



Universiteit  
Leiden  
The Netherlands

# Bachelor Bioinformatics

Analysis of the 3D Distribution of Granulomas in adult Zebrafish

A.M.L. Laauwen  
(s2960710)

Supervisors:  
F.J. Verbeek  
S. Javanmardi

BACHELOR THESIS

Leiden Institute of Advanced Computer Science (LIACS)  
[www.liacs.leidenuniv.nl](http://www.liacs.leidenuniv.nl)

05/08/2023

## Abstract

Tuberculosis (TB), caused by *Mycobacterium tuberculosis*, remains a leading global cause of death ([World Health Organization, 2022](#)). Consequently, research into its infection mechanisms and preventive strategies is of paramount importance. *Mycobacterium tuberculosis* is the bacterium that causes tuberculosis in humans. For research on the infection, the zebrafish has successfully been used as a model organism. For studies in zebrafish, *Mycobacterium marinum* is used; this bacterium resembles the human counterpart. Both infections caused by *Mycobacterium marinum* and *Mycobacterium tuberculosis* share a common feature, which is the presence of clusters of macrophages containing the bacteria. These clusters are referred to as granulomas. In general, the formation of granulomas is studied in larval stages of zebrafish. In this study we analyze the distribution of granulomas in adult zebrafish. The starting point is serial sections in which the granulomas can be well observed. The sections are digitized as Whole Slide Images (WSIs), which results in very high-resolution images. To analyze the 3D distribution of the granulomas, a 3D reconstruction from these WSIs needs to be realized. We have designed a protocol to preprocess the WSIs and submit them to a sequential registration procedure. The resulting set of registered slides was used for annotation to a 3D reconstruction highlighting the different instances of granulomas in the zebrafish. Furthermore, we employed the labels from this 3D reconstruction in a segmentation neural network. This strategy paves the way for future reconstructions to harness this trained network in pinpointing granulomas.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Thesis overview . . . . .	2
<b>2</b>	<b>Background</b>	<b>3</b>
2.1	Zebrafish as a Model Organism . . . . .	3
2.1.1	Zebrafish as a Model Organism for Tuberculosis Research . . . . .	3
2.1.2	Tuberculosis Pathogenesis . . . . .	4
2.2	Whole Slide Images . . . . .	5
2.2.1	BigTIFF format for WSI . . . . .	5
2.2.2	Serial Sections . . . . .	6
2.2.3	Staining of Slides . . . . .	6
2.3	Image Registration . . . . .	7
2.4	3D Reconstruction from Serial Sections . . . . .	8
2.5	Image Segmentation using Deep Learning . . . . .	8
2.5.1	Architecture of Image Segmentation Networks . . . . .	8
2.5.2	Applications of Image Segmentation using Deep Learning . . . . .	9
2.5.3	Performance Measures for Neural Networks . . . . .	9
<b>3</b>	<b>Materials &amp; Methods</b>	<b>11</b>
3.1	Zebrafish Material . . . . .	11
3.2	Image Preprocessing . . . . .	11
3.3	Image Registration . . . . .	12
3.3.1	VALIS: Virtual Alignment of pathoLogY Image Series . . . . .	12
3.3.2	Registration Error . . . . .	14
3.3.3	Simultaneous Registration of Images . . . . .	14
3.3.4	Group-wise Registration of Images . . . . .	14
3.4	3D Reconstruction . . . . .	15
3.5	Analysis with Segmentation Network . . . . .	16
<b>4</b>	<b>Results</b>	<b>19</b>
4.1	Workflow . . . . .	19
4.2	Cropped Images . . . . .	21
4.3	Image Registration . . . . .	21
4.3.1	Simultaneous Registration of Serial Sections . . . . .	21
4.3.2	Group-wise Registration of Serial Sections . . . . .	22
4.4	3D Reconstruction . . . . .	27
4.5	Analysis with Segmentation Network . . . . .	29
<b>5</b>	<b>Conclusion &amp; Discussion</b>	<b>31</b>
5.1	Conclusion . . . . .	31
5.2	Discussion . . . . .	32
5.3	Future Research . . . . .	33
<b>6</b>	<b>Acknowledgements</b>	<b>34</b>
	<b>References</b>	<b>35</b>

# 1 Introduction

Tuberculosis (TB) remains a significant worldwide health challenge, necessitating improvement in diagnostic methods and treatment approaches. Given the complicated nature of TB’s pathogenesis and its widespread effects on people across the globe, continuing the research regarding its development is crucial. Tuberculosis is caused by a *Mycobacterium tuberculosis* infection. A distinctive feature of this infection is the development of granulomas, which are immune cell aggregates that form around the site of infection. This study uses zebrafish as a model organism for tuberculosis to investigate the three-dimensional distribution of granulomas. By infecting adult zebrafish with *Mycobacterium marinum*, a bacterium that closely mimics tuberculosis progression, valuable insights into disease dynamics and host-pathogen interactions can be obtained.

The zebrafish has already played an important role as model organisms in numerous tuberculosis studies. Previous studies predominantly concentrated on zebrafish during their early developmental stages, making use of their optical transparency in these life stages (Clay et al., 2007; Tobin & Ramakrishnan, 2008; Ordas et al., 2015). However, these studies were limited to observing only the initial stages of disease progression. In contrast, this study focuses on adult zebrafish as the model organism, allowing for the examination of tuberculosis at later stages of progression. During earlier life stages, zebrafish exclusively rely on innate immunity; however, as they mature into adults, they develop adaptive immunity as well, similar to that of mammals (Bouz & Al Hasawi, 2018). While the use of adult zebrafish as tuberculosis models has been explored, most research in this field has primarily adopted a two-dimensional perspective (Myllymäki, Bäuerlein, & Rämets, 2016; Cheng, Kam, Johansen, & Oehlers, 2020).

By investigating the disease in three dimensions, valuable insights can be gained into the spatial relationships between granulomas and other organs, thereby advancing comprehension of the disease. The aim is to answer the research question: “Can a reliable analysis be performed on the distribution of granulomas in a 3D reconstruction of adult zebrafish?” By investigating this, insights can be gained into the intricate patterns of granuloma distribution and their significance in tuberculosis progression. The research process involves image registration, 3D reconstruction, and granuloma distribution analysis, which can be further divided into three sub-questions that contribute to answering the main research question:

1. Can a reliable registration be performed on WSIs of adult zebrafish using VALIS software?
2. Can a proper 3D reconstruction of granulomas be made from the registered WSIs of adult zebrafish?
3. Can a segmentation network be used to identify the granulomas in adult zebrafish?

## 1.1 Thesis overview

This bachelor thesis was conducted at the Leiden Institute of Advanced Computer Science (LIACS) under the supervision of Fons J. Verbeek and Shima Javanmardi. The objective of this study is to investigate the three-dimensional distribution of granulomas in adult zebrafish infected with *Mycobacterium marinum*, utilizing zebrafish for tuberculosis research. The thesis is divided into several sections to provide a comprehensive understanding of the topic.

Chapter 1 introduces the thesis, emphasizing zebrafish's significance in tuberculosis research. This section establishes the groundwork for the following chapters by setting up the research objectives.

Chapter 2 explores the background information necessary to comprehend the research. It begins by discussing zebrafish as a model organism and highlights its relevance in tuberculosis research. Whole slide imaging techniques and staining methods in pathology are discussed, along with the BigTIFF format used for handling large-scale images. Image registration and the error measures for registration are discussed. The concept of 3D reconstruction from serial sections is explored, followed by an overview of the TDR-3Dbase software. Lastly, the application of deep learning in image segmentation is covered, addressing different architectures, various applications, and performance measures.

Chapter 3 provides detailed information on the materials and methods employed in the research. It covers the zebrafish material used and the preprocessing techniques applied to the serial section images of these zebrafish. The image registration process is discussed, specifically the Virtual Alignment of pathoLogY Image Series (VALIS) pipeline, and the two different approaches used in this research were detailed. Additionally, the 3D reconstruction process and the analysis using segmentation networks are explained. Chapter 4 presents the results obtained from the conducted experiments.

Chapter 5 concludes the thesis, offering a concise summary of the research findings and their significance. It also discusses the challenges within the research and potential avenues for future studies.

## 2 Background

To understand the significance and relevance of the research, it is crucial to establish a well-rounded background that discusses the essential aspects. The essential aspects concerning this research include highlighting the importance of zebrafish as a prominent model organism, examining the biology of granulomas in zebrafish, and understanding image registration and 3D reconstruction from serial sections. Moreover, this background section will provide insights into using a deep learning segmentation network, which will be used to analyze the distribution of granulomas in zebrafish.

### 2.1 Zebrafish as a Model Organism

The zebrafish (*Danio rerio*) is a versatile model organism used across various research fields, including developmental biology, genetics, immunology, and drug discovery (Thisse & Zon, 2002; Trede, Langenau, Traver, Look, & Zon, 2004; Meijer & Spaink, 2011). Numerous advantages make the zebrafish an ideal model for studying various vertebrate processes (Meijer & Spaink, 2011). With an adult size between 3 to 5 centimeters, zebrafish can achieve a high population density, housing up to 5 fish per liter. Furthermore, their rapid reproduction cycle permits a single female to lay up to 300 eggs per week. The optical transparency of zebrafish embryos and larvae enables non-invasive imaging techniques, providing high-resolution views of cellular and subcellular processes. Additionally, zebrafish exhibit genetic and physiological similarities to mammals, surpassing invertebrate models in this regard (Thisse & Zon, 2002). These unique attributes collectively make the zebrafish a valuable and highly relevant model for diverse scientific investigations.

#### 2.1.1 Zebrafish as a Model Organism for Tuberculosis Research

One specific area of research that has used zebrafish as a model organism is the field of bacterial infection studies. Zebrafish models in bacterial infection research can reveal novel host-pathogen interactions and aid in developing innovative treatments. Zebrafish have already been used to study a range of bacterial infections, including *Salmonella typhimurium*, *Streptococcus*, and *Mycobacterium marinum* infection (van der Sar, 2003; Neely, Pfeifer, & Caparon, 2002; Stoop et al., 2011). In this research specifically, the *M. marinum* infection will be studied. The significance of *M. marinum* infection in zebrafish lies in its resemblance to the infection caused by *Mycobacterium tuberculosis* in humans, the pathogen responsible for tuberculosis (TB) (Stoop et al., 2011). In 2021, TB caused approximately 1.4 million deaths (World Health Organization, 2022). Despite progress in mitigating its global impact, TB remains a significant public health concern, especially in countries with limited resources and lower income levels.

In the past, zebrafish have already been used as a model organism in numerous tuberculosis studies. These studies have mainly relied on using zebrafish in their earlier life stages as model organisms, also making use of the optical transparency characteristic they possess in these life stages (Clay et al., 2007; Tobin & Ramakrishnan, 2008; Ordas et al., 2015). The scope of these studies was primarily focused on the initial phases of disease progression, and the examination of zebrafish during these early life stages only encompassed their innate immune response. In order to

investigate the later stages of disease progression and the additional impacts of adaptive immunity, adult zebrafish serve as a suitable model. Adult zebrafish possess an immune system similar to that of mammals, having both innate and adaptive immunity (Bouz & Al Hasawi, 2018). The adult zebrafish has already been used in studies regarding this topic (Myllymäki et al., 2016; Cheng et al., 2020). In these studies it has been shown that the adult zebrafish infected by *M. marinum* closely mimics the disease phases observed in human tuberculosis, this includes latency and reactivation. This makes the adult zebrafish an excellent model for studying different stages of TB, as well as the host immune responses and bacterial adaptations that occur during each phase. The zebrafish model also holds promise for identifying biomarkers for diagnosing different stages of TB and can contribute to the development of new drugs and vaccines.

### 2.1.2 Tuberculosis Pathogenesis

Understanding TB infection's mechanisms is crucial for enhancing its treatment. Humans typically contract *M. tuberculosis* when the bacterium is inhaled, predominantly infecting the lungs. In nature, zebrafish can become infected with *M. marinum* through exposure to water contaminated by infected fish or biofilms in their environment (Swaim et al., 2006). To artificially infect a zebrafish, researchers typically administer the pathogen, such as *Mycobacterium marinum*, through an intraperitoneal injection. (van Leeuwen, van der Sar, & Bitter, 2015). Intraperitoneal refers to a route or method of administration in which a substance is directly injected or introduced into the peritoneal cavity. The peritoneal cavity is the space within the abdomen that surrounds the abdominal organs, such as the stomach, liver, intestines, and spleen.

Once the pathogen enters a macrophage and triggers the activation of the immune response within the host tissue, the host responds by initiating the formation of a cellular mass called a granuloma (Russell, 2007). Granulomas are a defining characteristic of tuberculosis infection and will form as a result of both *M. marinum* and *M. tuberculosis* (Myllymäki et al., 2016). In granuloma formation, macrophages are recruited to the site of infection. After this, other immune cells accumulate and surround the offending agent as well, forming a central core. Granulomas act as a physical barrier, confining pathogens within a localized area. The formation of granulomas in zebrafish as an effect of infection by *M. marinum* has a lot of genetic components similar to granuloma formation in humans as a result of *M. tuberculosis*, and are structurally similar as well. The main distinction lies in the location of the granulomas (Stoop et al., 2011). In humans, *M. tuberculosis* typically leads to the formation of granulomas mainly in the lungs. Zebrafish lack lungs, and granulomas resulting from *M. marinum* primarily form in fatty tissues such as the pancreas, liver, spleen, and gonads. But since the processes of formation are still highly similar, the zebrafish model offers the opportunity to study tuberculosis pathogenesis, providing valuable insights into the disease progression and potential treatment strategies. The parallels between the two infections make *M. marinum* a valuable tool for advancing our understanding of tuberculosis and developing new remedies.

## 2.2 Whole Slide Images

In the past, digital imaging of microscopy slides relied on static cameras capturing images of these slides, resulting in only specific regions of a glass slide being captured (Farahani, Parwani, Pantanowitz, et al., 2015). This method was inconvenient due to the potential loss of important details since it was impossible to see the image as a whole.

As a solution for this limitation, robotic microscopy was introduced. This type of microscopy automated the process of capturing a whole microscopy slide instead of having to take multiple pictures one by one manually. In this process, multiple images of the same slide were captured and put together, resulting in the composite slide image. This approach enables the user to view the whole slide at high resolution, including regions that were previously missed when using static cameras.

More progress and developments in this area finally led to whole slide image scanners, capable of producing high-resolution digital slides in a short period of time. This method also captures multiple pictures of the slide and finally puts them together to create one larger image. This technique of whole slide imaging is now a well-received technique, aiming to emulate conventional light microscopy digitally (Pantanowitz et al., 2011).

### 2.2.1 BigTIFF format for WSI

Different WSI systems support various image formats for acquiring WSIs, with the possibilities depending on the model of the specific WSI system (Farahani et al., 2015). One such format is the BigTIFF format, the image format used in this research. The BigTIFF format is an extension of the TIFF (Tagged Image File Format) format. One of the strengths of using the TIFF format is its organization of image data, which allows for the preservation of high-quality images and their metadata in a lossless manner. Whereas the TIFF file format uses 32-bit offsets, limiting the file to 4 gigabytes, the BigTIFF format uses 64-bit offsets instead, allowing for bigger files (Aware Systems, 2008). Offsets refer to the positions or addresses within a file where specific data or information is located.

TIFF images are valuable because they organize image data in a way that allows for quick and easy random access (Pillay, 2023). This benefit is achieved through tiling, a technique used to divide images into smaller tiles, and the creation of multi-resolution image pyramids, of which a schematic representation can be seen in Figure 1. Tiling allows for quick decoding of specific regions within an image without the need to decode the entire image. Furthermore, the use of multi-resolution pyramids enables storing multiple resolutions of an image in a hierarchical structure, facilitating rapid access to images of any size. This powerful combination of tiling and multi-resolution pyramids grants efficient and speedy random access to any part of an image, even with extremely large images.

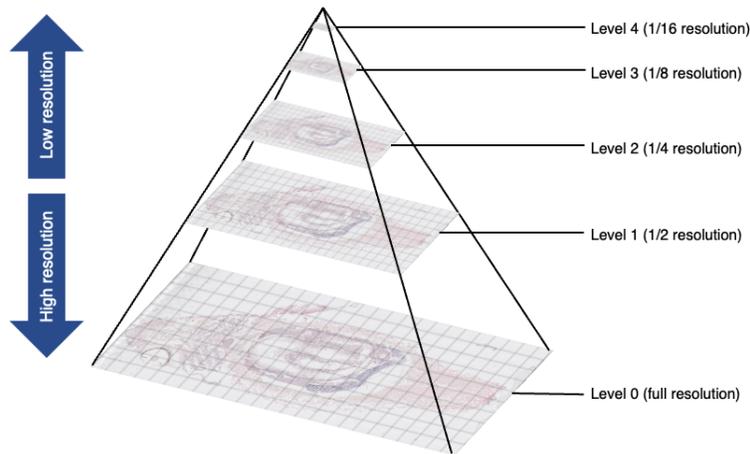


Figure 1: An illustration of how digital WSI files in TIFF format are stored in a pyramid structure. In this example, there are five resolution levels (levels 0-4). Lower resolutions are a down-sampled copy of the highest resolution at level 0.

### 2.2.2 Serial Sections

Serial sectioning refers to the process of obtaining a series of thin slices from a physical object or tissue sample (Verbeek, 1999). These sections allow for a detailed analysis of the internal structures by capturing individual layers of the object. Slides containing serial sections can also be converted into WSIs for further digital analysis. However, a common issue that arises when physically dividing an object into sections is the possibility of compromising its structural integrity in the third dimension, as well as possible deformations of the object in the section.

### 2.2.3 Staining of Slides

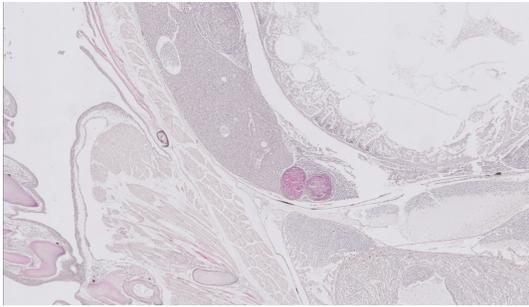
When creating WSIs, it is customary to stain the slides before they are digitized. There are several techniques for staining that can improve the visibility of tissue structures and cellular components. Once the slides are appropriately stained, they can be scanned using specialized digital scanners to generate high-resolution WSIs. Two different staining methods were used for acquiring the WSIs in this research: a Ziehl-Neelsen (ZN) stain and a hematoxylin and eosin (HE) stain.

The ZN stain is a type of acid-fast stain developed to exploit the acid fastness of the mycobacterial genus (Talbot & Raffa, 2015). This means it can be used for *M. marinum* and *M. tuberculosis*, giving the bacteria a red or pink color. This staining has been widely used as a test for the diagnosis of tuberculosis in resource-limited settings, as the staining is affordable and results are rapidly available.

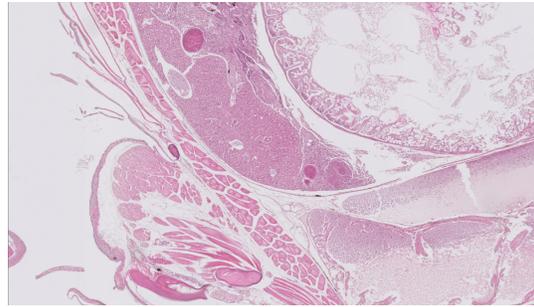
The HE stain is a crucial tissue stain in histology, widely used in medical diagnosis (Titford, 2005). This stain provides a general differentiation between tissue types and cellular components. It consists of two components: hematoxylin and eosin. Hematoxylin is a basic dye that stains acidic structures like cell nuclei, presenting them in shades of blue or purple. It selectively binds to DNA and RNA in the nucleus, allowing the visualization of individual cell locations and morphology. On the other hand, eosin is an acidic dye that stains basic tissue components,

such as cytoplasm and extracellular matrix, in shades of pink or red. This differentiation aids in distinguishing various tissue types and cellular components based on their staining characteristics.

In Figure 2 examples of the two different staining methods can be seen. The images were taken from the dataset used in the experiments. Figure 2a shows a WSI with a ZN stain and in Figure 2b a WSI with a HE stain can be seen. In both images the granulomas can be seen as circular pink structures. It is important to note that these figures are mere examples, in the actual dataset the intensity of the coloring may differ from image to image.



(a) Ziehl-Neelsen (ZN) stain



(b) Hematoxylin and Eosin (HE) stain

Figure 2: Serial sections of zebrafish infected with *Mycobacterium marinum*. In both the ZN stain (a) and the HE stain (b), the granulomas are prominently visualized as circular pink structures.

## 2.3 Image Registration

Image registration is the process of aligning multiple images of the same scene, which may have been captured at different times or from different perspectives (Zitová & Flusser, 2003). This alignment process aims to match the images geometrically. Usually, a reference image is used to align the other images. When the images have a specified order, they can be aligned serially, which means that the previous image in the dataset will be used as a reference image for the next.

We distinguish between two different registration methods, rigid and non-rigid, which differ in the types of transformations allowed to match the images to each other (Zhang, Dong, Gao, & Zhao, 2020). When performing rigid image registration, only rotation and translation operations are used to modify the images. Using these types of transformations causes the relationships between pixels to be preserved before and after the transformation. On the other hand, non-rigid image registration enables non-uniform mapping between the images, which results in changing the shape of the images. In this case, the registration will change the pixel-wise relations.

Both types of image registration play a crucial role in various applications, including medical imaging, computer vision, and more (Zitová & Flusser, 2003). It facilitates tasks such as image fusion, change detection, object tracking, and image analysis, enabling a comprehensive understanding and integration of multiple images. In the context of this research, registered serial section images are especially valuable, as they can be used to create a three-dimensional reconstruction.

## 2.4 3D Reconstruction from Serial Sections

Three-dimensional (3D) reconstruction is a valuable tool and is used to study the spatial structure, morphology, and organization of objects in three-dimensional space. One approach to creating such reconstructions involves using 2D serial sections, which capture consecutive slices or perspectives of an object or specimen (Verbeek, 1999). By drawing and layering contours from these serial sections, a comprehensive understanding of complex structures can be obtained.

However, a disadvantage to this method is that physical serial sectioning can be invasive and may result in a loss of integrity in the third dimension due to potential deformations in the object. To address this issue, image registration becomes crucial. Image registration techniques align the 2D serial sections, ensuring accuracy and precision in the 3D reconstruction process. Combining image registration with 2D serial sections can result in highly detailed and reliable 3D reconstructions. This approach enables thorough exploration and analysis of intricate features and characteristics present in various objects and specimens. The resulting 3D reconstructions offer a comprehensive perspective, allowing for the visualization and interaction with the structures from different angles, aiding in more detailed analyses.

## 2.5 Image Segmentation using Deep Learning

Image segmentation is a type of computer vision task that involves partitioning images into multiple segments, usually to isolate certain objects in relation to other parts of the image (Minaee et al., 2022). In essence, the concept of image segmentation can be seen as a pixel classification task. In this formulation, each pixel in an image is treated as an individual sample, and the goal is to classify each pixel into specific classes or categories.

### 2.5.1 Architecture of Image Segmentation Networks

There are many different methods for segmentation using deep learning (Minaee et al., 2022). These methods exhibit diverse architectures that can be based on convolutional neural networks (CNN), fully-connected neural networks (FCNN), recurrent neural networks (RNN), or other specialized structures. While the architecture is not confined to these categories alone, the primary objective remains to capture spatial dependencies and acquire meaningful representations from input images. A popular approach for segmentation networks is the encoder-decoder framework. In this structure, the encoder extracts high-level features and spatial information from the input data, typically through a series of convolutional layers, while the decoder reconstructs the image from the encoded features, performing upsampling operations to restore its original spatial dimensions and refining the segmentation output. This combination enables the network to effectively learn and predict pixel-level segmentation maps.

For a network trained through supervised learning, the network is fed with input images and their corresponding pixel-wise labels (Hesamian, Jia, He, & Kennedy, 2019). The network learns to predict the segmentation map by minimizing a loss function, such as cross-entropy or dice loss, which measures the discrepancy between the predicted segmentation and the ground truth. Cross-entropy is commonly used in image segmentation because it quantifies the dissimilarity between the predicted probability distribution and the true labels, encouraging the network to

assign high probabilities to correct class labels. On the other hand, dice loss is preferred in scenarios where class imbalance is a concern, as it focuses on the intersection between the predicted and true masks, making it particularly effective when dealing with imbalanced foreground and background classes. During inference, the trained segmentation network takes an input image and produces a pixel-level segmentation map, where each pixel is assigned a class label. This map highlights the boundaries and regions of different objects or areas of interest within the image. It should be noted that the training of segmentation networks is not limited to supervised learning. Unsupervised learning techniques have also found application in extracting more dependable information from inadequately labeled data, using the derived annotated data for network training purposes (Anirudh, Thiagarajan, Bremer, & Kim, 2016).

### 2.5.2 Applications of Image Segmentation using Deep Learning

Segmentation can be used in a variety of applications. For example, it can be used in autonomous vehicles to detect pedestrians and roads, for video surveillance to track objects (Ha, Watanabe, Karasawa, Ushiku, & Harada, 2017; Ojha & Sakhare, 2015; Hesamian et al., 2019). In the field of medical image analysis, segmentation plays a crucial role, particularly in delivering essential and accurate information regarding the intricate shapes, volumes, and spatial relationships of organs, cells, and lesions (Hesamian et al., 2019). This may help healthcare professionals with their tasks related to diagnoses, treatments, and monitoring.

A more specific example is how segmentation networks have been used to identify tumors within medical images (Jiang, Diao, & Yao, 2022). With the achievements in the field of deep learning and natural image segmentation came the U-Net architecture (Ronneberger, Fischer, & Brox, 2015). What sets U-Net apart is its unique combination of an encoder-decoder architecture with skip connections. The U-Net architecture, commonly used for image segmentation, employs skip connections to connect corresponding encoder and decoder layers. These connections facilitate the transfer of feature maps from the encoder to the decoder, enabling the conservation of detailed spatial information during the upsampling process. Overall, the success of deep learning has established it as the prevailing method for automated tumor segmentation, setting a standard in this field. (Jiang et al., 2022).

In addition to tumor identification, segmentation networks have also shown promising applications in the detection and analysis of granulomas (Y. Wang et al., 2019; Alilou et al., 2017). Neural networks offer a valuable tool for accurately identifying these inflammatory regions to diagnose and monitor diseases causing granulomas, including tuberculosis. Segmentation networks can capture the unique features and spatial patterns associated with granulomas, enabling automated segmentation. This automated segmentation may help in evaluating the progression of the disease.

### 2.5.3 Performance Measures for Neural Networks

In machine learning, evaluating the accuracy of models, including neural networks, is a crucial step in assessing their effectiveness and reliability, meaning this is important for an image segmentation network as well. To do this, performance measures are essential to understand the performance of the network against the expected outcomes. These performance measures are used during the

training, validation, and testing stages. This research concentrates on two pivotal performance metrics among the many available: the F1 score and the mean squared error (MSE).

The F1 score is a metric commonly used in classification tasks to evaluate the performance of a model. It is the harmonic mean of precision and recall and provides a balanced measure of both metrics (Derczynski, 2016). Precision is the ratio of accurate positive predictions within the set of all positive predictions, whereas recall involves the ratio of correct positive predictions among all existing positive instances. Here true positive predictions refer to the cases where the model correctly identifies a positive instance, while true negatives are where the model correctly identifies a negative instance. The formula can be seen in Equation 1. The value falls between 0 and 1, and a higher F1 score indicates a good balance between precision and recall.

$$\text{F1 score} = \frac{2 * \textit{Precision} * \textit{Recall}}{\textit{Precision} + \textit{Recall}} \quad (1)$$

MSE is an evaluation metric that quantifies the average squared difference between the predicted values and the true target values in a dataset (F. Wang & Zai, 2023). A lower MSE value indicates that the model’s predictions are close to the true values, while a higher value suggests a larger discrepancy between predictions and reality. The formula can be seen in Equation 2, where  $x_i$  is the predicted value and  $y_i$  is the target value. During training, the model aims to minimize the MSE on the training dataset, which measures how well the model fits the training data. However, to evaluate the model’s generalization capability, the MSE is also computed on a separate validation dataset that the model has not seen during training. The difference between training and validation MSE helps identify potential overfitting issues, where the model performs well on the training data but fails to generalize effectively to new, unseen data.

$$\textit{MSE} = \frac{1}{n} \sum_{i=1}^n (x_i - y_i)^2 \quad (2)$$

## 3 Materials & Methods

This section will describe the materials & methods used in this study. It will be explained how the images are prepared for the registration process. It will then clarify how image registration was performed to align multiple images of the zebrafish, and how 3D reconstruction was used to create a 3D image. Finally, it will be explained how the analysis of these images was done using a segmentation neural network to identify and locate granulomas in the zebrafish.

### 3.1 Zebrafish Material

This project used slides containing coronal serial sections of adult zebrafish infected with *Mycobacterium marinum*. Hematoxylin and eosin staining was applied to the majority of slides, while the Ziehl-Neelsen method was used for some. These slides were digitally captured as whole slide images (WSIs) and they were in BigTIFF format, offering ten levels of resolution. The initial dataset comprised 92 zebrafish whole slide images, each containing two serial sections.

### 3.2 Image Preprocessing

To be able to perform registration of the serial sections, the original images containing two serial sections were cropped to isolate individual zebrafish serial sections. The cropping of the BigTIFF images was done using pyvips, a Python binding for libvips (Cupitt & Martinez, 1996). The original images contained resolution levels 0 through 9, where 0 is the highest resolution and 9 is the lowest. The resulting images were PNG images at the resolution level 3. This resolution was chosen as it balances image clarity and the need to maintain relatively small file sizes to keep the computation cost lower. It is worth noting that this decision was empirical and based on practical considerations.

In order to crop the images, the coordinates of the top left corner and the size of the resulting image have to be given as input. Entering the inputs manually was essential to ensure that each resulting image effectively captured the entire fish within its boundaries. This was particularly crucial due to the diverse sizes and compositions of the fish in each image, which presented a challenge in the cropping process. To find the coordinates of the top-left corner in the WSIs, ASAP software was used to visualize the images (Computational Pathology Group, 2018). This made it possible to measure the pixels in the image to find the right coordinates. The size of the cropped images was set to 5500x11500 pixels. The code used for this cropping can be found in the technical document associated with this thesis.

Not all images in the original dataset were used for further registration, as some images contained poor-quality sections that could potentially compromise the accuracy of the registration. To make sure all images used had a proper quality, they were reviewed after being cropped. The images that contained damaged fish or images where the complete fish wasn't captured during image acquisition were removed.

### 3.3 Image Registration

In this research, the alignment of coronal serial sections of zebrafish was performed using VALIS, an automated pipeline for whole slide image registration with rigid and non-rigid transformations (Gatenbee et al., 2021). The choice to perform non-rigid registration was made due to challenges presented by potentially deformed images on the slides and occasional damage. Accurate alignment was crucial for subsequent 3D reconstruction. Non-rigid registration allowed the images with more complex deformations and local irregularities caused by tissue sectioning to be aligned more precisely. Various approaches were explored with different parameters to optimize alignment accuracy and minimize registration errors. Details of these approaches can be found in Sections 3.3.3 and 3.3.4, and further explanation of registration errors is provided in Section 3.3.2.

#### 3.3.1 VALIS: Virtual Alignment of pathoLogY Image Series

In this project, Virtual Alignment of pathoLogY Image Series (VALIS) was used for the registration of the WSIs. VALIS is a fully automated pipeline that automatically registers whole slide images, both rigidly and non-rigidly (Gatenbee et al., 2021). To better understand how VALIS achieves the registration of images, the methodology it uses will be summarized in the remainder of this section. The summary is based on the paper introducing the VALIS software (Gatenbee et al., 2021). A schematic overview of the process can be seen in Figure 3. The code used for the registration is located in the technical document associated with this thesis.

During the automated registration process, VALIS reads the original images in small tiles using Bio-formats and OpenSlide (Linkert et al., 2010; Goode, Gilbert, Harkes, Jukic, & Satyanarayanan, 2013). These tiles are converted to libvips images and combined to rebuild the entire image as a single whole slide libvips image (Cupitt & Martinez, 1996). These images use a “lazy evaluation” that allows the warping of whole slide images without requiring them to be loaded into memory. This feature is beneficial for large images since it efficiently avoids memory limitations that commonly occur when dealing with such larger image files.

Following the conversion of the input images, the slides are preprocessed. A lower-resolution version of the image is used to perform feature detection and to find the displacement fields, representing the spatial shifts within an image. These displacement fields serve as maps that show how individual points in the lower-resolution image have moved or shifted relative to a higher-resolution reference image, helping us understand the spatial transformations occurring within the image. In the context of registration, feature detection refers to the process of identifying and extracting distinctive points or features from images that can be used to align or match corresponding points between multiple images. As the input images in this study were PNGs of a single resolution level, the lower-resolution image is obtained by resizing the whole slide image using libvips. Since the images need to be as similar as possible for registration, VALIS applies a default method that adjusts the hue and colorfulness of each image, making them more comparable. Ensuring image similarity is crucial for the registration process because it allows VALIS to align and analyze corresponding features accurately across multiple images. This is achieved by converting the RGB image to a specific color space, performing transformations to manipulate the hue and colorfulness components, and converting back to RGB. This is important because if the images are significantly dissimilar, challenges can arise, including the misalignment of features and reduced accuracy in spatial analysis.

Once these steps are completed, VALIS performs rigid registration by aligning a large set of images through feature-based algorithms. It uses hierarchical clustering to establish the alignment sequence, ensuring similar images are placed together. If the image order is known, this step can be skipped. During the serial alignment process of VALIS image registration, a reference image is used for alignment, with the central image as the default if no reference image is provided. Common features between neighboring images are detected, and a transformation is applied to align each image with the reference, this process continues iteratively throughout the complete dataset.

Optionally non-rigid registration can be performed after the rigid registration has finished. Performing rigid registration before non-rigid registration ensures that the images are somewhat aligned, which is required for non-rigid registration methods. For this, the images are aligned towards the center image in the series. This process involves constructing displacement fields that adjust the images locally for better alignment. Once the transformation parameters are determined, the full-resolution images are warped using libvips.

The warped images can be saved as OME-TIFF files, which preserve the metadata if the original images contained this. OME-TIFF stands for Open Microscopy Environment Tagged Image File Format. It is a file format specifically designed for storing and exchanging microscopy imaging data. OME-TIFF combines the advantages of the TIFF file format with the ability to store rich metadata associated with the images (Leigh et al., 2016).

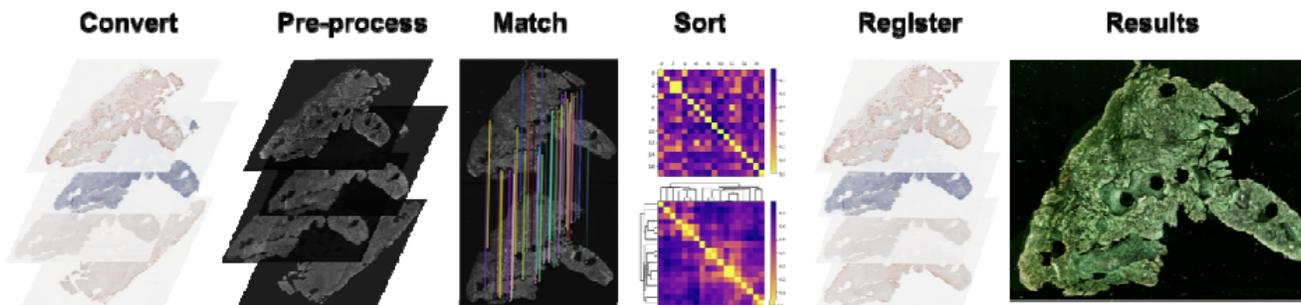


Figure 3: The VALIS alignment pipeline for slide image processing and registration. Slide images are converted, processed, and matched for feature detection. If the order is not known, the images are then ordered based on similarity and registered serially using rigid and optional non-rigid transformations. Processed images are overlaid to visualize registration results. Finally, the aligned slides are warped and saved as OME-TIFF. Sourced from paper introducing the VALIS software (Gatenbee et al., 2021).

### 3.3.2 Registration Error

In the field of image registration there are a few different types of error measures (Fitzpatrick, West, & Maurer, 1998). In the VALIS software one error measure is given to quantify the error of the registration, this measure is called Target Registration Error (TRE). The TRE of a point is the distance between that point and the corresponding point in the other space after registration has been performed (Wodzinski & Müller, 2021). Typically the average of the TRE for all features found in an image during registration is averaged. In order to make the errors comparable when dealing with image pairs of varying resolutions, the TRE is standardized by dividing it by the diagonal of the image, when this is done it will be called rTRE.

$$rTRE = \frac{TRE}{\sqrt{w^2 + h^2}} \quad (3)$$

In Equation 3,  $TRE$  denotes the target registration error, the  $w$  the image width and the  $h$  denotes the image height (Wodzinski & Müller, 2021).

### 3.3.3 Simultaneous Registration of Images

For the first registration attempt, all the selected preprocessed images were used as input for the automated VALIS pipeline. The registration process within VALIS started with a rigid registration step. This step involved applying translations and rotations to align the images based on global geometric changes. Following the rigid registration, a non-rigid registration was performed within the VALIS pipeline. This step aimed to further refine the alignment by allowing for deformable transformations that could account for local variations within the images. The output of the registration process was a set of non-rigidly registered images in the OME.TIFF format. These registered images represented the result of the combined rigid and non-rigid transformations, ensuring the alignment of features and structures within the images. As this experiment did not yield the desired results, different approaches were taken to attempt to improve these results.

### 3.3.4 Group-wise Registration of Images

The second approach aimed to improve the similarity among the images before their simultaneous registration, in an attempt to improve the final result of the registration. To accomplish this, the dataset was divided into seven distinct groups, with each group containing approximately 20-30 images. The exact number of images varied based on the removal of faulty images during the preprocessing stage. Notably, each group, except for the first one, included an additional set of five images. These additional images constituted the last few images within the range that had been registered in the preceding group. The intention of incorporating this overlap was to create a more gradual transition between the registered images of consecutive groups, contributing to a more cohesive and continuous alignment throughout the dataset.

After the registration of the first seven groups, the registered images were divided into two new distinct groups. One of these contained 84 and the other 85 images, these groups underwent individual registration processes. Additionally, the second of these two groups also contained the last 10 registered images from the first group, as the images are still ordered in these groups. This was again done with the intention of creating a more gradual overlap between the groups.

This process was repeated, differing the `max_processed_image_dim_px` parameter, which determines the image size used to find the rigid registration parameters. The choice of using the 850px value was driven by its status as the default parameter within VALIS. Moreover, a higher value of 2000px was deliberately selected to accommodate more intricate image details during the registration parameter calculation process, at the expense of extended runtime. To create the 3D reconstruction, the registration results achieved with a larger value of 2000px were selected since these results had the lowest tTRE value. After completing the grouped registration process, the images obtained were used to do a registration of the total dataset. The images resulting from this final registration can then be used to create the 3D reconstruction of elements within the serial sections, specifically the granulomas.

### 3.4 3D Reconstruction

The 3D reconstruction process involved several steps and specialized software called TDR-3Dbase (Baldoock, Verbeek, & Vonesch, 1997). Since TDR-3Dbase only supports PNG or BMP image formats, the registered serial sections were converted from OME.TIFF to PNG format. This conversion was done using the `libvips` library.

In the TDR-3Dbase system, a digitizer tablet was used to generate contour models of structures from each serial section. These contours were stored in a geometric database, allowing conversion into various geometric representations. The TDR-3Dbase had an input module for contour tasks and a module for alternative geometric manipulations. An input session involved establishing a structural and geometric database. The digitizer tablet was used for contour creation in each section. By stacking these contours, a 3D structure was constructed. Important to note is that a selection of serial sections was made after image cropping. For the 3D reconstruction sections of insufficient quality were substituted with duplicates of the preceding serial section, guaranteeing the preservation of the correct stack height for the 3D reconstruction. This process helped ensure the accuracy of the final reconstructed model.

In this research, a 3D reconstruction was performed to visualize and analyze the granulomas within the zebrafish. To facilitate error correction, individual granulomas were annotated and treated as distinct structures during the annotation process. This approach allowed for targeted and specific adjustments, making the overall error correction more manageable and efficient. Additionally, the ribs, spine, and intestines were reconstructed in 3D to provide spatial reference points for the granulomas. By reconstructing these additional structures in 3D, a fixed frame of reference can be established within the zebrafish's anatomy. This reference point helps in determining the position of the granulomas more precisely. In Figure 4 the interface of the TDR-3Dbase software can be seen. In the serial section currently displayed the granulomas have a red contour drawn around them, and the ribs have a green contour drawn around them.

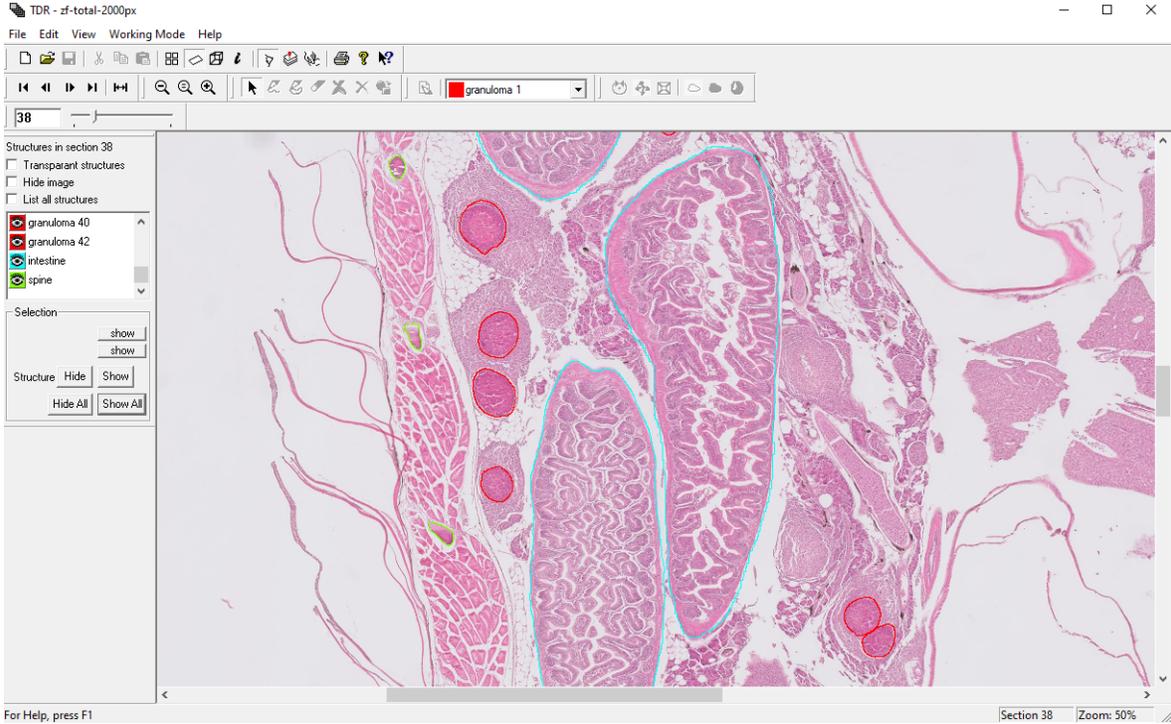


Figure 4: The TDR-3Dbase software interface used for 3D reconstruction of the granulomas, intestines, ribs, and spine of the zebrafish. In this image, one of the serial sections can be seen with red contours drawn around the granulomas, cyan around the intestines, and green around the ribs.

### 3.5 Analysis with Segmentation Network

In our quest to identify granulomas, we employed a deep neural network. Specifically, our choice was the U-Net model, renowned for its segmentation capabilities. The U-net model used features five hierarchical "floors" In the context of the U-Net model, the term "floors" refers to the distinct hierarchical levels within the neural network architecture. The network contains a total of 7,774,049 parameters. Among these parameters, 7,767,779 were trainable, with the remaining 6,848 being non-trainable. The model architecture, depicted in Figure 5, involved a sequence of convolutional layers, batch normalization, and activation functions (Buda, Saha, & Mazurowski, 2019).

The U-Net architecture begins with an input layer, and it consists of blocks that progressively downsample the input image while increasing the number of filters. Each block comprises two 3x3 convolutional layers, followed by batch normalization and Rectified Linear Unit (ReLU) activation functions. The ReLU activation function is a mathematical function used in neural networks that outputs zero for negative input values and passes positive input values directly. This design allows for efficient feature extraction, crucial for accurate segmentation. To refine the segmentation and perform upsampling, transpose convolutional layers (deconvolutional layers) are used. These layers gradually upsample the feature maps and are concatenated with corresponding block outputs from the downsampling path. Each upsampling block consists of a transpose convolutional layer, followed by batch normalization and ReLU activation, progressively restoring the resolution of the feature maps. The upsampling process is repeated four times to achieve precise segmenta-

tion by restoring the resolution. Finally, a 3x3 convolutional layer with a sigmoid activation function is applied to generate the segmentation output. The resulting output is represented as a binary mask, where pixel intensities indicate the likelihood of belonging to the segmented granuloma.

For model training, the Adam optimizer with a learning rate of 0.005 is used, along with a binary cross-entropy loss function, which is well-suited for binary segmentation tasks. In addition to that, during the training process, metrics such as mean squared error (MSE) and accuracy are computed to monitor and evaluate the model’s performance effectively. These metrics serve as valuable indicators of how well the network is learning to identify and segment granulomas in the serial sections.

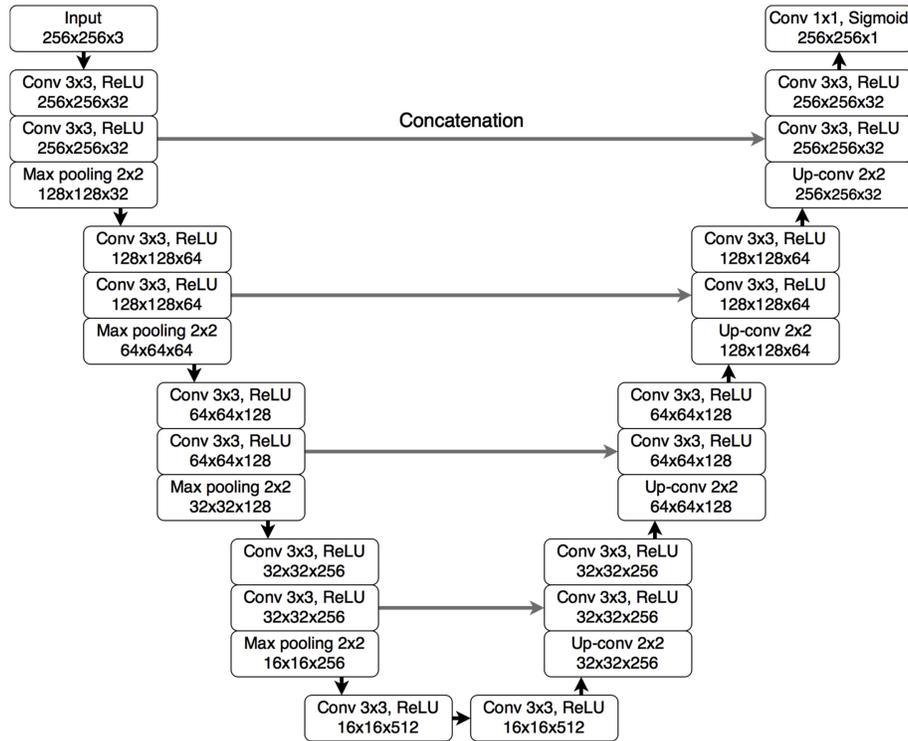
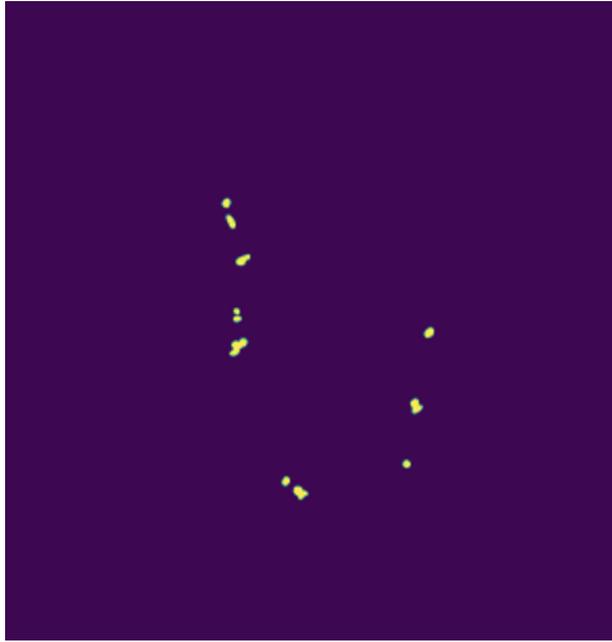
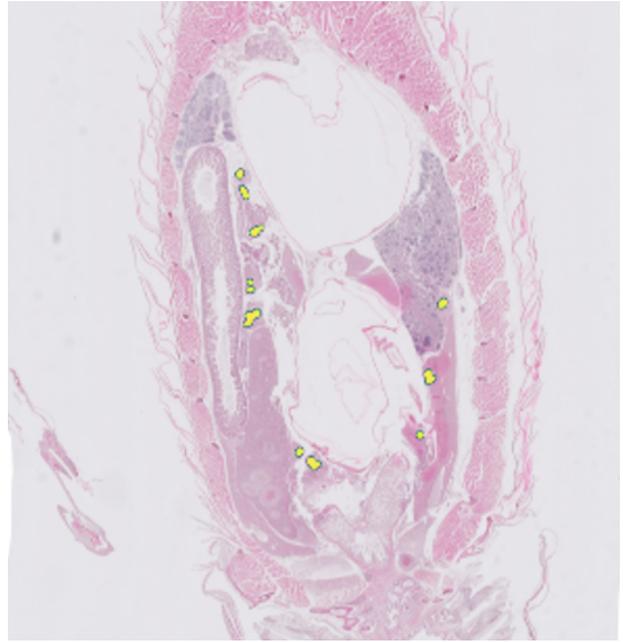


Figure 5: Architecture of U-net model used for segmentation (Buda et al., 2019).

A crucial phase in deploying the neural network was its training process, where it was exposed to annotated serial sections. During the creation of the 3D reconstruction, manual contouring was performed around the granulomas present in each section. These contours served as annotations, providing the network with labeled data for the granulomas. With this annotated dataset, the neural network could learn the distinctive features of granulomas and gain the ability to accurately detect and analyze similar structures in new, unseen data. The network underwent training for 100 epochs, using a dataset consisting of 92 images, which was only half of the available images, this allowed for the other half to be used in validation and testing. Figure 6 provides a depiction of the training input. Figure 6a displays only the labels of the granulomas, while Figure 6b shows the labels overlaid on the original serial section. Once trained, the neural network should be able to recognize granulomas even in unlabeled data.



(a) Labels for granulomas



(b) Labels overlaid on registered serial section

Figure 6: Training input for granuloma detection segmentation network. The neural network was trained using registered serial sections with manual contouring around granulomas, providing annotated data for learning distinctive features. The labels are shown (a), as well as the labels overlaid on the original serial section (b).

## 4 Results

Understanding the intricate nature of zebrafish granulomas within Whole Slide Images (WSIs) is paramount in advancing our comprehension of granuloma biology, driving the development of an efficient and precise workflow presented in this thesis. This thesis introduces a complete workflow for analyzing zebrafish granulomas in WSIs. The workflow includes several key steps, including preprocessing of the WSIs, sequential registration, and 3D reconstruction of granulomas. In addition to this, a segmentation neural network was trained to recognize the granulomas, enabling automated granuloma detection in WSIs. This approach not only enhances understanding of granuloma biology but also establishes a valuable protocol for future research to analyze zebrafish granulomas efficiently and accurately.

### 4.1 Workflow

Figure 7 presents an in-depth depiction of the comprehensive workflow, showcasing the primary outcome of this research. The workflow is structured into four distinct protocols, each designed for specific steps within the process:

1. **Preprocessing**

To isolate individual serial sections, the original images, containing two sections each, were cropped. Subsequently, any images deemed to be of insufficient quality were discarded from the dataset.

2. **Registration**

The cropped images were non-rigidly registered using the VALIS automated pipeline. This process involved registering the images individually within specific groups and subsequently aligning the resulting registered images from each group together.

3. **3D Reconstruction**

After registering the images, a 3D reconstruction was generated using the registered data. For the granulomas, a contour was drawn around the structures present in each serial section. By stacking these annotated contours, a 3D reconstruction of the granulomas was created. This reconstruction could be displayed in various views, one with simply the stacked contours, another with stacked filled contours, and a third with a triangulated continuous surface. Similarly, the same process was applied to the spine, ribs, and intestines, using them as spatial reference points for more precise localization of the granulomas.

4. **Segmentation Network**

Finally, the contours drawn to create the 3D reconstruction of the granulomas were utilized as labels for training a U-net segmentation network. This network is designed to recognize and segment granulomas in unlabeled pictures, enabling automated detection and identification.

In the upcoming sections, a more detailed description of the results achieved through the implementation of this workflow will be presented. The outcome of each individual protocol will be examined, providing insights into the effectiveness and success of the approach.

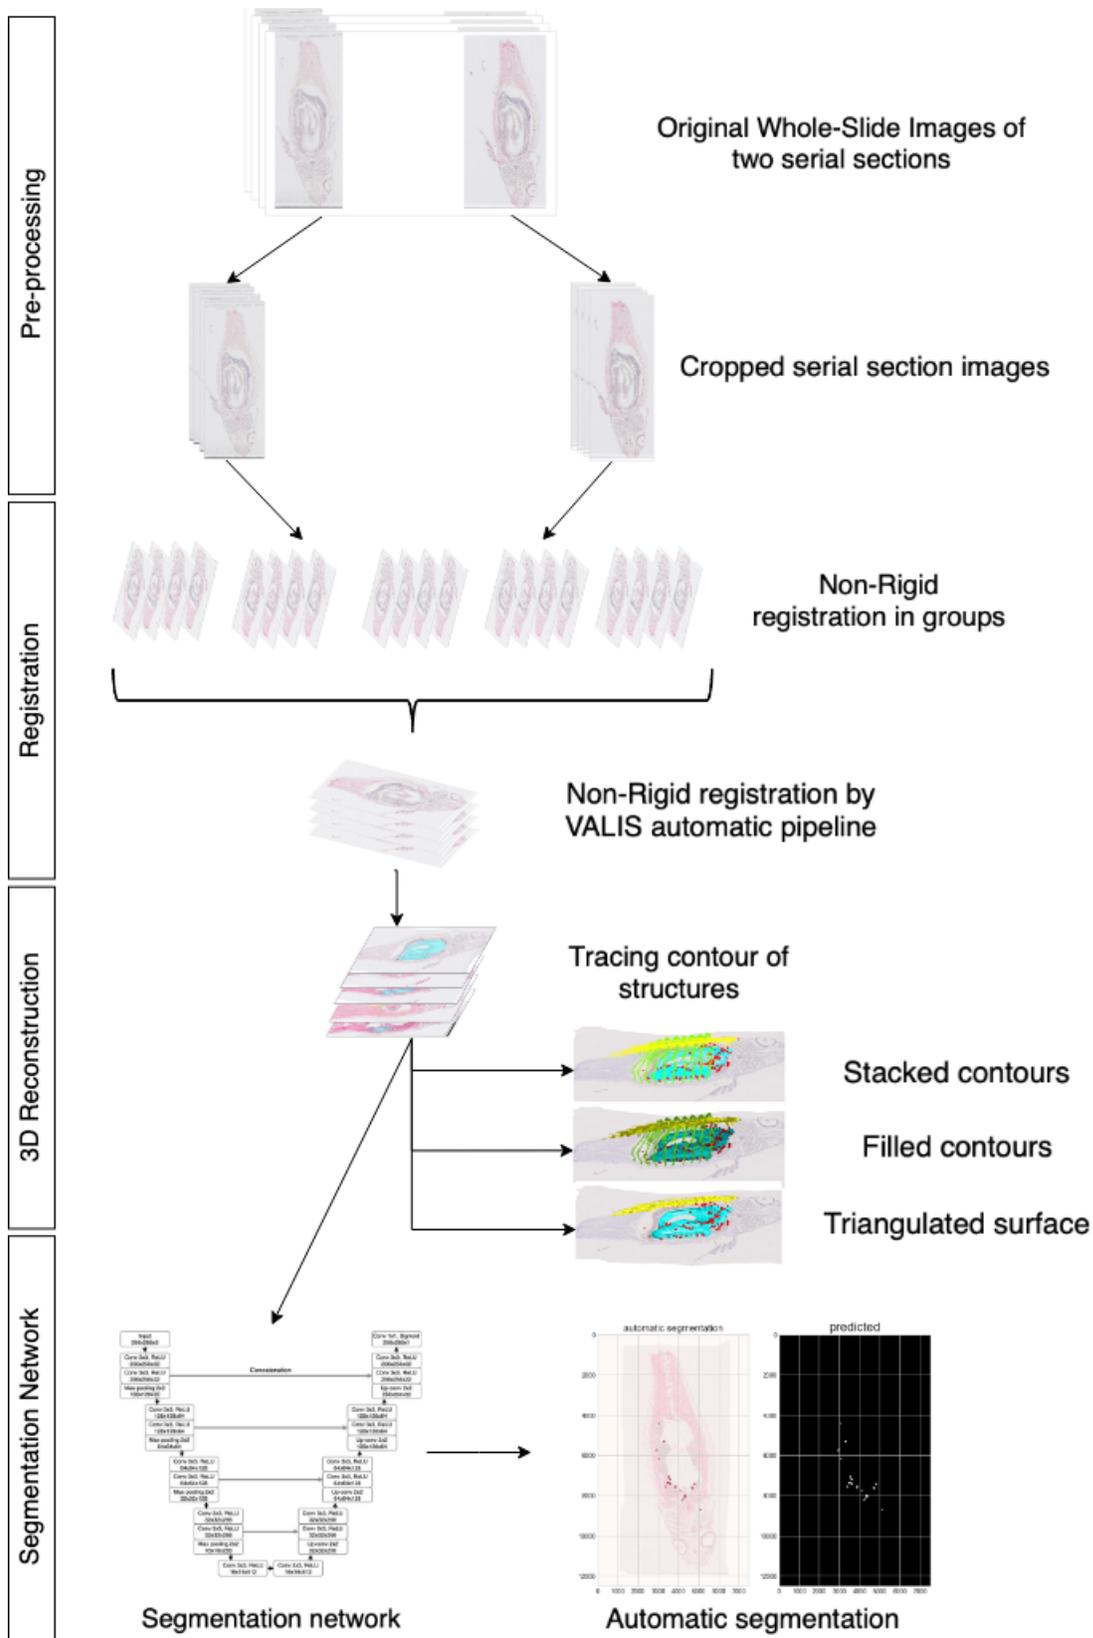


Figure 7: Workflow resulting from this research. The workflow can be divided into four protocols: preprocessing, registration, 3D reconstruction, and the use of a segmentation network.

## 4.2 Cropped Images

The 92 original BigTIFF images containing two serial sections were cropped to produce PNG images of 5500x11500 pixels, each containing one serial section. The images were cropped at level 3 of the resolution pyramid of the BigTIFF images, as it was a good balance between image quality and file size. This resulted in 184 new images, some of which were not used for the registration as some images were found to be damaged, with visible artifacts or distortions. Other images were excluded because the original image did not capture all parts of the fish, which would be challenging for the following steps, specifically the registration process. After this selection procedure, 167 images were left to be used as input for the registration.

## 4.3 Image Registration

The registration process was aimed at addressing the challenge of aligning images obtained from cropping the original images. Various approaches were explored to find an effective method to accurately align the images. This section presents the results of these registration attempts, highlighting the strengths and limitations of each approach, and discussing the overall success in achieving proper alignment.

The distance between features in different images, as well as a numerical error measure, are provided by VALIS to demonstrate alignment accuracy between the two images, this error measure is the rTRE described in Section 3.3.2. All values given by VALIS, including distance and error measures, are rounded to three significant figures after the decimal point within this result section.

### 4.3.1 Simultaneous Registration of Serial Sections

After cropping the images and making a selection of the images that were of sufficient quality, they were collectively given as input into VALIS for non-rigid registration. In Figure 8 it can be seen how the registered images overlap at different stages of the registration process. In Figure 8a the overlap of the original images can be seen before they were registered. In Figure 8b the overlap of the rigidly registered images is depicted, this is an intermediate step for achieving the non-rigid registration. Figure 8c displays the overlap of the non-rigidly registered images, which was the final result of this registration.

However, the experiment encountered an unfortunate setback as this first registration attempt was not successful, as evidenced by the noise present in the overlap images shown in Figures 8b and 8c. Upon closer inspection of the individual registered slides, it could be seen that some of the serial sections were not correctly aligned. This was likely due to the input images being too dissimilar for accurate feature matching. In Table 1 the mean rTRE error measure and the mean distance are displayed along with the standard deviation.

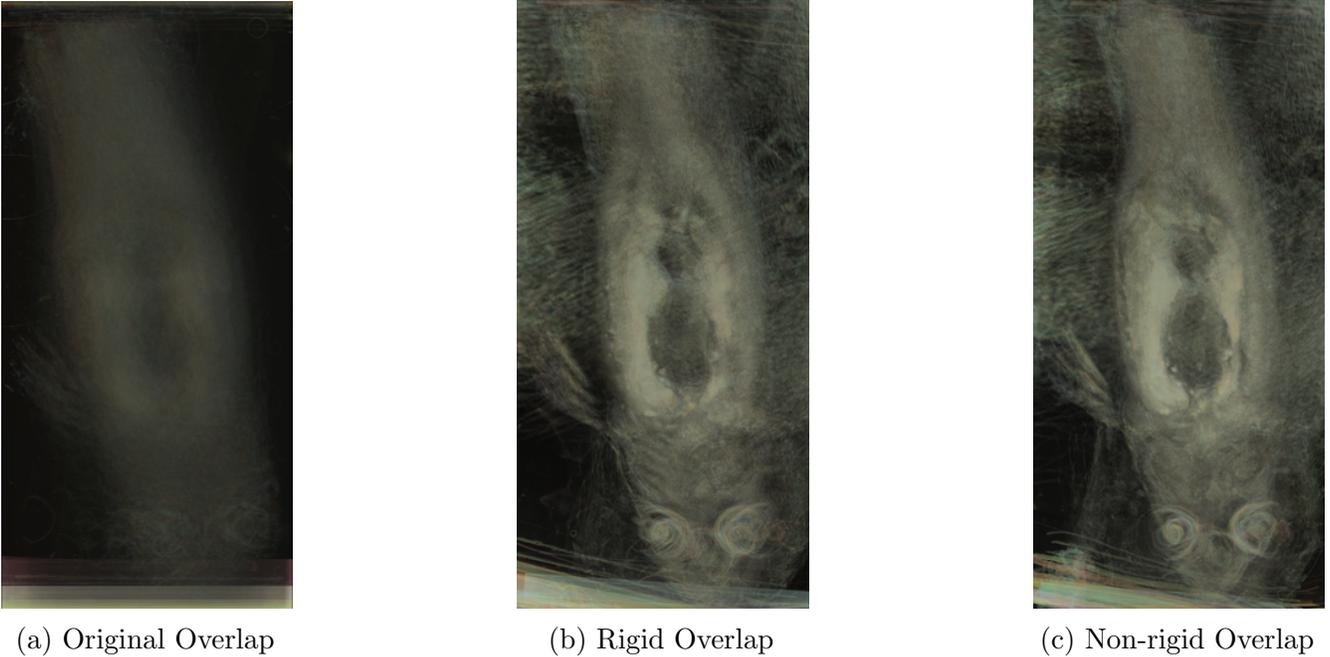


Figure 8: The overlaps of the original images (a), the rigidly registered images (b), and the non-rigidly registered images (c) resulting from the VALIS registration of all images in the dataset simultaneously.

	<b>Original</b>	<b>Non-Rigid</b>
Mean rTRE (SD)	0.607 (0.396)	0.125 (0.568)
Mean Distance (SD)	571.808 (372.772)	139.449 (548.400)

Table 1: Numerical results of simultaneous non-rigid registration for the complete dataset. The mean registration Target Registration Error (rTRE) and distance were computed from VALIS-generated output values, along with their corresponding Standard Deviation (SD).

### 4.3.2 Group-wise Registration of Serial Sections

To overcome the issue of improper registration in the first experiment with all images, a different approach was used, whereby images were divided into smaller groups. Each of these groups was registered separately, which allowed for a more precise alignment by making the images more similar before they were registered as one big group. This process was repeated twice, with only differing the variable that determines the image size used to find the rigid registration parameters.

First, the variable in question was set to its default value of 850px, this is the standard setting in the VALIS software. In Figure 9 the overlaps of the serial sections are shown, again at different stages of the registration process. In Figure 9a it can be seen that the images look similar as they had been registered in groups beforehand. Both the rigid registration in Figure 9b, as well as the non-rigid registration in Figure 9c show that the images have been accurately aligned. Upon closer inspection of the individual registered images, it can be seen that the non-rigid alignment is better than the rigid alignment in terms of aligning the details within the serial sections.

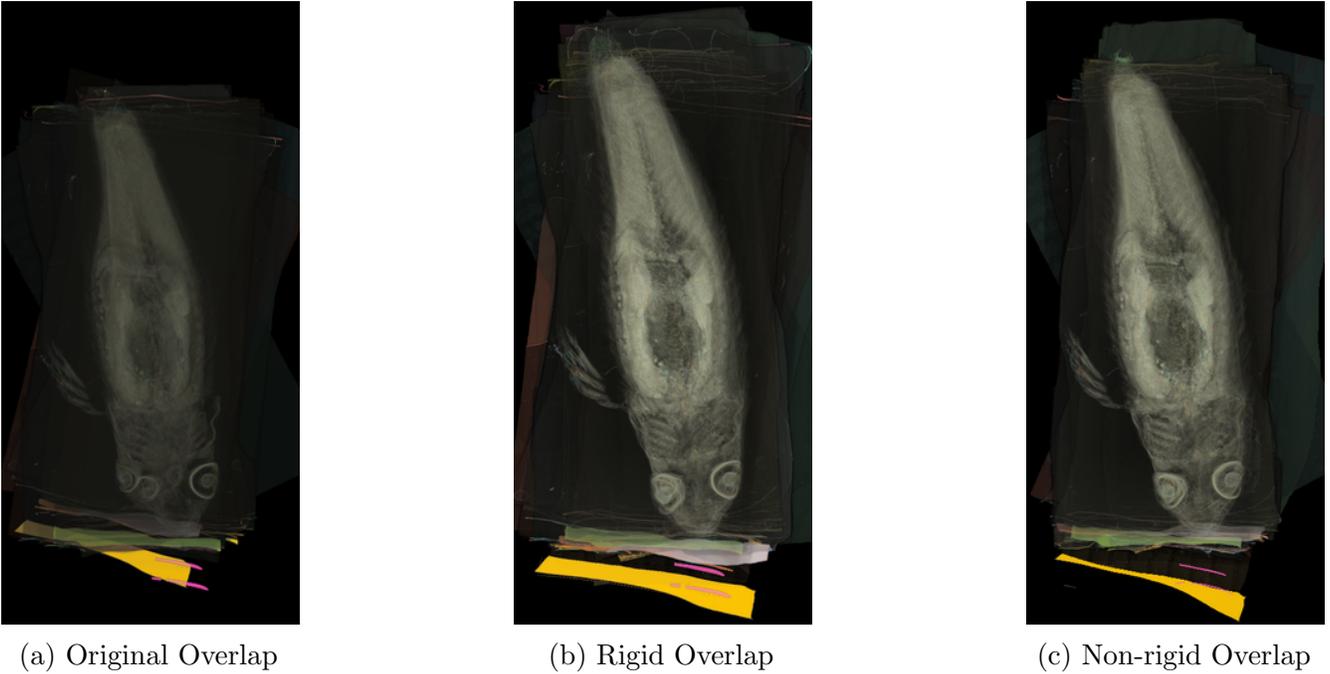


Figure 9: The overlaps of the original images (a), the rigidly registered images (b), and the non-rigidly registered images (c) resulting from the VALIS registration of all images in the dataset after being registered in groups. The variable in VALIS that determines the image size used to find the rigid registration parameters was set to 850px.

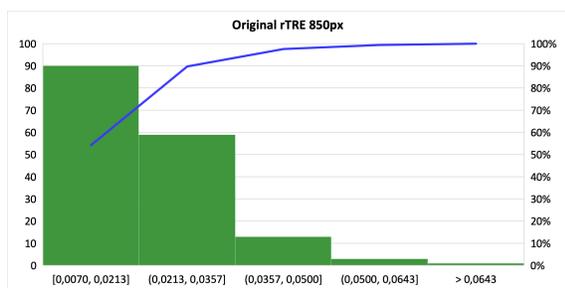
In addition to the visual evidence that the image alignment was successful, the quantitative validation through rTRE values confirms the accuracy of the alignment process. When comparing the data of the previous experiment performing registration of the initial serial sections simultaneously to the values in Table 2, it can be seen that the latter exhibits a notable improvement with lower values. Interesting about Table 2 is that the rTRE increased after registration, indicating larger errors on average. However, the reduced standard deviation suggests a more consistent and predictable performance. Despite the compromise in mean accuracy, the registration method’s enhanced reliability and stability should lead to a more robust and coherent registration.

	<b>Original</b>	<b>Non-Rigid</b>
Mean rTRE (SD)	0.0245 (0.0351)	0.0247 (0.0125)
Mean Distance (SD)	25.512 (36.399)	24.691 (11.361)

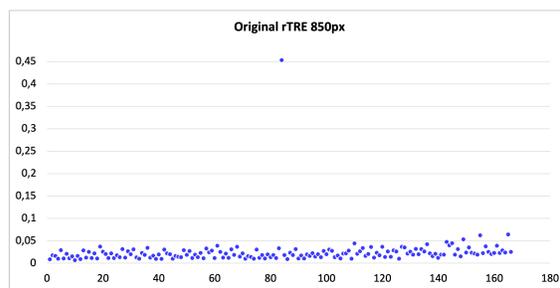
Table 2: Numerical results of group-wise non-rigid registration where the image size used to find the rigid registration parameters was set to 850px. The mean registration Target Registration Error (rTRE) and distance were computed from VALIS-generated output values, along with their corresponding Standard Deviation (SD).

In Figure 10 the distribution of the error values for each serial section can be seen. Figures 10a and 10c show histograms of the rTRE metric, with Pareto lines, before and after registration. This provides a visual representation of the rTRE intervals for the serial sections. Evidently, the majority of serial sections, both before and after the non-rigid registration, exhibit low rTRE values. However, before the registration the error was overall lower.

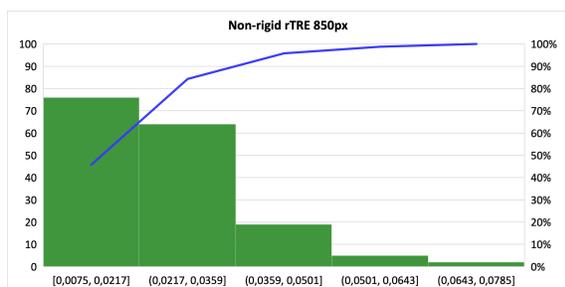
Figures 10b and 10d display scatterplots illustrating the rTRE values corresponding to each serial section before and after the registration. Figure 10b showcases a notable outlier where, before the registration process, one of the serial sections displayed a significantly higher error compared to the rest. After the registration, the error for this particular serial section has significantly reduced, bringing it much closer to the average error. It seems while this outlier value was improved, there is a bigger variance within the main group of data points. Nevertheless, upon visual inspection of the dataset, it became evident that the dataset was better aligned after the registration process compared to the original input. Despite the presence of the outlier mentioned earlier, the overall improvement in alignment was noticeable, indicating the effectiveness of the non-rigid registration.



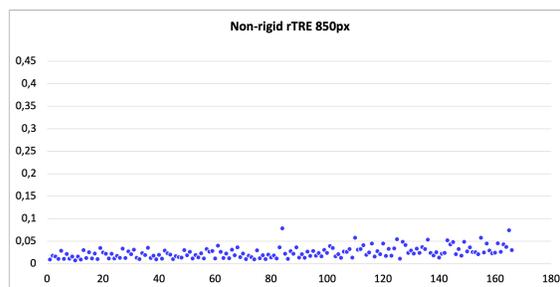
(a) Original rTRE 850px histogram with Pareto line



(b) Original rTRE 850px scatterplot



(c) Non-rigid rTRE 850px histogram with Pareto line



(d) Non-rigid rTRE 850px scatterplot

Figure 10: Visualizations of rTRE values before and after non-rigid registration with the image size for finding rigid registration set to 850px. Histogram with a Pareto line is shown before registration (a) and after registration (c). Additionally, scatter plots for before registration (b) and after registration (d) are shown.

The second time this group registration was done, the variable in VALIS that determines the image size used to find the rigid registration parameters was set to 2000px, this serves as the recommended upper limit, as specified in the VALIS documentation (Gatenbee, 2023). The larger image size enables better visibility of details during the registration process, at the expense of increased computation time.

In Figure 11 the overlaps for this experiment can be seen during different stages of the registration. This figure is quite similar to Figure 9, as the input data was already somewhat aligned as can be seen in 11a. Additionally, Figures 11b and 11c show a proper alignment of the serial sections similar to the rigid and non-rigid overlap images that can be seen in Figures 9b and 9c. Visual assessment alone was insufficient to determine the superior alignment between the two, which is why the rTRE values played an important role in making a more informed comparison.



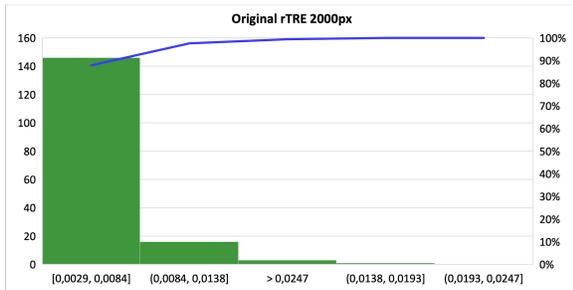
Figure 11: The overlaps of the original images (a), the rigidly registered images (b), and the non-rigidly registered images (c) resulting from the VALIS registration of all images in the dataset after being registered in groups. The variable in VALIS that determines the image size used to find the rigid registration parameters was set to 2000px.

In Table 3, we observe lower rTRE and distance values in comparison to Tables 1 and 2. Notably, the non-rigid registration process resulted in a reduction in the mean rTRE. However, it is essential to consider the observed increase in the standard deviation, which suggests a slightly higher variability in registration errors. Nevertheless, these results suggest improvement in registration by setting the parameter that decides the image size for finding rigid registration parameters to 2000px, as it positively influenced the non-rigid registration of the serial sections.

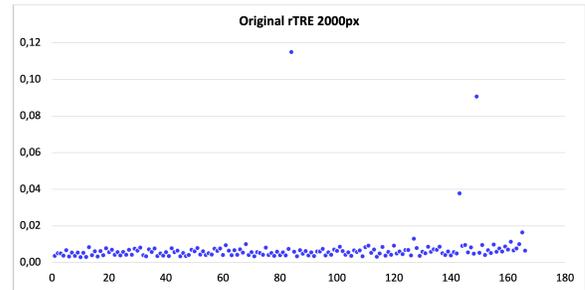
	<b>Original</b>	<b>Non-Rigid</b>
Mean rTRE (SD)	0.00723 (0.0111)	0.00678 (0.00330)
Mean Distance (SD)	17.773 (27.449)	16.826 (8.516)

Table 3: Numerical results of group-wise non-rigid registration where the image size used to find the rigid registration parameters was set to 2000px. The mean registration Target Registration Error (rTRE) and distance were computed from VALIS-generated output values, along with their corresponding Standard Deviation (SD).

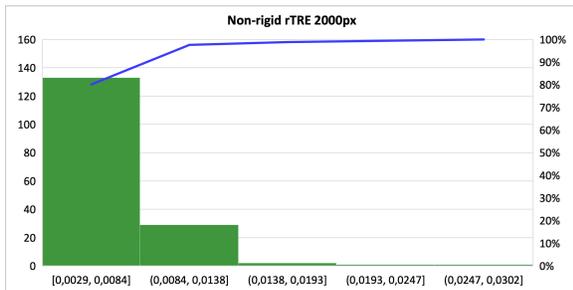
The distribution of rTRE values was visualized using histograms with Pareto lines and scatter plots before and after the registration process as can be seen in Figure 12. In Figures 12a and 12c the histograms can be seen, in these figures it seems that the majority of the serial section had a low error, even more so than in the previous experiment. Figures 12b and 12d demonstrate a noticeable reduction in outlier error values after the non-rigid registration process. Additionally, these values exhibit considerably lower variance compared to the previous experiment, as evidenced by the more compact distribution. It is important to note the difference in the y-axis range in this context.



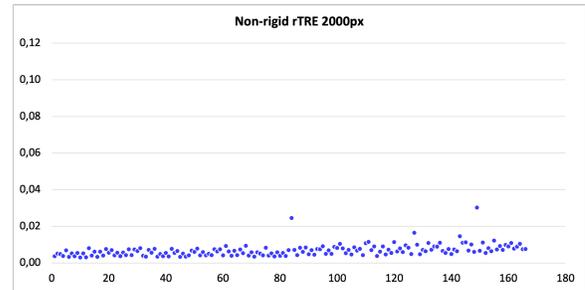
(a) Original rTRE 2000px histogram with Pareto line



(b) Original rTRE 2000px scatterplot



(c) Non-rigid rTRE 2000px histogram with Pareto line



(d) Non-rigid rTRE 2000px scatterplot

Figure 12: Visualizations of rTRE values before and after non-rigid registration with the image size for finding rigid registration set to 2000px. Histogram with a Pareto line is shown before registration (a) and after registration (c). Additionally, scatter plots for before registration (b) and after registration (d) are shown.

## 4.4 3D Reconstruction

During the 3D reconstruction process, the set of registered images with the smallest average rTRE was utilized. Granulomas were reconstructed using the TDR-3DBase software. The ribs, spine, and intestines, were also reconstructed to act as spatial reference points, this was to provide insights into the distribution and localization of the granulomas relative to surrounding structures in the zebrafish. In Figure 13 an example can be seen of how a serial section is annotated. This specific slide shows the annotation of granulomas in red, intestines in cyan, and ribs in green.

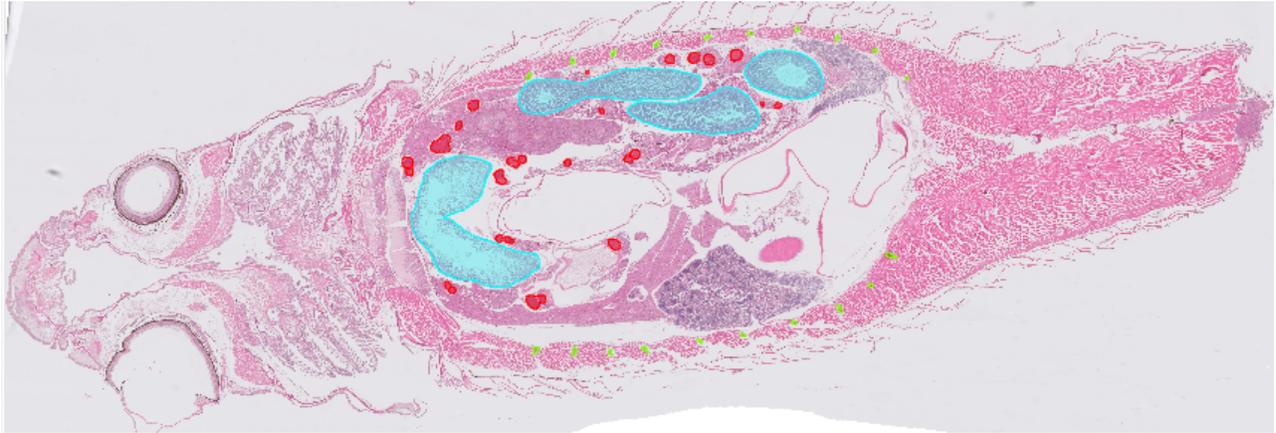
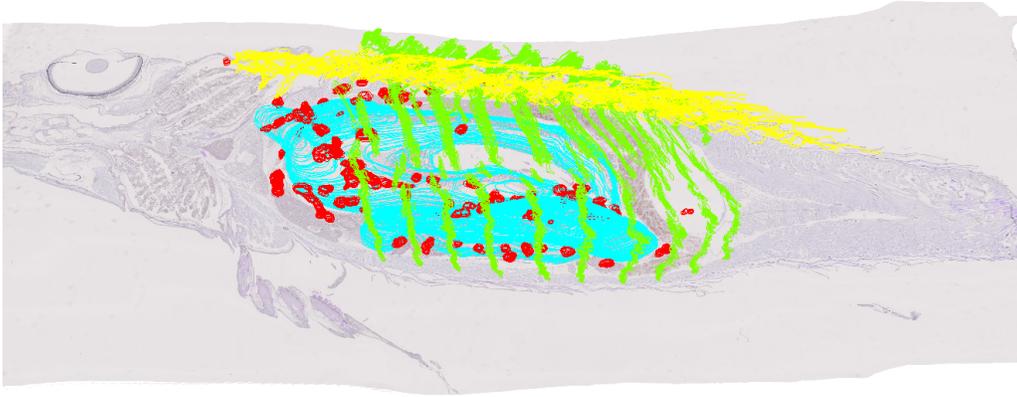


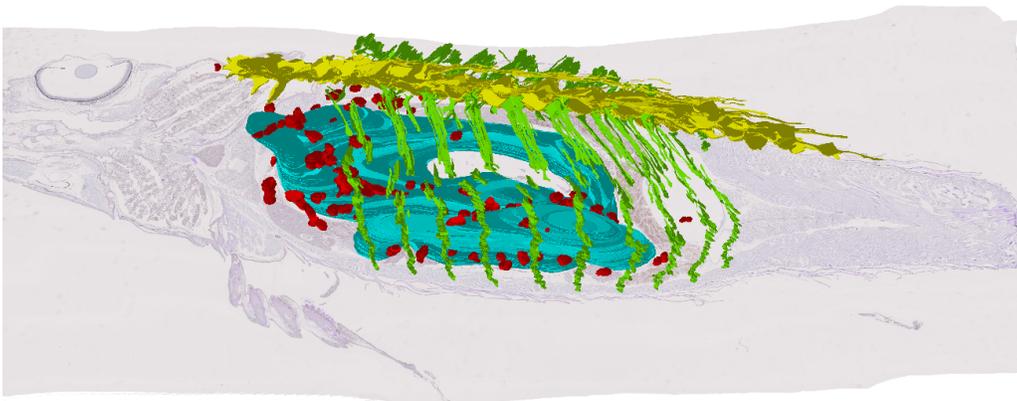
Figure 13: Example of an annotated serial section. Granulomas are highlighted in red, ribs in green, and intestines in cyan.

The reconstruction results are presented in Figure 14 as stacked drawn contours and in Figure 14b as filled planes. The structures were triangulated using an algorithm included in TDR-3DBase, creating continuous structures as shown in Figure 14c. TDR-3DBase software allows convenient 3D reconstructions viewing from different angles and displaying distinct serial sections for analysis. The 3D reconstruction revealed significant granuloma presence around the intestines, with no granulomas found within them. The liver displayed a concentrated number of granulomas, appearing in both solitary and clustered configurations. Clusters are defined as two or more granulomas directly adjacent to each other. Most clusters consisted of exactly two adjacent granulomas, but larger clusters with up to six adjacent granulomas were also observed. Slightly over half of the granulomas appeared as individual formations, while slightly less than half were detected in clustered configurations.

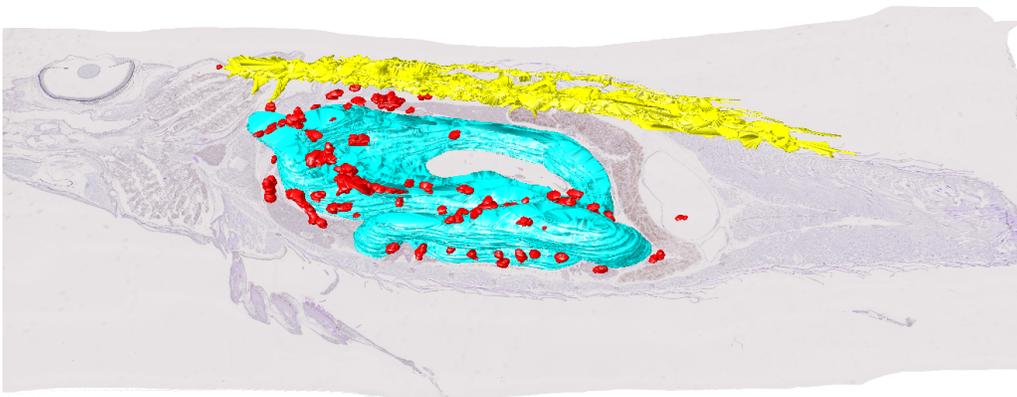
Another intriguing observation was the notable asymmetry in the distribution of granulomas within the adult zebrafish model. A higher occurrence of granulomas was observed on one side of the zebrafish compared to the other side. These findings provide insights into the spatial patterns and clustering tendencies of granulomas in the zebrafish model, contributing to a better understanding of the disease's spatial distribution.



(a) Stack of contours in three dimensions.



(b) Stack of filled planes in three dimensions.



(c) Objects displayed as triangulated surfaces.

Figure 14: Different forms of displaying the structures as geometric objects overlaid on a serial section image: The objects displayed as contours (a), the objects displayed as filled planes (b), and the objects displayed as triangulated continuous surfaces (c). The granulomas are colored red, the spine yellow, the ribs green and the intestines are cyan.

## 4.5 Analysis with Segmentation Network

The U-net model was trained to segment granulomas using labels derived from the contours drawn during the creation of the 3D reconstruction over 100 epochs. An example of a prediction performed by the network can be seen in Figure 15. The leftmost panel displays the original serial section. The next panel shows a mask representing manually drawn labels for the granulomas. Continuing to the right, the results of the automatic segmentation process are displayed overlaid on the serial section. Finally, the rightmost panel exhibits a mask displaying the predicted locations of the granulomas.

Across various serial sections, the network’s predictions closely match the manually drawn granuloma labels. The network also displayed the ability to discover previously unidentified potential granulomas, thus expanding its detection capabilities. However, the network did have problems in detecting very small annotated granulomas, indicating some limitations in detecting extremely small structures.

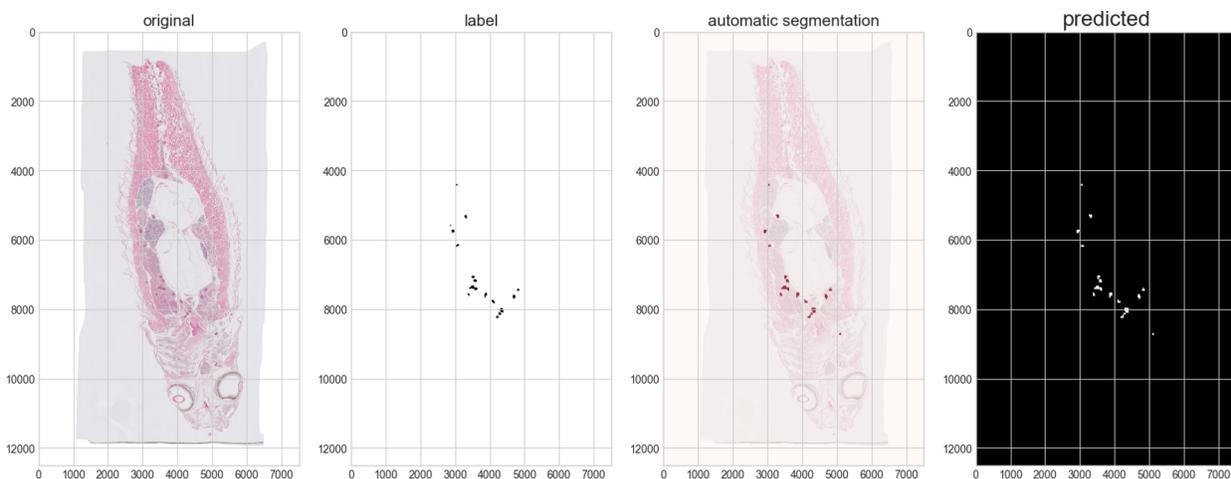
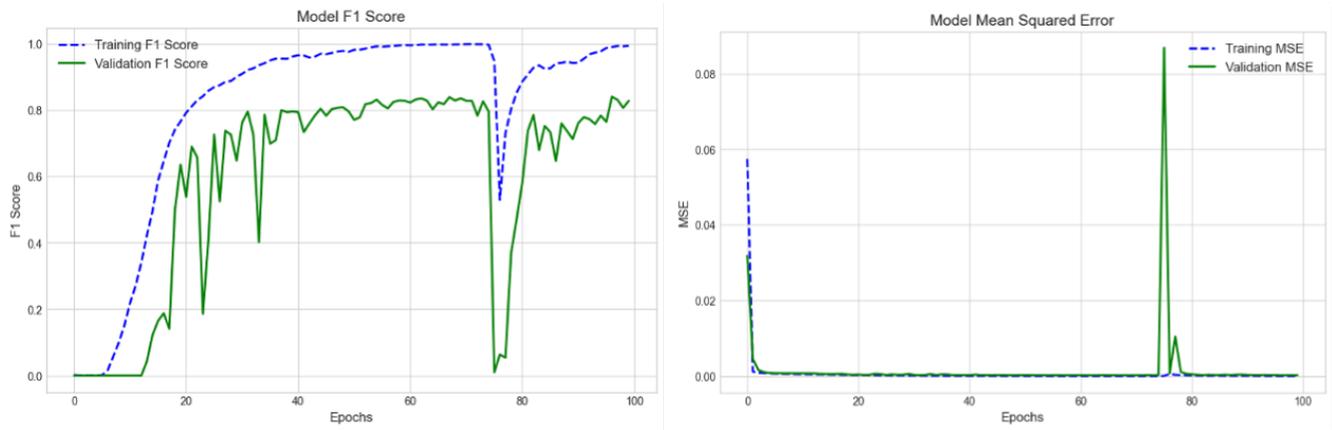


Figure 15: The original serial section, mask derived from manually drawn granuloma labels, a prediction overlaid on the serial section, and the corresponding mask revealing the predicted granuloma areas (left to right).

Figure 16 provides additional insights through quantitative results. In Figure 16a the F1 score over the 100 epochs is displayed. The training F1 score measures the accuracy of the model on the training dataset, while the validation F1 score evaluates the ability of the model to generalize unseen data on a separate validation dataset. The training line exhibits a consistent upward trend, indicating that the model progressively improves its performance and approaches an F1-score of 1, signifying high precision and recall balance. However, at approximately 75 epochs, the training score suddenly drops to around 0.5, suggesting a potential issue or obstacle in the learning process. This could be due to images of lower quality being used in training. Despite this, the training line manages to recover and continues climbing. Further analysis might be required to identify specific causes, such as the possibility of introducing a problematic training sample or potential model instability. On the other hand, the validation line also shows improvement over epochs but is more susceptible to fluctuations, indicating a less stable performance. At around 75 epochs, the validation score nearly reaches zero, revealing a severe decline in performance. Nevertheless, the validation score rebounds and settles slightly above 0.8, suggesting the model’s ability to generalize reasonably well on unseen data.

In Figure 16b the Mean Squared Error (MSE) over the 100 epochs is displayed. Initially, the training line quickly decreases to zero and stays there, indicating that the model rapidly learns to fit the training data. The validation line shows a different behavior, shooting up to a value above 0.08 around 75 epochs, indicating that the model is struggling to generalize to unseen data, leading to overfitting. After this initial peak, there is another much shorter peak, suggesting that the model is still trying to adapt to more complex patterns in the validation data. However, after these fluctuations, the validation line stabilizes and stays close to zero for the remainder of the training, which is a positive sign, again indicating that the model has achieved reasonable generalization and performs well on unseen data.



(a) F1 score

(b) Mean Squared Error

Figure 16: Graphs showing two different metrics of measuring the network accuracy. The F1 score is shown (a), as well as the Mean Squared Error (b).

## 5 Conclusion & Discussion

In this concluding section, the key insights gained from the research will be summarized. Deeper insights as well as the challenges encountered throughout the study will be addressed within the discussion section. Lastly, suggestions for future research will be provided.

### 5.1 Conclusion

In this research, the main objective was to investigate the feasibility of performing a reliable analysis of the distribution of granulomas in a 3D reconstruction of an adult zebrafish. Three sub-questions were formulated to guide the investigation. These questions will be answered in this section, each addressing a specific aspect of the overall research question.

1. **Can a reliable registration be performed on WSIs of adult zebrafish using VALIS software?**

Multiple attempts were made to discover an accurate non-rigid registration method for WSIs. The most reliable approach involved dividing the WSIs into groups and registering them individually within each group before registering all WSIs together. Setting the image size to 2000px in VALIS for finding rigid registration parameters yielded the best performance. These findings highlight the capability of the VALIS software to achieve a sufficiently reliable registration of adult zebrafish WSIs, providing a foundation for subsequent analyses of zebrafish tissue sections.

2. **Can a proper 3D reconstruction of granulomas be made from the registered WSIs of adult zebrafish?**

The results show that a proper 3D reconstruction of granulomas can be successfully accomplished from the registered WSIs of adult zebrafish. In addition to this the ribs, spine, and intestines were also reconstructed to use as spatial reference points. From this 3D reconstruction, it was observed that granulomas occur in both singles as well as clusters and occur mostly within the liver.

3. **Can a segmentation network be used to identify the granulomas in an adult zebrafish?**

The study demonstrated the efficacy of a segmentation network in identifying granulomas in adult zebrafish. The network exhibited a promising level of accuracy, offering an automated approach for granuloma detection.

The answers to the three sub-questions confirm that a reliable analysis can be performed on the distribution of granulomas in a 3D reconstruction of an adult zebrafish, supporting the main research question. This is a step forward in understanding granuloma distribution and offers valuable insights into the spatial organization of granulomas, contributing to the field of zebrafish analysis. Moreover, these findings have the potential to lead to new strategies for future investigations and applications within the broader scientific community, promising advancements in biomedical imaging and analysis as they add to the existing knowledge in zebrafish research.

## 5.2 Discussion

As the conclusion stated, the research has created a reliable workflow to create a 3D reconstruction from WSIs of serial sections of zebrafish. Nevertheless, each of the individual steps in the workflow presented its own challenges. The difficulties encountered during the research will be further discussed in this section.

There were a few difficulties during the registration of the cropped serial sections. While this step resulted in serial sections that were aligned properly enough for the purpose of this study, the alignment was certainly not perfect. Especially the heads and tails of the fish occasionally contained unnatural deformations after the registration process. In addition to this, smaller features occasionally suffered from misalignment, likely attributed to the damage on the original zebrafish slides. Another setback was that VALIS lacked the capability to register all serial sections simultaneously. To properly align the images, they had to be registered in groups first before they could be registered simultaneously. Moreover, it appeared that the VALIS output of the rTRE exhibited some form of bug or anomaly. As the rTRE value was given as the same for both rigid as well as non-rigid registration while the distance values were different. As the rigid registration is an intermediate step in getting the non-rigid registration this measure was still useful, as something was needed to decide the best registration. However, this did introduce a vulnerability regarding the selection of the best registration.

Despite the difficulties in achieving a good registration, the structures were still aligned well enough to result in a proper 3D reconstruction. The most issues were found with creating a 3D reconstruction of the ribs and spine, as these structures were already damaged on the original slides, which resulted in an alignment that was not entirely accurate for some slides and thus also had negative effects on the 3D reconstruction. The granulomas and intestines seemed to be aligned better and thus resulted in more accurate 3D reconstructions.

Thus far the segmentation network has shown promising results during the training and validation, however the network has not been tested yet. Interesting about the neural network was that it also showed to be able to already predict the location of granulomas that were missed during the annotation process. A limitation of the network so far seemed to be that it had trouble finding very small granulomas in some cases.

A significant issue encountered throughout the workflow was the utilization of excessively large images. These oversized images posed a considerable challenge, pushing the limits of the processing capabilities of the computer that was used. Consequently, certain steps in the workflow suffered from longer processing times and sometimes required a lot of computer memory.

Overall, the workflow offers a way in which a 3D reconstruction can be created from WSIs of serial sections in order to analyze three-dimensional structures. But other than that the workflow could be generalized to analyze other structures besides granulomas as well. However, there were still some parts that could be improved with further research.

### 5.3 Future Research

The workflow created in this research has shown to be effective, despite its limitations. However, to enhance the reliability and accuracy of the findings, future research endeavors could concentrate on overcoming the limitations of the study. In this section, some suggestions for further research will be discussed.

Firstly, in the registration process, different ways of grouping the WSIs can be explored in order to improve registration accuracy. Implementing micro-registration techniques may also prove beneficial in achieving more precise alignments, this is an available option in VALIS where it performs a second non-rigid registration to align detailed features better. This micro-registration does require a lot of computer memory. A possible solution is to attempt to allocate this process to the GPU instead of the CPU. As GPUs are designed for parallel processing, and tasks that can be parallelized, such as image processing.

Furthermore, for 3D reconstruction, additional structure could be incorporated into the current model to create a more comprehensive representation. This could lead to a more detailed and informative reconstruction. For example, if one would create a 3D reconstruction of the liver it would be easier to see how many of the granulomas are located within the liver.

Lastly, in the segmentation network, using a more complex neural network with transformer architecture could improve the accuracy of the segmentation process. It could potentially also be trained on different structures within the fish. However, for future research, it should be kept in mind that due to the registration, the granulomas are mostly located in very similar locations within the image. It might be useful to check if the network can also recognize granulomas if they are in different locations to make sure the network does not base its prediction on the location of the structures. Eventually, the segmentation of structures may even be used to generate contours which can then, in turn, be used to create 3D reconstructions, if the network is accurate enough, it could save a lot of time compared to manual annotation. By addressing these points, future studies can make steps in refining the analysis process and potentially improve the generalizability of the process.

## 6 Acknowledgements

To end my thesis, I want to express my gratitude to everyone who played a role in the completion of this work. Firstly, I extend my thanks to my supervisor, Fons Verbeek, for his guidance and insightful feedback throughout the research process.

I would like to thank Shima Javanmardi for her work on the segmentation network. Her contributions allowed me to showcase more of what was possible with the results of the 3D reconstruction. The results from the segmentation network allow for new possibilities in future research.

Finally, I would like to thank Jia Li for her assistance in helping me get started with Pyvips and VALIS, which facilitated my initial experimentation with the capabilities of this software.

## References

- Alilou, M., Beig, N., Orooji, M., Rajiah, P., Velcheti, V., Rakshit, S., . . . others (2017). An integrated segmentation and shape-based classification scheme for distinguishing adenocarcinomas from granulomas on lung ct. *Medical physics*, 44(7), 3556–3569.
- Anirudh, R., Thiagarajan, J. J., Bremer, T., & Kim, H. (2016). Lung nodule detection using 3d convolutional neural networks trained on weakly labeled data. In *Medical imaging 2016: Computer-aided diagnosis* (Vol. 9785, pp. 791–796).
- Aware Systems. (2008). *Bigtiff, tiff breaking the 4 gigabyte boundary*. <https://www.awaresystems.be/imaging/tiff/bigtiff.html>.
- Baldock, R., Verbeek, F., & Vonesch, J. (1997). 3-d reconstructions for graphical databases of gene expression. In *Seminars in cell & developmental biology* (Vol. 8, pp. 499–507).
- Bouz, G., & Al Hasawi, N. (2018). The zebrafish model of tuberculosis-no lungs needed. *Critical reviews in microbiology*, 44(6), 779–792.
- Buda, M., Saha, A., & Mazurowski, M. A. (2019). Association of genomic subtypes of lower-grade gliomas with shape features automatically extracted by a deep learning algorithm. *Computers in Biology and Medicine*, 109, 218–225. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0010482519301520> doi: <https://doi.org/10.1016/j.compbimed.2019.05.002>
- Cheng, T., Kam, J. Y., Johansen, M. D., & Oehlers, S. H. (2020). High content analysis of granuloma histology and neutrophilic inflammation in adult zebrafish infected with mycobacterium marinum. *Micron*, 129, 102782.
- Clay, H., Davis, J. M., Beery, D., Huttenlocher, A., Lyons, S. E., & Ramakrishnan, L. (2007). Dichotomous role of the macrophage in early mycobacterium marinum infection of the zebrafish. *Cell host & microbe*, 2(1), 29–39.
- Computational Pathology Group. (2018). *ASAP: Automated slide analysis platform*. <https://computationalpathologygroup.github.io/ASAP/#contact>.
- Cupitt, J., & Martinez, K. (1996). Vips: an image processing system for large images. In *Very high resolution and quality imaging (01/01/96)* (pp. 19–28). Retrieved from <https://eprints.soton.ac.uk/252227/>
- Derczynski, L. (2016, May). Complementarity, F-score, and NLP evaluation. In *Proceedings of the tenth international conference on language resources and evaluation (LREC'16)* (pp. 261–266). Portorož, Slovenia: European Language Resources Association (ELRA). Retrieved from <https://aclanthology.org/L16-1040>
- Farahani, N., Parwani, A. V., Pantanowitz, L., et al. (2015). Whole slide imaging in pathology: advantages, limitations, and emerging perspectives. *Pathol Lab Med Int*, 7(23-33), 4321.
- Fitzpatrick, J. M., West, J. B., & Maurer, C. R. (1998). Predicting error in rigid-body point-based registration. *IEEE transactions on medical imaging*, 17(5), 694–702.
- Gatenbee, C. D. (2023). *Examples*. Retrieved from <https://valis.readthedocs.io/en/latest/examples.html>
- Gatenbee, C. D., Baker, A.-M., Prabhakaran, S., Slebos, R. J., Mandal, G., Mulholland, E., . . . others (2021). Valis: Virtual alignment of pathology image series. *bioRxiv*, 2021–11.
- Goode, A., Gilbert, B., Harkes, J., Jukic, D., & Satyanarayanan, M. (2013). Openslide: A vendor-neutral software foundation for digital pathology. *Journal of pathology informatics*, 4(1), 27.

- Ha, Q., Watanabe, K., Karasawa, T., Ushiku, Y., & Harada, T. (2017). Mfnet: Towards real-time semantic segmentation for autonomous vehicles with multi-spectral scenes. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (p. 5108-5115). doi: 10.1109/IROS.2017.8206396
- Hesamian, M. H., Jia, W., He, X., & Kennedy, P. (2019). Deep learning techniques for medical image segmentation: achievements and challenges. *Journal of digital imaging*, *32*, 582–596.
- Jiang, H., Diao, Z., & Yao, Y.-D. (2022). Deep learning techniques for tumor segmentation: a review. *The Journal of Supercomputing*, *78*(2), 1807–1851.
- Leigh, R., Gault, D., Linkert, M., Burel, J.-M., Moore, J., Besson, S., & Swedlow, J. R. (2016). Ome files—an open source reference library for the ome-xml metadata model and the ome-tiff file format. *BioRxiv*, 088740.
- Linkert, M., Rueden, C. T., Allan, C., Burel, J.-M., Moore, W., Patterson, A., ... others (2010). Metadata matters: access to image data in the real world. *Journal of Cell Biology*, *189*(5), 777–782.
- Meijer, A. H., & Spaink, H. P. (2011). Host-pathogen interactions made transparent with the zebrafish model. *Current drug targets*, *12*(7), 1000–1017.
- Minaee, S., Boykov, Y., Porikli, F., Plaza, A., Kehtarnavaz, N., & Terzopoulos, D. (2022). Image segmentation using deep learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *44*(7), 3523-3542. doi: 10.1109/TPAMI.2021.3059968
- Myllymäki, H., Bäuerlein, C. A., & Rämets, M. (2016). The zebrafish breathes new life into the study of tuberculosis. *Frontiers in immunology*, *7*, 196.
- Neely, M. N., Pfeifer, J. D., & Caparon, M. (2002). Streptococcus-zebrafish model of bacterial pathogenesis. *Infection and immunity*, *70*(7), 3904–3914.
- Ojha, S., & Sakhare, S. (2015). Image processing techniques for object tracking in video surveillance—a survey. In *2015 International Conference on Pervasive Computing (ICPC)* (p. 1-6). doi: 10.1109/PERVASIVE.2015.7087180
- Ordas, A., Raterink, R.-J., Cunningham, F., Jansen, H. J., Wiweger, M. I., Jong-Raadsen, S., ... others (2015). Testing tuberculosis drug efficacy in a zebrafish high-throughput translational medicine screen. *Antimicrobial agents and chemotherapy*, *59*(2), 753–762.
- Pantanowitz, L., Valenstein, P. N., Evans, A. J., Kaplan, K. J., Pfeifer, J. D., Wilbur, D. C., ... Colgan, T. J. (2011). Review of the current state of whole slide imaging in pathology. *Journal of pathology informatics*, *2*(1), 36.
- Pillay, R. (2023). *IIPImage Images* — [iipimage.sourceforge.io](https://iipimage.sourceforge.io). <https://iipimage.sourceforge.io/documentation/images>. ([Accessed 17-Jun-2023])
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—miccai 2015: 18th international conference, munich, germany, october 5-9, 2015, proceedings, part iii 18* (pp. 234–241).
- Russell, D. G. (2007). Who puts the tubercle in tuberculosis? *Nature Reviews Microbiology*, *5*(1), 39–47.
- Stoop, E. J., Schipper, T., Rosendahl Huber, S. K., Nezhinsky, A. E., Verbeek, F. J., Gurcha, S. S., ... van der Sar, A. M. (2011). Zebrafish embryo screen for mycobacterial genes involved in the initiation of granuloma formation reveals a newly identified esx-1 component. *Disease models & mechanisms*, *4*(4), 526–536.
- Swaim, L. E., Connolly, L. E., Volkman, H. E., Humbert, O., Born, D. E., & Ramakrishnan, L.

- (2006). Mycobacterium marinum infection of adult zebrafish causes caseating granulomatous tuberculosis and is moderated by adaptive immunity. *Infection and immunity*, *74*(11), 6108–6117.
- Talbot, E. A., & Raffa, B. J. (2015). Chapter 92 - mycobacterium tuberculosis. In Y.-W. Tang, M. Sussman, D. Liu, I. Poxton, & J. Schwartzman (Eds.), *Molecular medical microbiology (second edition)* (Second Edition ed., p. 1637-1653). Boston: Academic Press. Retrieved from <https://www.sciencedirect.com/science/article/pii/B9780123971692000925> doi: <https://doi.org/10.1016/B978-0-12-397169-2.00092-5>
- Thisse, C., & Zon, L. I. (2002). Organogenesis—heart and blood formation from the zebrafish point of view. *Science*, *295*(5554), 457-462. Retrieved from <https://www.science.org/doi/abs/10.1126/science.1063654> doi: 10.1126/science.1063654
- Titford, M. (2005). The long history of hematoxylin. *Biotechnic & histochemistry*, *80*(2), 73–78.
- Tobin, D. M., & Ramakrishnan, L. (2008). Comparative pathogenesis of mycobacterium marinum and mycobacterium tuberculosis. *Cellular microbiology*, *10*(5), 1027–1039.
- Trede, N. S., Langenau, D. M., Traver, D., Look, A., & Zon, L. I. (2004). The use of zebrafish to understand immunity. *Immunity*, *20*(4), 367-379. Retrieved from <https://www.sciencedirect.com/science/article/pii/S1074761304000846> doi: [https://doi.org/10.1016/S1074-7613\(04\)00084-6](https://doi.org/10.1016/S1074-7613(04)00084-6)
- van der Sar, A. (2003). Zebrafish embryos as a model host for the real time analysis of salmonella typhimurium infections. *cell microbiol. Cellular Microbiology*, *5*, 601–611.
- van Leeuwen, L. M., van der Sar, A. M., & Bitter, W. (2015). Animal models of tuberculosis: zebrafish. *Cold Spring Harbor perspectives in medicine*, *5*(3).
- Verbeek, F. J. (1999). Three-dimensional (3d) reconstruction from serial sections. *Image Processing and Analysis: A Practical Approach*, *219*, 153.
- Wang, F., & Zai, Y. (2023). Image segmentation and flow prediction of digital rock with u-net network. *Advances in Water Resources*, *172*, 104384. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0309170823000192> doi: <https://doi.org/10.1016/j.advwatres.2023.104384>
- Wang, Y., Chen, Y., Yang, N., Zheng, L., Dey, N., Ashour, A. S., ... Shi, F. (2019). Classification of mice hepatic granuloma microscopic images based on a deep convolutional neural network. *Applied Soft Computing*, *74*, 40–50.
- Wodzinski, M., & Müller, H. (2021). Deephistreg: Unsupervised deep learning registration framework for differently stained histology samples. *Computer Methods and Programs in Biomedicine*, *198*, 105799. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0169260720316321> doi: <https://doi.org/10.1016/j.cmpb.2020.105799>
- World Health Organization. (2022). Global tuberculosis report 2022. In *Global tuberculosis report 2022*.
- Zhang, X., Dong, H., Gao, D., & Zhao, X. (2020). A comparative study for non-rigid image registration and rigid image registration. *arXiv preprint arXiv:2001.03831*.
- Zitová, B., & Flusser, J. (2003). Image registration methods: a survey. *Image and Vision Computing*, *21*(11), 977-1000. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0262885603001379> doi: [https://doi.org/10.1016/S0262-8856\(03\)00137-9](https://doi.org/10.1016/S0262-8856(03)00137-9)