Okinawa Institute of Science and Technology
Graduate University

Thesis submitted for the degree

Doctor of Philosophy

Automated segmentation of micro-CT images by deep learning and its application to comparative morphology

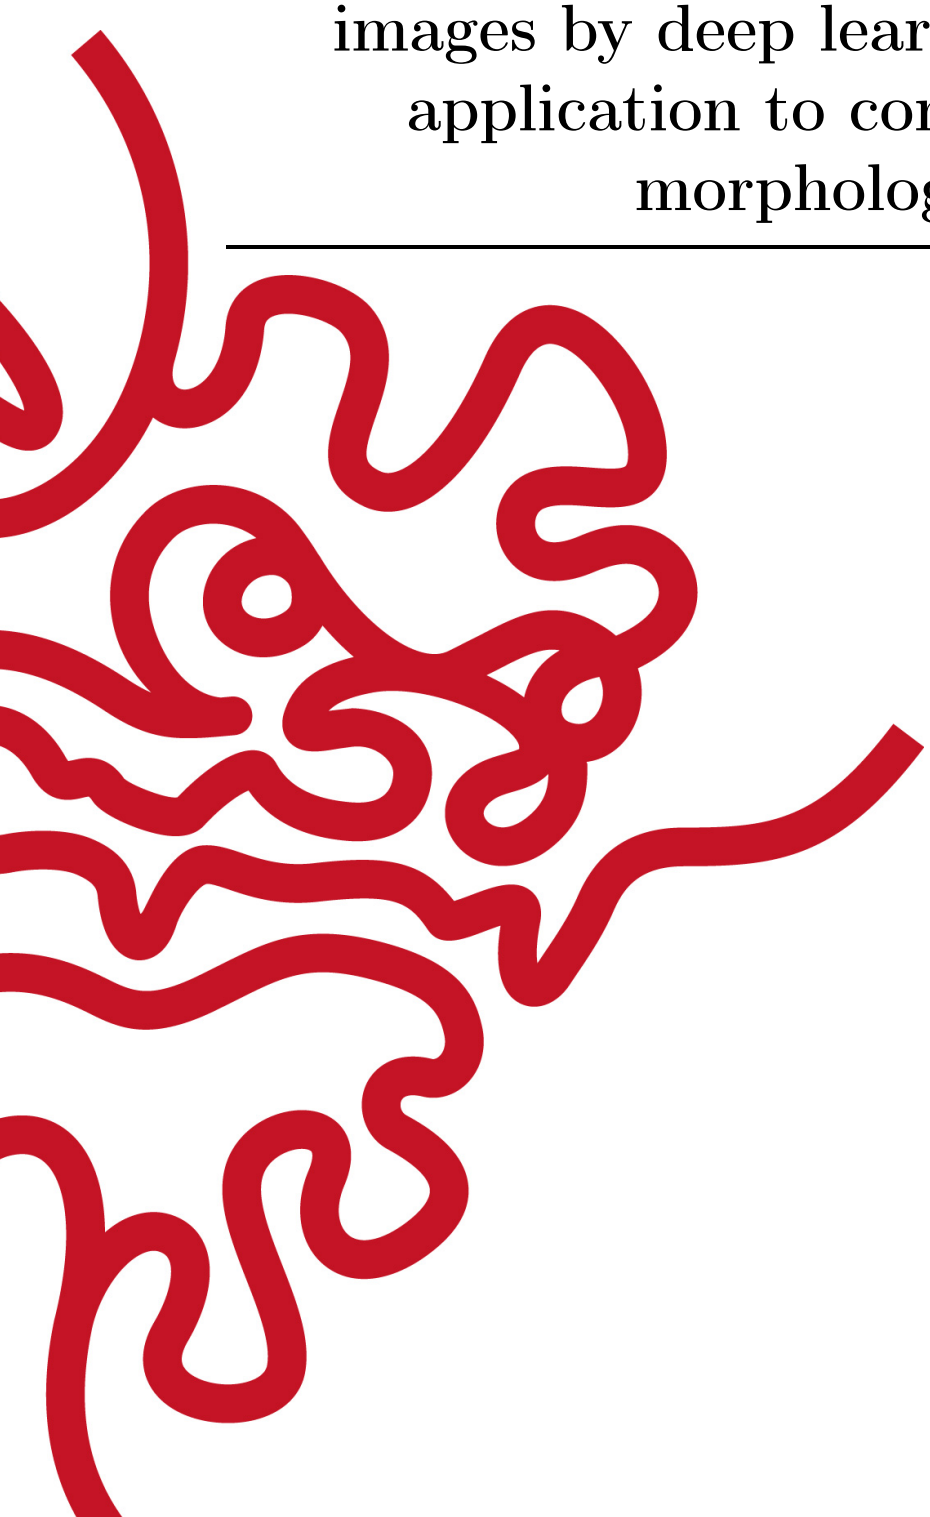

by

Evropi Toulkeridou

Supervisor: **E. Economo**
Co-Supervisor: **K. Doya**

December 2021

Declaration of Original and Sole Authorship

I, Evropi Toulkeridou, declare that this thesis entitled *Automated segmentation of micro-CT images by deep learning and its application to comparative morphology* and the data presented in it are original and my own work.

I confirm that:

- No part of this work has previously been submitted for a degree at this or any other university.
- References to the work of others have been clearly acknowledged. Quotations from the work of others have been clearly indicated, and attributed to them.
- In cases where others have contributed to part of this work, such contribution has been clearly acknowledged and distinguished from my own work.

Date: December 2021
Signature:

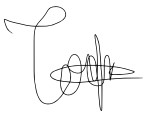

Abstract

Automated segmentation of micro-CT images by deep learning and its application to comparative morphology

Image segmentation is one of the most fascinating challenges of computer vision. A field of potential application is organismal biology, where researchers are increasingly using three-dimensional (3D) scanning which produces data-rich volumetric images for precise and comprehensive anatomical characterization. To date, the segmentation of anatomical structures remains a bottleneck to research, as it is commonly performed with highly tedious and time-consuming manual work. During recent years, however, machine learning methods are an emerging approach to overcoming this limitation, especially with the use of deep learning techniques such as convolutional neural networks (CNNs), which proved to be very efficient and, as such, promising candidates for image segmentation. The main objective of this PhD project was to develop a pipeline for the fully-automated segmentation of anatomical structures in micro-computed tomography (micro-CT) images of insects using state-of-the-art deep learning methods. The restricted number of available high-resolution 3D labeled images necessitated the use of a CNN architecture that performs segmentation satisfactorily even with limited data; the U-Net architecture is such a CNN that has shown good performance in medical images using few annotated images. Ant brains were selected as the test case. Since no dataset of micro-CT images of ant brains existed for the current case study, a new extensive dataset was created across a wide variation of 94 ant species. Its existence can be of importance, as brain images of ants are similar to those of other insects; therefore, our dataset can act as a starting point for the development of a substantial library of micro-CT images of insects, and work as a pre-training dataset for future CNNs. Also, our network is generalizable for segmenting the whole neural system in full-body scans, and works in tests of distantly related and morphologically divergent insects (e.g., fruit flies). The latter suggests that algorithms such as our network can be applied generally across diverse taxa. The chosen species set was designed to be interesting for further evolutionary morphology analysis. Therefore, we used it to test the social brain hypothesis for ants, i.e., whether there is a connection between the brain investment and the sociality of each species. Volumetric statistical analysis was performed, also considering phylogenetic data; its results, however, did not validate the hypothesis.

Acknowledgment

Before starting the discussion about the research conducted and results obtained within the framework of this PhD project, I wish to thank my supervisor, Prof. Evan Economo, as well as my co-supervisor, Prof. Kenji Doya. I consider myself really privileged to have worked under such supportive supervisors, both on the professional and the personal level.

I would especially like to thank my academic advisor, Dr. Daniel Baum of Zuse Institute of Berlin (ZIB), who oversaw the main body of my research and guided me with fruitful suggestions and revisions. I will always be grateful for his supervision over the past four years.

My thanks also goes to all past and present members of the Biodiversity and Bio-Complexity and the Neural Computational Units at OIST. Especially, I would like to thank Dr. Carlos Enrique Gutierrez for discussions and input on deep learning aspects of my project, Dr. Alexandre Casadei Ferreira and Dr. George Fischer, who had the patience to explain to a mathematician about the importance of ants in ecology, and for teaching me how to process ant specimens and use the micro-CT scanner. Also, I am indebted to Dr. Henrik Skibbe for useful discussion, and Drs. Noushin Hajarolasvadi, Dan Warren, and Jamie Kass for their valuable advice, both scientific and technical, while putting my thesis together. Further, my colleagues Mr. Shubham Gautam for insightful discussions, input, and help regarding the brain evolution aspect of my project, as well as for personal support, Mr. Julian Katzke and Ms. Fumika Azuma for their help with scanning, Dr. Francisco Hita Garcia for his help with organizing the specimens, Prof Timothy Linksvayer and Dr. Chao Tong for providing the specimens and the phylogenetic tree used for this study, Dr. Ales Buček and Prof. Todd A. Schoborg for providing fly and termite scans. Finally, I would like to specially acknowledge Mr. Pavel Puchenkov from the Scientific Computing and Data Analysis section at OIST for assistance relaed to graphic design.

I would like to express my gratitude to my husband, Panagiotis, who is always there to support, love, and encourage me, as well as to our magnificent son, Iraklis, who tolerated mum's busy schedule. I would also like to thank my family in Greece (my parents Dimitris and Sofia), my family in Cyprus (Giorgos, Marina, Dimitris and Sergios) and "family" in Okinawa (Ami, Jeremie, Shinji and Rei), and my friends (especially Domna, Shoko, Angela and Albert, Kosmas, and Theo) who believed in me and supported me in many ways all these years.

Finally, I would like to thank Ms. Loretta White from Ganjuu Wellbeing Service at OIST, and Ms. Chiaki Chibana, Ms. Kozue Higashiona, and Dr. Harry Wilson from OIST Graduate School for their support during the hardest personal times during this

journey.

Abbreviations

AI	Artificial Intelligence
AIC	Akaike Information Criterion
BM	Brownian Motion
CNN	Convolutional Neural Network
CT	Computed Tomography
F1	Dice Coefficient
FCN	Fully Convolutional Network
FPN	Feature Pyramid Network
FPR	False Positive Rate
GAN	Generative Adversarial Network
GPU	Graphics Processing Unit
IoU	Intersect over Union (Jaccard Index)
LR	Linear Regression
ML	Maximum Likelihood
MRI	Magnetic Resonance Imaging
OLS	Ordinary Least Squares
OU	Ornstein-Uhlenbeck
PaNet	Path Aggregation Network
PGLS	Phylogenetic Generalized Least Squares
RELU	Rectified Linear Unit
ROI	Regions of Interest
RNN	Recurrent Neural Network
RPN	Region Proposal Network
SVM	Support Vector Machine
SGD	Stochastic Gradient Descent
TPR	True Positive Rate

To Iraklis

Contents

List of Figures

List of Tables

Chapter 1

Introduction

"in biology... we are drowning in a sea of data and starving for knowledge"
Sydney Brenner

Artificial intelligence (AI) and machine learning have become very popular concepts not only in computer science but in almost every scientific field over the past two decades [1], even though they were conceptualized more than half a century ago. Having gone through two dark ages (or "AI winters" as they came to be known, alluding to a decline in popularity and stagnation), AI has been experiencing a spring of rapid growth since the 1990s, after IBM's Deep Blue supercomputer beat Soviet chess grand-master Gary Kasparov at a six-game match, as shown in Figure 1.1. Even though there have been difficulties since then, the field of AI is currently growing faster than ever before. The reasons behind this paradigm shift are the availability of massive experimental data that cannot be appreciated by human eye or conventional statistics, and the development of increasingly powerful computers and fast hardware that are both easier to use and more accessible to everyone.

Computers are harnessed to build AI by simulating the real world; however, despite the progress in understanding how the human brain works and comprehends the world, the development of learning algorithms that enable AI to recognize semantic concepts, as brains do, still remains challenging [3]. Yet, under the machine learning general umbrella term, deep learning emerged, utilizing artificial neural networks with multiple layers (some of which are hidden), as shown in Figure 1.2. Inspired by the human brain, deep learning methods gained more interest in recent years, especially with the advent of the 4^{th} scientific paradigm and the rise of the big-data era.

Big datasets are essential for deep learning methods to perform sufficiently. However, for many biological studies either the required massive datasets needed for training are not attainable or the tools to analyze them are not sufficient. This comes as a result of the inherent complexity of biological datasets, which need to capture a multitude of variable facets of biological systems. Moreover, through misevaluation of the needs and/or tools provided by different fields, biological datasets are often orders of mag-

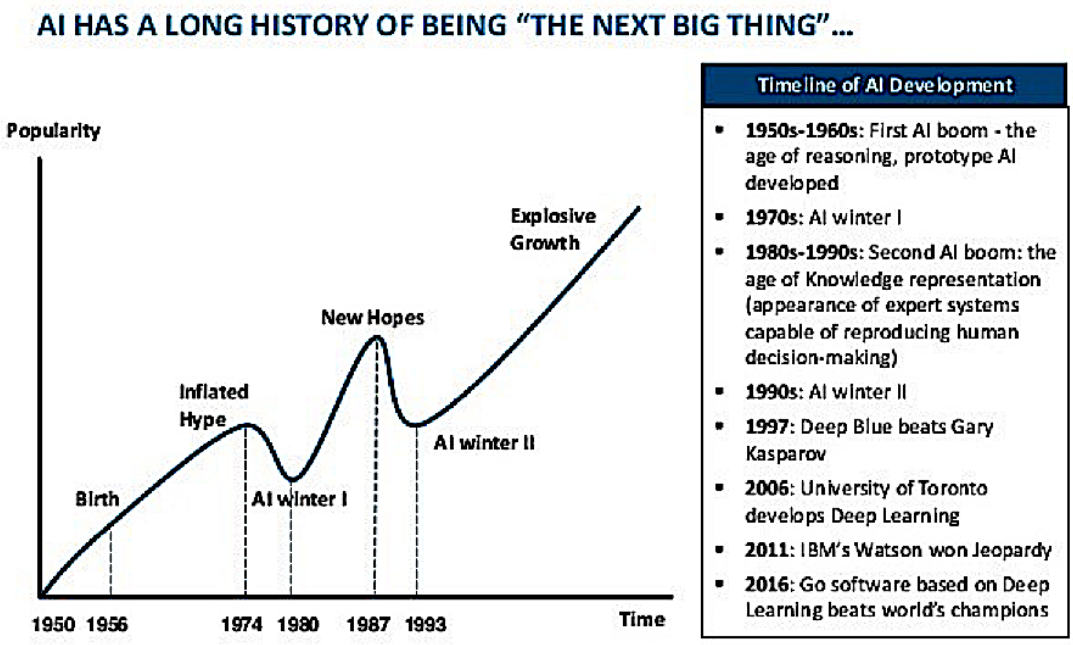

Figure 1.1: The boom and bust cycle of AI research. Reproduced from [2]

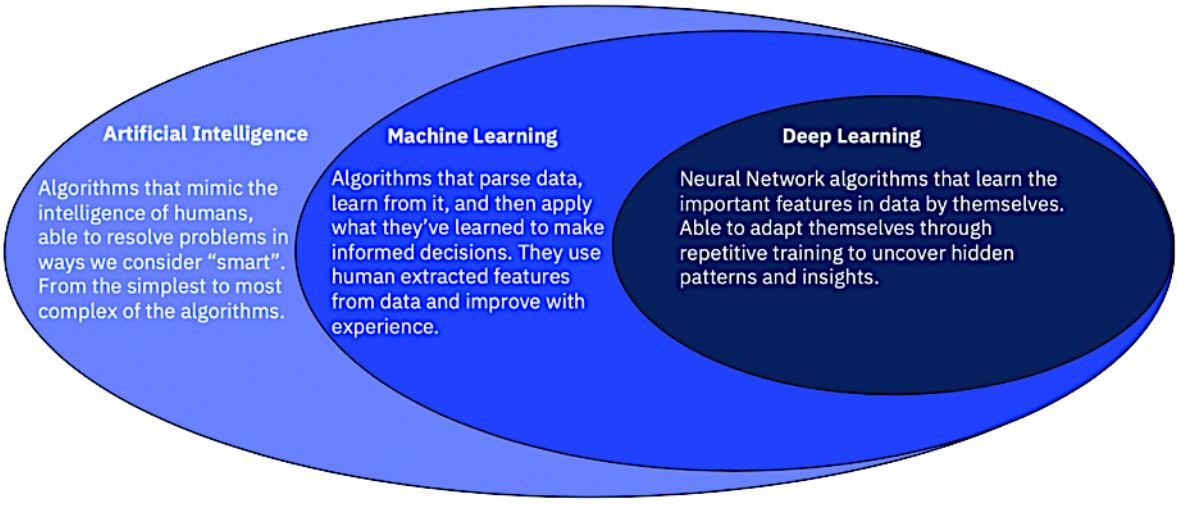

Figure 1.2: Specialization of AI algorithms. Reproduced from [4]

nitude too small for training deep learning networks. Considering that annotation of the raw data, although essential, can also be extremely time consuming, it comes as no surprise that usable biological datasets are being developed only during the last decade.

Deep learning based segmentation methods mostly use supervised learning; therefore, the unavailability of large amounts of labeled training data for medical images has been a major bottleneck for research. Indeed, for classification projects most investment has been made on natural image datasets instead, simply because they are easier to annotate. Moreover, the old computer science motto *"garbage in, garbage out"* holds undeniably true when we talk about deep learning methods. The performance of any deep learning technique depends on the training data as much as on the model itself; therefore, the quality of the data, as well as their quantity, can change the prediction process dramatically. As a result, proper pre-processing is a critical step in addition to the collection of high-quality images. Investing in the creation of well-annotated, large, and clear datasets can lead to state-of-the-art studies, also producing specifically designed deep learning algorithms.

In addition to an increase in the number of useable datasets during recent years, there has also been a significant boost in the number and efficiency of deep learning methods used for the analysis of biological data, as shown in Figure 1.3. As the number of studies focused on classification and segmentation of medical images keeps increasing, more datasets and deep learning methods emerge; from a simple literature search on Scopus, only in the last two years the number of Computer Vision and Pattern Recognition (CVPR) conference papers doubled from ∼2,000 to ∼4,000, while the number of Medical Image analysis studies in 2021 have exceeded the total number of the studies in the field over the previous three years. From these simple considerations, it is clear that the fields of computer vision and medical imaging are growing, attracting thousands of scientists from various fields and getting funded for interdisciplinary projects.

However, this is not the case for studies focusing on ecological data and, more specifically, on the segmentation of vertebrate and insect images. Although digitalization of most physical specimens with the use of micro-computed tomography (micro-CT) has led to an explosion of imaging data in ecology, neither a dataset of annotated pre-processed images existed, nor any other study had focused on the automatic analysis of the data prior to this study.

1.1 Ants as model organisms

Ants are well-known for their significance in ecology as they are extraordinarily diverse and combine both collective behavior and division of labor [6]. Though numerous studies have focused on their behavior and anatomy, their evolution leading to both still remains largely unclear. Micro-CT scanning of ants has enabled a more thorough view of their morphological and anatomical characteristics, which has led to the discovery and definition of new ant species [7] and the analysis of their evolution. The creation of a 3D, bias-free atlas of ants will provide a major boost to evolutionary studies. As such, ants were chosen as the case study for the present segmentation neural

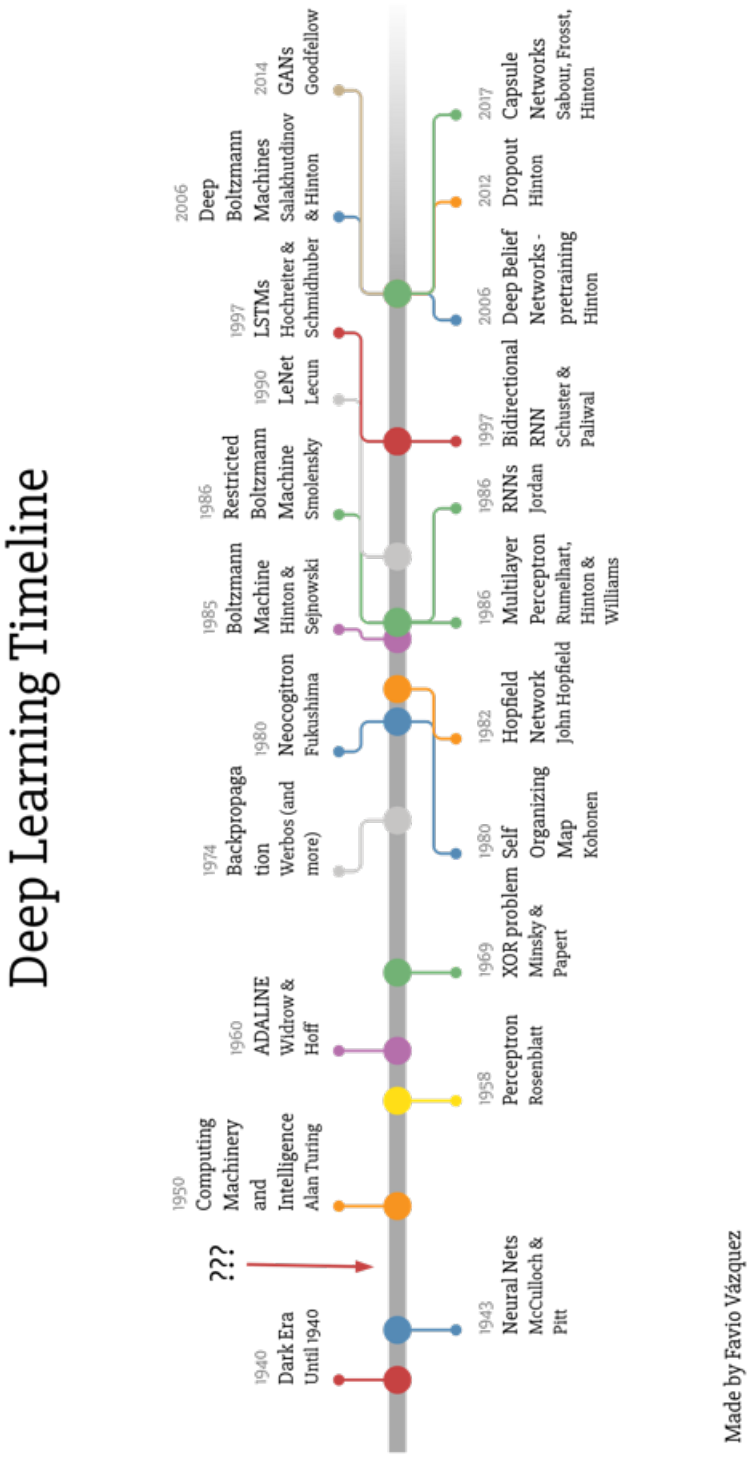

Figure 1.3: Timeline of deep learning. Reproduced from [5]

network. Furthermore, one of the most interesting questions in ecology is the evolution of the brain [8], especially in social insects, and its relationship to their behavior. It is commonly formulated as the "social brain hypothesis", which postulates that the social behavior of animals adapts to changes in their ecosystem and/or social structure with modifications in their brain investment. Thus, the exploration of a potential connection between the ant brain size variation and the colony size and structure is a problem of high importance, since it may elucidate how the environment affects the organism's most important organ, the brain.

1.2 Objectives

The current PhD study focused on two main questions:

- Does brain size change with colony size across ant species as predicted by the social brain hypothesis?
- How can we automate the segmentation of micro-CT images of insects, using ants as a test case?

To tackle both problems, a database of manually segmented micro-CT head images of ants was created, which was used as morphological set for the first part of the study and as a training set for the second part. As a next step, after the volumes of all brains and bodies of the tested ants were calculated, their morphological and ecological data were statistically analyzed, providing an answer to the first question. Finally, using deep learning methods, a convolutional neural network (CNN) [9] was created, which was trained and tested with the aforementioned dataset, and which could segment automatically the brain area of the micro-CT images of ants.

1.3 Thesis structure

The thesis is structured as follows:

Chapter 2 summarizes the methodology and theoretical foundation needed to appreciate the main body of work. Starting with an introduction to micro-CT and statistical analysis used in ecology, the main focus of Chapter 2 is on deep learning methods for biomedical image analysis.

Chapters 3, 4 and 5 constitute the main body of this research work. All chapters are heavily based on publications under preparation, with Chapter 5 having already been uploaded to BioRxiv; as a result, some degree of redundancy is expected. Chapter 3 focuses on the creation of the dataset that was used in both Chapters 4 and 5, providing details for the ant species that were used in this study. Chapter 4 focuses on the morphological study of the ant brain and explores how its evolution may be connected with the social structure/complexity of the colony. After the said connection is hypothesized, the results of its statistical evaluation are investigated. Chapter 5 focuses on the development of an automatic segmentation framework for the analysis of micro-CT images of ants.

Chapter 6 concludes this thesis by summarizing the key points of the whole work and projecting into the future.

Chapter 2

Technichal and Theoretical Background

2.1 Introduction

In this chapter the experimental technicalities and theoretical concepts used in later chapters are introduced. Moreover, a brief review of related literature is presented; although by no means exhaustive (especially considering the overwhelming amount of new publications reported practically every day), it aims to outline a main body of knowledge that formed the basis for the research of the current PhD project.

To present the methodologies employed for this study in the most intuitive way, we opted for a rather chronological approach reflecting each individual aspect of the project. For this reason, micro-CT scanning is introduced first, as it was a necessary prerequisite for both main research investigations (elaborated in Chapters 4 and 5). Subsequently, statistical methods pertinent to ecology are touched upon, as the first research investigation analyzed the acquired micro-CT scans to focus on the morphology and evolution of the ant brain (Chapter 4), followed by specific reference on the social brain hypothesis and its applicability on ants. The main body of the current PhD project, however, revolved around automating the segmentation of the brain area in ant micro-CT images (Chapter 5); therefore, fundamental concepts of machine learning are introduced last, with a special focus on the specific deep learning techniques and architectures used here.

Our aspiration is that the utilization of two primarily different statistical and computational tools (or families of tools) combined with hands-on sample preparation and imaging provide a broad view of how mathematics can be applied in evolutionary biology and a roadmap for the conduct of cross-disciplinary research in the field.

2.2 Micro-CT scanning

X-ray computed tomography (CT) is a non-invasive biomedical imaging technique that can accurately and efficiently provide detailed structural information from a series of projections of an X-ray beam passing through a specimen. More specifically relevant to the current PhD project is micro-CT, i.e., high-resolution CT that can provide phe-

notypic data of small specimens such as insects [10, 11]. Due to its non-invasiveness and high-resolution 3D imaging of small features the technique quickly gained popularity with evolutionary biologists, availing new opportunities for in-depth morphological studies [12]. As a result, micro-CT imaging enabled the development of extensive libraries of 3D scans of vertebrates and insects [13]. Recently, modern synchrotron facilities provide high-flux beams, high brilliance, and superb stability, and allow the user to select the wavelength of the emitted radiation; they can, thus, yield micro-CT imaging data at higher resolution and fewer artifacts than conventional CT scanners. As the need for such high-resolution imaging keeps rising, the number of facilities around the world that can provide micro-CT imaging (including synchrotron facilities) is growing at a fast pace, too [14].

All micro-CT systems consist of an X-ray source formed as a small-sized microfocus tube, a detector or camera, and a holder for the specimen being imaged. Scanner functions can range from extremely simple to very complex depending on the scanner geometry [11]. Using micro-CT requires accuracy and precision in positioning the object at the right distance from both the detector and the source, rotating the specimen about the central axis of view, or choosing vertical stitch (i.e., translating the object vertically) to spread the slices when the specimen is larger than a certain threshold size. Moreover, the user chooses the maximum energy and the tube current for each scan, depending on the absorption of the material. Specimen preparation also requires meticulousness since good stability is imperative for high quality image acquisition. Holders are typically provided by the instrument developer, but in the present case a new customized holder was commissioned and fabricated by OIST's imaging section, since ants can be quite small and fragile, which minimized the shifting and trembling of the specimen during scanning. Compensatory algorithms can, in principle, further minimize mechanical motion control errors, but were avoided as they are not always reliable. Ultimately, the utility of a micro-CT scan is judged by the possibility to discern the features of interest (e.g., the surface of dry specimens or the inner parts of wet specimens), which is determined both by technical specifications of the instrument and the expertise and diligence of the user.

For the current PhD project, scans were performed using a ZEISS Xradia 510 Versa 3D X-ray microscope, schematically shown in Figure 2.1, and ZEISS Scout and Scan Control System software (version 10.7.2936) at OIST. Scan settings were determined by (either dry or wet) specimen size with a view to optimizing scan quality. In order to assess the importance of individual setting values, it was essential to systematically scrutinize how scanning parameters affect the quality, resolution, and contrast of output images during the acquisition process. For the 3D image to be used for both manual (Chapter 4) and automated segmentation (Chapter 5), the output required the highest possible resolution and contrast. 3D reconstruction of the resultant scans was performed with XMReconstructor (version 10.7.2936) and saved in DICOM file format (default settings; USHORT 16-bit output data type). These were processed with Amira software (version 6.0) and ITK- SNAP (version 3.6.0) in order to be used for manual segmentation and shape analysis.

Figure 2.1: Schematic of OIST's ZEISS Xradia 510 Versa 3D X-ray microscope micro-CT scanner. The pipette containing a wet specimen and resultant head scan are shown as insets.

2.2.1 Micro-CT scanning of dry specimens

Before starting work on the more demanding scans of wet specimens, it was deemed expedient to obtain experience by scanning dry ant specimens first. Dry specimens were initially glued to their paper point. Eighty different species of various sizes belonging to the ant genus *Malagasy Stumigenys* were subsequently scanned; this provided enough data to contribute to the analysis of the morphological evolution of adaptive radiation of five different genera in Okinawa. The scans were processed for surface rendering using Amira software; 42 landmarks were placed in the body using MeshLab and Stratovan Checkpoint software. The morphological diversity of their mandibles was thus readily identifiable from their 3D images, paving the way toward understanding mandible evolution (ongoing investigation, in parallel to current PhD project).

2.2.2 Micro-CT scanning of wet specimens

Wet specimens were first stained with iodine and then placed in pipettes of variable sizes. More than 100 ant heads and 80 bodies were scanned. An average head scan lasted for $\sim$ 10 hours, in addition to an average of $\sim$3 hours per body scan. Subsequently, all head scans were manually or semi-automatically segmented with the use of Amira software, and $\sim$50,000 raw images of heads were created, with masks annotating the brain areas. More details on scanning and processing of wet specimens, which were at the core of this PhD project, can be found in Chapters 4 and 5.

2.3 Statistical analysis models for phylogenetic signal

Whether phenotype is shaped by phylogenetic history is one of the quintessential questions in ecology and evolutionary biology [15–17]. The fact that closely related species and taxa are biologically and statistically dependent often leads to the assumption that ecological data are phylogenetically connected [18]. Nevertheless, when we are comparing traits across taxa the assumption of phylogenetic dependence should first be scrutinized. Since residual errors in ordinary least squares (OLS) regression are assumed to be independent among observations, results of naive OLS analysis of trait data are error prone. To overcome this obstacle, several statistical methods have been developed that take phylogeny into account for the statistical analysis of traits among taxa, with the most common being the phylogenetic generalized least squares (PGLS) regression [16, 19].

As a first step, before the analysis of any ecological data one needs to evaluate which model the evolution traits follow. Within the framework of this PhD project, we tested three different models designed for ecological data. Which model best fitted our data was ultimately determined with the use of the Maximum Likelihood (ML) test [20] and the Akaike Information Criterion (AIC) [21], two of the most commonly used mathematical methods for model evaluation.

The simplest way to describe evolution is by assuming it follows a trajectory of Brownian motion (BM), i.e., a random walk [22, 23]. Under Brownian evolution, trait values change randomly due to genetic drift. Although Brownian evolution is a popular model due to its simplicity and convenience for calculations, trait values can in fact evolve under various scenarios. To include these variations in trait evolution, Pagel introduced a statistical model [24, 25] that can be determined by three branch length transformations (i.e., lambda, λ; delta, δ; and kappa, κ) of the phylogenetic tree, with the most common being lambda. Lambda values range between 0 and 1, measuring the phylogenetic signal: when λ equals 1, Pagel's model is equivalent to the Brownian model; when lambda equals 0, the species are statistically independent, i.e., they lack phylogenetic signal. Additionally, to include the effect rates of evolution changing over time, Pagel designed the δ transformation. Finally, the κ transformation is focused on speciation events of the phylogenetic tree: when $\kappa = 0$, all change is associated with speciation events and there is no anagenetic change. Most ecological studies restrict themselves to exploring the effects of the lambda parameter. Finally, an improved Brownian model is the Ornstein-Uhlenbeck (OU) model, which describes the pattern of traits that evolve under stabilizing selection but with constant optimum [26].

After identifying the phylogenetic model that best fits the pattern of trait evolution, PGLS can be performed to judge whether there is an association between the dependent/response variable/trait and the predictor variables/traits over evolutionary time. In practice, it applies the correlation structure determined by the phylogeny and the model of character evolution to a generalized least square regression; thus, if there is no phylogenetic signal in the data, the PGLS will return the same result as OLS. Therefore, identifying the evolutionary model of the traits is important, as it can lead to statistical errors if misclassified.

Logarithmic transformation is commonly used in continuous traits and comparative data in ecology before any analysis is performed, since values typically follow distri-

butions that are skewed to the right. This transformation often results in variables with normal error distributions, which helps to meet the assumptions of many statistical methods. As an additional advantage, the distance between trait values can be measured more easily, as taking logarithms improves linearity. In ecological data, logarithmic transformations can equalize the relative weights of common and rare species in addition to highlighting the informative species [27].

2.4 Social brain hypothesis for ants

The brain is one of the most expensive investments in every animal's body [28]; as a result, investing in a larger brain can have a pronounced impact on other physiological traits [29]. More importantly, when we discuss the ant brain we need to distinguish whether we are referring to the investment of individual ants or of the entire colony. Individual workers in ant societies have often been thought of as sensory units of a "collective brain" [30]. Ant colonies, on the other hand, have been described as superorganisms which represent societies varying from simple-structured (one queen and few workers) to more complex (multiple queens and thousands of workers) [31]. For more complex structures, each member has a single task, although this can be modified according to the dynamic needs of the colony [32]. In simpler colonies, each member performs numerous tasks to ensure survival and growth of the society [33]. In small colonies, ants are obliged to rely on centralized control, adapt to changes in colony structure and size, and perform tasks correctly and efficiently. Moreover, individual workers of small colonies need to coordinate their behavior with that of other individuals more efficiently than in large colonies, which are more homeostatic [34]. In contrast, in large colonies there is always redundancy through parallel operation: more ants are responsible for the same task, and thus the probability of success in the collective task is high even if some individuals are less competent. It is therefore reasonable to assume behavioral routines are connected to social structure. Moreover, brain size, neurophysiology, and body size are correlated with behavioral routines especially in polymorphic ant colonies [35]. However, it is currently unknown whether brain evolution responds to the social complexity of ant colonies. The great number and diversity of different ant species provide an excellent opportunity to address these questions statistically.

The social brain hypothesis [36] connects the adaptability of the social behavior of animals to modifications in their ecology and social structure with changes in brain investment [37]. This hypothesis implies that selection pressures associated with social living favor larger brains to navigate social interactions [38]. Originally it was proposed for mammals, where the community (family, herd, pride, etc.) is formed by members who need to gain individual recognition, understand social hierarchy, cooperate and/or compete for social rank, etc. As a result, most studies on brain size and sociality have been performed on vertebrates [39]. However, in social insects, societies are gigantic families formed by clones or siblings, whose well-developed communication systems are based on chemical mechanisms; these societies function through division of labor and cooperation. Because of these marked differences from mammal societies, in recent years there is an increasing number of studies on social insects, and more specifically on how their neuropils (i.e., the parts of the brain responsible for memory, learning, and

sensory functions) and their cognitive capacities are connected to ecological traits [40].

Neuroethological studies have shown that simple recognition of colony mates is connected to simple cognitive operations, and there is no need for higher integration centers or acquisition of a long-term memory [41]. Thus, ants' mushroom bodies (i.e., their neuropils responsible for memory) do not show any change. However, ants can overcome limitations in information processing through communication, with their eyes and antennas playing a significant role in their adaptation and survival. This can lead to changes in their optical or antennal lobes (i.e., the neuropils that receive environmental information through eyes or antennas) and to their overall brain size [42]. In small ant colonies, the cognitive demands that individuals face differ from those in large colonies. The need for larger eyes (for example to identify other members of the colony or to enable more efficient foraging) is imperative in small colonies. Moreover, workers are responsible for different social roles and experience their environment differently, and as a result their sensory processing ability and eye structure have evolved in different ways. As an example of how the environment can affect their morphology, bull ants *Myrmecia pyriformis* are active mainly during the dawn and dusk twilight and their eyes are extremely large so as to collect the most possible information for their navigation. Moreover, as colonies can be considered organisms themselves with individual workers behaving as organs, investment in individual brain and eye size rather than overall colony size can benefit the functionality of the colony as a whole. As the social brain hypothesis suggests, behavioral problems such as foraging are solved socially rather than at the level of single individuals, which may impact brain size evolution. Most of the studies find a connection between brain and body sizes, as well as a connection between the environmental changes and the brain investment in different neuropils [40, 43]. Support for the social brain hypothesis is still mixed, however; there are a number of studies that come to contradictory conclusions, and the majority of such studies present data from only a few ant species each time [44, 45]. The connection between brain size and sociality in ants is complex and requires a broad approach, since each colony is unique in structure and adaptable to environmental changes. Additionally, brain plasticity in queens and the size difference between soldiers and workers in polymorphic colonies should be considered in choosing study species.

2.5 Computational methods for automated image segmentation

2.5.1 Computer Vision

While strongly connected with AI, computer vision is a distinct, interdisciplinary scientific field that focuses on understanding how computers can learn several concepts from digital images or videos. The three main steps to simulate human vision are (i) image acquisition, (ii) image processing, and (iii) image analysis/understanding. Recent technological accomplishments have rendered the first step an easy task, resulting in the acquisition of high-quality images and video [46]. Image processing algorithms, on the other hand, keep increasing and, by taking advantage of new available software and hardware, improving [47]. The general term "image analysis" includes segmenta-

tion, semantic segmentation, classification, and object detection, the most challenging one being the latter, as object recognition and tracking are the most complicated tasks. Noisy or scarce data, real-time processing, and limited machine power are some of the challenges computer vision needs to overcome. In recent years CNNs [9] have gained popularity by solving computer vision problems like object recognition and semantic segmentation [48]. The term semantic segmentation designates the segmentation of an image in non-overlapping pixelwise common areas belonging to the same objects, such as cars or humans in scene analysis, and cells or organs in biomedical image analysis. This means that after semantic segmentation, an image is divided into areas similar to those identified by the human eye.

2.5.2 Machine Learning

Machine learning is a research area which deals with the computer learning process. Learning is built on statistical and mathematical rules with a view to developing predictive models based on the statistical associations among features of a given dataset. The learned model can be used to predict categorical labels, continuous values and binary responses. As amounts of data are accumulating fast due to the increasing computer power available, machine learning aims to make their processing automated. Mapping input data or hand-designed features with new, unlabeled output data is its basic rule; therefore, depending on how much feedback the user gives to the computer, machine learning classes can be categorized as [49]:

a. Supervised learning. In supervised learning the input/training data are paired with target output. From them the machine learns how to predict the output for any new, testing data. There are two main categories of supervised learning problems: classification, where the machine tries to identify to which category the testing data belong, and regression, where the machine tries to predict a continuous output value for the given input data.

b. Unsupervised learning. In unsupervised learning all data are unlabeled. The machine discovers the structure in the data distribution and creates a low-dimensional representation; the most common approach is clustering. Since no annotated data are needed, unsupervised learning can be applied even when sample labels are incorrect or missing. At the end of the learning/clustering process it is hard to evaluate the result, as the data are all unlabeled.

c. Human in the loop – semi-supervised learning. Semi-supervised learning is a combination of supervised and unsupervised learning, as the training data are divided into labeled and unlabeled. Labeled data are used to identify the data classes, and unlabeled data give additional information from the overall data distribution. In the "human in the loop" scenario unlabeled data that are hard to classify are presented to human experts for improvement of the boundaries between classes.

2.5.3 Machine learning techniques for image segmentation

Image segmentation algorithms divide a given image into non-overlapping homogeneous and uniform areas, each having the same features. Semantic segmentation algorithms have evolved over the past decades from split-and-merge [50], and region growing [51]

in the 1970s, to clustering algorithms in the 1980s [52, 53], and to Markov Random Fields [54, 55] and other models eventually [56]. Grayscale and color images have been processed with the most commonly used clustering algorithms such as K-Means clustering [57, 58], watershed segmentation [59], and Expectation-Maximization (EM) clustering [60]. As image data increased and images became more detailed, the necessity for new algorithms emerged. Machine learning techniques were soon adopted by computer vision in order to make segmentation more efficient. Supervised learning by Support Vector Machines (SVMs) is one of the most common image segmentation machine learning architectures [61]. Two main learning approaches have gained attention in the past decades: probabilistic graphical models [62] and neural networks [63]; their difference lies in their use of either probabilistic graphical models or computation graphs, respectively. Data features are the determining factor for the performance of each method. Therefore, pre-processing pipelines and data transformations are very important and lead to more effective learning. On the other hand, the effort required for such "feature engineering" is often too costly, and, as such, it is the biggest drawback of conventional machine learning techniques. An AI that can learn fast and also understand the world is the desirable outcome which forced novel applications to emerge [64].

2.5.4 Deep Learning

Deep learning belongs to the machine learning family of methods. The method is quite old, though it became trendy only during the past decades. The architecture of deep learning algorithms is similar to that of the human brain (cerebral cortex) in that it is built in layers. In deep learning the computer combines simple concepts to create more complex ones that can help solve complicated representation learning problems. The term "representation learning" describes the machine learning methods that use raw data as input and let the machine find the representations required for classification or detection [65]. The machine extracts from these data the important features and creates learning layers for each of them. Complicated structures of high-dimensional data can be thus simplified with the use of deep learning architectures. Like their machine learning counterparts, deep learning algorithms can also be supervised, unsupervised, or semi-supervised. Deep supervised learning outperforms simple supervised learning, as non-linearity between input and output is not a concern. The loss function, i.e., the function that measures the error between machine output and expected output, has real numbers as parameters (called weights) which adjust during learning, and depends on the features of input. Backpropagation [66], a quite old method which uses the chain rule to calculate the gradient descent of the loss function in every layer, identifies the optimal values of weights for minimum error, thus accelerating the procedure of learning. Backpropagation is commonly used for training multi-layer artificial neural networks in various fields like computer vision [67], speech recognition [68], video games [69], and medical diagnosis [70]. As most deep learning methods construct hidden layers, they can recognize patterns missed by other techniques. Thus, their prediction is the best, compared with those of other methods, making them a very powerful tool especially when the task at hand is feature extraction from complex biological datasets. However, because of these hidden layers deep learning models are

often regarded as "black boxes" whose results are impenetrable for interpretation from a biological perspective. Though this is not always the case, great attention is currently being paid on how models understand the input features, how they determine the important ones, and how they correlate them, positively or negatively, with the output. The development of new means to transform the "black boxes" to open, "white boxes" and the interpretation of their results from a biological perspective form a new approach for the future of deep learning. Among various families of neural networks, the most commonly used are CNNs and Recurrent Neural Networks (RNNs) [71]. Their difference lies mainly in the format of the processed data, as RNNs are focused on sequential data whereas CNNs are specialized for grid values such as those comprising images. It should then come as no surprise that CNNs were deemed more suitable and were utilized in this work; as a result, henceforth we will focus solely on CNNs, as RNNs are out of the scope of this thesis.

Convolutional Neural Networks

CNNs were first introduced by LeCun et al. in 1989 [56]. Their structure is based on the organization found in neurons located in visual cortexes [72, 73]. They are easy to train, although they process multiple-array data. As a result of their simplicity and efficiency they have become the most common network architecture for image processing and analysis. They consist of convolutional transformations, as well as pooling and non-linear operations (Figure 2.2).

More specifically, their *building blocks* are:

- Convolutional layers: CNNs consist of sets of convolutional layers, i.e., in each layer unstructured filters (kernels) which create convolutional transforms are applied in the input data, x, and produce an output usually referred to as feature map, activation map, or convolved feature. Moreover, each kernel's values represent the weights. The number of kernels and their size are two hyperparameters of the CNN. Additionally, before applying the convolution, two more hyperparameters need to be decided: stride and padding. Stride is the number of shifting over pixels during the convolution, and padding refers to the number of pixels that can be considered as borders around the image. If we assume that x is the input and W the kernel, the feature map s is:

$$s(t) = (x * W)(t) + b(t)$$

  where b is the bias of the kernel in layer t.

- Pooling layers: The pooling layers are similar to convolutional layers as they are applied in each layer to reduce spatial size in the individual network layers. The input of the pooling layer is the feature map of the convolutional layer. The feature map is divided in patches and each operation is applied to each patch. If we assume that the patches $s^{nxn}_{(i,j)}$ are the nxn squares centered in (i, j) location of the input feature map, some examples of pooling operations are:

  1. Sub-sampling: $Z = \beta \sum(s^{nxn}_{ij}) + b$

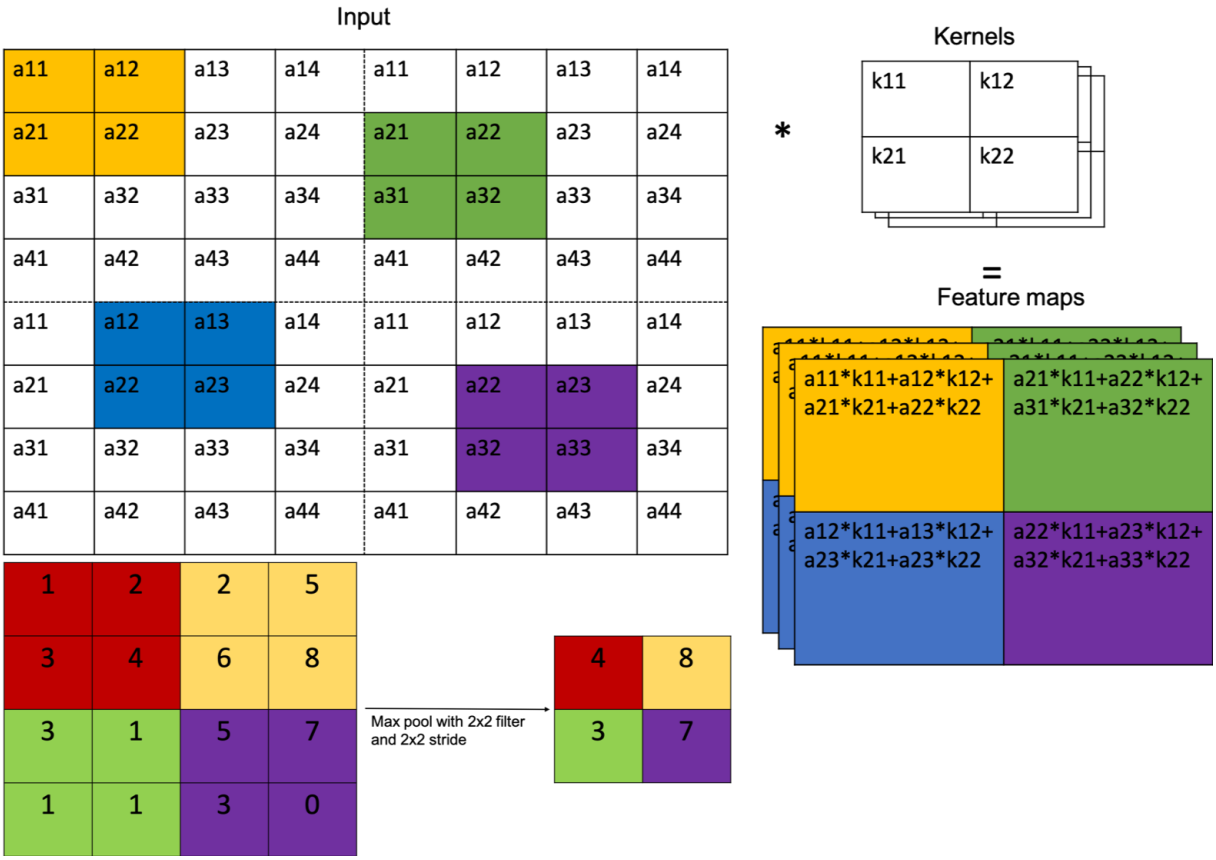

Figure 2.2: CNN operations. In a convolution layer, an input patch is multiplied by kernels, producing feature maps. Pooling layer results in the max value of the convolution layer moving to the next layer.

2. Average pooling: when $b = 0$ and $\beta = \frac{1}{nxn}$
3. Max pooling: $Z = max(s_{ij}^{nxn})$

- Activation functions: Non-linear layers in all deep learning frameworks are formed by activation functions and are combined with other layers. As they simulate the non-linear transformation from input to output, it is important to select the appropriate one for better feature extraction. Some commonly used activation functions are:

1. Sigmoid function: $\frac{1}{1+e^{-1}} \in [0, 1]$
2. Hyperbolic tangent: $\frac{1-e^{-2x}}{1-e^{-2x}} \in [-1, 1]$
3. Softmax: $\frac{e^x}{\sum e^x}$
4. Rectified linear unit (ReLU): $\max(0,\, x) \in [0, \infty]$

Training

The goal is to find the parameters that minimize the loss function of the network. The loss function measures the fit between the output and the true observations. One of the most commonly used loss function for regression tasks is cross-entropy [74]. However, a good network performs as satisfactorily on the test data as on the training data. To achieve this, regularization terms apply penalties to avoid overfitting. A recent regularization method is Dropout, that adds noise in the hidden layers of the CNN. This is achieved by multiplying each hidden layer with a certain probability [75]. Moreover, the L2 regularization method is often used. L2 adds a squared magnitude of coefficient penalty to the loss function [76].

Optimization

Optimization is an important step to achieve the least loss, i.e., to maximize the performance of the network. Before introducing some of the optimization algorithms that are commonly used, three hyperparameters of the network should be defined:

- Learning rate: determines the step size at each iteration, i.e., each pass while minimizing the loss function
- Batch size: the number of the group of examples that are used during training
- Epoch: one forward and backward pass of the full training set

Optimization algorithms

- Stochastic gradient descent (SGD): one of the most commonly used optimization algorithm. It is a variant of gradient descent but it updates the network's parameters after computation of loss on each example. [77, 78]
- Momentum: accelerates the convergence to minimum by introducing one more hyperparameter, which is considered velocity [79].

Another simple way to increase the size of the dataset and to avoid overfitting of the model is by data augmentation. Image augmentation usually consists of (i) random translation, (ii) rotation, and (iii) deformation. [80] Increasing the diversity across the dataset helps generalize the model toward unseen features. More specifically, for the micro-CT images of ants we used crop, resize, and histogram equalization to focus on the head area of interest and to enhance the contrast of the image.

CNN architectures for medical image segmentation

Over the past decade the number of highly complex datasets generated by biological experiments has increased dramatically. For the raw biological data to provide answers that lead to better understanding of biological systems they need to be analyzed and interpreted with tools and techniques specifically tailored for complex systems. Image segmentation has been one of the most useful tasks in medical image analysis because it can provide valuable information in various applications. Depending on the application and the data, many CNN architectures have been (and are still being) developed. Regardless of architecture, hyperparameter (i.e., number of layers, dropout rate, regularization coefficient) tuning is an important task required for all deep learning methods. Even though the weights are determined by back propagation, hyperparameters are often empirically chosen. This can lead to potential limitations of the models, as small changes in hyperparameter values can lead to drastic changes in a network's performance. Some of the most well-known architectures are AlexNet [81], ResNet [82] and GoogleLeNet [83]. Next, the CNN architectures most commonly used in medical imaging are briefly summarized. For a more comprehensive review, the interested reader may wish to refer to [84].

Fully Convolutional Network (FCN)

First introduced in 2015 by Long et al. [85], the Fully Convolutional Network (FCN) is a CNN architecture that contains convolutional layers and fully connected layers. As its connections are local, the network can segment images of variable sizes. Moreover, using convolutional layers in addition to fully connected layers, it is faster to train and requires fewer parameters. The FCN consists of two parts: convolution (encoder) and deconvolution (decoder). During convolution it outputs a heatmap similar to that of classification nets, which is subsequently used as input for decoder. In the latter part, it obtains information about the location of different classes using unpooling by targeting maximum activations, thus resulting in a high-resolution segmentation map.

U-Net

One superior extension of FCNs is U-Net, first applied for biological microscopy image segmentation [86]. It consists of a conventional FCN but also includes skip connections between the encoder and decoder parts, transferring information of locality. The encoder part (or down-sampling), consists of 3×3 convolutions. The decoder part (or up-sampling, deconvolution) uses up-convolution increasing the image resolution. As a last step, a 1×1 convolution is applied to generate the final segmentation map. Since there is no fully connected layer, the parameters of the network are reduced, making

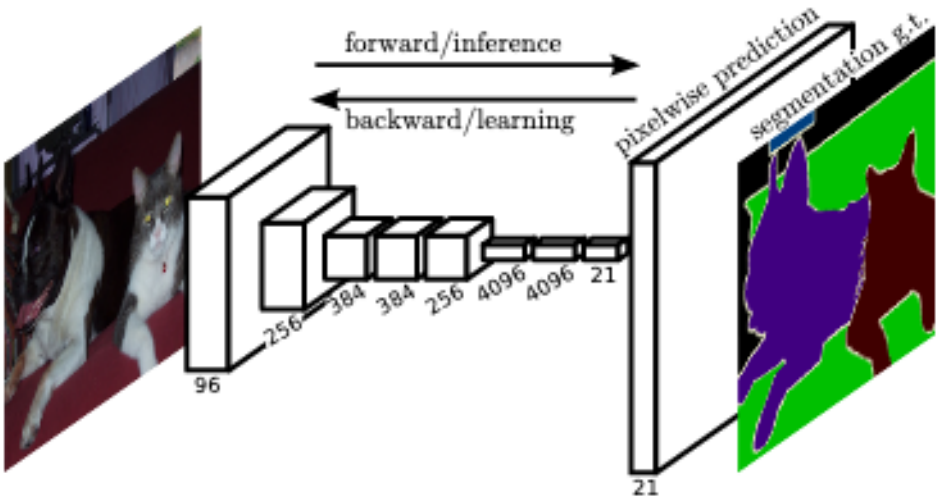

Figure 2.3: Fully convolutional networks. FCNs can be efficiently trained to perform pixelwise predictions such as image segmentation. Reproduced from [85]

the network easy to train with only a small dataset. The performance of the network in medical image segmentation has been outstanding, winning two ISBI challenges in 2015 [87], and leading to the development of more architectures built on top of U-Net, such as the 3D U-Net (which processes volumes) [88] and the V-Net (for 3D MRI volumes) [89].

Other methods

The Mask R-CNN model, winner of COCO Challenges in 2017 for object detection [90], is an extension of Faster R-CNN [91] architecture. Used for a semantic segmentation and classification, Mask R-CNN is a combination of an FCN (such as ResNet [82]) that extracts features and a Region Proposal Network (RPN) that proposes object bounding boxes, i.e., Regions of Interest (ROI) where the object might be located within the image. In both Faster and Mask R-CNN their RPN part uses a specific layer type called ROIPool to put bounding boxes around the potential objects. Later, the FCN part classifies these objects to the corresponding classes. As it combines two different networks, it results in classification and segmentation of the image, making it ideal for complex images.

Since the application of CNNs in medical imaging segmentation becomes more popular, more architectures are currently being developed such as the Feature Pyramid Network (FPN) [92], DeepLabv3 [93] and Path Aggregation Network (PaNet) [90], to name some. Moreover, combining CNNs and RNNs is an alternative method for achieving higher quality classification.

Methods employed in the current project

For the current PhD project, we opted for U-Net [86] as the most suitable CNN architecture. As the images have diffused edges between the inner parts of the head, and the annotated dataset was small compared with other biological datasets (refer

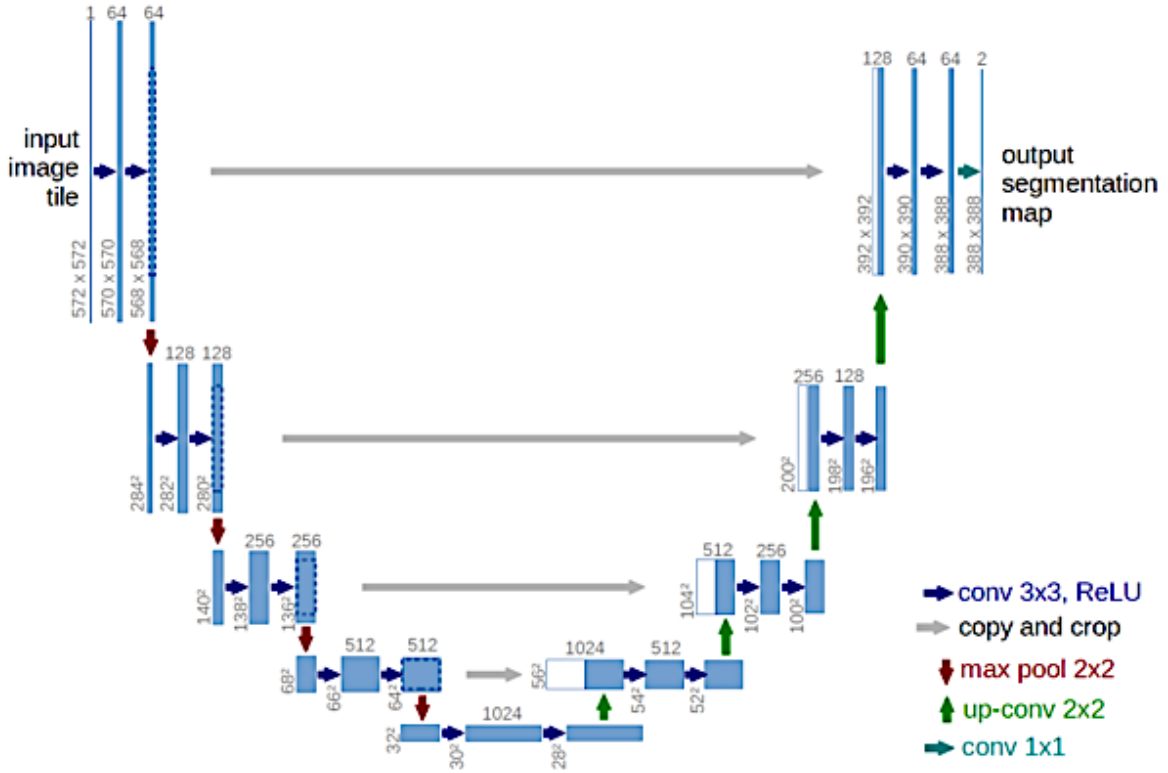

Figure 2.4: U-Net architecture. U-Net consists of two parts: encoder and decoder, with skip connections transferring locality information from encoder layers to the decoder layers. Reproduced from [86].

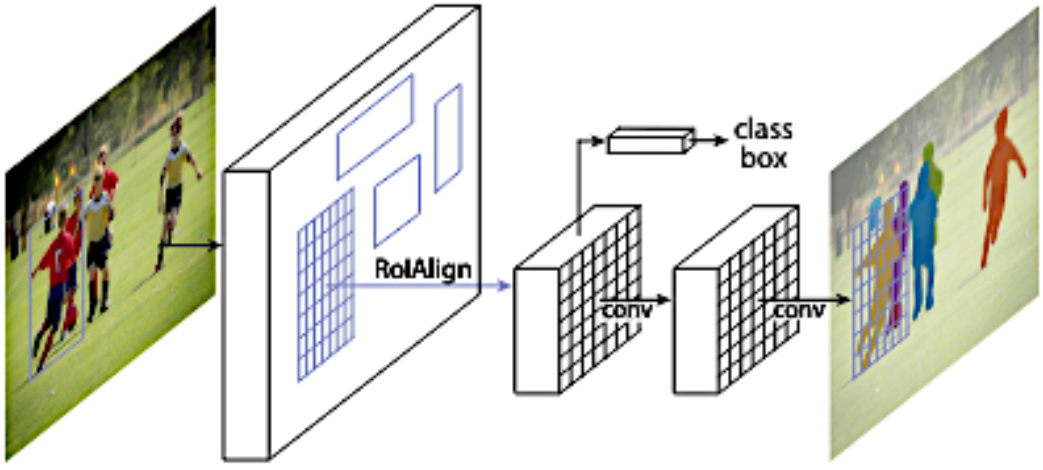

Figure 2.5: Mask R-CNN architecture. Mask R-CNN uses a convolutional network, usually ResNet, to extract features from the input, and then process them using a PRN and a mask classifier to result in both object detection and segmentation of the image. Reproduced from [90].

to paragraph 2.5.7 for more details), U-Net was the ideal architecture since it kept information about locality (thus, the ages where identified without loss) and moreover could produce excellent results even with limited datasets. Specific details about its implementation for the automated segmentation of micro-CT images of ants with their brain as the target area are elaborated in Chapter 5. Moving to full head segmentation as a future direction stemming from this study, Mask R-CNN [90] would make a great candidate since it combines segmentation and classification, both needed for multi-organ segmentation (even by non-experts) in ant biology.

2.5.5 Deep learning frameworks and libraries

With an increasing number of deep learning networks continuously being developed, the need for dedicated software frameworks and libraries has become more apparent during recent years. As a response to this emergent need their number has also grown remarkably. Below, the most well-known and commonly used open source frameworks and libraries are listed:

- **TensorFlow** is one of the most broadly used frameworks, developed by Google Brain based on Python language [94].
- **Caffe**, developed by Berkeley Vision and Learning Center, is one of the first frameworks, written in C++ with Python and MATLAB bindings [95].
- **Pytorch**, based on Torch framework, is an extension for Python language developed by Facebook AI [96].
- **Keras** is one of the most developed and used open-source libraries. It is written in Python and is mostly used in combination with different frameworks such as TensorFlow [97].

For this study, TensorFlow and Keras were the two frameworks utilized for building our segmentation CNN, and also for the pre-possessing and post-processing steps. The language used was Python in Jupyter environment.

2.5.6 Evaluation methods

To assess whether a CNN is accurate in segmenting a given set of images, some form(s) of reliable evaluation criteria are necessary. Since the input images may be numerous, complex, and dissimilar from one another, there is no one-size-fits-all solution to this task, and more than one evaluation metrics are routinely employed to validate or detect potential flaws in the segmentation model. Two of the most intuitive and widely used evaluation metrics for segmentation (also used in the current study, see Chapter 5) are the Intersection-over-Union (IoU, also known as the Jaccard Index) and the Dice Coefficient (F1 Score) [98].

IoU is defined as the area of overlap between the predicted segmentation and the ground truth divided by the area of union between them. Its values range between 0-100%, where zero indicates no overlap and 100 indicates full overlap.

$$IoU = \frac{TP}{TP + FN + FP}$$

F1 Score, on the other hand, is defined as twice the area of overlap between the predicted segmentation and the ground truth divided by the sum total of the areas (in number of pixels) of both images:

$$F1 = \frac{2TP}{2TP + FN + FP}$$

where:

- True Positive (TP): Correct annotations
- True Negative (TN): Correct not annotated
- False Negative (FN): Missing annotations
- False Positive (FP): Incorrect annotations

2.5.7 Deep learning databases

Many datasets of images (such as MNIST [99], CIFAR10 [100], COCO [101] and ImageNet [102]) have been developed over the past decades, giving scientists the opportunity to test their learning methods. Around the mid-2000s deep learning methods were applied on the MNIST database digital image classification problem [103] showing better results than the hitherto leading SVMs, kept improving [104], and moved from the digital image classification problem to that of natural images. The main breakthrough came with the application of Deep CNNs in the ImageNet dataset in 2012 [81], even though the concept of CNNs had been developed many years earlier [9]. Moreover, the increasing amount of data led to the creation of huge voxel-based datasets that are used in volumetric models. Several publicly available biological datasets are widely used for biomedical image segmentation, such as Brain Tumor Segmentation (BRATS) [105], ischemic stroke lesion segmentation (ISLES) [106], and Fetal Tissue Annotation Dataset [107], to name a few. Yet, to the best of our knowledge no database for ecological data exists, since prior to this work no other project was focused on segmenting inner parts in micro-CT images of insects. A major obstacle for the use of micro-CT lies in the analysis of these data, since even with powerful state-of-the-art hardware and the corresponding software packages data exploration is extremely tedious work.

The invention of new medical imaging methods marked the beginning of more morphological and anatomical research on microscopically small animals. It should be emphasized that most insect images have low contrast and segments with the same pixel intensity which are separated by diffused edges. As a result, the most common segmentation method to date is by manual processing, which is extremely time-consuming, with the existing software that is used for the segmentation of micro-CT images using statistical methods to semi-automate the process [108].

As a way to overcome this obstacle the networks are often pre-trained on different but related datasets; their weights are transferred and they are fine-tuned while the

classifier is retained. Though transfer learning has shown good results in the past few years [109, 110], it is still an active area that needs to be explored. Another way to generate data that is becoming dominant in the deep learning field (which can be successfully extended to the multi-omics datasets) is by using Generative Adversarial Networks (GANs) [111], which create synthesized datasets similar to the training ones. GANs are a deep learning model that has opened the way to new approaches for the required large amounts of manually segmented data.

2.5.8 Deep Learning in medical imaging

One of the first applications of deep learning in medical image analysis was image classification [112]. Through image classification one gives images as input and expects a diagnostic variable as output, such as whether a disease is present or not. The most commonly used networks for classification problems are CNNs. During recent years, networks have reached such a great accuracy that they are challenging even human experts. Among the most researched classification problems is the lesion classification and lesion detection [113]. Additionally, there exist a great number of studies in organ localization and segmentation, cell segmentation, protein structure, and gene sequencing [114].

Nevertheless, practically all existing studies are focused on CT images of human parts and organs, with only a handful of studies applying deep learning methods on animal and insect datasets for their posture quantification and behavior. Thus, no previous studies on segmentation of micro-CT images of insects or small animals exist, making this PhD study unique and the produced dataset presented here first of its kind.

Chapter 3

Creation of Manually Annotated Dataset

"Data is food for AI"
Andrew Ng

3.1 Introduction

Nowadays information is being constantly created at unprecedented, staggering rates. In fact, it was recently estimated that every two days the production of new data reaches more than five exabytes; to put this into perspective, it corresponds roughly to the information humanity generated between the beginning of time and 2003! However, to fill the *"black hole between data and knowledge"* in what came to be called information hierarchy (or DIKW pyramid, the acronym standing for data, information, knowledge, and wisdom), data need to be organized, labeled, and given shape [115]. Creating structured databases of images has been the focus of many studies, since over the past decade the rise of large-scale datasets has been the driving force for the development of the field of computer vision.

The continuous emergence of new deep learning methods that show great performance in image classification and segmentation necessitates the creation of more and better datasets in order to push forward their development. ImageNet is famously the first such dataset to enable breakthroughs in object classification [102], paving the way for numerous other available datasets of objects, faces, animals, and sceneries, with millions of annotated instances in hundreds of images [101]. With respect to biomedical imaging, digitalization of MRI data has led to an explosion in biomedical data where deep learning techniques have been successfully applied for image segmentation and classification, enabling faster and more accurate diagnosis of almost all types of cancer [116]. Moreover, new segmentation architectures have recently been developed based on tissue, organ, and nuclei image datasets [117].

However, this is not necessarily the case when it comes to ecological data. To the best of our knowledge, there is no large-scale dataset available of micro-CT images of inner parts of any insect species. Even though there exist platforms that enable

the semi-automated segmentation of inner parts of micro-CT images, typically the annotated data are not publicly accessible but restricted to the actual user. As a result, the first step of this study involved the development of a new dataset of annotated micro-CT images which was used in both research projects discussed in Chapters 4 and 5. The dataset consists of more than 50,000 annotated micro-CT images of ant heads, created from 76 different ant scans, each from a different species, making it the largest dataset of micro-CT images of insects available to date.

Selecting a wide spectrum of diverse species across their phylogeny was a necessary prerequisite for performing an in-depth comparative study, following ecological traits, and testing the social brain hypothesis in ants (Chapter 4). On the other hand, the great variety in specimen morphologies posed a great challenge regarding the consistency of the acquired images, as it required tedious manual or semi-automated pre-processing and post-processing. How this challenge was addressed to produce a database of high-quality annotated images is the subject of the current Chapter.

3.2 Construction

3.2.1 Specimen selection

One worker from each of the 76 different species was collected from various locations across the USA and Asia. The species were chosen to maximize phylogenetic coverage. Moreover, the species were selected as pairs from the same genus but with different colony size (small or large) allowing us to isolate social complexity as the primary differentiating parameter between the species in each pair. All specimens were preserved in 97% ethanol solution to avoid shrinkage of their inner parts.

3.2.2 Collection/Scanning

The collection of a highly accurate micro-CT based dataset relies on the expertise and meticulousness of the users during manual segmentation; to this end, acquiring experience with practice samples in advance proved beneficial. Eventually, all available ant specimens were scanned using ZEISS Xradia 510 Versa 3D X-ray microscope and ZEISS Scout and Scan Control System software (version 10.7.2936) after at least two weeks of staining in iodine solution. 3D scans were reconstructed using XMReconstructor (version 10.7.2936) and saved in DICOM file format (default settings; USHORT 16-bit output data type). Scanner parameters were selected based on the size and condition of each specimen. Table 3.1 shows the genera and species of used ants as well as the number of slices and voxel sizes of their respective scans. Micro-CT scanning methodology and details on its specific implementation for the current project can be found in Chapters 2 and 4, respectively.

Table 3.1: Taxon names of all species that were used for this study. brain vol., body vol., and eye size columns show the brain and body volumes, and the size of eye area of each specimen, calculated based on their micro-CT scans. Colony size data were provided by their collectors. The last four columns indicate the number of slices per direction that were used from each specimen's scan and the corresponding voxel size of the scan.

Taxon code	brain vol. (μm^3)	body vol. (μm^3)	eye area (μm^2)	colony size	xy slices	xz slices	yz slices	voxel size
Acanthomyrmex_glabfemoralis	36648258.71	1739308306	25882.02050	small	300	150	100	1.84105
Acromyrmex_versicolor	78176936.82	7772172863	94817.08320	large	130	200	330	3.21392
Aenictus_paradentatus	35026151.08	1353925732	0.00000	large	250	200	140	1.09831
Anoplolepis_gracilipes	32152348.62	794729496.6	48087.31500	large	370	220	180	1.00562
Atta_texana	36983510.58	1919327575	38930.04660	large	250	300	200	1.51302
Camponotus_nearcticus	96305181.09	3997122934	107480.46790	small	500	300	500	1.62596
Camponotus_hyatti	62220147.58	5471886110	105785.61600	small	300	250	230	1.84968
Camponotus_modoc	177584966.1	1.67E+10	199156.07690	large	200	400	300	2.53308
Camponotus_vicinus	155978090.9	1.34E+10	183008.71120	large	350	350	200	2.25045
Carebara_affinis	14306858.12	261688278	1865.70000	large	300	250	330	0.914802
Carebara_atoma	2045403.622	21237672.11	149.08720	small	300	200	200	0.51953
Carebara_diversa	10344835.94	425755012.7	11975.04000	large	270	100	130	0.990045
Cephalotes_atratus	162318862.1	1.30E+10	111285.37360	large	200	450	450	2.18409
Cephalotes_minutus	41622884.79	1486331971	96651.81400	small	400	250	400	1.38875
Crematogaster_lineolata	29154497.66	758527756.4	25328.90250	small	280	400	350	1.06144
Crematogaster_pinicola	24970928.93	929717003.6	32894.45260	large	320	400	430	1.01303
Daceton_armigerum	167114492.9	1.96E+10	373322.37940	large		200	250	3.52417
Dolichoderus_pustulatus	28379616.07	1076489651	36940.37120	small	400	350	420	0.93929
Dorylus_kohli	28876169.1	648968585.8	0.00000	large	400	400	400	1.05723
Dorymyrmex_insanus_small	29148510.25	836705696.3	68450.35860	small	400	300	400	1.12528
Dorymyrmex_insanus_large	23057672.69	974085009.5	41134.84830	large	250	250	250	1.30802
Eciton_burchellii	108266654.1	4853084680	21821.31210	large	300	300	300	2.67454
Formica_pallidefulva	82405544.62	7572850857	143386.76100	small	300	400	200	1.82821
Formica_exsectoides	79959436.59	4945148572	143716.49600	large	300	200	400	2.04899
Formica_gnava	106269566.9	4277058266	142072.02150		250	350	380	1.62671
Formica_neogagates	54211673.18	1671037590	69960.76790	small	150	50	150	2.43199
Formica_obscuripes	139845748.6	1.05E+10	190131.30800	large	320	300	280	2.3236
Formica_polyctena	120522711.8	5524471586	140312.98710	large	450	300	400	1.79283
Formica_rufa	79610745.29	5857953172	171070.48200	large	230	450	400	1.52425

Gesomyrmex_howardi	33522278.78	648618857.7	102600.39630	small	430	400	200	1.03468
Gnamptogenys_sp	61245456.78	2422913708	84902.96360	small	300			1.94183
Labidus_praedator	44583433.33	1090357104	2107.00000	large	300	300	300	1.38717
Lasius_fuliginosus	66816296.05	3523012974	72445.41590	large	400	200	150	2.18947
Leptogenys_peuqueti	60923038.8	3569816364	75353.09600	small	250	100	100	3.0069
Leptomyrmex_darlingtoni	96878058.05	5985781446	104531.48950	small	430	230	300	1.55856
Linepithema_humile	18471779.96	341297591.3	25059.07040	large	430	250	300	0.892673
Liometopum_apiculatum	60706911.91	5184338441	53831.02670	large	200	350	330	1.66387
Monomorium_floricola	4954962.296	52479065.71	3201.40500	small	250	100	200	0.613733
Monomorium_pharaonis	7835526.897	209168110	7272.11800	large	350	350	380	0.673107
Myrmelachista_ramulorum	13654481.02	653636138.7	13169.13180	large				1.35418
Mystrium_camillae	26484392.42	963663559.8	0.00000	small	220			1.79928
Neivamyrmex_nigrescens	39905346.27	1824483127	6049.26980	large				1.55809
Nylanderia_parvula	12280979.82	247021584.5	13950.84320	small				1.35025
Ooceraea_biroi	9388855.466	230714491.7	0.00000	small	300	200	200	0.870043
Orectognathus_versicolor	29870071.36	685552517.5	42488.15400	small	350	150	200	1.35041
Parasyscia_cribrinobis	25619325.57	1141417199	34367.27480	small	310	310	350	1.15305
Patagonomyrmex_angustus	37148238.99	1690960005	45030.36640	small	250	350	350	1.38728
Pheidole_bicarinata	12317993.28	307433835.8	10923.31500	small	460			0.783321
Pheidole_rhea	26912463.58	901055690.2	22703.82310	large	300	250	300	1.11148
Pogonomyrmex_badius	72786171.02	9177575546	108845.15850	large	250	150	300	2.24449
Pogonomyrmex_brevispinosus	45244375.72	4167972840	92423.27640	small	250	250	250	1.941
Pogonomyrmex_desertorum	87526556.69	4169870326	88771.96590	small	250	150	250	2.19439
Pogonomyrmex_magnacanthus	47707336.38	3642140540	148115.29060	small	150	200	230	1.82453
Pogonomyrmex_occidentalis	64005886.94	8041478490	110467.31010	medium				2.59613
Pogonomyrmex_pima	36145908.05	1298855487	36843.43320	small	350	350	200	1.36688
Pogonomyrmex_rugosus	77760055.2	9093456012	145375.79520	large	200	300	350	2.30809
Pogonomyrmex_schmitti	25125098.41	1249934595	28252.07100	small	300	300	300	1.34886
Pogonomyrmex_subdentatus	91962754.7	4578092547	102300.59400	small				1.92925
Polyrhachis_bihamata	343336805.1	3.40E+10	265188.82240	small				3.81426
Pristomyrmex_profundus	1.13E+07	209607952.7	10908.00000	small	480			0.742557
Pseudomyrmex_ejectus	28958446.55	753676719	119739.11040	small	350	300	330	1.11057
Pseudomyrmex_ferrugineus	63707531.22	1895670612	148780.50480	large	400			1.29843
Pseudomyrmex_gracilis	170939920.1	7857841605	555020.79240	small			300	1.82826
Pseudomyrmex_triplarinus	80923055.22	3175493676	187264.92960	large	300	300		1.61201
Solenopsis_geminata	20595232.03	1192929865	15844.54440	large	250	250	270	1.34886

Stenamma_heathi	3.29E+07	817607982	7181.16000	small	390	250	350	1.27715
Stigmatomma_sp	80653212.46	7646731977	40198.72530	small		230		2.44965
Syscia_augustae	19092311.06	555952712.7	0.00000	small	420			0.787728
Temnothorax_curvispinosus	11182082.53	170521049.5	12687.87520	small	450	350	400	0.656251
Tetramorium_hispidum	26775740.47	1119204030	49727.32080	small	250	200	150	1.31344
Tetramorium_pacificum	21648982	809933620.3	26384.89600	small	430			1.37859
Tetraponera_sp	34786266.18	804806566.6	114633.83640	small	310			1.12522
Trachymyrmex_carinatus	38696589.11	1800962450	36085.86170	small	100	100	220	2.04899
Veromessor_andrei	87046838.15	6786375486	93062.63040	medium	200	200	200	2.25027
Veromessor_lariversi	42534464.98	2554671888	102225.73050	small	250	250	100	1.70412
Veromessor_pergandei	41428810.18	1144110538	57887.27340	large		100		1.67449
Veromessor_smithi	54155021.2	3497873175	145753.54650	small	200	250	200	1.99029
Brachymyrmex_depilis	5307189.444	36164609.13		small	380	380	405	0.621502
Cyphomyrmex_flavidus	36259857.87	543562964.6		small	350	300	300	1.02553
Pseudomyrmex_veneficus	56915813.79	1001114276		large	370			1.2192
Dolichoderus.mariae	20788310.74	609267428.1	4564.98250	large	300	200	300	1.20906
Tetramorium.immigrans	33973982.67	1085211149	44649.70800	large	290	240	200	1.35523
Myrmelachista_nodigera	6386479.863	missing body scan	144269.47250	small	330			0.708583

3.2.3 Semi-automated segmentation

Next, the brain area in each scan was segmented using the seed-based watershed method [59] within the Amira [118] software environment. For better results three or four 2D slices of each plane were manually segmented before applying the method; however, in most cases, the software overpredicted brain areas. As the prediction depends on the number and the suitable selection of seed positions, more precise manual processing of more slices in all three planes resulted in more accurate segmentation. Even then, however, the nerves connected to the brain were still misidentified as brain. Exemplar results of semi-automated segmentation are shown in Figure 3.1(A-C), clearly indicating falsely identified parts of the head as well as rough borders of selected areas.

3.2.4 Manual processing/cleaning

As a result, it became obvious that thorough cleaning of the data was imperative. Smoothing and deleting of wrongly annotated pixels was performed manually, pixel by pixel, in each slice, for each plane, leading to more accurate brain segmentation. Results of this manual processing are shown in Figure 3.1 (D-F), for comparison with those of semi-automated segmentation. To readily identify the regions which were removed during this process, Figure 3.1(G-I) shows an ensuing 3D brain reconstruction (seen along three different directions) before and after smoothing; cleared-out regions are shown in blue, whereas correctly identified brain areas are depicted red. The resultant brain volume of each specimen was calculated with Amira, also shown in Table 3.1. Moreover, after reconstructing the surface of each specimen, each body was measured using its voxel size and the eye area was measured using Amira software.

3.2.5 Data processing and consistency

To be able to train a CNN efficiently with as limited a dataset as possible, the consistency of the data is of paramount importance. Following the brain segmentation of each specimen, we extracted the 2D tiff raw images of each head as well as their binary masks using an in-house built Python script in Jupyter notebook. Every scan was unique and the species were purposefully selected to show great diversity in head and brain volume; as a result, the numbers of extracted slices were vastly different and had to be input manually. Moreover, to significantly boost the number of images and to maintain the information on the locality of the brain, we extracted images from all three plane normals by rotating the scan along with its labels. The number of brain images per specimen for each plane normals is also tabulated in Table 3.2.2. In every scan, the void area of the vial containing the specimen generated totally empty images. These images were manually removed after regular dataset quality checks. To enhance the contrast of each image and to remove areas of no interest, the 'imaug' library in Python was used, performing histogram equalization and zoom. By enhancing the contrast, the texture and inner parts of the dense brain area were easily recognized. The resulting images were of 520 $\times$ 520 pixels. Figure 3.2 shows the pixel histograms of the images of one specimen, before and after equalization, clearly indicating the homogenization of the image. As a final step, all images were reevaluated and compared

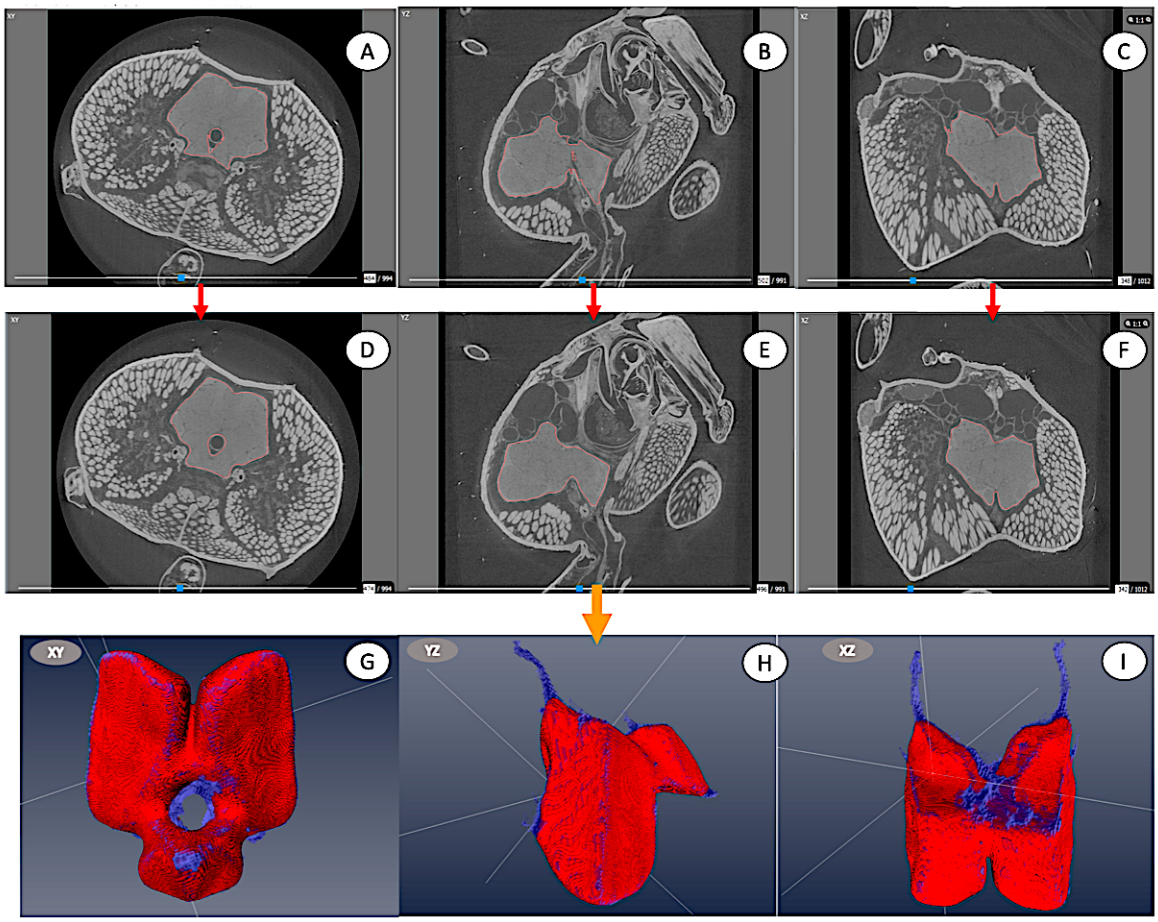

Figure 3.1: Segmentation of brain area of *Atta texana* micro-CT images. Panels A, B, and C show segmentation results with the watershed method in Amira, along three perpendicular directions. The red outline in each figure indicates the identified brain area. Panels D, E, and F show the corresponding segmentation results after smoothing and "clearing". In the bottom row, panels G, H, and I show the resultant 3D images of the segmented brain, with areas depicted red indicating the brain after smoothing, and blue indicating the pixels that were removed.

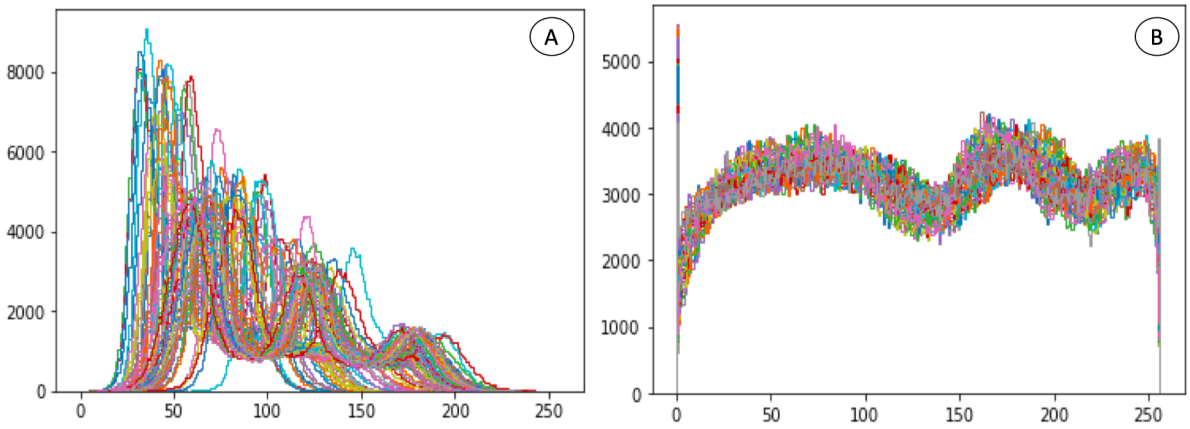

Figure 3.2: Exemplar Pixel Histograms before and after equalization for *Atta texana* ant. Different colors correspond to different images. The image processing resulted in a more homogeneous histogram, which enhanced their contrast.

with their masks, and named according to the specimen's species.

3.3 Applications

The created dataset was used as training and testing set of U-Net to automate the segmentation process. More details about the network and the project are given in Chapter 5. Moreover, the same dataset was used for the morphological study of the ant brain, which is detailed in Chapter 4.

As pre-training has become a first step in most deep learning segmentation methods and the use of a different but related network is a requirement, an extended dataset of micro-CT images of ants can become a key aspect. Moreover, our dataset can become the first step for the creation of a bigger dataset of micro-CT images of insects, in general, as digitalization of specimens becomes a fact.

Chapter 4

Ant Brain Evolution

"The brain of an ant is one of the most marvelous atoms of matter in the world . . . perhaps more so than the brain of man"
Charles Darwin

4.1 Introduction

Due to the complex relationships between the animal brain and the biology, ecology, and behavior of the individual, it is important to explore the parameters that have effects on its evolution. A particularly interesting question in ecology is how social structures and brain evolution are related. Studies focused on this question can shed light on mechanisms relevant to understanding how intelligence and cognition evolve. Ants (family Formicidae) are social insects that form rich, complex societies and have outsize ecological importance in the diverse environments in which they are found. Therefore, they are an ideal candidate taxon among other insects to investigate the connection between social structure and colony complexity, as well as the morphological evolution of the ant brain [119]. Here we investigate this question by using micro-CT imaging of ant brains to quantify the relationship between social complexity and the morphology of the brain and assess the validity of the social brain hypothesis on a representative taxon of social insect. We conduct the first comparative study of the brains of ant species which form societies of different structural complexities (also considering their phenotypic variation) using a large number of species ($n = 76$) across their phylogeny in order to test the social brain hypothesis. This analysis formed the first part of the current PhD project and involved (i) a good understanding of the morphology of ant brains; (ii) hands-on experience with micro-CT imaging (also necessary for the segmentation automation project presented in Chapter 5), and (iii) a main focus on the application of classical statistical methods and analysis (as elaborated in Chapter 2).

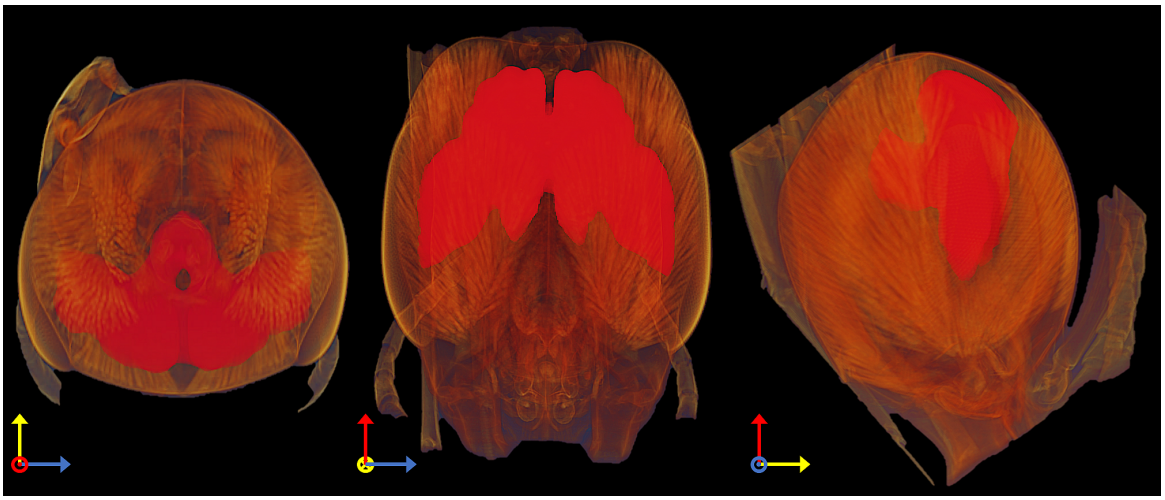

Figure 4.1: *Atta texana* head scan viewed along all directions with the segmented brain depicted red. The three views are to scale.

4.2 Methods

As a first step, we generated a unique dataset of micro-CT scans of ant workers of these 76 species, calculated their body and brain volumes (Figure 4.1) and eye sizes, and combined the morphological results with ecological data. We used these data to determine how social structure and ecological traits affect the evolution of the ant brain. More details about the construction of the dataset are provided in Chapter 3.

4.2.1 Sociality trait and phylogenetic tree

The phylogenetic tree of our specimens was constructed within the framework of a collaborative project focused on the genome analysis of the specific species (publication in preparation by *Tong et al.*) and the status of their colonies was defined to be used in the current statistical analysis. A threshold value of 3000 individuals per colony was selected to distinguish between small and large colonies, creating two sets of almost equal number of species. Both higher and lower values were also tested indicating insignificant sensitivity of the results below to the exact threshold value. As shown in Figure 4.2, species that are closely related to each other form sets of small (regular typeface) and large (bold typeface) colonies (i.e., showing similar morphology but different social structure). As an example, a pair of ants from the same genus (*Pseudomyrmex*) but from colonies of different size is shown in Figure 4.3. Moreover, our specimens were diverse regarding their eye size. Figure 4.4 shows images of 3 different ant head scans, each with different eye size. In some cases, eyes can cover a substantial part of the head, along with the presence of ocelli, whose function is akin to that of a compass, facilitating the ants' navigation [120]. Eyes and ocelli are connected with the optic lobes, the part of the brain where visual information is being processed.

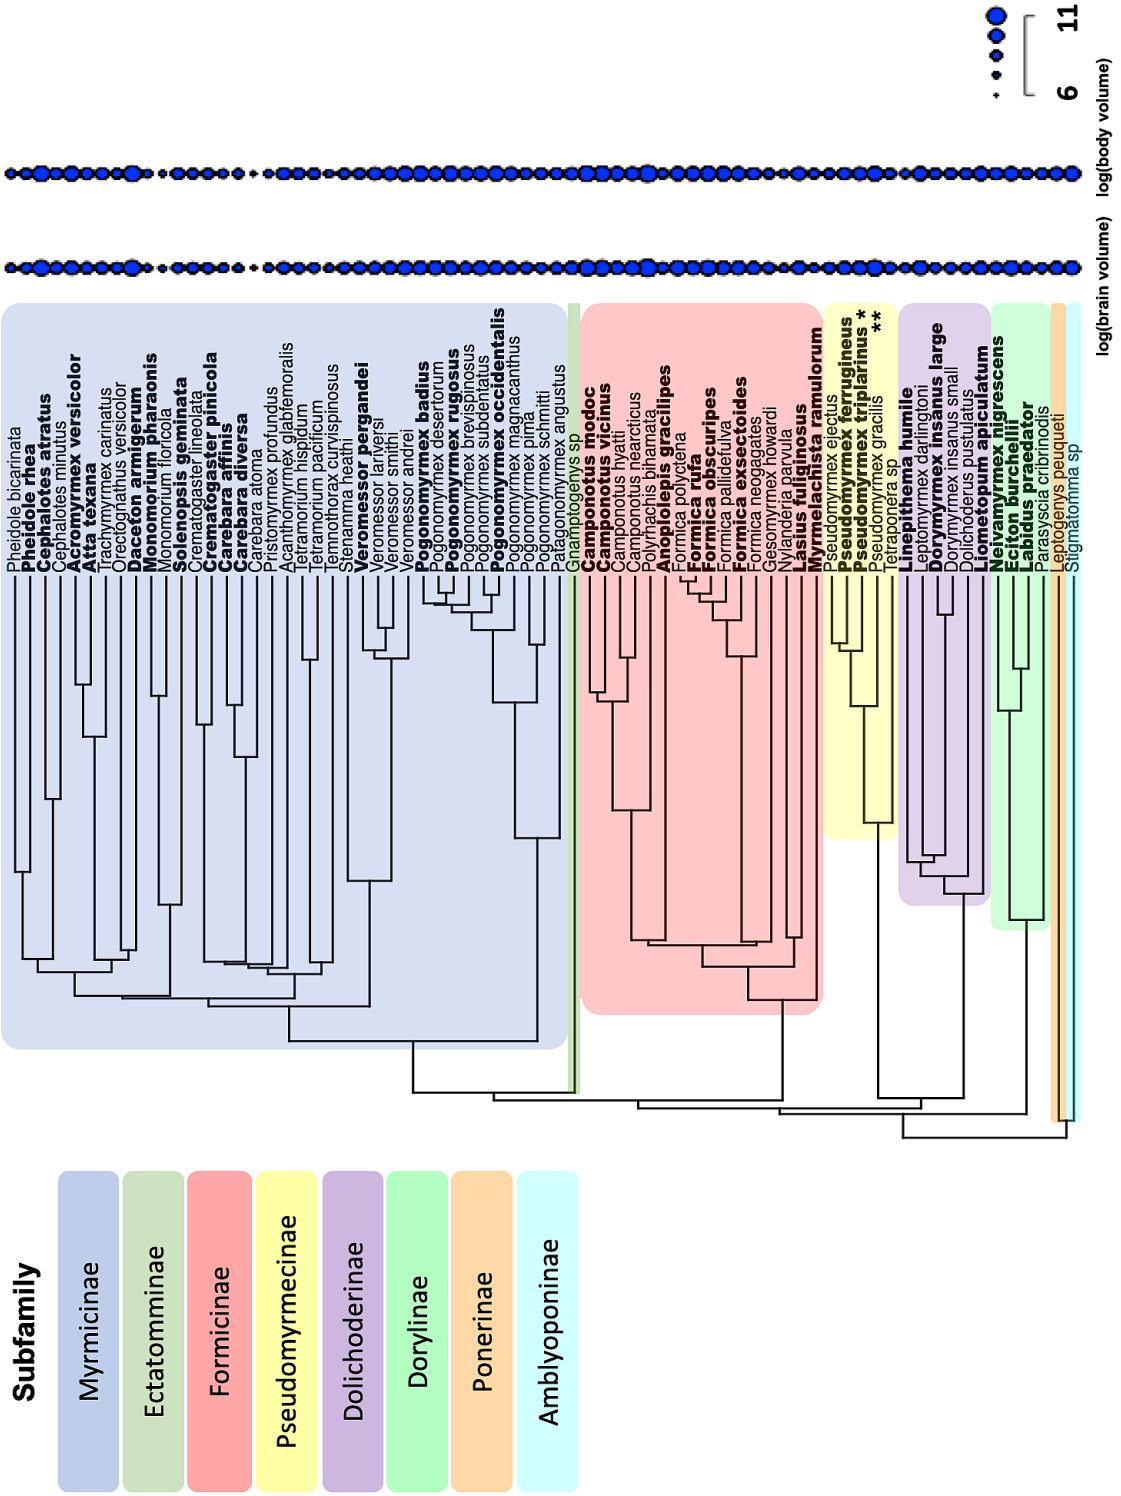

Figure 4.2: Phylogenetic tree of 76 ant species used for this study. Different colors correspond to different ant subfamilies. Species from large colonies are in bold typeface. Specimens of species *Pseudomyrmex triplarinus* and *Pseudomyrmex gracilis*, denoted with single and double asterisks, respectively, are shown in Figure 4.3.

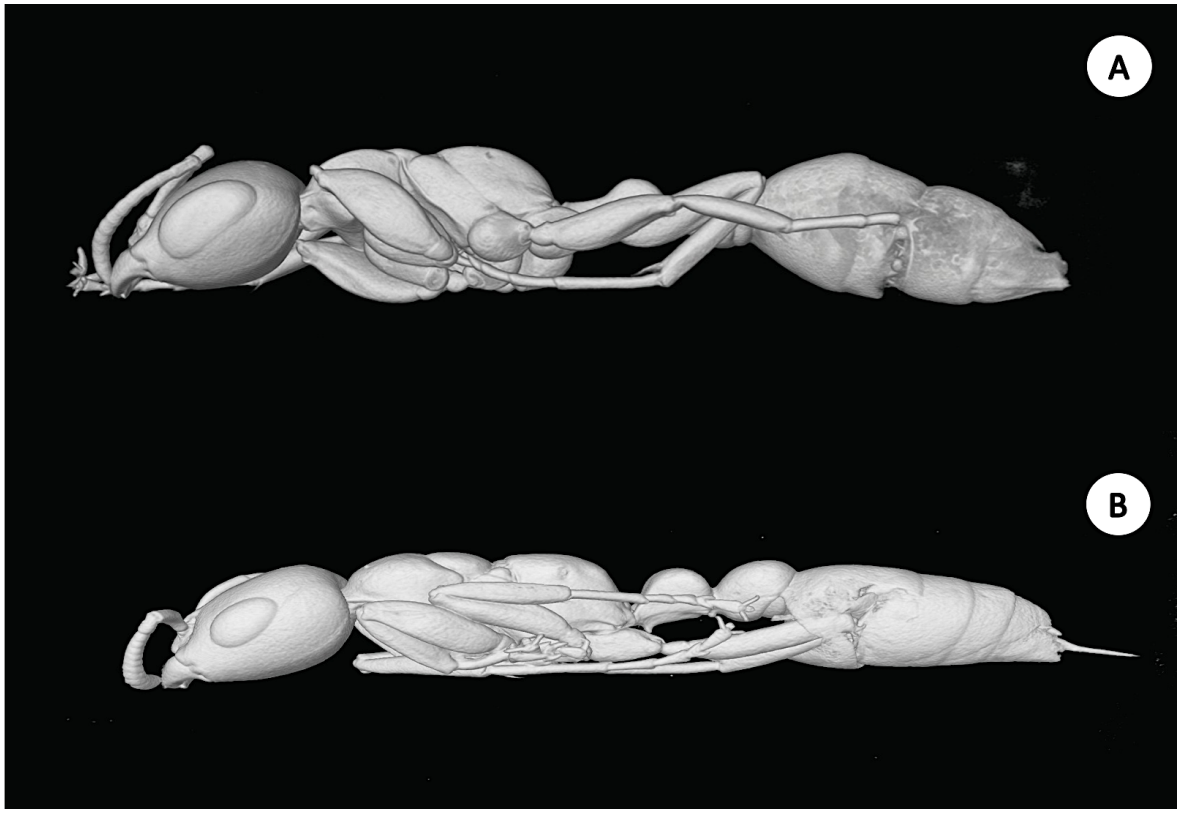

Figure 4.3: Body scans. A. *Pseudomyrmex gracilis* (small colony), and B. *Pseudomyrmex triplarinus*(large colony). Despite being congeners and sharing many similar characteristics, the species also exhibit significant differences, most remarkably in eye size.

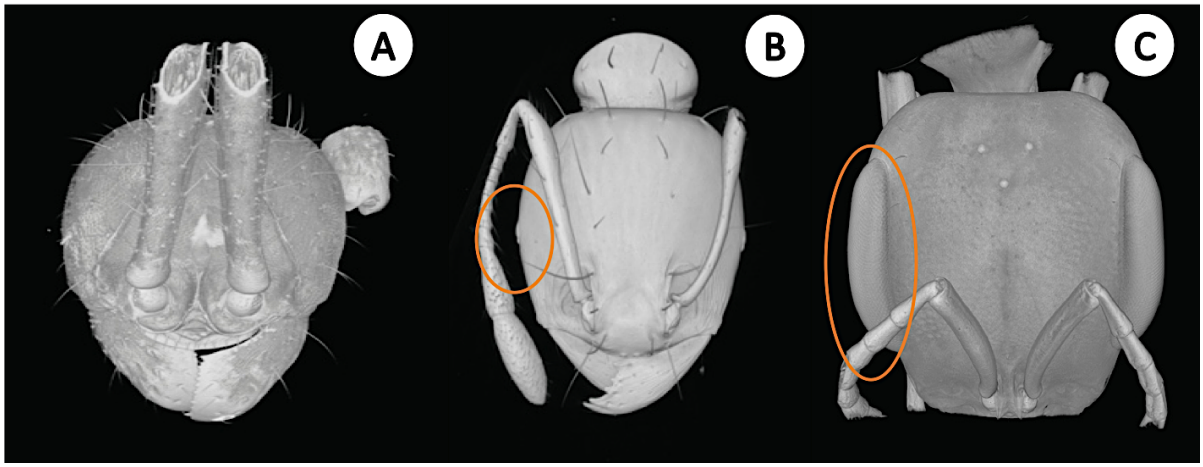

Figure 4.4: Head scan images. A. *Aenictus paradentatus*, B. *Carebara affinis*, and C. *Pseudomyrmex gracilis* ants. Heads are to scale. The three different species are exemplary cases exhibiting eyes which are: A. absent, B. small, and C. large (as indicated by the orange ovals).

Table 4.1: AIC and ML rates. AIC and ML rates for different models fitted to the residuals of the regression analysis of brain-size modifications. As the lambda model is maximizing ML and has the lower AIC, it is the best fitted model to our data.

	OU model	BM model	lambda model ($\lambda \sim 0.26$)
AIC	11.88	96.28	5.78
ML	-2.76	-46.05	0.28

4.2.2 Phylogenetic analysis/statistics

We statistically analyzed and evaluated the correlation between brain volume and body volume, eye size, and colony size to determine which of these traits affect the brain size. As a first step, we fitted three different evolutionary models, Brownian, OU, and Pagel's to the residuals of the linear regression of brain and body volume to test the effect of the phylogeny in brain changes. The model that best describes brain-size modifications (while considering the phylogenetic factors) is Pagel's λ [24, 25]. Lambda was calculated so as to maximize the likelihood and minimize the AIC value [21], as shown in Table 4.1. As a next step, we performed both linear regression and PGLS regression between the relative brain size (response variable) and the sociality and eye size traits (predictors). All statistical analyses were conducted in R (R Core Group 2021) using the 'lm' function of stats (version 3.6) package for linear least squares regression and 'pgls' function of Comparative Analyses of Phylogenetics and Evolution (CAPER) (version 1.0.1) package for PGLS regression [121]. Body volume was included in all calculations as a covariate to correct for allometric effects [27]. Finally, we examined direct and indirect correlations between all traits using the 'phylopath' package (version 1.1.2) in R [122–124].

4.3 Results

Consistent with previous studies [125], the brain size of each species was positively correlated with body size, as shown in Figure 4.5. The linear regression between log brain and log body volume shows their strong correlation ($R^2 \sim 0.9$) for both small and large colonies.

The performed OLS and PGLS regression between all four traits indicates that body volume is the strongest predictor of brain volume, but highlights two additional factors that may affect brain size: namely, eye size or the potential absence of eyes. Of the 76 species analyzed, 5 were eyeless; these species were considered extreme cases, and as such were excluded from subsequent analyses. Additionally, both linear and PGLS regression revealed no correlation between brain volume and colony size or eye size, as shown in Table 4.2. However, excluding the extreme cases of eyeless species led to a clear correlation between brain size and eye area. The results of linear regression and PGLS analysis of brain volume and eye size are shown in Table 4.3. The analysis indicates that ants with larger eyes have greater brain volume. Brain size may be positively correlated with eye size in ants because species with larger eyes have the

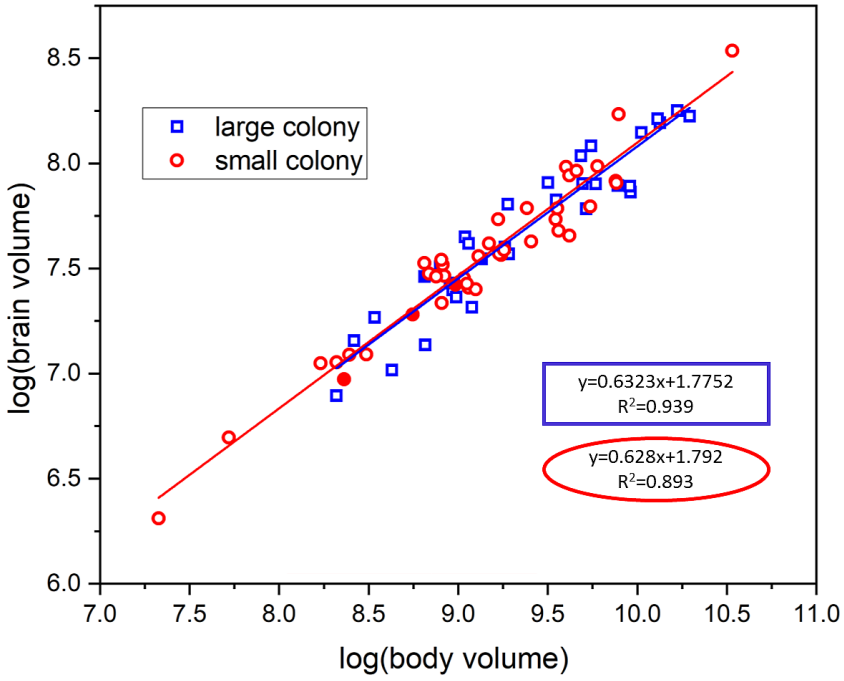

Figure 4.5: Linear correlation of log brain and log body volumes. The linear relationship between the log volumes indicates that a definite power relationship exists. Full symbols indicate eyeless species from both large and small colony sets.

Table 4.2: Results of linear and phylogenetic linear regression of brain volume, eye size and colony size for all our dataset. There is no significant correlation, as verified by the p-values.

Brain volume - Eye size - Colony Size	F-statistic	p-value for Eye/Colony size	R^2
Linear Regression	285.2	0.533/0.452	0.919
PGLS	255.8	0.427/0.2545	0.9107

Table 4.3: Results of linear and phylogenetic linear regression of brain volume and eye size for ants which have eyes. The correlation is significant, as verified by the p-value.

Brain volume - Eye size	F-statistic	p-value	R^2
Linear Regression	416.8	0.021	0.9213
PGLS	366.7	0.05	0.9115

need to process more visual information and thus require larger optic lobes and/or mushroom bodies (i.e., neuropils responsible for memory).

Next, using a least-squares linear regression we found evidence that eye size and sociality influenced brain evolution, and we investigated the connection between eye size and sociality of each species. The PGLS and linear regression analyses both show strong correlation between the eye size and sociality traits (once more when we excluded the species without eyes), as shown in Table 4.4. To sum up, small colonies were correlated with larger eyes in ant workers.

As eye size is correlated with both brain volume and sociality, we investigated the connection between brain volume and sociality using graph theory ('phylopath' package in R) that can reveal weak connections between traits (which can be overlooked when stronger traits interfere). The relationship between all traits was examined for eyeless ants, as shown in Figure 4.6. As expected, the very weak correlation between sociality and brain investment was ignored when using linear regression, as the correlation with the other traits is stronger, leading to a hardly detectable connection.

Table 4.4: Linear and phylogenetic analysis of eye size and colony size in non-eyeless ants. The results show correlation between the two traits. Log body volume was also included in the calculations.

Eye size - Colony size	F-statistic	p-value	R^2
Linear Regression	69.62	0.0237	0.6622
PGLS	55.04	0.020934	0.6069

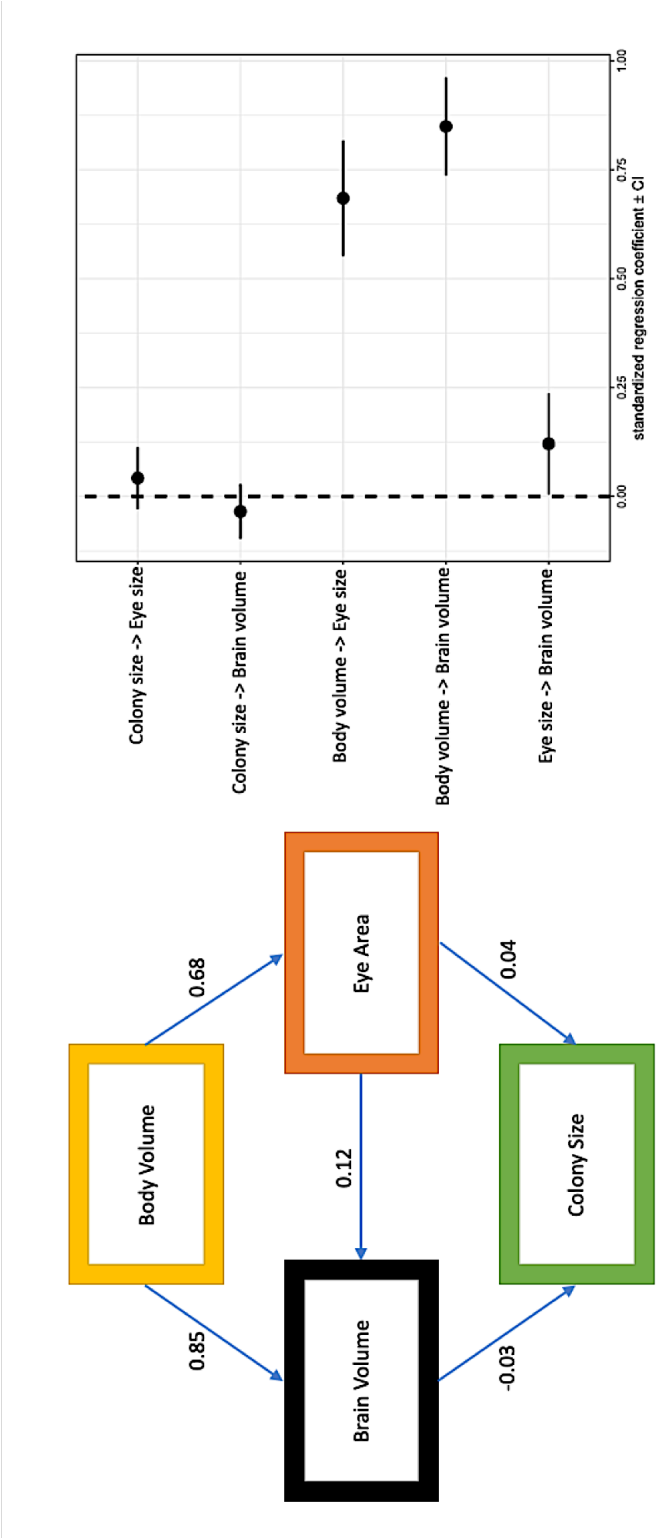

Figure 4.6: Results of phylogenetic path analysis for all traits for non-blind ants including the standardized regression coefficient ± CI. Low values of beta coefficients between colony size and brain volume show that the correlation between the two traits is weak.

4.4 Discussion

We found no support for the social brain hypothesis in ants, as there was practically no evidence that colony size is correlated with brain size. We did recover a strong brain-volume / body-volume allometry with power exponents (~ 0.6) comparable to previous studies. After correcting for body size scaling there were still strong correlations between eye area and brain volume, and between eye area and sociality.

The lack of correlation between brain volume and colony size found in this study, aligns with most previous studies on the social brain hypothesis on insects [126]. Totally social and totally non-social species should be included, however, to maximize their potential to highlight any relationships that might exist. Nevertheless, the connection between eye area and sociality is revealed as quite significant. Eye presence and size depends on many factors; a prevalent one being the interaction between individuals (e.g., for the protection of the colony from other species). Thus, the correlation between larger eyes and smaller colonies may suggest that in such colonies the superorganism (i.e., the colony) can invest in individual ants' traits to protect itself from outsiders. Of course, we cannot exclude other environmental and ecological factors, like habitat type and nocturnal foraging, as they were out of the scope of this study. The need for primary visual information processing, though, is connected with larger brain size, and is task related [127]. It is possible that ecological factors such as labor division may be related with brain development because of their connection to eye size and social structure.

The indirect weak connection of sociality and brain volume might also be a result of other ecological factors that drive the brain size evolution, whose strong connections swamp any social effects. These ecological factors need to be tested as future work. Moreover, what remains to be answered in a future study is whether the morphological adaptation is a preadaptation or an outcome of sociality.

4.5 Conclusion

Connecting ant brain investment with the social structure of the colony has frequently been investigated with consideration of both the ant nervous system and foraging and diet routines [126, 128]. However, most past studies have been limited to only a few species, and as a result it is difficult to make a general conclusion about the drivers of brain evolution across the ant phylogeny. In the current study we explored 76 different ant species spread across this phylogeny and analyzed this connection with the largest such dataset to date. In our study we found that the eye size of ant workers is correlated with the social structure of their colony. Moreover, their brain size is affected by the changes in eye size or presence. However, the connection between sociality and brain size remains unclear as the analysis showed only a weak correlation between them. Naturally, there are topics that warrant further investigation, including additional relevant physiological traits such as brain density and nervous system structure, as well as exploration of changes in neuropils, especially those which process visual and olfactory information. Nevertheless, this study helps provide a foundation for investigating variation of ant brain size. Our findings are important not only for evolutionary

biology, but also because they help propose a preliminary (if yet inconclusive) answer to the social brain hypothesis for ants.

Chapter 5

Automated Segmentation of Micro-CT Images of Ants

"Easy things are hard (especially in vision)"
Artificial Intelligence: A guide for thinking humans

5.1 Introduction

As micro-CT imaging is becoming readily accessible for morphological biology studies, the amount of generated data is piling higher than ever before at an overwhelming rate. However, in order to have significant ecological discoveries, 3D images of animals produced by micro-CT scanning need to be processed and analyzed efficiently. One of the most time-consuming processes along this line of investigation is the segmentation of the inner parts of the scanned specimens. To overcome this hindrance, automation of the process by deep learning techniques is proposed as an ideal candidate for an easy and efficient solution.

In this Chapter, the general segmentation problem of micro-CT images of insects is introduced, and a proposed pipeline for its automation using ants as model organisms is described. The following sections are currently deposited as a pre-print in the BioRxiv repository. As this chapter is based on this pre-print, it concludes with an explicit statement of the author's own contribution in this project.

5.2 Automated brain segmentation of ant micro-CT images

Three-dimensional (3D) imaging of animals by micro-CT has become popular in morphological biology as a non-destructive method to acquire high-precision data on organismal anatomy [10, 129–132]. The high-resolution 3D data enable the users to visualize and quantify internal and external structures, forming the basis for a wide range of biological applications.

A key challenge for the use of micro-CT lies in the analysis of huge amounts of acquired data. In particular, while 3D images are reconstructed shortly after scanning, segmentation of the images into specific body parts is often a necessary step for quantification and visualization of particular structures.

The most common segmentation method to date is manual processing, which is extremely time-consuming and compromises reproducibility [12]. This limits the number of samples that can be included in a given study, and thus the scientific applications of 3D scanning. For example, developmental biologists may want to analyze large numbers of experimental treatments and replicates. Or, in comparative biology, we may seek to analyze the evolution of a body part across hundreds or thousands of species. The recent emergence of large databases and coordinated projects to scan many species in specific taxonomic groups offers rich opportunities for new research directions if limitations on segmentation can be overcome.

In the medical literature, image segmentation methods have recently become more powerful and efficient due to significant developments in machine learning algorithms. To date, the main focus of automated segmentation methods has been on cells and human organs (e.g., human CT or MRI image segmentation for cancer detection [133, 134]). However, there is broad potential for automated segmentation to accelerate biological research on organisms across the tree of life [7, 135, 136].

New software for biomedical image analysis has steadily progressed during recent years, with the capability for analysis and segmentation of 2D or 3D biological images and the capability to build one's own data processing pipelines [137]. However, despite the unconstrained accessibility to free general-purpose software tools, the development of specific segmentation algorithms is essential to achieve high accuracy, objectivity, and reproducibility. Recently, deep learning methods like CNNs have been successfully applied in numerous image classification and semantic segmentation problems [67, 75].

CNNs have recently become widely used in image processing due to their high performance, the efficiency of GPUs, and the availablity of free software platforms and pre-trained networks [138]. Toolsets and pipelines that use classical statistical methods such as ANTs [139], Biomedisa [140], and Freesurfer [141] are accessible and accurate for the segmentation of high-resolution images. However, these are either not fully automated and still require an expert user and considerable amounts of time and effort [142], and/or require training examples within the same scan, and/or are not adaptable to diversity and complexity in the target set. On the other hand, accurate and general toolkits and application frameworks that use machine learning techniques such as SlideCam have been successfully used for medical image segmentation as well as computer-aided diagnosis and analysis of images spanning from human brain segmentation to cancer detection [143]. However, to date no toolkit has been designed to recognize homologous parts across a wide diversity of animal species, which would require an appropriate choice of network architecture, fine-tuning of hyperparameters, and the production and curation of substantial, high-quality datasets. When it comes to analyzing such images, segmentation remains a most challenging task, and often manual or semi-automated- segmentation is still the only way.

U-Net is a CNN architecture that has shown high accuracy and robustness for biomedical image segmentation [86]. It uses relatively small number of training images to accurately segment areas even with unclear borders. The simple architecture of U-

Net makes it easy to develop and very fast to train. Once a U-Net is trained, the acceleration of the segmentation is extreme: for example, the segmentation time for one ant brain, which may be up to a whole day's work if performed manually, is reduced to merely 1-2 minutes by automatic segmentation.

In this Chapter we present an automated pipeline for segmentation of different inner parts of insects in volumetric data using micro-CT scans, and specifically ant brains across a diverse set of different ant species, as a test case. A basic question for such studies is how general algorithms can be applied across the tree of life. Can an algorithm trained to recognize a part in one type of organism be used on more distant relatives, or do they break down once applied outside the group for which they were developed? Ants are a well defined clade following a similar overall body plan, but reflect $>$ 100 million years of diversification and a large range in ecological, sensory, and behavioral modes [6, 144]. We expect ant brains to have an intermediate level of diversity and thus be a reasonable test case: they will change in size and shape across species, while the general organization and tissue composition should be conserved [145]. As a secondary experiment, we assess whether the ant brain segmenting algorithm we developed can be applied with minimal modification to recognize brains in distantly related insects.

Overview of the segmentation pipeline

Our micro-CT image segmentation pipeline is composed of multiple modules, as illustrated in Fig 5.1.

- **Sample preparation:** Before scanning, specimens were stained in iodine for an average of two weeks to enhance tissue contrast in the raw images.
- **Image acquisition and reconstruction:** An X-ray micro-CT image dataset was acquired from 76 species of ants. The acquired images were reconstructed along all three perpendicular directions that comprise a Cartesian system forming a detailed cross-section dataset.
- **Volume rendering:** The reconstructed raw images were used for creating a 3D model for volume rendering, to be used for visual inspection and future morphological studies.
- **Semi-automated segmentation:** Raw images of heads were segmented semi-automatically using the seed-based watershed tool of the Amira software. Labels were assigned to areas of interest, starting with the brain. The output of this process was used as an input for the study of Chapter 4. The databases of both raw and labeled images were then pre-processed to enhance their homogeneity and used as training and validation data.
- **CNN development:** An implementation of the U-Net architecture was built for automated segmentation.
- **Training:** 60% of the acquired segmented brain images (46 species) were used for network training; the remaining 40% (30 species) was reserved for testing.

- **Pixel island detection and post-processing:** After segmentation by U-Net, pixel island detection was used to identify the largest continuous areas and to remove isolated segments.

5.3 Materials and methods

5.3.1 Image acquisition

In total, we collected one head scan per species from 76 different ant species using a ZEISS Xradia 510 Versa 3D X-ray micro-CT microscope, and ZEISS Scout and Scan Control System software (version 10.7.2936). The scanner settings were determined by the specimen size (e.g., voltage: 30 keV and exposure time: 3-10 s) resulting in 5- to 20-hour scans (12 hours on average).

With a view to expanding our dataset in order to eventually enhance robustness and suppress overfitting during network training (see below), we used 2D cross-sections of planes along all three directions of our 3D brain scans. To highlight the morphological diversity of the scanned specimens, we also performed 3D reconstruction of the resulting scans with XMReconstructor (version 10.7.2936). The output images comprised 1000×1000×1000 px, on average, with resolution down to 1 μm. Exemplaries raw images of full-body scans from different ant species are shown in Fig 5.2.

5.3.2 Image processing

Generation of data for training and validation

We processed the data with the Amira software (version 6.0), and semi-automatically segmented the brain areas (on average, 300×400×600 px per brain) using the seed-based watershed method [146] in the volumetric data, as shown in Fig 5.3. Each semi-automatically segmented 3D brain was dissected into 2D slices on planes along all three directions. To eliminate the empty space of the image and zoom in on the region of interest, the image was cropped and rescaled into 520×520 px to show only the brain area and its adjacent muscles and fibers. The texture of the brain is unique within the whole image of the head, which facilitates its identification. However, the borders are more challenging to classify as numerous nerves branch out from the brain connecting it with the rest of the ant's body; these nerves had to be removed manually, as Amira's watershed tool typically mistakes them for brain areas. This makes semi-automated segmentation challenging and considerably time-consuming. Eventually, this process resulted in an average of 1000 2D brain images per specimen at an estimated average time cost of 5 hours per specimen.

Data pre-processing

As a pre-processing method, we chose histogram equalization [147] using the "imaug" library [148] in Python. Since the images were collected from different samples and

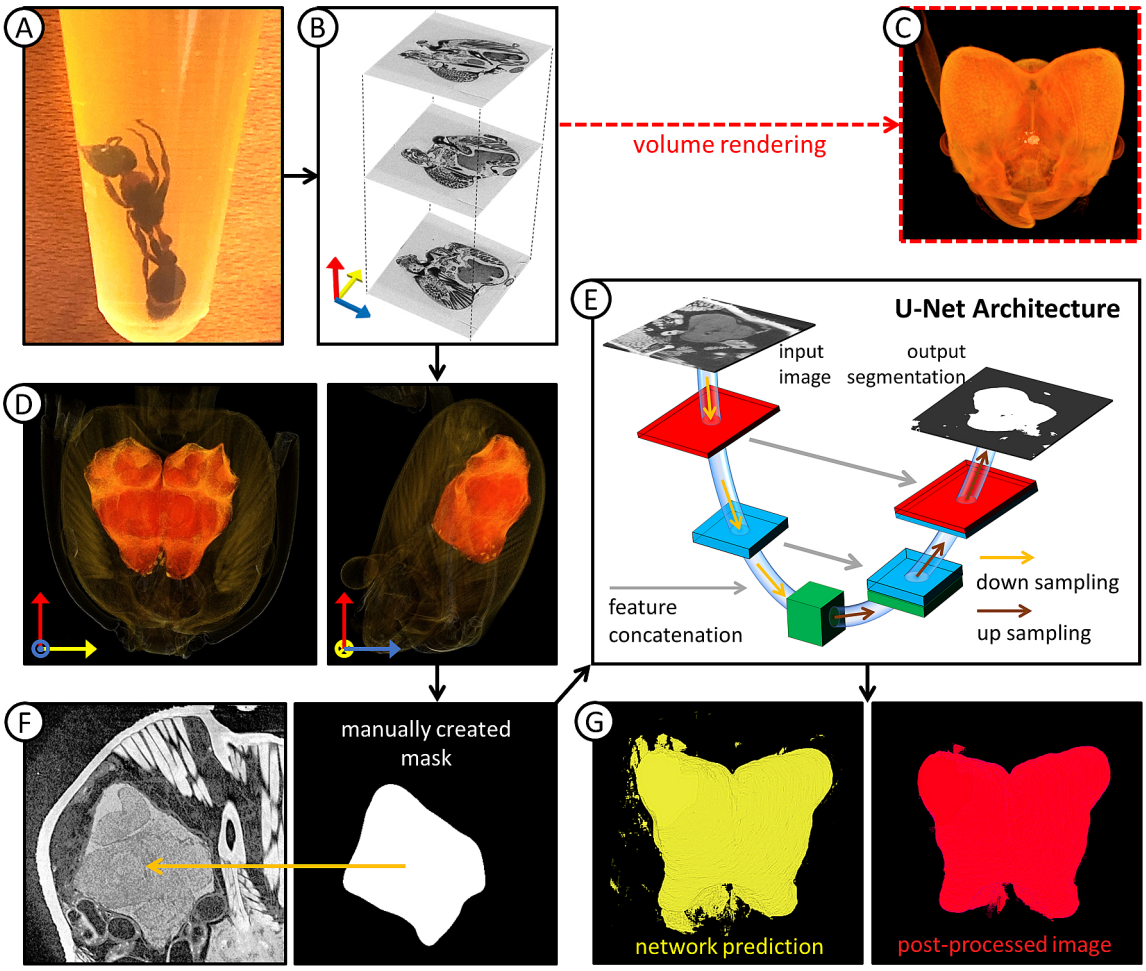

Figure 5.1: Segmentation pipeline overview. (A) Specimens are placed in iodine for staining for two weeks and then placed in small vials containing 99% ethanol to prevent them from moving during scanning. (B) The scanner acquires images along all three axes, and, using a user-defined reference image, automatically reconstructs the whole volume of the scanned specimen. (C) Volume rendering for future morphological studies is performed using Amira software. (D) Semi-automated segmentation of the brain volume of each scan (in orange) using the watershed method in Amira. (E) Schematic representation of the U-Net architecture used as the core of the pipeline for the development of a fully automated brain segmentation method. (F) The acquired brain images are used for training after pre-processing augmentation and manual creation of masks. (G) The network's prediction (in yellow) is post-processed for smoothing out over-predicted areas (in red).

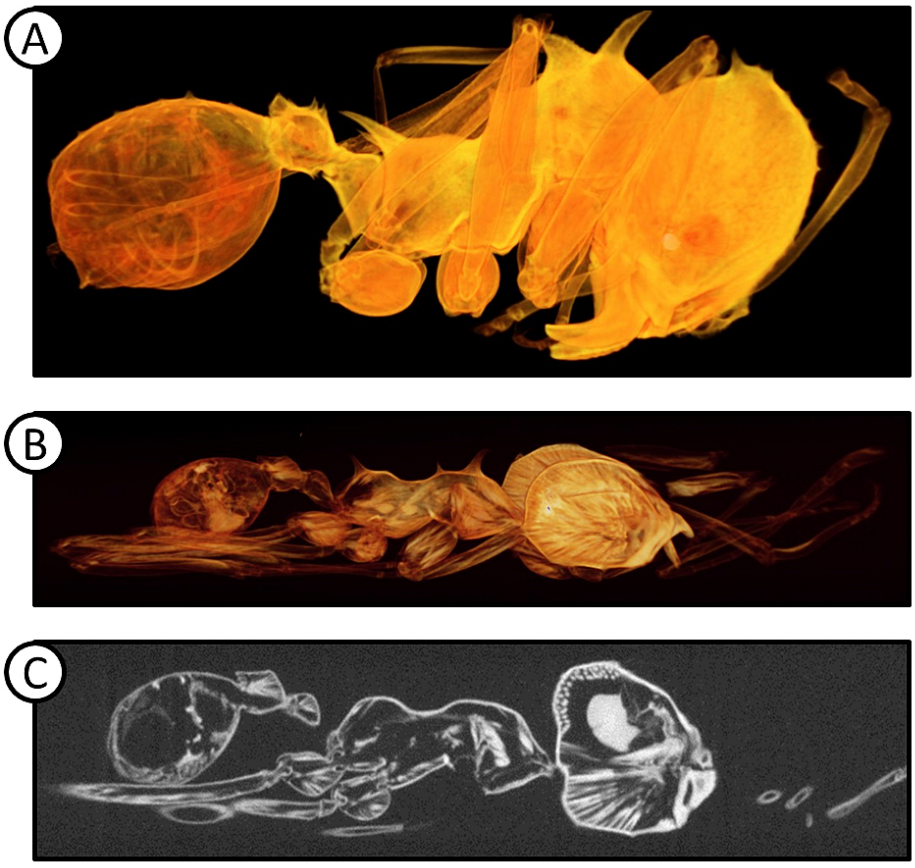

Figure 5.2: Exemplar raw images of full-body scans from different ant species. 3D reconstructed micro-CT image of (A) *Acromyrmex versicolor* and (B) *Atta texana* worker specimens, using volume rendering in Amira. (C) 2D micro-CT full body image of *Atta texana* specimen. The brain area is the densest, most uniform area in the whole body, which makes it easy to recognize in most high-quality scans.

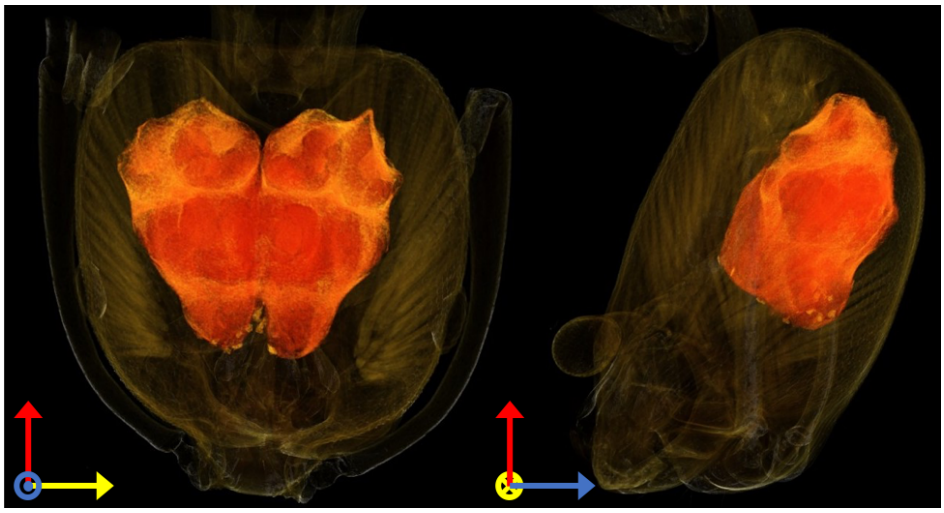

Figure 5.3: Example of semi-automated brain image segmentation. The brain area (in orange) of an *Atta texana* ant specimen was segmented using the watershed method in Amira; the image was manually post-processed by smoothing and cropping over-segmented areas.

scans, and their contrast was not optimized during scanning, pixel density equalization improved the network's performance remarkably. After histogram equalization, the contrast was improved, accentuating the texture of the brain and, thus, making it easier to identify. An exemplar result of data pre-processing is shown in Fig 5.4.

U-Net structure

We chose the U-Net architecture because it has been the most successful CNN for CT image segmentation to date. The U-Net is not a conventional CNN architecture, in the sense that it extends the contracting path of a typical CNN by a symmetrical expansive path [86]. For optimal efficiency, our code uses the open source GPU-TensorFlow library [111] and the TensorFlow U-Net implementation, as described in Akeret et al. [149], utilizing Jupyter notebook and Python. Our network consists of a five-fold repetition of two 3×3 convolutions followed by a rectifier linear unit (ReLU) and a 2×2 max pooling. Starting with 64 features, each layer doubles their number resulting in 256 features before starting the expansive path which consists of two 3×3 convolutions followed by ReLU and 2×2 up-convolutions. The number of features is halved with each up-convolution but the result is concatenated with the features from the matching down-sampling layer. Finally, 1×1 convolution is applied to map each feature to the number of classes, i.e., two. The batch size used is 1∼4, the stride is 1, and zero padding is used for the max pool (Figure 5.5).

Training

To assess the effect of various parameters on the performance and the processing time, different batch sizes, numbers of initial features, epochs, and iterations were tested. We selected a random 60% of data (46 species) for training and used the rest (30

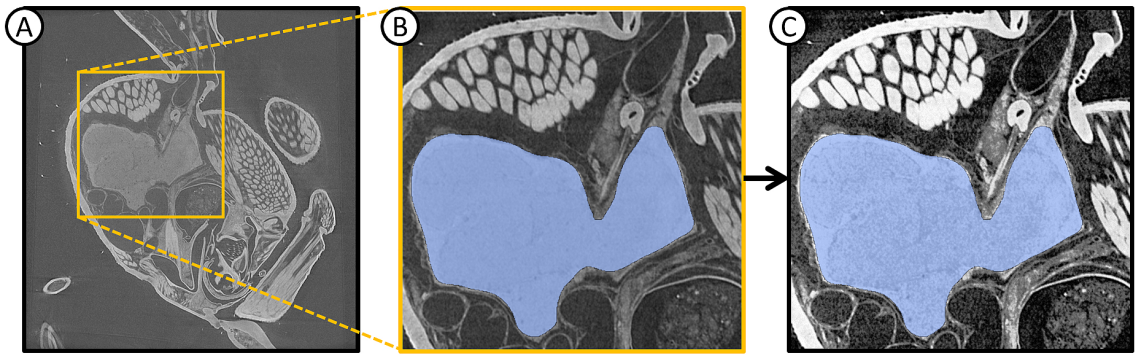

Figure 5.4: Data augmentation. (A) Initial 2D image of a full head scan of an *Atta texana* ant specimen. Pre-processing is performed in two steps: (B) The image is cropped around the brain area, keeping some of the muscles, nerves, and fibers that are close (or even attached) to the brain. The manual segmentation of the brain is indicated in blue. (C) Histogram equalization is used for additional augmentation, which enhances the contrast and projects the inner parts of the brain more clearly.

species) for testing. The optimal parameter values were chosen on grounds of low computational cost and high classification accuracy for training data. We trained the network by optimizing the binary cross-entropy function with L2 regularization using stochastic gradient descent with the momentum of 0.8. The initial weights were selected by using a Gaussian distribution, in agreement with Ronneberger et al. [86]. Batch normalization was added in the first 3 layers to avoid overfitting as well as to accelerate training [49]. Finally, we added dropout equal to 0.5 in the first 3 layers also to avoid overfitting. We trained our network for 10 epochs, with mini-batch 32 on a 520×520 pixel image, costing 120 hours in our workstation using a GeForce GTX TITAN Xp and a GeForce GTX 1080 graphics cards.

Post-processing

We post-processed our network's prediction by using pixel island identification and isolation [150]. After predicting the brain area along all three planes, the biggest pixel island was chosen as the brain area. This process boosted by almost 10% on average our prediction success rate of both the Jaccard Index (IoU) and Dice Coefficient (F1 score) [98].

5.4 Results

5.4.1 Segmentation of ant brains

First, we applied our method to our primary taxonomy group of choice, i.e., ants, and trained our network to segment the brain areas in micro-CT scans from different ant species.

Our processed data of 38,000 of the size of 520×520 pixel images from 46 species were used for training and validation (randomly split into 80% for training and 20%

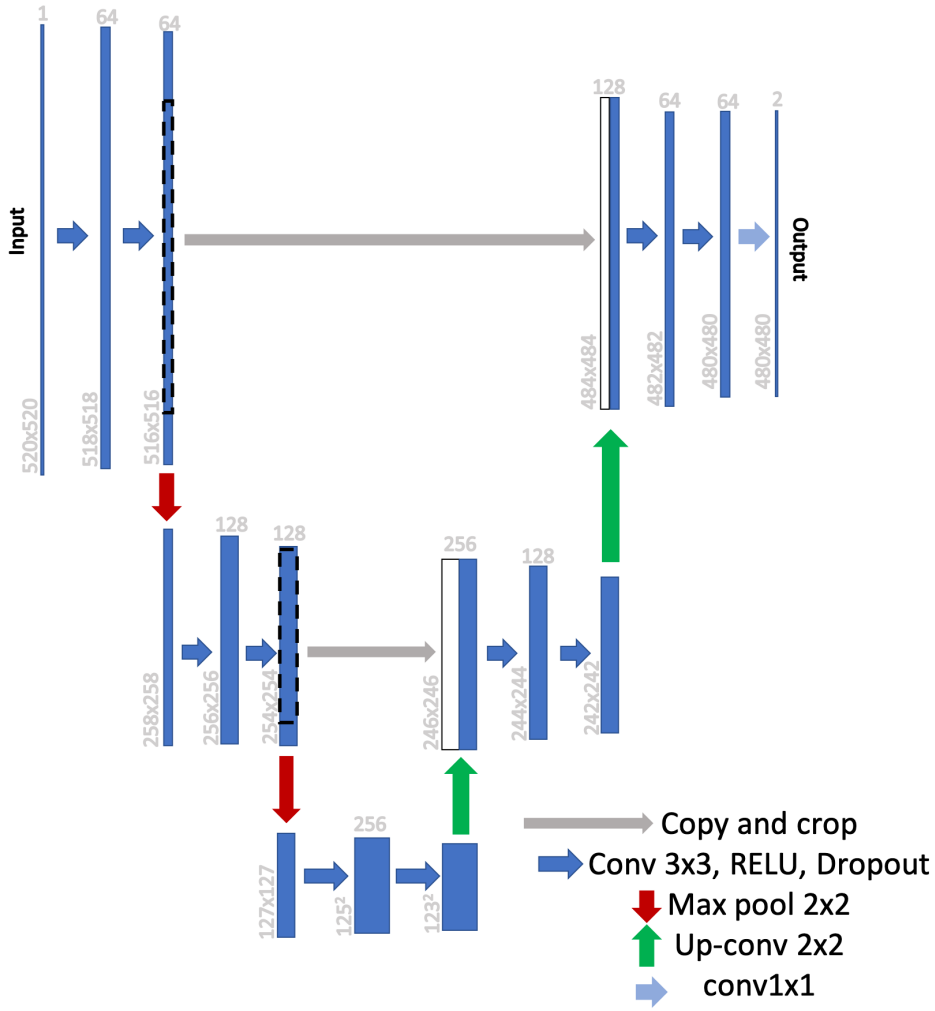

Figure 5.5: U-Net implementation. The architecture of the used CNN is an implementation of U-Net. It consists of two parts: a. two 3×3 convolutions followed by 2×2 max pooling and b. two 3×3 convolutions followed by 2×2 up-convolutions. Dropout was added to avoid ovefitting. As a final step, a 1×1 convolution is applied, resulting in an output map with two classes.

Table 5.1: Accuracy scores. Performance evaluation of our proposed pipeline. Both performance descriptors studied (IoU and F1 scores) increase steadily with increasing number of images and post-processing.

Number of images of training set	IoU	F1
3,500 - xy plane	50%	62%
10,000 - xy plane	63%	71%
38,000 - along all three directions - no post-processing	72%	80%
38,000 - along all three directions - after post-processing	80%	90%

for validation) and the remaining 20,000 of the same size images from 30 species were used for testing.

As shown in Table 5.1, both IoU and F1 scores were steadily increased as we added more 2D images from planes along the same xy directions of different species, and even more so after we included reconstructed 2D images from planes along all three directions of our 3D brain scans. To estimate the generalized performance of our network, we calculated the true positive rate (TPR) and false positive rate (FPR) values of our images by changing the discrimination threshold of our network [58], shown in Fig 5.6. All values are close to 1 for both training and testing sets while the FPR values remain less than 0.4 for most cases, indicating that our network predicted the brain region and its border accurately and without over-predicting. Results for test and training images are similar, suggesting good generalization capabilities for the optimized hyperparameters of our network.

Finally, a post-processing step also boosted the performance of our network, yielding even more satisfactory results. Example results of our network's performance on validation and testing data are shown in Fig 5.7 demonstrating a predicted area in good agreement with the ground truth; our automated segmentation pipeline achieves an approximate maximum of 80% IoU and 90% F1 score. Prediction times were in the order of only a few minutes, significantly lower than for the semi-automated segmentation commonly used to this day.

5.4.2 3D volume rendering

After segmenting the 2D slices, the 3D brain volume was readily computed by loading the stack of images in Amira or ITK-snap. Thus, using a 2D network allowed us to maintain high accuracy, performing 3D segmentation in a faster and easier to train way. An exemplar predicted brain area is shown in Fig 5.8; 3D volume was reconstructed from the 2D predicted images with Amira software. The switch from 2D to 3D is straightforward, giving the user of our pipeline the ability to adapt it to their own dataset circumventing the complications of using an actual 3D CNN.

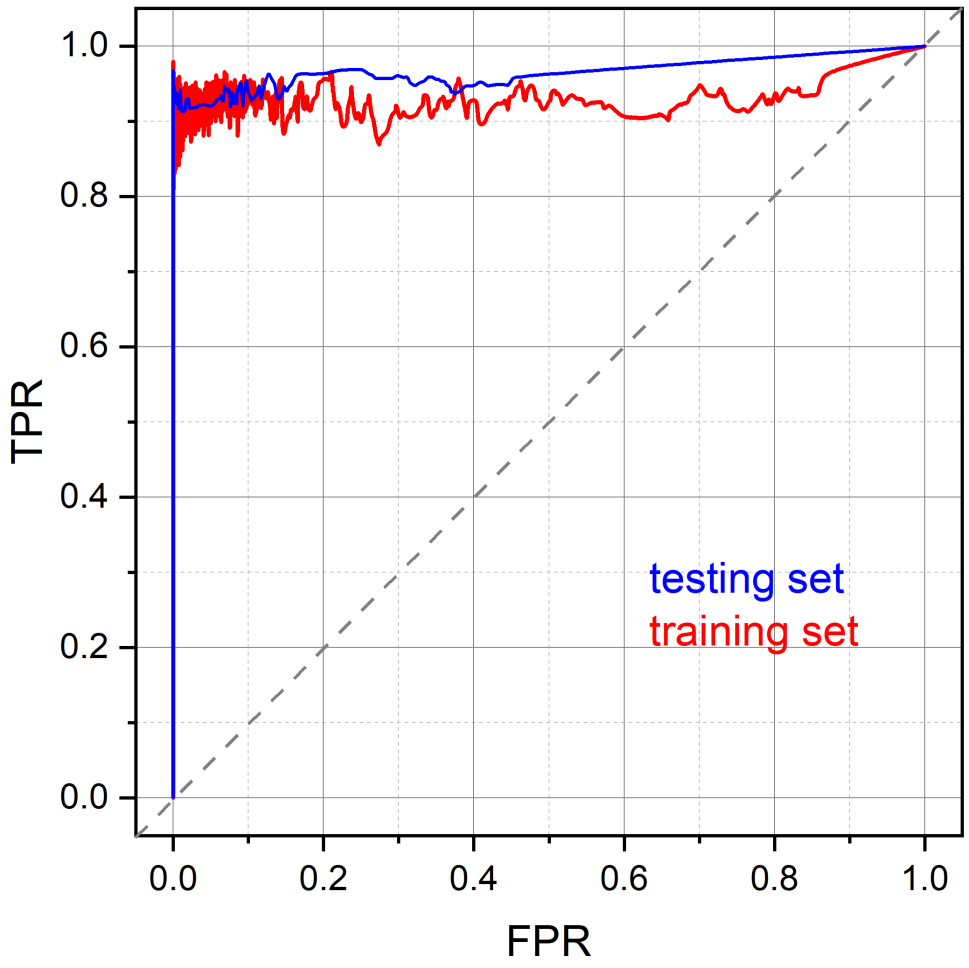

Figure 5.6: Network performance evaluation. High TPR and low FPR values for training (red) and testing data (blue) indicate the network's high generalizability.

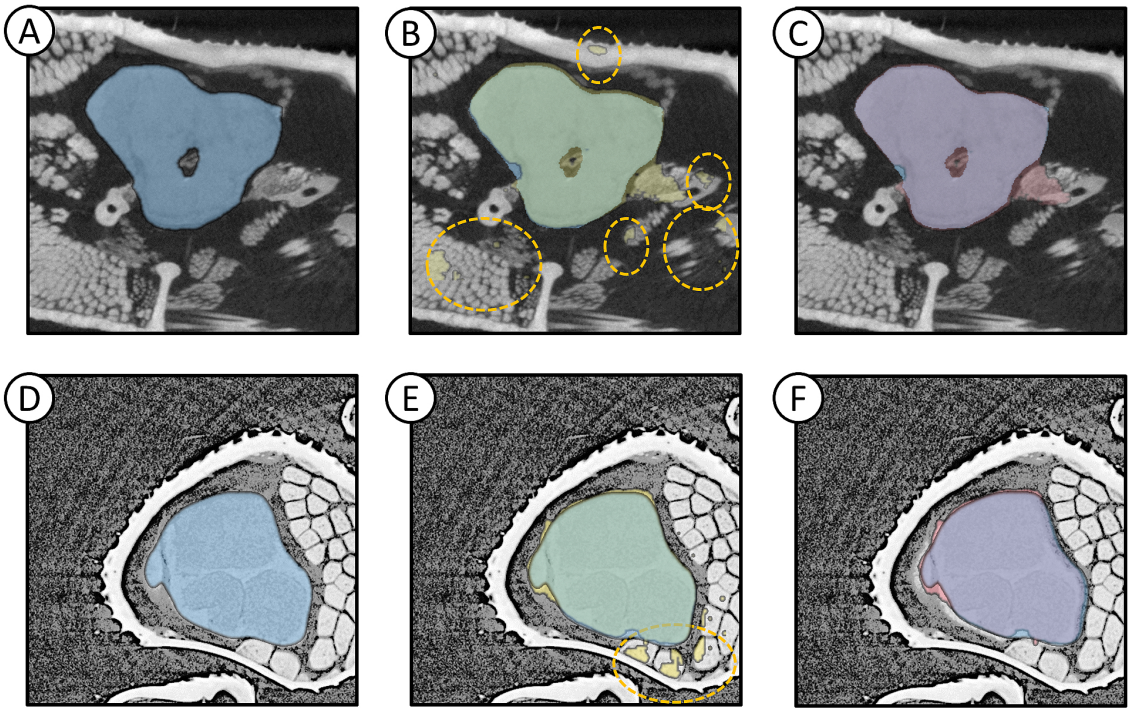

Figure 5.7: Pipeline performance calculated both for validation (top row) and testing (bottom row) sets. (A, D) Raw images of head of *Atta texana* ant and *Carebara atoma* specimens, cropped along the x-y axes. The manually segmented brain areas are indicated in blue. (B, E) Network predictions before post-processing (in yellow). Areas in yellow dotted circles are pixel islands not connected to the brain area that were overpredicted. (C, F) Predictions after post-processing (in red). The borders of the predicted areas show good agreement with the manual segmentation in both sets. Note that in overlapping manually and automatically segmented areas in B, C, E, and F, colors appear green or purple.

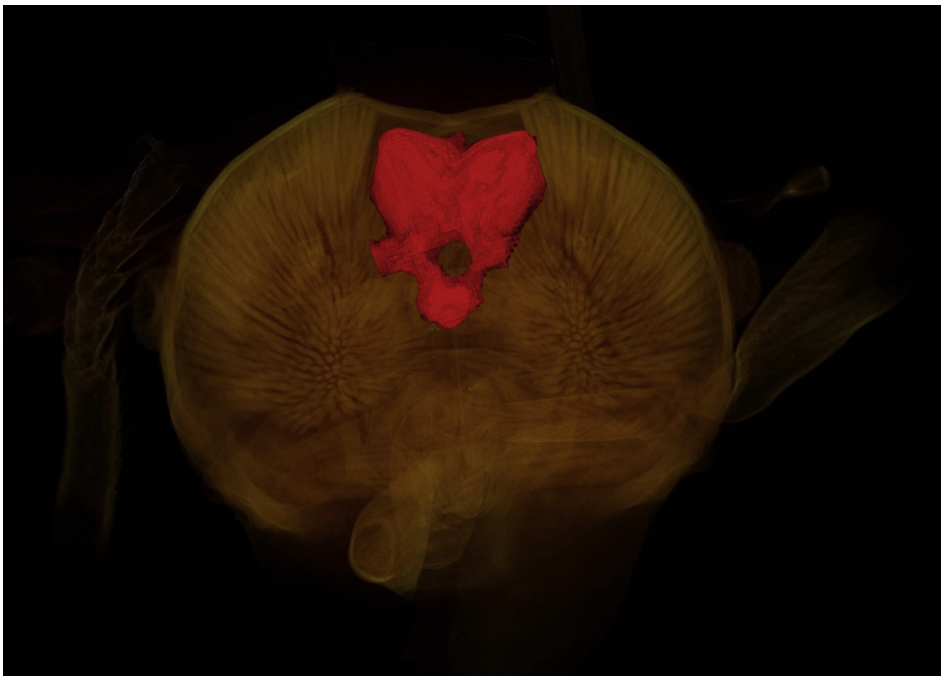

Figure 5.8: 3D volume of ant brain reconstructed from 2D images predicted from the algorithm. 3D reconstructed brain prediction of an *Atta texana* worker.

5.4.3 Generalization to other neural systems and other insects

The U-Net step appears to be largely driven by textures, with the pixel island detection step used to isolate the brain. Even though our customized U-Net was designed for the segmentation of ant brains, it was also successfully applied for the segmentation of neural tissue in other parts of ants and works on distantly related species. Our network was able to predict the whole neural system in full-body scans of ants, as shown in Fig 5.9, being able to predict the same texture as the brain in different ganglia in the thorax (called mesosoma in ants).

Our network also gave good prediction for the brain area in scans of various different distantly related insect species. We used our pre-trained (on ant brains) network to segment the brain areas of micro-CT scans of model organisms such as flies (*Drosophila*) and honey bees (*Apis mellifera*) as well as closely related insects such as praying mantises (*Leptomantella*) and termites. Since its prediction capability relies mainly on identifying the texture of the brain area, which does not differ significantly among different insect species, our pre-trained network was able to perform satisfactorily without further adaptation on the data. Exemplar results are shown in Fig 5.10 for (A-B) wasp and (C-D) praying mantis brain prediction, respectively: remarkably, our network was successful in segmenting the brains of different insects without any prediction accuracy losses (when compared to predictions for ants), indicating its flexibility and its lack of necessity for training on each specific distinct species.

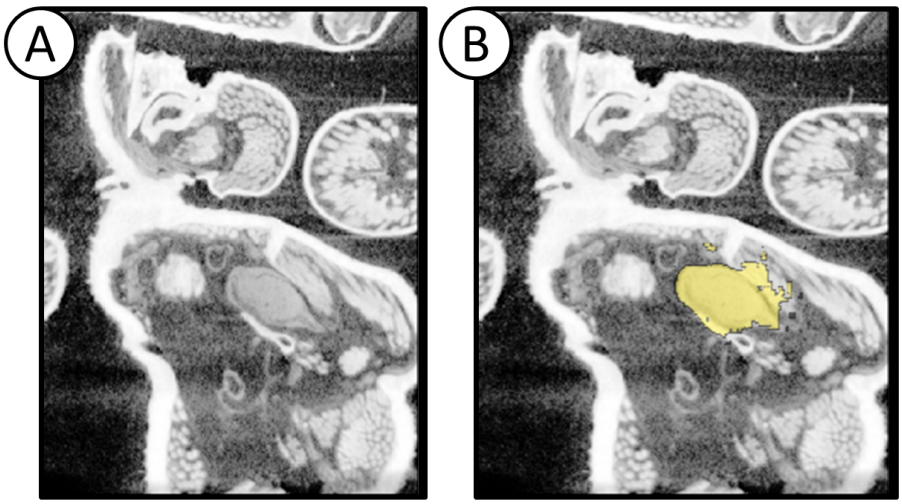

Figure 5.9: Prediction of ganglia in the thorax. As the tissue texture in the image is similar with that of the brain, the network accurately predicts other areas of nervous tissue in the organism. The pixel island detection step isolates the brain, but without this step neural tissue can be isolated.

5.5 Discussion

To bring morphology fully into the big data era, we need automated methods to retrieve biological meaning from large volumes of images. The proposed automated pipeline is a step in this direction, presenting considerable advantages over other standard methodologies. First of all, automated segmentation is achievable within a few minutes for each specimen, producing faster and more accurate results than semi-automated or manual segmentation. A noteworthy additional advantage is that once algorithms have been trained, advanced expertise in morphology is not required, while manual and semi-automated segmentation usually require advanced knowledge [151]. In fact, during testing our network often outperformed even experienced users and compensated for their oversights or misjudgments, predicting correctly brain areas that were accidentally missed out during manual segmentation.

The two approaches in our method, U-Net and pixel-island detection, represent two complementary steps which suggest a path forward for automated segmentation of structures in complex organisms. U-Net was efficient at retrieving tissue with similar properties in the image, but in our implementation did not make use of shape and position. Thus, we found it retrieved all the structures of neural tissue across the body, even though it was trained on the brain alone. The brain was then isolated with the pixel-island detection, which isolated the largest structure in the head. In general, we expect a combination of tissue-level identification followed by other methods that make use of size and spatial organization to be a powerful combination that should generalize to a wide range of anatomical tissues and parts.

During testing with other insect species, we used both high and low resolution/quality images acquired from different laboratory and synchrotron-based micro-CT scanners. Our results showed that our segmentation pipeline can perform without losing its ac-

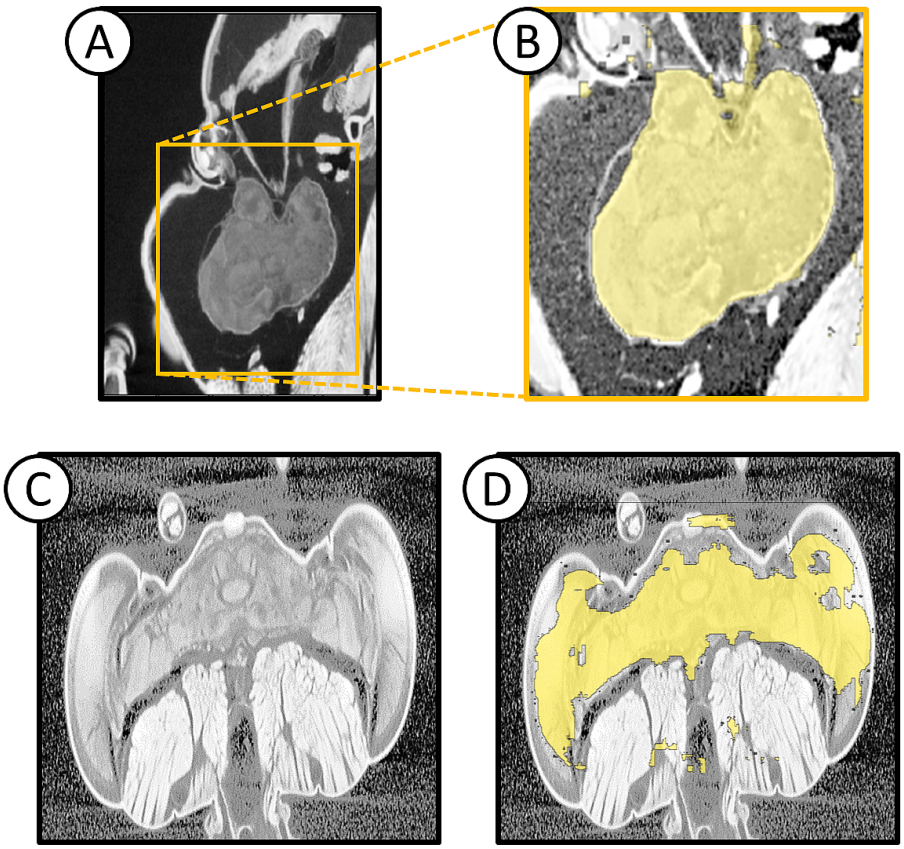

Figure 5.10: Application of pipeline for other insect species. The brain textures of various insect species can be very similar with those of ants, facilitating the prediction by the network even without pre-training on specific insect brain scans. (A) Raw image of wasp head and (B) its prediction without post-processing, indicating satisfactory identification of the borders of the brain area. (C) 2D image of praying mantis head and (D) the prediction of its brain area without post-processing. Even though the network over-predicts some small pixel islands, it excludes from its prediction areas of the muscles, fibers and cuticle.

curacy to predict the brain area across highly divergent arthropod species and across scanning methods. Finally, the prediction performance of low-resolution images indicates that there is a threshold in the image resolution below which our network is not performing well. Our network's generalizability is high and it can be widely used not only for head but also for whole-body scans of ants and other insects. More importantly, it shows that a similar approach could be used to build a suite of trained networks that can segment anatomy across a wide variety of organisms.

Last, it should be noted that both automated classification and segmentation tasks typically require big datasets for training and validation, which can be a challenge for researchers to produce for any given application. Since no publicly available dataset of micro-CT images of ant brains existed for our case study, we created a new, extensive dataset across a wide variety of ant species. Since neural anatomy across insects share features that make them targets for segmentation, our dataset can act as a starting point for the development of an even bigger library of micro-CT images of insects, and work as a pre-training dataset for future CNNs [152].

5.6 Conclusion

In this Chapter, we employed the U-Net based CNN for the fully-automated segmentation of micro-CT images of insects. We also present an extensive dataset of manually segmented brain images that can be used to pre-train other networks of interest. Our trained network predicted the brain area in ant images fast and with high accuracy. Further, our network was able to generalize and predict the whole neural system in full-body scans, as well as to predict ganglion areas that were missed by manual segmentation. After training, the network's performance was tested on training and testing data showing good agreement between prediction and mask scoring 90% F1 and 80% IoU. Our pipeline allows successful segmentation in only a few minutes instead of hours which are typically required for manual segmentation.

One of the most important features of the framework described here is that it can be applicable to other anatomical features. Preliminary results on other organs have shown that it can be easily tuned and trained to predict muscles as well as the cuticle of the insect bodies. Moreover, for future studies, segmenting sub-components of the brain will also be tested as it is of great importance. Specific attention was paid so that the application of the pre-trained network is straightforward and user-friendly, which we aspire will enable the community to adopt it as a valuable resource.

The development of large-scale 3D datasets across phylogenetically diverse taxa (e.g., OVert [13]) opens up new vistas for comparative research. Likewise, developmental biologists may want to use high-throughput scanning to image hundreds or thousands of specimens as part of an experiment. However, just as DNA sequence data needs bioinformatic algorithms to process massive datasets, large scale image collections require algorithms to digest and extract biologically meaningful data. Algorithms such as this one offer a way forward for powering a "big data" approach to organismal morphology.

5.7 Contributions

For this project the author (i) scanned the specimens used as training and testing data; (ii) manually segmented the brains of all specimens; (iii) created the training and testing sets; and (iv) wrote the code for the pre-processing, training network, and post-processing. Additionally, the author prepared a GitHub page for the code, and wrote the first draft (and contributed to the editing of subsequent drafts) of an article currently deposited as pre-print at bioRxiv (to be submitted for publication in an international peer-reviewed journal soon).

Chapter 6

Conclusion

"...at a speed of three hundred million facts per second, and mile after mile of tape coiled out and gradually buried the Ph.D. pirate beneath its windings, wrapping him, as it were, in a paper web..."
Cyberiad

This Chapter summarizes the challenges that defined the body of research that comprises this PhD project. The main findings and output (presented in Chapters 3, 4 and 5) are also reviewed and put into a broader context, with a few remarks about their novelty and potential significance for future research. This chapter, and the whole thesis, concludes with a short, slightly more personal, epilogue.

6.1 Premise

Ants hold a prominent position in ecological studies. They are eusocial insects, forming colonies which display an excessive size range (from few dozen to millions of individuals) with the resultant variability in division of labor, level of communication, and problem-solving protocols. Indigenous ant species are found in most part of the world, thriving in diverse ecosystems, and providing ample information for numerous behavioral and anatomical studies. However, various aspects of their evolution still remain unclear.

Micro-CT scanning, on the other hand, is a non-invasive technique which can produce data-rich volumetric images of insects. As such, it has availed a detailed view on the morphology and anatomy of ants, leading to the discovery of new species and helping decipher their evolution. Supported by phylogenetic considerations, morphological comparison between species enables the investigation of fundamental biological questions, such as the social brain hypothesis which connects the adaptability of the social behavior of animals to modifications in their ecosystem and social structure with changes in their brain investment.

However, a key challenge for the use of micro-CT lies in the analysis of the extensive amounts of 3D data it produces: indeed, the most common segmentation method

to date is still by manual processing, which is extremely time-consuming and often inaccurate. This work proposed a novel, cross-disciplinary approach to solve this problem: the goal was to create new, deep-learning based software tools that automate the segmentation of micro-CT images of ants, making morphological quantification and comparison easier and more efficient, and paving the way for new discoveries in evolutionary studies.

6.2 Main results

Within the framework of this PhD project:

1. A dataset of manually segmented micro-CT scans of ant heads and bodies was created and used in both main research directions (points 2 and 3, below). Both scanning and manual segmentation of the resulting scans were time-consuming and required expertise, which was gained with time.

2. A potential connection was investigated between the brain size evolution and the sociality of ants via detailed statistical analysis of corresponding traits.

3. A workflow was developed to automate the brain segmentation of micro-CT images of ants (that can be generalized for other insects or body parts).

6.2.1 Correlation of brain size and sociality trait in ants

In Chapter 4, we conducted detailed statistical analysis to investigate potential correlations between the brain size of ants and the social structure/complexity of their colonies. To conclude in a relationship between traits as general as possible, species of numerous genera and subfamilies were included in our study; yet, no direct correlation was found. However, a connection was actually indicated between the sociality and the size of the eyes, instead, in concord with previous studies utilizing different approaches [37]. Our unprecedentedly diverse dataset created more robust results, paving the way for a general resolution of the social brain hypothesis debate for insects.

6.2.2 Automation of micro-CT image segmentation

In Chapter 5, we acquired micro-CT scans using ant brains as a target area; our data were both high-resolution and 3D, thus enabling us to visualize both the morphology and the anatomy of each specimen. We created an automated segmentation workflow for the micro-CT images using a U-Net based CNN as the main building block (in addition to pre-processing and post-processing steps). The training set consisted of 38,000 annotated images, making it the largest dataset of micro-CT images of insects available to date. Our workflow achieved high accuracy in ant images but also in images of other insects, making it more general and robust than any other existing segmentation method in the field.

6.3 Outlook

Future research directly at the wake of the results of the current PhD project can include:

1. A generalization of the automated segmentation workflow: The generalizability of our network, which was developed specifically with ants' brains in mind, was tested on images of different insects and body parts, and on images of ants produced by a different type of source (i.e., synchrotron). In both cases, the results were of high accuracy; however, admittedly, the numbers of tested samples were quite limited, so further testing would be expedient for the network to be adopted as an indispensable tool by the relevant scientific community.

2. Application to additional datasets: The scarcity of annotated datasets (and the time required to produce them) meant we could only apply our CNN to a single dataset. This however is likely to change in the near future, as the possibility to produce literally thousands of scans of different species using synchrotron facilities is realized. Our aspiration is for our trained network to be used for the segmentation of the brains of these species as soon as such datasets become available, promoting a big-data approach to organismal anatomy.

3. The generalization of the social brain hypothesis investigation: Even though a diverse dataset of ants was used for the analysis of the connection between brain size and sociality, species demonstrating totally social or totally non-social features were not included in the current study. The influence of these two extreme edges should be considered closely in future studies. Further, as a ripple effect of points 1 and 2, above, the statistical analysis performed in a relatively limited dataset here can be applied to the massive new datasets, shedding light on the brain size evolution more conclusively.

4. Finally, our dataset can be used as a pre-training set for the brain segmentation of micro-CT images of other insects. As the number of projects focused on the digitalization of specimens is rising, a training dataset and a tested CNN can be of great value to ecologists.

Looking further into the future, we aspire that the outcome of our current project can be of importance, as no similar software exists that uses deep learning techniques to segment micro-CT images of any taxa. Exploration of features of the U-Net CNN and modifications in various parameters can lead to the expansion of the developed software toward fast, automated segmentation of images of numerous species. Therefore, the proposed software could also be used for segmentation of any micro-CT images with similar patterns.

6.4 Epilogue

To conclude this thesis, it might be appropriate to zoom out of the specific research subject and consider the role of AI and deep learning in the broader domain of organismal anatomy or ecology as a whole.

There is no doubt that the trending paradigm in current scientific research is AI-driven. Following rapid advances in computer science and the expansion of computational power over the last decades, scientists across various disciplines now have the unprecedented possibility to apply brute force approaches, producing and digging into such huge loads of data which would have been prohibitive even in the near past. The crucial point here is that efficient AI systems need not run the risk of being buried under endless rolls of tape, like the literary interstellar pirate Pugg from Stanislaw Lem's "Cyberiad", or getting forever lost in Jorge Luis Borges' "Library of Babel". Instead, AI can filter, sort, and scan the data for patterns. Therefore, AI is an indispensable tool for the manipulation and analysis of big data, with new mathematical and computational tools (machine and deep learning included) enabling it to overcome obstacles of previous years. Simultaneously, improvements in software and hardware are currently being combined, as there is a trend for the fabrication of integrated circuits that can perform logical functions on the device level, further boosting computational speed. The driving force for such a development is the new, holistic approach to various scientific disciplines, biology being a characteristic example.

As such, it may be interesting to note that the evolutionary debate of interest in this project, i.e., the social brain hypothesis in insects, is not new. In fact, the first association of the social behavior of some apocritan Hymenoptera species with their mushroom body calyx sizes was proposed in the mid-19^{th} century [153]. From an applied mathematician's point of view, it only feels appropriate that 21^{st} century tools are employed to resolve such age-old debates, underlying the maturity biology has reached through inoculation with computer science. Amusingly, computer science also draws inspiration from biology and the ultimate optimizer which is Nature; as an example, the ant colony optimization, ACO, algorithm springs to mind: a probabilistic problem-solving method via finding optimized paths through graphs, which models the pheromone-based communication of real, biological ants.

Bibliography

[1] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 3 edition, 2010.

[2] Lim Milton. History of ai winters, 2018. URL https://www.actuaries.digital/2018/09/05/history-of-ai-winters/.

[3] Yoshua Bengio. Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, 2(1):1–127, 2009. ISSN 1935-8237. doi: 10.1561/2200000006.

[4] Rodrigo Ceron. Ai, machine learning and deep learning: What's the difference?, 2019. URL https://www.ibm.com/blogs/systems/ai-machine-learning-and-deep-learning-whats-the-difference/.

[5] Vazquez Favio. A "weird" introduction to deep learning, 2018. URL https://towardsdatascience.com/a-weird-introduction-to-deep-learning-7828803693b0.

[6] Edward O. Wilson. Causes of ecological success: The case of the ants. *Journal of Animal Ecology*, 56(1):1–9, 1987. ISSN 00218790, 13652656. URL http://www.jstor.org/stable/4795.

[7] Michael Staab, Francisco Hita Garcia, Cong Liu, Zheng-Hui Xu, and Evan P. Economo. Systematics of the ant genus proceratium roger (hymenoptera, formicidae, proceratiinae) in china with descriptions of three new species based on micro-ct enhanced next-generation-morphology. *ZooKeys*, 770:137–192, 2018. doi: 10.3897/zookeys.770.24908. URL https://doi.org/10.3897/zookeys.770.24908.

[8] R I M Dunbar and Susanne Shultz. Why are there so many explanations for primate brain evolution? *Philos Trans R Soc Lond B Biol Sci*, 372(1727), Aug 2017. ISSN 1471-2970 (Electronic); 0962-8436 (Print); 0962-8436 (Linking). doi: 10.1098/rstb.2016.0244.

[9] Yann Lecun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Handwritten digit recognition with a back-propagation network. In D. S. Touretzky, editor, *Advances in Neural Information Processing Systems 2*, pages 396–404. Morgan Kaufmann, 1990.

[10] Ryuta Mizutani and Yoshio Suzuki. X-ray microtomography in biology. *Micron*, 43(2-3):104–115, Feb 2012. ISSN 1878-4291 (Electronic); 0968-4328 (Linking). doi: 10.1016/j.micron.2011.10.002.

[11] Peter Hawkes and John Spence. *Handbook of Microscopy*. Springer, 01 2019. ISBN 978-3-030-00068-4. doi: 10.1007/978-3-030-00069-1.

[12] Sebastian Schmelzle, Thomas Van De Kamp, Michael Heethoff, Vincent Heuveline, Philipp Lösel, Jürgen Becker, Felix Beckmann, Frank Schluenzen, Jörg U. Hammel, Andreas Kopmann, Wolfgang Mexner, Matthias Vogelgesang, Nicholas T. Jerome, Oliver Betz, Rolf Beutel, Benjamin Wipfler, Alexander Blanke, Steffen Harzsch, Marie Hörnig, and Tilo Baumbach. The NOVA project: maximizing beam time efficiency through synergistic analyses of SRμCT data. In Bert Müller and Ge Wang, editors, *Developments in X-Ray Tomography XI : [Proceedings] - SPIE, 2017. - ISBN 9781510612396 - doi:10.1117/12.2275121*, volume 10391, page 103910P, Bellingham, Wash., 2017. SPIE. (c) SPIE.

[13] Gabriel S. Yapuncich, Addison D. Kemp, Darbi M. Griffith, Justin T. Gladman, Erin Ehmke, and Doug M. Boyer. A digital collection of rare and endangered lemurs and other primates from the duke lemur center. *PLOS ONE*, 14(11):1–47, 11 2019. doi: 10.1371/journal.pone.0219411. URL https://doi.org/10.1371/journal.pone.0219411.

[14] wikipedia. List of synchroctron facilities, 2021. URL https://en.wikipedia.org/wiki/List_of_synchrotron_radiation_facilities.

[15] S. J. Gould, R. C. Lewontin, J. Maynard Smith, and Robin Holliday. The spandrels of san marco and the panglossian paradigm: a critique of the adaptationist programme. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 205(1161):581–598, 1979. doi: 10.1098/rspb.1979.0086. URL https://royalsocietypublishing.org/doi/abs/10.1098/rspb.1979.0086.

[16] Joseph Felsenstein. Phylogenies and the comparative method. *The American Naturalist*, 125(1):1–15, 1985. ISSN 00030147, 15375323. URL http://www.jstor.org/stable/2461605.

[17] Mary C. McKitrick. Phylogenetic constraint in evolutionary theory: Has it any explanatory power? *Annual Review of Ecology and Systematics*, 24(1):307–330, 1993. doi: 10.1146/annurev.es.24.110193.001515. URL https://doi.org/10.1146/annurev.es.24.110193.001515.

[18] David R. Roberts, Volker Bahn, Simone Ciuti, Mark S. Boyce, Jane Elith, Gurutzeta Guillera-Arroita, Severin Hauenstein, José J. Lahoz-Monfort, Boris Schröder, Wilfried Thuiller, David I. Warton, Brendan A. Wintle, Florian Hartig, and Carsten F. Dormann. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, 40(8):913–929, 2017. doi: https://doi.org/10.1111/ecog.02881. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ecog.02881.

[19] Theodore Garland, Paul H. Harvey, and Anthony R. Ives. Procedures for the analysis of comparative data using phylogenetically independent contrasts. *Systematic Biology*, 41(1):18–32, 1992. ISSN 10635157, 1076836X. URL http://www.jstor.org/stable/2992503.

[20] Alexander McFarlane Mood, Franklin Arno Graybill, and Duane C. Boes. *Introduction to the theory of statistics*. McGraw-Hill international editions. McGraw-Hill, Auckland [u.a.], 3. ed. edition, 1974. ISBN 0070428646. URL http://gso.gbv.de/DB=2.1/CMD?ACT=SRCHA&SRT=YOP&IKT=1016&TRM=ppn+021835462&sourceid=fbw_bibsonomy.

[21] Hirotogu Akaike. *Information Theory and an Extension of the Maximum Likelihood Principle*, pages 199–213. Springer New York, New York, NY, 1973.

[22] Karl Pearson. The problem of the random walk. *Nature*, 72(1865):294–294, 1905. doi: 10.1038/072294b0. URL https://doi.org/10.1038/072294b0.

[23] Edward A. Codling, Michael J. Plank, and Simon Benhamou. Random walk models in biology. *Journal of the Royal Society Interface*, 5(25):813–834, August 2008. doi: 10.1098/rsif.2008.0014.

[24] Mark Pagel. Inferring the historical patterns of biological evolution. *Nature*, 401 (6756):877–884, 1999. doi: 10.1038/44766. URL https://doi.org/10.1038/44766.

[25] Mark Pagel. The Maximum Likelihood Approach to Reconstructing Ancestral Character States of Discrete Characters on Phylogenies. *Systematic Biology*, 48 (3):612–622, 07 1999. ISSN 1063-5157. doi: 10.1080/106351599260184. URL https://doi.org/10.1080/106351599260184.

[26] G. E. Uhlenbeck and L. S. Ornstein. On the theory of the brownian motion. *Phys. Rev.*, 36:823–841, Sep 1930. doi: 10.1103/PhysRev.36.823. URL https://link.aps.org/doi/10.1103/PhysRev.36.823.

[27] Phylogenetic Comparative Methods (Harmon), 1 2021. URL https://bio.libretexts.org/@go/page/21571. [Online; accessed 2021-10-05].

[28] R I M Dunbar and Susanne Shultz. Evolution in the social brain. *Science*, 317 (5843):1344–1347, Sep 2007. ISSN 1095-9203 (Electronic); 0036-8075 (Linking). doi: 10.1126/science.1145463.

[29] Leslie C. Aiello and Peter Wheeler. The expensive-tissue hypothesis: The brain and the digestive system in human and primate evolution. *Current Anthropology*, 36(2):199–221, 1995. ISSN 00113204, 15375382. URL http://www.jstor.org/stable/2744104.

[30] Mathieu Lihoreau, Tanya Latty, and Lars Chittka. An exploration of the social brain hypothesis in insects. *Frontiers in Physiology*, 3:442, 2012. ISSN 1664-042X. doi: 10.3389/fphys.2012.00442. URL https://www.frontiersin.org/article/10.3389/fphys.2012.00442.

[31] George F. Oster and Edward O. Wilson. Caste and ecology in the social insects. *Acta Biotheoretica*, 28(3):234–235, 1979. doi: 10.1007/BF00046355. URL https://doi.org/10.1007/BF00046355.

[32] Bert Holldobler and Edward O. Wilson. *The superorganism : the beauty, elegance, and strangeness of insect societies*. W.W. Norton & Company, New York, 2009. ISBN 9780393067040 0393067041.

[33] R. Keating Godfrey and Wulfila Gronenberg. Brain evolution in social insects: advocating for the comparative approach. *Journal of Comparative Physiology A*, 205(1):13–32, 2019. doi: 10.1007/s00359-019-01315-7. URL https://doi.org/10.1007/s00359-019-01315-7.

[34] CARL ANDERSON and DANIEL W. McSHEA. Individual versus social complexity, with particular reference to ant colonies. *Biological Reviews*, 76(2): 211–237, 2001. doi: https://doi.org/10.1017/S1464793101005656. URL https://onlinelibrary.wiley.com/doi/abs/10.1017/S1464793101005656.

[35] Andreas Simon Brandstaetter, Wolfgang Rössler, and Christoph Johannes Kleineidam. Friends and foes from an ant brain's point of view – neuronal correlates of colony odors in a social insect. *PLOS ONE*, 6(6):1–9, 06 2011. doi: 10.1371/journal.pone.0021383. URL https://doi.org/10.1371/journal.pone.0021383.

[36] Robin I. M. Dunbar. The social brain hypothesis. *Evolutionary Anthropology: Issues, News, and Reviews*, 6(5):178–190, 2021/10/05 1998. doi: https://doi.org/10.1002/(SICI)1520-6505(1998)6:5<178::AID-EVAN5>3.0.CO;2-8. URL https://doi.org/10.1002/(SICI)1520-6505(1998)6:5<178::AID-EVAN5>3.0.CO;2-8.

[37] Sarah M Farris. Insect societies and the social brain. *Curr Opin Insect Sci*, 15: 1–8, Jun 2016. ISSN 2214-5753 (Electronic). doi: 10.1016/j.cois.2016.01.010.

[38] F Javier Pérez-Barbería, Susanne Shultz, and Robin I M Dunbar. Evidence for coevolution of sociality and relative brain size in three orders of mammals. *Evolution*, 61(12):2811–2821, Dec 2007. ISSN 0014-3820 (Print); 0014-3820 (Linking). doi: 10.1111/j.1558-5646.2007.00229.x.

[39] Karin Isler and Carel P van Schaik. The expensive brain: a framework for explaining evolutionary changes in brain size. *J Hum Evol*, 57(4):392–400, Oct 2009. ISSN 1095-8606 (Electronic); 0047-2484 (Linking). doi: 10.1016/j.jhevol.2009.04.009.

[40] R Keating Godfrey and Wulfila Gronenberg. Linking colony size with foraging behavior and brain investment in odorous ants (formicidae: Dolichoderinae). *Brain Behav Evol*, 95(1):15–24, 2020. ISSN 1421-9743 (Electronic); 0006-8977 (Linking). doi: 10.1159/000504643.

[41] Nick Bos and Patrizia d'Ettorre. Recognition of social identity in ants. *Frontiers in Psychology*, 3:83, 2012. ISSN 1664-1078. doi: 10.3389/fpsyg.2012.00083. URL https://www.frontiersin.org/article/10.3389/fpsyg.2012.00083.

[42] Nathalie Stroeymeyt, Fernando J. Guerrieri, Jelle S. van Zweden, and Patrizia d'Ettorre. Rapid decision-making with side-specific perceptual discrimination in ants. *PLOS ONE*, 5(8):1–8, 08 2010. doi: 10.1371/journal.pone.0012377. URL https://doi.org/10.1371/journal.pone.0012377.

[43] J Frances Kamhi, Iulian Ilieş, and James F A Traniello. Social complexity and brain evolution: Comparative analysis of modularity and integration in ant brain organization. *Brain Behav Evol*, 93(1):4–18, 2019. ISSN 1421-9743 (Electronic); 0006-8977 (Linking). doi: 10.1159/000497267.

[44] Sara Arganda, Andrew P. Hoadley, Evan S. Razdan, Isabella B. Muratore, and James F. A. Traniello. The neuroplasticity of division of labor: worker polymorphism, compound eye structure and brain organization in the leafcutter ant atta cephalotes. *Journal of Comparative Physiology A*, 206(4):651–662, 2020. doi: 10.1007/s00359-020-01423-9. URL https://doi.org/10.1007/s00359-020-01423-9.

[45] R. Keating Godfrey, Mira Swartzlander, and Wulfila Gronenberg. Allometric analysis of brain cell number in hymenoptera suggests ant brains diverge from general trends. *Proceedings of the Royal Society B: Biological Sciences*, 288(1947):20210199, 2021. doi: 10.1098/rspb.2021.0199. URL https://royalsocietypublishing.org/doi/abs/10.1098/rspb.2021.0199.

[46] David A. Forsyth and Jean Ponce. *Computer Vision: A Modern Approach*. Prentice Hall Professional Technical Reference, 2002. ISBN 0130851981.

[47] James Parker. *Algorithms for Image Processing and Computer Vision*. Wiley, 01 1997. ISBN 978-0-471-14056-6.

[48] Richard Szeliski. *Computer Vision: Algorithms and Applications*, volume 5. Springer, 01 2011. ISBN 978-1-84882-934-3. doi: 10.1007/978-1-84882-935-0.

[49] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. http://www.deeplearningbook.org.

[50] Steven L. Horowitz and Theodosios Pavlidis. Picture segmentation by a tree traversal algorithm. *J. ACM*, 23(2):368–388, 1976. ISSN 0004-5411. doi: 10.1145/321941.321956.

[51] Robert M. Haralick and Linda G. Shapiro. Image segmentation techniques. *Computer Vision, Graphics, and Image Processing*, 29(1):100–132, 1985. URL http://dblp.uni-trier.de/db/journals/cvgip/cvgip29.html#HaralickS85.

[52] Liu D.and Soran B.and Petrie G and Shapiro L. A review of computer vision segmentation algorithms. *Lecture Notes*, 53, 2012.

[53] Mojtaba Seyedhosseini, Mehdi Sajjadi, and Tolga Tasdizen. Image segmentation with cascaded hierarchical models and logistic disjunctive normal networks. In *ICCV*, pages 2168–2175. IEEE Computer Society, 2013. ISBN 978-1-4799-2839-2. URL http://dblp.uni-trier.de/db/conf/iccv/iccv2013.html#SeyedhosseiniST13.

[54] Shi Jianbo and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000. ISSN 0162-8828. doi: 10.1109/34.868688.

[55] M. Ranzato, V. Mnih, J. M. Susskind, and G. E. Hinton. Modeling natural images using gated mrfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(9):2206–2222, 2013. ISSN 0162-8828. doi: 10.1109/TPAMI.2013.29.

[56] Bruce Fischl, David H. Salat, André J. W. van der Kouwe, Nikos Makris, Florent Ségonne, Brian T. Quinn, and Anders M. Dale. Sequence-independent segmentation of magnetic resonance images. *NeuroImage*, 23:S69–S84, 2004. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2004.07.016. URL http://www.sciencedirect.com/science/article/pii/S1053811904003817.

[57] Nameirakpam Dhanachandra, Khumanthem Manglem, and Yambem Jina Chanu. Image segmentation using k -means clustering algorithm and subtractive clustering algorithm. *Procedia Computer Science*, 54:764–771, 2015. ISSN 1877-0509. doi: https://doi.org/10.1016/j.procs.2015.06.090. URL https://www.sciencedirect.com/science/article/pii/S1877050915014143. Eleventh International Conference on Communication Networks, ICCN 2015, August 21-23, 2015, Bangalore, India.

[58] Jiawei Han, Micheline Kamber, and Jian Pei. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 3rd edition, 2011. ISBN 0123814790.

[59] Laurent Najman and Michel Schmitt. Watershed of a continuous function. *Signal Processing*, 38(1):99–112, 1994. ISSN 0165-1684. doi: https://doi.org/10.1016/0165-1684(94)90059-0. URL https://www.sciencedirect.com/science/article/pii/0165168494900590. Mathematical Morphology and its Applications to Signal Processing.

[60] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B*, 39(1):1–38, 1977.

[61] Gabriele Steidl. *Supervised Learning by Support Vector Machines*. Springer New York, New York, NY, 2011. ISBN 978-0-387-92920-0. doi: 10.1007/978-0-387-92920-0_22. URL https://doi.org/10.1007/978-0-387-92920-0_22.

[62] Leo Grady. Random walks for image segmentation. *IEEE Trans Pattern Anal Mach Intell*, 28(11):1768–1783, Nov 2006. ISSN 0162-8828 (Print); 0098-5589 (Linking). doi: 10.1109/TPAMI.2006.233.

[63] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1:541–551, 1989.

[64] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013. ISSN 0162-8828. doi: 10.1109/TPAMI.2013.50.

[65] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521 (7553):436–444, 2015. doi: 10.1038/nature14539. URL https://doi.org/10.1038/nature14539.

[66] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986. doi: 10.1038/323533a0. URL https://doi.org/10.1038/323533a0.

[67] Yann Lecun, Leon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.

[68] Alex Graves, Abdel-Rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6645–6649, 2013. doi: 10.1109/ICASSP.2013.6638947.

[69] Niels Justesen, Philip Bontrager, Julian Togelius, and Sebastian Risi. Deep learning for video game playing, 2019.

[70] Ming Y. Lu, Tiffany Y. Chen, Drew F. K. Williamson, Melissa Zhao, Maha Shady, Jana Lipkova, and Faisal Mahmood. Ai-based pathology predicts origins for cancers of unknown primary. *Nature*, 594(7861):106–110, 2021. doi: 10.1038/s41586-021-03512-4. URL https://doi.org/10.1038/s41586-021-03512-4.

[71] Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, 11 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL https://doi.org/10.1162/neco.1997.9.8.1735.

[72] Kunihiko Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36(4):193–202, 1980. doi: 10.1007/BF00344251. URL https://doi.org/10.1007/BF00344251.

[73] D H Hubel and T N Wiesel. Receptive fields and functional architecture of monkey striate cortex. *J Physiol*, 195(1):215–243, Mar 1968. ISSN 0022-3751

(Print); 1469-7793 (Electronic); 0022-3751 (Linking). doi: 10.1113/jphysiol.1968.sp008455.

[74] Ch Bishop. *Neural Networks For Pattern Recognition*, volume 227. Oxford, 01 2005.

[75] Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *CoRR*, abs/1207.0580, 2012. URL http://arxiv.org/abs/1207.0580. cite arxiv:1207.0580.

[76] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996. ISSN 00359246. URL http://www.jstor.org/stable/2346178.

[77] Herbert Robbins and Sutton Monro. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400 – 407, 1951. doi: 10.1214/aoms/1177729586. URL https://doi.org/10.1214/aoms/1177729586.

[78] Yann A. LeCun, Léon Bottou, Genevieve B. Orr, and Klaus-Robert Müller. *Efficient BackProp*, pages 9–48. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. ISBN 978-3-642-35289-8. doi: 10.1007/978-3-642-35289-8_3. URL https://doi.org/10.1007/978-3-642-35289-8_3.

[79] B.T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964. ISSN 0041-5553. doi: https://doi.org/10.1016/0041-5553(64)90137-5. URL https://www.sciencedirect.com/science/article/pii/0041555364901375.

[80] Connor Shorten and Taghi M. Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):60, 2019. doi: 10.1186/s40537-019-0197-0. URL https://doi.org/10.1186/s40537-019-0197-0.

[81] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. URL https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf.

[82] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. doi: 10.1109/CVPR.2016.90.

[83] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2015. doi: 10.1109/CVPR.2015.7298594.

[84] Shervin Minaee, Yuri Boykov, Fatih Porikli, Antonio Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image segmentation using deep learning: A survey, 2020.

[85] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3431–3440, 2015. doi: 10.1109/CVPR.2015.7298965.

[86] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer*, pages 234–241, 2015.

[87] IEEE. Isbi2015, 2015. URL https://biomedicalimaging.org/2015/program/isbi-challenges/.

[88] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: Learning dense volumetric segmentation from sparse annotation. In Sebastien Ourselin, Leo Joskowicz, Mert R. Sabuncu, Gozde Unal, and William Wells, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*, pages 424–432, Cham, 2016. Springer International Publishing. ISBN 978-3-319-46723-8.

[89] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. *CoRR*, abs/1606.04797, 2016. URL http://arxiv.org/abs/1606.04797.

[90] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. *CoRR*, abs/1703.06870, 2017. URL http://arxiv.org/abs/1703.06870.

[91] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *CoRR*, abs/1506.01497, 2015. URL http://arxiv.org/abs/1506.01497.

[92] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. *CoRR*, abs/1612.03144, 2016. URL http://arxiv.org/abs/1612.03144.

[93] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *CoRR*, abs/1706.05587, 2017. URL http://arxiv.org/abs/1706.05587.

[94] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems*

Design and Implementation (OSDI 16), pages 265–283, 2016. URL https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf.

[95] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.

[96] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf.

[97] Keras. Keras, 2015. URL https://keras.io/.

[98] Alaa Tharwat. Classification assessment methods. *Applied Computing and Informatics*, 17(1):168–192, 04 2020. doi: 10.1016/j.aci.2018.08.003. URL https://doi.org/10.1016/j.aci.2018.08.003.

[99] Yann LeCun and Corinna Cortes. MNIST handwritten digit database. http://yann.lecun.com/exdb/mnist/, 2010. URL http://yann.lecun.com/exdb/mnist/.

[100] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research), 2010. URL http://www.cs.toronto.edu/~kriz/cifar.html.

[101] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollar, and Larry Zitnick. Microsoft coco: Common objects in context. In *ECCV*. European Conference on Computer Vision, September 2014. URL https://www.microsoft.com/en-us/research/publication/microsoft-coco-common-objects-in-context/.

[102] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.

[103] M.A. Ranzato, Fu Huang, Y-Lan Boureau, and Yann Lecun. Unsupervised learning of invariant feature hierarchies with applications to object recognition. In *Proceedings / CVPR, IEEE Computer Society Conference on Computer Vision and Pattern Recognition. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1 – 8, 07 2007. doi: 10.1109/CVPR.2007.383157.

[104] Dan C. Ciresan, Ueli Meier, and Jürgen Schmidhuber. Multi-column deep neural networks for image classification. *CoRR*, abs/1202.2745, 2012. URL http://arxiv.org/abs/1202.2745.

[105] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, Levente Lanczi, Elizabeth Gerstner, Marc-André Weber, Tal Arbel, Brian B Avants, Nicholas Ayache, Patricia Buendia, D Louis Collins, Nicolas Cordier, Jason J Corso, Antonio Criminisi, Tilak Das, Hervé Delingette, Çağatay Demiralp, Christopher R Durst, Michel Dojat, Senan Doyle, Joana Festa, Florence Forbes, Ezequiel Geremia, Ben Glocker, Polina Golland, Xiaotao Guo, Andac Hamamci, Khan M Iftekharuddin, Raj Jena, Nigel M John, Ender Konukoglu, Danial Lashkari, JoséAntonió Mariz, Raphael Meier, Sérgio Pereira, Doina Precup, Stephen J Price, Tammy Riklin Raviv, Syed M S Reza, Michael Ryan, Duygu Sarikaya, Lawrence Schwartz, Hoo-Chang Shin, Jamie Shotton, Carlos A Silva, Nuno Sousa, Nagesh K Subbanna, Gabor Szekely, Thomas J Taylor, Owen M Thomas, Nicholas J Tustison, Gozde Unal, Flor Vasseur, Max Wintermark, Dong Hye Ye, Liang Zhao, Binsheng Zhao, Darko Zikic, Marcel Prastawa, Mauricio Reyes, and Koen Van Leemput. The multimodal brain tumor image segmentation benchmark (brats). *IEEE Trans Med Imaging*, 34(10): 1993–2024, Oct 2015. ISSN 1558-254X (Electronic); 0278-0062 (Print); 0278-0062 (Linking). doi: 10.1109/TMI.2014.2377694.

[106] Sook-Lei Liew, Julia M. Anglin, Nick W. Banks, Matt Sondag, Kaori L. Ito, Hosung Kim, Jennifer Chan, Joyce Ito, Connie Jung, Nima Khoshab, Stephanie Lefebvre, William Nakamura, David Saldana, Allie Schmiesing, Cathy Tran, Danny Vo, Tyler Ard, Panthea Heydari, Bokkyu Kim, Lisa Aziz-Zadeh, Steven C. Cramer, Jingchun Liu, Surjo Soekadar, Jan-Egil Nordvik, Lars T. Westlye, Junping Wang, Carolee Winstein, Chunshui Yu, Lei Ai, Bonhwang Koo, R. Cameron Craddock, Michael Milham, Matthew Lakich, Amy Pienta, and Alison Stroud. A large, open source dataset of stroke anatomical brain images and manual lesion segmentations. *Scientific Data*, 5(1):180011, 2018. doi: 10.1038/sdata.2018.11. URL https://doi.org/10.1038/sdata.2018.11.

[107] Kelly Payette, Priscille de Dumast, Hamza Kebiri, Ivan Ezhov, Johannes C. Paetzold, Suprosanna Shit, Asim Iqbal, Romesa Khan, Raimund Kottke, Patrice Grehten, Hui Ji, Levente Lanczi, Marianna Nagy, Monika Beresova, Thi Dao Nguyen, Giancarlo Natalucci, Theofanis Karayannis, Bjoern Menze, Meritxell Bach Cuadra, and Andras Jakab. An automatic multi-tissue human fetal brain segmentation benchmark using the fetal tissue annotation dataset. *Scientific Data*, 8(1):167, 2021. doi: 10.1038/s41597-021-00946-3. URL https://doi.org/10.1038/s41597-021-00946-3.

[108] David Knötel, Ronald Seidel, Steffen Prohaska, Mason N. Dean, and Daniel Baum. Automated segmentation of complex patterns in biological tissues: Lessons from stingray tessellated cartilage. *PLOS ONE*, 12(12):1–24, 12 2017. doi: 10.1371/journal.pone.0188018. URL https://doi.org/10.1371/journal.pone.0188018.

[109] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning

for fast adaptation of deep networks. *CoRR*, abs/1703.03400, 2017. URL http://arxiv.org/abs/1703.03400.

[110] Yoshua Bengio, Yann Lecun, and Geoffrey Hinton. Deep learning for ai. *Commun. ACM*, 64(7):58–65, June 2021. ISSN 0001-0782. doi: 10.1145/3448250. URL https://doi.org/10.1145/3448250.

[111] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'14, page 2672–2680, Cambridge, MA, USA, 2014. MIT Press.

[112] Qing Li, Weidong Cai, Xiaogang Wang, Yun Zhou, David Dagan Feng, and Mei Chen. Medical image classification with convolutional neural network. In *2014 13th International Conference on Control Automation Robotics and Vision, ICARCV 2014*, pages 844–848. Institute of Electrical and Electronics Engineers Inc., March 1997. ISBN 9781479951994. doi: 10.1109/ICARCV.2014.7064414. 2014 13th International Conference on Control Automation Robotics and Vision, ICARCV 2014 ; Conference date: 10-12-2014 Through 12-12-2014.

[113] Maayan Frid-Adar, Idit Diamant, Eyal Klang, Michal Amitai, Jacob Goldberger, and Hayit Greenspan. Gan-based synthetic medical image augmentation for increased CNN performance in liver lesion classification. *CoRR*, abs/1803.01229, 2018. URL http://arxiv.org/abs/1803.01229.

[114] Stanislav Nikolov, Sam Blackwell, Ruheena Mendes, Jeffrey De Fauw, Clemens Meyer, Cían Hughes, Harry Askham, Bernardino Romera-Paredes, Alan Karthikesalingam, Carlton Chu, Dawn Carnell, Cheng Boon, Derek D'Souza, Syed Ali Moinuddin, Kevin Sullivan, DeepMind Radiographer Consortium, Hugh Montgomery, Geraint Rees, Ricky Sharma, Mustafa Suleyman, Trevor Back, Joseph R. Ledsam, and Olaf Ronneberger. Deep learning to achieve clinically applicable segmentation of head and neck anatomy for radiotherapy. *CoRR*, abs/1809.04430, 2018. URL http://arxiv.org/abs/1809.04430.

[115] Alberto Cairo. *The functional art : an introduction to information graphics and visualization*. Berkeley, CA : New Riders, 2013. URL https://search.library.wisc.edu/catalog/9910126981802121.

[116] Amirreza Mahbod, Gerald Schaefer, Benjamin Bancher, Christine Löw, Georg Dorffner, Rupert Ecker, and Isabella Ellinger. Cryonuseg: A dataset for nuclei instance segmentation of cryosectioned h&e-stained histological images. *Computers in Biology and Medicine*, 132:104349, May 2021. ISSN 0010-4825. doi: 10.1016/j.compbiomed.2021.104349. URL http://dx.doi.org/10.1016/j.compbiomed.2021.104349.

[117] Noah F. Greenwald, Geneva Miller, Erick Moen, Alex Kong, Adam Kagel, Christine Camacho Fullaway, Brianna J. McIntosh, Ke Leow, Morgan Sarah

Schwartz, Thomas Dougherty, Cole Pavelchek, Sunny Cui, Isabella Camplisson, Omer Bar-Tal, Jaiveer Singh, Mara Fong, Gautam Chaudhry, Zion Abraham, Jackson Moseley, Shiri Warshawsky, Erin Soon, Shirley Greenbaum, Tyler Risom, Travis Hollmann, Leeat Keren, Will Graf, Michael Angelo, and David Van Valen. Whole-cell segmentation of tissue images with human-level performance using large-scale data annotation and deep learning. *bioRxiv*, 2021. doi: 10.1101/2021.03.01.431313. URL https://www.biorxiv.org/content/early/2021/03/02/2021.03.01.431313.1.

[118] ThermoFisher. Amira, 2020. URL https://www.thermofisher.com/jp/en/home/electron-microscopy/products/software-em-3d-vis/amira-software.html.

[119] Bert Holldobler and Edward O. Wilson. *The ants*. Belknap Press of Harvard University Press, Cambridge, Mass., 1990. ISBN 0674040759 9780674040755.

[120] Sebastian Schwarz, Antoine Wystrach, and Ken Cheng. A new navigational mechanism mediated by ant ocelli. *Biology Letters*, 7(6):856–858, 2011. doi: 10.1098/rsbl.2011.0489. URL https://royalsocietypublishing.org/doi/abs/10.1098/rsbl.2011.0489.

[121] David Orme, R.P. Freckleton, G.H. Thomas, T. Petzoldt, S.A. Fritz, and Nick Isaac. Caper: comparative analyses of phylogenetics and evolution in r. *Methods in Ecology and Evolution*, 3:145–151, 01 2013.

[122] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2013. URL http://www.R-project.org/. ISBN 3-900051-07-0.

[123] Achaz von Hardenberg and Alejandro Gonzalez-Voyer. Disentangling evolutionary cause-effect relationships with phylogenetic confirmatory path analysis. *Evolution*, 67 - 2:378–387, 2013. doi: 10.1111/j.1558-5646.2012.01790.x.

[124] Wouter van der Bijl. phylopath: Easy phylogenetic path analysis in r. *bioRxiv*, 2017. doi: 10.1101/212068. R package version 1.0.0.

[125] Rüdiger Wehner, Tsukasa Fukushi, and Karin Isler. On being small: brain allometry in ants. *Brain Behav Evol*, 69(3):220–228, 2007. ISSN 1421-9743 (Electronic); 0006-8977 (Linking). doi: 10.1159/000097057.

[126] Ofer Feinerman and James F. A. Traniello. Social complexity, diet, and brain evolution: modeling the effects of colony size, worker size, brain size, and foraging behavior on colony fitness in ants. *Behavioral Ecology and Sociobiology*, 70(7):1063–1074, 2016. doi: 10.1007/s00265-015-2035-5. URL https://doi.org/10.1007/s00265-015-2035-5.

[127] Darcy Greer Gordon, Alejandra Zelaya, Ignacio Arganda-Carreras, Sara Arganda, and James F. A. Traniello. Division of labor and brain evolution in insect

societies: Neurobiology of extreme specialization in the turtle ant cephalotes varians. *PLOS ONE*, 14(3):1–16, 03 2019. doi: 10.1371/journal.pone.0213618. URL https://doi.org/10.1371/journal.pone.0213618.

[128] Zach N. Coto and James F. A. Traniello. Brain size, metabolism, and social evolution. *Frontiers in Physiology*, 12:184, 2021. ISSN 1664-042X. doi: 10.3389/fphys.2021.612865. URL https://www.frontiersin.org/article/10.3389/fphys.2021.612865.

[129] Francisco Hita Garcia, Georg Fischer, Cong Liu, Tracy L. Audisio, Gary D. Alpert, Brian L. Fisher, and Evan P. Economo. X-ray microtomography for ant taxonomy: An exploration and case study with two new terataner (hymenoptera, formicidae, myrmicinae) species from madagascar. *PLOS ONE*, 12(3):e0172641, 2017. doi: 10.1371/journal.pone.0172641. URL https://doi.org/10.1371/journal.pone.0172641.

[130] Thomas van de Kamp, Patrik Vagovič, Tilo Baumbach, and Alexander Riedel. A biological screw in a beetle's leg. *Science*, 333(6038):52–52, 2011. doi: 10.1126/science.1204245. URL http://science.sciencemag.org/content/sci/333/6038/52.full.pdf.

[131] Willi Ribi, Tim J. Senden, Arthur Sakellariou, Ajay Limaye, and Shaowu Zhang. Imaging honey bee brain anatomy with micro-x-ray-computed tomography. *Journal of Neuroscience Methods*, 171(1):93–97, 2008. ISSN 0165-0270. doi: https://doi.org/10.1016/j.jneumeth.2008.02.010. URL http://www.sciencedirect.com/science/article/pii/S0165027008001179.

[132] Kleoniki Keklikoglou, Sarah Faulwetter, Eva Chatzinikolaou, Nikitas Michalakis, Irene Filiopoulou, Nikos Minadakis, Emmanouela Panteri, George Perantinos, Alexandros Gougousis, and Christos Arvanitidis. Micro-ct(vlab): A web based virtual gallery of biological specimens using x-ray microtomography (micro-ct). *Biodiversity Data Journal*, 1 Nov 2016(4):e8740, 2016. ISSN 1314-2828. doi: 10.3897/BDJ.4.e8740. URL http://www.ncbi.nlm.nih.gov/pmc/articles/PMC5139143/.

[133] Hoo-Chang Shin, H. R. Roth, M. Gao, L. Lu, Z. Xu, I. Nogues, J. Yao, D. Mollura, and R. M. Summers. Deep convolutional neural networks for computer-aided detection: Cnn architectures, dataset characteristics and transfer learning. *IEEE Transactions on Medical Imaging*, 35(5):1285–1298, 2016. ISSN 0278-0062. doi: 10.1109/TMI.2016.2528162.

[134] Dan C. Cireşan, Alessandro Giusti, Luca M. Gambardella, and Jürgen Schmidhuber. Mitosis detection in breast cancer histology images with deep neural networks. In Kensaku Mori, Ichiro Sakuma, Yoshinobu Sato, Christian Barillot, and Nassir Navab, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2013*, pages 411–418. Springer Berlin Heidelberg, 2013. ISBN 978-3-642-40763-5.

[135] Claire McQuin, Allen Goodman, Vasiliy Chernyshev, Lee Kamentsky, Beth A. Cimini, Kyle W. Karhohs, Minh Doan, Liya Ding, Susanne M. Rafelski, Derek Thirstrup, Winfried Wiegraebe, Shantanu Singh, Tim Becker, Juan C. Caicedo, and Anne E. Carpenter. Cellprofiler 3.0: Next-generation image processing for biology. *PLOS Biology*, 16(7):1–17, 07 2018. doi: 10.1371/journal.pbio.2005970. URL https://doi.org/10.1371/journal.pbio.2005970.

[136] Chensi Cao, Feng Liu, Hai Tan, Deshou Song, Wenjie Shu, Weizhong Li, Yiming Zhou, Xiaochen Bo, and Zhi Xie. Deep learning and its applications in biomedicine. *Genomics, Proteomics & Bioinformatics*, 16(1):17–32, 2018. ISSN 1672-0229. doi: https://doi.org/10.1016/j.gpb.2017.07.003. URL https://www.sciencedirect.com/science/article/pii/S1672022918300020.

[137] Hao Chen, Qi Dou, Lequan Yu, Jing Qin, and Pheng-Ann Heng. Voxresnet: Deep voxelwise residual networks for brain segmentation from 3d mr images. *NeuroImage*, 170:446–455, 2018. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2017.04.041. URL http://www.sciencedirect.com/science/article/pii/S1053811917303348.

[138] R. Vaillant, C. Monrocq, and Y. Le Cun. Original approach for the localization of objects in images. *IEE Proceedings: Vision, Image and Signal Processing*, 141 (4):245–250, August 1994. ISSN 1350-245X. doi: 10.1049/ip-vis:19941301.

[139] Brian B. Avants, Nicholas J. Tustison, Gang Song, Philip A. Cook, Arno Klein, and James C. Gee. A reproducible evaluation of ants similarity metric performance in brain image registration. *NeuroImage*, 54(3):2033–2044, 2011. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2010.09.025. URL https://www.sciencedirect.com/science/article/pii/S1053811910012061.

[140] Philipp D. Lösel, Thomas van de Kamp, Alejandra Jayme, Alexey Ershov, Tomáš Faragó, Olaf Pichler, Nicholas Tan Jerome, Narendar Aadepu, Sabine Bremer, Suren A. Chilingaryan, Michael Heethoff, Andreas Kopmann, Janes Odar, Sebastian Schmelzle, Marcus Zuber, Joachim Wittbrodt, Tilo Baumbach, and Vincent Heuveline. Introducing biomedisa as an open-source online platform for biomedical image segmentation. *Nature Communications*, 11(1):5577, 2020. doi: 10.1038/s41467-020-19303-w. URL https://doi.org/10.1038/s41467-020-19303-w.

[141] Bruce Fischl. Freesurfer. *NeuroImage*, 62(2):774–781, 2012. ISSN 1053-8119 1095-9572. doi: 10.1016/j.neuroimage.2012.01.021. URL http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3685476/.

[142] Konstantinos Kamnitsas, Christian Ledig, Virginia F. J. Newcombe, Joanna P. Simpson, Andrew D. Kane, David K. Menon, Daniel Rueckert, and Ben Glocker. Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation. *Medical Image Analysis*, 36:61–78, 2017. ISSN 1361-8415. doi: https://doi.org/10.1016/j.media.2016.10.004. URL http://www.sciencedirect.com/science/article/pii/S1361841516301839.

[143] Daniël M. Pelt and James A. Sethian. A mixed-scale dense convolutional neural network for image analysis. *Proceedings of the National Academy of Sciences*, 2017. doi: 10.1073/pnas.1715832114. URL http://www.pnas.org/content/pnas/early/2017/12/21/1715832114.full.pdf.

[144] I. Arganda-Carreras, D. G. Gordon, S. Arganda, M. Beaudoin, and J. F.A. Traniello. Group-wise 3d registration based templates to study the evolution of ant worker neuroanatomy. In *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pages 429–432, 2017. doi: 10.1109/ISBI.2017.7950553.

[145] Sabrina Amador-Vargas, Wulfila Gronenberg, William T. Wcislo, and Ulrich Mueller. Specialization and group size: brain and behavioural correlates of colony size in ants lacking morphological castes. *Proceedings of the Royal Society B: Biological Sciences*, 282(1801):20142502, 2015. doi: 10.1098/rspb.2014.2502. URL https://royalsocietypublishing.org/doi/abs/10.1098/rspb.2014.2502.

[146] Daniel Baum, James C Weaver, Igor Zlotnikov, David Knötel, Lara Tomholt, and Mason N Dean. High-throughput segmentation of tiled biological structures using random-walk distance transforms. *Integrative and comparative biology*, 59 (6):1700–1712, 12 2019. doi: 10.1093/icb/icz117. URL https://pubmed.ncbi.nlm.nih.gov/31282926.

[147] Rafael C. González and Richard Woods. *Digital Image Processing, Third Edition*. Pearson, 2008.

[148] Alexander B. Jung. imgaug. https://github.com/aleju/imgaug, 2018. [Online; accessed 30-Oct-2018].

[149] Joel Akeret, C. Chang, A. Lucchi, and A. Refregier. Radio frequency interference mitigation using deep convolutional neural networks. *Astronomy and Computing*, 18:35–39, 2017. ISSN 2213-1337. doi: https://doi.org/10.1016/j.ascom.2017.01.002. URL https://www.sciencedirect.com/science/article/pii/S2213133716301056.

[150] Lee R. Dice. Measures of the amount of ecologic association between species. *Ecology*, 26:297–302, 1945.

[151] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen A.W.M. van der Laak, Bram van Ginneken, and Clara I. Sanchez. A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42:60–88, 2017. ISSN 1361-8415. doi: https://doi.org/10.1016/j.media.2017.07.005. URL https://www.sciencedirect.com/science/article/pii/S1361841517301135.

[152] Diogo M. Camacho, Katherine M. Collins, Rani K. Powers, James C. Costello, and James J. Collins. Next-generation machine learning for biological networks. *Cell*, 173(7):1581–1592, 2018. ISSN 0092-8674. doi: https://doi.org/10.1016/

j.cell.2018.05.015. URL https://www.sciencedirect.com/science/article/pii/S0092867418305920.

[153] Felix Dujardin. Memoire sur le systeme nerveux des insectes. *Ann Sci Nat Zool*, 14:195–206, 1850.