Cilia Segmentation Using U-Net with Gabor Filter

by

Chenxiao Li

(Under the Direction of Shannon Quinn)

Abstract

Cilia are hairlike organelles found on the surface of most cell types. Cilia segmentation is fundamental to many biological studies and primary ciliary dyskinesia (PCD) diagnosis. Previously, biologists and clinicians detected and classified cilia manually, which was time consuming and error-prone. Previous studies have been using traditional and deep learning-based image segmentation methods for segmenting cilia. In this study, we propose to use Gabor filters (GFs) to perform feature extraction and train the U-Net model for cilia segmentation. We show that the composite models with the combination of Gabor filtered features improve the performances of the U-Net base model. Our best composite model with Gabor 25 ($\theta = \pi/2$, σ=1, $\lambda = \pi*3/4$, γ = 0.05, ϕ=0) achieved an IoU of 0.37, almost 8% of improvement of the performance of the base model. By comparing the performances of the GF composite model with previous studies analyzing the cilia video set, we show that the presented framework outperformed previous models in terms of IoU score. Using GF as a data augmentation tool can help to enhance the robustness of features, and achieve a better performance.

Index words: [Image Segmentation, Cilia, Gabor Filter, U-Net]

Cilia Segmentation Using U-Net with Gabor Filter

by

Chenxiao Li

B.S., Fudan University, China, 2013

A Thesis Submitted to the Graduate Faculty of the
University of Georgia in Partial Fulfillment of the Requirements for the Degree.

Master of Science

Athens, Georgia

2023

Cilia Segmentation Using U-Net with Gabor Filter

by

Chenxiao Li

Major Professor:	Shannon Quinn
Committee:	Frederick Maier
	John A. Miller

Electronic Version Approved:

Ron Walcott
Dean of the Graduate School
The University of Georgia
May 2023

Dedication

This work is dedicated to my beloved aunt, Ying Liu.

Acknowledgments

I would like to express my gratitude to my advisor, Dr. Shannon Quinn, for his support and assistance throughout the course of my studies. I am incredibly grateful and honored to have worked with Dr. Quinn since 2020. I still remember three years ago when my previous PI left UGA, I came to Dr. Quinn and was excited at the potential to apply my prior life sciences experience towards an AI focused project. His knowledge, enthusiasm on research deeply impressed me. I have received lots of help, encouragement from him as well.

Special thanks are given to my committee members, Dr. Miller and Dr. Maier, for their time and service. I sincerely thank all members from Quinn Research Group.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Cilia are hair-like structures found on the surface of nearly all eukaryotic cells in vertebrates. Cilia play an important role in many biological processes such as fluid movement and cell locomotion. In the airway, motile cilia beat in a coordinated manner to propel the pathogens and debris out of the airway. Figure 1.1 shows an example airway cilia image. The coordinated beating of cilia plays a key role in maintaining effective mucociliary clearance (MCC), which is the primary innate defense mechanism of the lung (Bustamante-Marin & Ostrowski, 2017). Dysfunction of cilia causes severe diseases known as ciliopathies, including primary ciliary dyskinesia (PCD) and polycystic kidney disease (PKD), which lead to chronic lung infections and fluid-filled cysts respectively (McConnachie et al., 2021; Waters & Beales, 2011). Therefore, the effective analysis of ciliary function is critical in the diagnosis of cilia-related diseases. Currently, the assessment of ciliary function and characteristics remains challenging, and a fundamental part of overcoming this challenge is effective cilia segmentation.

1.1 Introduction to Cilia

Cilia can be broadly divided into two types: motile and non-motile (also known as primary cilia) (Mitchison & Valente, 2017). Primary cilia have sensory functions and play an important role in signal transduction (Goetz & Anderson, 2010). Motile cilia are mostly found on the apical surface of epithelial cells in the airways, brain, and oviduct (Brightman & Palay, 1963; Dirksen, 1971; Jeffery & Reid, 1975).The coordinated beating pattern of motile cilia sweeps the luminal contents over the epithelial surface. In the human respiratory tract, a ciliated epithelium has 200-300 motile cilia with a surface density around 5-8 cilia/μm (Salathe, 2007; Spassky & Meunier, 2017; Tilley et al., 2015). Motile cilia in the airways are the major functional components of MCC, which is the the primary innate defense mechanism of the lung (Bustamante-Marin & Ostrowski, 2017). In the airways, cilia coordinate with mucus to transport foreign particles out the the airway.

Dysfunction of cilia causes severe diseases known as ciliopathies, including PCD and PKD, which lead to chronic lung infections and fluid-filled cysts respectively (Hildebrandt et al., 2011; McConnachie et al., 2021; Waters & Beales, 2011). PCD is an autosomal recessive genetic disease characterized by ciliary

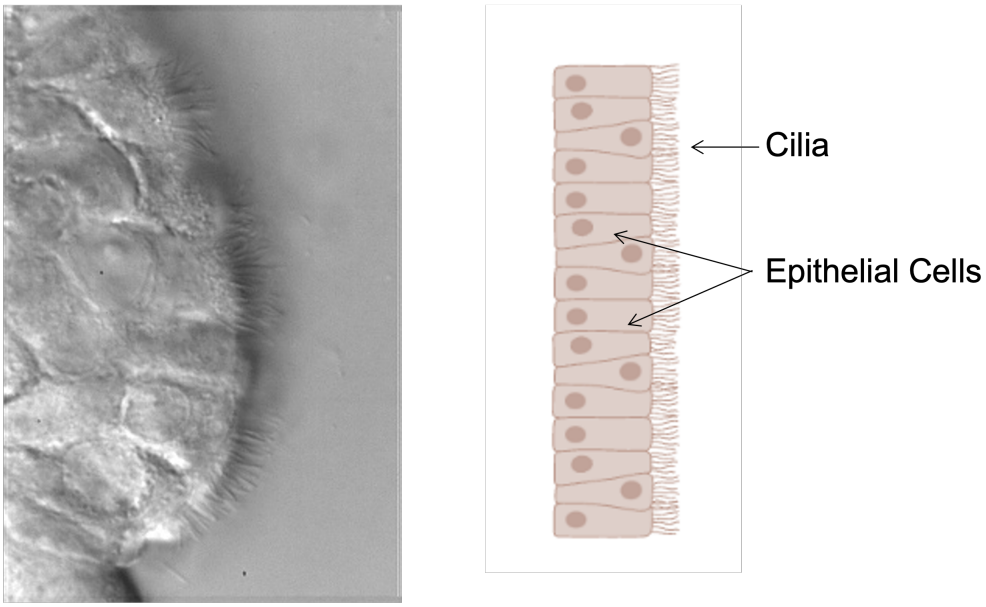

Figure 1.1: Airway cilia sample image. Left, a sample frame image of a cilia video. Right, diagram showing ciliated epithelial cells.

dysfunction and clinical features including chronic respiratory infection, heterotaxy and infertility (Horani et al., 2016; Knowles et al., 2013). Mutations in cilia-related genes result in abnormalities in ciliary structure and function in PCD patients. Defects in ciliary structure can be detected using transmission electricity microscope (TEM). For example, DNAH5 mutations are associated with outer dynein arm (ODA) defects (Omran et al., 2000). However, TEM analysis can only confirm ciliary structure using fixed samples. In fact, 30% PCD patients have normal or near-normal ciliary ultrastructure but abnormal ciliary function (Greenstone et al., 1983; Knowles et al., 2012). The analysis of cilia motion using high-speed video microscopy aids in the diagnosis of PCD. Reduced ciliary beat frequency (CBF) and abnormal ciliary beating such as immotile, stiff or circular can result in compromised MCC function, leading to chronic airway infection (Chilvers et al., 2001). CBF can be calculated using fast Fourier transform (FFT) or optical flow-based methods (Eshel et al., 1985; Kim et al., 2011; Smith et al., 2012). The assessment of ciliary function still largely relies on expert reviewers, which is time-consuming and error-prone. Cilia are tiny structures and only appear in a small percentage of an image, therefore the segmentation of cilia can help to reduce noisy background and mitigate human errors.

1.2 Image Segmentation

Image segmentation is a computer vision technique of splitting image data into multiple image segments, which reduces the data complexity and allows for further processing and analysis. Image segmentation has been used in many fields including biomedical imaging, robotics, and autonomous vehicles.

In biomedical studies, cell segmentation is fundamental for image-based cell biology research, which drives important fields such as diagnostics and drug-discovery. Cell segmentation, or identifying and splitting an image domain into cell instances, helps scientists to quickly analyze and evaluate biological features such as cell size, count, and shape. Traditionally, biologists or researchers had to manually label these targets, which is time-consuming and error-prone. In recent years, many tools have been developed to automatically segment targets of interest, which help researchers to acquire annotated data more efficiently. Those tools are built based on either traditional image processing methods, such as thresholding and clustering (Cheriet et al., 1998; Lauring et al., 2019), or deep learning-based methods such as Labkit (Arzt et al., 2022) and Cellpose (Stringer & Pachitariu, 2022; Stringer et al., 2021). However, those tools often provide models which are pre-trained on specific datasets, limiting their application to other biomedical tasks.

Cilia segmentation is used to detect and segment cilia from the images. Segmenting cilia can help to reduce background information and filter useful information for future quantitative cilia analysis or cilia motion phenotyping. Cilia are tiny compared to the other cellular compartments, therefore it is difficult to segment cilia from differential interference contrast (DIC) images. To detect cilia, many approaches have been developed (e.g. ACDC (Lauring et al., 2019), CiliaQ (Hansen et al., 2021)), however, those tools either rely on user's manual correction or fluorescent labels to identify cilia, and therefore, would not work on DIC images for cilia segmentation. Previous studies have used traditional image processing and deep learning-based methods for cilia segmentation of DIC images (Lu et al., 2018; Quinn et al., 2015; Zain et al., 2022; Zain et al., 2020).

1.2.1 Traditional image segmentation techniques

Traditional image segmentation techniques are low-level operations utilizing classic image processing methods and optimization algorithms. Those methods take a local or global view of the features and gain a segment map by comparing pixel values. Common image processing techniques include edge-based, threshold-based, region-based, cluster-based and watershed segmentation.

For cilia segmentation, the majority tools utilizing traditional techniques are designed to analyze fluorescent images, where cilia are stained with fluorescent labels. For example, automated cilia detection in cells (ACDC) utilized adaptive thresholding and Gaussian filter to detect fluorescence-labeled cilia (Lauring et al., 2019). CiliaQ is a software with a module utilizing various traditional techniques such as standard intensity thresholding and canny edge detection for cilia segmentation in fluorescent images (Hansen et al., 2021). For segmenting cilia in DIC images, Quinn et al. used intensity-based thresholding (Quinn et al., 2015) to segment cilia area from frame images. In this study, the authors calculated the standard deviation of the time-varying intensity changes at each pixel value and constructed a histogram.

The distribution peak is set as the pruning threshold value to segment cilia. This method achieves simple and fast segmenting map, however, the threshold value is specific to each experiment setting.

1.2.2 Gabor filter

Gabor filters (GFs), named after Dennis Gabor (Gabor, 1946), are bandpass filters applied to many image processing tasks, such as texture segmentation (Dunn et al., 1994; Weldon et al., 1996), edge detection (Mehrotra et al., 1992), and image representation (Porat & Zeevi, 1988). It is essentially a Gaussian modulated by a sinusoidal plane with frequency and orientation that can be used to separate edges and textures in an image. The advantage of using GFs is that the representations are similar to those of the human visual cortical cells. In addition, GF-based feature extraction has been extensively explored for texture analysis, face recognition, and fingerprint identification (Grigorescu et al., 2002; Idrissa & Acheroy, 2002; Kong et al., 2003; Lee, Wang, et al., 1999; Rahman & Bhuiyan, 2008). GFs are highly flexible in the function shape because of a set of parameters. By manipulating the parameters of GF, a set of filters can be generated with combinations of multiple wavelengths and orientations. Cilia are hair-like structure, which can be viewed as specialized texture with certain frequency and orientation, so we propose to apply Gabor filter for cilia feature extraction, which has not yet been explored.

1.2.3 Deep learning-based methods

With the rapid development of artificial intelligence (AI) and deep learning algorithms, many approaches have been established for the study of cilia (Bansal et al., 2021; Djenoune et al., 2023; Lu et al., 2018; Quinn et al., 2015; Zain et al., 2020). Bansal et al. used AI-based module to detect cilia and measure length, fluorescence intensity and co-localization (Bansal et al., 2021). However, this module only works on fluorescent images, where cilia are labeled in green or red channel. To segment cilia, methods such as densely-connected convolutional networks (DenseNets) (Lu et al., 2018), fully connected DenseNets (FC-DenseNets) (Zain et al., 2020) have been proposed to segment cilia on DIC images. In recent years, the U-Net (Ronneberger et al., 2015) has become a dominate deep learning architecture for biomedical image segmentation tasks. U-Net consists of two parts: the encoder, which is used to extract relevant features from images, and the decoder, which takes the extracted features and reconstructs a reconstruction a segmentation mask. After the convolutions in the contracting path, a rectified linear unit and a 2x2 max pooling are applied for downsampling. Although U-Net is able to perform feature extraction, the modification of a specialized architecture is a difficult task. Zain et al. used the U-Net backbone for segmenting cilia and achieve an Intersection over Union (IoU) of 0.399 (Zain et al., 2022). However, the performance of the U-Net architecture on the cilia dataset is still not ideal. The authors propose to use two low-level feature extraction, zero-phase PCA sphering (ZCA) and sparse autoencoder (SAE), to aid the U-Net base model. ZCA can normalize the data using the significant spacial features across the dataset and 'rotate' the data through the dataspace. SAE is a deep learning model with an additional constraint to reconstruct images as low-level representation. Both ZCA and SAE are used as data augmentation tools that can improve the performances of the U-Net segmentation model by 10%.

1.3 Cilia Data

For cilia segmentation, we used cilia data from our previous studies (Lu et al., 2018; Quinn et al., 2015) containing 325 videos of ciliary motion and their corresponding ground-truth masks. Briefly, the cilia videos were obtained using nasal epithelial tissues from patients with ciliopathies and healthy controls. DIC videos were recorded at 200 frames per second for 1.25 seconds using a Phantom v4.2 high speed camera. Manual masks were drawn using ITK-SNAP software with four annotation classes: side-view cilia, top-down cilia, cell body, and background. For cilia segmentation tasks, we simplify the masks into cilia and non-cilia. For image preprocessing, we cropped the first frame of each video and its manual masks into 128x128 patches and only selected the patches with more than 5% cilia area.

1.4 Contributions

Gabor filter (GF) is widely used for recognizing patterns with different orientations and frequencies for texture analysis, therefore we propose to use GF for highlighting and extracting cilia features. Since the performances of current deep learning-based approaches for cilia segmentation are not ideal, we propose to combine GF with U-Net to improve model performances. Although U-Net itself can perform feature extraction and localization, U-Net architecture capability is limited in localizing objects with non-standard shape (Shahedi et al., 2020). Additionally, modifying the architecture or layer is an intricate task.

To improve the performances of the U-Net model in cilia segmentation, we proposed to use GFs for feature extraction and combine them with the U-Net framework for generating cilia segmentation maps. We generated a set of GFs and used random forest to select GFs with high feature importance based on mean decrease in impurity. The Gabor filtered feature maps are concatenated to the original image as an additional channel to train the U-Net model. The addition of features extracted by our selected GFs improved the performances of the base model. Our best composite model with Gabor 25 ($\theta = \pi/2$, σ=1, $\lambda = \pi^*3/4$, γ = 0.05, ϕ=0) achieved an IoU of 0.37, almost 8% of improvement.

1.5 Structure of the thesis

The organization of the thesis is as follows. Recent and related work in image segmentation is reviewed in Chapter 2. In Chapter 3 we describe the cilia dataset and the methodology of the proposed study. In Chapter 4, we present the experimental results. Finally, in Chapter 5 conclusions are drawn and future work is proposed.

Chapter 2

Related Work

Image segmentation has many practical applications including biomedical imaging, autonomous vehicles, facial recognition, and traffic control systems. There are two types of image segmentation: semantic segmentation, which detects multiple objects within a category as one entity, and instance segmentation, which distinguishes between different instances of the same category. Semantic image segmentation is important for image analysis in removing background noises, improving precision and area detection. It is also a difficult task in computer vision and digital image processing. Traditional approaches such as thresholding and watershed compare pixel values in order to achieve the segment map. With the development of deep learning and neural networks, there are many tools and architectures developed for image segmentation.

Our group is interested in cilia segmentation, which is to detect and segment cilia from images. Cilia segmentation is the fundamental step for many biological studies and PCD diagnosis. Visualization of ciliary movement is critical in diagnosing cilia-related diseases. Previously, biologists and clinicians used to detect and classify cilia manually, which is time consuming and error-prone. In recent years, there have been attempts to segment cilia using traditional and deep-learning based methods (Hansen et al., 2021; Lauring et al., 2019; Lu et al., 2018; Quinn et al., 2015; Zain et al., 2022; Zain et al., 2020).

This chapter is organized as follows: 1) traditional image segmentation methods, 2) deep learning-based methods.

2.1 Traditional image segmentation methods

Traditional image segmentation techniques are based on threshold, edge, region, and cluster to divide an image into areas of similar nature. Conventional methods produce fast and simple segmentation masks. To detect features or characteristics in images, a large number of studies have reported numerous manually-designed, or handcrafted, features with variations in scale and illumination. Handcrafted algorithms such as Gabor (Gabor, 1946), Histogram of Oriented Gradients (HOG) (Dalal & Triggs, 2005), and Scale-invariant feature transform (SIFT) (Lowe, 2004) are commonly used to extract features from images. The drawbacks of image segmentation with traditional methods are that they are not flexible and exhibiting

low accuracy. Moreover, they often require manual labor to fine-tune parameters for specific segmenting tasks.

For cilia segmentation, Zhang et al. utilizes two thresholding-based methods to segment fluorescent signal-labelled cilia (Zhang et al., 2019). Specifically, they used the Transflour module and MultiWavelength Scoring (MWS) module within MetaXpress software. Both modules require the users to set parameters such as size, intensity and area to threshold cilia signal from the fluorescent images. Automated cilia detection in cells (ACDC) is a software for automated cilia detection in fluorescent images (Lauring et al., 2019). To detect cilia, the authors applied the Gaussian filter as an adaptive thresholding method to binarize the cilia image. Then the directional score parameter is applied as a threshold for optimizing cilia detection. This method achieved a high F1 score of 0.96 on fluorescently-labeled cilia images. CiliaQ is an open-source software for automatic quantification of ciliary function in 2D and 3D images (Hansen et al., 2021). The first step of CiliaQ workflow is to segment images into cilia and background using classic algorithms such as Renyi Entropy or Canny3D. However, this approach can only provide a coarse segmentation and require user's input for manual correction. Therefore, the above-mentioned thresholding methods have very limited application for the study of cilia. They can only be applied to fluorescent images where cilia are stained with a specialized antibody against Arl13b or acetylated alpha-tubulin. Those labeled cilia signals are acquired in a channel different from the rest of cells. For a live cell setting such as differential interference contrast (DIC), thresholding methods work poorly to segment cilia. In a study analyzing cilia video data, Quinn et al. (Quinn et al., 2015) designed a pruning method to segment cilia and discard the noisy background. The prune algorithm is an intensity-based thresholding method by calculating the standard deviation of the time-varying changes at each pixel and setting an adaptive threshold to filter out background pixels. This method achieves a cilia segmented map quickly with decent accuracy. However, the threshold of the prune method is adaptive and specific to each video set, therefore it is not ideal for batch work.

Gabor filters (GFs), named after Dennis Gabor, are bandpass filters used in image processing. It is essentially a Gaussian modulated by a sinusoidal plane with frequency and orientation (Gabor, 1946). After being convolved with GFs, patterns in images are highlighted at edges and points where the texture changes. GFs are extensively used for pattern analysis and feature extraction due to its reliability in localizing spatial-frequency domain properties. There are several advantages of GFs: first, the frequency and orientation representations of GFs are similar to those of the human visual system; second, GF can achieve fast and robust segmentation results by varying parameters according to the specific visual objects (Tadic et al., 2021); third, GFs specialize in detecting small, localized stretches or blobs.

Traditional methods are often combined with classifiers such as random forest and Support Vector Machine (SVM) for image segmentation to increase accuracy and flexibility. Nguyen-Thanh et al. proposed to use an orientation adaptive GF to extract features and implement with k-nearest neighbor (k-NN), SVM and neural network-based classifiers for mitochondria segmentation (Nguyen-Thanh et al., 2019). Their proposed method showed high accuracy in segmenting mitochondria. Since cilia can be viewed as specialized hair-like texture with certain frequency and orientation, we propose to use GF for cilia feature extraction.

2.2 Deep learning-based methods

With the rapid advancements in deep learning, great efforts have been made in segmenting biological images. For cell segmentation, several bio-image segmentation tools including StarDist (Schmidt et al., 2018), PlantSeg (Wolny et al., 2020), Cellpose (Stringer et al., 2021), and LABKIT (Arzt et al., 2022) have achieved great performances. However, most of the state-of-the art segmentation methods require a large dataset with human-labelled ground-truth data for training.

Fully Convolutional network (FCN) was first introduced by Long and colleagues (Long et al., 2015). FCN can train end-to-end and use skip connections to enhance image segmentation. U-Net (Ronneberger et al., 2015) contains encoder layers and decoder layers. The encoder block reduces the spatial dimensions of the image, while the decoder repairs the details in the spatial dimension of the image. The architecture of standard U-Net is shown in figure 2.1. However, the drawback of U-Net is that it fails to separate objects when they are crowded or overlapped. The usage of skip connections would need redundant information, causing training overhead.

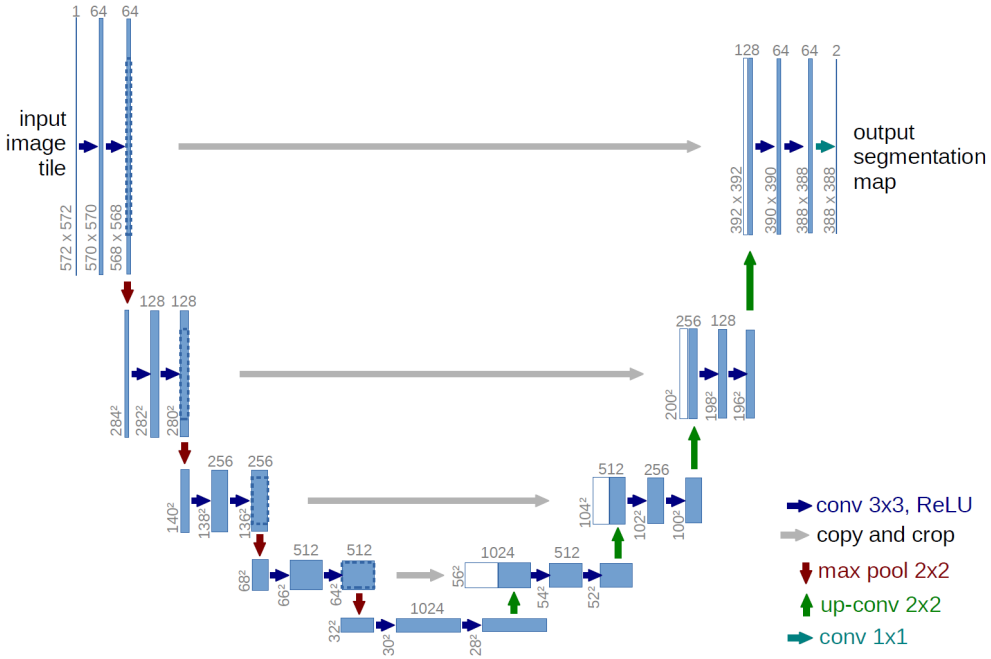

Figure 2.1: Schematic of a U-Net architecture (Ronneberger et al., 2015)

In the StarDist method, the cell nuclei is localized via star-convex polygons instead of bounding boxes, then the U-Net-based framework is trained to predict cell instances as polygons (Schmidt et al., 2018). The advantage of Stardist is that it can segment crowded cells with high accuracy. However, StarDist works for segmenting roundish-shape targets, thus it can not be used to segment rod- or hair-like structures. PlantSeg uses a 3D U-Net to predict cell boundaries of dense plant tissues on images from confocal and light sheet settings (Wolny et al., 2020). It achieves high accuracy in segmenting plant and non-plant samples. Cellpose is a deep learning-based pipeline utilizing a modified U-Net as a backbone for cell segmentation

(Stringer et al., 2021). The pre-trained models provided by the Cellpose platform were trained on a broad set of images with cells of different shapes and types and thus can not provide a generalized high quality cell segmentation. Cellpose 2.0 (Stringer and Pachitariu, 2022) solves this problem by providing a zoo of pretrained models and a human-in-the-loop pipeline. This approach significantly reduced manual segmentation work to 100-200 regions of interest (ROIs) while achieved near-maximal performance.

The segmentation of biological images is highly diverse due to different imaging techniques, cell shapes and structures. For cilia segmentation, Ram et al. introduced the convolutional neural random forest classifier to classify the primary cilia in fluorescence images. Input images are fed in a CNN model with the U-Net architecture, and then the learned features are passed to the random decision forests classifier. Their approach achieved the F_β - score of 0.9102, outperforming the k-means and multi-layer perceptron (MLP) classifiers (Ram et al., 2018).

Our group is interested in segmenting motile cilia, which are highly overlapping hair-like structures with low convexity and hence extremely difficult to segment. Besides, cilia motility videos are usually recorded as DIC images, which often achieves poor segmenting results due to low contrast. Quinn et al. developed a computational pipeline to classify cilia motion (Quinn et al., 2015). In this framework, ciliary motion is considered as dynamic texture, which can be estimated by using autoregressive (AR) model with the optical flow features. Lu et al. developed an end-to-end pipeline utilizing a densely-connected convolutional network (DenseNet) with 74 layers to automatically recognize cilia area (Lu et al., 2018). This proposed framework achieved an accuracy of 86.2% in cilia segmentation. Zain et al. modified Lu's model and built a fully Connected DenseNet with 103 layers. This model achieved a IoU score of 33.06% and a accuracy of 88.3% (Zain et al., 2020). In a most recent work, Zain and Miller combined zero-phase PCA sphering (ZCA) and Sparse Autoencoders (SEA), which served as low level feature extraction, with U-Net. This new model achieved a 10% improvement with an IoU of 0.441 and an accuracy of 0.77 (Zain et al., 2022). Table 2.1 summaries the above mentioned cilia segmentation techniques.

To improve the performances of deep learning models, several studies incorporate low-level feature extraction to aid deep learning frameworks. Luan et al. proposed to incorporate Gabor filters into the deep convolutional neural networks (DCNNs) (Luan et al., 2018). By manipulating GF parameters, the proposed Gabor Convolutional Networks (GCN) framework effectively reduced the training complexity and improved performances over several benchmarks. Reyes and colleagues proposed an approach combining GF with U-Net for image segmentation (Reyes et al., 2022). They applied the framework on two benchmark datasets, the ISLES 2018 dataset and the 2018 Atrial Segmentation Challenge dataset, and outperformed U-Net and other state-of-the-art architectures in mIoU and Dice scores. However, this work is relatively computational expensive and not efficient because the input contains the original dataset plus additional features extracted by a whole set of GFs. Some of the filters extract meaningless feature information compared to the other filters in the bank. An optimized GF bank with selected parameters is critical in the performance of the GF-embedded U-Net architecture for specific segmentation tasks.

Segmentation Technique	Description	Results	Advantages	Disadvantages	Reference
Intensity-based thresholding	Separates images into cilia and non-cilia regions based on threshold values	-	Simple and fast calculations.	Threshold is adaptive and specific to each video	Quinn et al., 2015
FC-Densenet 109	Fully Convolutional Dense Net with 109 layers	Accuracy: 0.862	Efficient in the parameter usage and feature reuse.	Data is replicated multiple times. High computational cost.	Lu et al., 2018
FC-DenseNet 103	Fully Convolutional Dense Net with 103 layers	IoU: 0.3306, Accuracy: 0.883, Precision: 0.5326	Efficient in the parameter usage and feature reuse.	Data is replicated multiple times. High computational cost.	Zain et al., 2020
U-Net	U-shaped architecture consists of a contracting path and an expansive path	IoU: 0.339, Accuracy: 0.759, Precision: 0.692, Recall: 0.501, F1: 0.529	Flexible and general approach. State-of-the-art image segmentation algorithm	High training time.	Zain et al., 2022
ZCA/SAE + U-Net	Adding ZCA and SAE for feature extraction to aid U-Net	IoU: 0.441, Accuracy: 0.767, Precision:0.661, Recall: 0.580, F1: 0.585	Adding low feature extraction improves U-Net model's performance	Difficult to find optimal parameters of ZCA and SAE	Zain et al., 2022

Chapter 3

Methodology

In this work, I propose to use Gabor filters (GFs) to extract cilia features and apply the features to U-Net for cilia segmentation. For chapter 3, I will introduce the details of the dataset, Gabor filters, U-Net and the training strategies.

3.1 Cilia Dataset

The cilia dataset is sourced from our previous study (Quinn et al., 2015). Briefly, the cilia dataset contains 325 videos of nasal cilia motion acquired from patients as well as healthy controls. Cilia videos were recorded at 200 frames per second (fps) in grayscale. The corresponding ground-truth masks are manually generated using ITK-SNAP from Lu's study (Lu et al., 2018). Figure 3.1 shows an example of original image patch and its manual mask patch. For data preprocessing, the frame images and their corresponding masks are cropped into 128x128 patches. Since cilia are tiny structures in the frame images, we exclude cropped patches with less than 5% cilia area. Overall, our cilia dataset contains 1267 patches and their corresponding masks. Images are normalized to [0, 1] and masks are also rescaled to 0 to 1.

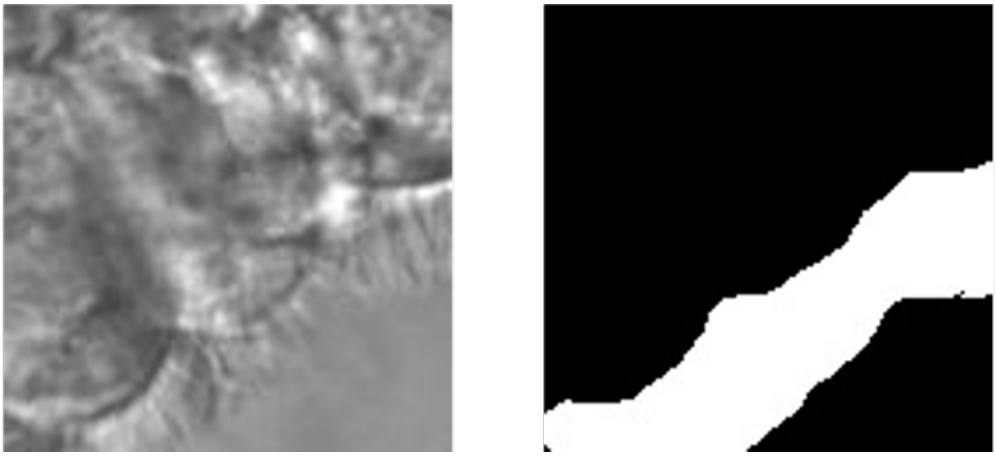

Figure 3.1: Example of a cilia image patch (left) and its ground-truth mask patch (right)

3.2 Gabor Filter

Gabor filter (GF), named after Dennis Gabor, is a linear filter commonly used for feature extraction, texture analysis, edge detection etc. The Gabor wavelets are defined as follows:

$$g(x, y; \lambda, \theta, \psi, \sigma, \gamma) = exp([-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2})exp[i(2\pi\frac{x'}{\lambda} + \psi)]$$

$$x' = x\cos\theta + y\sin\theta$$

$$y' = -x\sin\theta + y\cos\theta$$

where λ represents the wavelength of the sinusoidal component, θ represents the orientation from normal to the parallel stripes of the Gabor function, ϕ represents the phase offset, σ is the standard deviation of the Gaussian envelope and γ is the spatial aspect ratio specifying the ellipticity of the support of the Gabor function.

A response matrix can be obtained by convolving the original image I(x,y) with the GF as follows:

$$R(x, y; \lambda, \theta, \psi, \sigma, \gamma) = \sum_{x'}\sum_{y'} I(x - x', y - y')g(x', y'; \lambda, \theta, \psi, \sigma, \gamma)$$

By manipulating the parameters λ, θ, ψ, σ and γ, a GF bank can be generated with a set of different GFs.

In this study, a set of 54 GFs with various parameters is generated for the cilia image so that the outputs highlight cilia patterns and edges. The generated Gabor bank of 54 filters is shown in Figure 3.2. This Gabor bank is achieved with the following parameters: θ: {π/4, π/2, π*3/4}; σ: {1,3}; λ: {π/4, π/2, π*3/4}; γ: {0.05, 0.25, 0.5}; ϕ: {0}. The kernel size is set as 9.

By applying the GF bank to a cilia image, we achieve a set of convoluted outputs. Figure 3.3 shows an example of the GF bank applied to a cilia sample image. The output figures showed a variety of highlighted patterns and edges based on the orientation of GFs. Some of the filtered images present highlighted cilia textures. There are also meaningless filtered outputs with all 0-value pixels.

3.3 Random Forest

Random forest (RF) is a popular machine learning tool for regression and classification introduced by Amit and German (Amit & Geman, 1997). RF consists of a multitude of decision trees that operate as an ensemble, thus handling a variety of visual features. RF has high computational efficiency in training and classification and avoids overfitting. During the training process, each tree randomly picks the features, which provides the opportunities to find the feature that influences the majority of trees. Therefore, we use RF to find the features and their corresponding GFs with maximum importance. Due to high computational cost, we did not apply all filtered images as input for the U-Net model. Instead, we use RF

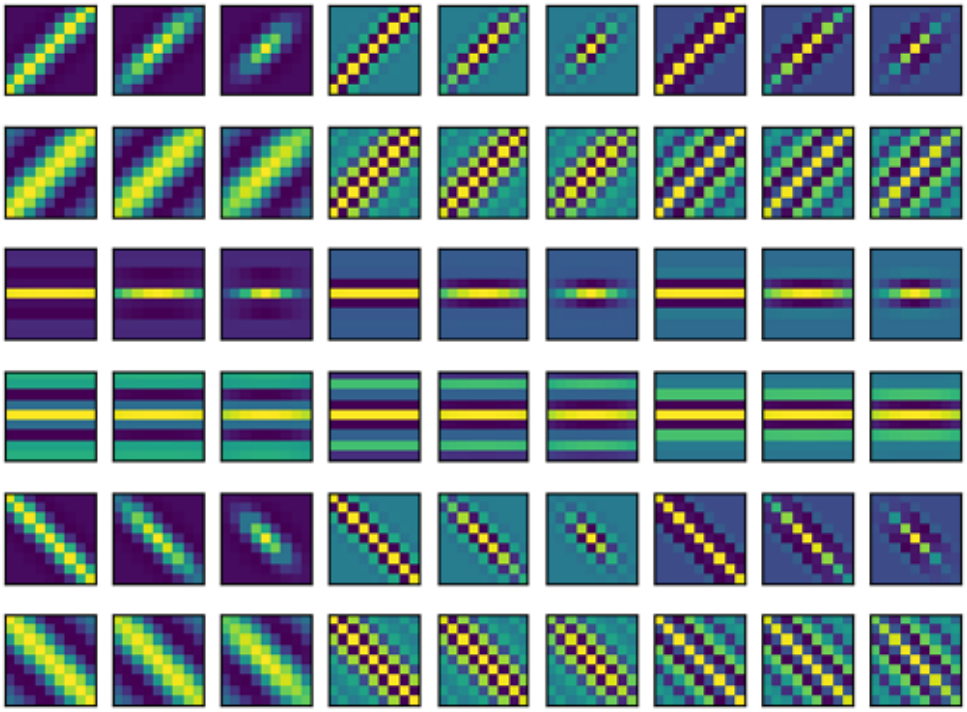

Figure 3.2: Gabor filter bank

for classifying the generated feature vectors and selecting the GFs that give us the distinguishing features of cilia at the spatial location.

Figure 3.4 shows the flow diagram of using RF to select GFs with high importance. Specifically, we extract features through the bank of 54 GFs with various parameters and then we train RF with the original image vector and 54 feature vectors. The most relevant features and their GFs that contribute most to the segmentation accuracy are selected for training the U-Net model.

3.4 U-Net for cilia segmentation

The standard U-Net is used as the backbone of the model. To begin with, we only train the U-Net model with the original image as a single channel input. For the composite models with feature extraction, we utilize the selected GFs with modified parameters. The Gabor filtered image is stacked to the original image as the second channel to train the U-Net model. The original U-Net paper (Ronneberger et al., 2015) used stochastic gradient descent optimizer, but here we use Adam (Kingma & Ba, 2014) optimizer and binary cross entropy as a loss function. The Adam algorithm is an extension to stochastic gradient descent for training deep learning models. Adam tends to converge faster and achieves better results in image

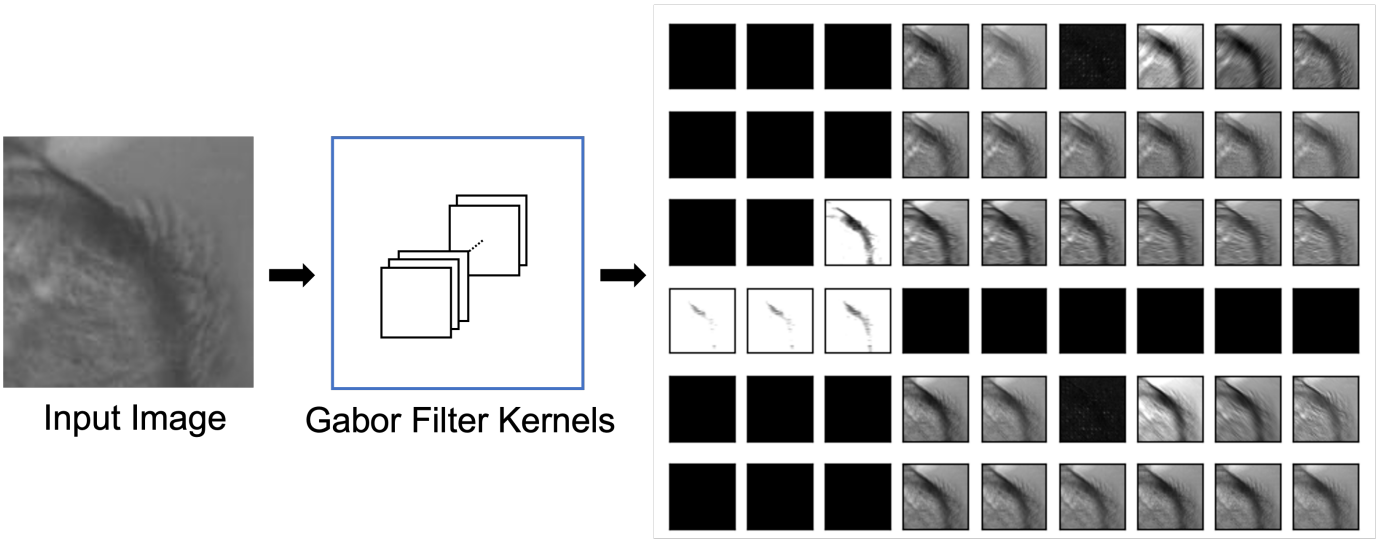

Figure 3.3: Illustration of Gabor filters applied to a cilia sample image

segmentation tasks (Anand et al., 2022; UYSAL et al., 2022). The dataset is split into 8:2 for training and testing. During the training process, the training set is randomly split into 8:2 for training and validation, respectively. We also use k-fold cross validation with k as 5.

3.5 Evaluation metrics

For image segmentation task, we use metrics including Intersection over union (IoU) or Jaccard similarity, accuracy, precision, recall and F1 or Dice similarity coefficient to evaluate our model performances of cilia segmentation.

IoU, also referred to as Jaccard similarity, can evaluate how similar a predicted mask is to the ground-truth mask. It measures the number of pixels common between the target and prediction masks divided by the total number of pixels present across both masks. IoU score ranges between 0 to 1, where 0 means no overlap and 1 indicates perfect overlap. IoU can be calculated as following:

$$IoU = \frac{\text{Area of Overlap}}{\text{Area of Union}}$$

Pixel accuracy is simple to report and commonly used for semantic segmentation. It is a metric to compare each pixel with the ground-truth mask. There are two drawbacks of accuracy for semantic segmentation. First, it can provide misleading results when there is a class imbalance. Second, pixel accuracy is informative as a standalone metric.

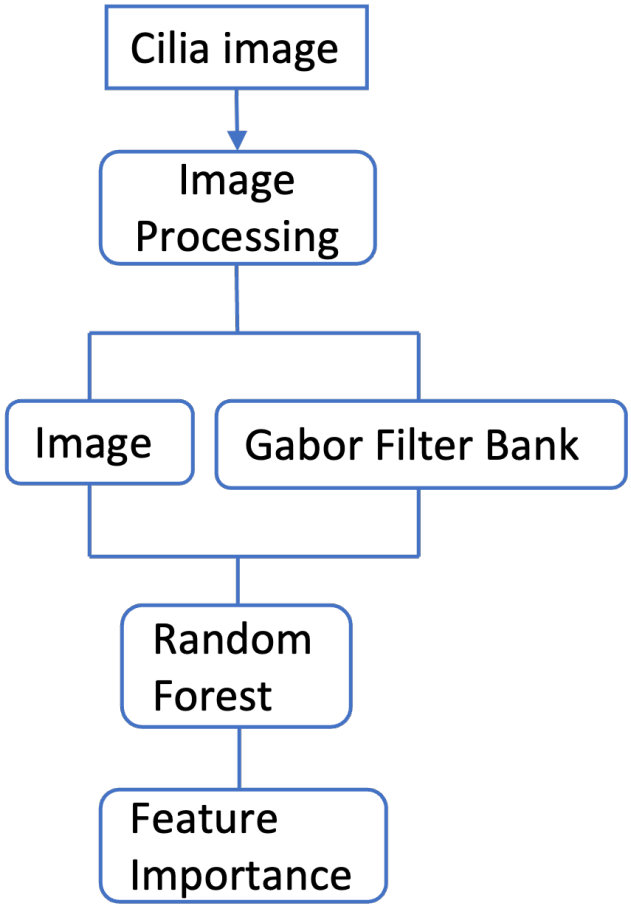

Figure 3.4: Random Forest flow diagram

We also use confusion matrix to evaluate the accuracy of the model's performance. A confusion matrix is generated using scikit-learn and presents four numbers: true negative (TN), false negative (FN), false positive (FP), true positive (TP) (Figure 3.5). From the confusion matrix, we can calculate metrics including precision, recall and F1. Precision is a measure of the quality of the predictions, indicating how precise the model is. Recall is a ratio of predictions of the number of true positives by ground truth to the total number of positive samples. Higher recall score indicates more positive samples the model has detected. F1 score is a harmonious mean of precision and recall. The accuracy, precision, recall, and F1 (or Dice) scores are computed as following:

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN}$$

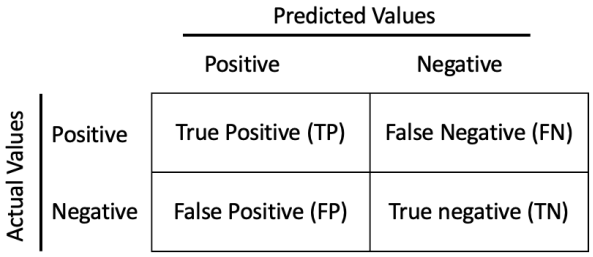

Figure 3.5: Confusion Matrix

$$Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}$$

$$F1 = \frac{2 * Precision * Recall}{Precision + Recall}$$

3.6 Code availability

Python code for Gabor filter generation and U-Net architecture is available at https://github.com/lizchenx/Gabor_UNet_Cilia_Segmentation.

Chapter 4

Experiments

4.1 Cilia Dataset

In this study, we use the cilia dataset containing 325 grayscale frame images and their corresponding masks with various sizes. All the images and ground-truth masks are from Lu et al., 2018; Quinn et al., 2015. For image processing, the images and their masks are cropped into small patches of 128x128. After cropping, we remove the patches containing less than 5% cilia area. Finally, the dataset has 1267 patches and corresponding mask patches. In addition, the dataset is split in a ratio of 8:2 for training and testing set correspondingly. During the training, 20% of the data is split for validation purpose. We also perform k-fold cross validation where k is set to 5. Table 4.1 shows the number of images and masks for training, validation and test.

Table 4.1: The number of images and masks for training, validation and testing

Type	Training (# of images)	Validation (# of images)	Test (# of images)	Total (# of images)
Images	810	203	254	1267
Masks	810	203	254	1267

4.2 Selected GFs

Previously, we generated a set of 54 GFs and used Random Forest (RF) to select a list of meaningful GFs. Figure 4.1 shows samples of applying RF for selecting GFs based on feature importance. Several of the GF feature vectors rank higher than the original image in feature importance for cilia segmentation. We select the GFs with high feature importance, reducing the set of GFs from 54 to 8. The list of GFs and their parameters is provided in Table 4.2. GFs with selected combination of parameters are applied to the cilia images for feature extraction. Figure 4.2 shows sample images filtered with the selected GFs. The extracted features are then concatenated with the original image as an additional channel for the input.

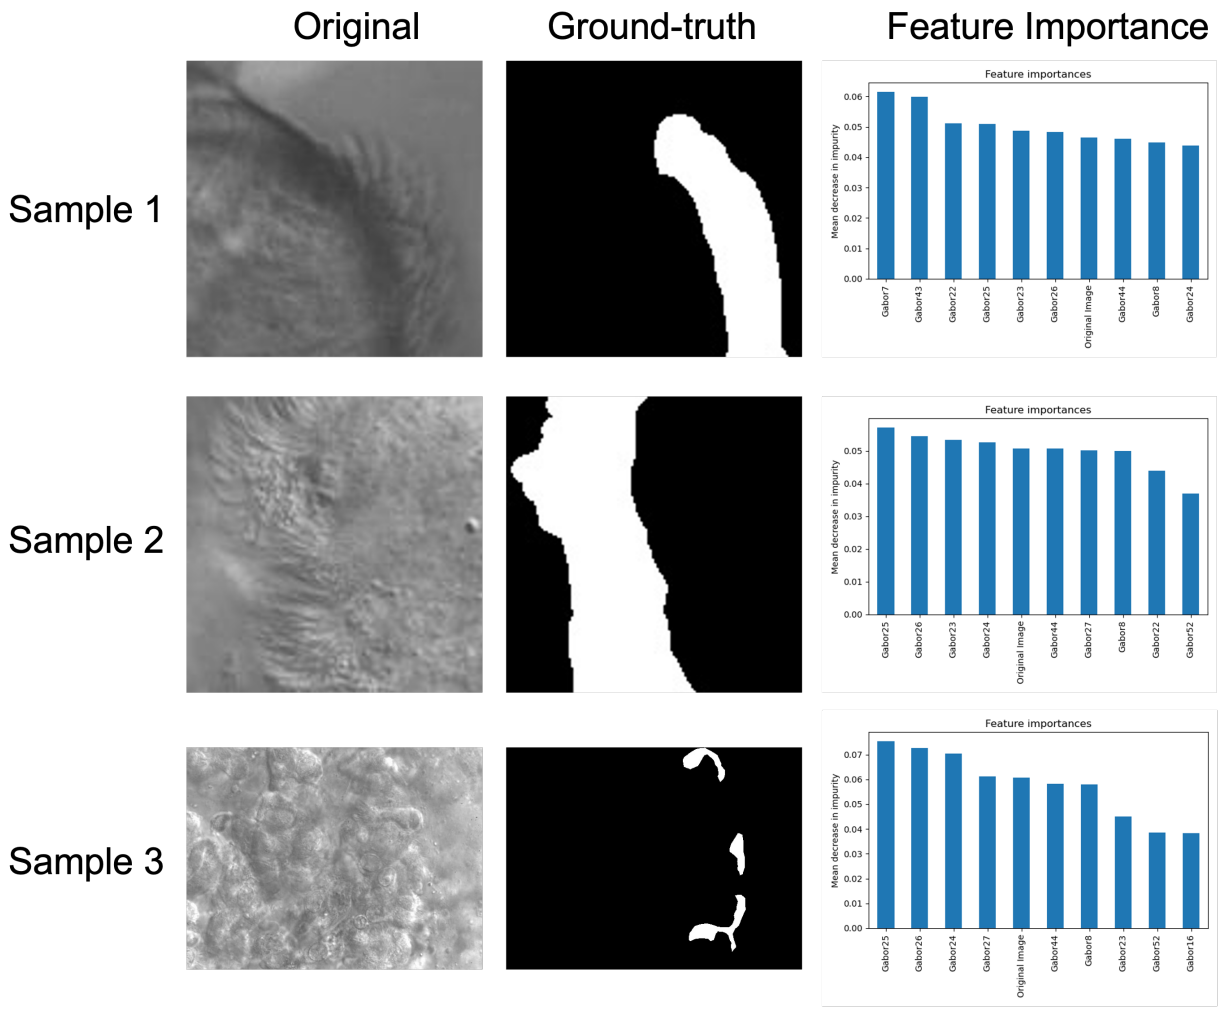

Figure 4.1: Feature importance of GFs on sample images

4.3 Experimental Setup

For the U-Net and composite models, we run the experiments for a total of 250 epochs with a batch size of 16. We use Adam optimizer and binary cross-entropy (BCE) loss function. For learning rate, we use learning rate monitor from PyTorch Lightning, which automatically monitor and logs learning rate for learning rate schedulers with a log rate of 10 during training. We train our models on 2x NVIDIA TITAN X GPU cards and NVIDIA Quadro RTX 5000.

For the FC-DenseNet model, we train the model for 100 epochs with a batch size of 12, learning rate of 0.001, learning rate decay of 0.995, and weight decay of 0.0001. The model was trained with the Adam

Table 4.2: Selected Gabor filters

Gabor filter names	θ	σ	λ	γ	ϕ
Gabor 7	$\pi/4$	1	π*3/4	0.05	0
Gabor 22	$\pi/2$	1	$\pi/2$	0.05	0
Gabor 23	$\pi/2$	1	$\pi/2$	0.25	0
Gabor 24	$\pi/2$	1	$\pi/2$	0.5	0
Gabor 25	$\pi/2$	1	π*3/4	0.05	0
Gabor 26	$\pi/2$	1	π*3/4	0.25	0
Gabor 27	$\pi/2$	1	π*3/4	0.5	0
Gabor 43	π*3/4	1	π*3/4	0.05	0

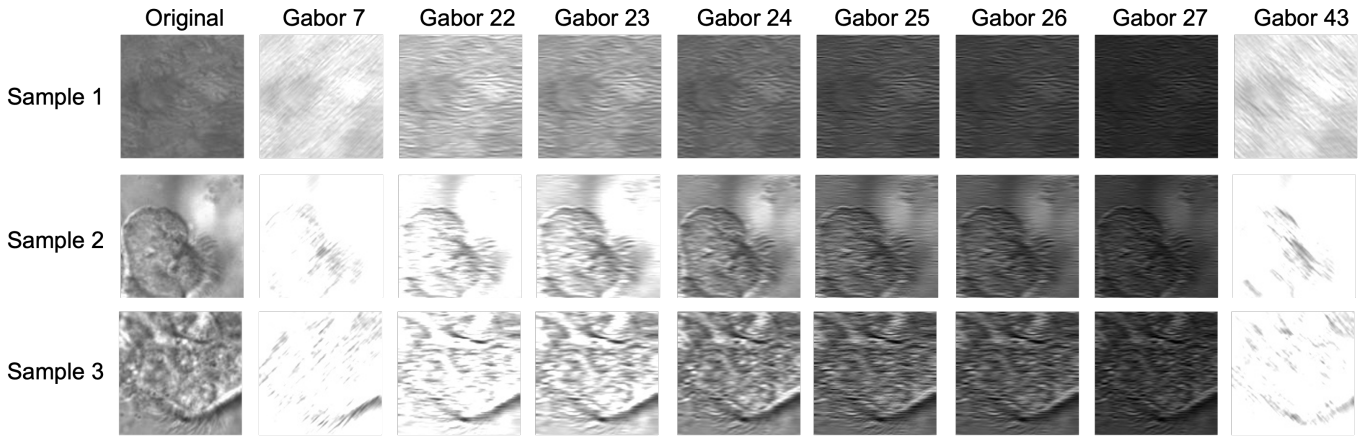

Figure 4.2: Gabor filter comparison examples.

optimizer and cross-entropy loss. To implement the ZCA/SAE composite model, we set the parameter ϵ of ZCA as 1e-4, α of SAE as 1e-3. During the training process, we use BCE loss function with a learning rate of 1e-3 and batch size of 12 for 225 epochs.

4.4 Results

The experiments are conducted on the training set from the cilia dataset. We first trained U-Net with original image patches, then we trained U-Net model with the original image concatenated with an additional Gabor-filtered image. We evaluated the models with five metrics: Intersection over Union (IoU) or Jaccard score, accuracy, F1 (or Dice) score, precision and recall. Table 4.3 provides a summery of experiments performed with selected GFs. The U-Net base model trained with the original image as a single channel achieved an IoU of 0.3461, accuracy of 0.8527, F1 of 0.4590, precision of 0.5976 and recall of 0.4394. Several of our GF models (e.g, Gabor 7, 22, 24, 25, 43) improved the performances with a higher IoU score

than that of the U-Net base model. The composite model with Gabor 25 achieves the highest scores of IoU (0.3726), F1 (0.4994) and recall (0.5161). The composite model with Gabor 26 achieves the highest precision of 0.6201. We also calculated the average scores of our composite models for each metric, the average performance of composite models outperformed the base model in terms of IoU, F1, and recall.

Table 4.3: Performances of Gabor U-Net models over the cilia dataset

Model	Evaluation Metrics				
	IoU	Accuracy	F1	Precision	Recall
U-Net	0.3461	**0.8527**	0.4590	0.5976	0.4394
Gabor7+U-Net	0.3528	0.8398	0.4709	0.5522	0.4795
Gabor22+U-Net	0.3553	0.8408	0.4795	0.5662	0.4852
Gabor23+U-Net	0.3423	0.8402	0.4605	0.5561	0.4635
Gabor24+U-Net	0.3519	0.8473	0.4751	0.5997	0.4648
Gabor25+U-Net	**0.3726**	0.8434	**0.4994**	0.5801	**0.5161**
Gabor26+U-Net	0.3286	0.8486	0.4468	**0.6201**	0.4173
Gabor27+U-Net	0.3433	0.8451	0.4620	0.5825	0.4531
Gabor43+U-Net	0.3532	0.8443	4804	0.5804	0.4798

Figure 4.3 shows a qualitative comparison of cilia segmentation results. The U-Net base model fails to predict the cilia area that is indicated in the ground-truth mask (Sample 1, 2 and 3). For Sample 4, the U-Net base model predicts non-cilia area as cilia. The composite models with our selected GFs tend to be more accurate in the location of the the segmentation mask. However, in some samples, a few composite models either include false positives or fail to predict the cilia areas.

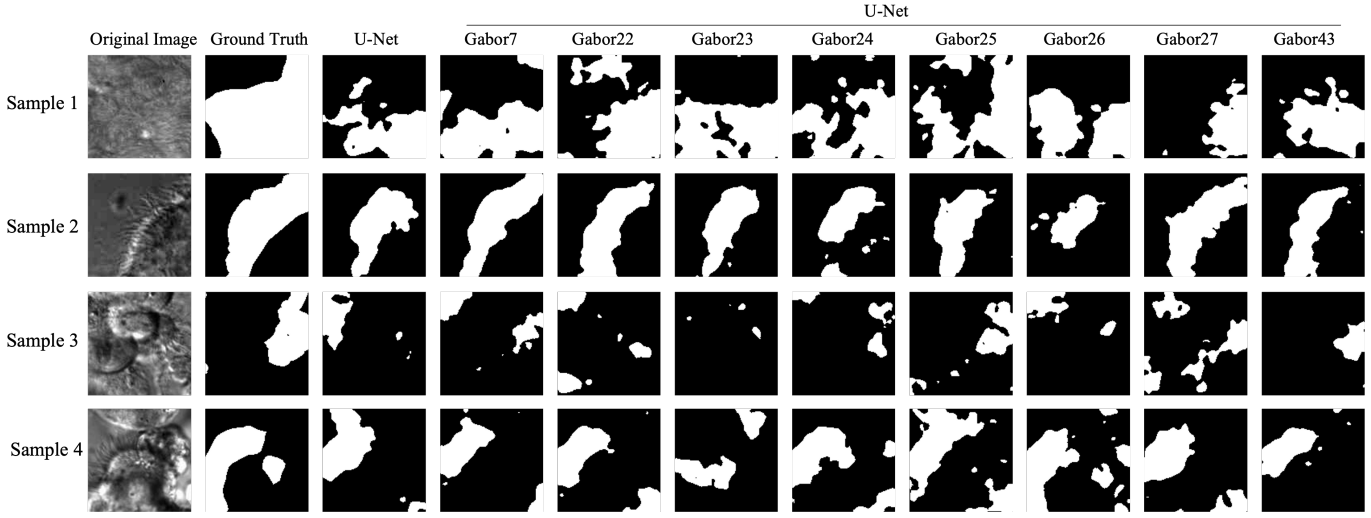

Figure 4.3: The visual comparison of cilia segmentation outputs

Previous models for segmenting cilia from DIC images are trained on different cilia sets and the data were pre-processed with different strategies. To compare our model performances with previous studies,

we also implemented previous studies utilizing threshold method (Quinn et al., 2015), FC-DenseNet (Lu et al., 2018; Zain et al., 2020), U-Net (Lu et al., 2018; Zain et al., 2022), and ZCA/SAE with U-Net (Zain et al., 2022) using the same training and testing sets. Table 4.4 summaries the comparison of segmenting the same cilia dataset with different algorithms. We compared the performances for segmenting cilia using non-deep learning method (intensity-based thresholding) and deep-learning methods, and we found that deep-learning methods outperformed the thresholding method. Gabor 25 composite model outperformed all the other deep-learning methods. Figure 4.4 shows a visual comparison of segmentation results achieved from above mentioned methods. For the intensity-based thresholding method, we set the threshold as 1.5 and applied the method to the testing set. ZCA (1e-4)/SAE (1e-3) composite model did not outperform U-Net base model because the parameters for ZCA and SAE from the original study might not be the optimal parameters for the current dataset. Overall, the Gabor 25 composite model outperforms the aforementioned methods by achieving an IoU of 0.3726, almost 8% improvement of the U-Net base model.

Table 4.4: Performances comparison over the cilia dataset

Model	Evaluation Metrics				
	IoU	Accuracy	F1	Precision	Recall
Intensity-based thresholding	0.3170	0.8239	0.4508	0.5757	0.4319
FC-DenseNet 103	0.3180	0.8034	0.4472	0.4732	0.5035
U-Net	0.3461	**0.8527**	0.4590	**0.5976**	0.4394
ZCA/SAE+U-Net	0.3185	0.8341	0.4388	0.5616	0.4382
Gabor25+U-Net	**0.3726**	0.8434	**0.4994**	0.5801	**0.5161**

4.5 Discussion

In this study, we proposed to use GF to extract features and then train the U-Net model for cilia segmentation. Overall, five out of our eight composite models achieved better performances in terms of the IoU score than the U-Net baseline model. Some of the composite models with GF extracted features predicted more accurate segmentation mask than the U-Net base model. Cilia segmentation is important in the studies of ciliogenesis and ciliopathies. Successful segmentation can detect cilia from background or cells. After segmentation, quantitative cilia analysis such as cilia length, density and cilia beat frequency can be estimated.

Ciliopathies are a group of disorders that can cause cilia dysfunction or malformation. Since cilia are important organelles on almost all vertebrate cells, cilia abnormalities are associated with many diseases including renal disease, diabetes, obesity and retinal degeneration. Primary ciliary dyskinesia (PCD) is an autosomoal recessive genetic disease characterized by ciliary dysfunction, which can lead to chronic lung and ear infection, infertility. So far, over 40 genes have been recognized to be linked to PCD (Zariwala et al., 2019). Current diagnosis tools of PCD includes nasal nitric oxide (nNO), genetic test, transmission electron microscopy (TEM), and high-speed video microscopy (HSVM) analysis. TEM can help

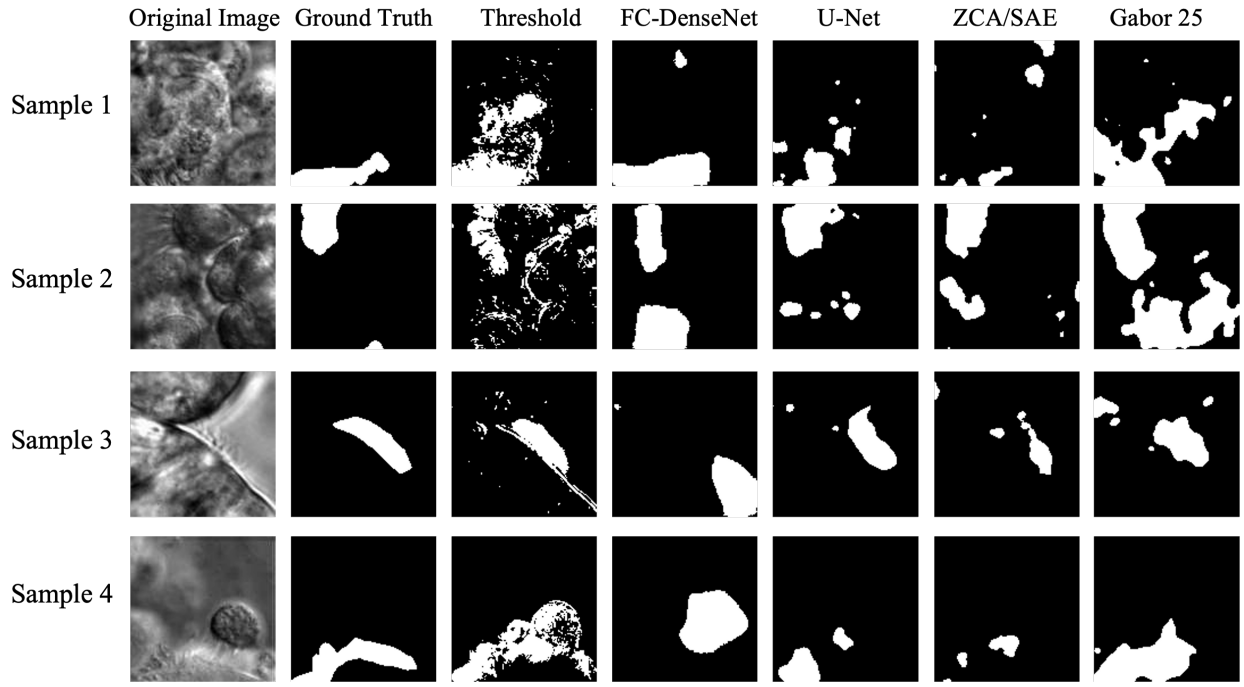

Figure 4.4: The visual comparison of cilia segmentation outputs using different methods

with detecting ultrastructural changes, however, 30% of PCD patients have normal ciliary ultrastructure (Knowles et al., 2012; Shoemark et al., 2012). As a diagnostic tool for PCD, HSVM is very efficient in visualizing cilia motility directly and detecting ciliary beat patterns with high sensitivity and specificity. However, the automation of cilia motion analysis is still under development. Many researchers still analyze cilia motion manually, which is time consuming and subject to selection bias. Software solutions such as CiliaFA (Smith et al., 2012), CiliarMove (Sampaio et al., 2021), CiliaQ (Hansen et al., 2021) were developed for quantitatively analyzing cilia. However, classifying cilia beat pattern remains a challenging task. Correctly segmenting cilia is a fundamental step of cilia quantification.

During the training process, overfitting can happen if the model learns the training dataset too well and performs well on the training set, but it does not perform well on the data outside of the training set. We trained for 250 epochs for all our models. To avoid overfitting, we used tensorboard to monitor the learning rate, training and validation loss.

4.5.1 Cilia dataset

One limitation of this study is that we only have a limited number of frame images and their ground-truth masks from our previous studies (Lu et al., 2018; Quinn et al., 2015). The original cilia videos were recorded

using samples from human nasal brush in DIC setting. Normally, each ciliated airway epithelium has 200-300 motile cilia, which can be presented in the videos in a lateral view or top-down view. This increases difficulties for both manual and generated masks. The current ground-truth masks were manually labelled a few years ago by a non-biologist (Lu et al., 2018). Most of the labels are accurate, however, there were a few masks mislabeled background into cilia or vice versa. Therefore, the current ground-truth masks limit the model performances to some extent. Building a validation dataset of cilia images with more accurate ground-truth masks can help the construction of a robust deep learning pipeline for cilia segmentation tasks.

4.5.2 Temporal information

Another limitation of the study is the exclusion of temporal information. We only include the first frame of each video in the dataset due to the high similarity of successive frames. We intend to compare the augmentation with selected GFs to the base U-Net model. Previously, Quinn et al. (Quinn et al., 2015) developed a prune method by setting a threshold to filter the background noise by calculating the standard deviation of pixel intensities across the time. This method gives quick and decent segmenting results with an IoU score of 0.317 using the dataset in the current study. Xu (Xu, 2019) explored Fourier Transform method introduced into convolutional neural network and proposed a Fourier W-Net for cilia segmentation. Although the model performance is not satisfactory, the proposed work utilized the frequency spectrum by considering signals in both spatial and temporal domain. Zain et al. (Zain et al., 2020) proposed an unsupervised spatiotemporal representation of cilia videos. In this pipeline, the authors proposed a VAE generative seq2seq module to embed the temporal dynamics. This work is still in early phase and more research will be done in utilizing temporal information in segmenting cilia.

4.5.3 Gabor filter implementation

In this study, we only use selected GF as a data augmentation tool instead of incorporating GF in the encoder part of U-Net. We generated a set of 54 GFs and used random forest to select meaningful GFs with high importance. We did not embed GFs into the U-Net encoder because modifying the U-Net architecture is an intricate task. For the future work, we will explore the GF-embedded U-Net for cilia segmentation.

4.5.4 Unbalanced data

Unbalanced data is critical issue in the field of machine learning due to unbalanced class distribution. Cilia segmentation can be viewed as a image classification task where cilia class is the is the minority class and non-cilia class is the majority class. The ratio between those two classes in our original video dataset is around 1:12. After image preprocessing, we selected patches with more than 5% cila area, cilia pixels account for 23% in our dataset. Figure 4.5 shows cilia and non-cilia distribution of the dataset in the current study.

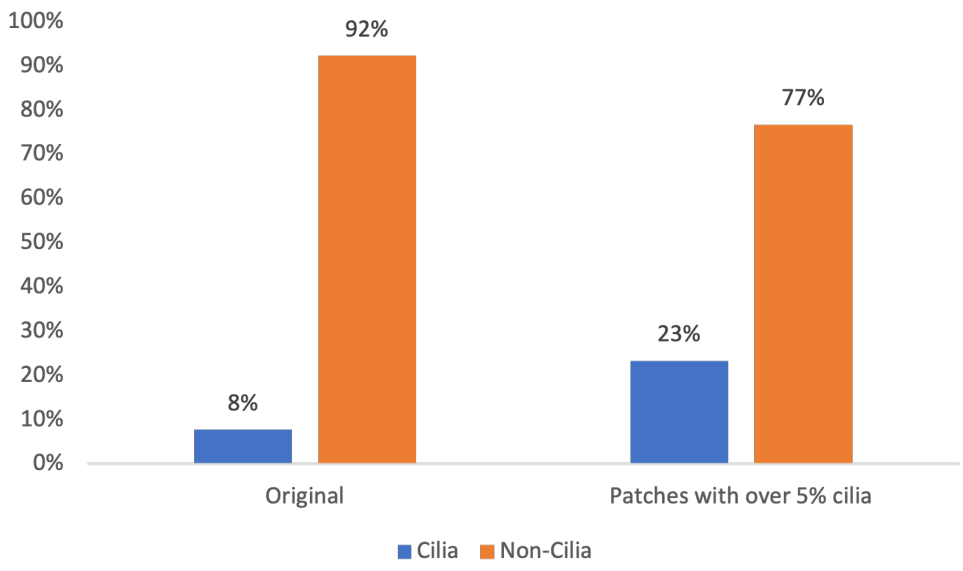

Figure 4.5: Cilia and non-cilia distribution

There are several techniques including traditional oversampling, undersampling techniques, and the Synthetic Minority Over-sampling Technique (SMOTE) to handle unbalanced data in image segmentation. Focal loss function (Lin et al., 2017) is a deep learning approach to address class imbalance issue. Currently, we use binary cross-entropy (BCE) loss function, which calculates the logarithm value of the output. However, BCE can not handle unbalanced dataset very well even with added weight. Focal loss overcomes the unbalanced dataset problem by adding an extra term to reduce the impact of correct predictions and focus on incorrect examples. The equation of focal loss is shown as following:

$$FL = -\alpha_c(1 - \hat{y}_i)^\gamma * log\hat{y}_i$$

where γ is a hyperparameter that can restrain the power of this reduction. We implemented IoU for evaluating our models, but IoU can also be used as a loss function for image segmentation tasks. Van Beers et al. (van Beers et al., 2019) tested IoU loss and BCE loss for 2 models on multiple datasets and found that IoU loss increases model performances significantly compared to training on BCE loss.

4.5.5 Cross validation

Cross validation (cv) is an important model validation technique in machine learning (ML) field to evaluate a model and test its performance. CV helps to select an satisfactory model for the specific modeling problem thus commonly used in ML tasks. There are several CV techniques: leave-one-out, hold-out, k-folds, stratified k-folds and so on. K-fold introduces a number of folds-k and split the dataset into k equal

parts, where k-1 folds are the training set while the remaining is the testing set. During each cv iteration, a new model is trained independently. Using k-fold technique can take full advantage of the dataset and give a stable and trustworthy result. However, k-fold method is not commonly used in deep learning due to its high cost and time-consuming.

Table 4.5: Performances comparison of FC-DenseNet with or without cv over the cilia dataset

Model	k-fold cv	Evaluation Metrics				
		IoU	Accuracy	F1	Precision	Recall
FC-DenseNet 103	No	0.3180	0.8034	0.4472	0.4732	0.5035
FC-DenseNet 103	Yes	0.3599	0.8164	0.4892	0.5232	0.5475

We trained the FC-DenseNet with k-fold cv where k is set to 5. Table 4.5 shows the performances of FC-DenseNet model with or without cv over cilia dataset. k-fold cv improves FC-DenseNet 103 model performance by 13% in terms of IoU score. We did not apply k-fold cv on U-Net-based models due to the fact that cv is computationally very expensive and time-consuming in the process of training k different models. Although FC-DenseNet model with k-fold cross-validation can improve the model performances, it did not outperform Gabor composite models.

Chapter 5

Conclusion and future work

In conclusion, this study presents a texture image segmentation using GFs and the U-Net model. We generated a set of 54 GFs and used random forest to select meaningful GFs with high feature importances. The Gabor filtered feature maps are concatenated to the original image input as an additional channel to train the U-Net model. The addition of features extracted by our selected GFs improved the performances of the base model. By comparing the GF composite model with the previous studies (intensity based-thresholding (Quinn et al., 2015), FC-DenseNet (Lu et al., 2018; Zain et al., 2020), ZCA/SAE with U-Net (Zain et al., 2022)), we showed that the presented framework outperformed the non-deep learning and deep-learning based models in the previous studies in terms of IoU score.

For the future work, more validated cilia videos with accurate ground-truth masks will be added to the current dataset to improve overall performances of the proposed framework. Currently, we use random forest to select meaningful GFs with high feature importance. For future direction, we will focus on developing an optimized GF bank with automatic variation of GF parameters. Also, we will introduce a reduced number of Gabor filter banks into the U-Net encoder. Although U-Net can perform feature extraction and localization, but the modification of the architectures or layers is often an intricate task. For the future direction, we will embed a multi-channel Gabor-filtering in the U-Net model for cilia segmentation.

Bibliography

Amit, Y., & Geman, D. (1997). Shape quantization and recognition with randomized trees. *Neural computation*, *9*(7), 1545–1588.

Anand, V., Gupta, S., Koundal, D., Nayak, S. R., Barsocchi, P., & Bhoi, A. K. (2022). Modified u-net architecture for segmentation of skin lesion. *Sensors*, *22*(3), 867.

Arzt, M., Deschamps, J., Schmied, C., Pietzsch, T., Schmidt, D., Tomancak, P., Haase, R., & Jug, F. (2022). Labkit: Labeling and segmentation toolkit for big image data. *Frontiers in computer science*, 10.

Bansal, R., Engle, S. E., Kamba, T. K., Brewer, K. M., Lewis, W. R., & Berbari, N. F. (2021). Artificial intelligence approaches to assessing primary cilia. *JoVE (Journal of Visualized Experiments)*, (171), e62521.

Brightman, M., & Palay, S. (1963). The fine structure of ependyma in the brain of the rat. *The Journal of cell biology*, *19*(2), 415–439.

Bustamante-Marin, X. M., & Ostrowski, L. E. (2017). Cilia and mucociliary clearance. *Cold Spring Harbor perspectives in biology*, *9*(4), a028241.

Cheriet, M., Said, J. N., & Suen, C. Y. (1998). A recursive thresholding technique for image segmentation. *IEEE transactions on image processing*, *7*(6), 918–921.

Chilvers, M., McKean, M., Rutman, A., Myint, B., Silverman, M., & O'Callaghan, C. (2001). The effects of coronavirus on human nasal ciliated respiratory epithelium. *European Respiratory Journal*, *18*(6), 965–970.

Dalal, N., & Triggs, B. (2005). Histograms of oriented gradients for human detection. *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, *1*, 886–893.

Dirksen, E. R. (1971). Centriole morphogenesis in developing ciliated epithelium of the mouse oviduct. *The Journal of cell biology*, *51*(1), 286–302.

Djenoune, L., Mahamdeh, M., Truong, T. V., Nguyen, C. T., Fraser, S. E., Brueckner, M., Howard, J., & Yuan, S. (2023). Cilia function as calcium-mediated mechanosensors that instruct left-right asymmetry. *Science*, *379*(6627), 71–78.

Dunn, D., Higgins, W. E., & Wakeley, J. (1994). Texture segmentation using 2-d gabor elementary functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *16*(2), 130–149.

Eshel, D., Grossman, Y., & Priel, Z. (1985). Spectral characterization of ciliary beating: Variations of frequency with time. *American Journal of Physiology-Cell Physiology*, *249*(1), C160–C165.

Gabor, D. (1946). Theory of communication. part 1: The analysis of information. *Journal of the Institution of Electrical Engineers-part III: radio and communication engineering*, *93*(26), 429–441.

Goetz, S. C., & Anderson, K. V. (2010). The primary cilium: A signalling centre during vertebrate development. *Nature reviews genetics*, *11*(5), 331–344.

Greenstone, M., Dewar, A., & Cole, P. (1983). Ciliary dyskinesia with normal ultrastructure. *Thorax*, *38*(11), 875.

Grigorescu, S. E., Petkov, N., & Kruizinga, P. (2002). Comparison of texture features based on gabor filters. *IEEE Transactions on Image processing*, *11*(10), 1160–1167.

Hansen, J. N., Rassmann, S., Stüven, B., Jurisch-Yaksi, N., & Wachten, D. (2021). Ciliaq: A simple, open-source software for automated quantification of ciliary morphology and fluorescence in 2d, 3d, and 4d images. *The European Physical Journal E*, *44*, 1–26.

Hildebrandt, F., Benzing, T., & Katsanis, N. (2011). Ciliopathies. *New England Journal of Medicine*, *364*(16), 1533–1543.

Horani, A., Ferkol, T. W., Dutcher, S. K., & Brody, S. L. (2016). Genetics and biology of primary ciliary dyskinesia. *Paediatric respiratory reviews*, *18*, 18–24.

Idrissa, M., & Acheroy, M. (2002). Texture classification using gabor filters. *Pattern Recognition Letters*, *23*(9), 1095–1102.

Jeffery, P., & Reid, L. (1975). New observations of rat airway epithelium: A quantitative and electron microscopic study. *Journal of anatomy*, *120*(Pt 2), 295.

Kim, W., Han, T. H., Kim, H. J., Park, M. Y., Kim, K. S., & Park, R. W. (2011). An automated measurement of ciliary beating frequency using a combined optical flow and peak detection. *Healthcare informatics research*, *17*(2), 111–119.

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Knowles, M. R., Daniels, L. A., Davis, S. D., Zariwala, M. A., & Leigh, M. W. (2013). Primary ciliary dyskinesia. recent advances in diagnostics, genetics, and characterization of clinical disease. *American journal of respiratory and critical care medicine*, *188*(8), 913–922.

Knowles, M. R., Leigh, M. W., Carson, J. L., Davis, S. D., Dell, S. D., Ferkol, T. W., Olivier, K. N., Sagel, S. D., Rosenfeld, M., Burns, K. A., et al. (2012). Mutations of dnah11 in patients with primary ciliary dyskinesia with normal ciliary ultrastructure. *Thorax*, *67*(5), 433–441.

Kong, W. K., Zhang, D., & Li, W. (2003). Palmprint feature extraction using 2-d gabor filters. *Pattern recognition*, *36*(10), 2339–2347.

Lauring, M. C., Zhu, T., Luo, W., Wu, W., Yu, F., & Toomre, D. (2019). New software for automated cilia detection in cells (acdc). *Cilia*, *8*(1), 1–21.

Lee, C.-J., Wang, S.-D., et al. (1999). Fingerprint feature extraction using gabor filters. *Electronics Letters*.

Lin, T.-Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2017). Focal loss for dense object detection. *Proceedings of the IEEE international conference on computer vision*, 2980–2988.

Long, J., Shelhamer, E., & Darrell, T. (2015). Fully convolutional networks for semantic segmentation. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3431–3440.

Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, *60*, 91–110.

Lu, C., Marx, M., Zahid, M., Lo, C., Chennubhotla, C., & Quinn, S. P. (2018). Stacked neural networks for end-to-end ciliary motion analysis. *arXiv preprint arXiv:1803.07534*.

Luan, S., Chen, C., Zhang, B., Han, J., & Liu, J. (2018). Gabor convolutional networks. *IEEE Transactions on Image Processing*, *27*(9), 4357–4366.

McConnachie, D. J., Stow, J. L., & Mallett, A. J. (2021). Ciliopathies and the kidney: A review. *American Journal of Kidney Diseases*, *77*(3), 410–419.

Mehrotra, R., Namuduri, K. R., & Ranganathan, N. (1992). Gabor filter-based edge detection. *Pattern recognition*, *25*(12), 1479–1494.

Mitchison, H. M., & Valente, E. M. (2017). Motile and non-motile cilia in human pathology: From function to phenotypes. *The Journal of pathology*, *241*(2), 294–309.

Nguyen-Thanh, N., Pham, T., & Ichikawa, K. (2019). Automated detection and segmentation of mitochondrial images based on gradient enhancement and adaptive gabor filter.

Omran, H., Haffner, K., Volkel, A., Kuehr, J., Ketelsen, U.-P., Ross, U.-H., Konietzko, N., Wienker, T., Brandis, M., & Hildebrandt, F. (2000). Homozygosity mapping of a gene locus for primary ciliary dyskinesia on chromosome 5p and identification of the heavy dynein chain dnah5 as a candidate gene. *American journal of respiratory cell and molecular biology*, *23*(5), 696–702.

Porat, M., & Zeevi, Y. Y. (1988). The generalized gabor scheme of image representation in biological and machine vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *10*(4), 452–468.

Quinn, S. P., Zahid, M. J., Durkin, J. R., Francis, R. J., Lo, C. W., & Chennubhotla, S. C. (2015). Automated identification of abnormal respiratory ciliary motion in nasal biopsies. *Science translational medicine*, *7*(299), 299ra124–299ra124.

Rahman, M. T., & Bhuiyan, M. A. (2008). Face recognition using gabor filters. *2008 11th International Conference on Computer and Information Technology*, 510–515.

Ram, S., Majdi, M. S., Rodriguez, J. J., Gao, Y., & Brooks, H. L. (2018). Classification of primary cilia in microscopy images using convolutional neural random forests. *2018 IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI)*, 89–92.

Reyes, A. A., Paheding, S., Deo, M., & Audette, M. (2022). Gabor filter-embedded u-net with transformer-based encoding for biomedical image segmentation. *Multiscale Multimodal Medical Imaging: Third International Workshop, MMMI 2022, Held in Conjunction with MICCAI 2022, Singapore, September 22, 2022, Proceedings*, 76–88.

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, 234–241.

Salathe, M. (2007). Regulation of mammalian ciliary beating. *Annu. Rev. Physiol.*, *69*, 401–422.

Sampaio, P., da Silva, M. F., Vale, I., Roxo-Rosa, M., Pinto, A., Constant, C., Pereira, L., Quintão, C. M., & Lopes, S. S. (2021). Ciliarmove: New software for evaluating ciliary beat frequency helps find novel mutations by a portuguese multidisciplinary team on primary ciliary dyskinesia. *ERJ open research*, *7*(1).

Schmidt, U., Weigert, M., Broaddus, C., & Myers, G. (2018). Cell detection with star-convex polygons. *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part II 11*, 265–273.

Shahedi, M., Devi, A., Dormer, J., & Fei, B. (2020). A study on u-net limitations in object localization and image segmentation. *Society for Imaging Informatics in Medicine (SIIM), Virtual meeting*.

Shoemark, A., Dixon, M., Corrin, B., & Dewar, A. (2012). Twenty-year review of quantitative transmission electron microscopy for the diagnosis of primary ciliary dyskinesia. *Journal of clinical pathology*, *65*(3), 267–271.

Smith, C. M., Djakow, J., Free, R. C., Djakow, P., Lonnen, R., Williams, G., Pohunek, P., Hirst, R. A., Easton, A. J., Andrew, P. W., et al. (2012). Ciliafa: A research tool for automated, high-throughput measurement of ciliary beat frequency using freely available software. *Cilia*, *1*(1), 1–7.

Spassky, N., & Meunier, A. (2017). The development and functions of multiciliated epithelia. *Nature reviews Molecular cell biology*, *18*(7), 423–436.

Stringer, C., & Pachitariu, M. (2022). Cellpose 2.0: How to train your own model. *BioRxiv*, 2022–04.

Stringer, C., Wang, T., Michaelos, M., & Pachitariu, M. (2021). Cellpose: A generalist algorithm for cellular segmentation. *Nature methods*, *18*(1), 100–106.

Tadic, V., Loncar-Turukalo, T., Odry, A., Trpovski, Z., Toth, A., Vizvari, Z., & Odry, P. (2021). A note on advantages of the fuzzy gabor filter in object and text detection. *Symmetry*, *13*(4), 678.

Tilley, A. E., Walters, M. S., Shaykhiev, R., & Crystal, R. G. (2015). Cilia dysfunction in lung disease. *Annual review of physiology*, *77*, 379–406.

UYSAL, N., YOZGATLI, T. K., YILDIZCAN, E. N., Emre, K., GEZER, M., & BAŞTU, E. (2022). Comparison of u-net based models for human embryo segmentation. *Bilişim Teknolojileri Dergisi*, *15*(1), 35–44.

van Beers, F., Lindström, A., Okafor, E., & Wiering, M. A. (2019). Deep neural networks with intersection over union loss for binary image segmentation. *ICPRAM*, 438–445.

Waters, A. M., & Beales, P. L. (2011). Ciliopathies: An expanding disease spectrum. *Pediatric nephrology*, *26*, 1039–1056.

Weldon, T. P., Higgins, W. E., & Dunn, D. F. (1996). Efficient gabor filter design for texture segmentation. *Pattern recognition*, *29*(12), 2005–2015.

Wolny, A., Cerrone, L., Vijayan, A., Tofanelli, R., Barro, A. V., Louveaux, M., Wenzl, C., Strauss, S., Wilson-Sánchez, D., Lymbouridou, R., et al. (2020). Accurate and versatile 3d segmentation of plant tissues at cellular resolution. *Elife*, *9*, e57613.

Xu, W. (2019). *Cilia segmentation in medical videos with fourier convolutional neural network* (Doctoral dissertation). University of Georgia.

Zain, M., Miller, E., Quinn, S., & Lo, C. (2022). Low level feature extraction for cilia segmentation. *Proceedings of the Python in Science Conference*.

Zain, M., Rao, S., Safir, N., Wyner, Q., Humphrey, I., Eldridge, A., Li, C., AlAila, B., & Quinn, S. (2020). Towards an unsupervised spatiotemporal representation of cilia video using a modular generative pipeline. *Proceedings of the Python in Science Conference*.

Zariwala, M. A., Knowles, M. R., & Leigh, M. W. (2019). Primary ciliary dyskinesia. *GeneReviews®[Internet]*.
Zhang, P., Kiseleva, A. A., Korobeynikov, V., Liu, H., Einarson, M. B., & Golemis, E. A. (2019). Microscopy-based automated live cell screening for small molecules that affect ciliation. *Frontiers in genetics*, *10*, 75.