring cells. Consequently, the cosine similarity score for the first patch (0.43) is lower than the second (0.75). This illustrates the potential of cell orientation to characterize local organization of cells in tissues.

5.5.2 Recognizing vascular regions and tears in tissues

From the computed similarity scores for each nuclei with respect to the orientation of its neighboring cells, I visualized this information to extend the patch-level comparison from the previous section. In Figure 5.12, an XZ plane with two visible holes is shown. We can see that in regions near the holes in tissues, the similarity scores of nuclei are lower (appearing as green) compared to the regions far from these structures (appearing as yellow).

The ability to recognize these structures offer several possibilities in research. The tissue holes that are highlighted using the method can serve as candidate regions that can be verified with additional markers or annotations by experts to confirm if they are of biological importance, particularly if they are part of the vascular network. The identification of the vascular network in tissues can help in modelling microenvironments and look at biological phenomena such as




Figure 5.10: Orientation map of cells in a 3D 4T1 histology volume. Using the nuclei segmentation from 4, the 3D orientation of each individual nucleus was computed and visualized. Lines in the figure represent the principal axes of individual nuclei. Yellow regions correspond to γ H2AX+ nuclei; red regions are CD8+. This orientation map is used to infer local structures in tissues. Visualization using AIDA [187].




Figure 5.11: Cosine similarity scores for patches based on nuclei orientation. Comparing two patches, one near and another far from a possible vasculature, the closer cell has a lower similarity score with its neighbors than the farther cell.




Figure 5.12: Grouping of nuclei based on orientation and locality segments holes in tissues. Using the computed orientation map, individual nuclei are assigned similarity scores with respect to their neighbors' orientation. Cells surrounding holes appear green (low similarity scores) while cells far from them appear yellow (higher similarity scores).

angiogenesis and hypoxia [188]. In the case that the tissue holes turn out to be tears through validation, the information can still be used as a form of quality assessment of handled tissues.

With the limited amount of data points, however, it would be ideal to demonstrate the generalizability of the method to recognize the tissue microstructures in subsequent image acquisitions of 4T1 and other tissue samples in future work. Additional markers are also needed to validate the highlighted regions in the proposed approach. Moreover, as the approach is reliant on the accuracy of the nuclei segmentation, ablation experiments that involve polarity information of

cells will be helpful. Nevertheless, the method in this section shows how cell-level information can be aggregated to characterize local neighborhoods to infer forms of organization—in this case, the orientation of nuclei. This provides for a multi-scale analysis that does not require multiple models trained on different scales. Instead, only a model working at the cellular level to extract information is trained, and on top of which, analysis on a higher biological scale can be built.

5.6 Summary

From the single cell analysis to differentiate between control and treated tissues in our 4T1 histology data, I presented in this chapter multiple ways to use tissue-level information and to aggregate cell-level information within tissue regions to achieve the same goal of characterizing DNA damage and immune cell activity after irradiation of 4T1 tissues. Additionally, I demonstrated how the orientation of nuclei extracted from individual cells in our 4T1 data can be used to recognize holes that can possibly be vascular structures and tears in tissues. These analyses have potential impact on tumor and experimental medicine studies as tissues can be characterized through enumeration of target cell types, regions of hypoxia can be estimated by inferring distance from vascular holes in tissues, and image volumes can be screened for tears. In particular, in future work in our collaborative project with the NIH for which the analysis pipelines described in this thesis were initially designed, we aim to look at using the information from these analyses to further study tissue damage and immune response in relation to hypoxia and cell death.

6

Discussion

In this chapter, I summarize the works presented in the preceding chapters, highlight the contributions of each work in relation to the research aims I laid out in Chapter 1, address the limitations in analysis and scope, and detail possible extensions and improvements in future work.

Chapter 3

In Chapter 3, I addressed the research question: *How can we build a model with fewer input and annotations and integrate temporal information with spatial features to visualize and to analyze the cell cycle?*. I looked at constructing a manifold of the cell cycle. I presented the developed approach involving a VAE-GAN model that incorporates not only the textural information from the cell imaging data but also the temporal information as the cells progress in their life cycle. I worked

under the constraint of using only a single fluorescence channel—following from previous works attempting to “do more with less”. In my attempt, I only used the DNA replication factory marker (PCNA) and modelled the textural changes in the different known cell cycle stages. Furthermore, I did not utilize any exhaustive form of annotation during training. The choice to employ self-supervised methods is inspired by recent developments and advancements in both machine learning and biological studies trying to disentangle model development from certain biases that are inherently present in data—annotation bias and injection of prior knowledge.

My work on the cell cycle demonstrated that despite the constraints imposed on model training, common tasks in the study of the cell cycle can still be reasonably achieved. In staging individual cells, the self-supervised model was able to correctly identify the cell cycle stage in the majority of the cells in our test data. Although the performance of the self-supervised model fails in comparison with its fully supervised counterpart, the potential applications that are now rendered possible due to simplicity and non-reliance on annotations can outweigh this difference in accuracy.

Another task that the model proved to be useful at is in the analysis of single cells to identify key cell cycle stages and to estimate cell cycle time in time-lapse imaging data. The effective demonstration on this task highlights the utility of spatio-temporal data in lieu of multi-channel high resolution images commonly used in cell cycle studies.

Furthermore, in Chapter 4, I have shown how the VAE-GAN model can be used for another kind of data. This demonstrates that the architecture can be adapted to fit different needs. In particular, I went from analyzing textures in 2D images to textures in 3D image volumes.

Chapter 4

In Chapter 4, I interrogated *how we can extend the representation learning pipeline from Chapter 3 for phenotyping cells in 3D histology*. I demonstrated other applications of our VAE-GAN model using a 4T1 dataset. Capturing the texture representations using VAE-GAN, as well as GLCM, can aid the DNA repair analysis on tissue microenvironments. My approach enabled to identify of γ H2AX+ from the segmented nuclei in the 3D histology data and to learn texture subtypes of DNA damage from which control and treated tissues can be profiled comparatively. Similarly, texture classes were derived from DAPI fluorescence to group nuclei into “cell cycle-like” classes. The joint derived classes from DAPI and γ H2AX provided a means to further subdivide the cell population into phenotypes. This profiling has the potential to enable a more precise quantification of experiments, particularly response to anti-cancer drugs and therapies.

Substantial amount of the work detailed in this chapter, as well as some parts from Chapter 5. were presented and published as part of the proceedings of the MOVI workshop at MICCAI [165]. In discussions from this workshop, together with comments from experts in various other settings, major points were raised particularly on the limited data. As this work is part of an ongoing exploratory projects, the acquisition of data is staggered and pivots quickly to accommodate new insights from experiments. While the limited amount of data points used does not necessarily invalidate the results presented in this chapter, it is ideal to work on more samples to increase the statistical power and to validate results. This, however, requires modification of the current preparation and acquisition setup which has been possible only recently. Despite the small sample size, the analysis presented here has supported the hypothesis of our collaborators and steered the direction of our next experiments. In the context of our ongoing project, the adaptability of methods and rapid demonstration of working hypotheses is important. The analysis

on DNA damage substantiated the hypothesis on the higher damage levels and increased repair capability of treated tissues. Even in the absence of additional markers to validate the constructed phenotypes, the profiles for treated and control tissues show significant differences attributable to tissue treatment. Acquiring additional tissue and tumor types, which we aim to accomplish in future works, can further strengthen this result and demonstrate the applicability of representation learning and our VAE-GAN architecture in phenotyping cells in 3D histology images.

Chapter 5

In Chapter 5, I looked into *how we can aggregate and utilize cell-level information to characterize and to differentiate higher-level structures in tissues*. I detailed how I addressed the intensity gradient across the image volume—decreasing contrast as we view deeper optical sections in the thick tissue samples. By stretching the contrast in these sections to match a reference z -slice, similar intensity ranges can be observed in different depths, correcting the problem brought about by the thickness of the tissue samples. I showed that while intensity values were adjusted by contrast stretching, this does not affect the comparative distribution of intensities of γ H2AX in control and treated samples. The mean intensities for treated tissue regions still show higher levels compared to control, following our expectations on tissues treated with irradiation.

Cells positive for γ H2AX and CD8 were also enumerated in control and treated tissue regions. Using the segmentation from Chapter 4, I performed the common task in histology that reports the number of the cell types within a field of view. In this particular work, as we are concerned with characterizing tissue regions for DNA damage and immune response, the number of γ H2AX+ and CD8+ cells after treatment serves as primary indicators of response to treatment. In my analysis, I substantiated this by showing that there is an increase in the number of both

γ H2AX+ and CD8+ cells in treated tissues compared to the baseline.

I also expanded on cell enumeration by looking at how the two cell types are spatially located relative to each other. By observing at different scales, I indirectly quantified the “interaction” between γ H2AX+ and CD8+ cells in the context of the irradiation treatment. While more validation methods that include additional markers to identify further cell subtypes like apoptotic cells and cytotoxic T cells, the multi-scale analysis of proximity of γ H2AX+ and CD8+ serves as another profiling mechanism to visualize differences in control and treated and to demonstrate the spatial clustering of these target cell types due to treatment.

As substantial parts of this thesis revolve within the context of a 3D histology workflow, I demonstrated an advantage of analysis on 3D data over conventional 2D (thin) slices. In comparison to the proximity values I presented done on 3D 4T1 image volumes, the same analysis done on $5\mu\text{m}$ thin slices resulted in less pronounced differences between treated and control tissue regions. This supports the hypothesis that significant amount of information is being discarded in 2D workflows due to its destructive nature and 3D workflows can offer better spatial analysis of tissues.

Another important contribution I presented is an approach to identify holes and possible vascular structures in tissues. In my analysis, I showed how the cell-level information on orientation can be used to infer structured patterns characteristic of the organization of cells to form tissues. The identification of these structures opens up several opportunities. In cancer contexts, we know that tumor microenvironments rely heavily on complex networks of blood vessels to promote their growth. The automatic identification of these vessels can aid studies on angiogenesis, hypoxia, and drug delivery [189–191].

Future Work

The works I presented in this thesis opened rooms for improvement and further and more specific biological inquiries. In particular, the works in Chapters 4 and 5, as part of our ongoing collaboration with the NIH, served as anchor points for the extension project aiming to interrogate additional tumor types. I envision to apply an approach similar to the ones described in this thesis to analyze and to learn representations from these additional 3D histology imaging data.

References

- [1] Mark P. Mattson. *Superior pattern processing is the essence of the evolved human brain*. 2014. DOI: 10.3389/fnins.2014.00265.
- [2] Stanley Greenspan and Stuart Shanker. “The developmental pathways leading to pattern recognition, joint attention, language and cognition”. In: *New Ideas in Psychology* 25.2 (2007). ISSN: 0732118X. DOI: 10.1016/j.newideapsych.2007.02.007.
- [3] John G. Daugman. “Two-dimensional spectral analysis of cortical receptive field profiles”. In: *Vision Research* 20.10 (1980). ISSN: 00426989. DOI: 10.1016/0042-6989(80)90065-6.
- [4] J. P. Jones and L. A. Palmer. “An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex”. In: *Journal of Neurophysiology* 58.6 (1987). ISSN: 00223077. DOI: 10.1152/jn.1987.58.6.1233.
- [5] Yong Cheol Peter Cho, Sungmin Bae, Yongseok Jin, Kevin M. Irick, and Vijaykrishnan Narayanan. “Exploring Gabor filter implementations for visual cortex modeling on FPGA”. In: *Proceedings - 21st International Conference on Field Programmable Logic and Applications, FPL 2011*. 2011. DOI: 10.1109/FPL.2011.63.
- [6] Yingjie Wang, Chin Seng Chua, and Yeong Khing Ho. “Facial feature detection and face recognition from 2D and 3D images”. In: *Pattern Recognition Letters* 23.10 (2002). ISSN: 01678655. DOI: 10.1016/S0167-8655(02)00066-1.
- [7] Xuwen Wang, Xiaoqing Ding, and Changsong Liu. “Gabor filters-based feature extraction for character recognition”. In: *Pattern Recognition* 38.3 (2005). ISSN: 00313203. DOI: 10.1016/j.patcog.2004.08.004.
- [8] Zehang Sun, George Bebis, and Ronald Miller. “On-road vehicle detection using evolutionary gabor filter optimization”. In: *IEEE Transactions on Intelligent Transportation Systems* 6.2 (2005). ISSN: 15249050. DOI: 10.1109/TITS.2005.848363.
- [9] Xinjian Zhu, Xuan He, Pin Wang, Qinghua He, Dandan Gao, Jiwei Cheng, and Baoming Wu. “A method of localization and segmentation of intervertebral discs in spine MRI based on Gabor filter bank”. In: *BioMedical Engineering Online* 15.1 (2016). ISSN: 1475925X. DOI: 10.1186/s12938-016-0146-5.

- [10] Rangaraj M. Rangayyan. “Detection of blood vessels in the retina with multiscale Gabor filters”. In: *Journal of Electronic Imaging* 17.2 (2008). ISSN: 1017-9909. DOI: 10.1117/1.2907209.
- [11] Shengyong Chen, Mingzhu Zhao, Guang Wu, Chunyan Yao, and Jianwei Zhang. “Recent advances in morphological cell image analysis”. In: *Computational and mathematical methods in medicine* 2012 (2012), p. 101536. ISSN: 1748-6718. DOI: 10.1155/2012/101536. URL: <https://doi.org/10.1155/2012/101536>.
- [12] D V A N ESSEN and J KELLY. “Correlation of Cell Shape and Function in the Visual Cortex of the Cat”. In: *Nature* 241.5389 (Feb. 1973), pp. 403–405. ISSN: 1476-4687. DOI: 10.1038/241403a0. URL: <https://doi.org/10.1038/241403a0>.
- [13] Céline Labouesse, Alexander B Verkhovsky, Jean-Jacques Meister, Chiara Gabella, and Benoît Vianay. “Cell Shape Dynamics Reveal Balance of Elasticity and Contractility in Peripheral Arcs”. In: *Biophysical Journal* 108.10 (May 2015), pp. 2437–2447. ISSN: 0006-3495. DOI: 10.1016/j.bpj.2015.04.005. URL: <https://doi.org/10.1016/j.bpj.2015.04.005>.
- [14] Juan C Caicedo, Shantanu Singh, and Anne E Carpenter. “Applications in image-based profiling of perturbations”. In: *Current Opinion in Biotechnology* 39 (2016), pp. 134–142. ISSN: 0958-1669. DOI: <https://doi.org/10.1016/j.copbio.2016.04.003>. URL: <http://www.sciencedirect.com/science/article/pii/S0958166916301112>.
- [15] P Dey. “{C}ancer nucleus: morphology and beyond”. In: *Diagn. Cytopathol.* 38.5 (May 2010), pp. 382–390.
- [16] W H Wolberg, W N Street, and O L Mangasarian. “{I}mportance of nuclear morphology in breast cancer prognosis”. In: *Clin. Cancer Res.* 5.11 (Nov. 1999), pp. 3542–3548.
- [17] Ines Orsolic, Deana Jurada, Nick Pullen, Moshe Oren, Aristides G Eliopoulos, and Sinisa Volarevic. “The relationship between the nucleolus and cancer: Current evidence and emerging paradigms”. In: *Seminars in Cancer Biology* 37-38 (2016), pp. 36–50. ISSN: 1044-579X. DOI: <https://doi.org/10.1016/j.semcancer.2015.12.004>. URL: <http://www.sciencedirect.com/science/article/pii/S1044579X15300043>.
- [18] Theresa J Reape, Elizabeth M Molony, and Paul F McCabe. “Programmed cell death in plants: distinguishing between different modes”. In: *Journal of Experimental Botany* 59.3 (Feb. 2008), pp. 435–444. ISSN: 0022-0957. DOI: 10.1093/jxb/erm258. URL: <https://doi.org/10.1093/jxb/erm258>.
- [19] Peter Goldsborough, Nick Pawlowski, Juan C Caicedo, Shantanu Singh, and Anne E Carpenter. “CytoGAN: Generative Modeling of Cell Images”. In: *bioRxiv* (2017). DOI: 10.1101/227645. URL: <https://www.biorxiv.org/content/early/2017/12/02/227645>.

- [20] Gregory R Johnson, Rory M Donovan-Maiye, and Mary M Maleckar. *Generative Modeling with Conditional Autoencoders: Building an Integrated Cell*. 2017.
- [21] K delas Peñas, M Dmitrieva, J Lefebvre, H Zenner, E Allgeyer, M Booth, D St Johnston, and J Rittscher. “Extracting Axial Depth and Trajectory Trend Using Astigmatism, Gaussian Fitting, and CNNs for Protein Tracking”. In: *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*. 2020, pp. 1634–1637.
- [22] Luca Rappez, Alexander Rakhlin, Angelos Rigopoulos, Sergey Nikolenko, and Theodore Alexandrov. “DeepCycle reconstructs a cyclic cell cycle trajectory from unsegmented cell images using convolutional neural networks”. In: *Molecular systems biology* 16.10 (2020), e9474.
- [23] Novanto Yudistira and Takio Kurita. “Gated spatio and temporal convolutional neural network for activity recognition: towards gated multimodal deep learning”. In: *Eurasip Journal on Image and Video Processing* 2017.1 (2017). ISSN: 16875281. DOI: 10.1186/s13640-017-0235-9.
- [24] Awet Fesseha, Shengwu Xiong, Eshete Derb Emiru, Moussa Diallo, and Abdelghani Dahou. “Text classification based on convolutional neural networks and word embedding for low-resource languages: Tigrinya”. In: *Information (Switzerland)* 12.2 (2021). ISSN: 20782489. DOI: 10.3390/info12020052.
- [25] Matthew C. Chen, Robyn L. Ball, Lingyao Yang, Nathaniel Moradzadeh, Brian E. Chapman, David B. Larson, Curtis P. Langlotz, Timothy J. Amrhein, and Matthew P. Lungren. “Deep learning to classify radiology free-text reports”. In: *Radiology* 286.3 (2018). ISSN: 15271315. DOI: 10.1148/radiol.2017171115.
- [26] S. Milyaev and I. Laptev. “Towards reliable object detection in noisy images”. In: *Pattern Recognition and Image Analysis* 27.4 (2017). ISSN: 15556212. DOI: 10.1134/S1054661817040149.
- [27] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. “Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39.6 (2017). ISSN: 01628828. DOI: 10.1109/TPAMI.2016.2577031.
- [28] Alex Krizhevsky, Ilya Sutskever, Geoffrey E Hinton, Sergey Levine, Chelsea Finn, Trevor Darrell, Pieter Abbeel, Peter Pastor, Alex Krizhevsky, Deirdre Quillen, Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, Martin Riedmiller, Richard S. Sutton, Andrew G. Barto, Sergey Ioffe, and Christian Szegedy. “ImageNet Classification with Deep Convolutional Neural Networks Alex”. In: *Proceedings of the 31st International Conference on Machine Learning (ICML-14)* 3.9 (2012). ISSN: 18726240.

- [29] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. "ImageNet: A large-scale hierarchical image database". In: 2010. DOI: 10.1109/cvpr.2009.5206848.
- [30] Hui Li, Jianfei Cai, Thi Nhat Anh Nguyen, and Jianmin Zheng. "A benchmark for semantic image segmentation". In: *Proceedings - IEEE International Conference on Multimedia and Expo*. 2013. DOI: 10.1109/ICME.2013.6607512.
- [31] Jun Lv, Ming Yang, Jue Zhang, and Xiaoying Wang. "Respiratory motion correction for free-breathing 3D abdominal MRI using CNN-based image registration: A feasibility study". In: *British Journal of Radiology* 91.1083 (2021). ISSN: 1748880X. DOI: 10.1259/bjr.20170788.
- [32] Ahmad Naeem, Muhammad Shoaib Farooq, Adel Khelifi, and Adnan Abid. "Malignant Melanoma Classification Using Deep Learning: Datasets, Performance Measurements, Challenges and Opportunities". In: *IEEE Access* 8 (2020). ISSN: 21693536. DOI: 10.1109/ACCESS.2020.3001507.
- [33] Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. "Dermatologist-level classification of skin cancer with deep neural networks". In: *Nature* 542.7639 (2017). ISSN: 14764687. DOI: 10.1038/nature21056.
- [34] Wenyuan Li, Jiayun Li, Karthik V. Sarma, King Chung Ho, Shiwen Shen, Beatrice S. Knudsen, Arkadiusz Gertych, and Corey W. Arnold. "Path R-CNN for Prostate Cancer Diagnosis and Gleason Grading of Histological Images". In: *IEEE Transactions on Medical Imaging* 38.4 (2019). ISSN: 1558254X. DOI: 10.1109/TMI.2018.2875868.
- [35] Marta Wojciechowska, Stefano Malacrino, Natalia Garcia Martin, Hamid Fehri, and Jens Rittscher. "Early Detection of Liver Fibrosis Using Graph Convolutional Networks". In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 12908 LNCS. 2021. DOI: 10.1007/978-3-030-87237-3_{_}21.
- [36] Fabian Heinemann, Gerald Birk, Tanja Schoenberger, and Birgit Stierstorfer. "Deep neural network based histological scoring of lung fibrosis and inflammation in the mouse model system". In: *PLoS ONE* 13.8 (2018). ISSN: 19326203. DOI: 10.1371/journal.pone.0202708.
- [37] Wenjian Wang, Huaying Wang, Shaokai Yang, Xiaolei Zhang, Xue Wang, Jieyu Wang, Jialiang Lei, Zijian Zhang, and Zhao Dong. "Resolution enhancement in microscopic imaging based on generative adversarial network with unpaired data". In: *Optics Communications* 503 (2022). ISSN: 00304018. DOI: 10.1016/j.optcom.2021.127454.
- [38] Amit Suveer, Anindya Gupta, Gustaf Kylberg, and Ida Maria Sintorn. "Super-resolution reconstruction of transmission electron microscopy images using deep learning". In: *Proceedings - International Symposium on Biomedical Imaging*. Vol. 2019-April. 2019. DOI: 10.1109/ISBI.2019.8759153.

- [39] Tzu An Song, Fan Yang, and Joyita Dutta. “Noise2Void: Unsupervised denoising of PET images”. In: *Physics in Medicine and Biology* 66.21 (2021). ISSN: 13616560. DOI: 10.1088/1361-6560/ac30a0.
- [40] Alexander Krull, Tomá Vičar, Mangal Prakash, Manan Lalit, and Florian Jug. “Probabilistic Noise2Void: Unsupervised Content-Aware Denoising”. In: *Frontiers in Computer Science* 2 (2020). ISSN: 26249898. DOI: 10.3389/fcomp.2020.00005.
- [41] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. “Noise2Noise: Learning image restoration without clean data”. In: *35th International Conference on Machine Learning, ICML 2018*. Vol. 7. 2018.
- [42] Mariia Dmitrieva, Helen L Zenner, Jennifer Richens, Daniel St Johnston, and Jens Rittscher. “Protein Tracking by CNN-Based Candidate Pruning and Two-Step Linking with Bayesian Network”. In: *IEEE International Workshop on Machine Learning for Signal Processing*. 2019.
- [43] Adrian Shajkofci and Michael Liebling. “Spatially-Variant CNN-Based Point Spread Function Estimation for Blind Deconvolution and Depth Estimation in Optical Microscopy”. In: *IEEE Transactions on Image Processing* 29 (2020). ISSN: 19410042. DOI: 10.1109/TIP.2020.2986880.
- [44] Yoshua Bengio, Aaron Courville, and Pascal Vincent. “Representation learning: A review and new perspectives”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35.8 (2013). ISSN: 01628828. DOI: 10.1109/TPAMI.2013.50.
- [45] Dor Bank, Noam Koenigstein, and Raja Giryes. *Autoencoders*. 2021.
- [46] Hiroshi Takahashi, Tomoharu Iwata, Yuki Yamanaka, Masanori Yamada, and Satoshi Yagi. “Variational autoencoder with implicit optimal priors”. In: *33rd AAAI Conference on Artificial Intelligence, AAAI 2019, 31st Innovative Applications of Artificial Intelligence Conference, IAAI 2019 and the 9th AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019*. 2019. DOI: 10.1609/aaai.v33i01.33015066.
- [47] Weonyoung Joo, Wonsung Lee, Sungrae Park, and Il Chul Moon. “Dirichlet Variational Autoencoder”. In: *Pattern Recognition* 107 (2020). ISSN: 00313203. DOI: 10.1016/j.patcog.2020.107514.
- [48] Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. *Importance Weighted Autoencoders*. 2016.
- [49] Tom Rainforth, Adam R Kosiorek, Tuan Anh Le, Chris J Maddison, Maximilian Igl, Frank Wood, and Yee Whye Teh. *Tighter Variational Bounds are Not Necessarily Better*. 2019.
- [50] Christopher P Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. “Understanding disentangling in β -VAE”. In: *arXiv preprint arXiv:1804.03599* (2018).

- [51] Arash Vahdat and Jan Kautz. “Nvae: A deep hierarchical variational autoencoder”. In: *arXiv preprint arXiv:2007.03898* (2020).
- [52] Alexej Klushyn, Nutan Chen, Richard Kurle, Botond Cseke, and Patrick van der Smagt. “Learning hierarchical priors in vaes”. In: *arXiv preprint arXiv:1905.04982* (2019).
- [53] Shengjia Zhao, Jiaming Song, and Stefano Ermon. “Learning Hierarchical Features from Deep Generative Models”. In: *Proceedings of the 34th International Conference on Machine Learning*. Ed. by Doina Precup and Yee Whye Teh. Vol. 70. Proceedings of Machine Learning Research. PMLR, 2017, pp. 4091–4099. URL: <http://proceedings.mlr.press/v70/zhao17c.html>.
- [54] Aäron van den Oord, Oriol Vinyals, and K Kavukcuoglu. “Neural Discrete Representation Learning”. In: *NIPS*. 2017.
- [55] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. “Generating diverse high-fidelity images with vq-vae-2”. In: *Advances in neural information processing systems*. 2019, pp. 14866–14876.
- [56] Anders Boesen Lindbo Larsen, S Kaae Sønderby, Hugo Larochelle, and Ole Winther. “Autoencoding beyond Pixels Using a Learned Similarity Metric”. In: *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*. ICML’16. JMLR.org, 2016, pp. 15581566.
- [57] Muta Tah Hira, M A Razzaque, Claudio Angione, James Scrivens, Saladin Sawan, and Mosharraf Sarker. “Integrated multi-omics analysis of ovarian cancer using variational autoencoders”. In: *Scientific Reports* 11.1 (Mar. 2021), p. 6265. ISSN: 2045-2322. DOI: 10.1038/s41598-021-85285-4. URL: <https://doi.org/10.1038/s41598-021-85285-4>.
- [58] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. “Generative Adversarial Nets”. In: *Advances in Neural Information Processing Systems 27*. Ed. by Z Ghahramani, M Welling, C Cortes, N D Lawrence, and K Q Weinberger. Curran Associates, Inc., 2014, pp. 2672–2680. URL: <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>.
- [59] Nobuyuki Otsu. “THRESHOLD SELECTION METHOD FROM GRAY-LEVEL HISTOGRAMS.” In: *IEEE Trans Syst Man Cybern SMC-9.1* (1979). ISSN: 00189472. DOI: 10.1109/tsmc.1979.4310076.
- [60] Le Hoang Son and Tran Manh Tuan. “A cooperative semi-supervised fuzzy clustering framework for dental X-ray image segmentation”. In: *Expert Systems with Applications* 46 (2016). ISSN: 09574174. DOI: 10.1016/j.eswa.2015.11.001.
- [61] Ahmad Chaddad and Camel Tanougast. “Quantitative evaluation of robust skull stripping and tumor detection applied to axial MR images”. In: *Brain Informatics* 3.1 (2016). ISSN: 21984026. DOI: 10.1007/s40708-016-0033-7.

- [62] Karen Simonyan and Andrew Zisserman. “Very Deep Convolutional Networks for Large-Scale Image Recognition”. In: *International Conference on Learning Representations*. 2015.
- [63] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. “Going Deeper with Convolutions”. In: *CoRR* abs/1409.4842 (2014). arXiv: 1409.4842. URL: <http://arxiv.org/abs/1409.4842>.
- [64] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. “Deep Residual Learning for Image Recognition”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016, pp. 770–778. DOI: 10.1109/CVPR.2016.90.
- [65] Andreas Hinterreiter, Marc Streit, and Bernhard Kainz. “Projective Latent Interventions for Understanding and Fine-Tuning Classifiers”. In: *Interpretable and Annotation-Efficient Learning for Medical Image Computing*. Ed. by Jaime Cardoso, Hien Van Nguyen, Nicholas Heller, Pedro Henriques Abreu, Ivana Isgum, Wilson Silva, Ricardo Cruz, Jose Pereira Amorim, Vishal Patel, Badri Roysam, Kevin Zhou, Steve Jiang, Ngan Le, Khoa Luu, Raphael Sznitman, Veronika Cheplygina, Diana Mateus, Emanuele Trucco, and Samaneh Abbasi. Cham: Springer International Publishing, 2020, pp. 13–22. ISBN: 978-3-030-61166-8.
- [66] Lei Xiang, Weitang Chang, Yong Chen, Weili Lin, Qian Wang, and Dinggang Shen. *Reconstruction in deep learning of highly under-sampled T2-weighted image with T1- weighted image*. 2018.
- [67] Yu Weng, Tianbao Zhou, Yujie Li, and Xiaoyu Qiu. “NAS-Unet: Neural architecture search for medical image segmentation”. In: *IEEE Access* 7 (2019). ISSN: 21693536. DOI: 10.1109/ACCESS.2019.2908991.
- [68] Carsen Stringer, Tim Wang, Michalis Michaelos, and Marius Pachitariu. “Cellpose: a generalist algorithm for cellular segmentation”. In: *Nature Methods* 18.1 (2021). ISSN: 15487105. DOI: 10.1038/s41592-020-01018-x.
- [69] Uwe Schmidt, Martin Weigert, Coleman Broaddus, and Gene Myers. “Cell detection with star-convex polygons”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 11071 LNCS. 2018. DOI: 10.1007/978-3-030-00934-2_{_}30.
- [70] Juan C. Caicedo, Jonathan Roth, Allen Goodman, Tim Becker, Kyle W. Karhohs, Matthieu Broisin, Csaba Molnar, Claire McQuin, Shantanu Singh, Fabian J. Theis, and Anne E. Carpenter. “Evaluation of Deep Learning Strategies for Nucleus Segmentation in Fluorescence Images”. In: *Cytometry Part A* 95.9 (2019). ISSN: 15524930. DOI: 10.1002/cyto.a.23863.
- [71] Weidi Xie, J. Alison Noble, and Andrew Zisserman. “Microscopy cell counting and detection with fully convolutional regression networks”. In: *Computer Methods in Biomechanics and Biomedical Engineering: Imaging and Visualization* 6.3 (2018). ISSN: 21681171. DOI: 10.1080/21681163.2016.1149104.

- [72] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. “Mask R-CNN”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42.2 (2020). ISSN: 19393539. DOI: 10.1109/TPAMI.2018.2844175.
- [73] Martin Weigert, Uwe Schmidt, Robert Haase, Ko Sugawara, and Gene Myers. “Star-convex polyhedra for 3D object detection and segmentation in microscopy”. In: *Proceedings - 2020 IEEE Winter Conference on Applications of Computer Vision, WACV 2020*. 2020. DOI: 10.1109/WACV45572.2020.9093435.
- [74] , Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. “3D U-net: Learning dense volumetric segmentation from sparse annotation”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 9901 LNCS. 2016. DOI: 10.1007/978-3-319-46723-8_{_}49.
- [75] Roberto A. Lotufo, Alexandre X. Falcao, and Francisco A. Zampiroli. “IFT-Watershed from gray-scale marker”. In: *Brazilian Symposium of Computer Graphic and Image Processing*. Vol. 2002-January. 2002. DOI: 10.1109/SIBGRA.2002.1167137.
- [76] María Casanova-Acebes, Erica Dalla, Andrew M. Leader, Jessica LeBerichel, Jovan Nikolic, Blanca M. Morales, Markus Brown, Christie Chang, Leanna Troncoso, Steven T. Chen, Ana Sastre-Perona, Matthew D. Park, Alexandra Tabachnikova, Maxime Dhainaut, Pauline Hamon, Barbara Maier, Catherine M. Sawai, Esperanza Agulló-Pascual, Markus Schober, Brian D. Brown, Boris Reizis, Thomas Marron, Ephraim Kenigsberg, Christine Moussion, Philippe Benaroch, Julio A. Aguirre-Ghiso, and Miriam Merad. “Tissue-resident macrophages provide a pro-tumorigenic niche to early NSCLC cells”. In: *Nature* 595.7868 (2021). ISSN: 14764687. DOI: 10.1038/s41586-021-03651-8.
- [77] Rémy Robinot, Mathieu Hubert, Guilherme Dias de Melo, Françoise Lazarini, Timothée Bruel, Nikaïa Smith, Sylvain Levallois, Florence Larrous, Julien Fernandes, Stacy Gellenoncourt, Stéphane Rigaud, Olivier Gorgette, Catherine Thouvenot, Céline Trébeau, Adeline Mallet, Guillaume Duménil, Samy Gobaa, Raphaël Etournay, Pierre Marie Lledo, Marc Lecuit, Hervé Bourhy, Darragh Duffy, Vincent Michel, Olivier Schwartz, and Lisa A. Chakrabarti. “SARS-CoV-2 infection induces the dedifferentiation of multiciliated cells and impairs mucociliary clearance”. In: *Nature Communications* 12.1 (2021). ISSN: 20411723. DOI: 10.1038/s41467-021-24521-x.
- [78] Steffen Grosser, Jurgen Lippoldt, Linda Oswald, Matthias Merkel, Daniel M. Sussman, Frederic Renner, Pablo Gottheil, Erik W. Morawetz, Thomas Fuhs, Xiaofan Xie, Steve Pawlizak, Anatol W. Fritsch, Benjamin Wolf, Lars Christian Horn, Susanne Briest, Bahriye Aktas, M. Lisa Manning, and Josef A. Kas. “Cell and Nucleus Shape as an Indicator of Tissue Fluidity in Carcinoma”. In: *Physical Review X* 11.1 (2021). ISSN: 21603308. DOI: 10.1103/PhysRevX.11.011033.

- [79] Ariel Waisman, Alessandra Marie Norris, M. Elías Costa, and Daniel Kopinke. “Automatic and unbiased segmentation and quantification of myofibers in skeletal muscle”. In: *Scientific Reports* 11.1 (2021). ISSN: 20452322. DOI: 10.1038/s41598-021-91191-6.
- [80] Shaohe Wang, Kazue Matsumoto, Samantha R. Lish, Alexander X. Cartagena-Rivera, and Kenneth M. Yamada. “Budding epithelial morphogenesis driven by cell-matrix versus cell-cell adhesion”. In: *Cell* 184.14 (2021). ISSN: 10974172. DOI: 10.1016/j.cell.2021.05.015.
- [81] Rongxin Fang, Chenglong Xia, Jennie L Close, Meng Zhang, Jiang He, Zhengkai Huang, Aaron R Halpern, Brian Long, Jeremy A Miller, Ed S Lein, and Xiaowei Zhuang. “Conservation and divergence of cortical cell organization in human and mouse revealed by MERFISH”. In: *Science* 377.6601 (2022), pp. 56–62. DOI: 10.1126/science.abm1741. URL: <https://www.science.org/doi/abs/10.1126/science.abm1741>.
- [82] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. “UMAP: Uniform Manifold Approximation and Projection”. In: *Journal of Open Source Software* 3.29 (2018), p. 861. DOI: 10.21105/joss.00861. URL: <https://doi.org/10.21105/joss.00861>.
- [83] Laurens van der Maaten and Geoffrey Hinton. “Visualizing Data using t-SNE”. In: *Journal of Machine Learning Research* 9.86 (2008), pp. 2579–2605. URL: <http://jmlr.org/papers/v9/vandermaten08a.html>.
- [84] Dmitry Kobak and Philipp Berens. “The art of using t-SNE for single-cell transcriptomics”. In: *Nature Communications* 10.1 (Nov. 2019), p. 5416. ISSN: 2041-1723. DOI: 10.1038/s41467-019-13056-x. URL: <https://doi.org/10.1038/s41467-019-13056-x>.
- [85] Grace X Y Zheng, Jessica M Terry, Phillip Belgrader, Paul Ryvkin, Zachary W Bent, Ryan Wilson, Solongo B Ziraldo, Tobias D Wheeler, Geoff P McDermott, Junjie Zhu, Mark T Gregory, Joe Shuga, Luz Montesclaros, Jason G Underwood, Donald A Masquelier, Stefanie Y Nishimura, Michael Schnall-Levin, Paul W Wyatt, Christopher M Hindson, Rajiv Bharadwaj, Alexander Wong, Kevin D Ness, Lan W Beppu, H Joachim Deeg, Christopher McFarland, Keith R Loeb, William J Valente, Nolan G Ericson, Emily A Stevens, Jerald P Radich, Tarjei S Mikkelsen, Benjamin J Hindson, and Jason H Bielas. “Massively parallel digital transcriptional profiling of single cells”. In: *Nature Communications* 8.1 (Jan. 2017), p. 14049. ISSN: 2041-1723. DOI: 10.1038/ncomms14049. URL: <https://doi.org/10.1038/ncomms14049>.
- [86] Damian Szklarczyk, Annika L Gable, David Lyon, Alexander Junge, Stefan Wyder, Jaime Huerta-Cepas, Milan Simonovic, Nadezhda T Doncheva, John H Morris, Peer Bork, and others. “STRING v11: protein–protein association networks with increased coverage, supporting functional discovery in genome-wide experimental datasets”. In: *Nucleic acids research* 47.D1 (2019), pp. D607–D613.

- [87] Payam Haftbaradaran Esfahani and Ralph Knöll. *Cell shape: effects on gene expression and signaling*. 2020. DOI: 10.1007/s12551-020-00722-4.
- [88] Miaomiao Zhang and Polina Golland. *Statistical shape analysis: From landmarks to diffeomorphisms*. 2016. DOI: 10.1016/j.media.2016.06.025.
- [89] Raj Bahadur Yadav, Naveen K. Nishchal, Arun K. Gupta, and Vinod K. Rastogi. “Retrieval and classification of shape-based objects using Fourier, generic Fourier, and wavelet-Fourier descriptors technique: A comparative study”. In: *Optics and Lasers in Engineering* 45.6 (2007). ISSN: 01438166. DOI: 10.1016/j.optlaseng.2006.11.001.
- [90] Saïd Mahmoudi and Mohammed Benjelloun. “Using global shape descriptors for content medical- based image retrieval”. In: *Handbook of Research on ICTs and Management Systems for Improving Efficiency in Healthcare and Social Care*. Vol. 1-2. 2013. DOI: 10.4018/978-1-4666-3990-4.ch025.
- [91] Li Shen, Hany Farid, and Mark A. McPeck. “Modeling three-dimensional morphological structures using spherical harmonics”. In: *Evolution* 63.4 (2009). ISSN: 00143820. DOI: 10.1111/j.1558-5646.2008.00557.x.
- [92] Dingwen Wang, Shilei Sun, Xi Chen, and Zhiwen Yu. “A 3D shape descriptor based on spherical harmonics through evolutionary optimization”. In: *Neurocomputing* 194 (2016). ISSN: 18728286. DOI: 10.1016/j.neucom.2016.01.081.
- [93] Marianne Chapleau, Christophe Bedetti, Gabriel A. Devenyi, Signy Sheldon, Howie J. Rosen, Bruce L. Miller, Maria Luisa Gorno-Tempini, Mallar M. Chakravarty, and Simona M. Brambati. “Deformation-based shape analysis of the hippocampus in the semantic variant of primary progressive aphasia and Alzheimer’s disease”. In: *NeuroImage: Clinical* 27 (2020). ISSN: 22131582. DOI: 10.1016/j.nicl.2020.102305.
- [94] Naofumi Kawahira, Daisuke Ohtsuka, Naoki Kida, Ken ichi Hironaka, and Yoshihiro Morishita. “Quantitative Analysis of 3D Tissue Deformation Reveals Key Cellular Mechanism Associated with Initial Heart Looping”. In: *Cell Reports* 30.11 (2020). ISSN: 22111247. DOI: 10.1016/j.celrep.2020.02.071.
- [95] Christian Waldenberg, Hanna Hebelka, Helena Brisby, and Kerstin Magdalena Lagerstrand. “MRI histogram analysis enables objective and continuous classification of intervertebral disc degeneration”. In: *European Spine Journal* 27.5 (2018). ISSN: 14320932. DOI: 10.1007/s00586-017-5264-7.
- [96] Ayushman Ramola, Amit Kumar Shakya, and Dai Van Pham. “Study of statistical methods for texture analysis and their modern evolutions”. In: *Engineering Reports* 2.4 (2020). ISSN: 25778196. DOI: 10.1002/eng2.12149.
- [97] Robert M. Haralick, Its’hak Dinstein, and K. Shanmugam. “Textural Features for Image Classification”. In: *IEEE Transactions on Systems, Man and Cybernetics* SMC-3.6 (1973). ISSN: 21682909. DOI: 10.1109/TSMC.1973.4309314.

- [98] C. C. Gotlieb and H. E. Kreyszig. “Texture descriptors based on co-occurrence matrices”. In: *Computer Vision, Graphics, & Image Processing* 51.1 (1990). ISSN: 0734189X. DOI: 10.1016/S0734-189X(05)80063-5.
- [99] Min Wang, Wenzhi Cao, Qingsong Guan, Gaojie Wu, Chao Jiang, Yiyun Yan, and Xiaoling Su. “Potential of texture metrics derived from high-resolution PLEIADES satellite data for quantifying aboveground carbon of *Kandelia candel* mangrove forests in Southeast China”. In: *Wetlands Ecology and Management* 26.5 (2018). ISSN: 15729834. DOI: 10.1007/s11273-018-9610-2.
- [100] Naveed Iqbal, Rafia Mumtaz, Uferah Shafi, and Syed Mohammad Hassan Zaidi. “Gray level co-occurrence matrix (GLCM) texture based crop classification using low altitude remote sensing platforms”. In: *PeerJ Computer Science* 7 (2021). ISSN: 23765992. DOI: 10.7717/PEERJ-CS.536.
- [101] Jagdish Lal Raheja, Sunil Kumar, and Ankit Chaudhary. “Fabric defect detection based on GLCM and Gabor filter: A comparison”. In: *Optik* 124.23 (2013). ISSN: 00304026. DOI: 10.1016/j.ijleo.2013.05.004.
- [102] T. R. Sivapriya, V. Saravanan, and P. Ranjit Jeba Thangaiah. “Texture analysis of brain MRI and classification with BPN for the diagnosis of dementia”. In: *Communications in Computer and Information Science*. Vol. 204 CCIS. 2011. DOI: 10.1007/978-3-642-24043-0{_}56.
- [103] Guang Ming Xian. “An identification method of malignant and benign liver tumors from ultrasonography based on GLCM texture features and fuzzy SVM”. In: *Expert Systems with Applications* 37.10 (2010). ISSN: 09574174. DOI: 10.1016/j.eswa.2010.02.067.
- [104] O. S. Al-Kadi. “Combined statistical and model based texture features for improved image classification”. In: *IET Conference Publications*. 540 CP. 2008. DOI: 10.1049/cp:20080455.
- [105] Paolo Napoletano. “Hand-crafted vs learned descriptors for color texture classification”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 10213 LNCS. 2017. DOI: 10.1007/978-3-319-56010-6{_}22.
- [106] Angel Cruz-Roa, Ajay Basavanahally, Fabio González, Hannah Gilmore, Michael Feldman, Shridar Ganesan, Natalie Shih, John Tomaszewski, and Anant Madabhushi. “Automatic detection of invasive ductal carcinoma in whole slide images with convolutional neural networks”. In: *Medical Imaging 2014: Digital Pathology*. Vol. 9041. 2014. DOI: 10.1117/12.2043872.
- [107] Alison O’Neil, Matthew Shepherd, Erin Beveridge, and Keith Goatman. “A comparison of texture features versus deep learning for image classification in interstitial lung disease”. In: *Communications in Computer and Information Science*. Vol. 723. 2017. DOI: 10.1007/978-3-319-60964-5{_}65.

- [108] Subbareddy Maddika, Sudharsana Rao Ande, Soumya Panigrahi, Ted Paranjothy, Kazimierz Weglarczyk, Anne Zuse, Mehdi Eshraghi, Kamala D Manda, Emilia Wiechec, and Marek Los. “Cell survival, cell death and cell cycle pathways are interconnected: Implications for cancer therapy”. In: *Drug Resistance Updates* 10.1 (2007), pp. 13–29. ISSN: 1368-7646. DOI: <https://doi.org/10.1016/j.drug.2007.01.003>. URL: <https://www.sciencedirect.com/science/article/pii/S1368764607000040>.
- [109] Jiri Lukas, Claudia Lukas, and Jiri Bartek. “Mammalian cell cycle checkpoints: signalling pathways and their organization in space and time”. In: *DNA Repair* 3.8 (2004), pp. 997–1007. ISSN: 1568-7864. DOI: <https://doi.org/10.1016/j.dnarep.2004.03.006>. URL: <https://www.sciencedirect.com/science/article/pii/S1568786404000655>.
- [110] Christina R Kahl and Anthony R Means. “Regulation of Cell Cycle Progression by Calcium/Calmodulin-Dependent Pathways”. In: *Endocrine Reviews* 24.6 (2003), pp. 719–736. ISSN: 0163-769X. DOI: 10.1210/er.2003-0008. URL: <https://doi.org/10.1210/er.2003-0008>.
- [111] B P Coe, W W Lockwood, L Girard, R Chari, C MacAulay, S Lam, A F Gazdar, J D Minna, and W L Lam. “Differential disruption of cell cycle pathways in small cell and non-small cell lung cancer”. In: *British Journal of Cancer* 94.12 (June 2006), pp. 1927–1935. ISSN: 1532-1827. DOI: 10.1038/sj.bjc.6603167. URL: <https://doi.org/10.1038/sj.bjc.6603167>.
- [112] *Cell Cycle*. 2021. URL: <https://www.genome.gov/genetics-glossary/Cell-Cycle>.
- [113] ROSS A MARKLEIN, MATTHEW W KLINKER, KATHERINE A DRAKE, HANNAH G POLIKOWSKY, ELIZABETH C LESSEY-MORILLON, and STEVEN R BAUER. “Morphological profiling using machine learning reveals emergent subpopulations of interferon- γ stimulated mesenchymal stromal cells that predict immunosuppression”. In: *Cytotherapy* 21.1 (2019), pp. 17–31. ISSN: 1465-3249. DOI: <https://doi.org/10.1016/j.jcyt.2018.10.008>. URL: <http://www.sciencedirect.com/science/article/pii/S1465324918306467>.
- [114] Ariel Waisman, Alejandro [La Greca], Alan M Möbbs, María Agustina Scarafia, Natalia L [Santín Velazque], Gabriel Neiman, Lucía N Moro, Carlos Luzzani, Gustavo E Sevlever, Alejandra S Guberman, and Santiago G Miriuka. “Deep Learning Neural Networks Highly Predict Very Early Onset of Pluripotent Stem Cell Differentiation”. In: *Stem Cell Reports* 12.4 (2019), pp. 845–859. ISSN: 2213-6711. DOI: <https://doi.org/10.1016/j.stemcr.2019.02.004>. URL: <http://www.sciencedirect.com/science/article/pii/S2213671119300529>.
- [115] G Gut, M D Tadmor, D Pe’er, L Pelkmans, and P Liberali. “Trajectories of cell-cycle progression from fixed cell populations”. In: *Nat Methods* 12.10 (Oct. 2015), pp. 951–954.

- [116] P Eulenberg, N Köhler, T Blasi, A Filby, A E Carpenter, P Rees, F J Theis, and F A Wolf. “{R}econstructing cell cycle and disease progression using deep learning”. In: *Nat Commun* 8.1 (2017), p. 463.
- [117] Michael Held, Michael H A Schmitz, Bernd Fischer, Thomas Walter, Beate Neumann, Michael H Olma, Matthias Peter, Jan Ellenberg, and Daniel W Gerlich. “CellCognition: time-resolved phenotype annotation in high-throughput live cell imaging”. In: *Nature methods* 7.9 (2010), pp. 747–754.
- [118] Dennis Bähr, Dennis Eschweiler, Anuk Bhattacharyya, Daniel Moreno-Andrés, Wolfram Antonin, and Johannes Stegmaier. *CellCycleGAN: Spatiotemporal Microscopy Image Synthesis of Cell Populations using Statistical Shape Models and Conditional GANs*. 2021. arXiv: 2010.12011 [eess.IV].
- [119] Diederik P Kingma and Max Welling. “An Introduction to Variational Autoencoders”. In: *CoRR* abs/1906.0 (2019). URL: <http://arxiv.org/abs/1906.02691>.
- [120] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. “Temporal cycle-consistency learning”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 1801–1810.
- [121] Jonne Rietdijk, Marianna Tampere, Aleksandra Pettke, Polina Georgiev, Maris Lapins, Ulrika Warpmann-Berglund, Ola Spjuth, Marjo Riitta Puumalainen, and Jordi Carreras-Puigvert. “A phenomics approach for antiviral drug discovery”. In: *BMC Biology* 19.1 (2021). ISSN: 17417007. DOI: 10.1186/s12915-021-01086-1.
- [122] Yushi Futamura, Makoto Kawatani, Sayaka Kazami, Kenichi Tanaka, Makoto Muroi, Takeshi Shimizu, Koji Tomita, Nobumoto Watanabe, and Hiroyuki Osada. “Morphobase, an Encyclopedic Cell Morphology Database, and Its Use for Drug Target Identification”. In: *Chemistry & Biology* 19.12 (2012), pp. 1620–1630. ISSN: 1074-5521. DOI: <https://doi.org/10.1016/j.chembiol.2012.10.014>. URL: <https://www.sciencedirect.com/science/article/pii/S1074552112004152>.
- [123] Miriam B. Brandl, Eddy Pasquier, Fuhai Li, Dominik Beck, Sufang Zhang, Hong Zhao, Maria Kavallaris, and Stephen T.C. Wong. “Computational analysis of image-based drug profiling predicts synergistic drug combinations: Applications in triple-negative breast cancer”. In: *Molecular Oncology* 8.8 (2014). ISSN: 18780261. DOI: 10.1016/j.molonc.2014.06.007.
- [124] Jungseog Kang, Chien Hsiang Hsu, Qi Wu, Shanshan Liu, Adam D. Coster, Bruce A. Posner, Steven J. Altschuler, and Lani F. Wu. “Improving drug discovery with high-content phenotypic screens by systematic selection of reporter cell lines”. In: *Nature Biotechnology* 34.1 (2016). ISSN: 15461696. DOI: 10.1038/nbt.3419.

- [125] Carmen Mirabelli, Jesse W. Wotring, Charles J. Zhang, Sean M. McCarty, Reid Fursmidt, Carla D. Pretto, Yuanyuan Qiao, Yuping Zhang, Tristan Frum, Namrata S. Kadambi, Anya T. Amin, Teresa R. OMeara, Jason R. Spence, Jessie Huang, Konstantinos D. Alysandratos, Darrell N. Kotton, Samuel K. Handelman, Christiane E. Wobus, Kevin J. Weatherwax, George A. Mashour, Matthew J. OMeara, Arul M. Chinnaiyan, and Jonathan Z. Sexton. “Morphological cell profiling of SARS-CoV-2 infection identifies drug repurposing candidates for COVID-19”. In: *Proceedings of the National Academy of Sciences of the United States of America* 118.36 (2021). ISSN: 10916490. DOI: 10.1073/pnas.2105815118.
- [126] Sai Rajeswar, Fahim Mannan, Florian Golemo, Jérôme Parent-Lévesque, David Vazquez, Derek Nowrouzezahrai, and Aaron Courville. “Pix2Shape: Towards Unsupervised Learning of 3D Scenes from Images Using a View-Based Representation”. In: *International Journal of Computer Vision* 128.10-11 (2020). ISSN: 15731405. DOI: 10.1007/s11263-020-01322-1.
- [127] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. “U-Net: Convolutional Networks for Biomedical Image Segmentation”. In: *CoRR* abs/1505.0 (2015). URL: <http://arxiv.org/abs/1505.04597>.
- [128] Udaranga Wickramasinghe, Edoardo Remelli, Graham Knott, and Pascal Fua. “Voxel2Mesh: 3D Mesh Model Generation from Volumetric Data”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 12264 LNCS. 2020. DOI: 10.1007/978-3-030-59719-1_{_}30.
- [129] Cheryl J. Aslakson and Fred R. Miller. “Selective Events in the Metastatic Process Defined by Analysis of the Sequential Dissemination of Subpopulations of a Mouse Mammary Tumor”. In: *Cancer Research* 52.6 (1992). ISSN: 15387445.
- [130] Daniel L. Dexter, Henryk M. Kowalski, Beverly A. Blazar, Zuzana Fligiel, Renee Vogel, H. Gloria, and H. Heppner. “Heterogeneity of Tumor Cells from a Single Mouse Mammary Tumor”. In: *Cancer Research* 38.10 (1978). ISSN: 15387445.
- [131] Barbara Schrörs, Sebastian Boegel, Christian Albrecht, Thomas Bukur, Valesca Bukur, Christoph Holtsträter, Christoph Ritzel, Katja Manninen, Arbel D. Tadmor, Mathias Vormehr, Ugur Sahin, and Martin Löwer. “Multi-Omics Characterization of the 4T1 Murine Mammary Gland Tumor Model”. In: *Frontiers in Oncology* 10 (2020). ISSN: 2234943X. DOI: 10.3389/fonc.2020.01195.
- [132] Nursyamirah Abd Razak, Swee Keong Yeap, Noorjahan Banu Alitheen, Wan Yong Ho, Chean Yeah Yong, Sheau Wei Tan, Wen Siang Tan, and Kamariah Long. “Eupatorin Suppressed Tumor Progression and Enhanced Immunity in a 4T1 Murine Breast Cancer Model”. In: *Integrative Cancer Therapies* 19 (2020). ISSN: 1552695X. DOI: 10.1177/1534735420935625.

- [133] Kathleen Bove, David W. Lincoln, and Min Fu Tsan. “Effect of resveratrol on growth of 4T1 breast cancer cells in vitro and in vivo”. In: *Biochemical and Biophysical Research Communications* 291.4 (2002). ISSN: 0006291X. DOI: 10.1006/bbrc.2002.6554.
- [134] Yuan Zou, Xiangtao Wang, Dongdong Bi, Jingxin Fu, Jianwei Han, Yifei Guo, Li Feng, and Meihua Han. “Pterostilbene nanoparticles with small particle size show excellent anti-breast cancer activity in vitro and in vivo”. In: *Nanotechnology* 32.32 (2021). ISSN: 13616528. DOI: 10.1088/1361-6528/abfdec.
- [135] Jared M. Newton, Jose H. Flores-Arredondo, Sarah Suki, Matthew J. Ware, Martyna Krzykawska-Serda, Mahdi Agha, Justin J. Law, Andrew G. Sikora, Steven A. Curley, and Stuart J. Corr. “Non-Invasive Radiofrequency Field Treatment of 4T1 Breast Tumors Induces T-cell Dependent Inflammatory Response”. In: *Scientific Reports* 8.1 (2018). ISSN: 20452322. DOI: 10.1038/s41598-018-21719-w.
- [136] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip H S Torr. “Conditional Random Fields as Recurrent Neural Networks”. In: *CoRR* abs/1502.0 (2015). URL: <http://arxiv.org/abs/1502.03240>.
- [137] Junxia Gao, Xianlin Yuan, Jia Yuan, and Liangping Li. “Complete rejection of large established breast cancer by local immunochemotherapy with T cell activation against neoantigens”. In: *Cancer Immunology, Immunotherapy* 70.11 (2021). ISSN: 14320851. DOI: 10.1007/s00262-021-02919-2.
- [138] Alan G. Fantel, N. Nekahi, T. H. Shepard, L. M. Cornel, A. S. Unis, and R. J. Lemire. “The teratogenicity of N(ω)-nitro-L-arginine methyl ester (L-NAME), a nitric oxide synthase inhibitor, in rats”. In: *Reproductive Toxicology* 11.5 (1997). ISSN: 08906238. DOI: 10.1016/S0890-6238(97)00033-6.
- [139] Jehad F. Alhmoud, John F. Woolley, Ala Eddin Al Moustafa, and Mohammed Imad Malki. *DNA damage/repair management in cancers*. 2020. DOI: 10.3390/cancers12041050.
- [140] Young Kwang Chae, Jonathan F. Anker, Benedito A. Carneiro, Sunandana Chandra, Jason Kaplan, Aparna Kalyan, Cesar A. Santa-Maria, Leonidas C. Plataniias, and Francis J. Giles. “Genomic landscape of DNA repair genes in cancer”. In: *Oncotarget* 7.17 (2016). ISSN: 19492553. DOI: 10.18632/oncotarget.8196.
- [141] Rebecca A. Burrell, Nicholas McGranahan, Jiri Bartek, and Charles Swanton. *The causes and consequences of genetic heterogeneity in cancer evolution*. 2013. DOI: 10.1038/nature12625.
- [142] Lan Ya Li, Yi Di Guan, Xi Sha Chen, Jin Ming Yang, and Yan Cheng. *DNA Repair Pathways in Cancer Therapy and Resistance*. 2021. DOI: 10.3389/fphar.2020.629266.

- [143] Kirk E. Hevener, Tatsiana A. Verstak, Katie E. Lutat, Daniel L. Riggsbee, and Jeremiah W. Mooney. *Recent developments in topoisomerase-targeted cancer chemotherapy*. 2018. DOI: 10.1016/j.apsb.2018.07.008.
- [144] Kamila Buzun, Anna Bielawska, Krzysztof Bielawski, and Agnieszka Gornowicz. *DNA topoisomerases as molecular targets for anticancer drugs*. 2020. DOI: 10.1080/14756366.2020.1821676.
- [145] Robert J. Kinders, Melinda Hollingshead, Scott Lawrence, Uping Ji, Brian Tabb, William M. Bonner, Yves Pommier, Larry Rubinstein, Yvonne A. Evrard, Ralph E. Parchment, Joseph Tomaszewski, and James H. Doroshow. “Development of a validated immunofluorescence assay for γ H2AX as a pharmacodynamic marker of topoisomerase I inhibitor activity”. In: *Clinical Cancer Research* 16.22 (2010). ISSN: 10780432. DOI: 10.1158/1078-0432.CCR-09-3076.
- [146] Christophe E. Redon, Asako J. Nakamura, Yong Wei Zhang, Jiuping Ji, William M. Bonner, Robert J. Kinders, Ralph E. Parchment, James H. Doroshow, and Yves Pommier. *Histone γ H2AX and poly(ADP-ribose) as clinical pharmacodynamic biomarkers*. 2010. DOI: 10.1158/1078-0432.CCR-10-0523.
- [147] Khanh Do, Deborah Wilsker, Jiuping Ji, Jennifer Zlott, Tomoko Freshwater, Robert J. Kinders, Jerry Collins, Alice P. Chen, James H. Doroshow, and Shivaani Kummar. “Phase I study of single-agent AZD1775 (MK-1775), a weel kinase inhibitor, in patients with refractory solid tumors”. In: *Journal of Clinical Oncology* 33.30 (2015). ISSN: 15277755. DOI: 10.1200/JCO.2014.60.4009.
- [148] Angie B. Dull, Deborah Wilsker, Melinda Hollingshead, Christina Mazcko, Christina M. Annunziata, Amy K. LeBlanc, James H. Doroshow, Robert J. Kinders, and Ralph E. Parchment. “Development of a quantitative pharmacodynamic assay for apoptosis in fixed tumor tissue and its application in distinguishing cytotoxic drug-induced DNA double strand breaks from DNA double strand breaks associated with apoptosis”. In: *Oncotarget* 9.24 (2018). ISSN: 19492553. DOI: 10.18632/oncotarget.24936.
- [149] Deborah F. Wilsker, Allison M. Barrett, Angie B. Dull, Scott M. Lawrence, Melinda G. Hollingshead, Alice Chen, Shivaani Kummar, Ralph E. Parchment, James H. Doroshow, and Robert J. Kinders. “Evaluation of pharmacodynamic responses to cancer therapeutic agents using DNA damage markers”. In: *Clinical Cancer Research* 25.10 (2019). ISSN: 15573265. DOI: 10.1158/1078-0432.CCR-18-2523.
- [150] L. J. Mah, A. El-Osta, and T. C. Karagiannis. *γ h2AX: A sensitive molecular marker of DNA damage and repair*. 2010. DOI: 10.1038/leu.2010.6.

- [151] Szilvia Brunner, Dániel Varga, Renáta Bozó, Róbert Polanek, Tünde Tőkés, Emília Rita Szabó, Réka Molnár, Nikolett Gémes, Gábor J. Szebeni, László G. Puskás, Miklós Erdélyi, and Katalin Hideghéty. “Analysis of Ionizing Radiation Induced DNA Damage by Superresolution dSTORM Microscopy”. In: *Pathology and Oncology Research* 27 (2021). ISSN: 15322807. DOI: 10.3389/pore.2021.1609971.
- [152] Dániel Varga, Hajnalka Majoros, Zsuzsanna Ujfaludi, Miklós Erdélyi, and Tibor Pankotai. “Quantification of DNA damage induced repair focus formation: Via super-resolution dSTORM localization microscopy”. In: *Nanoscale* 11.30 (2019). ISSN: 20403372. DOI: 10.1039/c9nr03696b.
- [153] Jennifer S. Dickey, Christophe E. Redon, Asako J. Nakamura, Brandon J. Baird, Olga A. Sedelnikova, and William M. Bonner. *H2AX: Functional roles and potential applications*. 2009. DOI: 10.1007/s00412-009-0234-4.
- [154] William M Bonner, Christophe E Redon, Jennifer S Dickey, Asako J Nakamura, Olga A Sedelnikova, Stéphanie Solier, and Yves Pommier. “GammaH2AX and cancer.” In: *Nat Rev Cancer* 8.12 (2008).
- [155] Oscar Fernandez-Capetillo, Arkady Celeste, and André Nussenzweig. *Focusing on foci: H2AX and the recruitment of DNA-damage response factors*. 2003. DOI: 10.4161/cc.2.5.509.
- [156] Emmy P. Rogakou, Chye Boon, Christophe Redon, and William M. Bonner. “Megabase chromatin domains involved in DNA double-strand breaks in vivo”. In: *Journal of Cell Biology* 146.5 (1999). ISSN: 00219525. DOI: 10.1083/jcb.146.5.905.
- [157] Pierre Golstein and Gillian M. Griffiths. “An early history of T cell-mediated cytotoxicity”. In: *Nature Reviews Immunology* 18.8 (2018). ISSN: 14741741. DOI: 10.1038/s41577-018-0009-3.
- [158] Bettina Weigelin, Annemieke Th den Boer, Esther Wagena, Kelly Broen, Harry Dolstra, Rob J. de Boer, Carl G. Figdor, Johannes Textor, and Peter Friedl. “Cytotoxic T cells are able to efficiently eliminate cancer cells by additive cytotoxicity”. In: *Nature Communications* 12.1 (2021). ISSN: 20411723. DOI: 10.1038/s41467-021-25282-3.
- [159] Jérôme Galon, Anne Costes, Fatima Sanchez-Cabo, Amos Kirilovsky, Bernhard Mlecnik, Christine Lagorce-Pagès, Marie Tosolini, Matthieu Camus, Anne Berger, Philippe Wind, Franck Zinzindohoué, Patrick Bruneval, Paul Henri Cugnenc, Zlatko Trajanoski, Wolf Herman Fridman, and Franck Pagès. “Type, density, and location of immune cells within human colorectal tumors predict clinical outcome”. In: *Science* 313.5795 (2006). ISSN: 00368075. DOI: 10.1126/science.1129139.
- [160] Wolf Herman Fridman, Franck Pagès, Catherine Sauts-Fridman, and Jérôme Galon. *The immune contexture in human tumours: Impact on clinical outcome*. 2012. DOI: 10.1038/nrc3245.

- [161] Tao Wan, Xu Liu, Jianhui Chen, and Zengchang Qin. “Wavelet-based statistical features for distinguishing mitotic and non-mitotic cells in breast cancer histopathology”. In: *2014 IEEE International Conference on Image Processing, ICIP 2014*. 2014. DOI: 10.1109/ICIP.2014.7025464.
- [162] Tao Wan, Wanshu Zhang, Min Zhu, Jianhui Chen, Alin Achim, and Zengchang Qin. “Automated mitosis detection in histopathology based on non-gaussian modeling of complex wavelet coefficients”. In: *Neurocomputing* 237 (2017). ISSN: 18728286. DOI: 10.1016/j.neucom.2017.01.008.
- [163] Zeynab Mousavikhamene, Daniel J. Sykora, Milan Mrksich, and Neda Bagheri. “Morphological features of single cells enable accurate automated classification of cancer from non-cancer cell lines”. In: *Scientific Reports* 11.1 (2021). ISSN: 20452322. DOI: 10.1038/s41598-021-03813-8.
- [164] Anabela Ferro, Tânia Mestre, Patrícia Carneiro, Ivan Sahumbaiev, Raquel Seruca, and João M. Sanches. “Blue intensity matters for cell cycle profiling in fluorescence DAPI-stained images”. In: *Laboratory Investigation* 97.5 (2017). ISSN: 15300307. DOI: 10.1038/labinvest.2017.13.
- [165] Kristofer E. delas Peñas, Ralf Haeusler, Sally Feng, Valentin Magidson, Mariia Dmitrieva, David Wink, Stephen Lockett, Robert Kinders, and Jens Rittscher. “Profiling DNA Damage in 3D Histology Samples”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 13578 LNCS. Springer Science and Business Media Deutschland GmbH, 2022, pp. 84–93. ISBN: 9783031169601. DOI: 10.1007/978-3-031-16961-8_{_}9.
- [166] Roghayyeh Baghban, Leila Roshangar, Rana Jahanban-Esfahlan, Khaled Seidi, Abbas Ebrahimi-Kalan, Mehdi Jaymand, Saeed Kolahian, Tahereh Javaheri, and Peyman Zare. *Tumor microenvironment complexity and therapeutic implications at a glance*. 2020. DOI: 10.1186/s12964-020-0530-4.
- [167] Hani A. Alturkistani, Faris M. Tashkandi, and Zuhair M. Mohammedsaleh. *Histological Stains: A Literature Review and Case Study*. 2015. DOI: 10.5539/gjhs.v8n3p72.
- [168] Jonathan T.C. Liu, Adam K. Glaser, Kaustav Bera, Lawrence D. True, Nicholas P. Reder, Kevin W. Eliceiri, and Anant Madabhushi. *Harnessing non-destructive 3D pathology*. 2021. DOI: 10.1038/s41551-020-00681-x.
- [169] Vincent M.J.I. Cuijpers, Hamdan S. Alghamdi, Natasja W.M. Van Dijk, Jakub Jaroszewicz, X. Frank Walboomers, and John A. Jansen. “Osteogenesis around CaP-coated titanium implants visualized using 3D histology and micro-computed tomography”. In: *Journal of Biomedical Materials Research - Part A* 103.11 (2015). ISSN: 15524965. DOI: 10.1002/jbm.a.35485.
- [170] Katrin Schweinzer, Lukas Kofler, Jürgen Bauer, Gisela Metzler, Helmut Breuninger, and Hans Martin Häfner. “Cytokeratin AE1/AE3 immunostaining and 3D-histology: improvement of diagnosis in desmoplastic squamous cell carcinoma of the skin”. In: *Archives of Dermatological Research* 309.1 (2017). ISSN: 1432069X. DOI: 10.1007/s00403-016-1700-5.

- [171] Predrag Jevtić, Lisa J. Edens, Lidija D. Vuković, and Daniel L. Levy. *Sizing and shaping the nucleus: Mechanisms and significance*. 2014. DOI: 10.1016/j.ceb.2014.01.003.
- [172] P. A. Humphrey. “Complete histologic serial sectioning of a prostate gland with adenocarcinoma”. In: *American Journal of Surgical Pathology* 17.5 (1993). ISSN: 01475185. DOI: 10.1097/00000478-199305000-00005.
- [173] David G. Bostwick and Hillel Kahane. “Adequate histologic sectioning of prostate needle biopsies”. In: *Annals of Diagnostic Pathology* 17.4 (2013). ISSN: 10929134. DOI: 10.1016/j.anndiagpath.2013.04.006.
- [174] Teresa Pusiol, Doriana Morichetti, Maria Grazia Zorzi, and Francesco Pisciolì. *Melanoma staging: implications of histologic sectioning procedures*. 2011. DOI: 10.1200/JCO.2011.37.5188.
- [175] Timo Dickscheid, Sarah Haas, Sebastian Bludau, Philipp Glock, Marcel Huysegoms, and Katrin Amunts. “Towards 3D reconstruction of neuronal cell distributions from histological human brain sections”. In: *Future Trends of HPC in a Disruptive Scenario*. 2019. DOI: 10.3233/APC190016.
- [176] Nick Y. Larsen, Xixia Li, Xueke Tan, Gang Ji, Jing Lin, Grazyna Rajkowska, Jesper Møller, Ninna Vihrs, Jon Sparring, Fei Sun, and Jens R. Nyengaard. “Cellular 3D-reconstruction and analysis in the human cerebral cortex using automatic serial sections”. In: *Communications Biology* 4.1 (2021). ISSN: 23993642. DOI: 10.1038/s42003-021-02548-6.
- [177] Xixia Li, Gang Ji, Xi Chen, Wei Ding, Lei Sun, Wei Xu, Hua Han, and Fei Sun. “Large scale three-dimensional reconstruction of an entire *Caenorhabditis elegans* larva using AutoCUTS-SEM”. In: *Journal of Structural Biology* 200.2 (2017). ISSN: 10958657. DOI: 10.1016/j.jsb.2017.09.010.
- [178] Ting Tian, Zhaoyang Yang, and Xiaoguang Li. *Tissue clearing technique: Recent progress and biomedical applications*. 2021. DOI: 10.1111/joa.13309.
- [179] Dian Jing, Shiwen Zhang, Wenjing Luo, Xiaofei Gao, Yi Men, Chi Ma, Xiaohua Liu, Yating Yi, Abhijit Bugde, Bo O. Zhou, Zhihe Zhao, Quan Yuan, Jian Q. Feng, Liang Gao, Woo Ping Ge, and Hu Zhao. “Tissue clearing of both hard and soft tissue organs with the pegasos method”. In: *Cell Research* 28.8 (2018). ISSN: 17487838. DOI: 10.1038/s41422-018-0049-z.
- [180] W. Spalteholz. *Über das Durchsichtigmachen von menschlichen und tierischen Präparaten*. 1914.
- [181] H. Steinke and W. Wolff. “A modified spalteholz technique with preservation of the histology”. In: *Annals of Anatomy* 183.1 (2001). ISSN: 09409602. DOI: 10.1016/S0940-9602(01)80020-0.
- [182] Lindsey A. Barner, Adam K. Glaser, Chenyi Mao, Etsuo A. Susaki, Joshua C. Vaughan, Suzanne M. Dintzis, and Jonathan T. C. Liu. “Multiresolution nondestructive 3D pathology of whole lymph nodes for breast cancer staging”. In: *Journal of Biomedical Optics* 27.03 (2022). ISSN: 15602281. DOI: 10.1117/1.jbo.27.3.036501.

- [183] Anthony M. Taylor and Bruno Bordoni. *Histology, Blood Vascular System*. 2020.
- [184] Evan P. Troendle, Ayesha Khan, Peter C. Searson, and Martin B. Ulmschneider. “Predicting drug delivery efficiency into tumor tissues through molecular simulation of transport in complex vascular networks”. In: *Journal of Controlled Release* 292 (2018). ISSN: 18734995. DOI: 10.1016/j.jconrel.2018.11.010.
- [185] Marvin Fritz, Prashant K. Jha, Tobias Köppl, J. Tinsley Oden, Andreas Wagner, and Barbara Wohlmuth. “Modeling and simulation of vascular tumors embedded in evolving capillary networks”. In: *Computer Methods in Applied Mechanics and Engineering* 384 (2021). ISSN: 00457825. DOI: 10.1016/j.cma.2021.113975.
- [186] Syed Ahmed Taqi, Syed Abdus Sami, Lateef Begum Sami, and Syed Ahmed Zaki. *A review of artifacts in histopathology*. 2018. DOI: 10.4103/jomfp.JOMFP{_}125{_}15.
- [187] Alan Aberdeen. *AIDA*. Oxford. URL: <https://github.com/alanaberdeem/AIDA>.
- [188] Sharath M. Bhat, Vaishnavi A. Badiger, Sampara Vasishta, Juhi Chakraborty, Seetharam Prasad, Sourabh Ghosh, and Manjunath B. Joshi. *3D tumor angiogenesis models: recent advances and challenges*. 2021. DOI: 10.1007/s00432-021-03814-0.
- [189] Stephen B. Fox. “Assessing tumor angiogenesis in histological samples.” In: *Methods in molecular biology (Clifton, N.J.)* 467 (2009). ISSN: 10643745. DOI: 10.1007/978-1-59745-241-0{_}3.
- [190] Stephen B. Fox and Adrian L. Harris. *Histological quantitation of tumour angiogenesis*. 2004. DOI: 10.1111/j.1600-0463.2004.apm11207-0803.x.
- [191] Mark Zaidi, Fred Fu, Dan Cojocari, Trevor D. McKee, and Bradly G. Wouters. “Quantitative Visualization of Hypoxia and Proliferation Gradients Within Histological Tissue Sections”. In: *Frontiers in Bioengineering and Biotechnology* 7 (2019). ISSN: 22964185. DOI: 10.3389/fbioe.2019.00397.