Learning to Study the Developing Brain: Characterising Brain Development with Machine Learning, Image Analysis and *Ex Vivo* Culture

Thesis submitted for the degree of:

DPhil in Biomedical Imaging

Martin Hailstone

Linacre College, University of Oxford

Trinity Term 2018

Under the supervision of:

Professor Ilan Davis

Dr Richard M. Parton

DECLARATION

I hereby declare that this thesis is my own work and the experiments described in the following pages were performed by myself unless otherwise stated. The experiments were carried out in the Department of Biochemistry, University of Oxford. This thesis has not been submitted for any other degree or professional qualification

The following parts of the thesis contain data from others:

1. Figure 3.4 B Imaging and Measurements of NB cell cycle by Lu Yang (Davis lab)
2. Figure 4.10 A,B,C,D Preparation and Imaging performed by Lu Yang.(Davis lab)
3. Figure 5.3 Preparation and Imaging performed by Tamsin Samuels (Davis lab)
4. Figure 6.1 Images from Eldred et al. 2017 under CC-BY license
5. Figure 6.2 Preparation and Imaging performed by Eldred et al. 2017
6. Figure 6.4 Preparation and Imaging performed by Ita Costello (Roberston lab)
7. Figure 6.5 Preparation and Imaging performed by Ben Kroeger (Wilson lab)

Martin Hailstone

04.10.2018

ABSTRACT

Brain development, and indeed development in general, is complex. Correct development requires tight regulation of cell numbers, movements and fates. However, decisions are taken by individual cells, integrating spatial and temporal cues. Understanding development, therefore, necessitates studying cells in the native tissue context. In this thesis, I focus on neurogenesis, studying neural stem cell division in cultured whole *Drosophila* brains. To facilitate this, I develop an approach based on live imaging, image analysis and the use of a novel software tool that is applicable to studying development in a range of systems. I start with optimised imaging, culture and image processing to study cell divisions in *Drosophila* brains. Building on previous work by L. Yang, I develop a pipeline allowing high resolution, medium-throughput imaging of multiple brains. Using patch-based denoising, and gradient-descent-based image registration, I follow divisions of stem cells and their daughters in an intact brain, in a way that was previously unachievable. This allowed me to record cell behaviours, such as cell cycle length, which I show are consistent with published results. However, analysing the resulting datasets was highly time-consuming, and impractical on larger scales. Existing automated analysis tools proved to be inadequate, likely due to signal-to-noise ratio and complex structures in these images. To overcome this, I developed CytoCensus, image analysis software for complex 4D data, in collaboration with D. Waithe. Unlike other software, which aims to identify whole cells and their boundaries, CytoCensus identifies only cell centres, allowing it to cope with neighbouring cells and low signal-to-noise. I show that CytoCensus identifies different cell types in the brain, performing better than competing approaches. Combining the aforementioned live imaging and culture methods with CytoCensus analysis, provided insights into normal and aberrant development. I characterise changes in the development mutant *syp*, in which both neuroblasts and daughter cells exhibit an increased division rate, accounting for the observed increase in brain size. I further demonstrate that this approach is generally applicable to studying developmental processes in other contexts by quantifying transcription factors in mouse embryos, counting organelles in spermatogenesis, and comparing cell distributions in zebrafish retinal organoids. Taken together, it is clear that this approach, and these tools can be used to study the development of different organisms, by studying the behaviour of cells in a relevant context.

TABLE OF CONTENTS

Acknowledgments

I would like to thank my supervisor Ilan, for his support and advice over the last four years. For giving me the freedom to research what I found interesting, and the guidance to make it into something useful.

Richard, for all your advice and for at various times being a supervisor, an editor, a colleague, and a friend. For struggling through what might one day be a published paper, and then doing it again.

The Davis lab, for their guidance, feedback and help at various times, especially Tamsin, MK, Ana, Josh, and Fran who have been there for my whole time (ish). And also for all the cake. Darragh for all the little things that make a lab run, and the important things like the Christmas party.

The Costello lab, for sitting through my lab meetings, and for all their feedback.

All my collaborators, without whom this work would not be possible. Dominic Waithe for all his work on QBrain/CytoCensus. Lu Yang for her work on live imaging. Ita Costello and Ben Kroeger for their interest in CytoCensus and their willingness to let me test it on their data. The Harris lab, for their images of organoids.

Micron, especially Alan, David, and Andrew for all their help with microscopes, and dealing with my stupid amounts of data and complaints about the air-conditioning.

All my friends here in Oxford: Katie, Anne, Cat (even Jonathan), for food, friendship, and many good times in the last 4 years. All my ONBI friends who made DTC courses fun and worthwhile, and who have been a constant throughout my research.

Cat, Hannah, and Sarah for reading what might one day be a thesis.

My family; Mum, for always being there, even when I'm focused on thesis, and always being an encouragement. Gem, for being my sister, and a welcome distraction from work. Dad, for many a barbecue, and for always wanting to discuss what I've been working on.

Finally, Sarah, for all the tea, hugs, being a great housemate, and support writing this thesis.

Abbreviations

A/H	Amacrine/Horizontal
AL	Antennal Lobe
BCM	Brain Culture Medium
BP	Bipolar
BrdU	5-bromo-2'-deoxyuridine
CB	Central Brain
COM	Centre of Mass
CT	Computed Tomography
DAPI	4',6-diamidino-2-phenylindole dihydrochloride
DetH	Determinant of the Hessian
DNA	Deoxyribose Nucleic Acid
EdU	5-ethynyl-2′-deoxyuridine
FACS	Fluorescence-Activated Cell Sorting
FLP	Flippase
FN	False Negative
FP	False Positive
FRST	Fast Radial Symmetry Transform
FRT	FLP Recognition Target
GFP	Green Fluorescence Protein
GGM	Gaussian Gradient Magnitude
GMC	Ganglion Mother Cell
GT	Ground Truth
GUI	Graphical User Interface
INP	Intermediate Progenitor
LMP	Low Melting Point
LoG	Laplacian of Gaussian
MARCM	Mosaic Analysis with a Repressible Cell Marker
MB	Mushroom Body

MRI	Magnetic Resonance Imaging
NB	Neuroblast
NN	Neural Network
OL	Optic Lobe
PBS	Phosophate Buffered Saline
PR	Photoreceptor
PS	Primitive Streak
RFP	Red Fluorescence Protein
RG	Retinal Ganglion
RNA	Ribose Nucleic Acid
RNAi	RNA interference
ROI	Region of Interest
RPE	Retinal Pigment Epithelium
SC	Secondary Cell
SNR	Signal to Noise Ratio
SoFa	Spectrum of Fates
SPIM	Selective Plane Illumination Microscopy
STE	Structure Tensor Eigenvalues
TAE	Tris-Acetate-EDTA
TBE	Tris-Borate-EDTA
TF	Transcription Factor
TP	True Positive
UAS	Upstream Activation Sequence
VE	Visceral Endoderm
VNC	Ventral Nerve Cord
WT	Wild Type

CHAPTER 1

INTRODUCTION

1.1. OVERVIEW OF DROSOPHILA BRAIN DEVELOPMENT AND NEUROGENESIS

1.1.1. *DROSOPHILA* AS A MODEL SYSTEM

The fruit fly *Drosophila melanogaster* has been used as a model system to understand cellular processes, and the process of development since the beginning of the 20th century (Stephenson and Metcalfe, 2013). It is small enough that the whole *Drosophila* can be imaged, but large enough that more complex patterning is required than for the nematode *Caenorhabditis elegans* (*C. elegans*). Furthermore, it is relatively easy to manipulate genetically. As a model organism, *Drosophila* is frequently used to understand mammalian processes. It has been variously used for the study of disease (two thirds of genes involved in disease in humans have an analogue in *Drosophila)*, cancer, and development both at the cellular and molecular level (Bellen et al., 2010; Henrique and Bally-Cuif, 2010; Rudrapatna et al., 2012; Wodarz and Gonzalez, 2006). Despite such interest in using *Drosophila* to model vertebrate processes, *Drosophila* development differs significantly from vertebrates, for instance in embryogenesis (**Figure 1.1**). Nevertheless, there are common features between *Drosophila* and vertebrates. These similarities exist both at the cellular level in terms of lineages and patterning mechanisms, and at the molecular level in terms of signalling pathways (Ammeux et al., 2016; Homem and Knoblich, 2012). Hence, understanding development in *Drosophila* has, and will continue to, provide further insights into mammalian development. Of all organs, the human brain remains the most mysterious; we do not fully

understand its function, its structure, or its development. In this thesis, I will focus on using *Drosophila* as a model of brain development. I specifically focus on using a live imaging approach to study the dynamics of cells within the brain, and on using image analysis tools in order to make automated analysis of such imaging practical.

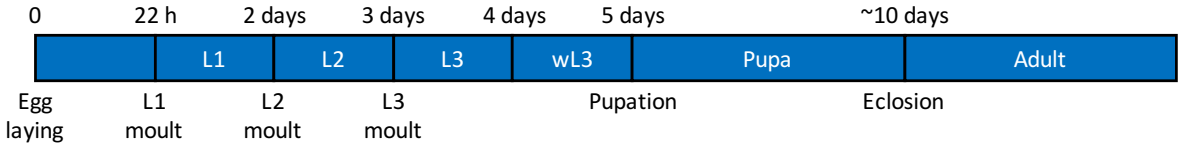

Figure 1.1: Development of *Drosophila*. In normal laboratory culture, eggs develop for around 22 h, before hatching into L1 larvae. L1 larvae feed for around 24 h before a pulse of ecdysone triggers a 'larval moult' into an L2 larva that can be seen as a morphological change. After a further 24 h, the larva moults into the L3 stage. It ceases feeding at around 24 h of L3, and climbs out of the food (at this stage it is considered a wandering L3 larva (wL3). After a further 24 h, a final ecdysone pulse triggers the larva to begin pupation, which lasts for around 4-5 days before finally eclosing to become an adult fly. Based on (Bainbridge and Bownes, 1981; Homem and Knoblich, 2012).

1.1.1.1. The Drosophila Brain as a Model for Brain Development

Historically, the study of human brain development was confined to post-mortem dissection and histology (Dobbing and Sands, 1973; Sidman and Rakic, 1973). More recently the widespread use of magnetic resonance imaging (MRI) and computed tomography (CT) has made measurements of internal structures more straightforward (Stiles and Jernigan, 2010), although not at a cellular level. The use of human cell cultures has furthered our understanding at the cellular level (Ahmed, 2009), but the clinical relevance of these has often been poor (Rudrapatna et al., 2012). Therefore, there is a need for the use of model systems. By understanding such simpler systems, we can hope to extract principles that we can then use to understand more complex systems.

The human brain consists of approximately 100 million neurons, whereas the adult *Drosophila* brain contains roughly one million neurons (Homem and Knoblich, 2012; Muotri and Gage,

2006). Nevertheless, this is significantly more complex than simpler models such as *C. elegans* which has a mere 302 neurons in the whole worm (Jabr, 2018). Collectively, these neurons make billions of specific synaptic connections in the human brain. This complexity has led to interest in studying the mechanisms of brain development in *Drosophila*, both to understand normal brain development, and also to understand where it goes wrong, resulting in disease and cancer (Miles et al., 2011).

In the context of development, *Drosophila* has proven to be a very useful model. The regulatory mechanisms that control development and patterning are largely conserved between humans and *Drosophila* (Noctor et al., 2004). Some vertebrate neural stem cells, and their *Drosophila* equivalent (known as neuroblasts) undergo analogous asymmetric divisions (Homem and Knoblich, 2012; Wodarz and Gonzalez, 2006). These cells produce similar lineages, although vertebrate lineages are typically larger and consist of more 'transit amplifying' lineages (i.e. there is a second level of propagating cells (Homem and Knoblich, 2012). Furthermore, *Drosophila* neural stem cells are regulated by orthologous transcription factors (Brand and Livesey, 2011), and show temporal determination of cell fate (Homem and Knoblich, 2012). As with all stem cell divisions, a balance between renewal and differentiation is required. When this balance is disturbed, it can result in underdevelopment or tumorigenesis, such as in *lethal giant larva*e (*lgl*) and *brain tumour* (*brat*) (Rudrapatna et al., 2012; Wodarz and Gonzalez, 2006). By studying these mutant phenotypes in *Drosophila*, we aim to gain an understanding of the mechanisms by which stem cell behaviour is regulated. Knowledge of these mechanisms can be used to understand the development of diseases, such as brain cancer, in humans.

1.1.2. Neuroblasts as a Model for Neural Stem Cells

In *Drosophila*, the majority of cells that make up the brain arise from self-renewing divisions of stem cells known as neuroblasts (NBs). "Stem cells" are cells with a potentially unlimited proliferative capacity, the progeny of which can also differentiate into multiple types of normal

functional cells (Weissman et al., 2001). Pluripotent stem cells, such as those from the inner cell mass of mammalian blastocysts, have the potential to form cells from the three germ layers (endoderm, ectoderm, and mesoderm) and, hence, a whole organism. Multipotent stem cells maintain the potential for unlimited self-renewal, but are more restricted in what cells they can form. Progenitor cells, by contrast, may have the potential to form multiple cell types but with limited ability to self-renew.

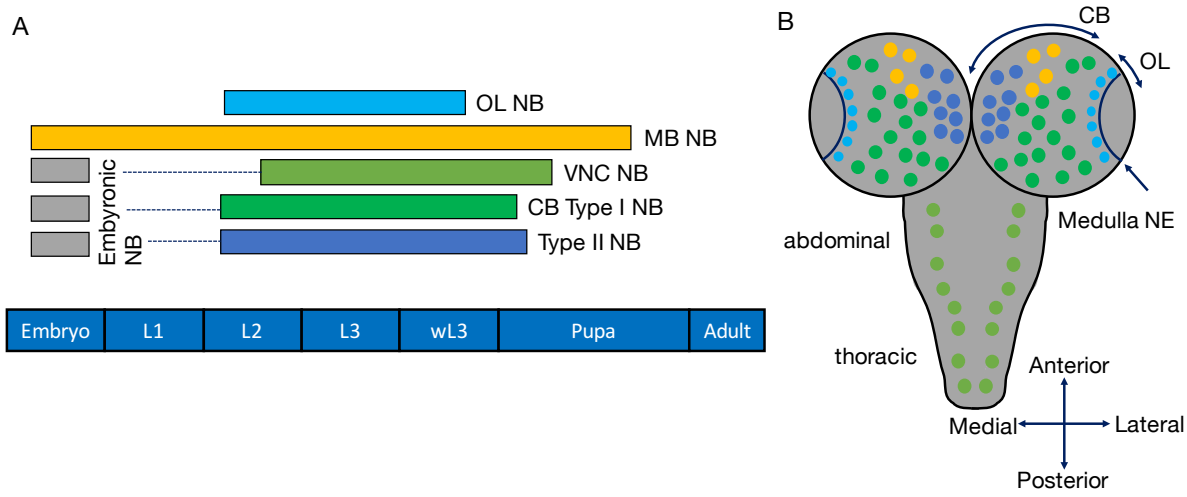

Figure 1.2: Embryonic and larval neuroblast divisions. A) Timings of NB divisions. Following the formation of NBs around embryo stage 9, embryonic NBs divide until L1, where they become quiescent until they resume division in L2. After puparium formation they cease division. Mushroom body (MB) NBs are independent of the normal signals for quiescence and continue dividing for much longer. B) 2D diagram of the Drosophila central nervous system (CNS) showing NB locations. 1) Location of the **Optic Lobe** (OL) NB; 2) Type I **Central Brain** (CB) NBs; 3) **Type II** NBs; 4) ventral nerve cord (**VNC)** NB (including abdominal, thoracic and gnathal NB) Medulla and Neuroepithelium (NE). Based on (Homem and Knoblich, 2012), (Pinto-Teixeira et al., 2016), (Bainbridge and Bownes, 1981).

Neuroblasts are often considered the stem cells of the *Drosophila* brain. There are three types that give rise to all the neurons and some glia that make up the adult brain: Types 0, I, and II (**Figure 1.2**). In *Drosophila* larvae they go through multiple rounds of self-renewing divisions, but they do not usually persist to adulthood. Embryonic NBs express a sequential series of transcription factors that define the neurons that they produce, and do not regrow between

divisions, which is atypical of a stem cell (Li et al., 2014). However, there is accumulating evidence that a similar, if not entirely analogous, process of sequential transcription factors happens in larval neuroblasts. They are therefore perhaps more correctly referred to as stem cell-like progenitors, as a reflection of their progressively restricted potential.

1.1.2.1. NEUROBLAST TYPES 0, I AND II.

Embryonic NBs arise in the embryo from the neuroepithelium by lateral inhibition mechanisms that are well studied (Artavanis-Tsakonas and Simpson, 1991; Skeath and Thor, 2003). They quickly begin a series of rounds of asymmetric division (Reichert, 2011). In Type I and II NBs, each division gives rise to a daughter ganglion mother cell (GMC), which itself will divide symmetrically to form neurons or glia; Type 0 NBs directly give rise to neurons (**Figure 1.3**). After each round of division, the neuroblast shrinks, and the transcription factors that are expressed change (**Figure 1.3E**). This is linked to a change in the identity of the GMC and the final fate of the neurons that are produced (Yu et al., 2010). These embryonic divisions give rise to around 10,000 neurons of the larval brain that are key to coordination of larval movement and response to external stimuli.

Following this rapid series of divisions, neuroblasts either undergo quiescence, terminally divide or undergo programmed cell death except for the MB neuroblasts (**Figure 1.2**, **Figure 1.3A**)(Pinto-Teixeira et al., 2016). Quiescent neuroblasts regrow and resume divisions at around L2 in response to nutrient signals. They continue dividing, but in a different regime, in which they regrow between divisions (**Figure 1.3C, D**)(Homem and Knoblich, 2012). These NBs can be divided into Type I and Type II NBs (**Figure 1.2**, **Figure 1.3D**). In this thesis I will focus on these larval NBs, particularly the Type I NBs of the CB (**Figure 1.2**, **Figure 1.3C**).

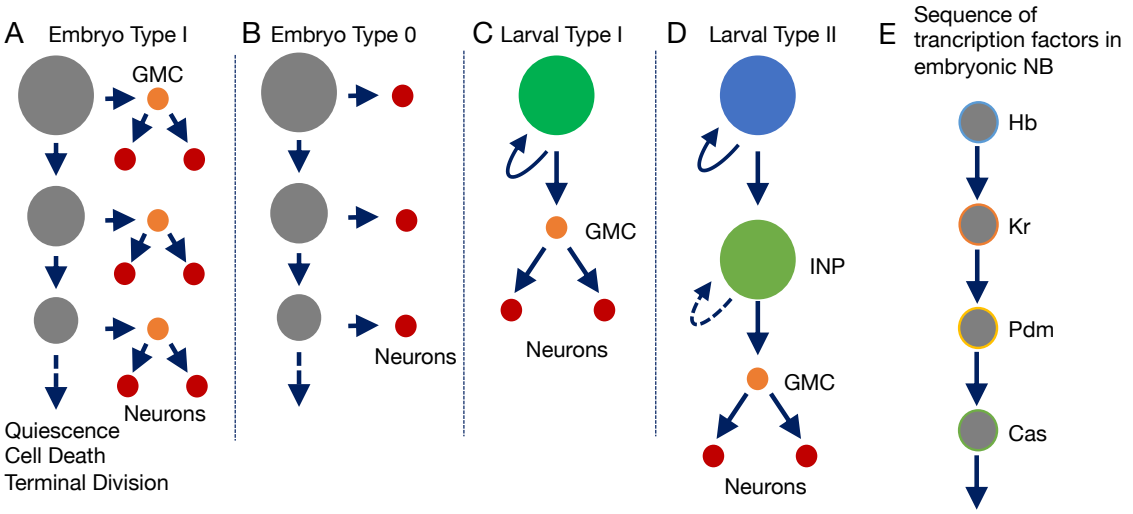

Figure 1.3: Transcription factor progression in NB divisions in *Drosophila*. A) Most embryonic NBs divide to give daughter GMCs but do not regrow, and ultimately enter quiescence, undergo programmed cell death, or terminal division. B) Some embryonic NBs switch to dividing as Type 0 NBs, directly giving rise to neurons. C) Most larval NBs divide as Type I NBs, similar to that of embryonic Type I NBs but self-renew, regrowing after each division. D) Type II NBs undergo self-renewing divisions giving rise to Intermediate Progenitor cells (INPs) a process that is similar to mammalian lineages with transit amplifying cells. E) Embryonic NBs undergo a sequential transition through a defined set of transcription factors (TFs): Hb (Hunchback), Kr (Kruppel), Pdm (POU domain), Cas (Castor). Based on (Homem and Knoblich, 2012).

Similar to the transcription factor cascade described above for embryonic NBs (**Figure 1.3E**), larval NBs appear to go through a series of changes in their expressed transcription factors. Although a transcription factor cascade has not been identified, recently, it has been shown that there are opposing temporal gradients of the RNA binding proteins Imp and Syncrip (Liu et al., 2015). Overexpression of Imp or Syncrip is sufficient to induce early or late fates of neurons, respectively, at least in two lineages (Mushroom Body (MB) and Antennal Lobe (AL) NBs), as well as controlling decommissioning of NBs (Liu et al., 2015; Yang et al., 2017). Subsequently, experiments explored this relationship in more detail, finding a transition from Chinmo/Imp/Lin-28+ to Syncrip/Broad/E93 expression, that appears to be regulated by the hormone ecdysone (Syed et al., 2017). Perhaps unlike the previous cascade of factors, cell fate

relies on gradients which change over time, as well temporally controlled responses to external factors.

Although great effort has been put into identifying these TFs, we have little direct knowledge of their effects on individual CB NB behaviour. For instance, what effect do these TFs have on the division rate of NBs? Similarly, our knowledge of the variation between NBs is primarily limited to easy to distinguish lineages (Type II, which are marked by Worniu, and MB and AL neuroblasts which can be marked by intersectional strategies), and without further markers it is difficult to explore the central brain (Homem et al., 2013; Liu et al., 2015; Yu et al., 2010). Although the MB and AL NBs can be individually followed, what of the remaining lineages? Do they fall into further groups? Further understanding of these differences could be aided by live imaging approaches.

1.2. Overview of Using Imaging to Study Development

In complex tissues, such as the *Drosophila* brain, the processes of cell proliferation and differentiation lead to a vast and diverse population of neurons and glia (Kohwi and Doe, 2013). Elucidating the molecular basis of such developmental processes is not only essential for basic neuroscience, but is also important for discovering new treatments for neurological diseases and cancer. Modern imaging approaches have proven indispensable in addressing the underlying molecular mechanisms of development, with the ability to image intact zebrafish and *Drosophila* tissues at the cellular level (Barbosa and Ninkovic, 2016; Cabernard and Doe, 2013; Dray et al., 2015; Graeden and Sive, 2009; Medioni et al., 2015).

In this section I will provide a brief overview and comparison of fixed imaging with *in vitro*, *in vivo*, and *ex vivo* live imaging methods. I will specifically focus on methods that are appropriate for studying brain development within *Drosophila*. I exclude those that are better suited to larger scale imaging such as MRI, CT and ultrasound, instead focusing on microscopy methods and protocols for cellular resolution imaging (Gregg and Butcher, 2012).

1.2.1. Fixed Tissue Imaging Approaches to Studying Development

Much of the imaging in *Drosophila* to study development has been of fixed tissues, whether those are directly dissected from larvae, or from *in vitro* culture. This has a number of advantages which make imaging more straightforward. Unlike in live imaging, phototoxicity is not an issue. Furthermore, the fixation process usually makes use of detergents that remove much of the fat from tissues. This can be advantageous as it reduces the scattering effect of tissues which is problematic for imaging thick specimens. Additionally, immunohistochemical staining can be used straightforwardly. The availability of antibodies against a huge number of different proteins allows a level of investigation into cell identity that is difficult to match with live imaging, which requires the more time consuming and specific production of genetic constructs and transgenic fly lines. Immunohistochemistry also allows the use of more stable fluorophores (such as Alexa dyes) that are more difficult to introduce in a live setting. As a result of these advantages, and also for historical reasons, fixed imaging has been key to our understanding of *Drosophila* brain development. In particular, a range of methods (genetic and otherwise) have been developed that allow dynamic processes such as cell division and migration to be recorded using fixed imaging approaches, although these are not without limitation. In the rest of this section, I will consider the contribution of fixed imaging to our understanding of *Drosophila* brain development, with a specific focus on methods that provide information on NB lineages and divisions, both of which are major processes that I explore further in this thesis.

1.2.1.1. MARCM Clones for Cell Lineage Analysis

Mosaic and clonal analysis methods were developed in *Drosophila* to allow the study of cell lineages. Although many variants of MARCM exist, one of the most common is MARCM (Mosaic Analysis with a Repressible Cell Marker), which allows the study of single cells by employing spatial and temporal control over expression (**Figure 1.4**), (Bowman et al., 2008; Lee and Luo, 2001; Yu et al., 2010).

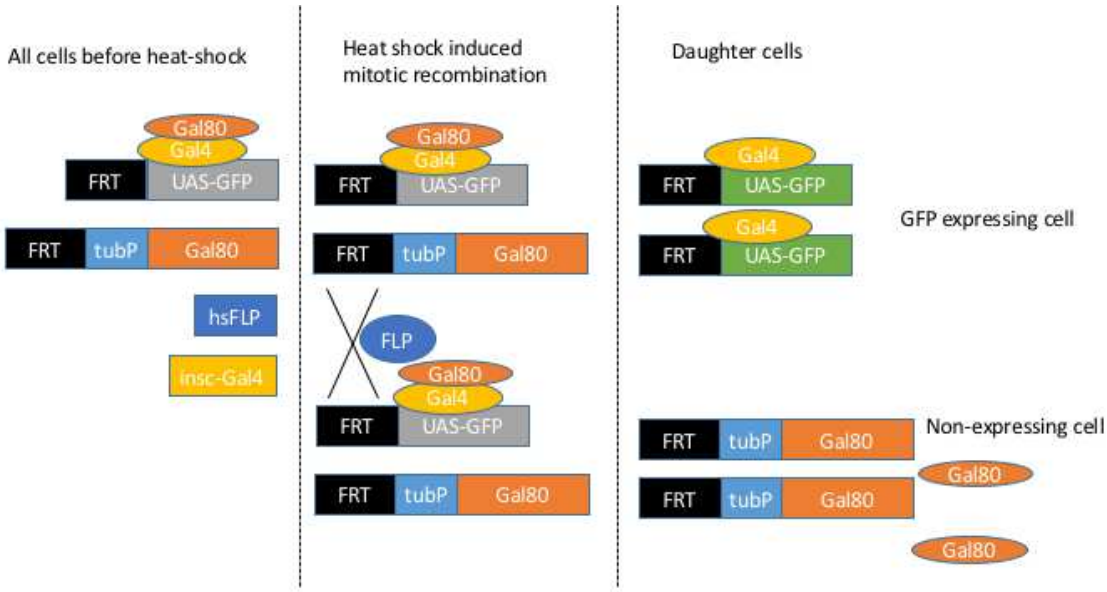

Figure 1.4: In MARCM, all cells are heterozygous for a UAS driven marker (here, GFP), a ubiquitously driven GAL80 repressor, heat-shock FLP recombinase (hsFLP) and a GAL4 driver (here the *inscuteable* promoter). The Gal80 represses the activation of GFP expression by GAL4. Upon heat-shock, FLP (flippase, a recombinase derived from *Saccharomyces cerevisiae* that recognises short flippase recognition target (FRT) sequences) is produced and induces mitotic recombination at the FRT sites. A proportion of the resulting daughter cells no longer express the GAL80 repressor, and therefore GAL4 expression activates the UAS and causes GFP expression, marking that cell and all of its progeny.

By marking all cells resulting from a lineage, MARCM was used to identify that there were two key classes of NBs (Type I and Type II, see Section 1.1) (Bowman et al., 2008). Additionally, MARCM can be used to make estimates of cell division rate, by inducing clones (through a temperature change) and then fixing, staining, and counting progeny. Using this approach, (Bowman et al., 2008) report estimates of the average Type I NB division rate between 1.36 h and 1.74 h, and GMC cell cycle length of 8-10 hours. By the nature of the experiment, however, this is an average. Additionally, in order to achieve temporal specificity of induction, temperature changes are required that may influence the processes under investigation as well as adding variability to measurements. I will explore measurement of division rate in **Chapter 5**.

1.2.1.2. Following Cell Cycle with PH3/EdU/BrdU

There are a number of markers for following the cell cycle in fixed tissue imaging. Phospho-Histone H3 marks chromosomes in mitosis. By calculating the percentage of cells expressing Histone H3 over a period, a mitotic index can be calculated (**Figure 1.5**). Based on the fact that the length of mitosis is generally considered fixed, this allows for an estimation of the division rate of cells. This measurement also assumes that the population of cells over which the average is calculated behave homogenously. In the case of NBs, it is known that this is not true (Bowman et al., 2008; Liu et al., 2015).

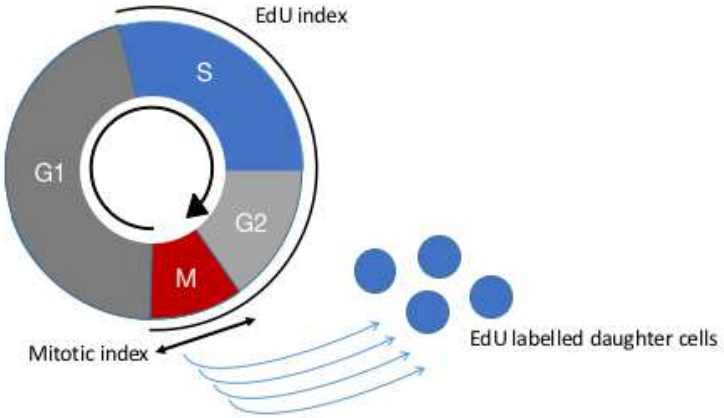

Figure 1.5: The cell cycle. G1- Growth phase 1. S- DNA Synthesis G1- Growth Phase 2 M- Mitosis EdU and BrdU are incorporated at the point of DNA synthesis (in S-phase). Hence all cells that pass through S-phase during the labelling period are labelled. This can be used to make an estimate of the mitotic index (technically not just of cells in mitosis), based on the proportion of cells that are labelled. Additionally, by using a longer pulse, the incorporated DNA is segregated into daughter cells. By using a pulse of known length, the number of daughter cells containing the labelled DNA can be measured, and the division rate estimated over this time.

EdU (5-ethynyl-2′-deoxyuridine) and BrdU (5-bromo-2'-deoxyuridine) incorporation can also be used to produce a mitotic index of the proportion of cells that have undergone DNA synthesis. Of course, such labels are known to be toxic, so there is some question of whether longer pulses are an accurate representation of the *in vivo* situation. While such approaches

are limited, EdU labelling has been used to measure clone sizes in (Carney et al., 2013), which can be used to estimate NB division rate. I will consider these measurements of division rate in **Chapter 5**.

1.2.2. Live Imaging

Live imaging has, in the past, provided significant insights into the development in a wide range of organisms (Gregg and Butcher, 2012; Pantazis and Supatto, 2014). In principle, live imaging can provide information on the dynamics of processes that occur during development that are difficult or impossible to achieve using fixed methods (Moraru et al., 2012). Live imaging can capture (natural) variation in the dynamics of cells, and answer questions of inter- and intra- sample variability that could not be determined otherwise. Live imaging can also be used to understand very fast processes such as calcium transients, and/or those processes without markers that can be easily fixed. Live imaging methods, like development, cover a range of scales: from *in vitro* imaging of single cells, to *in vivo* imaging of whole organisms over hours and days (Huang et al., 2018; Lerit et al., 2014; Pantazis and Supatto, 2014).

1.2.2.1. Imaging *In Vitro*

Live imaging of individual cells is typically done *in vitro* using either primary cells isolated from tissues or from immortalised cell lines. The ability to make large numbers of cells easily makes immortalised cell lines attractive, but they are often poorly representative of the *in vivo* situation (Carter, 2015). Primary cells isolated from tissues provide a more realistic representation of the *in vivo* situation, as well as the advantages of being an isolated population of cells. They can be sorted by FACS for transcriptomic studies or imaged at high resolution on glass slides, for instance (Berger et al., 2012; Harzer et al., 2013). However, primary cells are still isolated from their native environment; on the one hand, this provides an excellent system in which to study the effect of external factors, but on the other, there are questions over how well this recapitulates *in vivo* development.

The ability to make primary cultures of neural lineages in *Drosophila* has existed since 1983. However, it was not until much more recently that live imaging measurements of the cell cycle were performed, with most measurements of mitotic index performed on whole populations from disaggregated cells (Ceron et al., 2006; Wu et al., 1983). Early measurements of the cell cycle, and in particular GMC cell cycle length, were observed to be aberrantly short (15 minutes) (Ceron et al., 2006; Savoian and Rieder, 2002). Nonetheless, using FACS-sorted NBs, Berger et al. measured the NB cell cycle to be 1.58 ± 0.25 hours, and the GMC cell cycle to be 4.6 ± 0.9 hours (Berger et al., 2012; Harzer et al., 2013). Subsequently, Homem *et al.* aimed to make measurements more representative of the *in vivo* environment. Using a 3D culture, imaging and automated analysis setup using disaggregated NBs, they measured cell cycle rates of 4.2 hours for GMCs, and 1.3 hours for Type I NBs (Homem et al., 2013).

The absence of native environment makes isolated cells something of an artificial context and raises questions over whether such experiments are representative of the *in vivo* situation or not (Pampaloni et al., 2007). These issues associated with studying cells *in vitro*, have made the study of cells in a more native-like environment highly attractive. Organoids are rapidly becoming a popular way of understanding the requirements of tissue and organ development (Eldred et al., 2017a; Simian and Bissell, 2017). They trade off some of the advantages of individual cells (e.g. not needing to deal with thick tissues), in exchange for a (hopefully) more *in vivo*-like environment, which retains the ability to dissect the essential components of development. Of course, developments in organoids would not be possible without first an understanding on how to culture cells, nor can we be sure they are accurate without a comparison to the *in vivo* situation.

1.2.2.2. IMAGING *IN VIVO*

Studying cells directly *in vivo* avoids the issues of *in vitro* systems. However, *in vivo* imaging presents a new set of difficulties. Imaging *in vivo* usually necessitates imaging at depth, which

is always a challenge; the movement of organisms makes this a more difficult task still. Embryo development has been imaged for a variety of organisms. This is by far the easiest stage to image, primarily due to its small size (which makes imaging more straightforward, see Section 1.2.2) and relatively simple culture (Gregg and Butcher, 2012; Keller et al., 2011; Moraru et al., 2012). However, by the time a zebrafish embryo has reached 24 hours old, it consists of more than 20,000 cells, by which point cellular resolution imaging is a challenge (Keller et al., 2011). Similarly, in *Drosophila*, since the discovery of GFP, cellular dynamics of the embryo have been imaged extensively (Davis et al., 1995; Parton et al., 2010). Its small size and unusual multinucleate situation makes imaging relatively straightforward during embryogenesis and bespoke light-sheet microscopy has allowed high resolution imaging until hatching when the emerging larva "walks off the coverslip" (Keller et al., 2011). This graphically illustrates the problems of *in vivo* imaging of organisms that move.

Various approaches have been developed that compensate for or isolate the movement of different organisms, as well as dealing with other problems, such as opacity. In mice, 'windows' in the skull are a popular way to image the brain. A similar idea has allowed imaging past the opaque exoskeleton of *Drosophila* and other insects while they respond to stimuli (Huang et al., 2018; Pilz et al., 2018). Conversely, anaesthesia can be used to immobilise organisms, in some cases for extended periods with relatively few ill effects. Of course, in mammals, the necessity of maintaining circulation means that some movement can remain. In *Drosophila*, it is even possible to stop the heart (albeit briefly) allowing exceptional imaging for a short time with full recovery (Zhang et al., 2010). Ultimately though, long-term anaesthesia prevents feeding and will cause the organism to die without additional input. Hence, *in vivo* imaging of larval brain development has been limited to either short timescales or not at cellular resolution.

1.2.2.3. IMAGING *EX VIVO*

Although *in vitro* approaches account for most live imaging, *ex vivo* approaches, i.e. the culture of a tissue or organ outside an organism, are also used for studying development. Depending on the method used, they have the 'cells-in-context' advantages of *in vivo* imaging, combined with some of the easier imaging and control over conditions that *in vitro* culture provides. Typically, however, *ex vivo* cultured tissue has limited viability, due to problems with culture conditions, including lack of circulation, inadequate nutrients and growth factors, or bacterial contamination.

Due to the opacity *Drosophila* pupae which makes *in vivo* imaging difficult, there has been much interest in *ex vivo* culture of the pupal brain. Initially, imaging of neuronal maturation in the *Drosophila* pupal brain was performed for only for short periods of time, imaging under oil under a transparent membrane (Medioni et al., 2015). Recent advancements have simplified the mounting process, removing the need for custom machined parts, and improved the viability of the brain using agar and a brain culture medium rather than oil. These approaches support imaging for longer periods, but only at coarse time intervals (4 hours) (Rabinovich et al., 2015). Similar protocols for imaging larval brains have been developed, although they typically require complex imaging protocols and use short imaging durations, and/or coarse time intervals that make the study of NB behaviours challenging (Cabernard and Doe, 2013; Lerit et al., 2014; Wu and Luo, 2006). A more detailed comparison of methods for larval *ex vivo* brain culture and imaging is covered in **Chapter 3.**

1.2.3. IMAGING TRADE-OFFS

Imaging is inherently a trade-off between the amount of light used, the resolution, the time for acquisition, image contrast and the signal to noise ratio (SNR).

1.2.3.1. Microscopy Methods – A Brief Overview

Different microscopy methods make these trade-offs in different ways, and even within an imaging setup, settings can be optimised for speed, sensitivity, resolution and SNR. Confocal imaging, for instance, typically allows for excellent lateral and good axial resolution, by excluding out-of-focus light, at a cost of long acquisition times and high light dose. Spinning disk confocal imaging can increased the speed of imaging compared to point-scanning confocal imaging, but usually achieves reduced resolution. SPIM (selective plane illumination microscopy) imaging allows for much reduced light dose, but at the cost of decreased resolution, and potential shadowing artefacts, particularly in live tissues (Goldman et al., 2010).

1.2.3.2. Trade-Offs in Live Imaging

In live imaging the trade-off is primarily between the optimal imaging parameters and the viability of the tissue. In fixed imaging of shallow tissues, the images are limited by the quantum yield of the fluorophore and its resistance to bleaching (Cranfill et al., 2016). However, in thicker tissues, they also become limited by the ability of light to penetrate through the layers of tissue. Whilst fixation and clearing methods can alleviate many of these penetration problems, in live imaging this is not possible (Richardson and Lichtman, 2015). Live tissues contain significant amounts of lipids, which can lens and scatter light. The effect of this is readily apparent in confocal imaging; it is generally difficult to image more than about 40 μm into tissues. Various methods have been developed to circumvent this problem. The two most common are light-sheet imaging approaches (although resolution is reduced) and multiphoton imaging (which requires expensive instrumentation and is limited in flexibility) (Dong et al., 2013; Huisken and Stainier, 2009). In addition to this imaging difficulty, live imaging presents another limiting factor: light can cause toxicity, either by the production of free radicals, a process catalysed by the fluorophores themselves or by other mechanisms (Gottschalk et al., 2015; Kim et al., 2015). This limits the dose that can be safely applied to

tissues while retaining biologically meaningful results. However, long term live imaging at low temporal resolution does have a minor advantage over fixed imaging: the inherent turnover of proteins means that bleached proteins are eventually degraded and replaced with new ones, limiting the effect of bleaching to some extent. I will explore these issues further in **Chapter 2**.

1.3. Overview of the Importance of Image Analysis in Dealing with Modern Live Imaging Data

1.3.1. Classical Image Analysis

The fields of image analysis, and subsequently computer vision have grown exponentially with the arrival of digital imaging approaches. The ability to collect, store and retrieve large amounts of imaging data, has both enabled, and necessitated the development of methods capable of measuring and interpreting them in a range of situations. Image analysis is now widespread in astronomy, photography, and in microscopy. In this section I will consider methods and tools for extracting useful information from images, with a focus on dealing with microscopy data, and, more specifically, live imaging time-series.

1.3.1.1. Image Segmentation and Detection

Image analysis methods can be broadly split into those performing one of two tasks: segmentation or detection. In biology, the focus has typically been on "segmenting" objects, that is defining their extent or boundary (**Figure 1.6B**)(Kang et al., 2009; Zhang et al., 2008). With appropriate imaging, segmentation can provide a wealth of information on cell shape, migration, membrane movements, and collective behaviours (Mikut et al., 2013; Stegmaier et al., 2016; Toyoshima et al., 2016). In contrast, "detection" of objects, seeks only to identify the location of an object (or a region containing the object) and provides more limited information about object location, and movement, without considering its extent (**Figure 1.6C**) (Gould et al., 2009). Cell detection in biological images, however, has usually been considered as a segmentation problem (Schmitz et al., 2014). In such cases cell/object detection is achieved

by first segmenting an image, and then identifying objects within the segmentation to make detections (Gould et al., 2009).

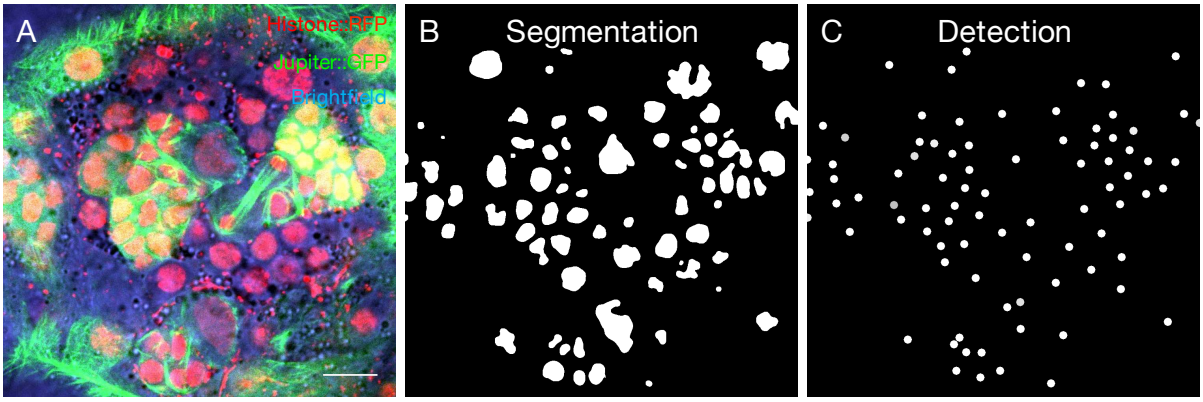

Figure 1.6: Image Segmentation and Detection. A) Image of a region of the Drosophila brain (Scale bar 15μm) B) Possible nuclear segmentation C) Detections of nuclei.

One of the early successes of computer vision was face detection. Face detection, which is a common first step in analysis of TV footage for instance, relies on finding only the location of a face. This task turns out to be simpler and achieves better accuracy than segmenting a complex object like a face (Dang and Sharma, 2017). Recent use of direct object detection methods to find nuclei have shown excellent performance on complex images (Shuvaev et al., 2017; Toyoshima et al., 2016). However, they have yet to be incorporated into biologically relevant analysis tools and, hence, are not in regular use.

1.3.1.2. Image Post-Processing

In high signal-to-noise situations, simple objects can be easily segmented by thresholding an image (Gonzalez and Woods, 2002). In fixed images, with bright, specific fluorophores, this is straightforward. However, the nature of imaging biological samples often makes this task more complex, particularly when it comes to live imaging (Stegmaier et al., 2016). On the one hand, imaging with high intensity lasers is inherently damaging to samples; on the other, higher intensity lasers typically give a better SNR (Gottschalk et al., 2015). When combined with the requirement to image 3D volumes over extended periods, it is necessary to trade off

some signal in return for reduced toxicity and a better representation of the biology (Stephens, 2003). In practice then, the ability of post-processing to deal with noisy images would allow one to deal with biology at better time and imaging resolution, while still producing a useful analysis.

Image post-processing steps are intended to improve image quality after acquisition by correcting for the limitations of different microscopy methods. For instance, the resolution of basic wide-field imaging systems can be dramatically improved by deconvolution (Kim and Naemura, 2015). Complex microscopy systems, for instance various super-resolution microscopy methods, may rely on built-in post-processing steps in order to form useful images (De Luca et al., 2013; Gustafsson, 2000). Here, I consider these image post-processing steps in preparation for further analysis (Dufour et al., 2017), to be distinct from processing that intends to quantitate aspects of the images (such as cell number and location, covered in Section 1.3.1.3) . Consideration must be given to the ultimate aims of the image analysis in all processing steps, including post-processing, or else risk losing important information.

Table 1.1: Examples of post-processing methods applied to live imaging.

Method	Aim	Organism/Tissue	Reference
Image Fusion	Cell tracking	Human Kidney Organoids	(Held et al., 2018)
Image Fusion + Deconvolution	Lineage Reconstruction	*Parhyale hawaiensis*	(Wolff et al., 2018)
Image Contrast Enhancement	Membrane contour tracking	*Drosophila* wing imaginal disc	(Heller et al., 2016)
Deconvolution	Cell cycle timing	*C. elegans* (whole)	(Keil et al., 2017)
Denoising	Lineage reconstruction	*C. elegans* embryo	(Kruger et al., 2015)
Denoising + Deconvolution	Neural activity measurement	Zebrafish	(Pnevmatikakis et al., 2016)
Deconvolution	RNA localisation	*Drosophila* embryo	(Trcek et al., 2017)

Post-processing methods generally aim to correct for errors due to physical phenomena in the imaging of specimens and can be roughly grouped into four categories (**Table 1.1**) (Dufour et al., 2017). (1) **Image fusion** deals with the reconciliation of multiple imaging modes or acquisitions. Image fusion is common in light-sheet microscopy because multiple-views circumvent some of the limitations of imaging thick and large specimens such as scattering and absorption (de Medeiros et al., 2015; Pitrone et al., 2013; Preibisch et al., 2014). (2) **Deconvolution** seeks to correct the optical limitations of microscopes by removing the blurring effects of the point spread function (either estimated from the data, or purely based on the optical system) (Griffa et al., 2018; McNally et al., 1999; Soulez et al., 2012). (3) **Image contrast enhancement** aims to correct for errors that result in a loss of contrast, usually due to the effects of scattering, especially in thick tissues. However, where possible, ramping laser intensity with depth is likely to be a more effective method of correcting for loss of signal (Heller et al., 2016; Kervrann et al., 2004; Rigaut and Vassy, 1991; Zuiderveld, 1994). (4) **Denoising** corrects for various sources of noise, including detector noise, and photon shot noise caused by the stochastic emission of photons from fluorophores (Boulanger et al., 2010; Meiniel et al., 2018). Simple image filtering, for instance with Gaussian or Median filters can

be used for denoising, as well as more complex methods (Gonzalez and Woods, 2002). Correct application of these post-processing methods for image fusion, deconvolution, contrast enhancement, and denoising can transform data that would otherwise be unsuitable for analysis. As a result, they are often incorporated into image analysis pipelines for segmentation or detection.

1.3.1.3. Software for Segmentation and Detection

Even after post-processing methods are applied, imperfections, noise and artefacts remain in images. Ideally, detection and segmentation methods need to be robust to these issues, although image quality may be a fundamental limitation. Hundreds of algorithms have been developed with the aim of allowing efficient processing of such images. It is less common, however, for such algorithms to be presented for convenient use. Some algorithms are patented, others are presented without example code, which makes implementation more difficult. Even among those that are available in software packages, most are unavailable to biologists who do not have prior programming and image analysis knowledge.

Various programs and packages are available that make image analysis more available and approachable. The FARSIGHT toolbox aims to bring together various methods for analysing images into a single package. However, it is primarily aimed at image analysts, and it still requires programming to use (Roysam et al., 2008). CellProfiler goes one step further, largely removing the need for programming (although not image analysis) skills, by presenting a graphical interface for making analysis pipelines (Stoter et al., 2013). ImageJ/FIJI takes a different approach, providing a graphical interface with which to directly run (and see the results of) image filters and analysis. This is ideal for those with limited image analysis knowledge, but also provides a basic scripting interface, suitable for larger scale analysis (Schindelin et al., 2012). FIJI also has a significant community of developers who develop plugins, providing additional functionality. On the basis that microscopy is now so common

that there is a need for image analysis in an everyday setting, by people who are not dedicated image analysts or programmers, I will focus on reviewing programs that, like FIJI, target this audience. Specifically, I look at those that provide a graphical interface, allow users with limited image analysis skills to perform analysis, and do so using conventional image analysis methods (see **Section 1.3.2** for a comparison to machine learning tools).

1.3.1.4. Limitations of Current Software for Segmentation and Detection

Software programs such DeadEasy seek to allow a user to apply a series of filters and methods (such as median filtering, watershed, and binary erosion) without the need for explicit programming, in order to identify cells of a particular type (Forero et al., 2009). However, DeadEasy, as with many conventional image analysis approaches, is somewhat 'fragile', requiring specific programs for different types of cells and labels (Forero et al., 2009). As a case in point, DeadEasy has 3 different versions available for characterising different tissue types (DeadEasy Caspase, DeadEasy Mito-Glia, DeadEasy Neurons) (Forero et al., 2010a; Forero et al., 2010b; Forero et al., 2009). In addition to the susceptibility of these approaches to changes in imaging conditions, this highlights a common issue with conventional, segmentation-based analysis methods: they require specific markers, or combinations thereof in order to successfully analyse an image.

Typically, algorithms rely either on clear nuclear markers to identify cells, or on cell membrane markers to define cell extents, although many cell markers may not meet such strict criteria. There are hundreds of algorithms available for nuclear segmentation and detection, but almost all of them rely on level set, watershed, region growing and morphological operations that assume ideal nuclear or boundary information within the images (Irshad et al., 2014; Xing and Yang, 2016). This focus on nuclear and boundary information is reflected in available tools such as MINS, RACE, CellSegm and DeadEasy (Forero et al., 2009; Hodneland et al., 2013; Lou

et al., 2014; Stegmaier et al., 2016), despite the fact that there are a wealth of other markers and information available (e.g. actin, cell cycle markers)(Sakaue-Sawano et al., 2008; Wernike et al., 2016). The complexity of other markers, however, makes them difficult to segment, and precludes their application to general purpose tools for standard analysis. Where such analysis is required, the application is typically highly specific, and requires highly specific custom code in order to analyse (Mikut et al., 2013).

When markers and imaging quality are appropriate, the specificity of conventional image analysis approaches allows for high precision, tailored segmentations that can be used to answer simple questions. (Lou et al., 2014). However, the specificity in image requirements, the fragility of these methods, and sensitivity to image quality, have led to a search for approaches that are more robust to changes, less specific to markers, and capable of analysing more complex problems. Breakthroughs in machine learning have the potential to make development of such analysis more straightforward, and allow the analysis of data with such limitations. I consider these approaches in the following section.

1.3.2. Machine Learning Concepts in Biological Image Analysis

1.3.2.1. An Overview

In this section I will consider machine learning-based approaches to image analysis in contrast to the conventional approaches I have just described. The origin of the term **machine learning** is claimed by A. L. Samuel, with reference to the idea of generalising learning in a game of checkers, and describes the concept of computers solving tasks without explicit programming (Samuel, 1959). Machine learning can be broadly divided into **unsupervised** and **supervised** methods. Unsupervised methods seek to find patterns and structures in data without external "supervision" from humans (Kotsiantis et al., 2007). Put simply, they can be used to answer questions of the kind "What groups/clusters/patterns exist in this data?". The identity of the groups in usually unknown, but can be assigned after-the-fact. Supervised

methods on the other hand, aim to predict an output based on (usually human) supervision. This supervision, and the corresponding output can take many forms: from simple objects counts to individual pixel labels (Kotsiantis et al., 2007). Supervised methods can answer many different questions depending on the training data, for instance "How many objects are there/Where are they in this data?". Supervised methods rely on having training data with corresponding annotations (labels).

Both supervised and unsupervised machine learning approaches have been applied to biology in various circumstances: for instance, unsupervised methods are relatively common in genetics and genomics (Libbrecht and Noble, 2015), whilst supervised approaches have been applied widely applied in image processing applications, from movies of rat behaviour to identifying cancer (Cruz and Wishart, 2017; Gris et al., 2017). In the image analysis field, DeepYeast (Pärnamaa and Parts, 2017), a supervised approach, has been used to classify the distribution of proteins (e.g. cytoplasmic, golgi-associated etc.) in images based on thousands of training examples.

1.3.2.2. LIMITATIONS

Whilst the applications of machine learning are broad, they are not without limitations. First, by definition, they are highly dependent on the data they learn from and hence susceptible to bias. For supervised learning, because of the reliance on user annotation, there is a high potential for bias: if the training set is unrepresentative, the machine learning may fail to generalise to the remaining data and give spurious results. Similarly, if the annotations contain systematic errors, then these will be perpetuated. As a result, it is important to make quality annotations in a representative dataset. For some problems, where an analytical solution exists, the analytical solution may be preferable because its properties are known and understood. For others, such as fluorescence quantification, it is crucial to be able to understand the transformations and assumptions made, something that machine learning

methods are currently poor at. Other methods require vast amounts of data to train, which in a biological setting may be difficult to achieve. Whilst these issues may be a problem, there is also a lack of tools in biology that allow biologists to make the most of machine learning in a flexible way.

1.3.2.3. Supervised Machine Learning in Image Analysis

In this thesis, I will focus on supervised approaches in image analysis, which are considered more mature, and easily allow specific questions to be asked of data (Sommer and Gerlich, 2013). Images can be very complex, and their interpretation may require knowledge from other experiments or imaging conditions. The ability of the user to assign importance or meaning in supervised approaches makes them better suited to analysis of typical microscopy images, although unsupervised approaches may be of more interest in large scale phenotypic screens because they can recognise previously unknown groups and patterns.

1.3.2.4. Evaluating Machine Learning Performance

Machine learning is closely associated with the idea of a "**ground truth**". The term comes from aerial and satellite imaging of forests. Verifying estimates of forest cover based on the satellite imagery requires someone to put boots on the ground and make direct counts of trees to verify the estimates. In this way "ground truth" has come to mean a verifiable truth of a situation. This often corresponds to human annotation of images, although it is sometimes based on the results of other approaches, and in the case of synthetic data, it is the known segmentation (or other basis for the synthetic data). The performance of machine learning is compared to this ground truth (usually assuming it is "perfect"), to determine "**accuracy**". In this way, it is possible to compare performances across different algorithms or approaches.

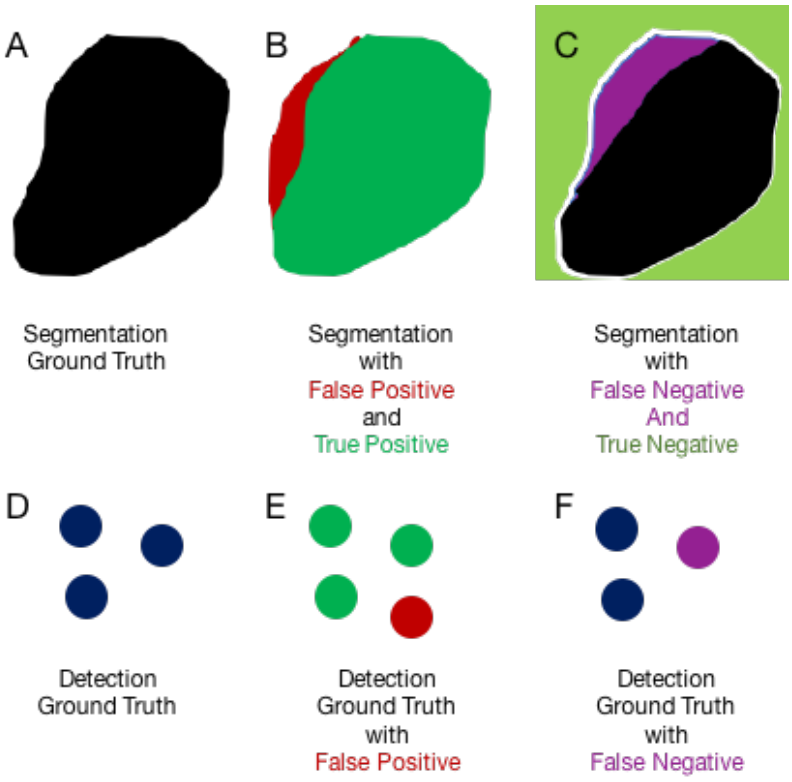

Figure 1.7: Ground truth, Segmentation, Detections. A) Segmentation Ground Truth; B) Segmentation with additional area compared to the ground truth (False Positive); C) Segmentation with missing area (False Negative) and correctly excluded area (True Negative); D) Detection Ground Truth; E) Additional Detection (False positive); F) Missing Detection (False Negative). For detections there is no true negative.

Machine learning approaches may be compared using different metrics of accuracy. In this thesis, I primarily consider **precision** and **recall** due to the nature of **cell detection** problems (**Figure 1.7**). To understand precision and recall, consider that, unlike in image segmentation, in cell detection problems, there are no true negatives. True negatives (TN) are correctly called results that are not positives. In segmentation, this is easily visualised as pixels that are correctly identified as not belonging to the image (**Figure 1.7C**).

$$(1) \qquad Precision = \frac{TP}{TP + FP}$$

$$(2) \qquad Recall = \frac{TP}{TP + FN}$$

$$(3) \qquad F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$$

Precision represents the fraction of objects that are correctly classified, as a proportion of the total identified. It therefore combines measurement of true positives (TPs) and false positives (FPs) (Equation 1). Recall, on the other hand, represents the fraction of objects that are correctly identified out of all correct objects (Equation 2). These two measurements can be combined into a score, the F1-score, that represents a weighted mean of the two measurements (Equation 3). The F1-score is a useful single measure for comparing approaches, although it assumes equal importance to both scores, which may or may not be desirable depending on the problem at hand.

1.3.2.5. Classification and Regression

Supervised machine learning for image analysis typically performs **classification**. In image classification tasks, the aim is to determine the class to which an image belongs. For examples, in the ImageNet challenge, classification can be to determine whether an image is of a car, a tree or another class (Russakovsky et al., 2014). To perform well, the classifier must learn a "**model**" that can cope with natural variations of the object, such as orientation and pose. This model can take different forms depending on the machine learning algorithm used. Classification is also typically used in image segmentation (**Section 1.3.1.3**) to determine the label (class) for each pixel. Depending on the task, these labels might refer to cell types (e.g. NB, GMC), cell components (e.g. membrane, nucleus, cytoplasm) or whole structures (e.g. central brain, optic lobe).

Classification problems contrast with regression problems, which are those with a continuous quantity that needs to be predicted. For instance, a patient's risk (a continuous quantity) of complications from a surgery or treatment can be estimated using regression (Cruz and Wishart, 2017; Kang et al., 2015). Regression is rarely used in image processing applications because most problems are considered as segmentation, for which regression is poorly suited. However, there are a number of problems that can be framed as regression. It is possible, for instance, to learn the distance to a boundary in a segmentation task (for instance, the distance to the nearest cell-edge), rather than directly segmentation. Some success has been found with considering problems in this manner. In particular, cell counting has been successfully considered as a regression task (Kainz et al., 2015). Here, a 'density map' is learned, in which each pixel represents the cell density (the number of cells per pixel). Most cells are represented by multiple pixels, so this is a fractional number. However, when summed up over the entire image, this gives the total number of cells. This has been shown to be more successful than segmenting and detecting each cell individually using traditional classification approaches, as it avoids problems of over- and under-segmentation (Kainz et al., 2015).

1.3.2.6. Machine Learning in Modern Microscopy

The scale of microscopy datasets is increasing. While it can be straightforward to look at a single image and say 'there are 3 cells here', when asked to do that 100's of times it can be very time consuming. This is an excellent application for supervised machine learning, because it is straightforward to ask a user to annotate a small number of images that can be used for training. The machine learning model trained on these images can then extrapolate the annotations to a much larger dataset. Thus, machine learning has the potential to greatly simplify dataset annotation.

Table 1.2: Machine learning based software for biological image analysis.

Program	Purpose	GUI	Language	Image Dimension	Reference
Ilastik (pixel classification)	Segmentation	YES	Python	4D	(Sommer et al., 2011)
FIJI WEKA	Segmentation	YES	C/Java	3D (4D with macros)	(Arganda-Carreras et al., 2017)
CellProfiler	Image analysis pipelines (Segmentation)	YES	MATLAB	3D or 2D + time	(Kamentsky et al., 2011)
CytoMine	Annotation/ Segmentation	Web server	HTML/ JavaScript	4D	(Rubens et al., 2018)
KNIME	Programming pipelines	NO	JAVA	N/A	(Dietz and Berthold, 2016)
CellExplorer	Modelling	YES	Python	4D	(Johnson et al., 2017)
CellCognition	Segmentation and lineages	NO	Python	2D + time	(Held et al., 2010)
SuRVoS	Segmentation	YES	Python	3D	(Luengo et al., 2017)
ACC	Phenotype characterisation	YES	MATLAB	3D	(Piccinini et al., 2017)

A number of programs have exploited machine learning for a variety of different purposes in microscopy (**Table 1.2**). Of those, only two can deal with large scale 3D time-lapse microscopy images: Ilastik and FIJI/WEKA. Ilastik provides a series of workflows for different tasks built around machine learning. Of these, the "pixel classification" workflow is appropriate for analysing 4D images. It uses a segmentation based approach, automatically predicting the "pixel identity" (which cell type a pixel belongs to) based on the given examples (Sommer et al., 2011). WEKA is a platform for image analysis and machine learning, which has a graphical user interface (GUI) built into the popular image processing program FIJI (Arganda-Carreras

et al., 2017). It utilises a similar segmentation-based approach to Ilastik, but has the advantage of being able to utilise the existing plugins and extensibility of FIJI. Although it does not natively support 4D images, with some knowledge of macros, it is not difficult to use FIJI to extend its native 3D functionality.

1.3.2.7. Comparison of Approaches for Analysing Live Imaging

As I have discussed, between machine learning and conventional analysis approaches there are a multitude of programs for analysing biological images (**Table 1.2**, Section 1.2). Live imaging data, however, presents problems of scale and dimension, which many programs are not able to cope with. Conventional image analysis programs, such as RACE, aim to plug this gap, allowing efficient segmentation of live imaging data. However, they can be inflexible, restricting the use of markers and imaging conditions. Machine learning approaches, such as Ilastik and WEKA, aim to circumvent these problems by making use of user supervision but these programs almost exclusively rely on a segmentation-based approach. As a result, they can struggle with complex markers and highly dense cells, like those found in complex tissues. I will explore this issue Chapters 4 and 5, detailing the development of accessible software that uses a supervised machine learning-based approach for direct object detection.

1.4. Aims

First, the work described in this thesis aims to develop ex vivo live imaging and analysis methods to study stem cells in the developing Drosophila brain (**Chapter 2, Chapter 3**).

Second this work aims to describe the development and validation of machine learning tools (CytoCensus) for cell detection and show that this supports the analysis of live imaging data in a user-friendly way (**Chapter 4**).

Third, in this thesis I aim to use CytoCensus to analyse live imaging data; using this analysis in order to describe neuroblast behaviour in normal development, and to understand how

such behaviour is controlled (**Chapter 5**). This work seeks to determine the role of the RNA binding protein Syncrip in neuroblast behaviour and in doing so understand the Syncrip phenotype (**Chapter 3, Chapter 5**). Finally, this work aims to illustrate that CytoCensus is a generally applicable program that can provide useful biological quantitation in a range of systems and tissues (**Chapter 6**).

CHAPTER 2

MATERIALS AND METHODS

2.1. SOLUTIONS AND REAGENTS

All solutions were prepared as in "Molecular cloning - A laboratory manual" (Sambrook et al., 1989), unless otherwise stated. Phosphate buffered saline (PBS), Tris-acetate-EDTA (TAE), Tris-Borate-EDTA (TBE), were prepared and autoclaved by the Department of Biochemistry media kitchen.

2.2. FLY STRAINS

Drosophila stocks were raised on standard cornmeal-agar medium at either 21 °C or 25 °C. To assist with determining larval age for studies of larvae at the beginning and at the end of L3, Bromophenol Blue was added to a final concentration of 0.05% in cornmeal-agar medium (Maroni and Stamey, 1983). Oregon-R was used as the Wild-Type (WT) strain. The *Drosophila* fly strains used in this study are detailed in **Table 2.1**.

For live imaging ;;Jupiter::GFP,Histone::RFP was used as the WT strain. For live imaging of *syp* RNAi ;insc-GAL4; Jupiter::GFP,Histone::RFP was crossed to *;syp*-RNAi; Jupiter::GFP, Histone::RFP. Virgins of either genotype were crossed to the corresponding males to yield ;insc-GAL4/*syp* RNAi; Jupiter::GFP,Histone::RFP Hatching rates were very low (<1 in 50), so tens of females were used per cross. Knockdown of Syncrip was confirmed by immunofluorescence imaging with anti-Syp.

Table 2.1: Fly strains used in this study.

Genotype	Source
w1118;;ase-GAL4	J. Knoblich Lab
w;inscl407-GAL4;	J. Knoblich Lab BL 8751
w*; P[His2Av-mRFP1] III	Bloomington 23650
syp -/- w[1118O]; PBac(PB)Syp[e00286]/TM6B	Harvard Exelixis Collection
w;;Jupiter::GFP;Histone::RFP (recombination on the third)];	Ephrussi
syp RNAi w1118O; P[GD9477]v33011	Vienna Drosophila Resource Center
w;;[*ase*-GAL4>UAS-MCP-mNeonGreen]/P[His2Av-mRFP1]III	This work
w;inscl407-GAL4; Jupiter::GFP;Histone::RFP	This work
w; *syp* RNAi; Jupiter::GFP;Histone::RFP	This work

2.3. Fixed Tissue Preparation and Labelling

2.3.1. Drosophila

Flies of both genders were raised as described above and larvae from second instar to pre-pupal stages were collected and dissected directly into fresh 4% EM grade paraformaldehyde solution (from a 16 % stock; Polysciences) in PBS with 0.3% TritonX-100 then incubated for 25 minutes at room temperature (RT). Following fixation, samples were washed 3 times for 15 minutes each in 0.3 % PBST (1x PBS containing 0.3% Tween) and blocked for 1 hour at RT in Immunofluorescence blocking buffer (1 % BSA prepared in 0.3% PBST). Samples were incubated with primary antibody prepared in blocking buffer for either 3 hours at RT or overnight at 4˚C. Subsequently, samples were washed 3 times for 20 minutes each with 0.3% PBST followed by incubation with fluorescent labelled secondary antibodies, prepared in blocking buffer, for 1 hour at RT. For nuclear staining, DAPI was included in the penultimate wash step. Samples were mounted on glass slides in VECTASHIELD (Vector Laboratories) for examination.

For details of preparation of *Drosophila* Secondary Cells (Chapter 6) see (Corrigan et al., 2014).

2.3.2. Mouse Embryos (Prepared and Imaged by Ita Costello)

Refer to (Simon et al., 2017) for details on mouse embryo preparation. Briefly, E6.5-7.5 embryos were fixed with 1% PFA overnight, washed in 0.1% Triton-X in PBS, permeabilised for 15 minutes in 0.5% Triton-X in PBS and blocked with 0.1% Triton-X, 0.2% donkey serum, 0.2% BSA. Following blocking, embryos were incubated with primary antibodies overnight at 4 °C, washed in block solution, and incubated with secondary antibodies in block solution for 2 hours at RT. Finally, embryos were washed, counterstained with DAPI, mounted in Vectashield (Vector Laboratories) on chamber slides (LabTek), and imaged on an Olympus FV1000 inverted confocal (see **Section 2.4.3**).

2.4. Live imaging

2.4.1. Culture of Live Explanted Larval Brains on the Microscope for NBs

Brains were dissected from third instar larvae in brain culture medium (BCM) according to https://www.youtube.com/watch?v=9WlIoxxFuy0 and placed inside the wells of a pre-prepared culturing chamber. To assemble the culturing chamber, 1% low melting point (LMP) agarose (ThermoFischer) was prepared as 1:1 *v/v* ratio of 1 × PBS and Schneider's medium (ThermoFisher 21720024), pipetted onto a 3 cm Petri dish (MatTek) dish and allowed to solidify. After solidification, circular wells were cut out using a glass capillary ~ 2 mm diameter. To secure the material in place, a 0.5% LMP agarose solution (1% LMP agarose solution brain diluted 1:1 with BCM) was pipetted into the wells to form a cap. Finally, the whole chamber was flooded with BCM. BCM was prepared by homogenising ten third instar larvae in 200 µl of Schneider's medium and briefly centrifuge to separate the fat and supernatant from the larval carcasses. This lysate was added to 10 ml of 80% Schneider's medium, 20% Fetal Bovine Serum (GibcoTM ThermoFisher), 10 µl of 10 mg/ml insulin (Sigma). A lid was used to reduce evaporation.

2.4.2. Culture of Live Explanted Larval Brains on the Microscope for Imaging of GMCs

For imaging of GMCs, brains were prepared as above, with some modifications. The culturing chamber was prepared using 2% LMP agarose in Schneider's medium (kept at 4 °C until use). A solid-agar cap (2% LMP agarose) was cut using a custom tool. One inch of Pasteur pipette was cut off, and a glass capillary tube placed inside with a short length extending out of the Pasteur pipette. This was sealed in place by either SYLGARD silicone elastomer (Sigma) or Parafilm (Sigma). By placing the glass capillary tube on the prepared dish, wells are easily cut and removed. Alternatively cut wells can be removed using tweezers. Following preparation of the culture chamber, the chamber was flooded with BCM (see above), and brains pipetted into the wells. A solid agar cap was then placed on each brain using tweezers. This was more consistent at holding brains against the coverslip than the earlier approach. Care must be taken not to flatten brains during this process, as it appears to result in a higher rate of stalled NB divisions which are likely to be artefacts. Brains were then imaged on Olympus FV3000 30xSI lens (**Section 2.43**) with 2 minute intervals, at 0.2 × 0.2 × 0.5 µm resolution.

2.4.3. Microscopy

Confocal, live imaging of *Drosophila* was performed using: either an inverted Olympus FV1200 six laser line spectral confocal with environmental chamber, fitted with high sensitivity gallium arsenide phosphide (GaAsP detectors), ×60 1.4 NA, or ×60 SI 1.3 NA lenses at room temperature (typically 21–24 °C); or an inverted Olympus FV3000 six laser line spectral confocal, fitted with high sensitivity gallium arsenide phosphide (GaAsP detectors), ×30SI 1.1 NA at room temperature (typically 21–22 °C). The confocal pinhole was set to one airy unit to optimise optical sectioning with emission collection. For FV1200 imaging, laser power was kept less than ~4 µW for both 488 nm and 559 nm excitation. Images were collected at 1024×124 pixels (pixel size 0.207 µm) with a scan dwell time of 2.0 µs/pixel (frame rate 6.5 s/frame, averaged ×2). The total exposure time per Z stack (35) frames was 3 min, 48 s.

For imaging the cell cycle of GMCs, modifications to the above protocol were made. First, the temporal and spatial resolution was increased to 2 min intervals, with 0.2 × 0.2 × 0.5 µm resolution. We, reduced the amount of light by approximately a factor of 10, using resonance scanning with a 2× average, and used post acquisition patch-based denoising to restore image quality (Carlton et al., 2010).

For brain size measurements, a GE Deltavision Core widefield system was used with a Lumencor 7-line illumination source. For live culture and imaging the culturing chamber containing the sample was closed with a lid and kept at 20 °C inside the microscope environmental chamber.

For imaging of fixed material, an Olympus FV1200 confocal was used with ×20 0.75 dry or ×60 1.4 NA. lenses. Settings were adjusted according to the labelling and were kept consistent within experiments. For details on imaging of mouse embryos refer to (Simon et al., 2017).

2.4.4. Image Analysis

All programs used for image analysis were installed and run on a MacBook Pro 15", 2015; Intel Core i7 2.88GHz; 16GB RAM. Basic image handling and processing was carried out in FIJI (ImageJ V1.51d) (Schindelin et al., 2012). The CytoCensus software was written in Python, available at github.com/hailstonem/CytoCensus. Scripts for the denoising pipeline and tracking are also available at github.com/hailstonem. The CytoCensus User guide is available in Appendix A. Server side processing was performed on the 'mprocessor1' biochemistry server with 32 core with 125GB memory running Ubuntu 14.04.5 LTS.

2.4.4.1. Quantifications of NB Divisions

CytoCensus models were trained and evaluated on 512×512, 12–24 hour live imaging datasets, with at least 100 timepoints, using the pyramid feature set, the value of sigma was set to 8. For NB identifications, typically training was performed on 5–10 planes containing neuroblasts

from the first time point, and a further 5–10 planes in another time point (usually the final time point). For dividing neuroblasts, training was performed on 5–10 planes in the first time point containing both dividing and non-dividing neuroblasts (selecting only the dividing NB), followed by training on 5–10 planes containing dividing neuroblasts in the following 10–20 time points. Training on NB from a single dataset could typically be applied to multiple brains from different imaging sessions, as long as imaging conditions were similar, with some reduction in the quality of detections. Training on dividing NBs was less robust; training on one image could typically be applied to other imaging from the same session for dividing NBs, but not across imaging sessions from different days.

NB division plots were created using a python script that updated NB locations (based on the CytoCensus NB detections) by nearest neighbour search within 4 μm. For each NB, the mean proximity score within a 10 μm radius was calculated for each timepoint, and each trace normalised to the maximum proximity score for that trace. A similar effect can be achieved in FIJI by 3D mean filtering and selecting a region of interest (ROI) in a single plane using the function "Plot Z Profile". Robustness is improved by low pass filtering the resulting signal, for instance by subtracting a moving average over 20 frames which removes spatial and temporal differences in the background of the probability density maps.

2.4.4.2. Image Registration

Images for measurement of NB division were registered using the FIJI Correct 3D drift plugin with multiscale calculation. Images were registered for measurement of GMC division using a custom Script based on SimpleITK, using conjugate gradient descent optimisation (learning rate 0.5, convergence threshold 1e-07), median filtering (10×10×1), Euler 3D transform (using Geometry for initialisation), linear interpolation for optimisation, and nearest neighbour for output.

2.4.4.3. Quantifications of GMC Divisions

CytoCensus models were trained and evaluated on 512×512, 12 hour live imaging datasets, using the pyramid feature set, the value of sigma was set to 2. For GMC candidate identifications, positive training was performed by identifying around 5 newly born GMCs, and GMCs that would divide in subsequent frames, and training in 5–10 frames containing each GMC. Negative training (i.e. selecting regions without positives) was performed by identifying cells and regions in which no cells were dividing (as these are likely to be neurons). GMC candidates were then tracked using the Crocker-Grier algorithm in trackpy, with a memory of 5, and a potential search radius of 4.5 μm. Evaluation of tracking errors were performed on tracks that were at least 50 time-points long.

2.4.4.4. Quantification of Cell Distribution in Zebrafish Organoids

As per **Chapter 5**, individual models were trained using CytoCensus for each of the cell classes using the pyramid feature set, and a value of 4 for sigma. Training was performed on cells from two images for each cell type and applied to the whole dataset. Using the cell centre output from CytoCensus, an automatic calculation of the organoid edge was made. In this case, the cells were first binned by radius. Then, the size of the organoid was calculated automatically from the 90th percentile of the layer of Dead cells, which allowed normalisation of organoid size to a fixed range (0–100). Cell distributions were then calculated relative to this range. Hence, a proportion of the dead cells may have relative radii larger than 100.

2.4.4.5. Quantification of Cell Distribution in Mouse Embryos

As per **Chapter 5**, individual models were trained using CytoCensus for each of the cell classes using the pyramid feature set, and a value of 2 for sigma. Two models were trained for each TF (Blimp1+TF+ cells, Blimp1+) using an empirically chosen threshold in the range 1–20 for each model. The model was then restricted to the region of the primitive streak using the ROI tool of CytoCensus, and the ratio of TF+ cells calculated.

2.4.4.6. File Transfer

OIR files (from Olympus microscopes) were transferred over 1 Gbit connection to the biochemistry server. Following transfer, they were converted to using Bioformats CLI tools (Linkert et al., 2010) into OMETIF files for upload to OMERO because OMERO (at the time of writing) does not accept OIR files (Goldberg et al., 2005).

2.4.4.7. Denoising

Denoising was performed using NDSAFIR (Boulanger et al., 2010), using 3 iterations, in zt mode (that is denoising in both Z and time). The NDSAFIR function was available in PRIISM (Carlton et al., 2010). Preparation of the MRC file from OMERO was modified with a custom script from the one provided by David Pinto (Micron Oxford) under GPL license (github.com/MicronOxford/scripts/).

2.4.5. Data and Software Availability

The following freely available image analysis tools were used: FIJI, ImageJ V1.51d (Schindelin et al., 2012), Ilastik (V1.17)(Sommer et al., 2011), RACE (Stegmaier et al., 2016). The CytoCensus software can be installed as a stand-alone program: full install available at www.github.com/dwaithe. Image data was archived in OMERO V5.3.5 (Goldberg et al., 2005). Image conversions were carried out using the BioFormats plugin in FIJI (Linkert et al., 2010).

CytoCensus is built in python with the pyqt, matplotlib, numpy, scipy, scikit-learn and scikit-image libraries (Hunter, 2007; Jones et al., 2014; Pedregosa et al., 2011; Van der Walt et al., 2014; Walt et al., 2011).

2.4.6. Quantification and Statistical Analysis

Mutant comparisons were performed using Student's T test, following the Shapiro-Wilk test to test for normal distribution of the data. A p-value of <0.05 was considered significant.

Numbers of replicates are detailed in the figure legends and main text. Unless otherwise stated, error bars shown are standard error of the mean (SEM). Where stated, Bonferroni correction of p-value significance for multiple comparisons was applied. Counts were compared using a χ-squared test. Power calculations were made using G*Power 3 (Faul et al., 2007).

2.4.7. KEY RESOURCES TABLE

Reagent or Resource	Source	Identifier/Link
Antibodies		
Guinea pig anti-Syncrip (use 1:100)	I. Davis Lab (McDermott et al., 2014)	N/A
Mouse anti-Prospero (use 1:100)	Abcam	ab196361
Guinea pig anti-Asense (use 1:200)	Gift from JA Knoblich	N/A
Rat anti-Deadpan (use 1:100)	Abcam	ab195173
Goat anti-Mouse Alexa Fluor 488 (use 1:250)	ThermoFischer	A-11001
Goat anti-Guinea Pig Alexa Fluor 647 (use 1:250)	ThermoFischer	A-21450
Goat anti-Rabbit Alexa Fluor 594 (use 1:250)	ThermoFischer	R37117
Goat anti-Mouse Alexa Fluor 647 (use 1:250)	ThermoFischer	A-32728
Chemicals, Peptides, and Recombinant Proteins		
VECTASHIELD Antifade Mounting Medium	VECTOR Laboratories	H-1000
Formaldehyde, 16%, methanol free, Ultra Pure	Polysciences, Inc.	18814-20
Low melting point agarose	ThermoFischer	
Foetal Bovine Serum (BSA)	GibcoTM, ThermoFischer	

Schneider's Medium	ThermoFischer	21720024
Bromophenol Blue	Sigma-Aldrich	116K3528

Reagent or Resource	**Source**	**Identifier/Link**
Software and Algorithms		
Fiji, ImageJ (V1.51d)	Schindelin *et al.,* (2012)	*http://imagej.nih.gov/ij*
Ilastik (V1.17)	Sommer *et al.*, (2011)	*ilastik.org*
CytoCensus	This Thesis	This article
SoftWoRx, Resolve3D	GE Healthcare	N/A
Microsoft Excel	Microsoft Cooperation	150722
OMERO V5.3.5	Allen *et al.,* (2012)	openmicroscopy.org/omero/
Bio-Formats	Linkert *et al.,* (2010)	openmicroscopy.org/bio-formats/
RACE	Stegmaier et al., (2016)	
Other		
Superfine Vannas dissecting scissors	WPI	501778
MatTek (or Eppendorf) 3 cm glass-bottom Petri-dish	MatTek (or Eppendorf)	P35G-1.5-14-C
Broad Bioimage Benchmark Collection Neutral Challenge Datasets	(Svoboda et al., 2009)	BBBC024v1

CHAPTER 3

STUDYING BRAIN DEVELOPMENT WITH LIVE IMAGING

3.1. INTRODUCTION

Drosophila larval brains develop for more than 120 hours, increasing the number of neurons in the brain by an order of magnitude over this period (Homem et al., 2013). These neurons are formed by repeated divisions of stem cell-like progenitors called neuroblasts (NBs) in *Drosophila*, forming tens of thousands of new neurons. NBs go through a series of divisions in the embryo, producing neurons that are required for larval behaviour such as locomotion and feeding, before becoming quiescent (see **Introduction, Section 1.1**). NBs undergo reactivation around late L1 –L2 (Tsuji et al., 2008), then continue dividing for over 100 hours until around pupation. At this point they cease to regrow between divisions and decrease in size before ultimately undergoing terminal division.

Live imaging of brain development in *Drosophila* has primarily used one of two different approaches. First, and most common, is the use of disaggregated brain cells in culture (Furst and Mahowald, 1985; Homem et al., 2013; Moraru et al., 2012; Savoian and Rieder, 2002), (**Figure 3.1**A). In particular, the use of NBs from disaggregated brains is one of the most popular methods for examining divisions of NBs, for instance mitotic spindle assembly, the segregation of cytoplasmic components, and measurements of division rate (Egger et al., 2013; Moraru et al., 2012). The relative ease of high resolution imaging, reduced background, and cellular complexity make it one of the only practical methods to study many problems (such

as the measurement of the cell cycle length). Indeed, some of our best estimates of NB and GMC division rate come from measurements of NBs in disaggregated culture (Homem et al., 2013). However, in the past, disaggregated culture has produced flawed measurements, unrepresentative of *in vivo* conditions. For instance, early measurements of NBs demonstrated a high rate of mitotic defects, (Savoian and Rieder, 2002), although subsequent experiments that avoid physical distortion (Lerit et al., 2014) avoid these issues. Nonetheless, the question of how representative *in vitro* culture is of the *in vivo* situation remains. The relatively low success rate of following NB division in some *in vitro* conditions (1 in 10, Dr Evanthia Zacharioudaki, private communication), highlights questions about the robustness of the method. Interestingly, although extensively studied, a range of division rates, from MARCM methods and *in vitro* imaging, are reported for both NB and GMCs (Bowman et al., 2008; Ceron et al., 2006; Homem et al., 2013). In this chapter, I will examine the development of *ex vivo* imaging of the intact larval *Drosophila* brain as an alternative approach to address these issues. I will also return to this theme in Chapter 5.

Most imaging of intact brain development at cellular resolution has tended to be for relatively short periods of around a few hours (Cabernard and Doe, 2013; Lerit et al., 2014; Prithviraj et al., 2012) (**Figure 3.1B**). Recent developments in imaging of pupal and adult brains has allowed for much longer imaging durations and higher resolutions, allowing the study of neuron maturation (Cabernard and Doe, 2013; Homem et al., 2013; Sabado and Nagoshi, 2018) (**Figure 3.1C**).

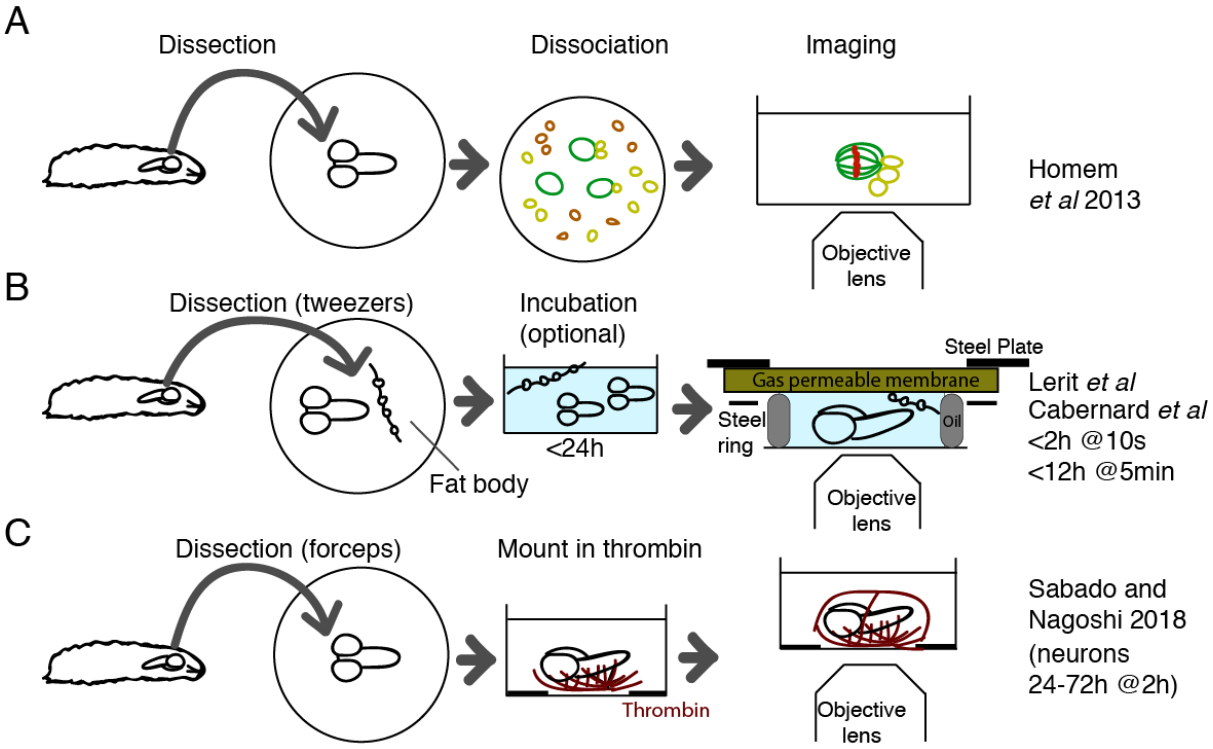

Figure 3.1: 3D culture and imaging methods. A) Culture of dissociated NB from Homem *et al*. 2013. B) Culture of whole larval brains from Lerit *et al*. 2014 and Cabernard *et al*. 2014 C) Culture of adult brains from Sabado and Nagoshi 2018.

In this chapter, I will discuss the development of a pipeline for culture and mounting of *ex vivo* brains, and imaging and processing of images to yield high-resolution images of NB and GMC divisions suitable for studying their development in an intact brain. The initial protocol for culture and mounting was developed by Lu Yang; modifications are stated in the text. All images and validations shown are my own unless otherwise stated.

3.1.1. Specific Aims

Develop an *ex vivo* imaging and culture method for analysing NB and GMC behaviour

Develop a pipeline for image post-processing to support the measurement of cell division

Identify and quantify cell divisions in long term live imaging

3.2. Development of Culture and Mounting

We first developed an isolation procedure, incorporating scissor-based dissection of second or third-instar larvae, in preference to solely tweezer or needle-based dissection. Lu Yang

optimised and simplified the culture medium (compared to Cabernard & Doe, 2013) and developed a convenient brain mounting technique (**Figure 3.2A**). Briefly, a well in which the brain will sit is made using a glass capillary tube in 1% agar on a glass-bottomed petri dish. We add the brains to the wells and then cap the wells with 0.5% low melting point agarose to prevent brain movement. We made use of bright, endogenously expressed protein traps Jupiter::GFP and Histone::RFP, marking microtubules and chromosomes respectively, to follow development (**Figure 3.3**). These generic cytological markers were chosen as they are more consistent across wild-type (WT) and various mutants than more specific markers, such as Deadpan (Dpn), Asense (Ase) or Prospero (Pros), which although commonly used to identify NBs, GMCs and progeny, are often disrupted in mutants. Finally, we went through an optimisation process to ensure our culture and imaging conditions provided 3D datasets of sufficient temporal and spatial resolution for us to follow cell proliferation over time, but did not compromise viability.

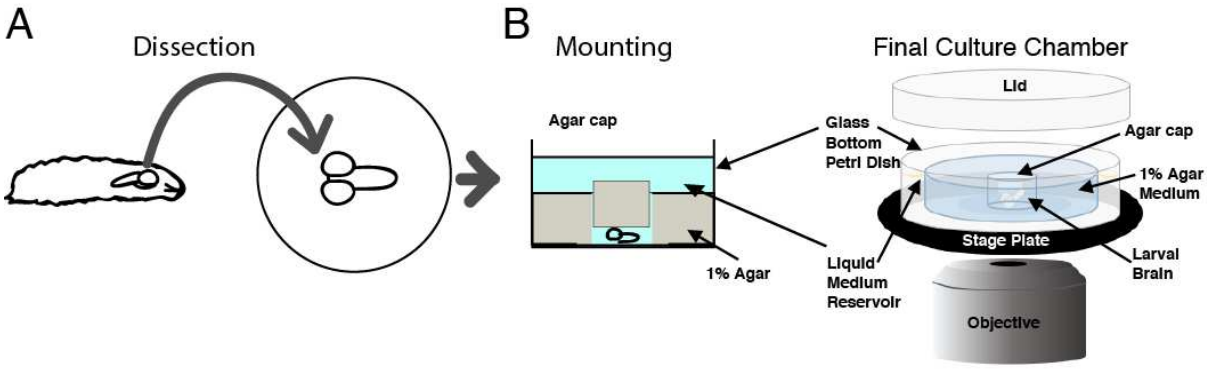

Figure 3.2: Isolation and brain mounting (see **Methods**). A) Larvae are dissected in culture medium, removing the eye imaginal discs. B) Brains are then mounted with the VNC and ventral CB against the glass coverslip (down) in wells in 1-2% agar in culture medium and imaged.

To ensure this *ex vivo* imaging is an accurate reflection of the *in vivo* situation, I examined a number of criteria and optimisations. Using our *ex vivo* conditions, brain development appeared normal or close to normal. First, we take care to dissect brains as carefully as possible. We find that the use of dissecting scissors helps to minimise damage compared to solely

tweezer or needle-based dissection. Second, we discard brains that show obvious morphological damage such as distorted brain lobes (Lerit et al., 2014) or a damaged ventral nerve cord (VNC) during preparation and those that, under fluorescence imaging, show signs of holes or lesions in the tissue (**Figure 3.3**). Although these brains may look externally viable, from experience, with time in culture, these holes and lesions expand to encompass large portions of the brain. Second, I find that our cultured larval brains consistently continued to develop and increase in size (**Figure 3.3C,D**), easily within the normal range of larval brain size. Although this general trend was positive, in some brains there was an initial pause in the growth, even shrinking of the brain, visible within the first hour. The cause of this unclear, but possibilities are a difference in the osmotic potential of the conditions (although Schneider's medium should be a good match for haemolymph), or an adaptation period for the brain to adjust to new conditions. Nevertheless, under these conditions no bleaching of fluorophores is seen. Due to the brightness of these markers, we can use a fraction of the light and hence have fewer concerns regarding phototoxicity compared to other protein traps (for instance I find Moesin::GFP to have approximately one tenth of brightness). Together this suggests that, at least at a coarse scale, our *ex vivo* imaging can sustain normal development of the central brain.

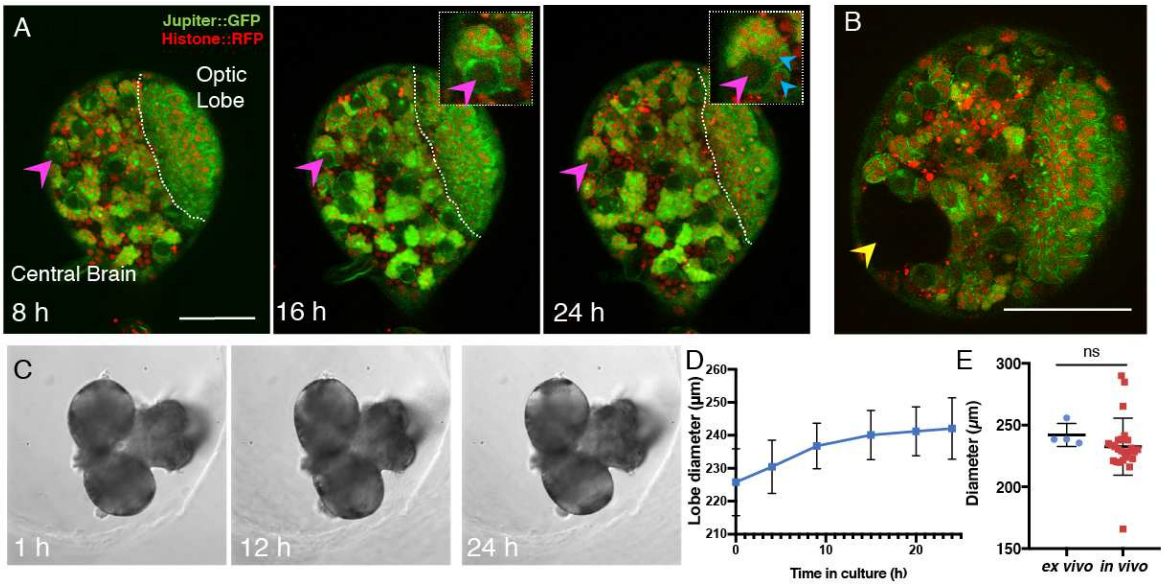

Figure 3.3: Brain development in culture A) WT brain labelled with Jupiter::GFP (green), Histone::RFP (red) showing a NB (magenta) which divides in culture to produce GMCs (blue arrow, inset) Scale bar 50 µm; B) Damaged brain showing lesion after 8h in culture (yellow arrow) Scale bar 50 µm; C) Brain growth in culture of WT brain over 24h; D) Plot of brain growth in culture in wL3 showing standard deviation (blue - average of 4 brains) E) Scatter dot chart comparing final brain diameters of isolated late stage wL3 brains (n=4) with wL3 (n=25) brains cultured to an equivalent age *in vivo.* Diameters are within the normal range (ns, t-test).

3.3. Imaging of NB to determine individual cell division rates and behaviour

The quality of imaging that the culture approaches detailed above support allows identification of individual NB within the brain, so we next examined the division of NB within the CB region (see **Introduction 1.1**). Using this approach, we recorded consistent Type 1 NB division rates of 60-90 min per cycle for wL3, as previously published (Bowman et al., 2008; Homem et al., 2013)(Figure 3.4; Movies S1 and S2). Type 1 NBs were identified by location according to Homem & Knoblich (2012). I rarely observed excessive lengthening or arrest of the cell cycle in NBs over a 22 h imaging period, which is approximately the length of the wandering third instar larval stage (wL3) (**Figure 3.4).** With longer duration culture and imaging, up to 48 h, I observe an increase in cell cycle length, which might be expected for wL3 brains transitioning to the pupal state (see **Discussion**). I conclude that these *ex vivo*

culture and imaging conditions accurately represent a model for studying NB division in the context of the whole *Drosophila* brain and allows high time- and z-resolution imaging for quantitation of cell numbers and division rates.

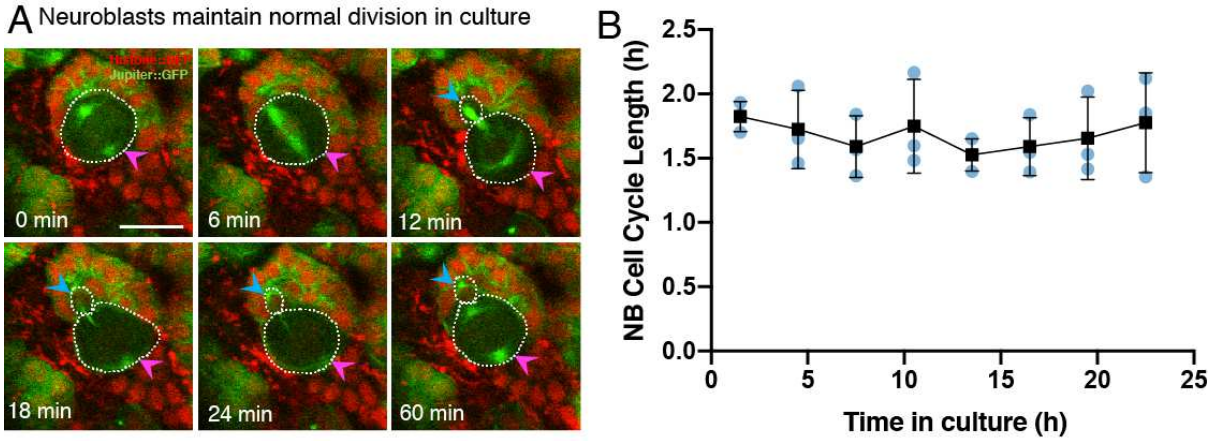

Figure 3.4: NB division in *ex vivo* culture is normal A) Imaging of Jupiter::GFP (Green), Histone::RFP (Red) labelled brain showing a single normal NB (Pink arrows) division in culture over 1 h (MH). GMC daughters (Blue). Scale bar 10 μm B) Analysis (MH) of NB cell cycle length over 22 h, average of 3 brains (black). Each point (blue) represents the average cell cycle length of all divisions that occurred within the corresponding 3 h window in a single brain. NB division rate is constant over the imaging period. Original Imaging and cell cycle measurements (L. Yang, unpublished work)

Despite attempts to immobilise the brain in this culture approach, brain movement was a recurring issue. The mounting protocol in **Figure 3.2** uses a soft agar cap to the 'wells' containing the brains. This is imperfect: there is residual movement of the brains (6.8 μm in **Figure 3.5**A-B), as well as some temperature-driven focus drift of the microscope over long periods of time (up to 0.3 μm/h in **Figure 3.5F,G**). In principle, focus drift can be corrected with Z-drift-compensation, which corrects for movement of the coverslip using reflections of an IR laser from the coverslip. Unfortunately, its performance was unpredictable over long imaging periods, and used without success. Temperature drift on the Olympus FV1200 was a particular issue due to reliance on AC for temperature control. Nonetheless, these movements were often (although not always), small enough to allow for correction using image registration

(**Figure 3.5C,G**). I used the image registration plugin Correct-3D-Drift in FIJI to correct for these drifts, which is straightforward to use and faster than iterative approaches (Parslow et al., 2014; Schindelin et al., 2012). This was usually sufficient to correct for most of this drift, although slow drifts were not always corrected (**Figure 3.5F,G**), although the plugin occasionally produced wildly inaccurate results (see **Section 3.5**), so I also produced a script using SimpleITK for this purpose. For both methods, I found it was crucial to image the edge of the brain for good quality registration (e.g. **Figure 3.5F**), presumably as this helped constrain the position of the brain in Z. The large size of NB (up to 15 μm diameter) meant that small errors in registration did not affect my ability to follow NBs over time. Together this culture, imaging and registration supports the measurement of NB division in *ex vivo* culture, although manual analysis of division rates is very time consuming. I will return to the measurement of NB divisions in **Chapter 5**.

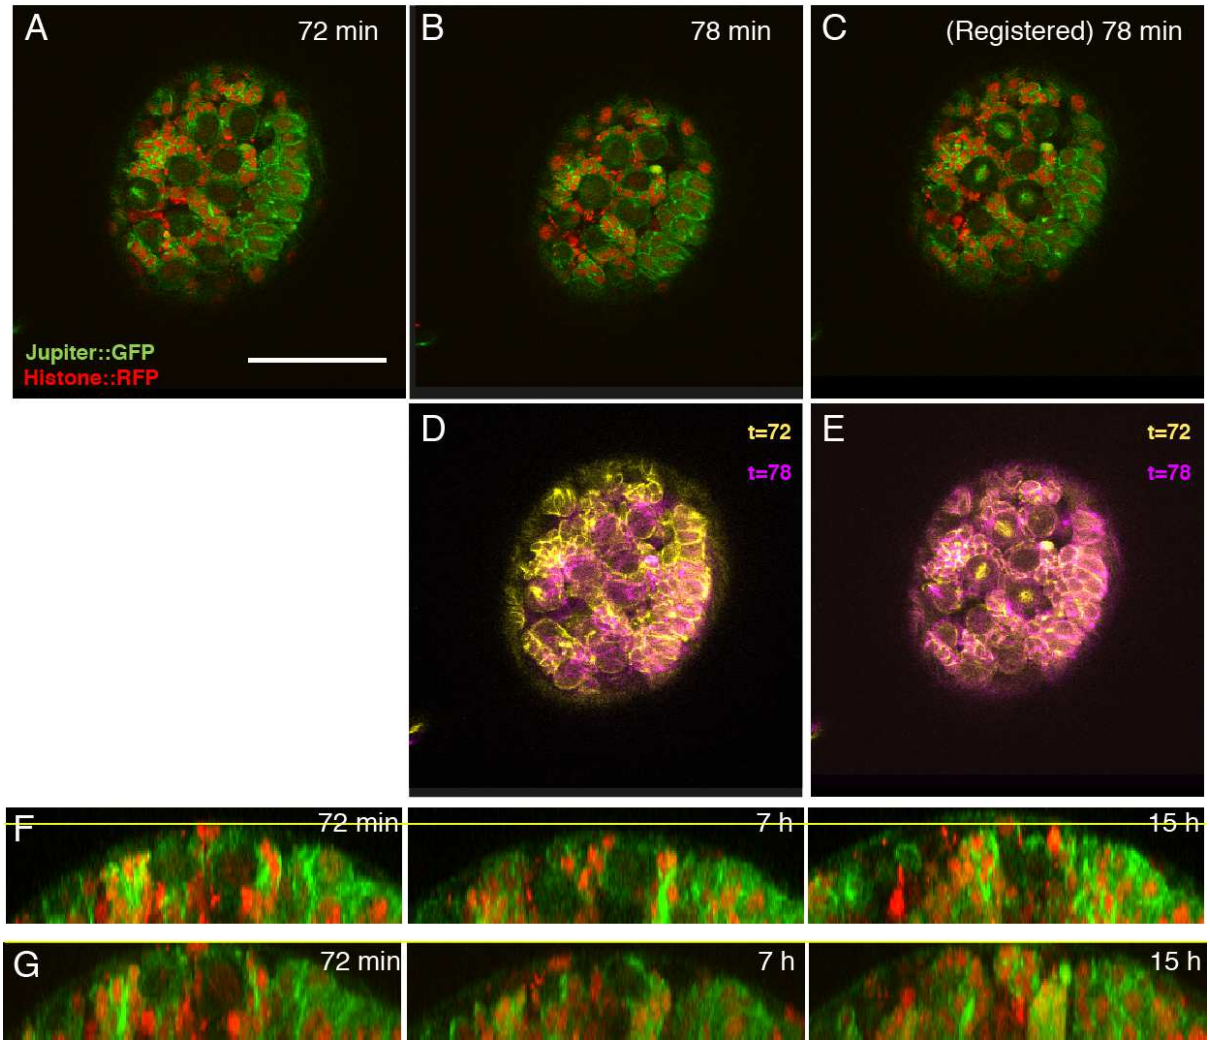

Figure 3.5: Brain movements can be corrected for by image registration. A-B) Sequential frames of brain imaging, showing significant (6.8 µm) movement, and (C) after image registration. Scale bar is 50µm. D-E) Overlay of frames showing imaging at 72 minutes (yellow) and at the next frame at 78 minutes (magenta), illustrating that movement is well corrected for. F) Movement of the brain over 15 h of imaging, show in in an yz-plane view. Initial position of the edge of the brain (yellow) indicates that the brain moves over time. Times are shown relative to the start of imaging. G) Corresponding, yz-plane of the brain after registration.

3.4. Imaging of GMCs

I have thus far detailed conditions for the culture and imaging of NB in the CB region of the *Drosophila* larval brain. Following GMCs, however turns out to be a significantly more difficult task, due to the variable appearance of GMCs over their long cell cycle length. Furthermore,

their similarity to newly-born neurons makes them indistinguishable using the Jupiter::GFP, Histone::RFP marker combination (see **Chapter 4**). Additionally, their long cell cycle length combined with growth and movement of the brains makes for extremely difficult manual tracking. In this section I will detail improvements in imaging that support the imaging of the GMC cell cycle.

First, imaging GMCs at sufficient resolution to follow their division required further optimisation. Since they are much smaller than NB (3-5 μm), GMCs were difficult to properly resolve using the imaging described above. Furthermore, whilst the xy resolution (0.41 μm) of the imaging of NB was sufficient to separate NB and GMCs, the reduced z resolution (1.15 μm) made them difficult to separate, especially given the low SNR of the imaging. Additionally, whilst registration of images with multiple-micron accuracy was sufficient for NB, these errors were enough to make following the smaller GMCs difficult. This difficulty was accentuated by the tendency of GMCs to make significant (~3 μm) movements whenever a NB divides (**Figure 3.6).** These rapid movements of GMCs are caused by the addition of a new GMC at NB division (**Figure 3.6A**), and can cause the entire niche to shuffle around (**Figure 3.6A,B**). To compensate for these challenges, I found that we needed much faster temporal and spatial resolution than previous imaging.

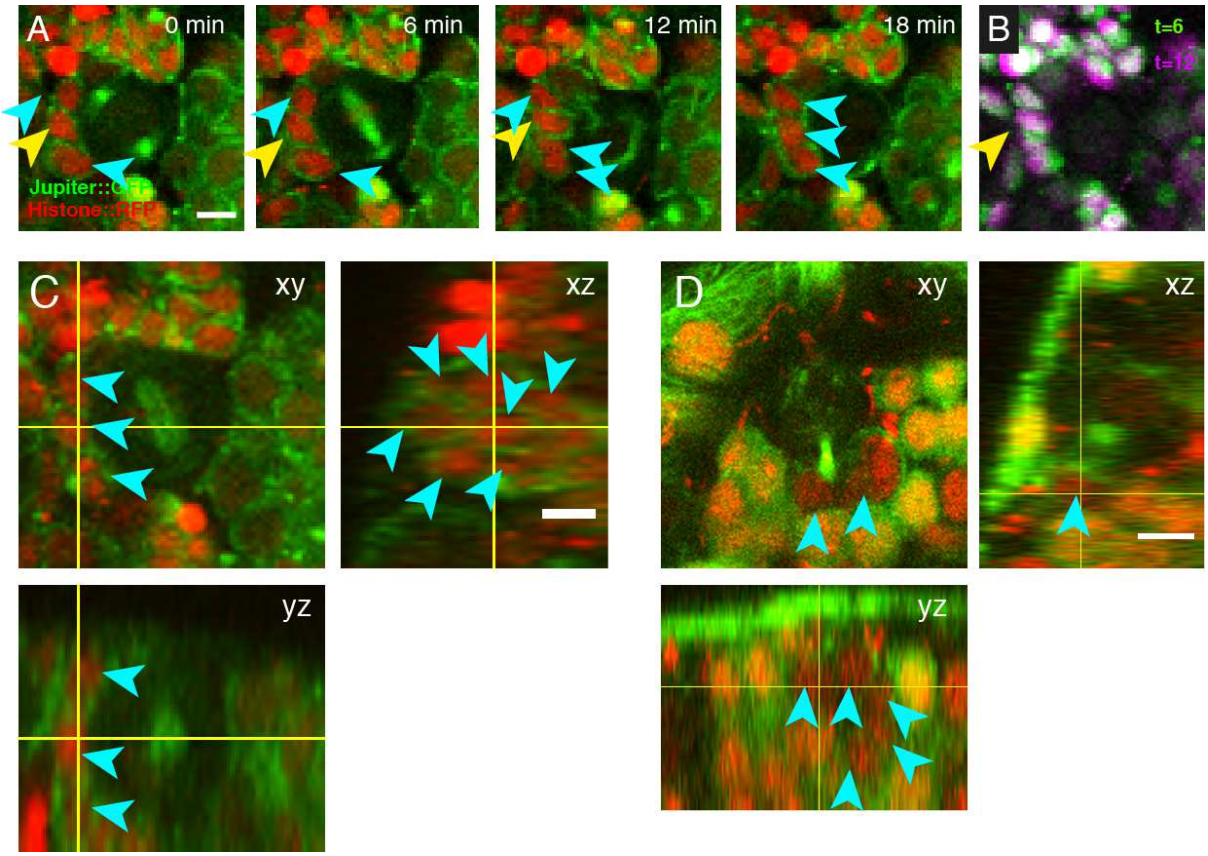

Figure 3.6: Imaging of GMCs A) When a NB divides, the niche reshuffles, and GMCs (cyan arrows) move significantly. Here, the labelled GMC (yellow arrow) is no longer visible following the division. Z-projection average over 3 frames for clarity. Scale bar is 5 µm. B) Comparison of the Histone channel at 6 minutes (green) and 12 minutes (magenta) showing the difference in the positions of the nuclei. Unmoved cells are shown as white. Yellow arrow: GMC with significant movement C) Orthogonal views of imaging GMCs (cyan) with NB protocol (6 minute intervals, 0.41 µm x 0.41 µm x 1.15 µm) results in clear identifications in XY, but lack of clarity in Z. Scale bar is 5 µm. D) Orthogonal views of imaging GMCs with modified protocol (2 minute intervals, 0.2µm x 0.2µm x 0.5 µm) more clearly separates GMCs (cyan) in Z. Scale bar is 5 µm.

To address the issues described, I increased the imaging resolution from 0.4 µm lateral (XY) and 1.15 µm axial (Z) resolution to 0.2 µm XY and 0.5 µm Z, to better resolve GMCs, and increased the temporal resolution from 6 minutes to 2 minutes, to better cope with movements. Although these conditions better support tracking of GMCs, these changes significantly increased the effective light-dose to the brains. Initial tests indicated that this increased light dosage caused significant bleaching (2.2-fold reduction in fluorescence

intensity over 6 hours of imaging) and seemed to reduce the number of GMC divisions over time, which suggests phototoxic effects. This level of imaging clearly does not support robust culture, so I therefore looked to post-processing methods in order to reduce the amount of light the tissues experience.

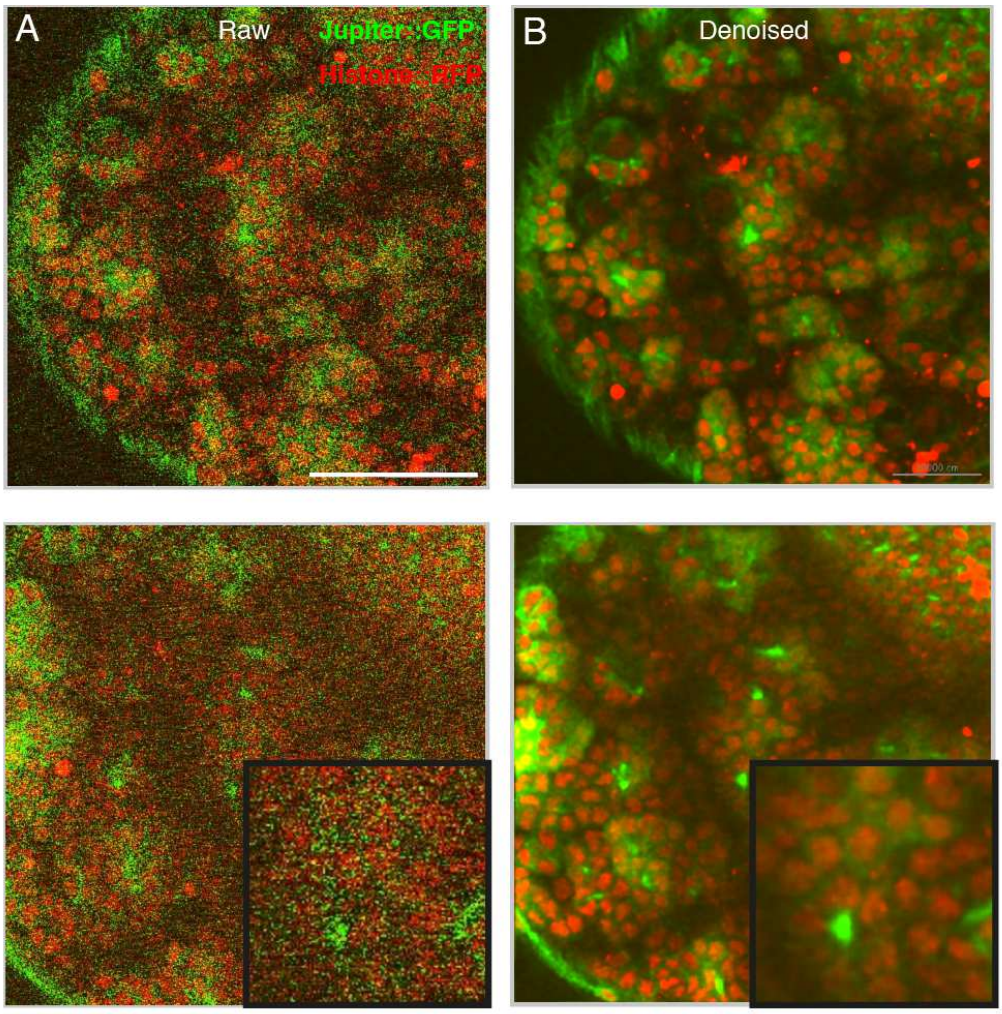

Figure 3.7: Comparison of raw imaging data, and denoised data A) Imaging (single plane) of Jupiter::GFP, Histone::RFP using resonance scanning with 2x average. Top: imaging near the surface and Bottom: imaging at depth (30µm) Scale bar is 50µm B) Denoised images using ND-SAFIR with 3 iterations

To reduce photo-damage I therefore reduced laser excitation levels by ~10 fold using resonance scanning. In order to do so I turned to an Olympus FV3000 microscope, over previous work on the FV2000, to make use of the resonance scanner. This allowed me to collect images 16x

faster than previously, with a corresponding decrease in image quality. The consequence of this is much noisier images (SNR_{RFP} = 0.5, SNR_{GFP} = 0.3), compared to previous imaging (SNR_{RFP} = 4.7, SNR_{GFP} = 9.8), so I investigated methods of restoring image quality. Using patch-based denoising, developed by Kervrann and Boulanger (2006), dramatically improved image quality (SNR_{RFP} = 9.0, SNR_{GFP} = 4.7) (**Figure 3.7**). This approach supported both NB and GMC divisions in the intact brain, imaging at subcellular resolution in 3D (**Figure 3.7, Figure 3.8**). I verified that this results in normal, symmetric GMC division and that the number of divisions per niche remains approximately constant over time (ns, p=0.5, n=5, χ^2-test, 1-β=0.8, w=0.35) (**Figure 3.8B**). This suggests that these conditions are sufficient to support the culture and imaging of GMCs.

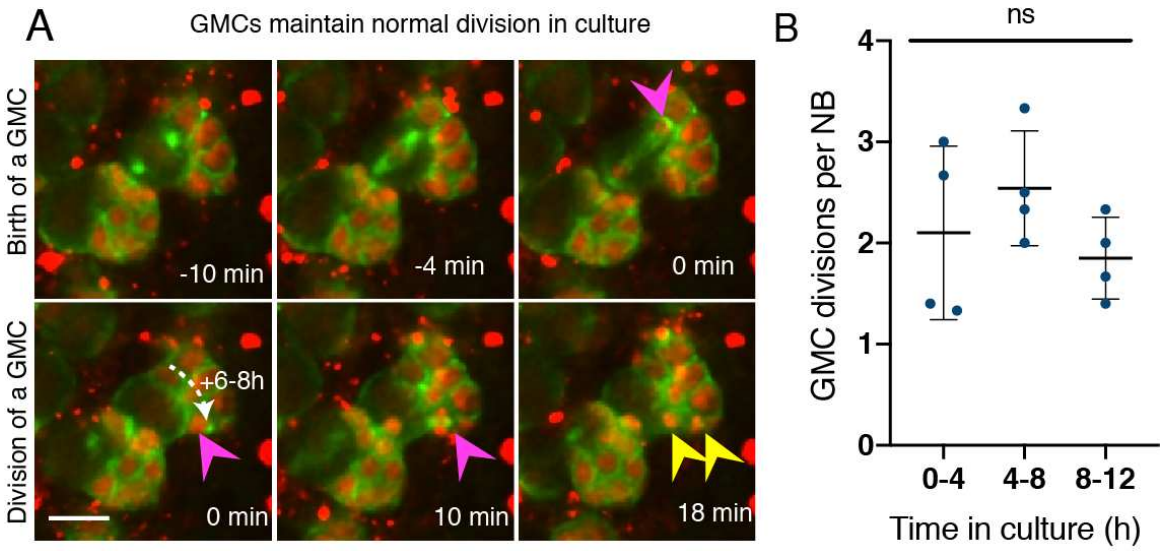

Figure 3.8: GMC division in culture. A) Newly born GMCs (magenta arrows) in culture take 6-8 h to undergo normal divisions to produce two daughter cells (yellow arrows) Scale bar is 5 µm. B) GMC divisions per NB remain approximately constant over time over 12 hours in culture (ns, one-way ANOVA, n=4 brains). Plot shows average and standard deviation over 5 brains, from at least 4 NB per brain

These changes did not improve the need to account for brain movements. I, therefore, further optimised our live imaging protocol, using a solid-agar cap (1-2% LMP agarose) placed, using tweezers, directly on top of the brains, which I found was more consistent at holding brains

against the coverslip than the earlier approach. This approach also minimises the exposure of brains to potentially damaging heat and drying due to the addition of LMP agar. I note that care must be taken not to flatten brains during this process, as it appears to result in a higher rate of stalled NB divisions which are likely artefacts. This approach reduced movement in brains significantly, but did not eradicate it, perhaps because the remaining movement is primarily the result of thermal drift of the microscope focus.

To correct this additional drift, I compared and optimised different approaches. I found that the Correct 3D drift plugin performed poorly, frequently resulting in worse registrations than the original (Parslow et al., 2014). Additionally, the large size of the images (60GB) did not lend itself to straightforward in-memory processing on standard computers. I therefore turned to using SimpleITK which is highly flexible in terms of methods and parameters, allowing control of processor and memory use. My existing scripts (**Section 3.2**) for NB registration produced highly variable results on this new data for a number of reasons. Although it is difficult to examine the cause of these variable results, I considered the following possibilities.

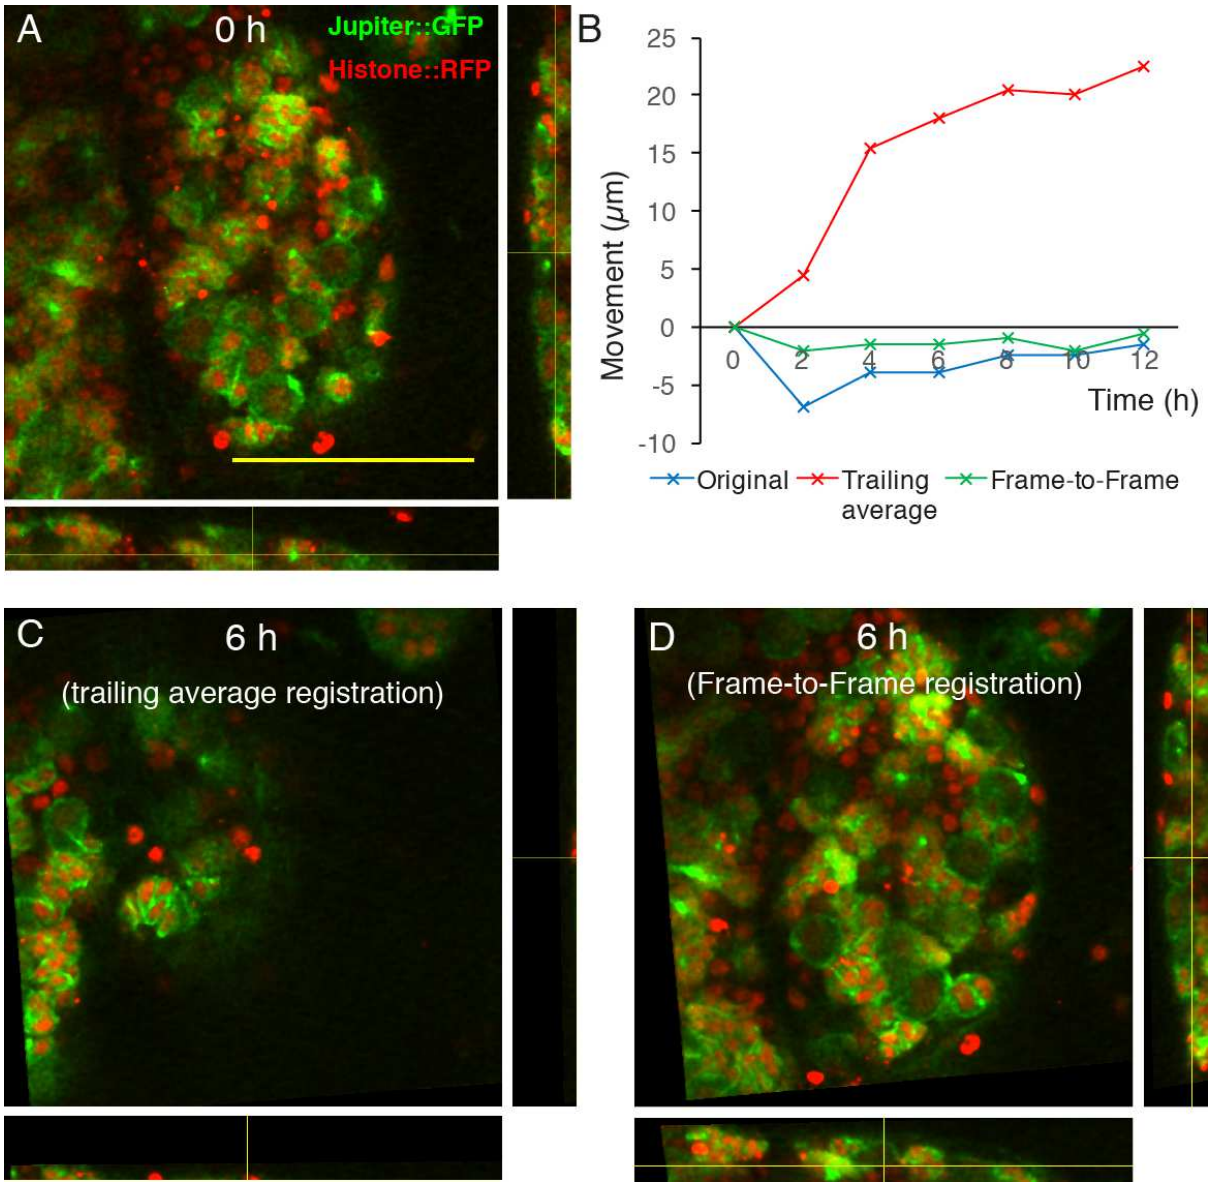

Figure 3.9: Optimisation of image registration. A) Orthogonal views of improved brain imaging (with denoising) at the start of imaging marker with Jupiter::GFP and Histone::RFP. Scale bar is 50μm. B) Movement of the brain edge over 12h imaging, showing the original movement, the slow drift of trailing average registration, and the improvement of frame-to-frame registration with median filtering. C) Result of trailing average registration after 6h. D) Result of Frame-to-Frame registration with median filtering after 6h is much more similar to the original time-point (A). Representative example shown (n=1).

First, the registration of a growing organ is complex, and there is no single registration that describes all the changes that occur. Growth of regions can be nonlinear and regions may move relative to each other. Nonlinear registration can correct for such changes (Kyriacou et al.,

1999). In this case I was not interested in correcting for these small scale changes, but for large scale movements of the brain. Therefore I used a rigid transform, using the conjugate gradient descent optimiser and the correlation metric, in an attempt to correct for large scale changes. Initially, based on my scripts for NB registration, I registered each frame to a smoothed version of previous frames of the brain (a trailing average). This resulted in excellent frame-to-frame registration, but slow, continuous drifts, likely caused by cumulative small errors due to growth of the brain that became increasingly large over time (**Figure 3.9B,C**).

Conversely, initial attempts at frame-to-frame registration resulted in sudden shifts between frames (likely caused by sub-optimal local minima in the registration). Since the registration in the current frame assumes that the previous frame is correct, accumulation of errors over time led to occasional large, inconsistent, shifts over time. These local minima can be an issue for any registration method. I investigated the possible causes for these minima. On examination of the Histone channel, I identified a possible cause: small, very bright Histone::RFP - marked foci that move very quickly (**Figure 3.10**) and accumulate in the spaces between NB niches that are not marked by Jupiter::GFP. Subsequent imaging with T. Samuels (Davis lab), suggests that these cells are glia, and that the foci are the result of glial processing of dead cells (see Discussion 3.6). The nature of these bright foci, particularly their brightness (2-4x that of nuclei) and fast movement (**Figure 3.10**), suggests that these could be disproportionately affecting the registration. To remove the effect of these objects, I turned to 3D median filtering (**Figure 3.10A-D**). By the nature of median filtering, unusually high-valued pixels (such as those belong to the foci) are replaced by the middle value of their surroundings, hence reducing their potential effect on the registration (**Figure 3.10E**). Indeed, when I use this approach, the robustness of the registration is significantly improved (**Figure 3.10E**). As an indication of the robustness, it is possible to use this registration successfully even with the raw noisy data, although the registration performance is not quite as good as using denoised data. This is in principle preferable to using it post-denoising, as the

denoising can better make use of similarities in temporally adjacent pixels, however, it requires some modification. The registration cannot make use of images with linear or higher order interpolation because doing so distorts the noise pattern that the denoising relies on, and this causes artefacts (**Figure 3.10F,G**).

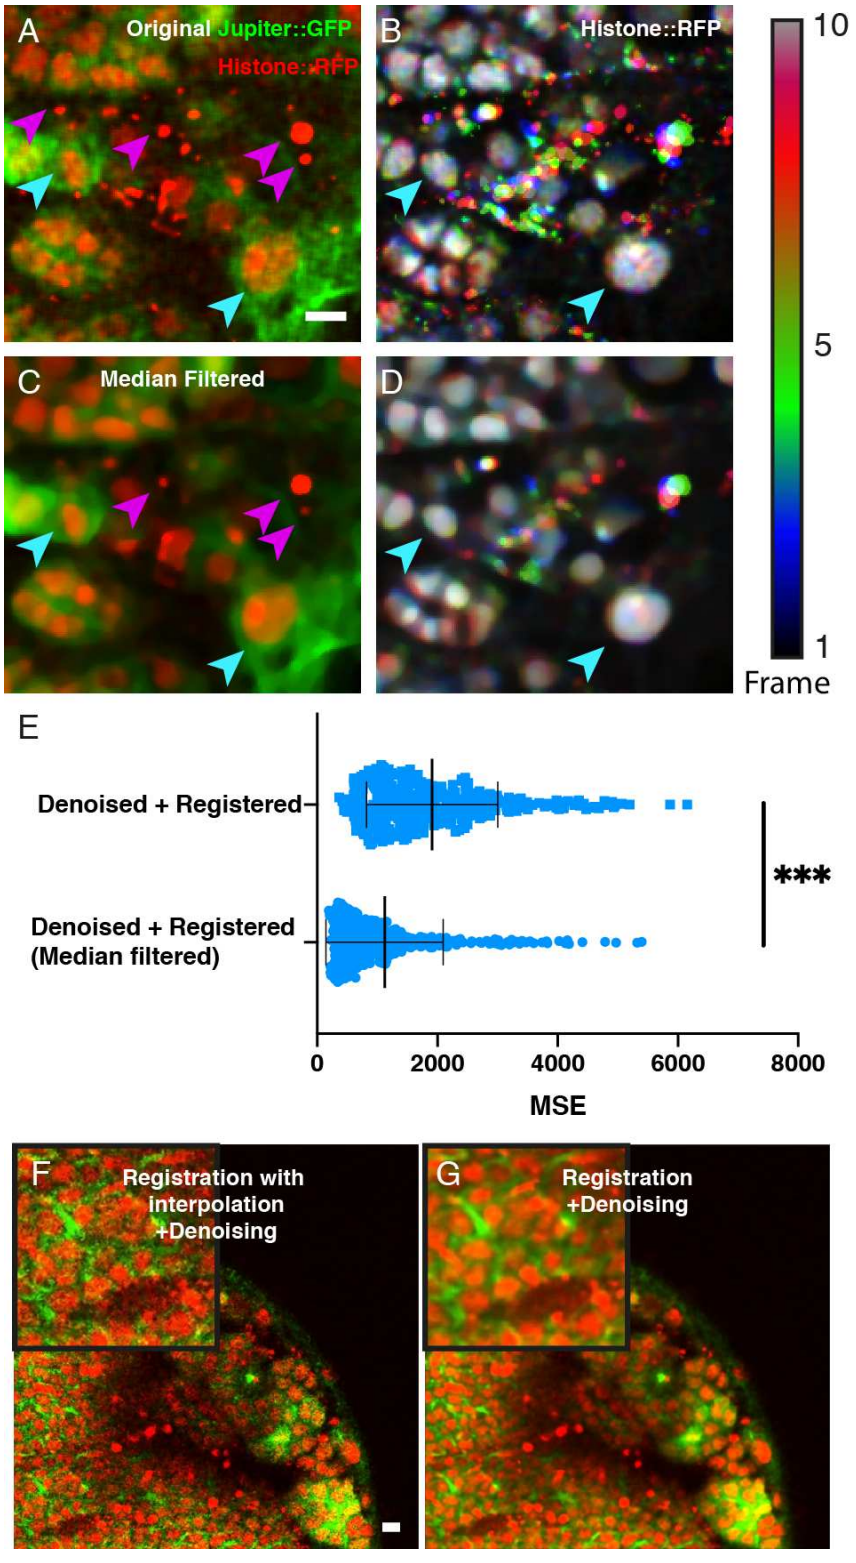
A
Original Jupiter::GFP
Histone::RFP
B
Histone::RFP
10
5
1
Frame
C
Median Filtered
D
E
Denoised + Registered
Denoised + Registered
(Median filtered)
***
0
2000
4000
6000
8000
MSE
F
Registration with
interpolation
+Denoising
G
Registration
+Denoising

Figure 3.10: Considerations in image registration. A-D) The effects of Median filtering on the contribution of small fast moving foci to a time sequence, B,D are temporally colour coded projections. E) Plot comparing Average Mean Squared Error (MSE) of the Histone channel for images with denoising with registration without Median filtering and registration with filtering. A lower MSE represents an improvement in registration quality ($p<0.0001$, t-test, n=360 frames). F) Noise-like artefacts in denoising caused by image registration with linear interpolation. G) Registering without interpolation results in the correct, smoother, denoising. Scale bar is 5μm.

I developed this combination of image registration into a processing pipeline (**Figure 3.11**). Briefly, this allows command-line image processing and denoising on a server of the .OIR (Olympus proprietary format) images into .TIFF images, and export into OMERO (Allan et al., 2012; Linkert et al., 2010). OMERO images are then exported using a custom script (credit D. Pinto for the original script) into MRC format, which is then denoised using NDSAFIR in PRISM (Carlton et al., 2010), before reupload to OMERO, and subsequent processing with my own SimpleITK based registration script to yield the final image. The scale of imaging means this process takes around 24 h per image.

Using this optimised imaging and image processing, I was able to follow the GMC cell cycle. First I followed the frequency of GMC division, which remained approximately constant over time (**Figure 3.8**). This indicates that GMC divisions were not significantly adversely affected by the imaging. Next, it was possible, although time consuming, to follow GMC divisions manually, from birth to death, and measure the cell cycle length. I will return to this idea in **Chapter 5**.

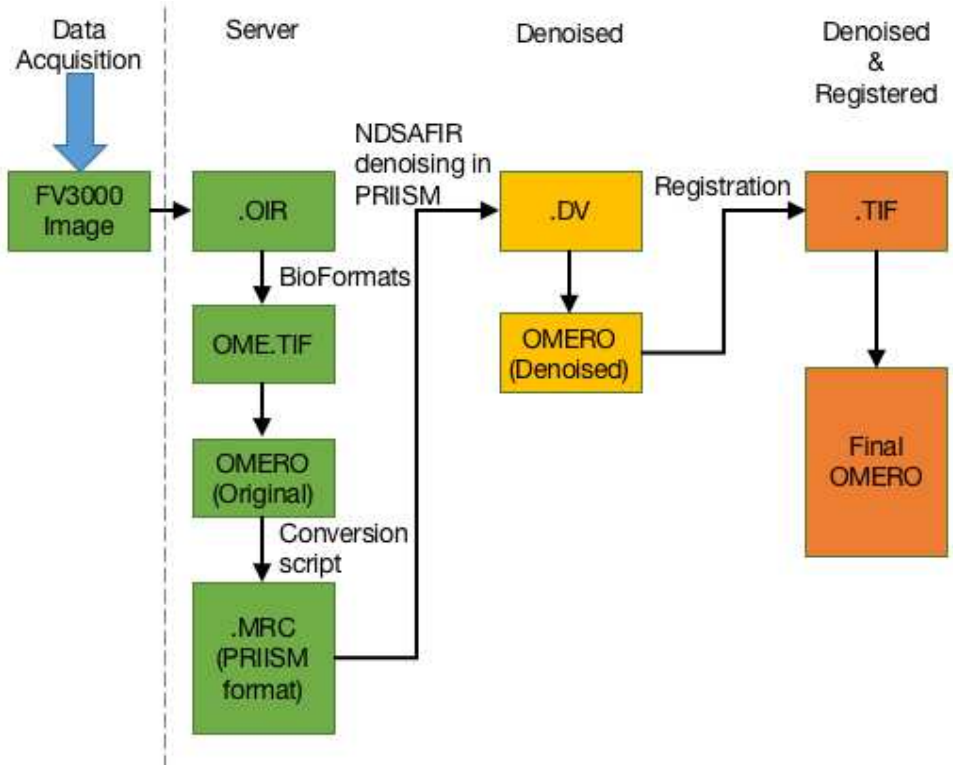

Figure 3.12: Scheme for image processing pipeline. The original image (green) is processed through a number of different formats so that it can be denoised in PRISM. The denoised image (yellow) is then registered using a custom script, and the resulting registered image (orange) is re-uploaded to OMERO.

3.5. Alternative imaging regime

Although the use of generic markers holds advantages in general applicability, allowing mutants to be studied with relatively minimal disruption, it is clear that the Jupiter::GFP, Histone::RFP marker combination is not ideal for following GMCs due to their similar appearance to neurons. Studying GMCs in live imaging is in general somewhat difficult, because, at least to my knowledge, there are no bright non-disruptive fluorescent protein traps that mark GMCs. However, the use of markers expressed under control of a GAL4/UAS system can allow better discrimination of GMCs and neurons. *ase*-Gal4 (see **Methods**) is in principle expressed only in NBs and GMCs, although in practice the lag-time due to the GAL4/UAS system and degradation of the protein means that other cells in addition to GMCs are marked. Using Ase-Gal4 to drive expression of mNeonGreen, I measured approximately 20

GMCs/Neurons marked per NB, with about half expressing mNeonGreen strongly. Nonetheless, imaging of *ase*-Gal4-driven mNeonGreen and Histone-RFP results in images that are much easier to follow and annotate (**Figure 3.13**). In principle, these could be used to follow GMC division more straightforwardly than previously documented, albeit with reduced ability to perform genetic manipulations (**Figure 3.13**).

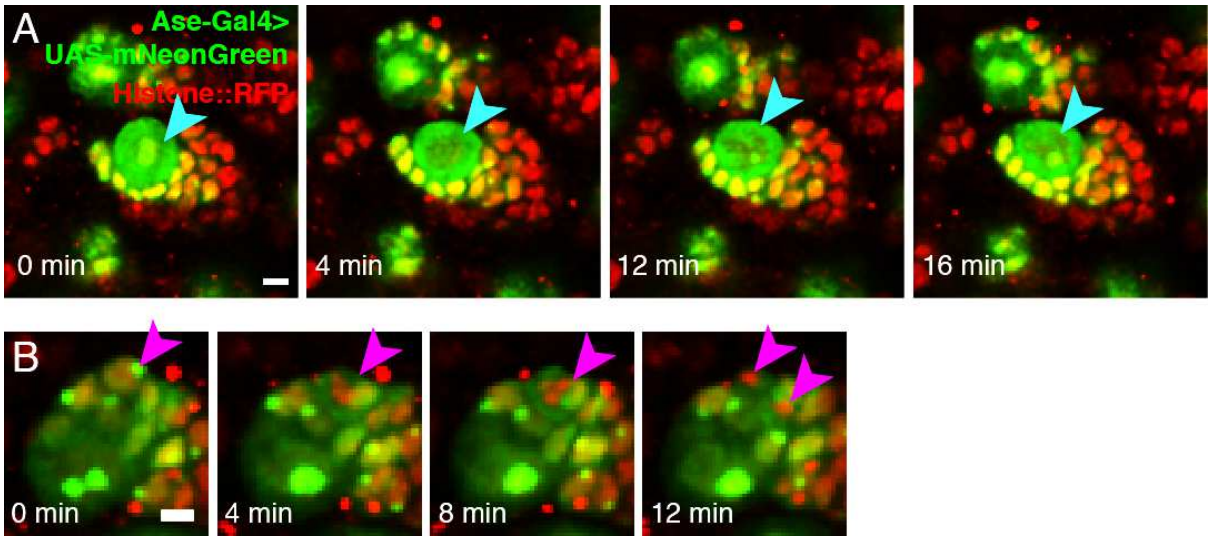

Figure 3.13: An example of imaging of alternative markers. A) Imaging of Ase-Gal4>UAS-mNeonGreen; Histone::RFP showing a single NB (Cyan) division (breakdown of nuclear envelope) over 20 minutes. Multiple daughter cells are marked by the mNeonGreen, but it is much more restricted than Jupiter::GFP (c.f. Figure 3.9), providing better distinction of cell types. Scale bar is 5 µm. B) Condensation of chromosomes (red) during mitosis of a GMC (magenta). Scale bar is 5 µm.

3.6. Discussion

I have developed and validated *ex vivo* brain culture imaging and processing for analysis of *Drosophila* neural lineages. I have shown that the live imaging and culture supports the measurement of NB division rate and the GMC cell cycle. These are the first reported measurements of the cell cycle of GMCs in an *ex vivo* system. The resulting datasets can be very large, rendering manual analysis time-consuming on individual movies, and is impractical at a larger scale. Developing an approach to automatically analyse these movies will be the subject of **Chapter 4**.

Although I have developed the approach for imaging of the central brain (which frequently gets the most interest due to the similarities to human neural lineages), there is nothing to prevent its use in analysing other parts of the brain such as the optic lobe and VNC. Similarly, although I have focussed on wL3, it is in principle possible to use this approach to examine other developmental stages. Of particular interest would be late L3 pupal transition, as well as the exit of NB from quiescence in early L2. Similarly, the study of other cell types and processes are possible. Indeed, imaging of the CB region using glial markers suggests that this method could be used to study cell death and glial processing (**Figure 3.14**). At this point it is unclear whether the cell death that occurs is natural or imaging-induced, but given that cell death is known to be a key part of tissue patterning (Pinto-Teixeira et al., 2016), it is tempting to speculate that this may be used to study programmed cell death.

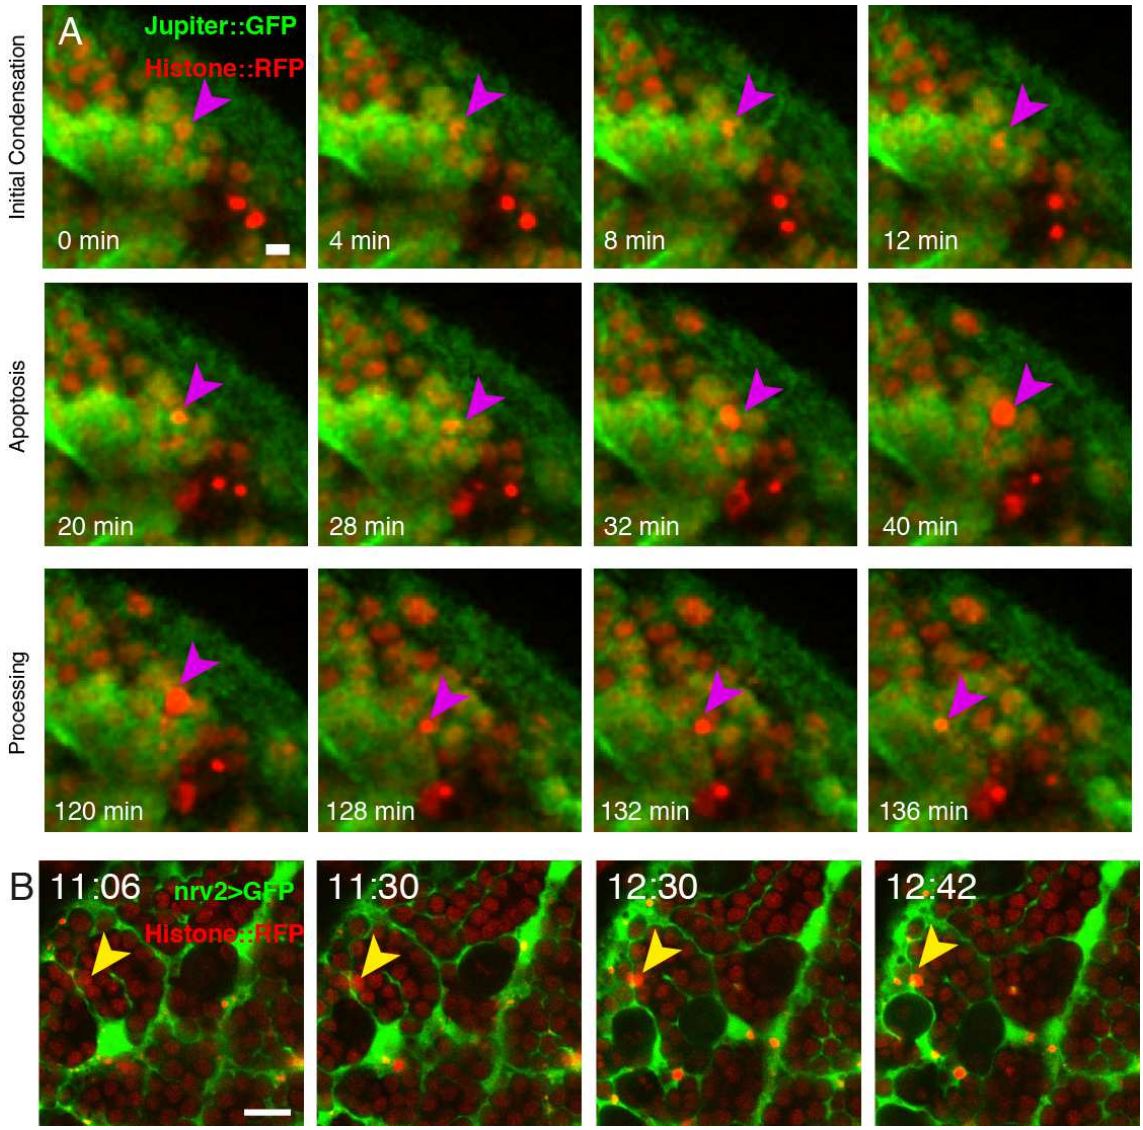

Figure 3.14: Cell death in the developing brain. A) Time series of neuronal cell death (initial condensation of chromosomes to bright red spot, followed by rounding up of the cell) (magenta arrow) and processing (presumable engulfment of cell debris, and movement away from site of cell death). Images are a maximum projection over 3 µm for clarity. Scale bar is 5 µm. B) Imaging of nrv2-GAL4>UAS-GFP (glial marker, see **Methods**), Histone::RFP, showing death of a cell (yellow arrow, indicated by rounding to bright red spot, and surrounding by glial projections (green)). Timestamp is hours: minutes. Scale bar is 5 µm

3.6.1. Limitations

The need to minimise the light dose for effective imaging, particularly of GMCs, means that live imaging can only be one tool among many. To follow dim markers (for instance tagged proteins that are natively expressed at low levels), the live imaging described above would

struggle to allow good enough time resolution to follow many processes. For this kind of imaging, isolated cells would be more appropriate (although, obviously, they may not represent the *in vivo* state so well). On the other hand, fixed imaging with antibodies against specific proteins can provide clarity of cell identity with an ease that is difficult to achieve with fluorescent markers, although it is more difficult to directly study cell behaviour.

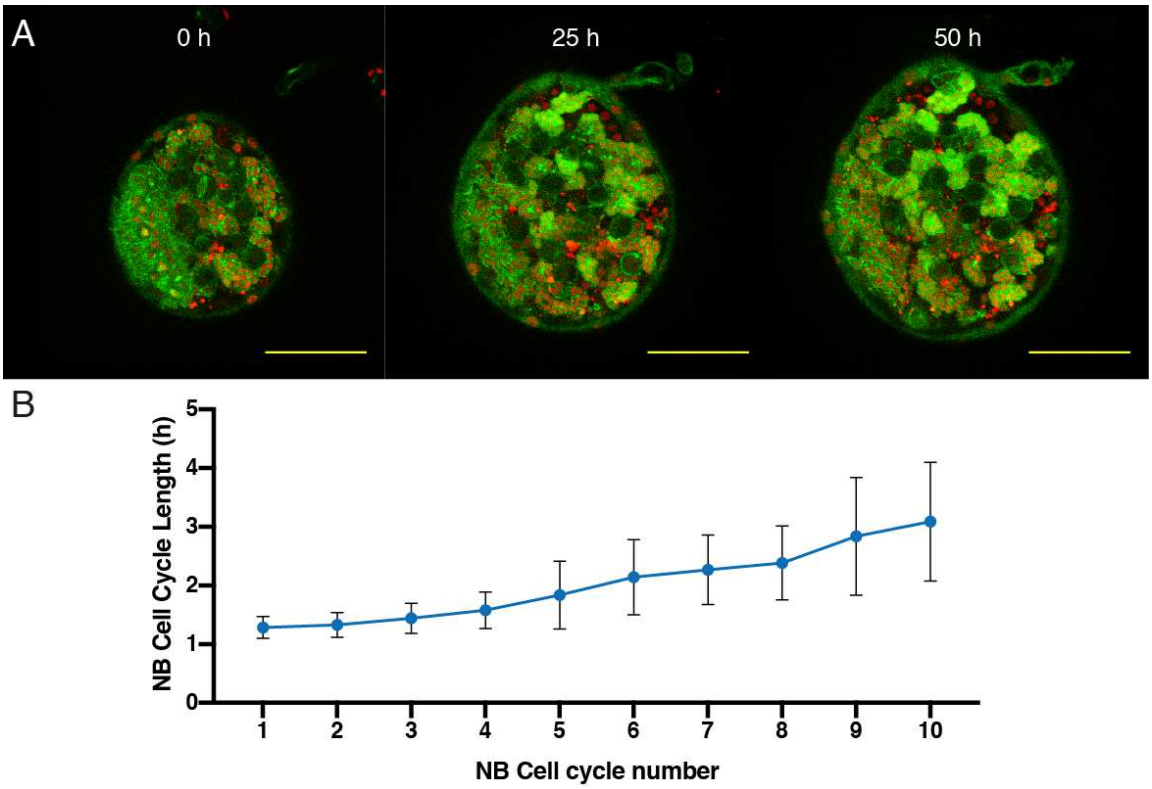

Figure 3.15: Imaging of wL3 larval brain for 50h. A) Images of the brain as it grows significantly over the 50h period. Scale bar is 50µm. B) Plot of NB Cell cycle length vs cell cycle number. Cell cycle length increased from a normal ~1.5h to over 3 hours after 10 cycles (~20h), after which many NB divisions stalled.

Although most of the imaging I have shown here is for 24 h or less, I have made attempts to image for longer periods. Imaging for 50 h starting with a wL3 brain highlights some interesting issues with the culture. The length of wL3 is typically around 24 hours. *In vivo* larvae pupate following wL3, and neuroblasts shrink and undergo terminal division. However, under imaging, no neuroblasts were observed to significantly shrink and undergo terminal division. Instead the rate of the division visibly slows, and eventually the NB stop dividing

(**Figure 3.15**). This is perhaps unsurprising, as it has been previously noted that ecdysone (Riddiford, 1993) is required for this transition. Nonetheless, this live imaging setup could be an interesting system with which to study the brain in the pupal transition.

Of interest is that the measurements of cell cycle length approximately tally with those in Bowman *et al.* (2015) but not with the later *in vitro* measurements (Bowman et al., 2008; Homem et al., 2013). This suggests that the *ex vivo* model of brain development recapitulates the *in vivo* situation, and that the *in vitro* culture measurements of isolated cells do not. It would be interesting to consider measuring the Type II NB lineage with this method, and to confirm whether existing measurements of this lineage can be replicated.

CHAPTER 4

ANALYSING LIVE IMAGING DATA: CYTOCENSUS DEVELOPMENT AND VALIDATION

In **Chapter 3** I established a method for live imaging of neuroblasts and GMCs within the *Drosophila* brain. This live-cell 3D imaging results in datasets consisting of thousands of images. The large scale of these datasets makes manual analysis time-consuming at best and impractical at worst. In this chapter I will address this problem, with the development and validation of a program, CytoCensus, for automatic analysis of large scale imaging data.

Practical analysis of large imaging datasets is a difficult challenge. Progress in understanding the mechanisms of brain development has depended on the characterisation and quantification of mutant phenotypes by painstaking manual image analysis (e.g. Neumüller *et al.*, 2011). Therefore, there is a general need for automated tools to assist in the quantification of this kind of data (Rittscher, 2010). Various tools based on conventional image analysis (e.g. MINS) exist for producing quantitative data from images, (see **Introduction 1.2**)(Lou et al., 2014). These programs can perform very well at specific tasks but are often complex, highly specific and/or require significant knowledge of image analysis and programming to use. To this end, new programs, for example, Ilastik, ACC and the FIJI-Weka plugin have been developed that apply machine learning (ML) to this problem (Eliceiri et al., 2012). Supervised ML reduces the need for image analysis and parameter optimisation, instead requiring user annotations to train the program. In the context of my *Drosophila* brain data,

these programs performed poorly, making the development of an appropriate program a necessity.

I address this issue with CytoCensus, a program intended to assist in the counting and identification of objects in 3D/4D images. CytoCensus is an easily deployed, machine learning based software. CytoCensus facilitates automated analysis of cell types and quantitative analysis of cell number, distribution and proliferation from complex tissues such as the *Drosophila* larval brain.

The original program (under the names Quantifly and QBrain) was designed and implemented by Dominic Waithe (Waithe et al., 2015). Implementation of CytoCensus to handle 3D and 4D datasets was primarily done by myself (stated in text), including the ROI tool, improvements to the DetH maxima finding approach and user interface changes. All validations and images shown here are my own.

4.1.1. Specific Aims

Develop an approach that uses the benefits of supervised machine learning for 3D cell detection.

Develop and validate a user-friendly program for analysing large scale live imaging data.

4.2. Theory of the Method and Existing State of the Art

4.2.1. Previous Work

In this section I detail previous observations and the work that CytoCensus builds on. CytoCensus builds on the concept of density learning, initially developed for pedestrian counting problems and more recently applied to microscopy images (Lempitsky and Zisserman, 2018). The original idea of density learning was to minimise the amount of effort required to count objects in images. In the case of (Lempitsky and Zisserman, 2018), they make the case that little additional effort is required for a human to make point annotations over

counting objects (**Figure 4.1**). They make use of this additional information (the cell centres) to improve the quality of whole-image *cell counts* by making a 'density map' (**Figure 4.1**).

The key feature of the density map (**Figure 4.1C**) is that the number of cells in a region (**Figure 4.1D**) is given by the sum over that region. Equally, the sum of the density map over the whole image gives the number of cells in that image (8 in **Figure 4.1**). Because the density-map only seeks to count cells within a region, and not identify their location, it is robust to poorly separated objects and occluded objects, providing accurate counts compared with segmentation, with a reduced requirement for annotation. Consider, for example, two cells that are touching. In segmentation the lack of a clear boundary means these cells are likely to be considered the same object, which complicates counting. In density learning, the fact they are touching is irrelevant, as the cells both contribute to the density map, and the sum over the region is largely unaffected by their nearness.

These advantages were noted, and applied to *cell detection* in Kainz *et al.*, who altered the map to form a 'proximity score map' which better represents the locations of cells (**Figure 4.1**)(Kainz et al., 2015). In this study, the cell centres are identified from the score map using non-maximum suppression, which suppresses maxima below a threshold. Unlike the previous work, this provides information on the locations of cell centres, rather than just counts.

The key difference between the cell detection and cell counting approaches is the way that interacting objects are considered. In the density map, pixels that are near to multiple objects have the contributions from each object added together (**Figure 4.1C**). In the proximity score map, the interacting object that is furthest away (and hence with the lowest proximity score) dictates the final proximity score. In effect this introduces a boundary (**Figure 4.1G**) between the two objects, and ensures that there is a local maximum that corresponds to each object.

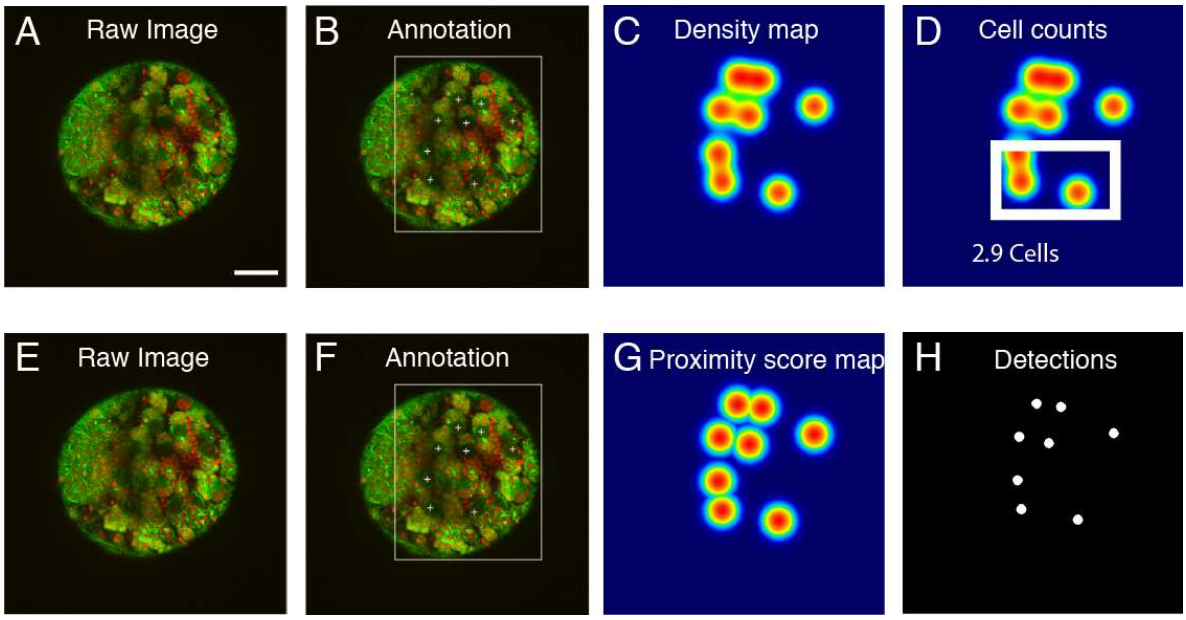

Figure 4.1: The difference between density counting and cell detection, illustrated with theoretical maps. A) Raw image; B) Image with cells (NB) annotated. C) Corresponding density map representing these annotations. D) Counting with the density map: the sum of the density map in the selected region reflects the number of cells in that region (~3). E,F) Raw image and annotations. G) Proximity score map (note the separation, rather than merging of objects). H) Detections (Maxima) based on the proximity score map. Scale bar is 50 µm.

In following section, I examine the application of this density-map based approach to cell detection in 3D and 4D microscopy images, including my contribution to (Waithe et al., 2016).

4.2.2. Application of the 'Proximity Score Map' in CytoCensus

The density counting approach described above can, in principle, be extended to 3D microscopy images. The direct application of this density estimation to 3D is detailed in (Waithe et al., 2016). However this process requires 3D annotations of cells. In microscopy images accurate 3D annotation can be highly complex and time-consuming. Without resorting to 3D virtual reality, annotations must be made in 2D, either using a 3D rendering, or by sequentially examining a series of 2D slices. It has previously been commented upon that this 2D stereoscopic approach is far from ideal (Schmitz et al., 2014). We examined approaches for making these kind of annotations, for instance 3D renderings in Imaris, and sequential 2D labelling using TrackMate in ImageJ. Both were time-consuming for even small datasets of 5

or so images, even with relatively uncommon cells. The difficulty of 3D annotation makes the direct use of the density counting approach less attractive, despite its useful properties like robustness to high cell densities.

4.2.3. APPLICATION TO 3D

Given the complexity of 3D point annotation we therefore considered other annotation methods. Segmentation of images can be performed using 2D annotations, but for good performance require extensive pixel accurate annotation, in effect 'painting' 2D images with labels. Additionally, segmentation based approaches struggle with the densely packed cells that are common in tissues. Instead we consider using 2D point annotations and attempt to extract the 3D cell centres. However, extracting 3D information from these annotations is not straightforward. We therefore adapt the density learning approach, which has the potential for better dense cell detection, and requires reduced annotation.

We depart from the assumptions of the density counting approach. Instead we use a centre-of-mass approach, similar to the proximity map of Kainz et al. (2018) creating a map that corresponds directly to the locations of the cell centres (**Equation 4.1**). In this map, the value of each pixel is given by the value of a normalised Gaussian centred on the nearest cell centre, with a user-specified 'sigma' parameter that corresponds to the size of the object of interest (**Figure 4.2**). This centre-of-mass (COM) map represents the proximity of pixels to the object centre.

$$\forall p \in R_j, \;\; F_j^0(p) = max\left[\mathcal{N}(p; pt, \sigma^2), \forall pt \in \mathrm{P_j}\right]$$

Equation 4.1 Equation describing the calculation of the COM map. p denotes a pixel in an image region used for training R. F(p), is the value of the COM map at p, defined as the the maximum value of all gaussian kernels N (of standard deviation σ) centred on each cell centre point pt (from the set of all points Pj) at point p.

We then use machine learning to learn the correspondence (i.e. the Model) between the image of interest and this proximity map (**Figure 4.2**). First we calculate a series of features based on the image (**Figure 4.2**). Next we subsample the features and the COM map (**Figure 4.2**), to reduce computation time and minimise overfitting (see **Discussion 4.6**). We then train an ensemble of decision trees (commonly referred to as a random forest) using regression to predict the COM map based on these features. We can then use this ensemble to predict the COM map on unseen images (**Figure 4.2**).

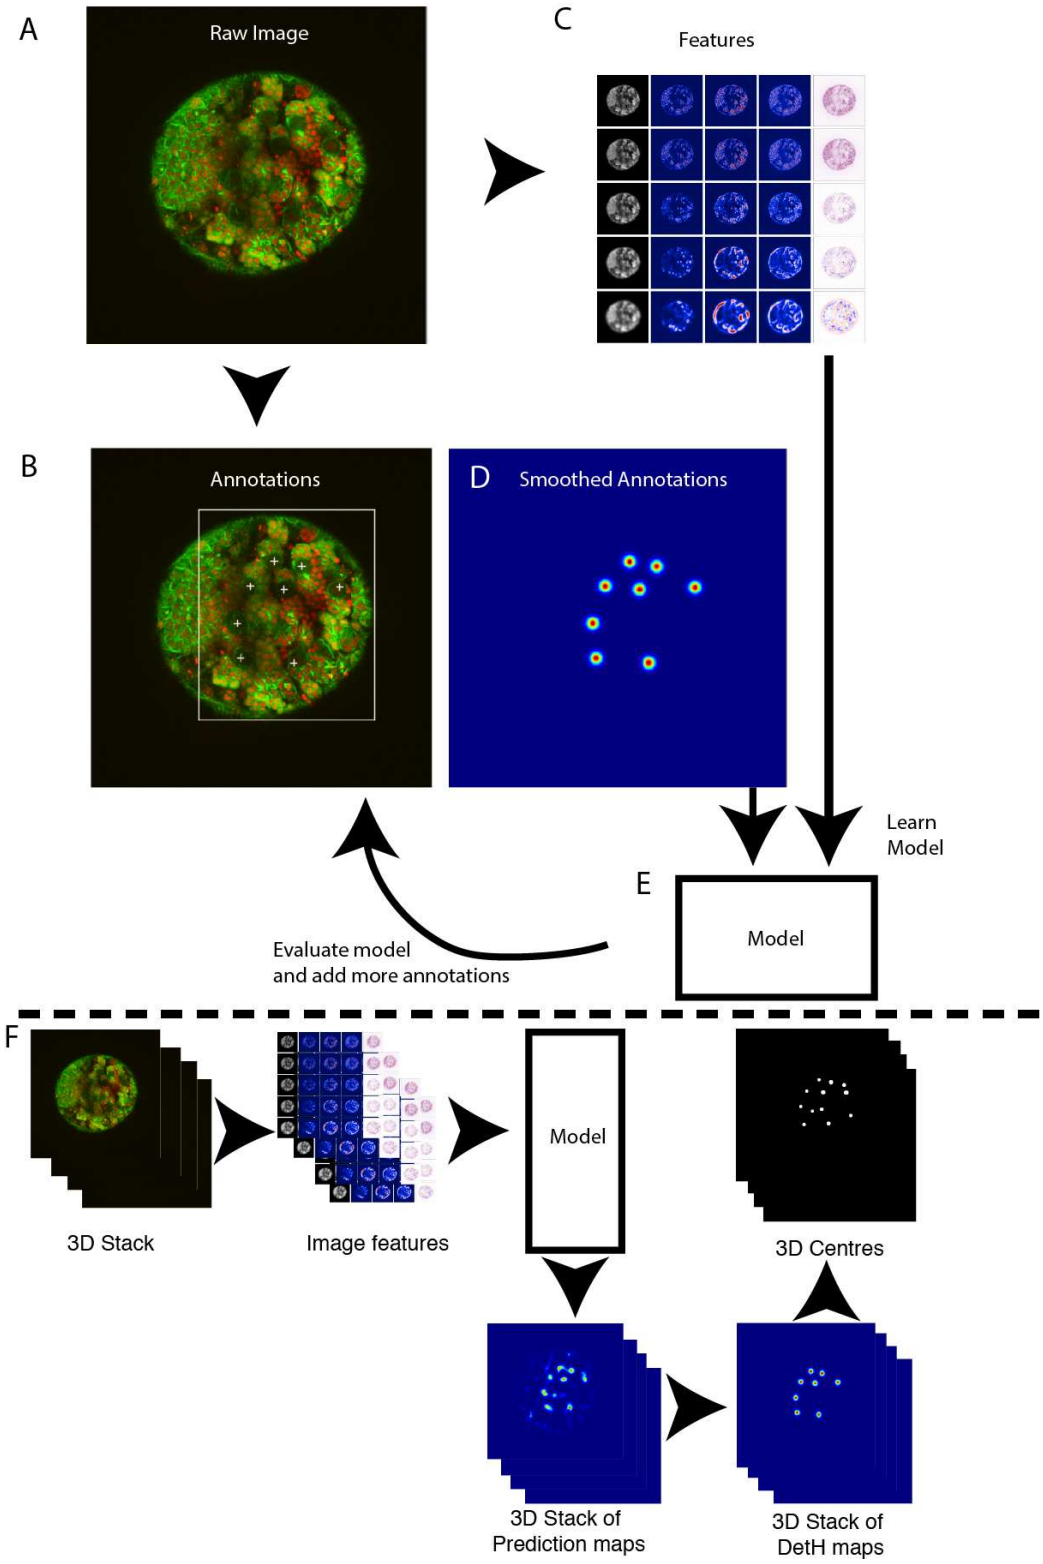

Figure 4.2: Scheme for 3D centre of mass (COM) finding. A) Raw image; B) User annotations of image; C) Features calculated from the raw image (A); D) Smoothed Annotations (COM/proximity map); E) Learned Model; F) Evaluating the Model on unseen data - features are calculated and the trained model applied to produce a set of prediction maps. Applying the Determinant of the Hessian (DetH) and finding the maxima then finds the 3D centres.

Given this COM map, we can then find the 3D centre of the cells. Other approaches have been taken to finding the centre or cells or objects previously, including non-maximum suppression (Kainz et al., 2015). The proximity map has an advantage over the density-map, in that each cell centre should, at least theoretically, follow a Gaussian distribution. In principle, a maximum filter would be sufficient to find the centres of the Gaussians as in Kainz et al., (2018) but this is susceptible to noise. Instead we use the Determinant of the Hessian matrix (DetH) applied at a given scale (sigma). This filter searches for regions of locally quadratic curvature that approximates a Gaussian of a particular size (defined by sigma). This effectively allows the region surrounding the cell centre to contribute to the detection of a maximum. This also adds a size selection aspect; large regions of high proximity score will not affect the local curvature (as they are largely flat), whereas they would be detected as maxima. Following the calculation of the DetH image, a maximum filter with a radius given by the sigma of the original Gaussian prevents multiple detections within an object, and a threshold value allows the user to supress objects of low quality. In principle, this filter then directly yields the 3D cell centres.

The DetH calculation in 3D relies on the assumption that in 3D the objects in the COM-prediction are approximately Gaussian shaped. However, since we use 2D annotations, this is hardly a guarantee. In fact, the 3D shape can vary significantly. The COM map on which the ensemble of trees is learned, if fully annotated, resembles a smoothed line in 3D (**Figure 4.3D**). However, the ensemble is never perfect, and so, in practice, the predicted COM tends to be smoothed somewhat (**Figure 4.3B**). The degree by which is smoothed depends on a number of factors, of which the degree of correlation between planes, and how features change in Z with relation to the object are the most important. In many cases, particularly for nuclear or cytoplasmic markers, features correlate in Z such that the 3D shape is not that dissimilar from a Gaussian (**Figure 4.3B**). These shapes can be detected by the DetH regardless of the true shape, although the smoothed line shape requires an increased sigma of the DetH

in Z for optimal detection (**Figure 4.3**D-F). This need for asymmetric detection in Z is reflected in the parameter choice in CytoCensus.

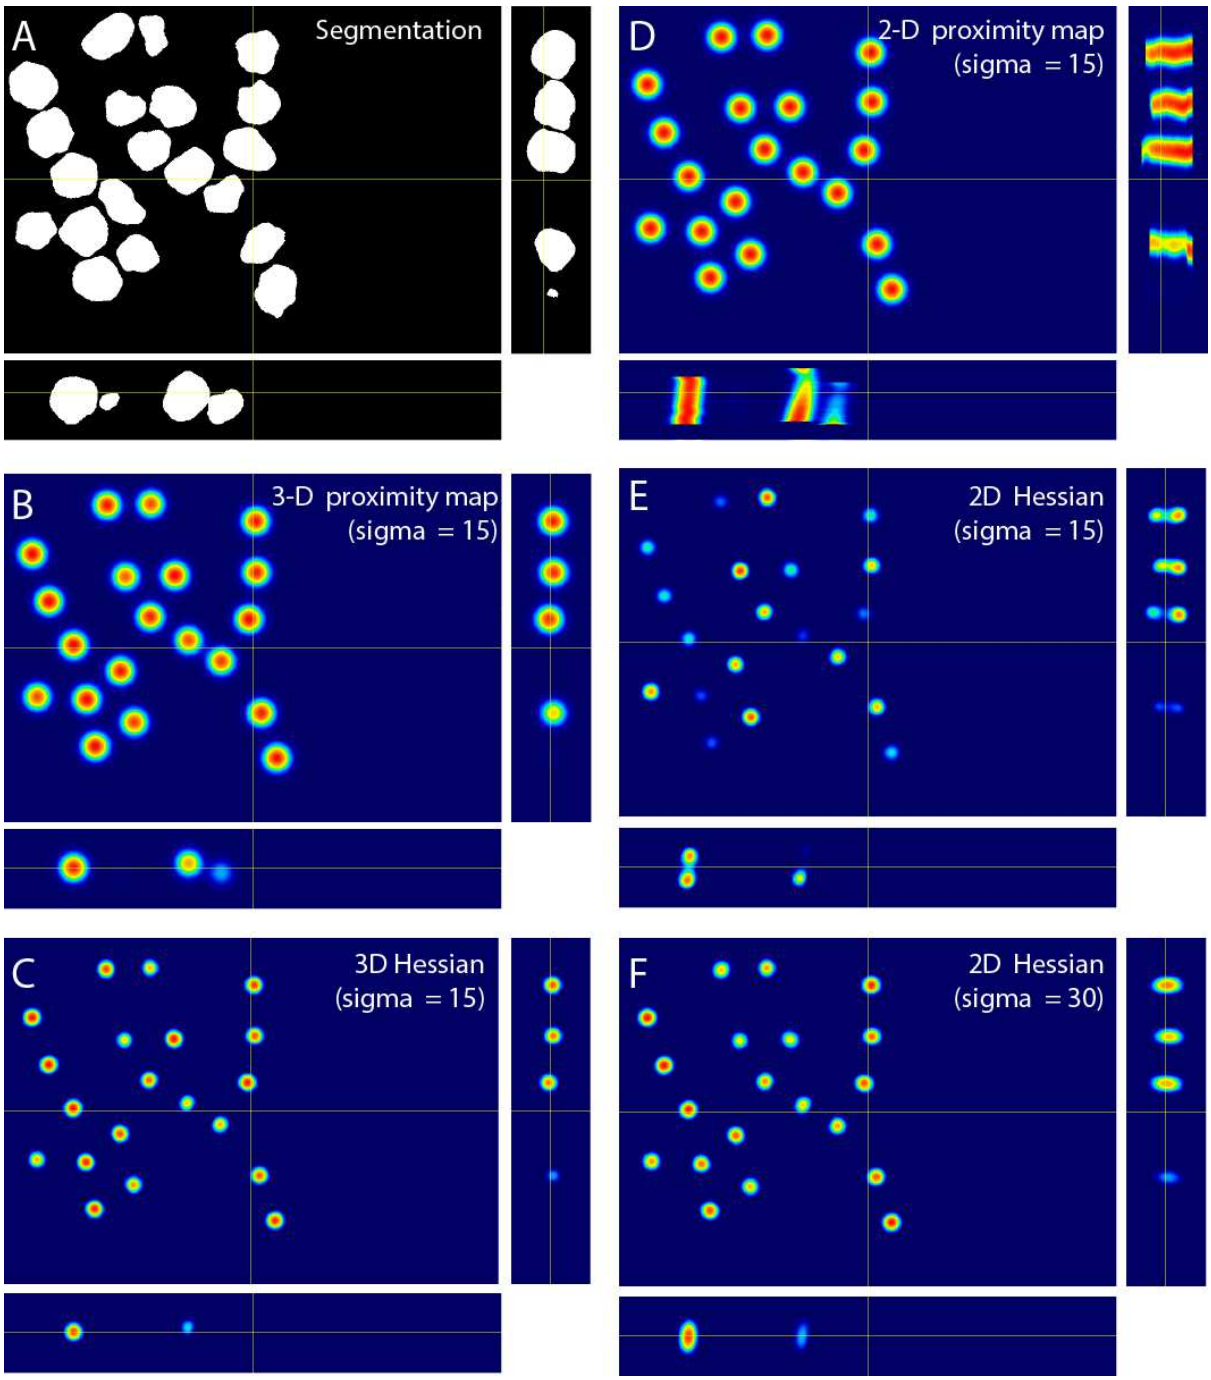

Figure 4.3: Shape of 3D proximity/COM map. A) COM map from 3D KDE. B) Shape of COM map C) Theoretical predicted COM map with no correlation. D) Theoretical predicted COM map with correlation. E) DetH of (C). F) DetH of (D).

4.2.4. LIMITATIONS

In effect, the consequence of using 2D annotations makes the problem of detecting cell centres in 3D poorly defined in Z. This lack of a defined shape means that optimal detection either requires image-specific optimisation by the user, or use of a sub-optimal sigma. It is worth noting, however, that the resolution of typical microscopy images is usually much lower in Z than XY anyway, both in terms of sampling resolution and image resolution, due to the limits of optical microscopy. This means that objects are not isometric (represented by an equal number of pixels in each dimension), and consequently this parameter may need optimising independently of the assumptions of this approach.

The use of a defined sigma limits the use of this approach to approximately similar sized objects: cells with highly dissimilar sizes may not be detected appropriately (either by under-detection of small objects or splitting of large objects, depending on the value of sigma). Theoretically, it could be possible to allow variable cell size, although this would require features to correlate with objects across different scales, which is not the case for the current feature set. This topic, including solutions and workarounds will be returned to in **Chapter 6**.

4.3. DEVELOPMENT OF CYTOCENSUS AS A USER FRIENDLY TOOL FOR OBJECT DETECTION IN 3D

4.3.1. MOTIVATION

Description of possible methods for cell detection in computer vision are common, but even for highly successful methods, few achieve frequent use in biological image analysis. One of the obstacles to this is the long lag time between published methods and translation into a format useable by biologists. In this light, we developed CytoCensus into a general-use program. The development of CytoCensus for use by biologists without knowledge of programming requires some additional implementation over the theory of the method described above. The program I present here, CytoCensus is built on Quantifly, an open-source

software developed by D. Waithe for counting *Drosophila* eggs in images. This program was adapted for 3D counting, with initial work by D. Waithe and subsequent development by myself, to make it appropriate for use with 3 and 4D datasets. This section details my main contributions to the development of the program CytoCensus.

We make use of the density-learning based approach described above. In practice, the user simply selects, with a single mouse click, the approximate centre of all examples of a cell type within small user-defined regions of interest in the image. To optimise training, users can apply the probability map output as an overlay to assess the accuracy of the prediction and if necessary, provide additional training. I find this supervision approach is more convenient and faster than other machine learning based approaches, such as FIJI-WEKA (Arganda-Carreras et al., 2017) or Ilastik, (Sommer et al., 2011), which, for optimal performance, require extensive labelling of areas at the pixel level that include both cells and background (**Figure 4.4** for details of the user interface). CytoCensus then learns a model based on this training that can be applied to unseen data. During the learning of the model, CytoCensus chooses from a collection of filters and scores and combines them to find features in the image that identify the user-defined cell centres. In this way, a series of transformations of the image data (referred to as the "Model") are learned. This model is applied pixel by pixel to further data sets, and CytoCensus then outputs both an estimated probability map of cell identification and the predicted coordinates of the cell centres. This probability map and the predicted locations of cell centres across the entire volume and time-series are saved and can be passed to ImageJ (FIJI) (Schindelin et al., 2012) for further manipulation. Thus, the program requires a minimal amount of manual user input to achieve accurate automated cell identification that is scalable to large image data sets, such as live imaging of the *Drosophila* brain. In the rest of this section I will examine various application-specific modifications and additions to this approach to allow user-friendly annotation of images.

4.3.2. THE *DROSOPHILA* CENTRAL BRAIN: A CASE STUDY FOR CYTOCENSUS

I will demonstrate the use of the program using an example from the live imaging data of NB division (**Chapter 3**). Using this as a case study, I aim to identify NBs and GMCs/neurons within the CB region of the *Drosophila* brain. I will return to a discussion of the biological aspects of this in **Chapter 5**.

Although I have outlined the method above, adapting it into an easy-to-use program requires some additional work. In this section, I describe the most important aspects of presentation and optimisation of CytoCensus, based on the motivation of creating a user-friendly program. In the subsequent section, I describe various optimisations to the approach and available parameters based on the central concept.

4.3.2.1. INTERFACE AND USE OF CYTOCENSUS

The processes of annotating images and of applying a learned model to large datasets have different aims. In order to reflect this CytoCensus consists of separate training and evaluation interfaces (**Figure 4.4**). The purpose of the training interface (**Figure 4.4**) is to allow the user to make annotations and evaluate the performance of the learned model. The evaluation interface then allows the user to apply this training to any number of additional files and or timepoints, and must potentially cope with large datasets.

The training interface therefore has somewhat different requirements to the evaluation mode. First, it contains the tools for annotation (**Figure 4.4**), similar to those found in the 2D Quantifly program (Waithe et al., 2015). These tools allow the user to select a region within the currently displayed image to annotate (**Figure 4.4A**). This makes annotation more flexible and less time consuming than forcing the user to annotate an entire image plane. Second, to minimise the time to evaluate the performance of the model under the current training, the model evaluation is restricted to the current time point and file. Third, features are, where possible, only calculated once. Keeping these features in memory means that evaluating the

model subsequently is very fast. However, it limits the number of annotated slices that can be held in memory. Given that the aim is to minimise user annotation this seems reasonable, but this could be problematic for images that are very large in xy or z. A simple workaround however, is to crop subregions of the image before passing to CytoCensus. Together these aspects speed and ease user training.

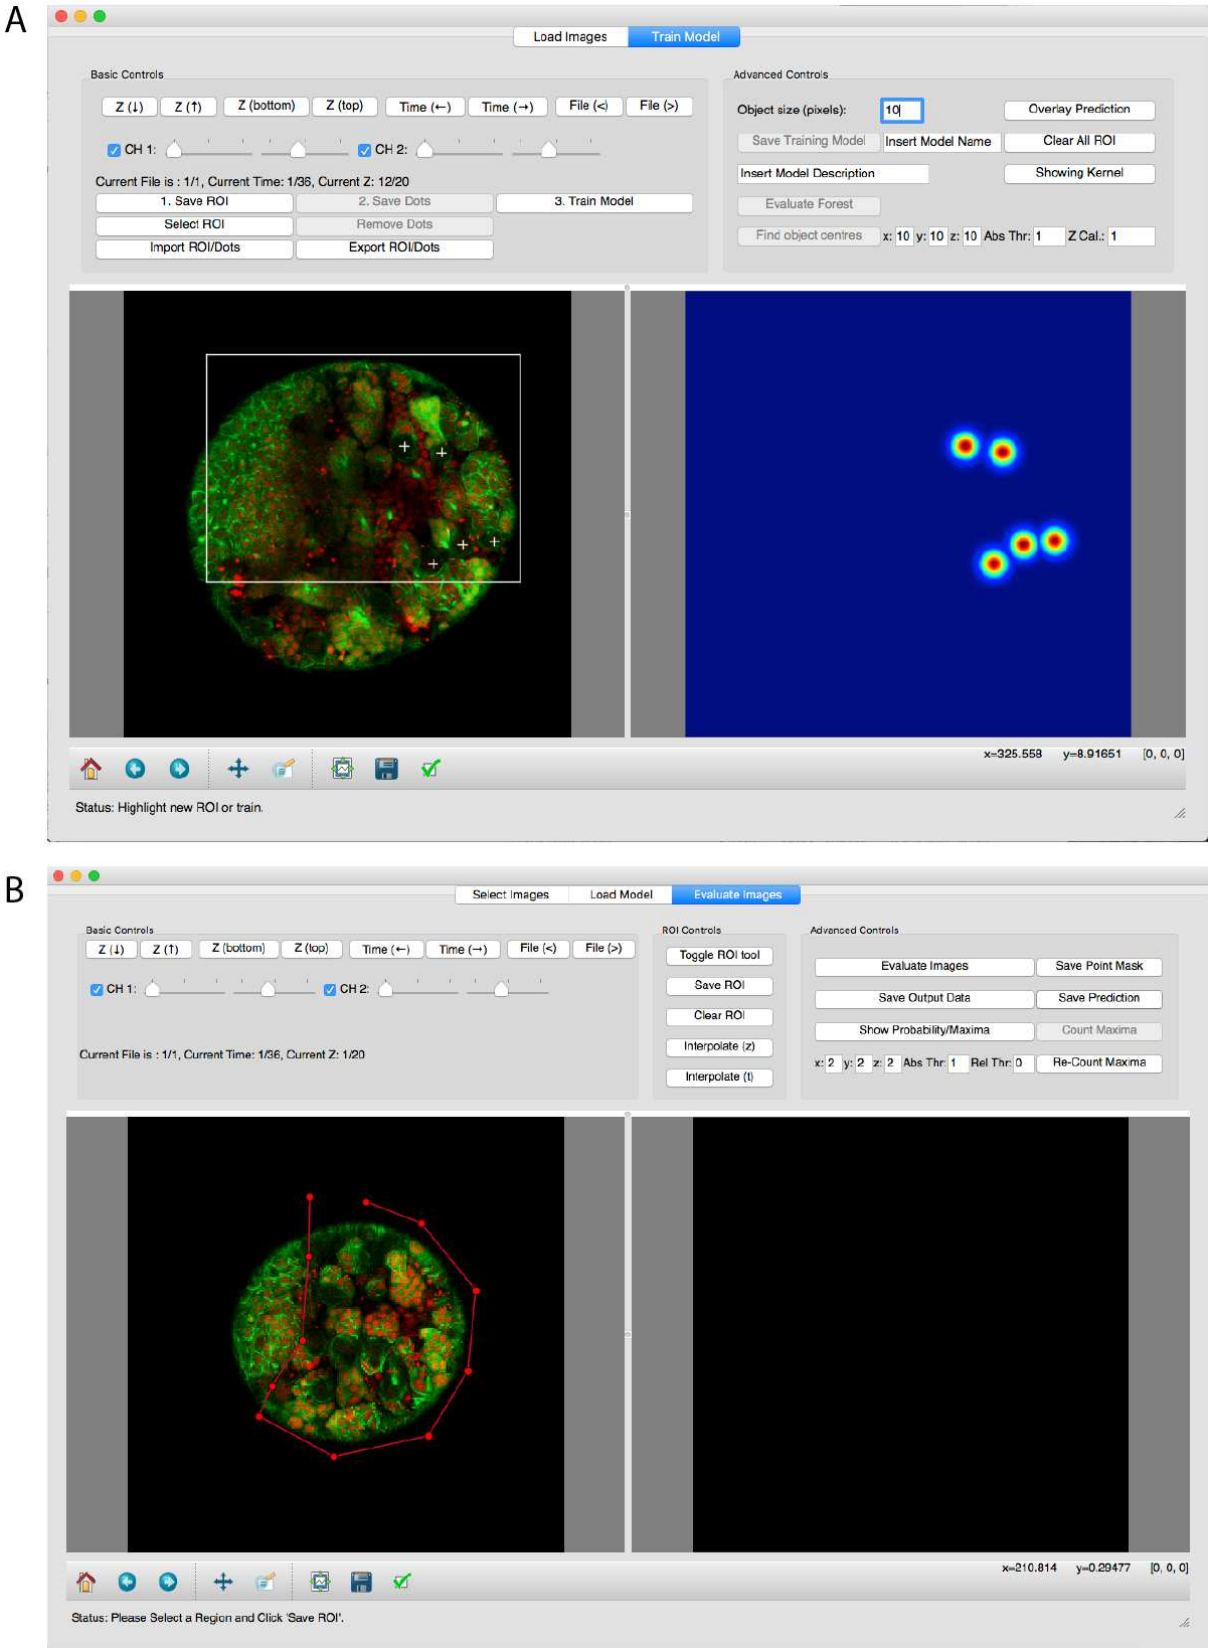

Figure 4.4: Training and evaluation interfaces of CytoCensus. A) Training interface, showing annotations (left), and COM map (right). B) Evaluation interface showing image with ROI

In the evaluation interface the aims are different (**Figure 4.4B**). Instead of training a model, a model is selected and evaluated on a dataset. At this stage, we are only interested applying the model to additional images (rather than training the model), so we need only save the prediction maps and final cell centres. User annotation is unnecessary, and features are only needed transiently for calculating the prediction, so need not be kept in memory. Because we train on a subset of the images, it is possible the remaining dataset has a different optimal threshold to the training set, so we allow the user to adjust the DetH parameters and threshold finding at this stage. In addition, outputs are provided in different formats, either as TIFF images or CSV files, for convenience of subsequent processing. TIFF images can be easily opened in FIJI (Schindelin et al., 2012) and most image processing programs, whilst CSV files of coordinates are more convenient input for tracking software or for analysis (**Chapter 6**).

4.3.2.2. ROI TOOL

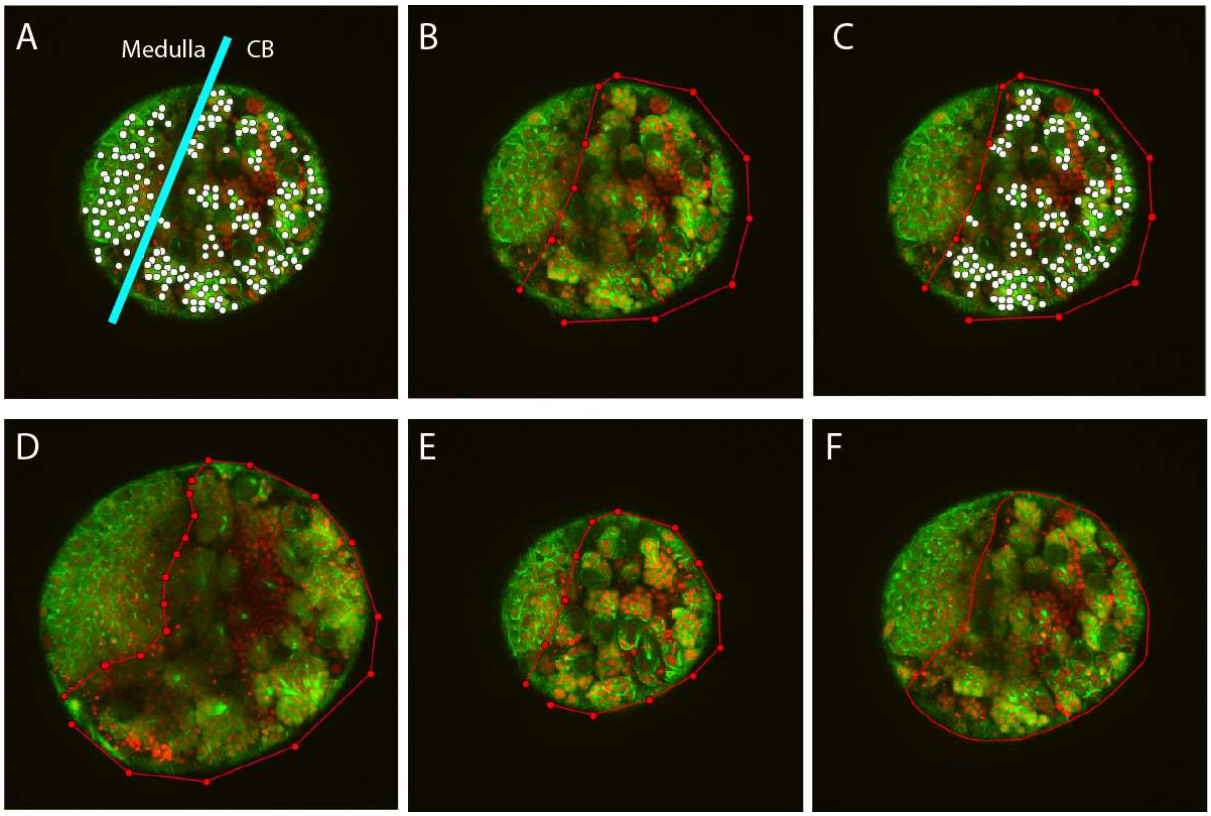

Figure 4.5: ROI tool. A) Drosophila brain labelled with Jupiter::GFP and Histone::RFP showing predicted neurons in the central brain (CB) and medulla regions. B) ROI defining CB region C) Predicted GMC and neurons in the brain restricted by ROI. D and E) Regions defining top and bottom of the CB region in these images. F) ROI tool interpolation over time and 3D.

Models do not take account of possible regional variation in the identity of cells that might otherwise be considered similar. A good illustration of this is the medulla and central brain regions of the *Drosophila* larval brain (**Figure 4.5**). Both contain GMCs and neurons, and there are no obvious distinguishing characteristics other than their location (**Figure 4.5A,B**). In this context I developed a 4D ROI tool based on D. Waithe's 2/3D version. The tool allows a user to annotate a series of points that define a region within an image (**Figure 4.5B**). From these annotations, a contour is defined by interpolation (**Figure 4.5B**), and cell centres are restricted to only this region (**Figure 4.5**). However annotating every slice would be time-consuming. In order to minimise the need for annotation, we instead linearly interpolate between slices

(**Figure 5D-F**). This means that, for simple shapes only the top and bottom in 3D need be annotated, drastically reducing the annotation burden. More complex shapes require more regions to be defined. In addition to 3D interpolation, these regions can be interpolated over time, meaning minimal annotation can be used to account for drift or growth of regions over time (**Figure 5F**). These regions can then be saved for subsequent use with other models. This ROI tool represents a useful aid to restrict cellular identity based on spatial location.

The regions found by the ROI tool could, in principle, be learned, rather than rigidly defined, perhaps using a segmentation based approach. This would allow further automation of the process of defining regions further. Depending on the implementation, this could use the pattern of detections, and/or the same image features used for calculation of cell centres.

4.3.2.3. DISPLAY

CytoCensus is based on PyQt4, using the matplotlib interface for displaying images and plotting. The matplotlib interface is reportedly one of the slower interfaces for display of images, however it has significant advantages in terms of the ease of annotation, and navigation within images compared to Qt native images, for instance. In order to increase the speed at which images can be browsed, I used the modest-image library (which is based on matplotlib), which speeds loading of large images, as it limits rendered image quality to the screen resolution. This, combined with optimisation of the way images are supplied to the plot, significantly increased the speed of display.

4.3.3. IMPROVING ON CYTOCENSUS

I have briefly described use of and development of the user interface of CytoCensus. In addition to these user-facing aspects, I also sought optimise a number of aspects of the density learning approach: the features calculated, the machine learning method, and the available parameters. Throughout this optimisation process I aimed to balance performance with

usability. I have considered the speed of the approach in training under different conditions, the complexity of available settings, and the choice of machine learning algorithm.

4.3.3.1. FEATURES

In this section, I aim to reduce the time taken to calculate features, and remove the need for the user to set a scale parameter for the features. The original Fiashi and Waithe papers on which this approach is based make use of a standard set of features at 3 different scales: GGM, LoG, STE features (**Figure 4.6**) (Fiaschi et al., 2012; Waithe et al., 2016). These features are calculated by applying a series of different filters and mathematical operations to the image, and, in general, have the advantage of being generally fast to calculate for small object sizes. Generally, the GGM can be thought of detecting edges within the image. The LoG detects blobs or foci of a particular size, and the STE eigenvalues correspond to 'texture' or local shape. Together, these features describe many of the important features in images (Gonzalez and Woods, 2002).

Of course there are many other possible features, and various feature sets that describe different aspects of images are available. SIFT is a highly popular feature detector, that calculates histograms of local orientation and magnitude (Lowe, 1999) but it is not designed to be calculated as a dense feature set (that is at every point in an image), and is therefore very time-consuming to calculate. The DAISY feature set is similar to SIFT (Lowe, 1999; Tola et al., 2008), but is intended to be used as a dense feature set and has shown good results with machine learning . Other approaches, such as GLCM, LBP and other texture filters make use of a different method, rather than creating features that describe aspects of the image, they describe the distribution of the direction of local image gradients (Mohanaiah et al., 2013). This makes for a somewhat complex feature, but various papers have shown good performance with these features (Chen and Agu, 2015). I compare the performance of these features in terms of time to calculate, in **Table 4.1**. Feature calculation is typically the most time

consuming part of running CytoCensus, and limits the speed at which the user can get feedback on their training, so speed of feature calculation is a key factor in ease of use of the program. Due to the length of time taken to calculate many of these features, I selected the set of GGM, LoG, STE features for use in CytoCensus, although I note that DAISY and LBP features are relatively quick to calculate, and perform well by comparison.

Table 4.1: Speed of feature calculation based on a 512 x 512 x 3 single channel image, showing best of 3. Computer specification details can be found in **Section 2.4.4**.

	GLCM	**DAISY**	**LBP**	**Fiashi (GGM, LoG, STE)**	**Pyramid Fiashi**
Time (s)	146.1	14.1	3.04	**2.06**	**0.844**

Given that feature set (GGM, LoG, STE) performed best, I reduced the time it takes to calculate these features. Using only small scale features when large objects are of interest is likely to lead to poor predictive performance because the features have little relation to the proximity map. However, calculating many features, particularly at large scales can be very time consuming. Therefore I used a pyramid based structure to reduce the computation time at a range of scales (**Figure 4.6A**). These image pyramids are increasingly common in image analysis (Gonzalez and Woods, 2002). In an image pyramid, a series of images are created at different resolutions, with each image smaller than the last (**Figure 4.6A**). This relies on the fact that fine details have little effect on large scale features, and can be safely lost. They have the advantage of minimising the computation time and memory requirements of feature calculation even when calculating features at large scales. To gain the most from a pyramid structure, features would be both calculated at different resolutions and stored in memory at these reduced resolutions. However, the prediction is calculated from the features on a pixel by pixel basis, so although the features are calculated in a pyramid they are subsequently upsampled to the original image resolution. Nonetheless calculating the features in an image pyramid represents a significant decrease in the time taken to calculate features. I therefore adapt this feature to cover a wider

range of scales by repeatedly downsampling the image 5 times, calculating features at each level (**Figure 4.6**). This wider range of scales is sufficient for most cells or objects of interest and is calculated in a reasonable time, without the need for a user to set a scale for the feature calculation.

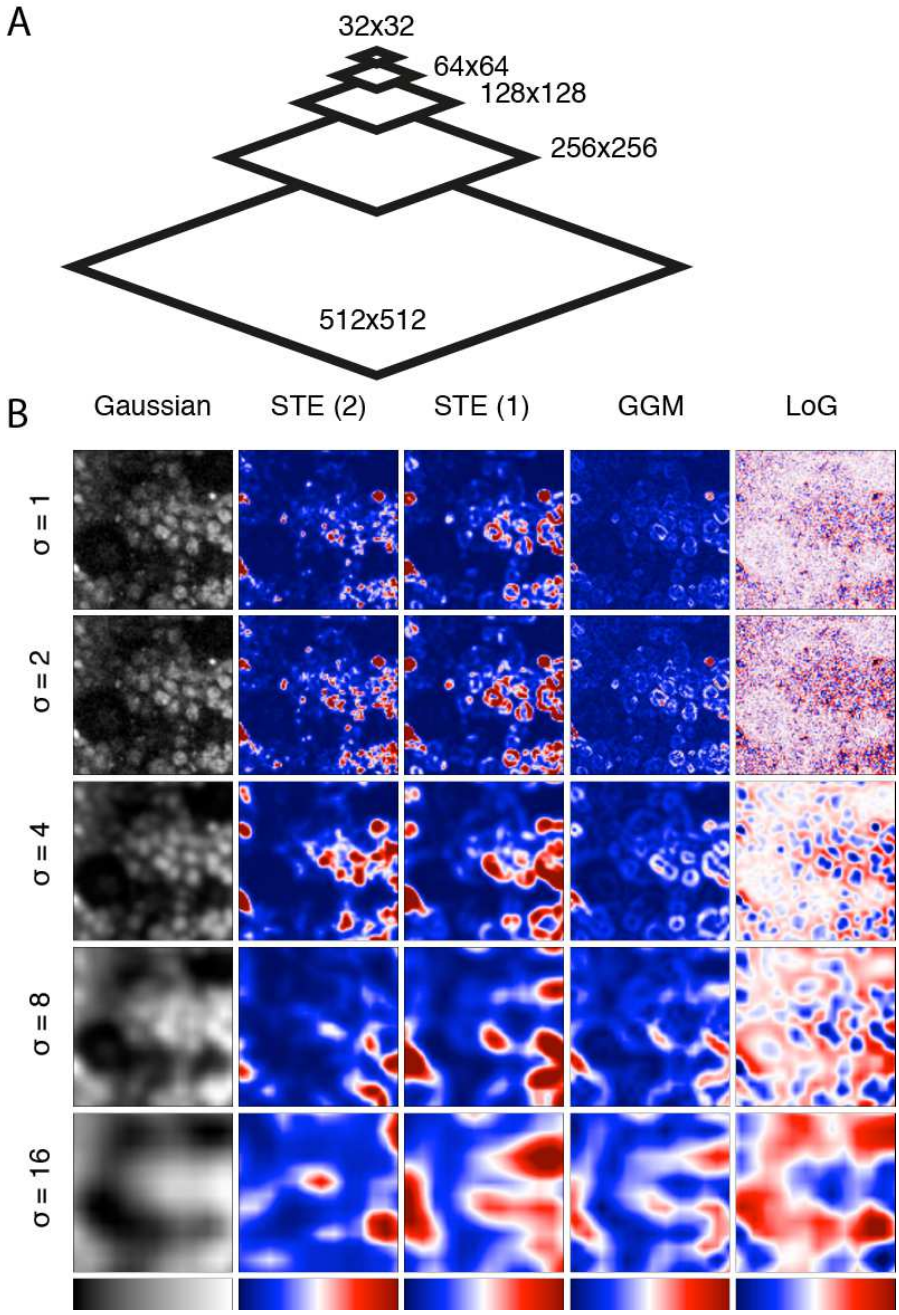

Figure 4.6: Features used in CytoCensus. A) Features are calculated in an image pyramid with 5 levels, shown here for a 512x512 image B) Five different features are calculated. First, a Gaussian smoothed version of the image is calculated at different scales (σ). Second, shape and texture features, the eigenvalues of the Structure Tensor are calculated at each scale: First Structure Tensor Eigenvalue (STE(1)), Second Structure Tensor Eigenvalue(STE(2)). Third, the edge detecting feature, the Gaussian Gradient Magnitude (GGM). Finally the 'blob-detecting' Laplacian of Gaussian (LoG) is calculated.

4.3.3.2. EXPLORING AND OPTIMISING IMPORTANT PARAMETERS

A key aspect of this density counting approach is that good performance can be achieved using relatively little training data. This justifies the use of an ensemble of decision trees, over neural networks (NN) for instance. In order to assess the impact of training on performance, I compared the performance of CytoCensus with different levels of training. In order to prevent bias, I generated an annotated dataset that contained a series of regions with individual NB. I used this dataset to automatically vary the number of NB containing regions used in training, creating an unbiased estimate of performance with a number of annotated NB. Although this is not entirely representative of normal training (which relies on user-feedback to improve training), it is sufficient to evaluate the performance of training. I assessed the performance of the algorithm using the F1 score, which weights false positive and false negatives into a single score (**Figure 4.7**)(Davis and Goadrich, 2006). It is clear from **Figure 4.7** that performance begins to plateau after 13 examples, with only small improvements after 17 NB.

In addition, I explored other parameters, that affect algorithm performance using this method. I find that the threshold value and sigma, as expected, are key that significantly affects performance, with a clear optimum for sigma corresponding to the object size (**Figure 4.7)**, indicating it is a key parameter and which need to be user-modified. In particular, this approach allowed us to assess the level of training required to achieve good detection of cell types of interest with different datasets (**Figure 3B**). In all cases tested, we found that successful identification of NBs or progeny required minimal user training (of the order of tens of examples on only a few image planes) and increasing training gave only marginal improvement (**Figure 4.7A,B**).

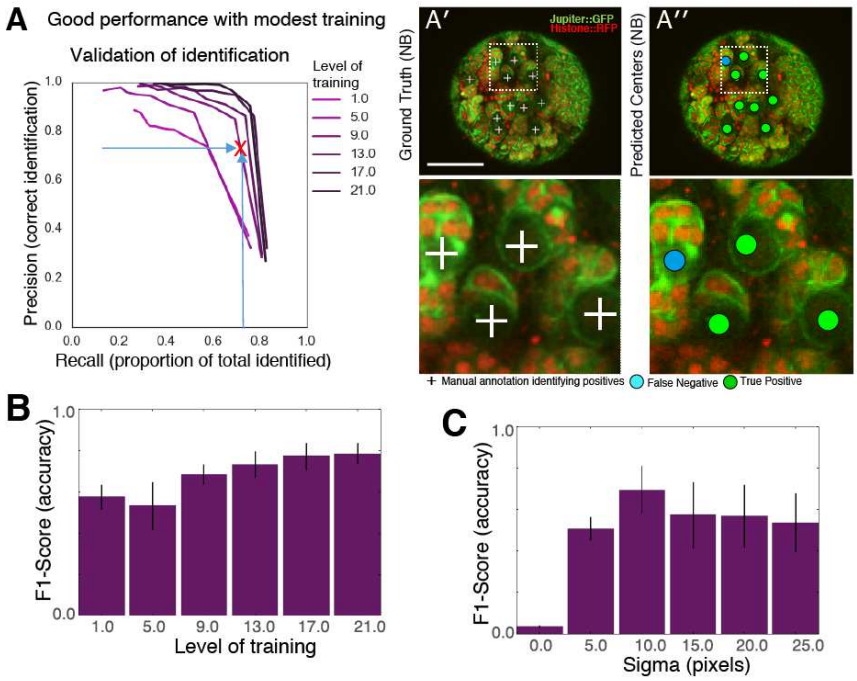

Figure 4.7: Effect of training and parameters on performance A) Precision recall graph reflecting degree of training (shades of magenta) A'-A") Corresponding identifications of NB, including False Negatives (Blue) and True Positives (Green) B) Summary of the number of NB trained on and performance (F1-score) C) Summary of the effect of varying sigma and performance (F1-score)

4.3.4. Detecting Cells: the Determinant of the Hessian Matrix

As discussed in **Section 4.2**, cell centres are detected by filtering the predicted proximity map. The method by which these centres are detected significantly impacts the quality of the detection. Achieving high quality detections is crucial to the program being useful. In particular, the way the Determinant of Hessian matrix (DetH) is implemented turns out to have a significant effect on the quality of the result. In a 1-D DetH, all points that are maxima correspond to real maxima in the image. However, in higher dimensions, not all points correspond to real maxima. In 2D, quadratic maxima can also correspond to minima (**Figure 4.8C**),and in 3D can be saddle points or minima. These minima may arise in the proximity map from image features which are anti-correlated with the object of interest, or from the shape formed from multiple interacting maxima. These must be filtered out

subsequently by assessing whether or not these points also appear as lower-dimensional maxima (i.e. in 3D, a point must also be a maximum in the 2D and 1D detH). The suppression of false maxima using the 1D Hessian (in **Figure 4.8**D), is illustrated in 2D, for simplicity, in **Figure 4.8**E. In addition to this suppression of maxima, CytoCensus also applies a local maximum filter that suppresses nearby maxima. It is crucial this step occurs before the calculation of local maxima, otherwise the local maximum filter may suppress true maxima in favour of minima or saddle points. Together these steps result in robust centre detections from the proximity map.

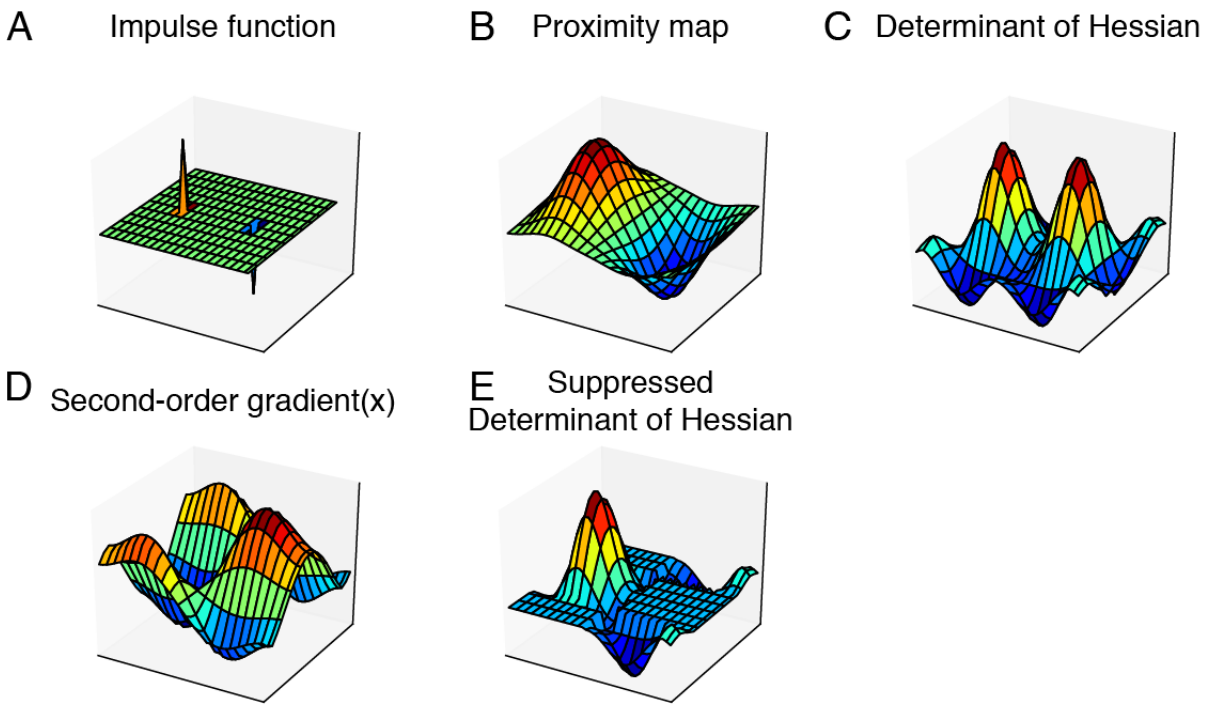

Figure 4.8: 2D Determinant of Hessian also detects minima. Minima are suppressed using the second order gradient. A) 2D Surface plot showing a synthetic starting function with a single maximum and single minimum (see Section 4.4); B) Gaussian blurred impulse representing a predicted proximity map; C) Determinant of Hessian of (B), which detects both maxima and minima; D) Second order gradient is positive only at minima; E) Determinant of Hessian with non-maxima suppressed.

4.4. Validating CytoCensus

4.4.1. Motivation

In the previous sections, I have described the method CytoCensus relies on, and the key details of its implementation for ease of use. Previously, I have demonstrated that CytoCensus can identify NB in test images. In this section, however, I will more rigorously examine and validate the performance of both the CytoCensus and the underlying algorithm in a range of examples. I compare CytoCensus to the state of the art and demonstrate excellent performance.

4.4.1.1. Neutral Challenge Dataset

To assess the performance of CytoCensus, I first evaluated its output for both an artificial 3D dataset with a known "ground truth" (**Figure 4.9**, see Introduction 1.3). A neutral challenge dataset of highly clustered synthetic cells in 3D with a low SNR was obtained from Broad Bioimage Benchmark Collection **(BBB)** (Svoboda et al., 2011; Svoboda et al., 2009) , which allowed me to perform validation of algorithm performance (**Figure 4.9**). Using this dataset I compared the abilities of Ilastik and CytoCensus to identify cell centres in 3D.

Appropriate quality 3D benchmark datasets for validation are not widely available. In the case of CytoCensus we are not performing segmentation, but nonetheless require fully segmented 3D datasets to fairly compare to segmentation based approaches. Additionally, acquiring sufficient data for robust analysis is not straightforward. Synthetic datasets such as the dataset described above have the advantage that they can in principle be arbitrarily large. However, there can be questions over how accurately they represent real life datasets. In my experience, real-life datasets can be significantly more varied, for instance in terms of lighting conditions and inhomogeneity of imaging fields, and hence represent more of a best-case imaging dataset. Nevertheless, synthetic datasets have their place for comparing the performance of different algorithms.

I applied CytoCensus and Ilastik to the BBB data, in order to compare performance. I find that CytoCensus (F1-score, 0.98±0.05) outperforms Ilastik (F1-score, 0.21±0.13) in the accuracy of cell centre detection (**Figure 4.9A**). The Ilastik result, in particular suffered from merging of objects due to their high density. Therefore, I also processed the Ilastik result with a simple post-processing step for separating objects: binary erosion (**Figure 4.9B**). I find that, using this step, the Ilastik result is significantly improved, although it remains significantly below the CytoCensus result (F1-score, 0.88±0.09). I expect that the Ilastik segmentation could be processed to yield better results. However, the aim is to produce a program that is simple to use without dedicated programming or image analysis knowledge. Therefore, I conclude that CytoCensus is significantly more accurate at identifying cells with this remit.

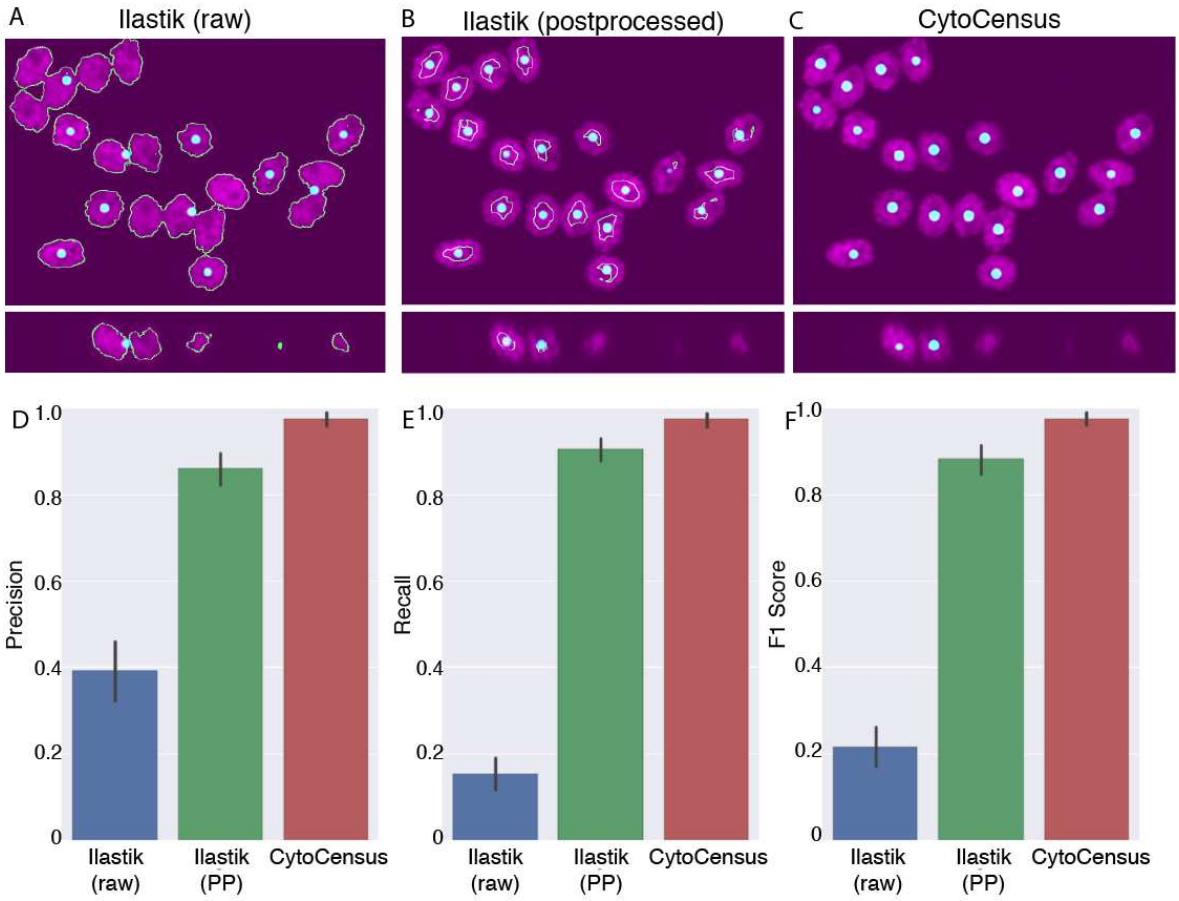

Figure 4.9: Comparison on neutral challenge dataset (30 3D images with 20 cells from Broad Bioimage Benchmark Collection BBBC024v1). A-C) Magenta: Centre plane in XY and ZY for benchmark image. A) Ilastik result showing segmentation outlines (cyan lines) and cell centre detections (cyan circles). Note the merging of touching cells. B) Ilastik result with binary erosion post-processing to aid separation of nearby objects, and corresponding centres (cyan circles). C) CytoCensus cell centre detections (cyan circles). Size of circles indicates depth. D) Precision (fraction of correct identifications) of the three approaches. E) Recall (fraction of true centres that detected) of the three approaches. F) F1-score (weighted average of precision and recall, n=30).

4.4.2. IDENTIFYING NEUROBLASTS IN FIXED IMAGING

Although identifying cells in artificial datasets is useful, they do not accurately represent the complexity of a live tissue. Therefore, I assessed the performance of CytoCensus on real data sets that were manually annotated by a user to generate ground truth results. In particular, I was interested to confirm whether CytoCensus can successfully detect NB and GMCs (See

Introduction 1.1, **Chapter 3**). I therefore established and validated CytoCensus could accurately identify NB's and GMC's on 3D images of fixed whole mounted (wL3) larval brains immuno-labelled against Ase and Dpn, and with DAPI (**Figure 4.10A-D**). The combination of markers allowed definitive identification of NB and GMC locations and numbers (Boone and Doe, 2008; Neumuller et al., 2011)

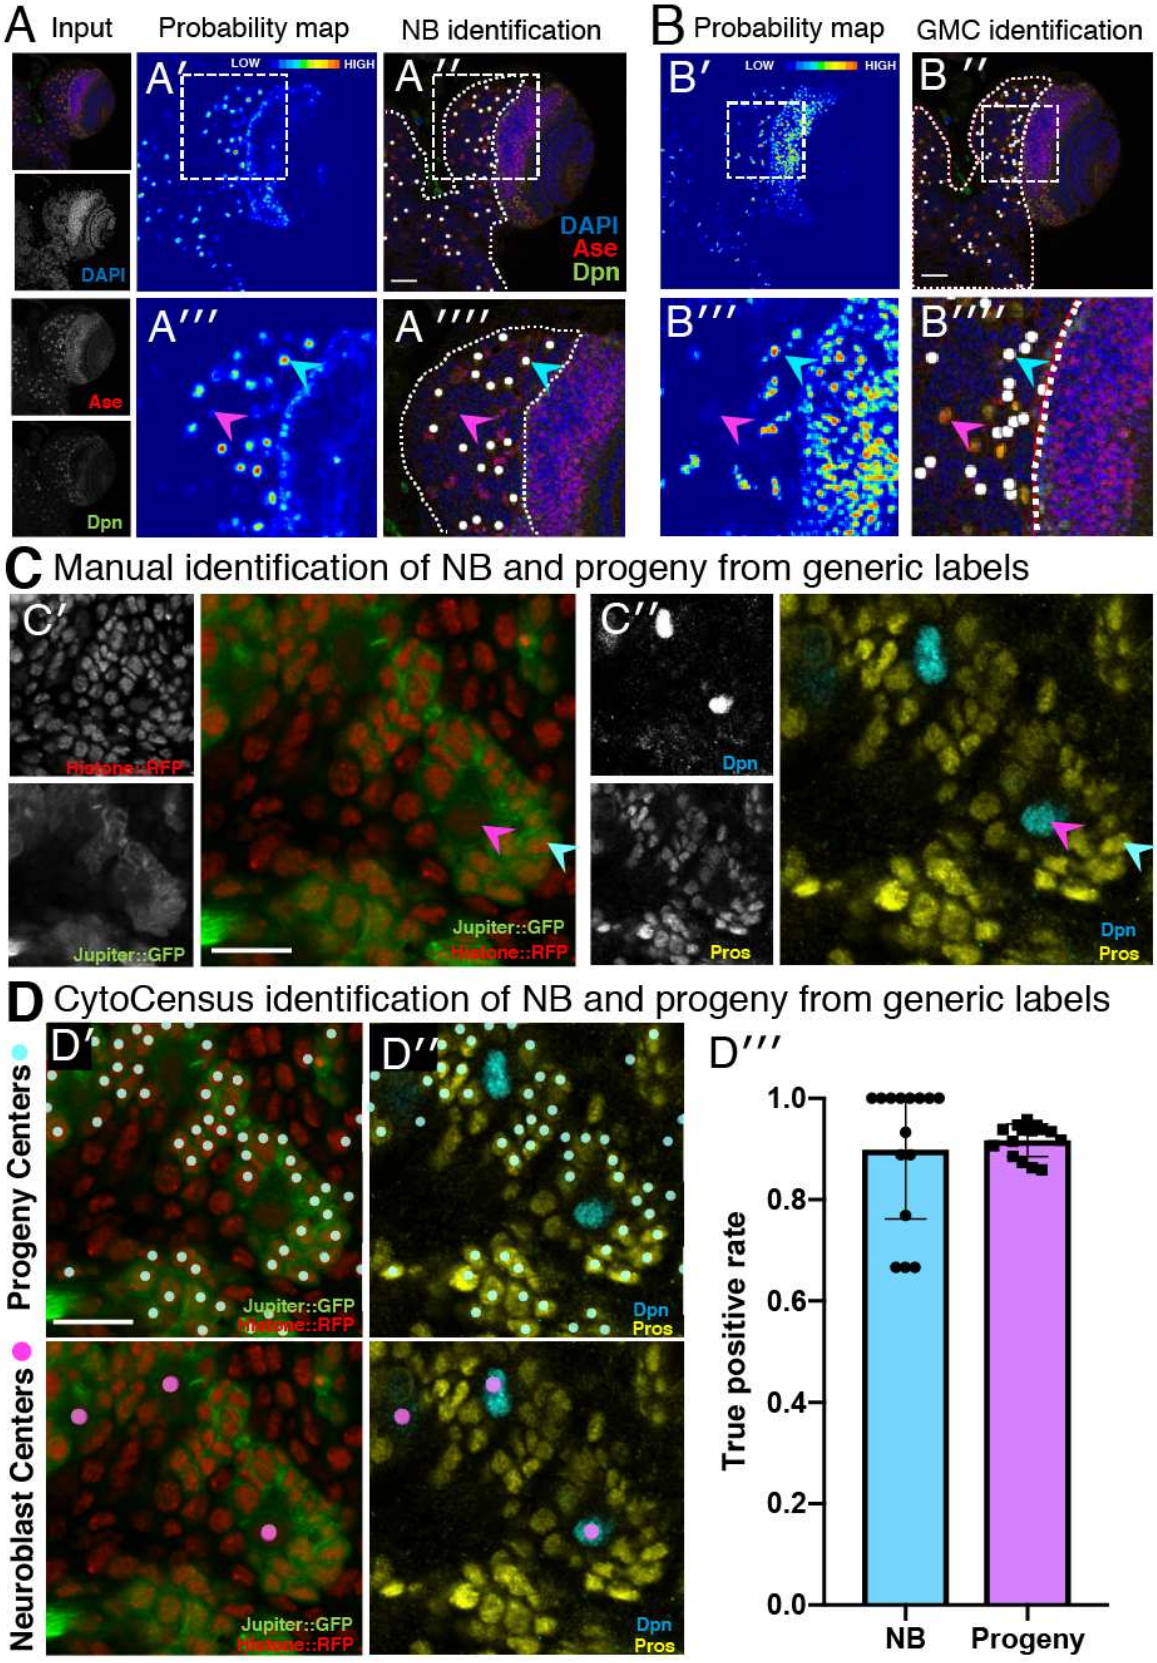
A Input
Probability map
NB identification
B Probability map
GMC identification
LOW
HIGH
DAPI
Ase
Dpn
C Manual identification of NB and progeny from generic labels
Histone::RFP
Jupiter::GFP
Pros
D CytoCensus identification of NB and progeny from generic labels
Progeny Centers
Neuroblast Centers
D'''
True positive rate
1.0
0.8
0.6
0.4
0.2
0.0
NB
Progeny

Figure 4.10: CytoCensus performance on fixed immunohistochemistry images. A) fixed Jupiter::GFP, Histone::RFP labelled brain (A′), immuno-labelled for Dpn and Ase (A′′) to permit unequivocal identification of NB and their progeny. B) Cell centre predictions are shown from CytoCensus analysis of the dataset from (A′) with generic markers. C) Manual identifications of NB (magenta arrow) and progeny (cyan arrow) based upon Dpn and Pros markers in 4 colour images shows they are recognisable in Jupiter::GFP, Histone::RFP based on size and amount of Jupiter. D) Images and plot showing cell centre identifications based upon Jupiter::GFP, Histone::RFP labelling alone effectively identifies NB and progeny compared to identification from Dpn and Pros labels: 90% ± 14 NB identification (n=15, with 2-3 NB/image) and 92% ± 3 progeny identification (n=15, with 60-70 progeny/image). Scale bars 20 μm.

I extended this analysis to determine how well CytoCensus performed cellular identification in the more challenging case of generic cytological markers, such as those found in my *Drosophila* live imaging datasets. To assess whether CytoCensus can effectively identify the NBs and their progeny from generic markers, L. Yang fixed Jupiter::GFP / Histone::RFP expressing brains and labelled them for Dpn and Pros to identify NBs and their progeny (**Figure 4.10C**). These markers allow unambiguous identification of NBs and their progeny but are not specific to GMCs or neurons unlike the previous set of markers. After training, based only on the Jupiter::GFP / Histone::RFP labels, CytoCensus successfully and accurately identified both NBs (90% ± 14 Dpn positive, n = 15) and progeny (n=15, with 60-70 progeny/image) (**Figure 4.10-D**). I conclude that CytoCensus can be used effectively to identify NBs and progeny from either specific cell type markers or from generic markers of cell components that are expressed in the brain.

4.4.3. Finding Neuroblasts in live imaging data

Finally, I assessed the applicability of CytoCensus on actual live imaging data of the *Drosophila* brain. I used a subset of live-imaging time-series of the generic cytological markers Jupiter::GFP / Histone::RFP as a test dataset to compare CytoCensus to other freely available programs (**Table 4.1**). I annotated the number and location of NBs identified in five time-

points from a movie sequence to create a ground truth (**Figure 4.11A**). I then trained and evaluated CytoCensus, and compared its performance to that of the ground truth (**Figure 4.11B**). I repeated this approach for Ilastik and FIJI-Weka, using their respective training interfaces. In each case I attempted to optimise the parameters used based, whenever possible, on the published information. In my tests I found that CytoCensus significantly outperformed these other freely available approaches in F1-score, speed and simplicity of use (**Figure 4.11, Table 4.2**). I conclude that CytoCensus represents a significant advance over the other current freely available state-of-the-art methods of analysis, especially in its ability to accurately and automatically analyse the large volumes of data that result from live imaging of an intact complex tissue such as a brain.

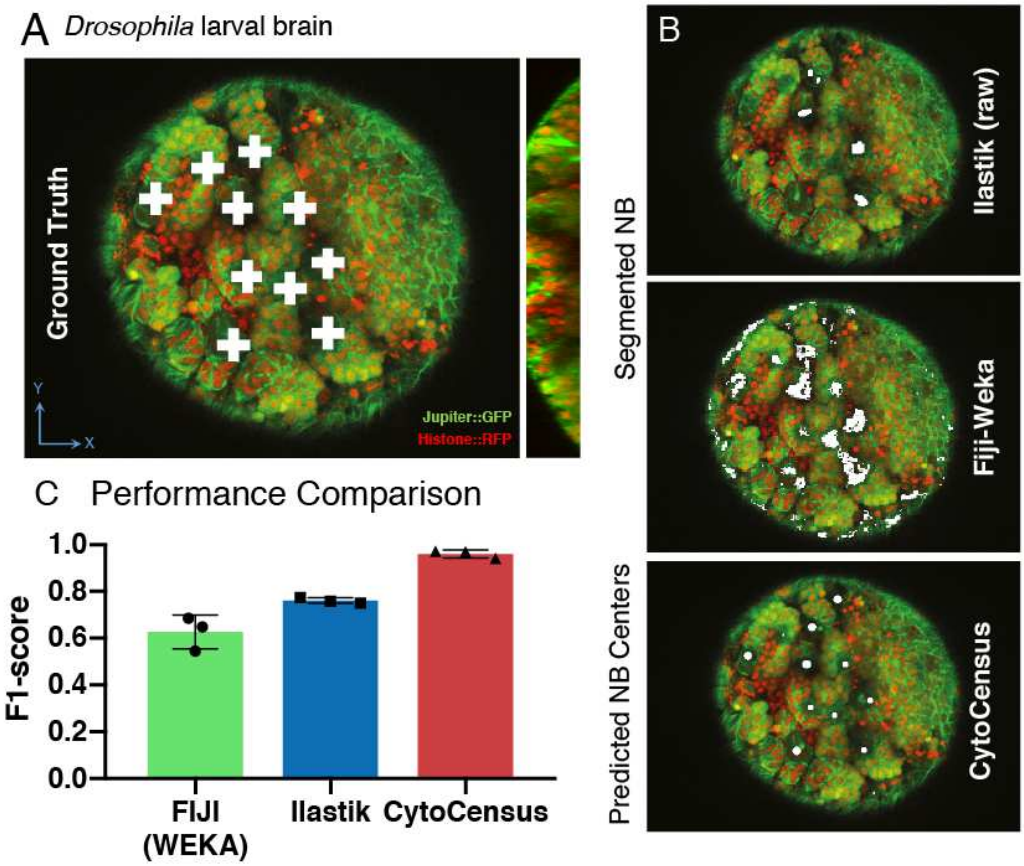

Figure 4.11 Performance of CytoCensus, WEKA, and Ilastik on live imaging data,. A) Ground truth of NB Cell Centres in the central brain B) Single Z slice showing output of of Ilastik, WEKA (after processing) and CytoCensus (white areas = detections). C) Performance comparison plot of Ilastik and CytoCensus, showing F1-score (n=3).

Table 4.2: CytoCensus outperforms other freely available programs for cell detection

	Manual	**Fiji/auto local threshold**	**Ilastik (1.17)**	**Fiji/Weka**	**CytoCensus V0.1**
Total Parameters to select	-	1	67	25	6
Handles 3-D easily	-	YES	YES	YES	YES
Handles 4-D easily	-	NO	YES	NO	YES
Time to Train model (min.)	-	N/A	15	18	6
Post processing required?	-	YES	YES	YES	INCLUDED
Time to Run (min. including postprocessing)	550 (equivalent)	5	70	105	19
% F1-score	-	FAIL	62.7±7.3	76.1±1.2	96.1±1.7

4.5. Applying CytoCensus to large data

The approach taken with CytoCensus is intended to cope with large imaging data. The calculation of 2D features reduces the computational cost compared with 3D features, at a potential cost in its ability to recognise complex 3D shapes. It is also convenient given that images are normally stored as a series of 2D image planes, although modern formats (.hdf5, ome.tif to an extent) allow for more complex structures and access to data in a relatively efficient manner.

Nonetheless, there are limitations in its practical implementation, particularly in terms of memory requirements. For instance, 26 features are created per channel, all of which are stored in memory. In 'evaluation' mode, features are calculated for each stack and evaluated in a multi-threaded fashion. This therefore requires 26x the size of a z stack to be stored in memory at a given time. In addition, the resulting prediction maps are also stored in memory. This generally causes few problems for multiple time-point images, where each Z-stack comfortably fits into memory, but causes issues when images are large in x,y or z.

However, it has become apparent that locating objects in large datasets in x and y, for example in whole slide datasets is desirable. In this section, I document some optimisation for the application to datasets of this size, as well as more clearly the limitations of CytoCensus as it is currently implemented.

First, and most straightforwardly, it is possible to save prediction maps to disk directly rather than keeping them in memory, as well as the hessian maps of local maxima. This was done using the 'tifffile' library (Christoph Gohlke, Laboratory for Fluorescence Dynamics, University of California), appending additional frames to the image each time. However, this adds a complication: namely that maps can no longer be calculated in a multi-threaded fashion without first establishing locks on file writing, and additionally, breaks the desirable, but not necessary, guarantee that images are written sequentially (which allows images to be sensibly opened in other programs, for instance, ImageJ). Removing the use of multi-threading, however, does significantly reduce the memory requirement, which in these circumstances is desirable.

Second, given the removal of multithreading of feature calculation, prediction can be made on a plane-by-plane basis. This reduces the need to store the entire Z-stack of features, effectively making the memory requirement limited to x*y*f rather than x*y*z*f. The reduced memory requirements of this stage now make 10000*10000 pixel image planes straightforward to work with 16GB RAM. However, the hessian calculation stage requires the calculation of gradients in multiple dimensions, including z. Whilst theoretically possible to calculate these gradients without storing multiple image planes in memory at once, it is most straightforward to do so. Instead, the datasets is split into chunks in xy of size 256x256, which can be simultaneously processed. This does not remove the requirement to store the detH, briefly, in memory, but does reduce the memory load by removing the requirement to have all the gradient images held in memory simultaneously. Thus, it, approximately limits memory

usage to x*y*z (although this may be transiently greater as the multiple threads process the images) whilst simultaneously speeding up calculation of the hessian images.

In summary, these optimisations make the batch evaluation of large images using CytoCensus practical on standard images.

4.6. Discussion

In this Chapter, I have outlined the development of a method, based on density counting, for quickly annotating 4D microscopy images. I have discussed the development of this method into the program CytoCensus with the aim of producing a user-friendly program. I have specifically detailed choices and optimisations in its development to reflect these values. I later extended the program, showing that the approach can be extended to cope with large datasets with a minimal memory footprint.

I then validated CytoCensus in a series of different tasks. In the first case, I showed that CytoCensus can achieve excellent performance on 3D challenge data from the Broad Collection, achieving better performance compared with the segmentation based approach of Ilastik (Sommer et al., 2011). I then applied CytoCensus to a variety of additional imaging datasets of complex tissues. I first showed that it can correctly identify NB with an ideal marker combination, and subsequently that it can identify with good accuracy NB and progeny GMCs and neurons on the Jupiter::GFP Histone::RFP generic marker combination. Finally I applied CytoCensus to a live imaging dataset from **Chapter 3**, comparing its performance on this dataset to Ilastik and FIJI-WEKA, achieving a significantly better F1-score (**Table 4.2**) and significantly reducing the time taken to annotate a dataset. This shows that CytoCensus can successfully identify cell types that are important in neurogenesis. I will expand on this idea in **Chapter 5**.

CHAPTER 5

STUDYING BRAIN DEVELOPMENT WITH LIVE IMAGING AND CYTOCENSUS

5.1. INTRODUCTION

In the previous two chapters I have detailed the development of culture methods and imaging to support analysis of cellular behaviour during brain development. In collaboration with D. Waithe I have developed and validated a supervised machine learning approach to analysing extensive 3D time-lapse datasets. In this chapter I will illustrate how CytoCensus can be applied to these datasets, and show that how analysing live imaging data with CytoCensus can be used to answer questions about brain development. Specifically, I use this approach to characterise the phenotypic effects of mutation in the key regulator of development, Syncrip, which turns out to be surprisingly complex. Studying such mutants can give insights into the molecular mechanisms that control stem cell behaviour and development.

To illustrate the use of CytoCensus in studying brain development, I analysed *Drosophila* brains lacking Syncrip expression. Syncrip[286] had been previously shown in the Davis lab to result in an enlarged brain phenotype by L. Yang and T. Dubrovski, contrary to the results of a previous RNAi screen of genes (Neumueller et al., 2007). Syp was initially identified by Neumüller *et al.* (2007) and was found to be a weak undergrowth phenotype (that is, smaller brains), and a small reduction in the number of GMCs (Neumuller et al., 2011). I will return to this later in this chapter. In this section, I characterise the brain overgrowth phenotype of the

central lobe of Syncrip in more detail than has previously been investigated. SYNCRIP/hnRNPQ, the mammalian homologue of Syp, has interesting roles in mammalian hippocampal cells. Syp also determines neuronal fate in the Drosophila brain (Liu et al., 2015; Ren et al., 2017) and is required for neuromuscular junction development and function (Halstead et al., 2014; McDermott et al., 2014)

I investigate possible causes of the Syp phenotype: I consider cell fate changes (reversion to stem cell fate), changes in cell behaviour, and, briefly, changes in whole organism development. I will primarily focus on understanding the Syp phenotype using CytoCensus, but in some cases will use other methods as appropriate. This Chapter builds on work on CytoCensus that was published in pre-print (Hailstone et al., 2018).

5.1.1. Specific aims

Determine the cause of the *syp* overgrowth phenotype

Demonstrate the use of CytoCensus in analysing live imaging data

Develop appropriate tracking and analysis to measure the GMC cell cycle

5.2. The role of Syncrip in the Central Brain

5.2.1. Motivation

Previous work by L. Yang, R. M. Parton and T. Dubrovski showed that Syncrip -/- brains were larger at wL3 compared to WT. They compared the length of the VNC to the width of the brain (in order to normalise for development), and found a significant increase in the width of the brain at wL3. Here I focussed my investigations on the CB region. The CB region has attracted significant interest because it is more heterogenous compared with other parts of the brain and contains more complex lineages (which are closer to that of mammalian lineages, see Introduction, **Section 1.1**) than in other parts of the brain. Specifically, wL3 central brain contains two main types of neuroblast: type I and II, that produce the neurons of the central

complex, which is of particular interest due to similarities between the mushroom body development and mammalian neurogenesis.

5.2.2. SYNCRIP -/- RESULTS IN AN ENLARGEMENT OF THE CENTRAL BRAIN

I first set out to determine whether the CB region has a key role in the enlarged brains of the Syp -/- phenotype. To characterise changes in the CB, I dissected early wL3 and late wL3 larval brains (See **Methods**), and imaged them using brightfield microscopy (**Figure 5.1A**). I find that in early wL3, the brains of *syp* loss-of-function mutants were significantly ($p<0.0001$, t-test, WT n=10, *syp* n=5) enlarged compared to WT larvae at the same stage of development (**Figure 5.1**). The brain lobes of *syp* larvae had 22% larger diameters compared with WT (WT 206.5 µm ± 5.0, n = 10, *syp* 253.7 µm ± 11.0, n = 5). This increase corresponds to an estimated 35% increase in estimated CB volume. By late wL3, the increase is even larger (24% increase in diameter) (**Figure 5.1C**). RNAi knockdown of Syncrip demonstrates a similar ($p=0.002$, t-test) although less pronounced increase in the diameter of the lobes (13% larger than WT; 234 µm diameter ± 17.0, n=12) (**Figure 5.1B**). This phenotype may be less extreme than that of Syncrip -/- due to residual Syncrip expression, or due to the restricted expression of the RNAi driver (in this case the Inscuteable promoter, which is expressed in NB, GMCs and neurons) (**Figure 5.1B**)

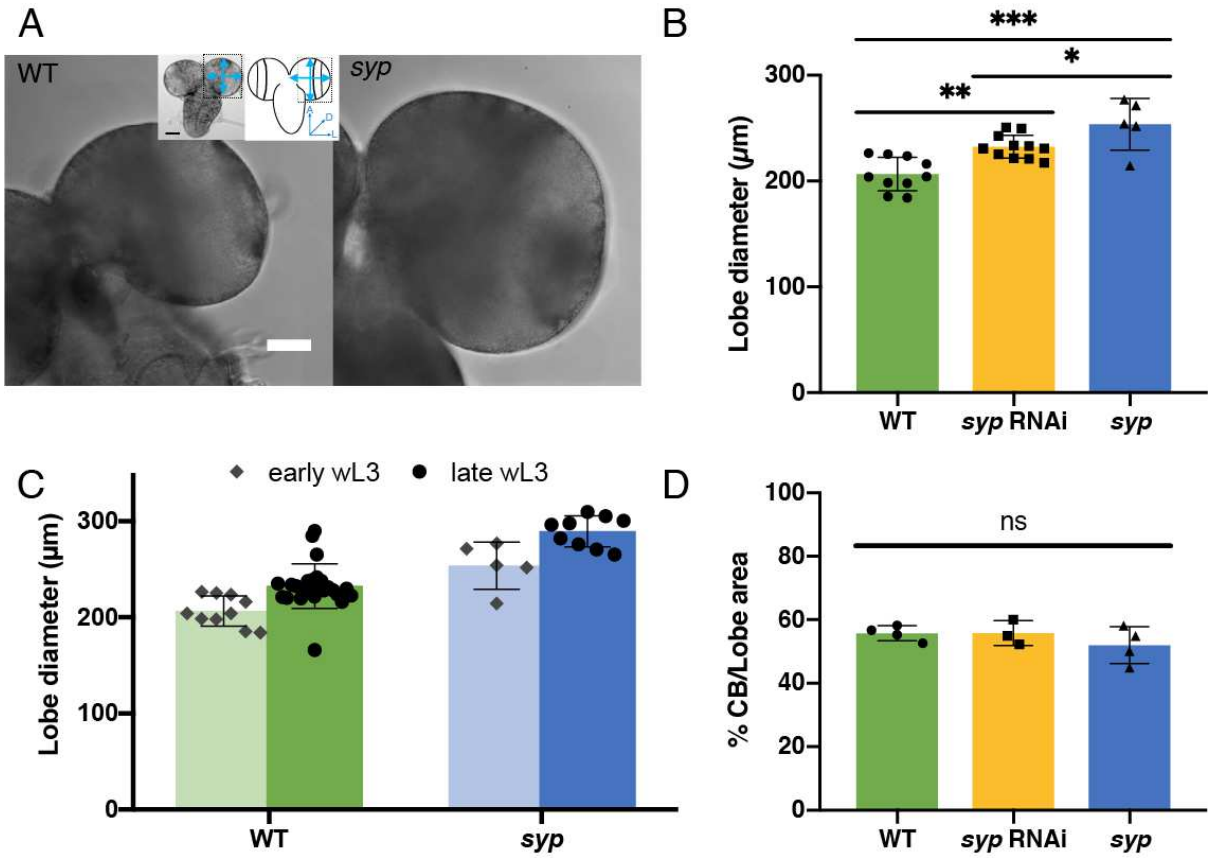

Figure 5.1 The CB region is enlarged in *syp*. A) Brightfield microscopy of WT and *syp* mutant brain lobes. Scale bar 50 µm (with thanks to R.M. Parton for assistance with imaging). B) Comparison of brain diameter shows significant differences ($p<0.0001$, one-way ANOVA) for WT (green, n=10), *syp* RNAi (yellow, n=11), and *syp* (blue, n=5). Each point is an average of 4 measurements/brain. ($p=<0.0001$ WT vs *syp*; $p=0.002$ WT vs *syp* RNAi; $p=0.04$ *syp* RNAi vs *syp*, t-test with Dunn-Sidak correction) C) Diameter (µm) of brain lobes in *syp*, and WT in early wL3 (lighter colour bars) and late wL3 (dark colour bars) C) The CB as a proportion of the whole brain lobe is the same (ns, one-way ANOVA) in *syp* (n=4), *syp* RNAi (n=3) and WT (n=4). (Error bars indicate standard deviation)

5.2.3. SYNCRIP AFFECTS NB NUMBER

Given this increase in the size of the CB, I proceeded to investigate possible causes. In *brat* mutants, brain overgrowth is caused by additional ectopic NBs. In *brat*, the progeny (INPs) of Type II NB revert to a NB-like identity, increasing the number of NB in the brain and leading to overgrowth (Bello et al., 2006; Betschinger et al., 2006; Lee et al., 2006). I tested whether additional ectopic NB were also found in *syp* RNAi brains (**Figure 5.2**). T. Samuels fixed and

stained wL3 brains (**Figure 5.2**). I used CytoCensus (**Chapter 3**) to accurately determine the total number of NBs in the CB of fixed *syp* mutant brains and compared them to the brains of WT wL3 larvae. My results showed that the number of NB in wL3 brains with *syp* RNAi knockdown is not significantly different from WT (**Figure 5.2**; WT 45.6 ± 1.3, n = 22, RNAi 44.1 ± 2.1, n = 15), however, *syp* mutant larvae do show a small but significant increase in the number of NB (p=0.004, t-test).

This increase in NB number is not sufficient to explain the observed phenotype; in the *syp* mutant the lobe volume is increased by 35% (**Section 5.22**). Since the production of neurons is linear with the number of NB, an increase of 10% in NB number does not adequately explain the magnitude of the increase in overall brain volume. Additionally, the *syp* RNAi brains exhibit larger brains but no significant change in the number of NBs. I conclude that the increase in brain size in the *syp* mutant is not primarily caused by this increase in NB number.

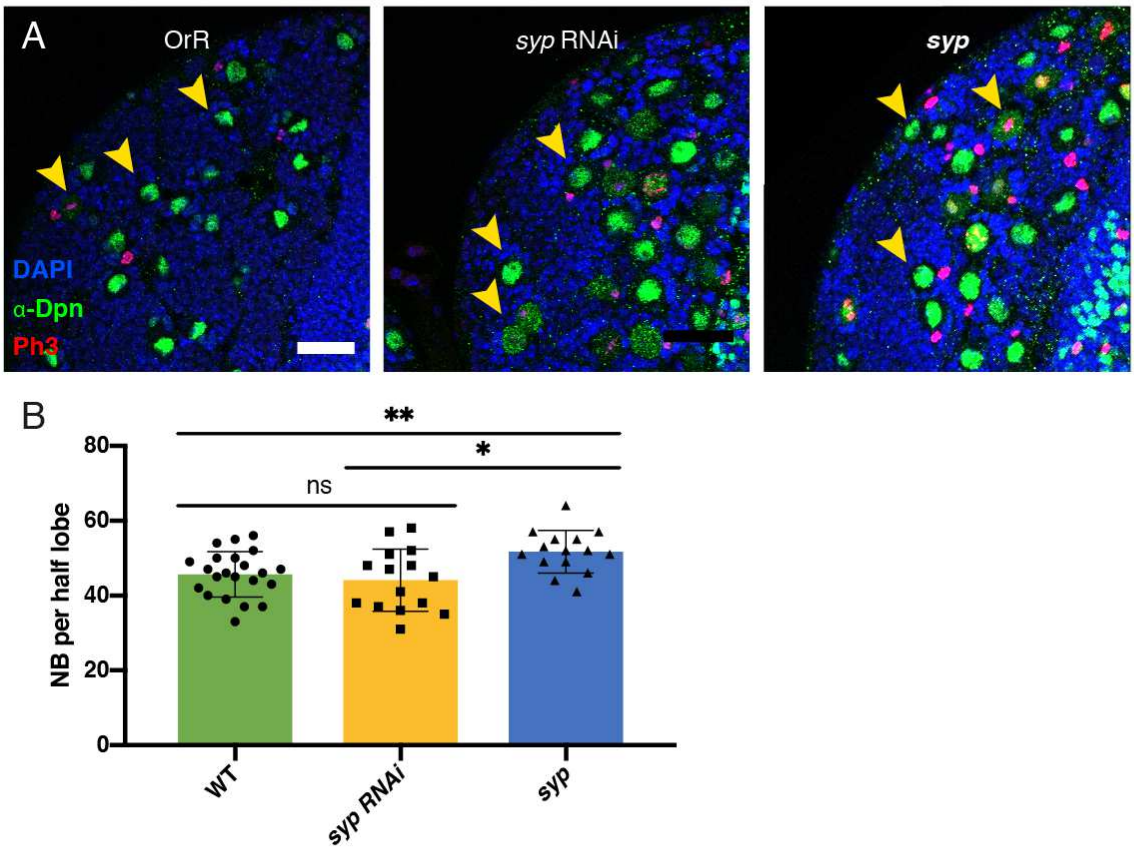

Figure 5.2: Type I NB number in the dorsal CB is increased in *syp* but not in *syp* RNAi. A) Brain lobe showing CB region stained with DAPI, Dpn, Ph3 (T. Samuels) to mark cells, NB (yellow arrows), and dividing cells respectively. Scale bar 50 µm. B) Plot of CytoCensus counts of NBs in the CB per brain lobe, trained on the images in (A), showing standard deviation. Significant differences are indicated ($p<0.0001$, one-way ANOVA, WT n=22, *syp* RNAi n=15, *syp* n=15). (ns *syp* RNAi vs WT; p=0.02 *syp* RNAi vs *syp;* p=0.009 *syp* vs WT)

Having identified ectopic Type I NBs in *syp*, I proceeded, in collaboration with T. Samuels, to investigate any difference in the number of Type II NBs. In a normal phenotype type II NB consist of 8 per brain lobe located medially on the dorsal posterior (**Figure 5.3**)(Bello et al., 2006; Boone and Doe, 2008; Bowman et al., 2008; Ren et al., 2017). We found no obvious change in the size or morphology of the Type II lineage in *syp* (**Figure 5.3**, T. Samuels). I conclude that, unlike in *brat*, there is no change in the Type II lineage that leads to ectopic NB. However, the source of the additional Type I NB in *syp* remains, and will be left to the discussion.

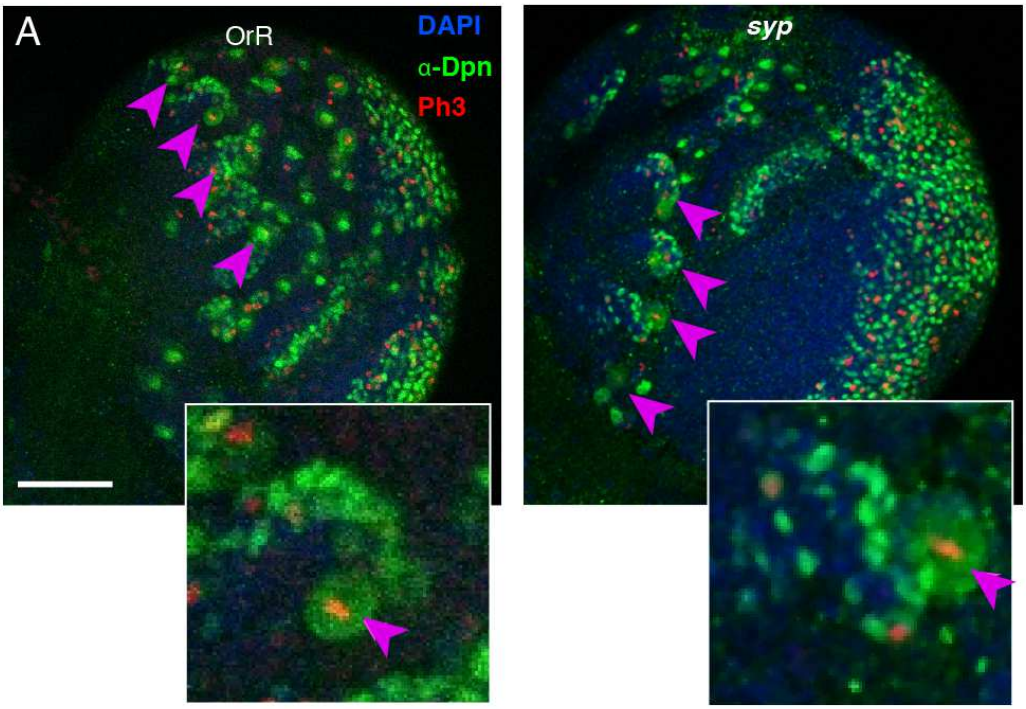

Figure 5.3: Type II NB in the CB region are unaffected by *syp*. A-B) The Type II Dpn-positive lineages are indicated with magenta arrows for WT (A) and syp mutant brains(B). Scale bar is 50 µm. Inset: Single Type II NB (magenta arrowhead) visible within the lineage. (Imaging T. Samuels).

Nevertheless, it is clear that the increase in NB number observed in *syp*, unlike the increase also seen in *brat,* does not fully account for the brain enlargement observed in *syp*/*syp* RNAi and therefore other causes must be contributing to brain overgrowth. I next investigated the role of cell proliferation and division rate in the Syncrip phenotype.

5.2.4. SYNCRIP IS INVOLVED IN CELL CYCLE REGULATION IN NB

An increase in the NB division rate during development could contribute significantly to brain overgrowth. To test this possibility, I examined the rate of NB division in cultured living brains using live imaging and culture (**Chapter 3**), analysing them with CytoCensus (**Chapter 4**).

5.2.4.1. CytoCensus can be used to follow the cell cycle

Although the primary aim of CytoCensus is to allow straightforward identification of cell centres the software can be utilised for a range of tasks. In this section I outline its use in monitoring the cell cycle, specifically following dividing NB. (I will consider the use of CytoCensus for other tasks further in **Chapter 6**). During the process of cell division, the appearance of NB marked by Jupiter::GFP, Histone::RFP changes dramatically (**Figure 5.4**). The chromosomes condense, leading to bright RFP spots corresponding to the mitotic chromosomes, which then segregate into the daughter and parent cell (**Figure 5.4**). Similarly, Jupiter::GFP associates with the mitotic spindle, forming an obvious spindle structure, in addition to marking the cell cortex (**Figure 5.4**). These changes sufficiently alter the appearance of NBs that it should be possible to train a model to specifically identify the dividing cells.

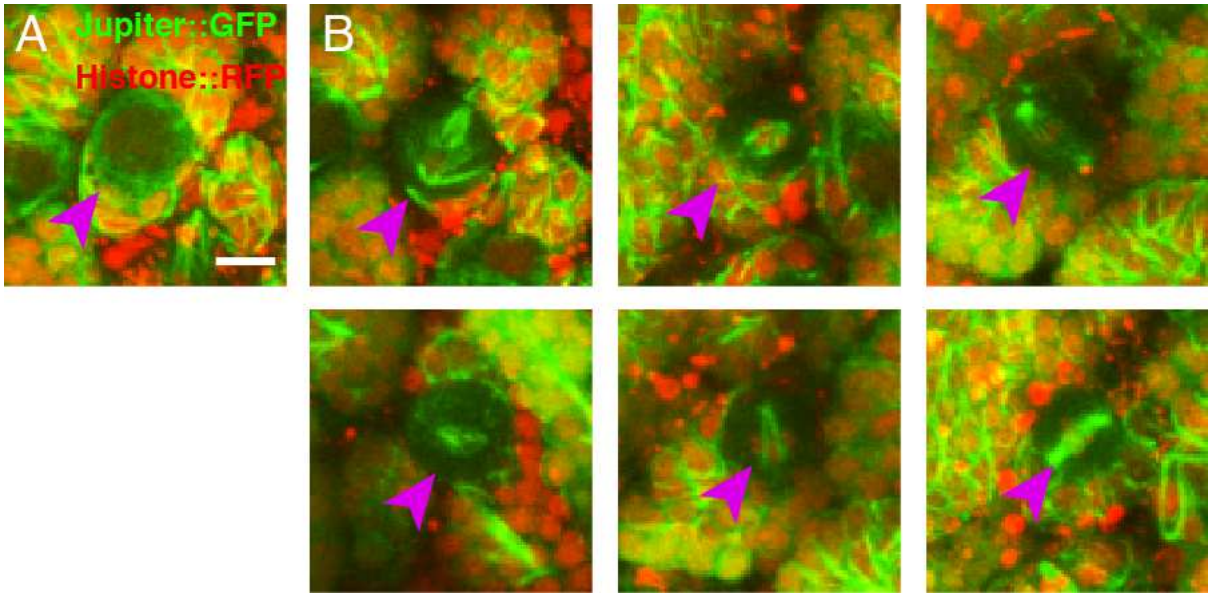

Figure 5.4: NB appearance varies with the cell cycle. A) Live imaging of NB (magenta arrowheads) marked with Jupiter::GFP (green), Histone::RFP (red) showing a maximum projection over 4 µm. Scale bar 5 µm. B) Multiple different dividing NB, at various different times throughout mitosis, with different spindle orientations.

Although dividing NBs do not quite meet the assumption of roundness from CytoCensus, nor does CytoCensus directly track cellular identity over time, I find it is possible to train on them and produce reasonable predictions of NB division rates. I therefore first trained on all NB, and then specifically on dividing NB, excluding non-dividing NB, using the training interface of CytoCensus. I then evaluated both of these models on the remaining image dataset to get the coordinates of dividing and non-dividing NB. This approach results in identifications of NB and dividing NB centres (**Figure 5.5**), not a direct measurement of division rate.

Dividing NB were not always well identified in these images, due to various confounding factors. First, they retain some features of non-dividing NB even when the spindle forms. This similarity of features means that some non-dividing NB are sometimes detected as dividing ones. Second, the apparent 2d shape of the spindle varies significantly between NB, primarily due to variable orientation of the spindle (**Figure 5.4**). Within a NB, spindle orientation remains fairly static over time, which is consistent with previous reports (Cabernard and Doe, 2009).

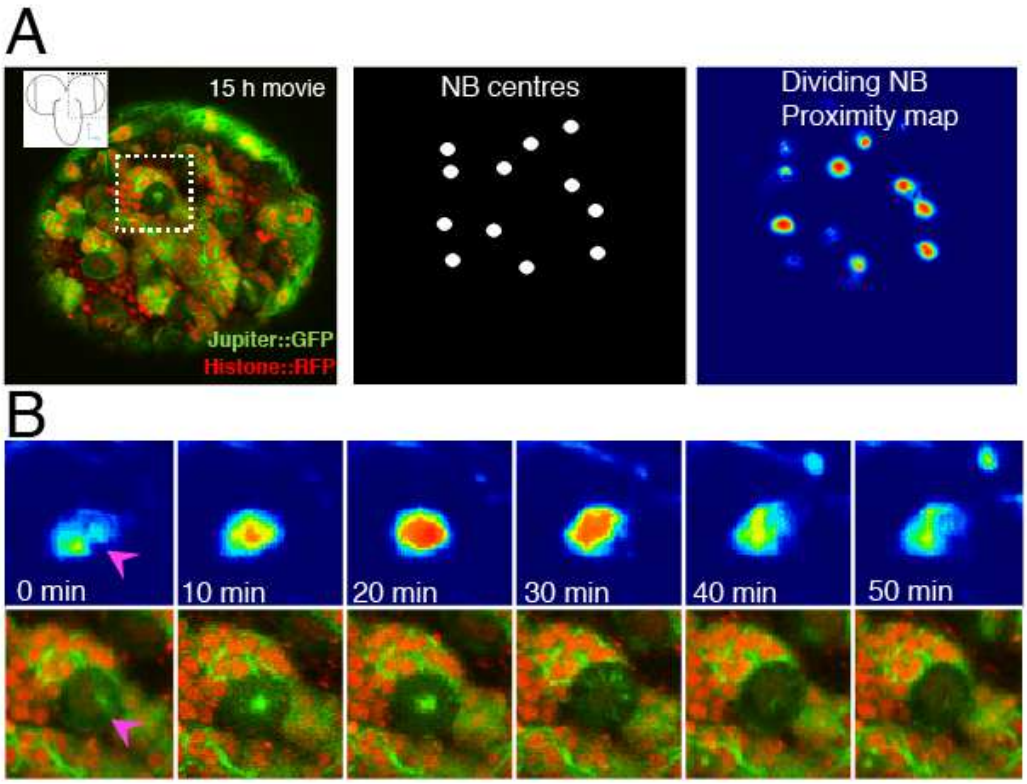

Figure 5.5: Detection of dividing NB in the brain. A) Live imaging of the central brain marked with Jupiter:GFP, Histone::RFP. Inset: position of imaging region. CytoCensus detections of NB centres within the larval brain, and the corresponding proximity map for dividing NB highlights those NB that show signs of division are to the right. B) Changes in single NB from the dashed box in (A) over 50 minutes of imaging as the NB goes through mitosis, with corresponding changes in the dividing NB proximity map above.

These complexities make it difficult to directly produce measurements of cell division. I therefore turned to the proximity map from the dividing NB (**Figure 5.5A**). Despite the problems outlined above, examination of these maps (**Figure 5.5B**) reveals that, as NBs divide, there are peaks in intensity in these proximity maps that correspond to NB division events (**Figure 5.5**). Therefore, rather than relying on the CytoCensus detections of dividing NB, I created plots that correspond to the 'probability' that a given cell is dividing over time (**Figure 5.6**). By following a NB detection, fluctuations in the proximity map for dividing NB can be measured, and these peaks correspond to division events. This results in plots (**Figure 5.6A**) showing a distinctive periodic pattern that corresponds to the NB division rate

(see **Methods**). By automatically generating plots of the changes in the proximity map over time, I was able to simplify the task of annotating division events in large datasets into detecting peaks in division plots (**Figure 5.6**). These plots greatly simplify the process of measuring NB division rate, removing the need to follow actual NB over time within movies, instead only requiring measurements of intervals corresponding to each division. In principle the peaks in these plots could automatically be converted into detections of division events, although here I rely on manual annotation of division events from these plots, primarily because they can be noisy, and intensity of the peaks can vary with NB movement. Additional filtering of these probability maps can improve robustness (**Methods**), and could be used for automatic detection.

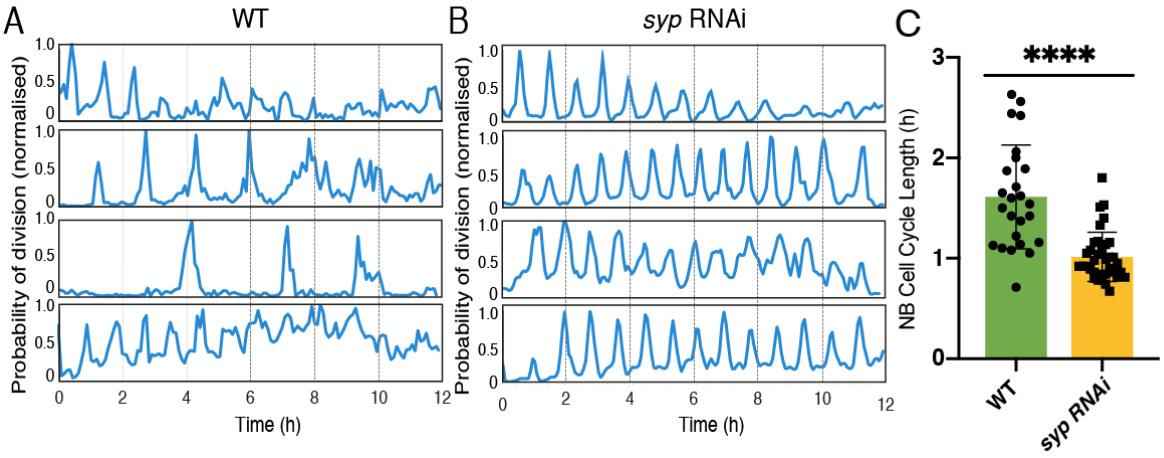

Figure 5.6: Comparison of WT and *syp* RNAi division rate. A) Examples of normalised proximity map changes ("probability of division") for 4 NB over 12 hours in WT. B) Examples of normalised proximity map changes ("probability of division") for 4 NB over 12 hours in *syp* RNAi. C) Average measurements of NB division rate for all NB that undergo at least 2 divisions within 12 h of imaging from 5 brains. Showing standard deviation. Significant differences indicated (p=<0.0001, t-test with Welch's correction, WT n=25, *syp* RNAi n=37 NB from 5 brains)

Using this approach, I can generate plots of NB division for each NB within the imaging area (I did so using a python script, but this can also be achieved with the FIJI 'plot profile' function). Generating these plots relies on the fact that NB are relatively large, meaning that

tracking is, in many cases, unnecessary, and where tracking is required a simple nearest-neighbour search was sufficient. I will return to a discussion of tracking with the discussion of GMCs (**Section 5.3**). I then used these plots to measure the division rate of all detected NB in a number of images (**Figure 5.6**). I compared *syp* RNAi brains to WT and found that the rate of division rate of NBs in *syp* RNAi brains was significantly faster than WT (*syp* RNAi 1.01±0.24 h, WT 1.62±0.5 h, $p<0.0001$ t-test with Welch's correction) (**Figure 5.6**C). I note that, in some brains a significant population of stalled NB was present (these are explicitly excluded from **Figure 5.6**), primarily in the syp RNAi population, which may be due to an increased sensitivity to phototoxicity (Ettinger and Wittmann, 2014) - indeed it appears that it is crucial to be aware of possible differential effects of phototoxicity for accurate assessment of division rates. Although best results were achieved using a model specifically trained on the dataset in question, good results were obtained on most datasets using the same model. It is essential to train a separate model for different genotypes, especially given that the NB in *syp* RNAi brains are significantly larger (Neumuller et al., 2011) which, unsurprisingly, reduces the accuracy of the WT model.

I conclude that average rate of NB division is significantly increased (by 62%) in the absence of Syp expression and, assuming this increase is constant throughout larval development, would result in 62% more neurons. This is more than enough to account for the 35% increase in *syp* brain volume. I therefore conclude this is the primary cause of the increased brain size at wL3 in *syp* RNAi. These results illustrate the potential of CytoCensus to assist in the analysis of the pattern of cell divisions in live imaging data of the *Drosophila* brain.

5.2.5. GMC CELL CYCLE LENGTH IS ALSO REGULATED BY SYNCRIP

Following on from the characterisation of the effect of *syp* on NB division rates, I also investigated the role of Syncrip in GMCs. Given the changes in division rate that we see in the NB, it is possible that similar changes in GMCs also contribute to the enlarged brain

phenotype. I therefore extended the analysis described above to GMCs. Unlike NBs which divide multiple times, GMCs undergo a single, symmetric division that gives rise to two daughter neurons. I confirmed that normal cell division was unaffected in *syp* RNAi brains, using the live imaging method described in **Chapter 3**, and saw no evidence for aberrant cell division as occurs for *brat* INPs. Because they undergo only a single division, I will discuss the GMCs 'cell cycle length' rather than division rate, even though they are synonymous.

Analysing the cell cycle of GMCs turns out to be notably more complex than the much larger NB (**Chapter 3**). It is apparent from live imaging that they also move around much more than their parent cells, moving quite significant distances in some cases (particularly during NB division). The most challenging aspect, however, is almost certainly the long cell cycle length. Based on my manual measurements in **Chapter 3**, the GMC cell cycle is approximately 8-10h in *ex* vivo culture conditions. This long cell cycle means that highly accurate tracking is required to successfully follow GMCs from birth until division. I have previously discussed imaging improvements that support following the GMC cell cycle (**Chapter 3**). Here, I describe a method for tracking GMCs over considerable lengths of time with the aim of determining the effect of *syp* on GMC cell cycle length.

The flexibility to use existing markers and imaging agents that are imperfect is one of the advantages of the supervised machine learning based approach in CytoCensus. However using existing markers (Jupiter::GFP, His::RFP) does not allow us to uniquely determine which cells are GMCs (**Figure 5.10, Chapter 4**). This makes the problem of tracking GMC cell cycle length considerably harder. Although using a different set of markers would make this problem simpler, creating new markers is not always an option. For instance it is desirable to be able to follow both NB and GMCs, which might not be possible with a GMC specific marker. Similarly expressed markers under the GAL4-UAS system could simplify the problem, but complicate

manipulations such as RNAi. Therefore, I will focus on using the existing Jupiter::GFP, His::RFP system.

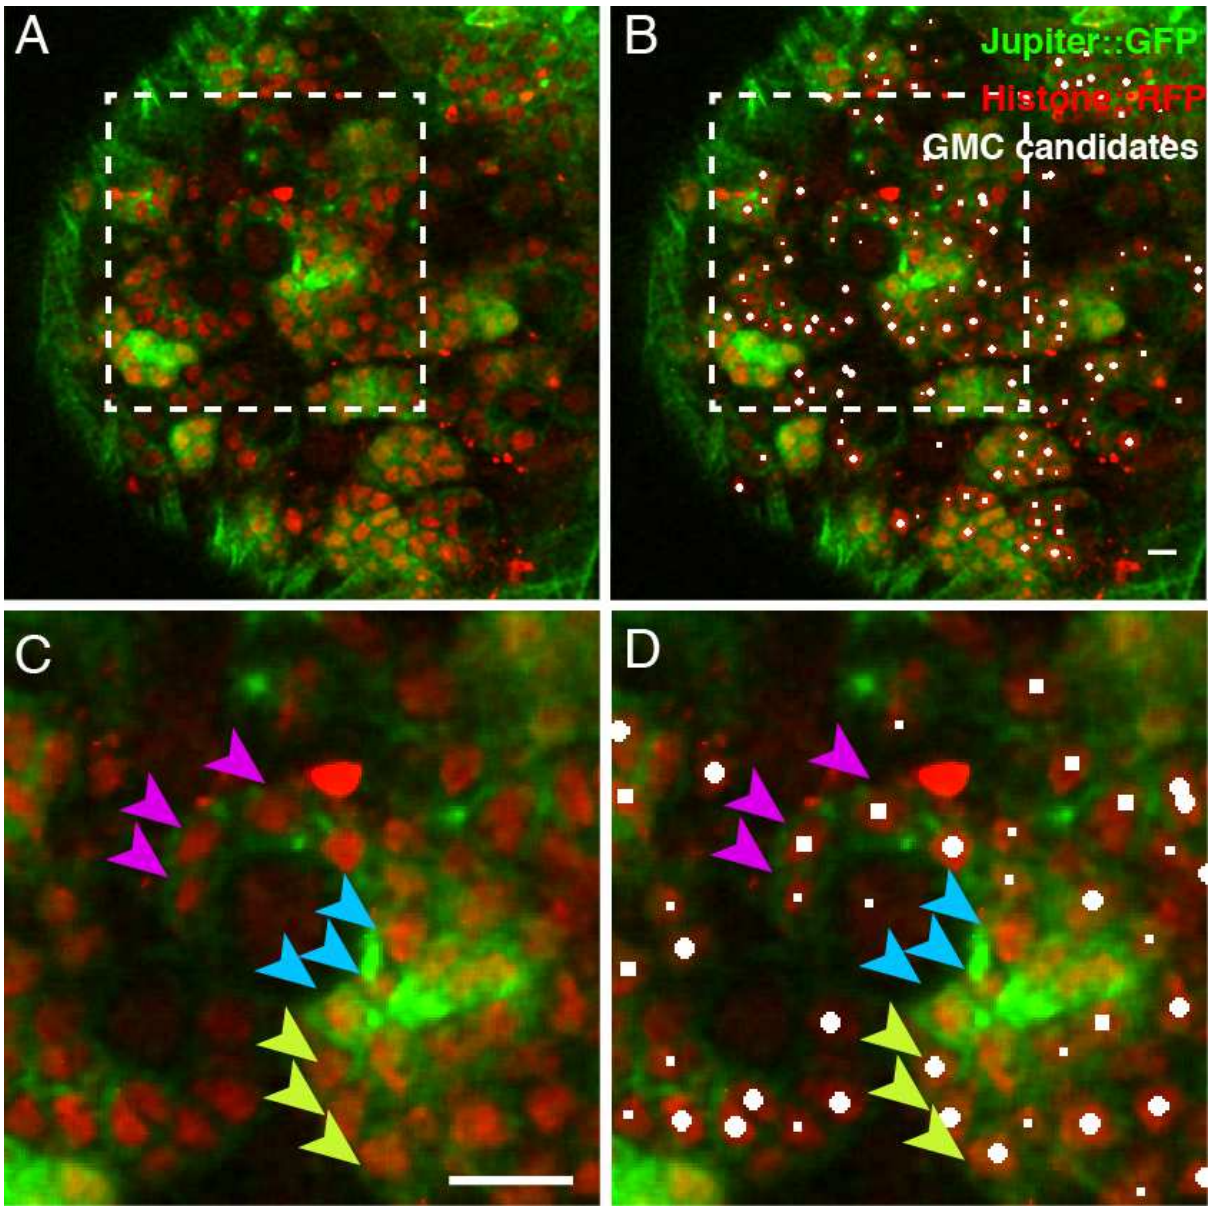

Figure 5.7: Identification of GMC candidates using CytoCensus. A) Improved imaging (with generic cytological markers Jupiter::GFP, Histone::RFP) to follow the GMC cell cycle, showing single plane. B) CytoCensus GMC candidates overlaid on the original image. Scale bar 5 µm. C-D) Identification of candidate GMC centres (magenta arrowheads) by CytoCensus. Correctly ignored neurons (Cyan arrowheads). Incorrectly identified GMC candidates (Green arrowheads) that are likely to be neurons. Scale bar 10 µm

Unlike NB, GMCs show a strong nuclear histone signal, that in principle could be tracked with existing software (for instance TrackMate in ImageJ or TGMM from the Keller lab)(Amat et al., 2014; Tinevez et al., 2017). However, the large number of cells present within the brain, as well as the bright cell debris (**Chapter 3**), and low SNR result in a high error rate of standard tracking. Therefore, I use CytoCensus as a pre-processing step to identify GMC candidates and exclude the majority of neurons, cell debris and NB from tracking (**Figure 5.7**). By training on known GMCs (those that divide or are born) and excluding those that are clearly neurons, it is possible to train a model that selects good candidate GMCs. This use of CytoCensus therefore simplifies the tracking problem.

Both these detections and the raw proximity maps can be used for tracking. The proximity map is compatible with TrackMate, (which finds objects using the LoG – See **Chapter 4**) and TGMM (which models cells as Gaussians). In principle this should work particularly well with TGMM, because CytoCensus and TGMM make similar Gaussian assumptions. Unfortunately, I ran into repeated problems getting the TGMM software to run. Instead I opted to use existing python tools (TrackPy) that makes use of the Crocker-Grier algorithm for tracking objects over time(Allan et al., 2018). I am able to directly use the detected centres from CytoCensus, which reduces the amount of processing required for tracking (compared to using the proximity map). The results of the tracking are shown in (**Figure 5.8**).

I next assessed the performance of tracking. The number of candidates that are tracked over varying amounts of time are shown in Figure. It is apparent that most candidates are tracked over relatively short periods (30% for less than 20 frames) (**Figure 5.8D**). Many of these represent isolated detections of non-GMCs, which are not surprising given the length of the imaging. It is possible to change the parameters of the tracking to be less stringent, leading to a higher proportion of long tracks; however in tests this results in increasingly large populations of cells that have one or more incorrect links, which are useless for calculating cell

cycle length. Nonetheless, once short tracks are excluded this represents a viable, if imperfect, method for tracking GMCs.

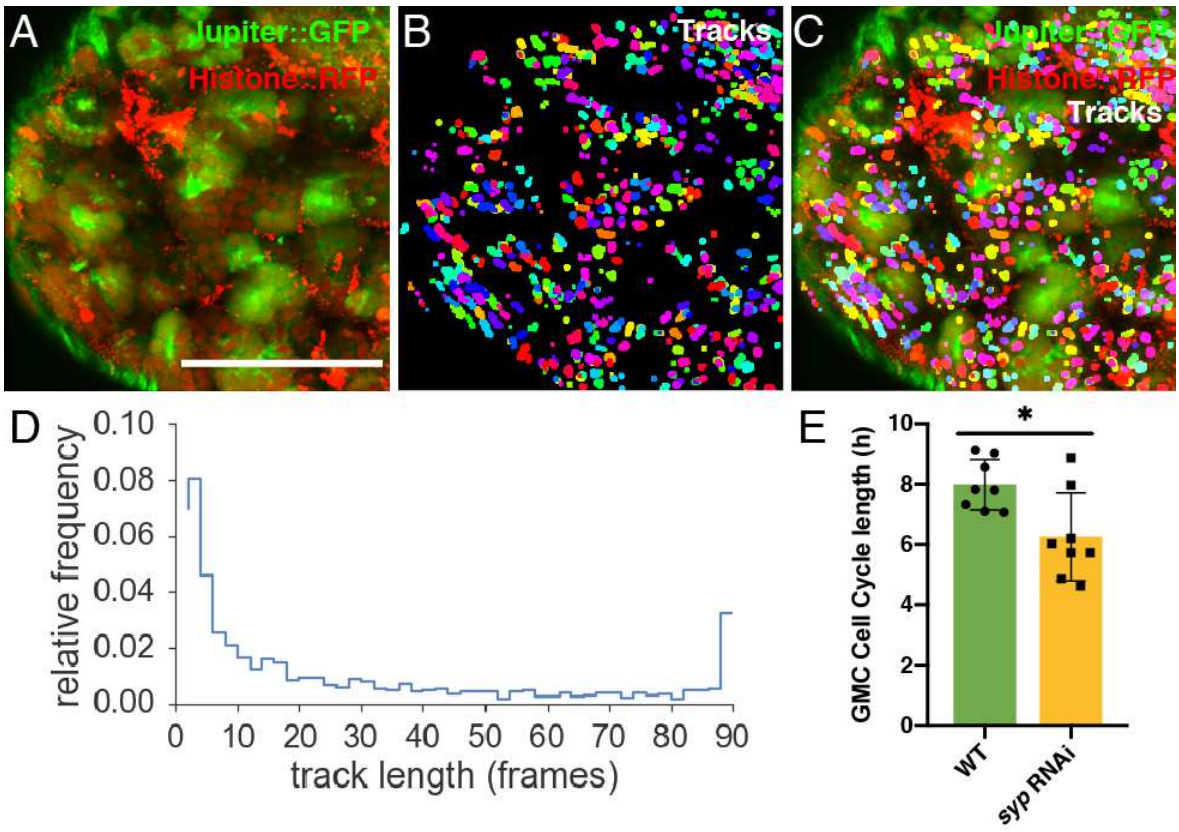

Figure 5.8: GMC tracking based on CytoCensus detections. A) Single plane (t=0) marked by Jupiter::GFP(green), Histone::RFP(red) Scale bar is 50μm. B) Corresponding tracks of GMC candidates over 100 frames (max projection over 5 μm). Each candidate track is randomly assigned a colour. Tracks that change colour represent errors in tracking. Tracks within constant colour represent successful stretches of tracking. C) Overlay of tracks with original image. Regions without any tracks represent areas that contain no GMCs. D) Proportion of tracks as a function of track duration. Note that only around 4% are 90+ frames. E) Plot of average measurements of GMC cell cycle length (hours) for both WT (n=8) and *syp* RNAi (n=8) showing standard deviation (p=0.04, Welch's t-test).

Using the methods described above, a number of GMCs were successfully tracked over 100s of time points (in 200 minutes, ¼ GMCs divided, ½ had at least one tracking error). This error rate means that, on average, multiple (~5) tracking errors are expected per GMC cell cycle. Although imperfect, this can be used as a processing step for simplifying the process of

tracking, which is very time consuming. I was able to use this method to simplify the measurement of GMC cell cycle length in both WT and *syp* brains, relying on automatic tracking for most timepoints, and manually correcting errors. The results (**Figure 5.8E**) show a significant decrease in the length of the cell cycle in GMCs in the absence of Syncrip.

These additional measurements confirm that RNAi knockdown of *syp* affects both NB and GMC cell cycle length. This suggests that common factors control the GMC cell cycle and the NB cell cycle. Given the small size of the change in GMC cell cycle length (25%), and the fact the GMC divisions remain normal (**Chapter 3**), it seems unlikely that GMCs significantly contribute to changes in brain size in *syp* RNAi, and hence the primary cause of the increase in brain size is an increase in the division rate of NB.

5.2.6. CONSISTENT HETEROGENEITY OF NB DIVISION RATES

One of the advantages of using live imaging and automatic CytoCensus analysis is that it allows us to examine the behaviour of multiple neuroblasts over time within the same brain. This is not possible using *in vitro* imaging. Using this approach I can examine the effects of *syp* RNAi on the distribution of NB division rate. Interestingly, I find that while each NB has a consistent cycle period, which matches observations made in vitro, there is considerable variation in cell cycle length between NB within the same brain without an obvious spatial correlation of cell division rate **(Figure 5A)**. By contrast, *syp* RNAi neuroblasts appear to have a more consistent division rate in addition to a shorter average rate (**Figure 5B, C**). This appears to hold true both for individual neuroblasts (**Figure 5B**), and between brains (**Figure 5C**). It is unclear what causes this variation, but it is tempting to speculate that this is the result of differing levels of rate-determining factors, Syp among them. I conclude that live imaging and CytoCensus can provide biological insights into development that would be difficult to obtain by other means.

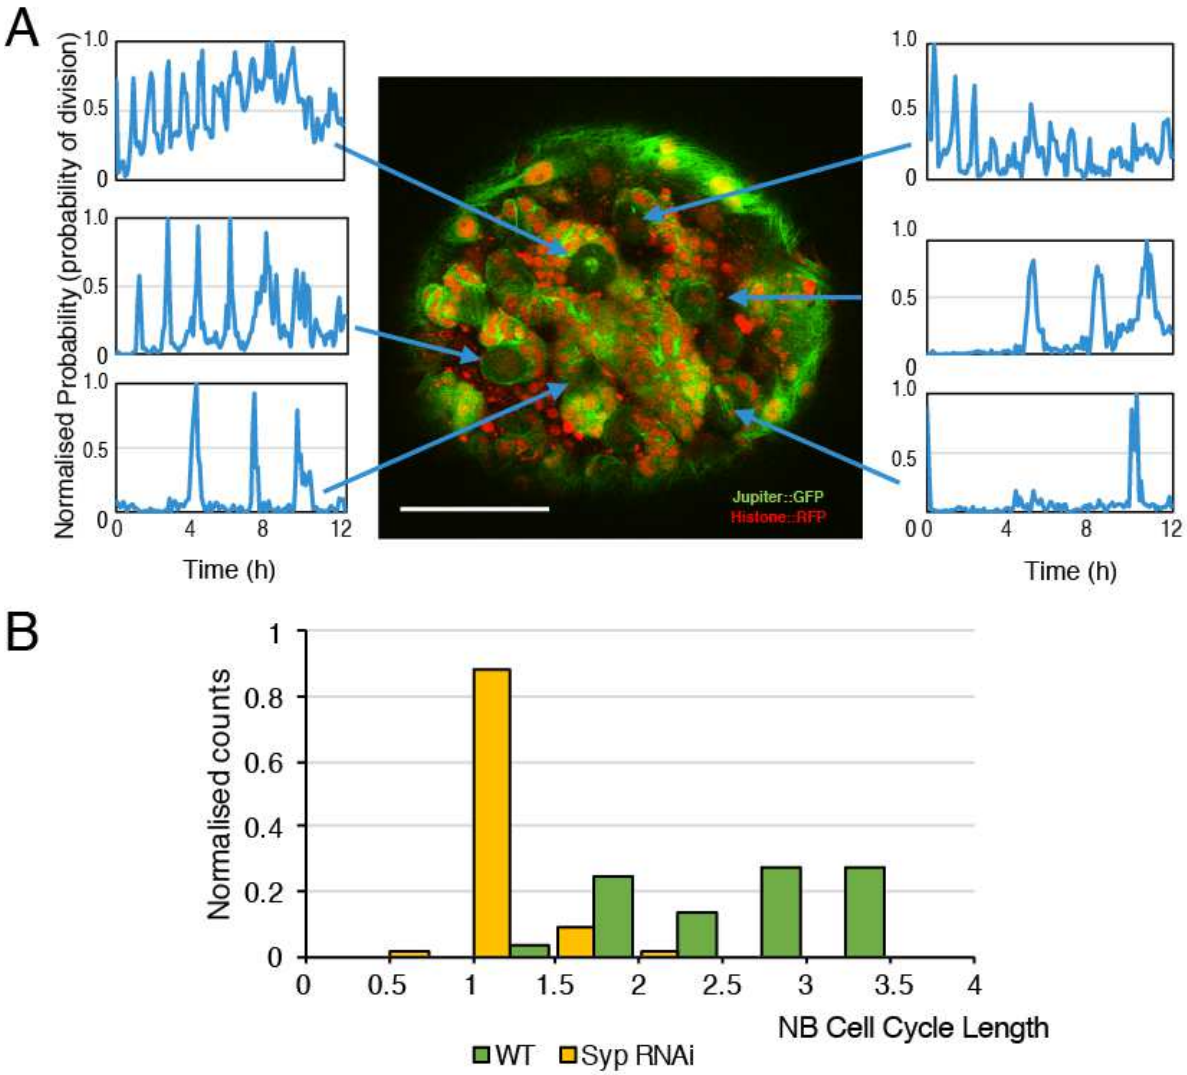

Figure 5.9: NB division rates are heterogeneous. A) Plots of NB division and their corresponding locations within the brain. Normalised probability based on proximity maps for dividing NBs. Scale bar 50µm. B) Histogram of normalised counts (sum of all counts = 1) of cell cycle lengths, and over time for *syp* RNAi and WT.

5.3. Discussion

In this chapter I have explored the *syp* phenotype in *Drosophila* larvae, using live imaging and culture (**Chapter 3**) and CytoCensus (**Chapter 4**). I show that in combination these approaches provide useful data on cell behaviours, including following division over long periods of time, and use this to characterise the *syp* phenotype, identifying a role for Syncrip

in regulating NB division rate. I show that in *syp* RNAi there is a significant reduction in the division rate of NB, with almost all neuroblasts dividing consistently at a cell cycle length of 50-60 minutes. I conclude this is the primary driver of the brain overgrowth phenotype, with a secondary effect of increasing the number of NB in *syp* -/-.

The difference between the phenotype of *syp* RNAi and the *syp* mutant (the increase in NB number, **Figure 5.2B**) suggests that an additional effect is at play in the *syp* mutant. This raises a couple of possibilities. First, the effect of the RNAi may be incomplete. Inscuteable-GAL4 is, in principle, widely expressed, however, the GAL4 driver is clearly more restricted in practice e.g. (Neumuller et al., 2011). This suggests the origin of the additional NB may be another cell type other than the NB, GMC or Neurons in which the GAL4 is expressed. Second, the presence of residual Syncrip in the RNAi knockdown could be sufficient to disrupt the unknown cause of these additional NBs. Further experiments are required to establish the origin of the effect, for instance using the ubiquitous actin–Gal4 driver to determine if it is specific to the inscuteable promoter. Perhaps the most interesting possibility, however, is that the NB are prevented from normal programmed cell death at the end of embryogenesis due to the absence of Syncrip (Pinto-Teixeira et al., 2016). Additional neuroblasts would therefore be the result of escaping the normal sequence of cell death, which would be consistent with the effects of Syncrip in the pupa (Yang et al., 2017).

In addition to measuring the cell cycle in NB, I used CytoCensus-assisted tracking to determine the length of the cell cycle in GMCs. Whilst this does reduce the amount of effort required to track GMCs, it still requires manual annotation and requires significant time (upwards of 50h) to process large images. Using smaller image datasets is in principle much more efficient, but may limit the processes that can be studied. Furthermore, I show that there is significant variation in the normal period of neuroblast division within a brain, something that is difficult to determine without *ex* vivo imaging. Although each NB has its own unique

period, which is approximately maintained throughout culture, variation occurs from one NB to the next within a brain. This is in fact roughly consistent with estimates of NB division rates based on clonal analysis which find significant variation in the rates, although clonal analysis does not determine whether this is an intra- or inter-brain effect (Bowman et al., 2008). The measurements give somewhat longer average rates compared to other approaches, perhaps, in part, due to temperature differences, although I cannot exclude the effects of culture or imaging toxicity. The estimated rates in Bowman in *et al*. (2008), however, significantly differ from those in *in vitro* culture, which are much faster on average. This favours *ex vivo* imaging as a better representation of the true *in vivo* state.

I hypothesise that this variation between individual NB may be due to differing levels of factors that influence division rate within the NB population. Under *syp* RNAi, this heterogeneity is almost completely abolished, suggesting that the absence of Syncrip disturbs the normal variation in factor distribution and levels. Given the consistency of NB division rates in *syp*, as well as the shorter cell cycle in GMCs, I suggest that *syp* is a key regulator of such factors, perhaps via its RNA binding activity. This is consistent with evidence that Imp and Syncrip vary throughout development (Yang et al., 2017). In fact, gradients of Imp and Syp have been shown to vary between NB, perhaps explaining the normal heterogeneity that I observed in division rate (Liu et al., 2015).

In this chapter, I have illustrated the use of CytoCensus in analysing live imaging data to characterise changes in NBs and GMCs in the *syp* mutant, revealing the cause of the enlarged brain phenotype. In future this approach could be applied to other mutants, such as *brat* or *imp*, to examine their role in the control of *Drosophila* neural lineages. In my final results chapter I will examine broader applications of CytoCensus beyond *Drosophila* and the brain.

CHAPTER 6

THE WIDER APPLICATION OF CYTOCENSUS: CHALLENGES AND MODIFICATIONS

6.1. INTRODUCTION AND AIMS

Previously I have validated the use of CytoCensus on synthetic and fixed images, and applied live imaging and CytoCensus to characterise normal and aberrant development in *Drosophila* brains. However, the challenges of quantifying cell number and behaviours that were faced in previous chapters are not unique to *Drosophila* or to the brain. Cellular resolution imaging of live and fixed tissues is increasingly widespread in vertebrates such as mouse and zebrafish, where it is used to investigate mechanisms of development and model disease states (Sullivan and Lundberg, 2018). With this kind of large scale imaging comes the same challenges of quantification of cell number and behaviour that I have addressed previously. In this chapter, I will illustrate the generality of CytoCensus, by applying it to a series of different biological exemplars from different organisms and at different scales, modifying it where necessary. In each case I outline a biological question for CytoCensus to address, and demonstrating that the approach simplifies the task, and provides useful quantitative data.

Historically, the most common approach to quantification in imaging is manual analysis. Hooke detailed his observations by drawing what he saw under the microscope; early embryology also relied on such drawings, counts and observations of aberrant features (Hooke, 1968). Today manual observations are made from digital images for small and simple

tasks. As the scale of imaging has increased, the development of analysis tools, such as ImageJ, has been necessary, and for more complex tasks enlisting the help of an image analyst is increasingly important (Schindelin et al., 2012). In this chapter I aim to illustrate that many of these tasks can be automated straightforwardly, using CytoCensus. In many cases, this can be achieved without dedicated image analysis or programming knowledge, using existing variable settings, and relatively straightforward data analysis where necessary. Even in images of more complex objects or features, CytoCensus can be straightforwardly modified to include additional information.

In this light, I apply CytoCensus to images of tissues from Zebrafish, mouse, and the *Drosophila* testis, and illustrate that CytoCensus can be used to answer questions about cell (and organelle) number, location, and size. In each case, existing analysis relied on manual counting for quantification, or bespoke analysis pipelines. This work shows CytoCensus is a practical and efficient alternative to existing methods. In one example, that of organelles in spermatogenesis, I determine that the organelles of interest are sufficiently different from the assumptions made by CytoCensus, that the problem is difficult to approach in the standard manner. Here, I illustrate the extensibility of CytoCensus, by defining an additional filter set (based on radial symmetry) that makes the problem tractable (Loy and Zelinsky, 2003). Using this new feature set (FRST-f) I show excellent results in locating organelles in both a WT and mutant setting.

6.2. Zebrafish Retinal Organoids

6.2.1. Motivation

Zebrafish development is a relatively well studied vertebrate system, and has many useful genetic and imaging tools with which to study development. One such tool is the Spectrum of Fates (SoFa) zebrafish (Almeida et al., 2014). This system places expression of distinct fluorophores under control of cell type specific promoters. The pattern of fluorophores allows direct identification of the different cell types that make up the zebrafish eye. Using the SoFa zebrafish, the Harris lab has developed methods to dissect important components of eye development in Zebrafish using an artificial organoid system (Eldred et al., 2017a; Eldred et al., 2017b). In these papers they show that, dissociated cells from a SoFa zebrafish can assemble into 'retinal organoids', that is, structures that resemble the Zebrafish eye. They demonstrate that these organoids show layers of different cell types, similar to those usually found in the eye. Unsurprisingly, they do not completely recapitulate zebrafish development: for instance, they appear to have a different structure from normal zebrafish eyes. Nonetheless, this represents a tractable system with which to follow the requirements for and dynamics of eye development over long periods of time. They show that this approach can be used to study how the structure of organoids is affected by the presence (or absence) of Retinal Pigment Epithelial (RPE) cells. In order to do so they developed, with the help of an image analyst, a MATLAB-based image analysis pipeline specific to their problem (Eldred et al., 2017b). Developing a pipeline is a common approach to this kind of problem, but they are typically inflexible, and cope poorly with imaging variation (Stegmaier et al., 2016). In this case, the pipeline produces only measurements of total fluorescence intensity as a function of distance from the organoid edge (Figure 6.2). This is an excellent example of the use of a bespoke image analysis pipeline being used to quantify changes in cell fate and behaviour; answering the question of whether RPE cells are important for the assembly of retinal organoids.

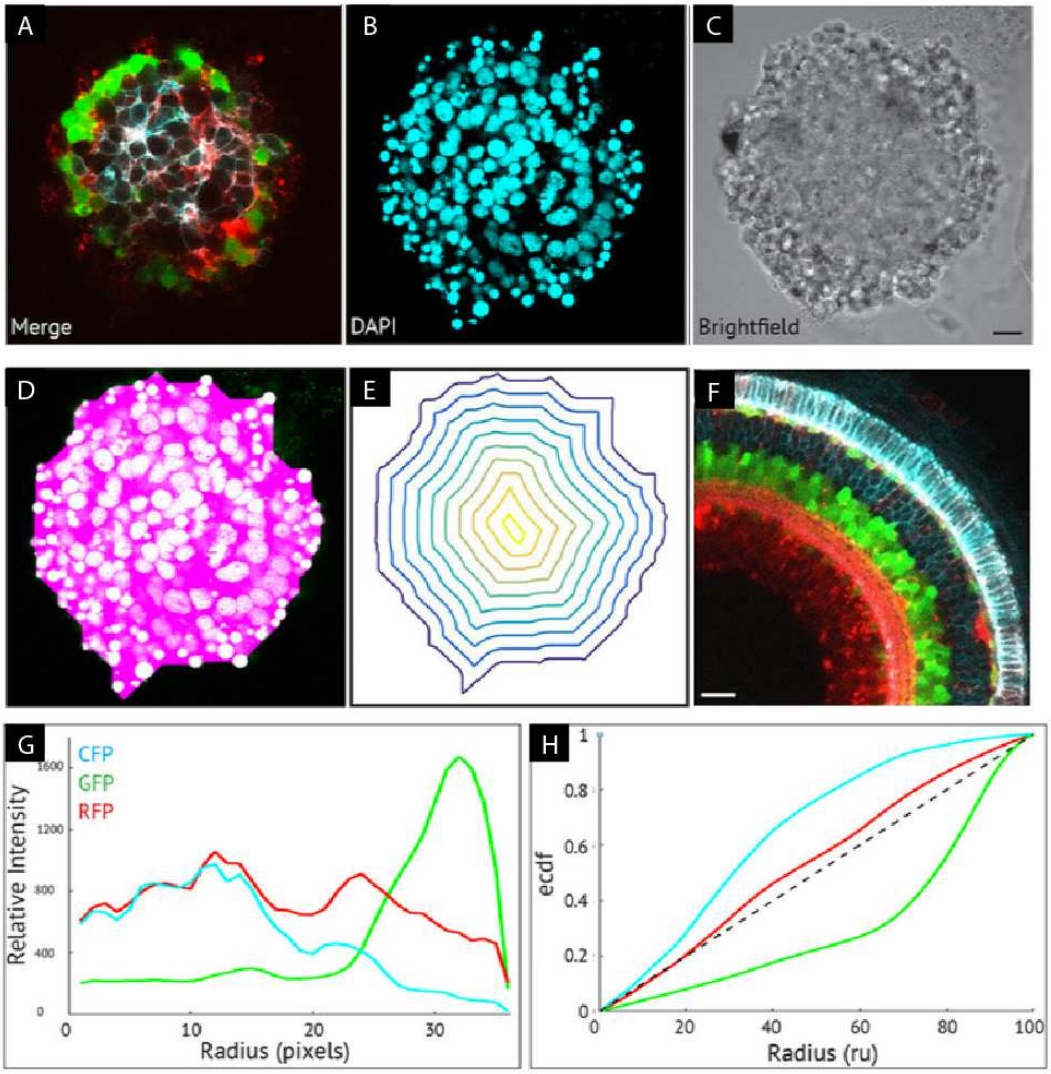

Figure 6.1: Fluorescence measurements of Zebrafish Retinal Organoids labelled with Spectrum of Fates fluorophores (Red:Atoh7:gapRFP, Green:Ptfla:cytGFP, Cyan, Crx:gapCFP) in the presence and absence of RPE cells. Adapted from Eldred et al. under CC-BY license Fig 2 A-C) Brightfield and Fluoresence images of organoids, showing single Z-Section D-E) Manual outline used for quantification of fluorescence intensity at different distances from the centre F) Zebrafish retina marked with Spectrum of Fates fluorophores G-H) Fluoresence as a function of intensity from the organoid centre

6.2.2. AIMS

As a test case for CytoCensus, I aim to illustrate that I can achieve similar quantification of the effect of RPE cells on retinal organoid development, but in more detail. I aim to answer whether CytoCensus is sufficient, using only user training, to determine whether the presence or absence of RPE cells significantly affects the distribution of cells within retinal organoids.

6.2.3. RESULTS

First, I tested whether CytoCensus can identify and quantify the different cell types that are found in the SoFa Zebrafish organoids. I trained multiple models on a set of images of zebrafish organoids (Fig 6.2A, 6.3B, gift from the Harris lab) to identify the 4 cell types defined by the SoFa cell line (A/H Amacrine/Horizontal, BP Bipolar, RG Retinal Ganglion, PR Photoreceptor) and to distinguish live and dead cells. I found that CytoCensus can successfully learn a set of models that correspond to these different cell types in the organoids when trained using two example images. I applied these models to the remainder of the dataset and used to extract the 3D cell centre locations of these different cell classes (**Figure 6.3**C). This achieved cellular level localisation (**Figure 6.3** B,C), and quantitation (6D) of each cell type.

I next examined whether these cellular localisations can be used to quantify changes in organoid size without the manual annotation of the organoid edge. I trained models on the dead cell layer at the edge of the organoid cells of the organoid, and all living cells (visible from DAPI and brightfield, Figure 6.3B,C). Living and dead cells can be distinguished by combination of brightfield and DAPI, which shows brighter spots characteristic of cell death, along with an uncharacteristic shape, and cannot be identified using total fluorescence measurements as in Eldred *et al.* 2017a. These localisations define the edge (Dead cells) and the extent of the organoid (Live cells), and allowed me to normalise the organoid size to the range (0-100) (see **Methods** for details). Although manual training on these classes was required, it was performed on only 2 images, and a small number of cells, and seemed to

generalise well to the rest of the data. In principle, this approach of automatically defining the extent of the organoid removes the need to annotate every image and could be very advantageous with much larger datasets, for instance in live imaging.

To determine whether these localisations and measurements of organoid size can be used to quantify the effect of RPE cells on the organisation of retinal organoids, I analysed the cell centre output from CytoCensus. By calculating the number of cells at different radii relative to the size of the organoid, these cellular locations were then used to calculate plots of the distribution of cell types in the RPE+ and RPE- conditions (**Figure 6.3**D). I find the distributions are very similar and no significant differences were found in the average (relative) positions of the 4 cell classes between the RPE+ and RPE- conditions (t-test with Dunn-Sidak correction for multiple comparisons). This showed that CytoCensus can produce similar results to Eldred et al. 2017a (**Figure 6.2**F,G) , but at cellular resolution, and without the need for a dedicated image analysis pipeline.

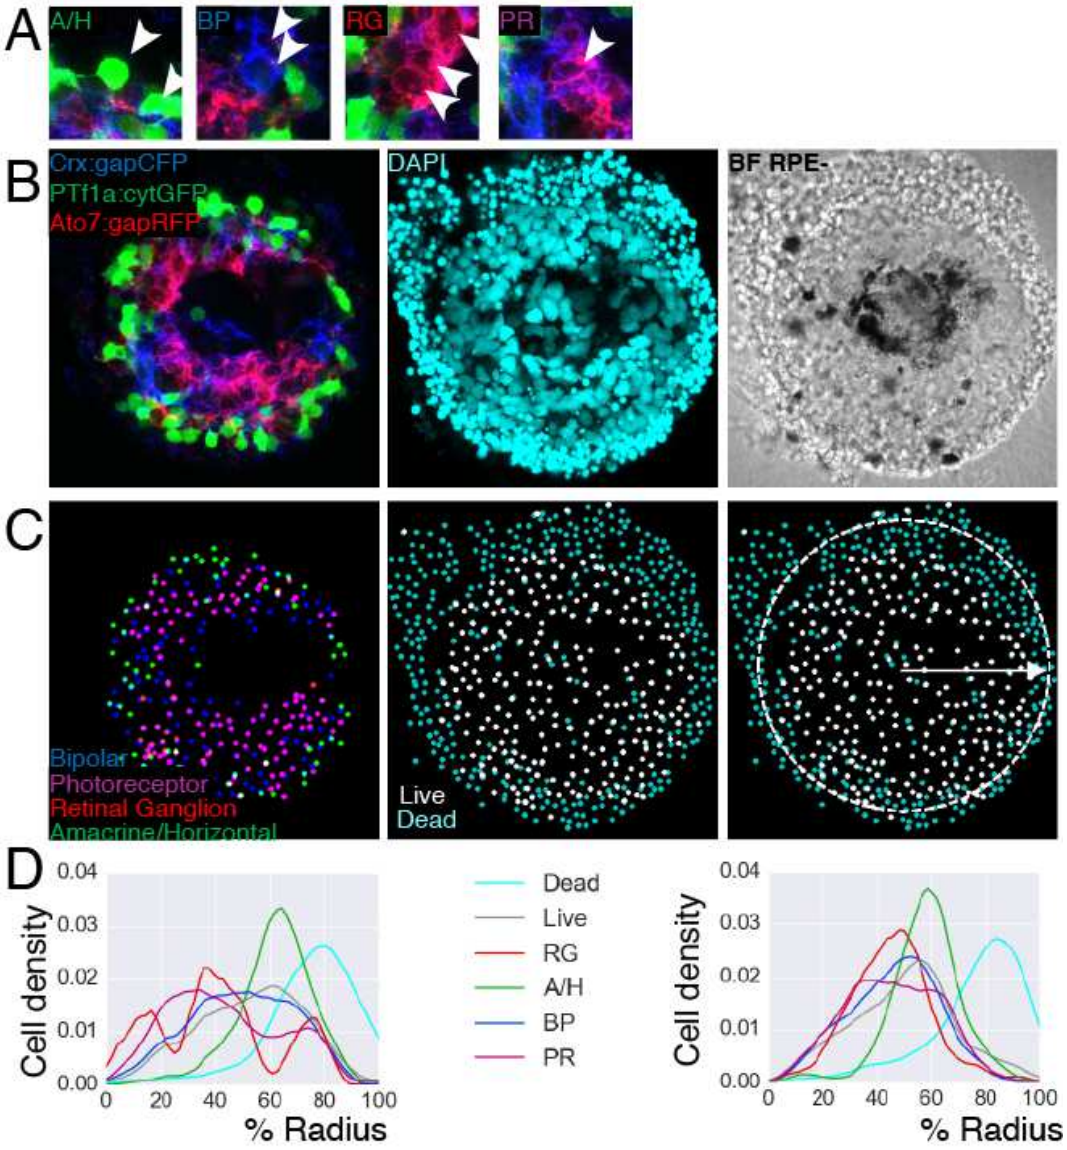

Figure 6.2: Analysis of Cell types from Organoid data using CytoCensus A) Cell types (white arrows; A/H Amacrine/Horizontal, BP Bipolar, RG Retinal Ganglion, PR Photoreceptor) found in the zebrafish eye and organoids marked using the SOFA markers B-C) CytoCensus identification of the different cell classes showing P/R, BP, A/H, Live, Dead cells, and the 'outside' of the organoid based on the Dead layer D) Cumulative numbers of different cell types plotted against distance from the organoid centre as a % of organoid radius in the presence and absence of the RPE cells

In order to determine if there were differences between the conditions that are visible only using quantitation at the cellular level, I investigated additional ways to present this data. CytoCensus results can be used to generate different ways of displaying the quantitation of the cellular level data (**Figure 6.3**). For instance, **Figure 6.3A** summarises the number of cells at a

given radius. It is clear from this figure that the distribution of cell classes is very similar between the two conditions (**Figure 6.3A**). In fact, no significant differences in the average position of each cell type were found in the RPE+ and RPE- conditions (ns, t-test with Dunn-Sidak correction). However, there were other differences between the conditions (**Figure 6.3B**). Organoids grown in the presence of RPE cells contain significantly more live cells than when they are grown without RPE cells (p=0.04, t-test with Dunn-Sidak correction, n=5 per condition). Interestingly, this change primarily seemed to be caused by an increase in the number of Retinal Ganglion cells, which were 3.7-fold more common in the RPE+ condition (p=0.047). Other cell classes were increased but not significantly. It is worth noting that these changes in live cell numbers were not possible to determine using the pipeline in Eldred et al. Similarly, it is possible to determine whether or not different cell types occupy different layers within the organoid. For instance, Bipolar cells are found significantly further from the organoid centre than Retinal Ganglion cells (p=0.037), but before the majority of A/H cells (p=0.011). This level of quantitation is impossible to do from total fluorescence measurements alone.

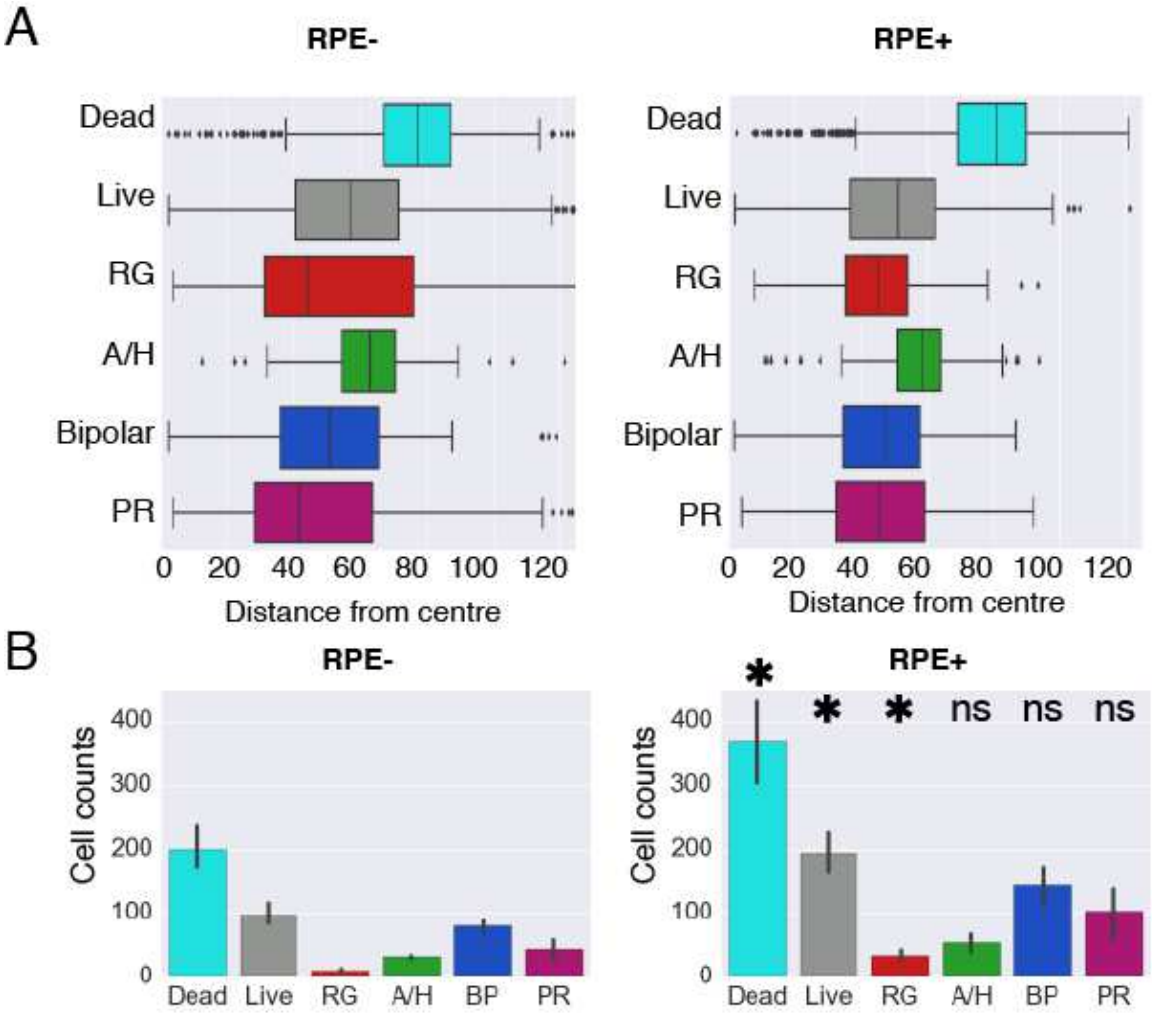

Figure 6.3: CytoCensus can be used for additional cellular-level analysis A) Box and whisker plots of cellular distributions plotted relative to total size of organoid, n=5 images per condition, no significant differences between conditions were found (t-test with Dunn-Sidak correction) B) Number of cells of each class in whole organoid in RPE- and RPE+ conditions. Significant differences in cell counts between the conditions are indicated (t-test with Dunn-Sidak correction, n=5 per condition)

6.2.4. Conclusions

In conclusion, CytoCensus showed that there was no significant effect of RPE cells on the relative position of cells in zebrafish retinal organoids. Building on previous analysis, CytoCensus produced cellular level quantitation of retinal organoids, and demonstrated there is a significant difference in total cell number between the conditions. More generally, CytoCensus can provide detailed quantification of cell numbers and types in different cell

classes in organoid development. This has the potential to increase the speed and ease with which organoid development can be followed. Although "still" images are analysed here, with appropriate imaging, it should be straightforward to quantitate and follow these cell types over time using CytoCensus and the approach detailed in Chapter 5. This could give interesting insights into how cell markers change over time, and/or how cells migrate to their respective cell layers.

6.2.5. LIMITATIONS

Although I used CytoCensus as the key image analysis step to successfully assess changes in cell type and distribution in retinal organoids, additional processing of the resulting data points is necessary. Specifically, to replicate the analysis performed in Eldred et al, the number of cells were calculated at different radii (e.g. Fig 6.4) using a python script. A similar analysis would be possible to do in Excel by use of the FREQUENCY array function (Microsoft® Excel for Mac 2016) if somewhat complex for a biologist without specific knowledge, or more straightforwardly in a suitable graphing program such as PRISM. It would be possible to incorporate this kind of analysis into CytoCensus, but given the range of possible analyses, it is preferable to make a more general-use program that relies on some degree of further data analysis for specific questions. Given that this analysis is possible using CytoCensus and other common programs, however, I nonetheless consider to be a significant simplification of the process compared with the published approach.

A key limitation of almost any supervised learning approach is that, at best, the approach can do no better than the data it was trained on. In this context, user training on the SoFa labels is subjective for some cell classes. Specifically, and in light of further experiments, the Harris lab raised questions on the identity of the photoreceptor (PR) cell class (private communication). Although the models produced by CytoCensus are reproducible, without a ground truth these models are subjective. User training is required to differentiate between cell types, and hence

the model is in principle no less subjective than the human that trained it. In this case, the difference between PR cells, BP cells and A/H cells, is based on the relative levels of CFP and RFP (Figure 6.3A), and, from this dataset, it is not possible to definitively determine their cell types. In fact, the combination of two membrane markers make the PR cell type particularly difficult to distinguish from RG and BP cells (Fig 6.3A, RG and BP vs PR). This is partially because CytoCensus does not directly work with object boundary information; it cannot distinguish cells that are near, but outside, a boundary from those that are inside a boundary. With more specific labels (e.g. cytosolic or nuclear) it should be possible to better identify the photoreceptor cell type with CytoCensus.

When dealing with multiple, potentially ambiguous or overlapping cell types like these, there is a need to reconcile multiple possible detections. In CytoCensus and normal regression, there is no enforcement of unique cell classes (that is each cell can only be assigned to a single class). Although approaches to enforce a unique identity exist, for instance by using multi-class (see Chapter 4) regression, this is not currently implemented in CytoCensus. Including this kind of approach would likely increase the training and prediction time for CytoCensus, which would make training more laborious. Although, by removing the need for multiple models, it would ultimately speed the analysis of multiple cell types. Using multi-output regression (based on multi-output decision trees) has the potential to improve discrimination of similar cell types, which would be an advantage. Instead, an approximation of this can be made by assuming that multiple detections within a given cell radius represent the same object. If multiple detections are made then the one with the highest value of the hessian is the best candidate for this. In the early versions of CytoCensus no confidence was given for candidate points in the exported coordinate files. In light of these problems, I have updated CytoCensus to include a confidence value for each point, allowing post-processing decisions to be made over the identity of a point in the multi-class case.

6.3. Mouse embryonic development

6.3.1. Motivation

The mouse blastocyst and early stages of embryogenesis are commonly used to study the spatial and temporal organisation of tissues during development, and the factors that dictate cell fate and differentiation (Arnold and Robertson, 2009; Piliszek et al., 2016). Imaging of different stages of mouse development, including of the late mouse blastocyst, is an important tool for examining these mechanisms. However, assessment of changes in cell numbers frequently relies on manual counting and estimation of phenotype, for instance in in-situ hybridisation and immunostaining of mouse blastocysts (Chazaud et al., 2006) Indeed in this field fluorescence imaging has not always been standard, and imaging has relied on, for instance, beta-galactosidase based reporting at resolutions that do not allow identification of individual cells (Sundararajan et al., 2012). Whilst this allows for excellent signal amplification, and is relatively straightforward, modern fluorescence microscopy can also provide good signal amplification and, in the right imaging conditions, have the potential to provide much more information at a cellular level (Saiz et al., 2015). Existing work in this direction (Saiz et al., 2015). shows the potential of imaging to add a temporal dimension to existing transcriptomic analysis but relies on a the somewhat inflexible MATLAB based MINS. Nevertheless, this combination of increasingly widespread cellular resolution fluorescence imaging and common manual analysis makes this an excellent target for an illustration of the use of CytoCensus in automating quantification of transcription factor expressing cells.

6.3.2. Aims

In this section I aim to quantify the number of cells in different compartments of the mouse embryo expressing different transcription factors.

In this section, I will consider three transcription factors (TF) that are important for the correct development of the murine embryo. Sox2, is a key regulator of self-renewal and pluripotency

with known roles in the neuroectoderm and mesoderm. Lhx1 is part of the Lim-1 family which are essential in the VE (Visceral Endoderm) and PS (Primitive Streak) for development of the head. Finally, Brachyury has wide ranging roles, including in the mesoderm, but is found in the PS around E6.

I aim to show that CytoCensus provides quantitative measurements of the proportion of cells for Brachyury, Sox2, and Lhx1 in the compartments of the endoderm and primitive streak.

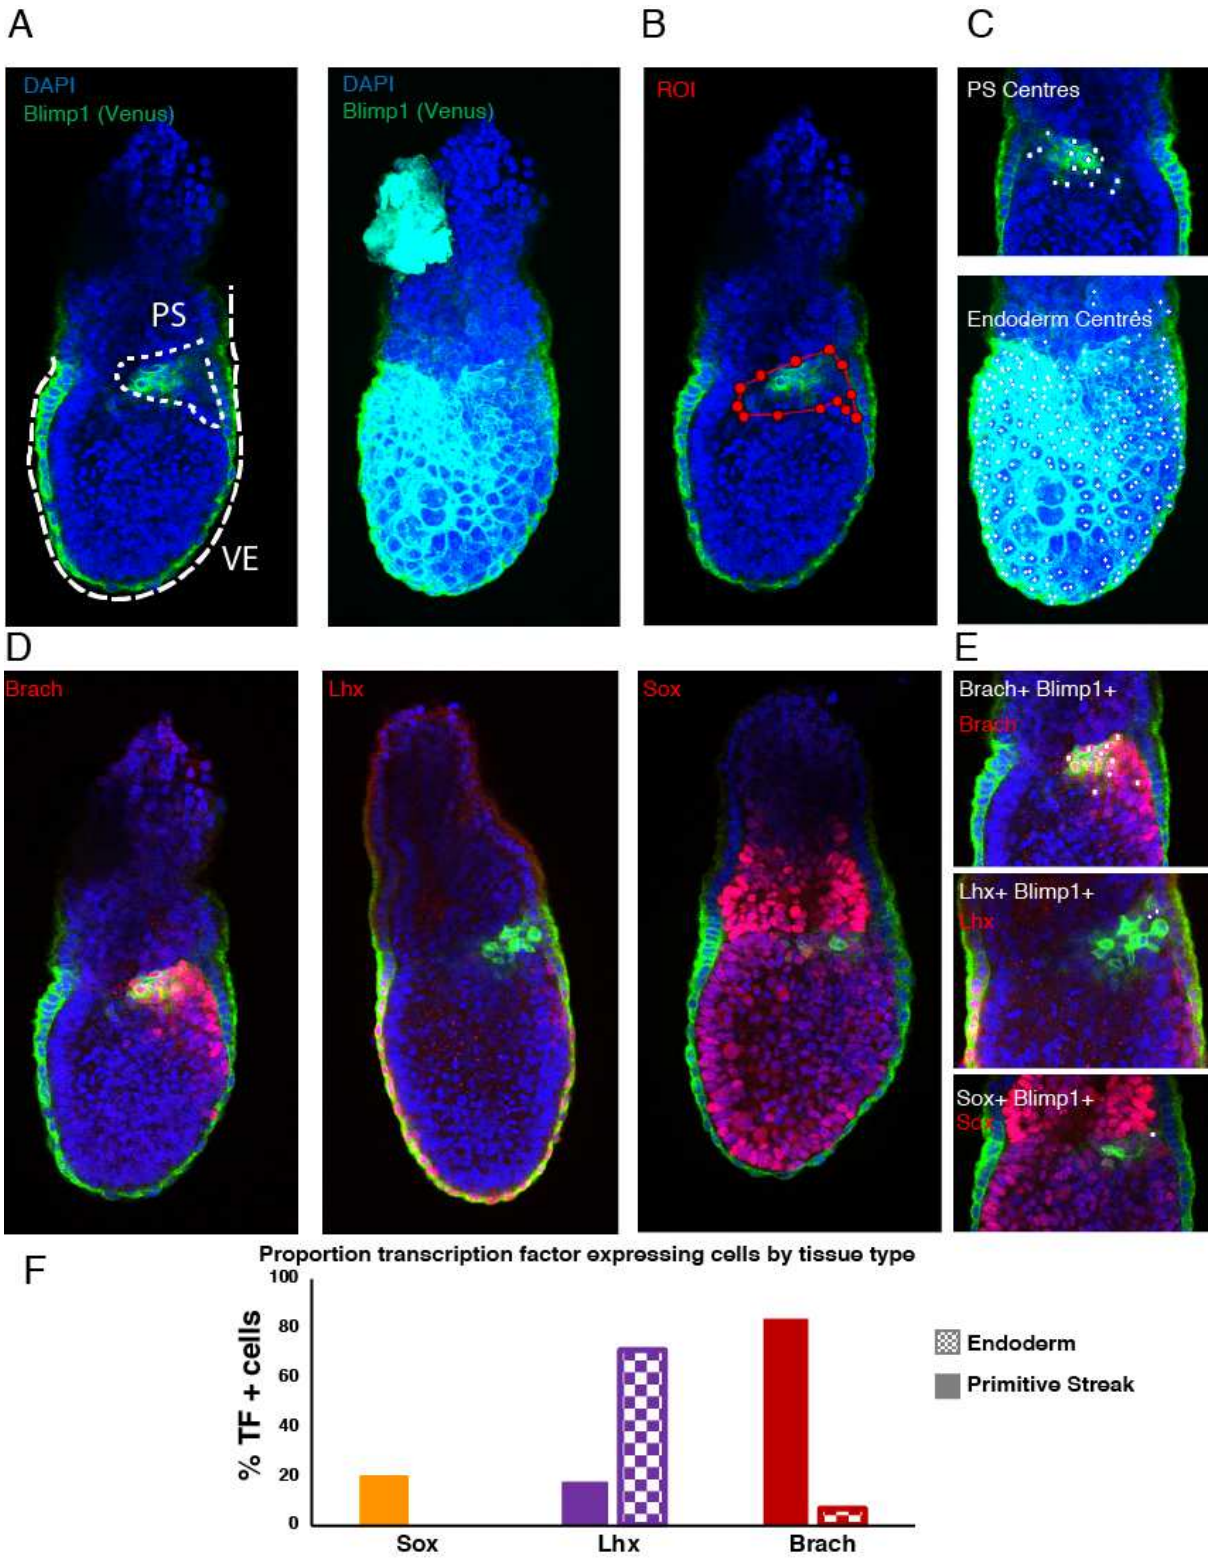
A
DAPI
Blimp1 (Venus)
PS
VE
DAPI
Blimp1 (Venus)
B
ROI
C
PS Centres
Endoderm Centres
D
Brach
Lhx
Sox
E
Brach+ Blimp1+
Brach
Lhx+ Blimp1+
Lhx
Sox+ Blimp1+
Sox
F
Proportion transcription factor expressing cells by tissue type
% TF + cells
100
80
60
40
20
0
Sox
Lhx
Brach
Endoderm
Primitive Streak

Figure 6.4: Distribution of TF+ cells in the E6.5-7.5 mouse embryo. Images with thanks to Ita Costello A) Single plane and maximum projection images of Blimp1 expression in the Primitive Streak (PS) and Visceral Endoderm (VE) B) Using the CytoCensus ROI tools to restrict analysis to the PS C) All Blimp1+ cells detected by CytoCenus C') PS Blimp1+ cells, as detected by CytoCensus with ROI D) Expression of Brachury, Sox and Lhx D' CytoCensus detected TF+ cells in the PS E) Plot of percentage of TF+ cells in the PS (Sox 9/40 PS, Lhx 21/115, Brach 36/43) and VE (Sox 0/505, Lhx 759/1063, Brach 22/317) for each TF.

6.3.3. RESULTS

First, I determined whether CytoCensus can successfully quantitate the number of cells in the mouse blastocyst in two regions of interest: the visceral endoderm (VE) and the primitive streak (PS) (**Figure 6.4**A). Blimp1-Venus is expressed in the membranes of both the VE and PS cells (**Figure 6.4**A,C). Since cells in the VE and PS do not have sufficiently different Blimp1-Venus morphology to train separate models, I instead used CytoCensus to first identify all Blimp1+ cells (**Figure 6.4**A marked regions and B), and refined that to the structures of interest using a 3D interpolated ROI (see **Chapter 4**). Using the ROI tool in CytoCensus, therefore, successfully allows the quantitation of the spatially distinct VE and PS populations of cells.

Using these regions I then tested whether I could identify the subpopulation of these cells that were positive for 3 different transcription factors., I trained an additional model for each TF based on antibody staining (Piliszek et al., 2016). Based on a these antibody markers I could identify Blimp1+TF+ cells for Brachyury Lhx1 and Sox2, (**Figure 6.4**D,E). However, the number of these cells varies significantly between embryos, primarily due to the stage – I verified these differences by counting the total number of cells in each embryo using CytoCensus (Sox2 2358, Lhx1 8396, Brachyury 1588), which roughly correspond to the differences in length of the embryos (Sox2 494um, Lhx1 550um ,Brachyury 440um). In order to compensate for this I counted the number of cells in the PS and VE as a proportion of the total Blimp1+ population. CytoCensus successfully identifies the subpopulations that are Blimp1+TF+ in each case.

Using these models, I determined the proportion of TF+ cells VE and PS populations using CytoCensus cell centre output for each TF. From this it is clear that Brachyury positive cells are highly enriched in the PS, and almost completely absent from the endoderm (**Figure 6.4**E, F). This analysis showed the expected pattern, although, in some cases, some over-counting of neighbouring cells occurred (primarily in VE Brachyury+ case). In this case, Brachyury should not be present in the VE, although CytoCensus counted 7% Brach+ cells, likely due to the over-counting discussed in 6.1. I repeated this analysis for the transcription factors Sox2 and Lhx1 (**Figure 6.4**F). I conclude that Sox2+ cells are absent from the VE, with a small proportion present in the PS, although the bulk of Sox2+ cells are clearly in the mesoderm. Lhx1+ cells are present in similar proportion to Sox2+ cells within the PS, but present in most of the VE, and absent from the mesoderm. CytoCensus was used to successfully quantitate Brachyury, Sox2 and Lhx1 cells in the PS and VE.

6.3.4. Conclusion

To conclude, CytoCensus provided detailed quantification of cell numbers and types defined by expression of particular transcription factors in the PS and VE in mouse embryogenesis. Although not explicitly shown here, this has the potential to allow changes in transcription factor levels, and numbers of cells positive for these TFs to be followed throughout both normal and aberrant development. Quantification of cell types in this context should make it easier to analyse subtle changes in embryonic mutants, and by following transcription factors, perhaps to identify early changes in cells before more obvious behavioural changes are seen.

6.3.5. Limitations

CytoCensus struggles with closely packed cells separated only by membrane markers. The limitations of the membrane based markers have been documented within Section 6.1. Although it is difficult to adjust imaging conditions, in some cases, the results could be improved with higher-resolution imaging, reduced antibody background, and the use of more

specific cytoplasmic or nuclear markers (see **Section 6.2.4** for a discussion on the use of membrane specific markers with CytoCensus).

6.4. Organelles in spermatogenesis

6.4.1. Motivation

The Drosophila accessory gland is a key tissue involved in spermatogenesis. It is used as a model for the human prostate, as well as for studying the processes of exosome and vesicle secretion, and is crucial for the correct development of sperm (Wilson et al., 2017). One of the key secretory cell types in the accessory gland is the secondary cell (SC). The SC contains several classes of unusually large organelles, which are an excellent target for studying the secretion and endosomal processing. Ben Kroeger (Clive Wilson lab) has followed the behaviour of these organelles using live imaging, in both wild-type and mutant background, but largely relies on manual analysis and observation for insights into the changes. The organelles are, for the most part, highly spherical, but of variable size, both within WT, and between WT and in mutants such as Shrb RNAi.

Instead of counting whole cells, in this section I will focus on using CytoCensus to locate and count organelles within these images of SCs. As discussed in this section, and unlike in previous sections, identifying these organelles is difficult using CytoCensus with the standard set of Fiaschi *et al.* based features. Therefore, I also created a custom feature set based on the Fast Radial Symmetry Transform which allowed accurate identification of organelles in these images.

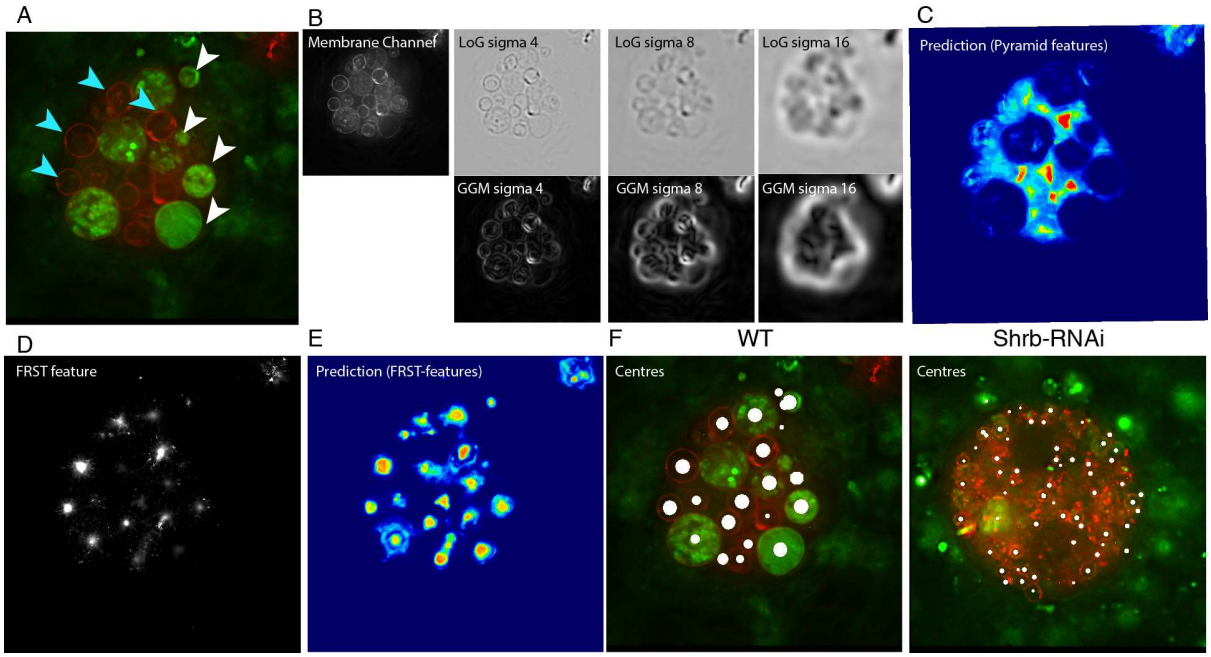

Figure 6.5: Finding Organelles with CytoCensus A) Example image of organelles marked by mGFP and Lysotracker Red. Acidic organelles (white arrows), other organelles (red arrows) B) Example features -Laplacian of Gaussian (LoG) and Gaussian Gradient Magnitude (GGM) at sigma=2,4,8 based on membrane channel C) Prediction map from CytoCensus trained using Pyramid feature set D) FRST based feature (unsmoothed) E) Prediction map using CytoCensus with FRST-features F) WT and Shrb-RNAi with CytoCensus FRST-f-based centres.

6.4.2. Aims

The objective was to count the organelles in these deconvolved widefield fluorescence images using CytoCensus (**Figure 6.5**A).

Given the large variation in size of the organelles, a secondary aim was to determine what efficient modifications would improve CytoCensus detection of these organelles.

6.4.3. Results

6.4.3.1. Organelle Identification

First I set out to determine whether the organelles of the SC were a suitable target for analysis by CytoCensus. Their highly spherical shape (**Figure 6.5**) is well suited to the assumptions of CytoCensus about object shape, although the use of a membrane outline makes for a more complex feature compared to an internal organelle marker. However, the main challenge is

that the significant differences in size make the objects difficult to identify using a combination of features at different scales. The features in CytoCensus are poorly suited to identify membrane bound organelles at a range of scales (**Figure 6.5**B), and even when feature is correlated with organelles of a particular size, it provides almost no information about organelles more than twice as large. Whilst it is possible, with considerable training, to make a model that detects most organelles (**Figure 6.5**), it is not especially robust, and it primarily detects organelles of a single size. For more reliable organelle detection it would be possible to train multiple models, each for a small size range of organelle. This approach, however, would be highly time consuming, somewhat defeating the point of the exercise. Similarly using a more complex system, for example a NN, that is capable of learning more complex decisions, could overcome this issue, but the requirement for much larger training datasets makes this prohibitive. This size-selectivity is partially by design: in many cases one feature that differentiates different object types is their size, and CytoCensus requires an object size estimate in order to work well. I conclude that using CytoCensus in its current form gives poor quality localisation of organelles.

To address this problem, I set out to identify an alternative feature set that was appropriate to identify organelles in the context of CytoCensus. To ensure efficient training (and evaluation), it was a prerequisite for the method to be fast (it should not take significantly longer that the existing features to calculate). Additionally, it must be calculated over a range of scales, to allow detection of the differently sized organelles. Ideally, it should also be straightforward to implement. I adapted the existing method by adding an additional feature set based on the Fast Radial Symmetry Transform (FRST), which was designed as a feature to detect small radially symmetric features such as eyes (Loy and Zelinsky, 2003) (**Figure 6.5**D). This modified set of features produced much more reliable prediction maps (**Figure 6.5**E) identifying organelles at a range of scales (**Figure 6.5**F) in both WT and RNAi conditions.

6.4.3.2. DETAILS OF FRST

First I explored possible features that were both efficient to compute and would respond strongly to organelles. I identified that the FRST was a good candidate for an additional feature set in CytoCensus because the organelles are highly symmetric. In the FRST paper (Loy and Zelinsky, 2003), they determine, in a single pass over the image for each radius of interest, the contribution that each pixel makes to the symmetry of the surrounding pixels.

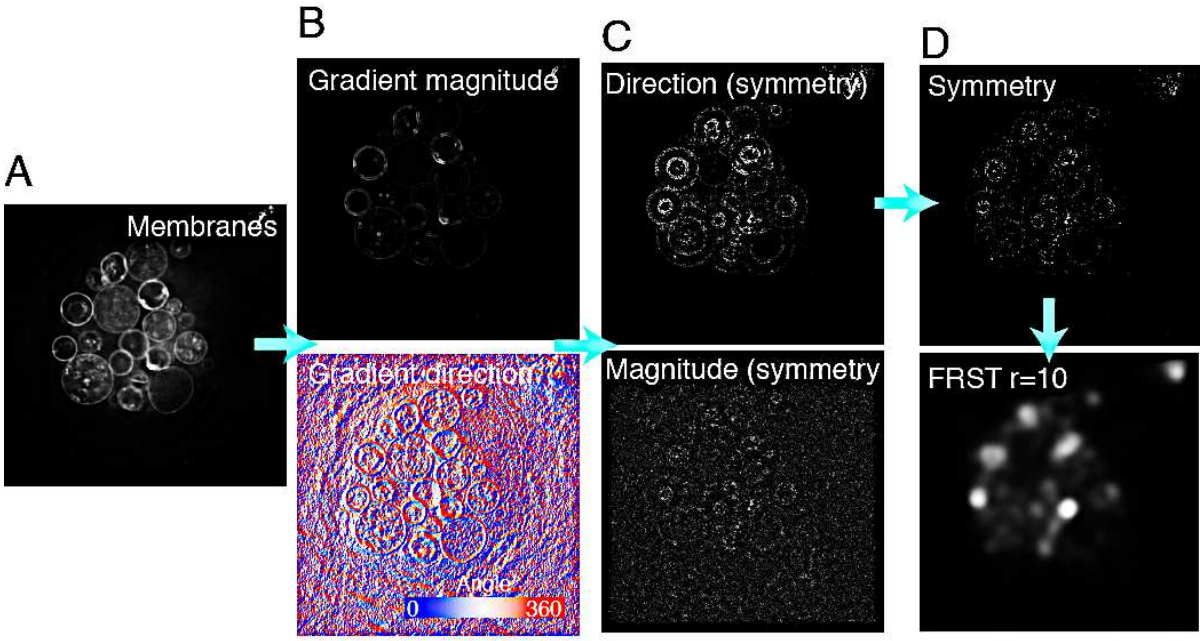

Figure 6.6: Calculation of FRST from Loy and Zelinsky, 2003 highlights some organelles A) Organelle membrane mage B) Magnitude (strength of gradient) and direction (colour represents the angle) images calculated from (A) C) Symmetry calculation components calculated at a radius of 10. The direction component shows rings corresponding to the organelles, but is independent on the strength of the gradient. The magnitude component is dependent only on the strength of the gradient, so is strongly affected by noise. D) The final symmetry transform combines the two components in (C), eliminating weak signals and those that are not symmetrical. The final FRST is calculated at a radius of 10 with an additional Gaussian smoothing sigma=10.

Briefly, gradients in the image are calculated over the image in the x and y directions, and from this the gradient magnitude and direction are computed (**Figure 6.6**B). Each pixel then adds a 'vote' at a given radius r, in the gradient direction (**Figure 6.6**C), and the gradient magnitude

of this vote is recorded (**Figure 6.6**C). When gradients converge, this is visible as a peak in this directional symmetry map (**Figure 6.6C**). From here, the symmetry map is squared, which acts as a factor to determine how strict the symmetry needs to be, and multiplied by the gradient map, to give the symmetry map in **Figure 6.6**D. Finally, this map is smoothed at radius r, to give the FRST in **Figure 6.6**C. The general form in Loy and Zelinsky is defined as:

$$p_{+ve} = p + round(\frac{g(p)}{||g(p)||})$$

$$p_{-ve} = p + round(\frac{g(p)}{||g(p)||})$$

$$O_r(p_{+ve}) = O_r(p_{+ve}) + 1$$

$$O_r(p_{-ve}) = O_r(p_{-ve}) - 1$$

$$M_r(p_{+ve}) = M_r(p_{+ve}) + ||g(p)||$$

$$M_r(p_{-ve}) = M_r(p_{-ve}) - ||g(p)||$$

$$S_r = ||O_n||^{\alpha} M_r * A_r$$

$$FRST = \sum_{r \in R} S_r$$

Equations 6.1: Definition of FRST. For a pixel p with gradient g(p) p_{+ve} is the positively affected pixel, and p_{-ve} is the negatively affected pixel. O_r is the orientation projection image, and M_n the magnitude projection image. A_r is a 2D Gaussian of radius r. Alpha is the radial strictness. S_r is the FRST at a given radius, and FRST is the sum over all radii R

This approach is much faster than the previous used approaches (Minor and Sklansky, 1981; Reisfeld et al., 1995) of considering the contribution of a region to a centre pixel, and scanning this region over the image. This makes the approach efficient O(KR) for an image of K pixels, a neighbourhood of RxR. In practice, R represents the radius of the object of interest. Hence in our case, with multiple objects (organelles) at different sizes, we need to perform this at multiple radii. In their paper, they use the 3 alternative radii (1,3,5) rather than the continuous set R = {r1,r2…rmax} increasing the efficiency further. Although this approach is useful, it was not originally used for microscopy images, or for use as a feature set. In this subsection, I will

detail the parameters and changes to the FRST detailed in Loy and Zelinsky that make it applicable to this challenge.

The original method returns a series of images representing the 'centres of symmetry' at different scales, smoothed to the appropriate scale (**Figure 6.7**A). I adapted this approach to be a better feature set, firstly by optimising parameters and secondly by combining these images into a single feature set representing a range of scales, and smoothing at a range of different scales (**Figure 6.7**B). I refer to this feature set as FRST-f and details of the parameter optimisation can be found in **Table 6.1**.

Table 6.1 Radii: The range FRST used was insufficient to capture the range of organelle sizes found in these images. As I a result I used set of radii N, with a step size of 3, but extending this to a maximum of 50 to represent a larger range of objects. This is an extension of their suggested 'fast' mode that uses a step size of 2 for increased speed. Alpha (radial strictness): in FRST-f is chosen to be 1.5 rather than 2. Due to the way that the radial strictness scales with the square of alpha, this provides an appropriate balance between detecting radial symmetry accurately and allowing for the significant variation of symmetry that is common within biological images. Gradient threshold: threshold for at which gradients are discarded is set to 1% of the maximum gradient rather than 2%. This avoids discarding gradients that might be important, but still reduces the number of points that are considered in my implementation, thus providing some speed up. Bright/dark mode: I use 'bright and dark' mode detailed in the paper. The bright mode, in effect, restricts the symmetry calculation to objects brighter than background, and vice-versa for dark mode. By using both, we make fewer assumptions about the objects of interest, and because we intend to use this as a feature, we can rely on the random forest to make choices over whether to include 'bright' or 'dark' objects, based on other features. Smoothing (sigma). This sets the multiplier for the radius of the Gaussian for smoothing at each radius.

Parameter	**FRST**	**FRST-f**
Radii range	Fast (1,3,5)	(1,4,7,10….50)
Alpha (radial strictness)	2	1.5
Gradient threshold %	2%	1%
Bright/dark mode	Bright	Bright and Dark
Smoothing (sigma)	1	0

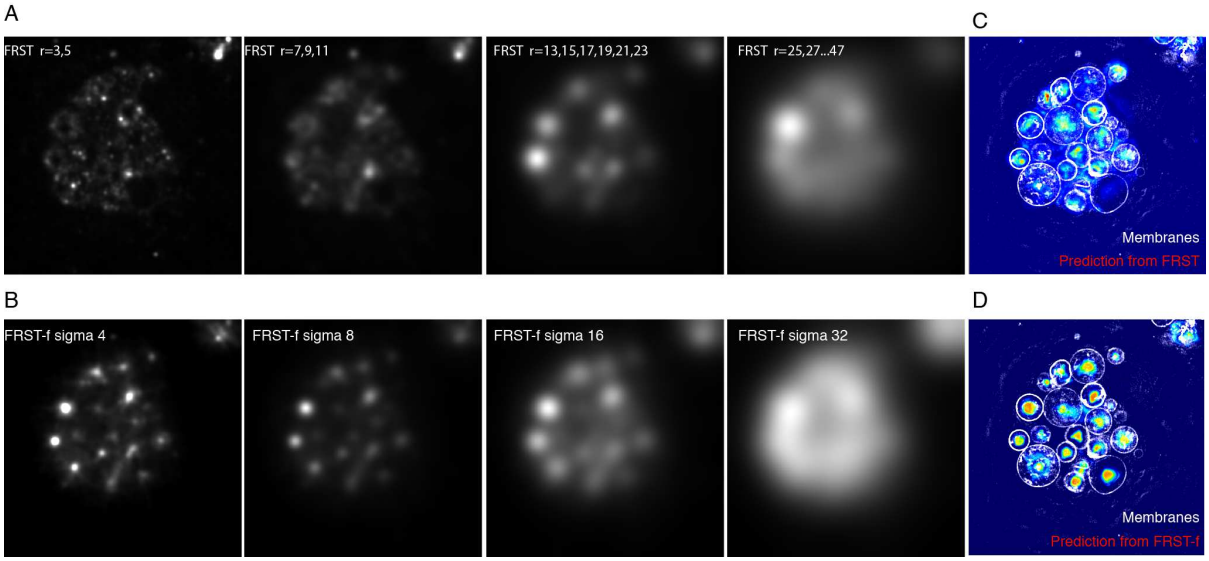

Figure 6.7: A modified FRST (FRST-f) makes a better feature set A) FRST calculated over multiple scales B) FRST-f at multiple scales C-D) Comparison of prediction maps from FRST and FRST-f using CytoCensus (heatmap), overlay membranes from Figure 6.5A (white). FRST-f results in much clearer peaks corresponding to organelle centres

In order to make FRST an appropriate feature set at a range of scales I used the existing pyramid structure within CytoCensus. Firstly, the radius of the final smoothing step in the FRST was set to 0 (**Table 6.1**). In the Loy and Zelinsky paper, they make use of Gaussian smoothing at each radius to spread information about the symmetry over an area that corresponds to the size of the object. This, therefore, requires multiple convolutions over the image when multiple radii are used, although they note in principle, a constant sized Gaussian kernel can be used for the same purpose. In fact, I found that these multiple convolutions are an unnecessary step. Nonetheless it is useful to spread the information about the radius over a range, so that the random forest can best make use of it as a feature. I used the existing pyramid structure (Chapter 4) of features, i.e. various Gaussian smoothed versions at different scales, to distribute information appropriately (**Figure 6.7**). This also ensured that the FRST features were appropriately represented in the random forest, and not dominated by other (less useful)

features. I found this to be a more effective feature and faster to calculate than the FRST over multiple sets of scales. This was then straightforward to incorporate into CytoCensus as an additional feature set, simply by supplying the method to calculate the feature, and adding it to the list of features. The prediction maps using CytoCensus with FRST and FRST-f are shown in **Figure 6.7**, and make it clear that FRST-f results in more robust detections (see **Figure 6.5**F for the resulting organelle detections).

To confirm that the FRST-f feature set was calculated in a reasonable time I compared it to FRST-f based features and to standard pyramid features. FRST-f in combination with the standard feature set did increase the time to calculate features, from 270ms to 1253ms, a significant increase which is not ideal for fast user feedback, and potentially problematic for large images, but still within a reasonable range, and significantly less than the 1950ms that FRST alone took. These results significantly improved the ability of CytoCensus to successfully detect organelles. The use of the modified FRST feature made training over a range of scales practical, detecting most organelles with minimal training (85%, **Figure 6.5**F).

6.4.4. Conclusions

To summarise, I defined an additional feature set based on the FRST that extends CytoCensus to cope with highly variable object sizes by relying on the radial symmetry of said objects at different scales. I successfully demonstrated its applicability to recognising organelles in SCs. This approach can be used to significantly speed up the process of determining differences in organelle numbers in mutants, for instance (as there can be hundreds of organelles). Although not shown here, in principle, organelle centres can be used as seeds for a marker based watershed approach to determine organelle size (see **Discussion**).

6.5. Discussion

I have shown in this Chapter that CytoCensus is useful for studying development in systems other than *Drosophila,* and can be used to quantify cell types using a range of different markers in multiple different tissues and organisms. I demonstrate that, in addition to cells, CytoCensus can be applied to subcellular organelles to quantify organelle number and location. In principle, this approach can be applied to objects that approximate the criteria of being round, and similar size in a range of different circumstances and at different scales. Using a modified version of the FRST I am also able to extend CytoCensus to relax the assumption of size, at least in the case of strong radial symmetry. Taken together this demonstrates CytoCensus is a broadly applicable tool for quantification of quantify cell and object numbers and locations in different systems.

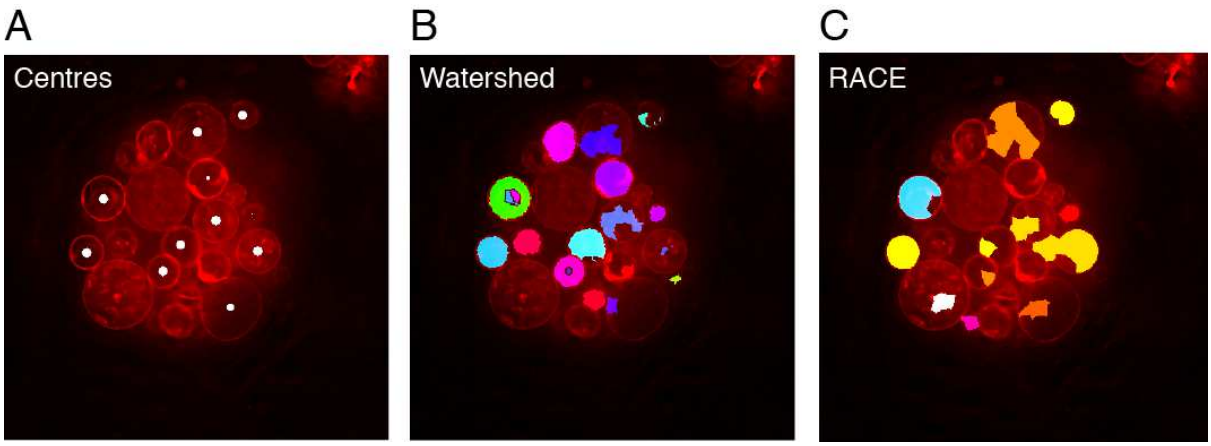

Figure 6.8: Illustration of possible methods for using cell centres to calculate object size A) Organelle image with Cell centres as calculated previously B) 'Catchment basins' from marker-based watershed to find C) Objects from RACE watershed (see **Chapter 3**)

For the tasks shown here, CytoCensus works well, however, it is worth noting that there are questions for which CytoCensus alone is not sufficient. Questions about cell contents, for instance, or cell shape and morphology are not approachable by CytoCensus alone because of the use of cell centres. However, it is possible to use CytoCensus as a pre-processing step to extract centres; for instance, Section 5.2 and 5.3 both mention the idea of using watershed as

a subsequent step. This allows 'seeds' to be grown from the CytoCensus centres until they encounter a barrier, or another seed. It is so named after the catchment basins for rivers separated by mountains and hills, and, by analogy, cell membranes from a suitable marker are the mountains, and by flooding from the cell centres until the waters meet, we find the line that represents the edge between cells. This allows us to turn centres into objects that represent cells. There are a number of other approaches that allow a similar kind of process to separate cells. RACE (Stegmaier et al., 2016) for instance, uses a special form of watershed with membrane enhancement to separate cells, and could be used to get more information about changes in cell morphology, shape and size than CytoCensus provides. An example of using RACE and marker-based watershed to find organelle boundaries is shown in **Figure 6.8**, with watershed outperforming RACE in this simple and imperfect example. These boundaries can in principle be used to ask questions about organelle content or changes in morphology, or detect merging compartments in live cell imaging.

In this chapter I have illustrated the use of CytoCensus on a range of examples, beyond the initial experiments on *Drosophila* in Chapter 5. I have shown that CytoCensus can be applied to provide cellular level quantitation of zebrafish organoids and to quantify the number of cells expressing different transcription factors in mouse. Where appropriate I have modified CytoCensus to enhance its general applicability. Furthermore, I have also shown that CytoCensus is not restricted to cells and can be applied at a subcellular scale, to locate and quantify organelles within *Drosophila* Secondary Cells in the accessory gland.

CHAPTER 7

DISCUSSION AND FUTURE WORK

In *Drosophila* and vertebrates alike, progress in understanding brain development has been limited by the lack of effective ways to culture and image whole brains for extended periods to monitor developmental behaviour, at the cellular level, in a native context. Here, I address this with the development of culture methods and high-resolution imaging of intact living explanted *Drosophila* larval brains over developmentally relevant time scales. In parallel, I developed CytoCensus, a convenient and rapid software, which employs a novel approach for cell detection based on supervised machine learning. CytoCensus was developed to identify neural stem cells and other cell types in 3D live imaging. I used this approach to quantitate cell numbers, distributions and rate of division at the individual cell level, in a range of 3D and 4D datasets including *Drosophila* brains and secretory glands, and mouse and zebrafish tissues. These uses of CytoCensus illustrate how this approach could be employed to enhance quantitative image analysis in a wide range of applications, including time-lapse movies of vertebrate tissues.

7.1. LIVE IMAGING

There were a number of key developments that contributed to this live imaging approach. Firstly, an optimised dissection and mounting protocol, with simplified culture medium extended the period of viability and mechanical stability of the specimens. Secondly, the use of very bright generic markers of cellular morphology, which offers advantages for tissue viability and phenotype analysis, allowed the use of low laser power to maximise viability.

Markers of general cell morphology can also be used in almost all mutant backgrounds, unlike specific markers of cell identity, whose expression is often altered in specific mutant backgrounds. These markers, therefore, offer a general approach for most model organisms. Lastly, the use of patch-based denoising combined with the aforementioned markers allowed unprecedented spatial and temporal live-imaging resolution with minimal cost to long term viability. I validated this culture and mounting method by following increases in brain size of *Drosophila* larvae, and by verifying normal NB and GMC division behaviour. Further work on live imaging and culture could seek to further optimise culturing conditions, as well as to examine the effects of additional factors, such as ecdysone, which is known to regulate brain development (Syed et al., 2017). Alternatively, the use of other fluorescent markers in live imaging, could be used to explore the dynamics of proteins in a tissue, for instance to quantify changes in protein levels. Imaging with other fluorescently tagged proteins such as Syncrip-GFP could allow the relationship between TF levels and the NB cell cycle to be explored in the context of brain development.

7.2. CytoCensus and Image Analysis

These long-term time-lapse experiments resulted in a huge volume of data, making manual image analysis impractical and certainly not scalable to include multiple conditions or genetic backgrounds (Peng et al., 2016). As microscopy methods improve, dealing with large scale data is likely to become increasingly common. The development of CytoCensus offers a generally applicable, automated analysis tool that can be applied to almost any equivalent kind of data set in any organism. While a variety of powerful image analysis tools already exist and are in common use for handling 3D and 4D datasets (Arganda-Carreras et al., 2017; Luengo et al., 2017; Sommer and Gerlich, 2013; Sommer et al., 2011; Stoter et al., 2013), CytoCensus offers a simple solution for fast and accurate cell detection. This enables analysis of subtle phenotypic changes without the need for bespoke programming and image analysis or intensive optimisation of many parameter sets. Furthermore, while machine learning approaches are

increasing in popularity, many that are currently freely available are limited to the analysis of 2D or 3D data (Held et al., 2010; Luengo et al., 2017). Moreover, existing conventional tools that handle 3D datasets are often inflexible, expensive, hard to use, or slow. I have demonstrated quantitatively that CytoCensus outperforms other freely available approaches. Therefore, developing supervised machine learning based cell detection represents a step change in the scale of datasets that can be easily and effectively analysed to obtain quantitative information on user defined cell types. CytoCensus directly provides simple information such as cell counts and locations. These can be processed to provide information on cell distributions and changes in cell behaviour such as mitosis, cell number, and movement. Crucially, CytoCensus analysis of extensive live imaging data allows such parameters to be determined for individual cells in a complex tissue, rather than conventional methods that provide snapshots or an ensemble view of average cell behaviour.

A technical limitation of the CytoCensus "point and click" strategy is that the program "assumes" a roughly spherical cell shape. This means that highly asymmetric cells, for instance those with cellular projections, e.g. neural extensions, would not be identified. Future work could examine methods to incorporate segmentation, which can identify these asymmetric cells, into the cell detection approach employed by CytoCensus. This would be analogous to the way that the human visual system has centres that separate different image processing types before combining them together. This, in turn, would make for a more general approach that can cope with both dense complex tissue as well as more complex objects and structures. A framework for an approach combining segmentation and detection has explored previously with some success (Glocker et al., 2012; Gould et al., 2009).

An alternative approach to improve detection performance would be to use more complex models for learning annotations. Neural networks are increasingly considered one of the most powerful ways to produce meaningful annotations of large scale datasets (e.g. ImageNet). In

the context of microscopy, their use has been limited, in part, by the need for extensive datasets. However, new state-of-the-art methods such as mixed-scale dense networks (Pelt and Sethian, 2018), can get impressive results with very little data. However, their use is currently very limited. Incorporating such a network into user-friendly tools such as CytoCensus could make for a powerful approach with which to analyse live imaging data.

7.3. The Role of Syncrip in the Brain

In the case of the *syp* RNAi phenotype, a previous analysis identified a weak under proliferation defect in NBs (Neumueller et al., 2007). I used live imaging and analysis to successfully characterise the cause of the enlarged brain phenotype in the Syncrip mutant to be due to an increase in NB division rate, thereby establishing the role of Syncrip in regulating the cell cycle of GMCs and NBs. Future work could examine the effect of other mutants on the cell cycle, and the role of such factors on other cell types, such as *brat* in the Type II NB, or in other regions of the brain such as the optic lobe. Additionally, this culture and imaging setup with existing markers would support a small scale RNAi screen for factors that influence division, and could be more sensitive to changes in NBs than previous screens (Neumueller et al., 2007). This setup could also be used to examine the effect of factors such as temperature on the development of the *Drosophila* brain. Furthermore, preliminary evidence suggests that the division rate of NBs slows over the course of larval development. The live imaging and CytoCensus analysis tools developed here, could be used to explore this observation further.

7.4. Tracking

The growing ability to image *ex vivo* or *in vivo* at cellular resolution is a powerful tool for understanding both the normal and aberrant development of organisms. However, it requires advances in image analysis to deal with data on this level. I present CytoCensus as a tool to answer questions about the locations and distributions of cells within complex tissues, and illustrate that CytoCensus can be used as an initial step for subsequent tracking of cells in live-

imaging data. This approach is generally applicable to live imaging data; here I apply it to follow GMCs over several hours. Specifically, CytoCensus is able to select a cell class of interest for tracking, thus reducing the burden on subsequent tracking software to track all cells within the tissue of interest. Furthermore, CytoCensus reduces false-positive identifications due to imaging artefacts or cell debris that may adversely affect tracking.. Although I have used a standard tracking approach (trackpy) in this work to allow tracking of cell types over time, incorporating more advanced methods, for instance a TGMM-based approach (Keller et al., 2008), which models cells as gaussians, or a joint tracking and detection approach (Akram et al., 2016), into CytoCensus has the potential to make both tracking and detection more robust.

7.5. Screening

Future applications of CytoCensus could include high throughput analysis of the effect of growth conditions on organoid development and drug discovery experiments using phenotypic characterisation of embryos or larvae rather than single cells (Walsh et al., 2017; Willoughby et al., 2013). Indeed, potential applications are not limited to large-scale quantitative screens: we envisage CytoCensus as a tool to assist everyday quantitation of imaging data. In this vein, CytoCensus is straightforward to install on standard operating systems, and can process reasonably large datasets even on a laptop.

7.6. Concluding Remarks

Live imaging approaches in *in vivo* settings and *ex vivo* models and organoids are increasingly being developed. These approaches have the potential to shed light on aspects of development that are otherwise difficult to study. This will aid our understanding of the importance of cell and tissue interactions, the natural variation within such tissues, and other highly dynamic processes. Here I show that live imaging, with suitable analysis, can be usefully applied even in complex situations, such as brain development in *Drosophila*, and use this to establish a role for Syncrip in the regulation of the cell cycle.

As such large live imaging becomes more commonplace, the need to develop tools to analyse such data is becoming increasingly necessary. Such tools can take a range of forms, from highly flexible programming pipelines, through to easy-to-use GUIs for those without specialist knowledge. In this context, CytoCensus adds another tool to the box, aiming specifically to bring benefits from supervised machine learning in order to make a user-friendly program for object detection.

REFERENCES

Ahmed, S. (2009). The culture of neural stem cells. J Cell Biochem *106*, 1-6.

Akram, S.U., Kannala, J., Eklund, L., Heikkilä, J., Gould, S., Gao, T., Koller, D., Glocker, B., Pauly, O., Konukoglu, E.*, et al.* (2016). Joint classification-regression forests for spatially structured multi-object segmentation. In Deep Learning and Data Labeling for Medical Applications (Springer), pp. 21-29.

Allan, C., Burel, J.-M., Moore, J., Blackburn, C., Linkert, M., Loynton, S., MacDonald, D., Moore, W.J., Neves, C., and Patterson, A. (2012). OMERO: flexible, model-driven data management for experimental biology. Nature Methods *9*, 245.

Allan, D.B., Caswell, T., Keim, N.C., and van der Wel, C.M. (2018). trackpy: Trackpy v0.4.1.

Almeida, A.D., Boije, H., Chow, R.W., He, J., Tham, J., Suzuki, S.C., and Harris, W.A. (2014). Spectrum of Fates: a new approach to the study of the developing zebrafish retina. Development *141*, 1971-1980.

Amat, F., Lemon, W., Mossing, D.P., McDole, K., Wan, Y., Branson, K., Myers, E.W., and Keller, P.J. (2014). Fast, accurate reconstruction of cell lineages from large-scale fluorescence microscopy data. Nat Methods *11*, 951-958.

Ammeux, N., Housden, B.E., Georgiadis, A., Hu, Y., and Perrimon, N. (2016). Mapping signaling pathway cross-talk in Drosophila cells. Proceedings of the National Academy of Sciences *113*, 9940-9945.

Arganda-Carreras, I., Kaynig, V., Rueden, C., Eliceiri, K.W., Schindelin, J., Cardona, A., and Sebastian Seung, H. (2017). Trainable Weka Segmentation: a machine learning tool for microscopy pixel classification. Bioinformatics *33*, 2424-2426.

Arnold, S.J., and Robertson, E.J. (2009). Making a commitment: cell lineage allocation and axis patterning in the early mouse embryo. Nature Reviews Molecular Cell Biology *10*, 91.

Artavanis-Tsakonas, S., and Simpson, P. (1991). Choosing a cell fate: a view from the Notch locus. Trends Genet *7*, 403-408.

Bainbridge, S.P., and Bownes, M. (1981). Staging the metamorphosis of Drosophila melanogaster. J Embryol Exp Morphol *66*, 57-80.

Barbosa, J.S., and Ninkovic, J. (2016). Adult neural stem cell behavior underlying constitutive and restorative neurogenesis in zebrafish. Neurogenesis (Austin) *3*, e1148101.

Bellen, H.J., Tong, C., and Tsuda, H. (2010). 100 years of Drosophila research and its impact on vertebrate neuroscience: a history lesson for the future. Nat Rev Neurosci *11*, 514-522.

Bello, B., Reichert, H., and Hirth, F. (2006). The brain tumor gene negatively regulates neural progenitor cell proliferation in the larval central brain of Drosophila. Development *133*, 2639-2648.

Berger, C., Harzer, H., Burkard, T., Steinmann, J., van der Horst, S., Laurenson, A.S., Novatchkova, M., Reichert, H., and Knoblich, J. (2012). FACS Purification and Transcriptome

Analysis of Drosophila Neural Stem Cells Reveals a Role for Klumpfuss in Self-Renewal. In Cell Rep, pp. 407-418.

Betschinger, J., Mechtler, K., and Knoblich, J.A. (2006). Asymmetric segregation of the tumor suppressor brat regulates self-renewal in Drosophila neural stem cells. Cell *124*, 1241-1253.

Boone, J.Q., and Doe, C.Q. (2008). Identification of Drosophila type II neuroblast lineages containing transit amplifying ganglion mother cells. Dev Neurobiol *68*, 1185-1195.

Boulanger, J., Kervrann, C., Bouthemy, P., Elbau, P., Sibarita, J.B., and Salamero, J. (2010). Patch-based nonlocal functional for denoising fluorescence microscopy image sequences. IEEE Trans Med Imaging *29*, 442-454.

Bowman, S.K., Rolland, V., Betschinger, J., Kinsey, K.A., Emery, G., and Knoblich, J.A. (2008). The tumor suppressors Brat and Numb regulate transit-amplifying neuroblast lineages in Drosophila. Dev Cell *14*, 535-546.

Brand, A.H., and Livesey, F.J. (2011). Neural stem cell biology in vertebrates and invertebrates: more alike than different? Neuron *70*, 719-729.

Cabernard, C., and Doe, C.Q. (2009). Apical/basal spindle orientation is required for neuroblast homeostasis and neuronal differentiation in Drosophila. Dev Cell *17*, 134-141.

Cabernard, C., and Doe, C.Q. (2013). Live imaging of neuroblast lineages within intact larval brains in Drosophila. Cold Spring Harb Protoc *2013*, 970-977.

Carlton, P.M., Boulanger, J., Kervrann, C., Sibarita, J.B., Salamero, J., Gordon-Messer, S., Bressan, D., Haber, J.E., Haase, S., Shao, L., *et al.* (2010). Fast live simultaneous multiwavelength four-dimensional optical microscopy. Proc Natl Acad Sci U S A *107*, 16016-16022.

Carney, T.D., Struck, A.J., and Doe, C.Q. (2013). midlife crisis encodes a conserved zinc-finger protein required to maintain neuronal differentiation in Drosophila. Development *140*, 4155-4164.

Carter, M. (2015). Guide to research techniques in neuroscience (Academic Press).

Ceron, J., Tejedor, F.J., and Moya, F. (2006). A primary cell culture of Drosophila postembryonic larval neuroblasts to study cell cycle and asymmetric division. Eur J Cell Biol *85*, 567-575.

Chazaud, C., Yamanaka, Y., Pawson, T., and Rossant, J. (2006). Early lineage segregation between epiblast and primitive endoderm in mouse blastocysts through the Grb2-MAPK pathway. Developmental Cell *10*, 615-624.

Chen, Q.W., and Agu, E. (2015). Exploring Statistical GLCM Texture Features for Classifying Food Images. In 2015 IEEE International Conference on Healthcare Informatics, P. Balakrishnan, J. Srivatsava, W.T. Fu, S. Harabagiu, and F. Wang, eds. (New York: IEEE), pp. 453-453.

Corrigan, L., Redhai, S., Leiblich, A., Fan, S.J., Perera, S.M., Patel, R., Gandy, C., Wainwright, S.M., Morris, J.F., Hamdy, F., *et al.* (2014). BMP-regulated exosomes from Drosophila male reproductive glands reprogram female behavior. J Cell Biol *206*, 671-688.

Cranfill, P.J., Sell, B.R., Baird, M.A., Allen, J.R., Lavagnino, Z., de Gruiter, H.M., Kremers, G.J., Davidson, M.W., Ustione, A., and Piston, D.W. (2016). Quantitative assessment of fluorescent proteins. Nat Methods *13*, 557-562.

Cruz, J.A., and Wishart, D.S. (2017). Applications of Machine Learning in Cancer Prediction and Prognosis:. Cancer informatics *2*.

Dang, K., and Sharma, S. (2017). Review and Comparison of Face Detection Algorithms. In Proceedings of the 7th International Conference on Cloud Computing, Data Science and Engineering.

Davis, I., Girdham, C.H., and O'Farrell, P.H. (1995). A nuclear GFP that marks nuclei in living Drosophila embryos; maternal supply overcomes a delay in the appearance of zygotic fluorescence. Dev Biol *170*, 726-729.

Davis, J., and Goadrich, M. (2006). The relationship between Precision-Recall and ROC curves. In Proceedings of the 23rd international conference on Machine learning (ACM), pp. 233-240.

De Luca, G.M., Breedijk, R.M., Brandt, R.A., Zeelenberg, C.H., de Jong, B.E., Timmermans, W., Azar, L.N., Hoebe, R.A., Stallinga, S., and Manders, E.M. (2013). Re-scan confocal microscopy: scanning twice for better resolution. In Biomed Opt Express, pp. 2644-2656.

de Medeiros, G., Norlin, N., Gunther, S., Albert, M., Panavaite, L., Fiuza, U.M., Peri, F., Hiiragi, T., Krzic, U., and Hufnagel, L. (2015). Confocal multiview light-sheet microscopy. Nat Commun *6*, 8881.

Dietz, C., and Berthold, M.R. (2016). KNIME for Open-Source Bioimage Analysis: A Tutorial. Adv Anat Embryol Cell Biol *219*, 179-197.

Dobbing, J., and Sands, J. (1973). Quantitative growth and development of human brain. Arch Dis Child *48*, 757-767.

Dong, C.Y., So, P.T., Mertz, J., Koenig, K., Xu, C., and Campagnola, P. (2013). Multiphoton microscopy: technical innovations, biological applications, and clinical diagnostics. J Biomed Opt *18*, 031101.

Dray, N., Bedu, S., Vuillemin, N., Alunni, A., Coolen, M., Krecsmarik, M., Supatto, W., Beaurepaire, E., and Bally-Cuif, L. (2015). Large-scale live imaging of adult neural stem cells in their endogenous niche. Development *142*, 3592-3600.

Dufour, A.C., Jonker, A.H., and Olivo-Marin, J.C. (2017). Deciphering tissue morphodynamics using bioimage informatics. Philos Trans R Soc Lond B Biol Sci *372*.

Egger, B., van Giesen, L., Moraru, M., and Sprecher, S.G. (2013). In vitro imaging of primary neural cell culture from Drosophila. Nat Protoc *8*, 958-965.

Eldred, M.K., Charlton-Perkins, M., Muresan, L., and Harris, W.A. (2017a). Self-organising aggregates of zebrafish retinal cells for investigating mechanisms of neural lamination. In Development, pp. 1097-1106.

Eldred, M.K., Muresan, L., and Harris, W.A. (2017b). Disaggregation and Reaggregation of Zebrafish Retinal Cells for the Analysis of Neuronal Layering. Methods Mol Biol.

Eliceiri, K.W., Berthold, M.R., Goldberg, I.G., Ibanez, L., Manjunath, B.S., Martone, M.E., Murphy, R.F., Peng, H., Plant, A.L., Roysam, B.*, et al.* (2012). Biological imaging software tools. Nat Methods *9*, 697-710.

Ettinger, A., and Wittmann, T. (2014). Fluorescence live cell imaging. Methods Cell Biol *123*, 77-94.

Faul, F., Erdfelder, E., Lang, A.G., and Buchner, A. (2007). G*Power 3: a flexible statistical power analysis program for the social, behavioral, and biomedical sciences. Behav Res Methods *39*, 175-191.

Fiaschi, L., Nair, R., Koethe, U., and Hamprecht, F.A. (2012). Learning to count with regression forest and structured labels. In Pattern Recognition (ICPR), 2012 21st International Conference on (IEEE), pp. 2685-2688.

Forero, M.G., Learte, A.R., Cartwright, S., and Hidalgo, A. (2010a). DeadEasy Mito-Glia: automatic counting of mitotic cells and glial cells in Drosophila. Plos One *5*, e10557.

Forero, M.G., Pennack, J.A., and Hidalgo, A. (2010b). DeadEasy neurons: automatic counting of HB9 neuronal nuclei in Drosophila. Cytometry A *77*, 371-378.

Forero, M.G., Pennack, J.A., Learte, A.R., and Hidalgo, A. (2009). DeadEasy caspase: automatic counting of apoptotic cells in Drosophila. Plos One *4*, e5441.

Furst, A., and Mahowald, A.P. (1985). Differentiation of primary embryonic neuroblasts in purified neural cell cultures from Drosophila. Dev Biol *109*, 184-192.

Glocker, B., Pauly, O., Konukoglu, E., and Criminisi, A. (2012). Joint classification-regression forests for spatially structured multi-object segmentation. In European Conference on Computer Vision (Springer), pp. 870-881.

Goldberg, I.G., Allan, C., Burel, J.M., Creager, D., Falconi, A., Hochheiser, H., Johnston, J., Mellen, J., Sorger, P.K., and Swedlow, J.R. (2005). The Open Microscopy Environment (OME) Data Model and XML file: open tools for informatics and quantitative analysis in biological imaging. Genome Biol *6*, R47.

Goldman, R.D., Swedlow, J.R., and Spector, D.L. (2010). Live cell imaging: a laboratory manual (Cold Spring Harbor Laboratory Press).

Gonzalez, R., and Woods, R. (2002). Digital Image Processing. Addison-Wesley Publishing Company.

Gottschalk, S., Estrada, H., Degtyaruk, O., Rebling, J., Klymenko, O., Rosemann, M., and Razansky, D. (2015). Short and long-term phototoxicity in cells expressing genetic reporters under nanosecond laser exposure. Biomaterials *69*, 38-44.

Gould, S., Gao, T., and Koller, D. (2009). Region-based segmentation and object detection. In Advances in neural information processing systems, pp. 655-663.

Graeden, E., and Sive, H. (2009). Live imaging of the zebrafish embryonic brain by confocal microscopy. J Vis Exp.

Gregg, C.L., and Butcher, J.T. (2012). Quantitative in vivo imaging of embryonic development: opportunities and challenges. Differentiation *84*, 149-162.

Griffa, A., Garin, N., and Sage, D. (2018). Comparison of Deconvolution Software in 3D Microscopy: A User Point of View—Part 1. GIT Imaging & Microscopy *12*, 43–45.

Gris, K.V., Coutu, J.P., and Gris, D. (2017). Supervised and Unsupervised Learning Technology in the Study of Rodent Behavior. Front Behav Neurosci *11*, 141.

Gustafsson, M.G. (2000). Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy. J Microsc *198*, 82-87.

Hailstone, M., Yang, L., Waithe, D., Samuels, T.J., Arava, Y., Dobrzycki, T., Parton, R.M., and Davis, I. (2018). Brain Development: Machine Learning Analysis Of Individual Stem Cells In Live 3D Tissue. bioRxiv, 137406.

Halstead, J.M., Lin, Y.Q., Durraine, L., Hamilton, R.S., Ball, G., Neely, G.G., Bellen, H.J., and Davis, I. (2014). Syncrip/hnRNP Q influences synaptic transmission and regulates BMP signaling at the Drosophila neuromuscular synapse. Biology open *3*, 839-849.

Harzer, H., Berger, C., Conder, R., Schmauss, G., and Knoblich, J.A. (2013). FACS purification of Drosophila larval neuroblasts for next-generation sequencing. Nat Protoc *8*, 1088-1099.

Held, M., Santeramo, I., Wilm, B., Murray, P., and Levy, R. (2018). Ex vivo live cell tracking in kidney organoids using light sheet fluorescence microscopy. Plos One *13*, e0199918.

Held, M., Schmitz, M.H., Fischer, B., Walter, T., Neumann, B., Olma, M.H., Peter, M., Ellenberg, J., and Gerlich, D.W. (2010). CellCognition: time-resolved phenotype annotation in high-throughput live cell imaging. Nat Methods *7*, 747-754.

Heller, D., Hoppe, A., Restrepo, S., Gatti, L., Tournier, A.L., Tapon, N., Basler, K., and Mao, Y. (2016). EpiTools: An Open-Source Image Analysis Toolkit for Quantifying Epithelial Growth Dynamics. Dev Cell *36*, 103-116.

Henrique, D., and Bally-Cuif, L. (2010). A cross-disciplinary approach to understanding neural stem cells in development and disease. In Development (England), pp. 1933-1938.

Hodneland, E., Kogel, T., Frei, D.M., Gerdes, H.H., and Lundervold, A. (2013). CellSegm - a MATLAB toolbox for high-throughput 3D cell segmentation. Source Code Biol Med *8*, 16.

Homem, C.C., and Knoblich, J.A. (2012). Drosophila neuroblasts: a model for stem cell biology. Development *139*, 4297-4310.

Homem, C.C., Reichardt, I., Berger, C., Lendl, T., and Knoblich, J.A. (2013). Long-term live cell imaging and automated 4D analysis of drosophila neuroblast lineages. Plos One *8*, e79588.

Hooke, R. (1968). Micrographia (Allestry).

Huang, C., Maxey, J.R., Sinha, S., Savall, J., Gong, Y., and Schnitzer, M.J. (2018). Long-term optical brain imaging in live adult fruit flies. Nat Commun *9*, 872.

Huisken, J., and Stainier, D.Y. (2009). Selective plane illumination microscopy techniques in developmental biology. Development *136*, 1963-1975.

Hunter, J.D. (2007). Matplotlib: A 2D graphics environment. Computing in science

engineering *9*, 90-95.

Irshad, H., Veillard, A., Roux, L., and Racoceanu, D. (2014). Methods for nuclei detection, segmentation, and classification in digital histopathology: a review-current status and future potential. IEEE Rev Biomed Eng *7*, 97-114.

Jabr, F. (2018). The Connectome Debate: Is Mapping the Mind of a Worm Worth It? Scientific American.

Johnson, G.R., Donovan-Maiye, R.M., and Maleckar, M.M. (2017). Generative Modeling with Conditional Autoencoders: Building an Integrated Cell. arXiv.

Jones, E., Oliphant, T., and Peterson, P. (2014). SciPy: open source scientific tools for Python.

Kainz, P., Urschler, M., Schulter, S., Wohlhart, P., and Lepetit, V. (2015). You Should Use Regression to Detect Cells. In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, N. Navab, J. Hornegger, W.M. Wells, and A.F. Frangi, eds. (Springer International Publishing), pp. 276-283.

Kamentsky, L., Jones, T.R., Fraser, A., Bray, M.A., Logan, D.J., Madden, K.L., Ljosa, V., Rueden, C., Eliceiri, K.W., and Carpenter, A.E. (2011). Improved structure, function and compatibility for CellProfiler: modular high-throughput image analysis software. Bioinformatics *27*, 1179-1180.

Kang, J., Schwartz, R., Flickinger, J., and Beriwal, S. (2015). Machine Learning Approaches for Predicting Radiation Therapy Outcomes: A Clinician's Perspective. Int J Radiat Oncol Biol Phys *93*, 1127-1135.

Kang, W.-X., Yang, Q.-Q., and Liang, R.-P. (2009). The comparative research on image segmentation algorithms. In Education Technology and Computer Science, 2009 (IEEE), pp. 703-707.

Keil, W., Kutscher, L.M., Shaham, S., and Siggia, E.D. (2017). Long-Term High-Resolution Imaging of Developing C. elegans Larvae with Microfluidics. Dev Cell *40*, 202-214.

Keller, P.J., Schmidt, A.D., Wittbrodt, J., and Stelzer, E.H. (2008). Reconstruction of zebrafish early embryonic development by scanned light sheet microscopy. Science *322*, 1065-1069.

Keller, P.J., Schmidt, A.D., Wittbrodt, J., and Stelzer, E.H. (2011). Digital scanned laser light-sheet fluorescence microscopy (DSLM) of zebrafish and Drosophila embryonic development. Cold Spring Harb Protoc *2011*, 1235-1243.

Kervrann, C., Legland, D., and Pardini, L. (2004). Robust incremental compensation of the light attenuation with depth in 3D fluorescence microscopy. J Microsc *214*, 297-314.

Kim, B., and Naemura, T. (2015). Blind Depth-variant Deconvolution of 3D Data in Wide-field Fluorescence Microscopy. Sci Rep *5*, 9894.

Kim, K., Park, H., and Lim, K.M. (2015). Phototoxicity: Its Mechanism and Animal Alternative Test Methods. In Toxicol Res, pp. 97-104.

Kohwi, M., and Doe, C.Q. (2013). Temporal fate specification and neural progenitor competence during development. Nat Rev Neurosci *14*, 823-838.

Kotsiantis, S.B., Zaharakis, I., and Pintelas, P. (2007). Supervised machine learning: A review of classification techniques. Emerging artificial intelligence applications in computer engineering *160*, 3-24.

Kruger, A.V., Jelier, R., Dzyubachyk, O., Zimmerman, T., Meijering, E., and Lehner, B. (2015). Comprehensive single cell-resolution analysis of the role of chromatin regulators in early C. elegans embryogenesis. Dev Biol *398*, 153-162.

Kyriacou, S.K., Davatzikos, C., Zinreich, S.J., and Bryan, R.N. (1999). Nonlinear elastic registration of brain images with tumor pathology using a biomechanical model. IEEE Trans Med Imaging *18*, 580-592.

Lee, C.Y., Wilkinson, B.D., Siegrist, S.E., Wharton, R.P., and Doe, C.Q. (2006). Brat is a Miranda cargo protein that promotes neuronal differentiation and inhibits neuroblast self-renewal. Dev Cell *10*, 441-449.

Lee, T., and Luo, L. (2001). Mosaic analysis with a repressible cell marker (MARCM) for Drosophila neural development. Trends Neurosci *24*, 251-254.

Lempitsky, V., and Zisserman, A. (2018). Learning To Count Objects in Images. In Advances in Neural Information Processing Systems, pp. 1324-1332.

Lerit, D.A., Plevock, K.M., and Rusan, N.M. (2014). Live imaging of Drosophila larval neuroblasts. J Vis Exp 89.

Li, S., Wang, H., and Groth, C. (2014). Drosophila neuroblasts as a new model for the study of stem cell self-renewal and tumour formation. Biosci Rep *34*.

Libbrecht, M.W., and Noble, W.S. (2015). Machine learning applications in genetics and genomics. Nat Rev Genet *16*, 321-332.

Linkert, M., Rueden, C.T., Allan, C., Burel, J.M., Moore, W., Patterson, A., Loranger, B., Moore, J., Neves, C., Macdonald, D., *et al.* (2010). Metadata matters: access to image data in the real world. J Cell Biol *189*, 777-782.

Liu, Z., Yang, C.P., Sugino, K., Fu, C.C., Liu, L.Y., Yao, X., Lee, L.P., and Lee, T. (2015). Opposing intrinsic temporal gradients guide neural stem cell production of varied neuronal fates. Science *350*, 317-320.

Lou, X., Kang, M., Xenopoulos, P., Munoz-Descalzo, S., and Hadjantonakis, A.K. (2014). A rapid and efficient 2D/3D nuclear segmentation method for analysis of early mouse embryo and stem cell image data. Stem Cell Reports *2*, 382-397.

Lowe, D.G. (1999). Object recognition from local scale-invariant features. In Proceedings of the Seventh IEEE International Conference on Computer Vision, pp. 1150-1157 vol.1152.

Loy, G., and Zelinsky, A. (2003). Fast radial symmetry for detecting points of interest. IEEE Transactions on Pattern Analysis, 959-973.

Luengo, I., Darrow, M.C., Spink, M.C., Sun, Y., Dai, W., He, C.Y., Chiu, W., Pridmore, T., Ashton, A.W., Duke, E.M.H., *et al.* (2017). SuRVoS: Super-Region Volume Segmentation workbench. J Struct Biol *198*, 43-53.

Maroni, G., and Stamey, S. (1983). Use of blue food to select synchronous, late third instar larvae. Dros Inf Serv *59*, 142-143.

McDermott, S.M., Yang, L., Halstead, J.M., Hamilton, R.S., Meignin, C., and Davis, I. (2014). Drosophila Syncrip modulates the expression of mRNAs encoding key synaptic proteins required for morphology at the neuromuscular junction. Rna *20*, 1593-1606.

McNally, J.G., Karpova, T., Cooper, J., and Conchello, J.A. (1999). Three-dimensional imaging by deconvolution microscopy. Methods *19*, 373-385.

Medioni, C., Ephrussi, A., and Besse, F. (2015). Live imaging of axonal transport in Drosophila pupal brain explants. Nat Protoc *10*, 574-584.

Meiniel, W., Olivo-Marin, J.C., and Angelini, E.D. (2018). Denoising of Microscopy Images: A Review of the State-of-the-Art, and a New Sparsity-Based Method. IEEE Trans Image Process *27*, 3842-3856.

Mikut, R., Dickmeis, T., Driever, W., Geurts, P., Hamprecht, F.A., Kausler, B.X., Ledesma-Carbayo, M.J., Maree, R., Mikula, K., Pantazis, P., *et al.* (2013). Automated processing of zebrafish imaging data: a survey. Zebrafish *10*, 401-421.

Miles, W.O., Dyson, N.J., and Walker, J.A. (2011). Modeling tumor invasion and metastasis in Drosophila. Dis Model Mech *4*, 753-761.

Minor, L.G., and Sklansky, J. (1981). The detection and segmentation of blobs in infrared images. IEEE Transactions on Systems, Man, and Cybernetics *11*, 194-201.

Mohanaiah, P., Sathyanarayana, P., and GuruKumar, L. (2013). Image texture feature extraction using GLCM approach. International Journal of Scientific and Research Publications *3*, 1.

Moraru, M.M., Egger, B., Bao, D.B., and Sprecher, S.G. (2012). Analysis of cell identity, morphology, apoptosis and mitotic activity in a primary neural cell culture system in Drosophila. Neural Dev *7*, 14.

Muotri, A.R., and Gage, F.H. (2006). Generation of neuronal variability and complexity. Nature *441*, 1087-1093.

Neumueller, R., Poernbacher, I., and Knoblich, J.A. (2007). The TRIM-NHL proteins Brat and Mei-P26 regulate cell growth and proliferation in Drosophila stem cell lineages. Febs Journal *274*, 8-8.

Neumuller, R.A., Richter, C., Fischer, A., Novatchkova, M., Neumuller, K.G., and Knoblich, J.A. (2011). Genome-wide analysis of self-renewal in Drosophila neural stem cells by transgenic RNAi. Cell Stem Cell *8*, 580-593.

Noctor, S.C., Martinez-Cerdeno, V., Ivic, L., and Kriegstein, A.R. (2004). Cortical neurons arise in symmetric and asymmetric division zones and migrate through specific phases. Nat Neurosci *7*, 136-144.

Pampaloni, F., Reynaud, E.G., and Stelzer, E.H. (2007). The third dimension bridges the gap between cell culture and live tissue. Nat Rev Mol Cell Biol *8*, 839-845.

Pantazis, P., and Supatto, W. (2014). Advances in whole-embryo imaging: a quantitative transition is underway. Nat Rev Mol Cell Biol *15*, 327-339.

Pärnamaa, T., and Parts, L. (2017). Accurate Classification of Protein Subcellular Localization from High-Throughput Microscopy Images Using Deep Learning. In G3 (Bethesda), pp. 1385-1392.

Parslow, A., Cardona, A., and Bryson-Richardson, R.J. (2014). Sample drift correction following 4D confocal time-lapse imaging. J Vis Exp 86.

Parton, R.M., Valles, A.M., Dobbie, I.M., and Davis, I. (2010). Live cell imaging in Drosophila melanogaster. Cold Spring Harb Protoc *2010*, pdb top75.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., and Dubourg, V. (2011). Scikit-learn: Machine learning in Python. Journal of machine learning research *12*, 2825-2830.

Pelt, D.M., and Sethian, J.A. (2018). A mixed-scale dense convolutional neural network for image analysis. Proc Natl Acad Sci U S A *115*, 254-259.

Peng, H., Zhou, J., Zhou, Z., Bria, A., Li, Y., Kleissas, D.M., Drenkow, N.G., Long, B., Liu, X., and Chen, H. (2016). Bioimage Informatics for Big Data. Adv Anat Embryol Cell Biol *219*, 263-272.

Piccinini, F., Balassa, T., Szkalisity, A., Molnar, C., Paavolainen, L., Kujala, K., Buzas, K., Sarazova, M., Pietiainen, V., Kutay, U., *et al.* (2017). Advanced Cell Classifier: User-Friendly Machine-Learning-Based Software for Discovering Phenotypes in High-Content Imaging Data. Cell Syst *4*, 651-655 e655.

Piliszek, A., Grabarek, J.B., Frankenberg, S.R., and Plusa, B. (2016). Cell fate in animal and human blastocysts and the determination of viability. MHR: Basic science of reproductive medicine *22*, 681-690.

Pilz, G.A., Bottes, S., Betizeau, M., Jorg, D.J., Carta, S., Simons, B.D., Helmchen, F., and Jessberger, S. (2018). Live imaging of neurogenesis in the adult mouse hippocampus. Science *359*, 658-662.

Pinto-Teixeira, F., Konstantinides, N., and Desplan, C. (2016). Programmed cell death acts at different stages of Drosophila neurodevelopment to shape the central nervous system. FEBS Lett *590*, 2435-2453.

Pitrone, P.G., Schindelin, J., Stuyvenberg, L., Preibisch, S., Weber, M., Eliceiri, K.W., Huisken, J., and Tomancak, P. (2013). OpenSPIM: an open-access light-sheet microscopy platform. Nat Methods *10*, 598-599.

Pnevmatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., *et al.* (2016). Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. Neuron *89*, 285-299.

Preibisch, S., Amat, F., Stamataki, E., Sarov, M., Singer, R.H., Myers, E., and Tomancak, P. (2014). Efficient Bayesian-based multiview deconvolution. Nat Methods *11*, 645-648.

Prithviraj, R., Trunova, S., and Giniger, E. (2012). Ex vivo Culturing of Whole, Developing Drosophila Brains. J Vis Exp.

Rabinovich, D., Mayseless, O., and Schuldiner, O. (2015). Long term ex vivo culturing of Drosophila brain as a method to live image pupal brains: insights into the cellular mechanisms of neuronal remodeling. Front Cell Neurosci *9*, 327.

Reichert, H. (2011). Drosophila neural stem cells: cell cycle control of self-renewal, differentiation, and termination in brain development. Results Probl Cell Differ *53*, 529-546.

Reisfeld, D., Wolfson, H., and Yeshurun, Y. (1995). Context-free attentional operators: the generalized symmetry transform. International Journal of Computer Vision *14*, 119-130.

Ren, Q., Yang, C.P., Liu, Z., Sugino, K., Mok, K., He, Y., Ito, M., Nern, A., Otsuna, H., and Lee, T. (2017). Stem Cell-Intrinsic, Seven-up-Triggered Temporal Factor Gradients Diversify Intermediate Neural Progenitors. Curr Biol *27*, 1303-1313.

Richardson, D.S., and Lichtman, J.W. (2015). Clarifying Tissue Clearing. Cell *162*, 246-257.

Riddiford, L.M. (1993). Hormone receptors and the regulation of insect metamorphosis. Receptor *3*, 203-209.

Rigaut, J.P., and Vassy, J. (1991). High-resolution three-dimensional images from confocal scanning laser microscopy. Quantitative study and mathematical correction of the effects from bleaching and fluorescence attenuation in depth. Anal Quant Cytol Histol *13*, 223-232.

Rittscher, J. (2010). Characterization of biological processes through automated image analysis. Annu Rev Biomed Eng *12*, 315-344.

Roysam, B., Shain, W., Robey, E., Chen, Y., Narayanaswamy, A., Tsai, C.L., Al-Kofahi, Y., Bjornsson, C., Ladi, E., and Herzmark, P. (2008). The FARSIGHT Project: Associative 4D/5D Image Analysis Methods for Quantifying Complex and Dynamic Biological Microenvironments. Microscopy and Microanalysis *14*, 60-61.

Rubens, U., Hoyoux, R., Vanosmael, L., Ouras, M., Tasset, M., Hamilton, C., Longuespée, R., and Marée, R. (2018). Cytomine: Toward an Open and Collaborative Software Platform for Digital Pathology Bridged to Molecular Investigations. PROTEOMICS–Clinical Applications, 1800057.

Rudrapatna, V.A., Cagan, R.L., and Das, T.K. (2012). Drosophila cancer models. Dev Dyn *241*, 107-118.

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., *et al.* (2014). ImageNet Large Scale Visual Recognition Challenge. arXiv.

Sabado, V., and Nagoshi, E. (2018). Single-cell Resolution Fluorescence Live Imaging of Drosophila Circadian Clocks in Larval Brain Culture. J Vis Exp.

Saiz, N., Plusa, B., and Hadjantonakis, A.-K. (2015). Single cells get together: High-resolution approaches to study the dynamics of early mouse development. Paper presented at: Seminars in cell & developmental biology (Elsevier).

Sakaue-Sawano, A., Kurokawa, H., Morimura, T., Hanyu, A., Hama, H., Osawa, H., Kashiwagi, S., Fukami, K., Miyata, T., Miyoshi, H., *et al.* (2008). Visualizing spatiotemporal dynamics of multicellular cell-cycle progression. Cell *132*, 487-498.

Sambrook, J., Fritsch, E.F., and Maniatis, T. (1989). Molecular cloning: a laboratory manual (Cold spring harbor laboratory press).

Samuel, A.L. (1959). Some Studies in Machine Learning Using the Game of Checkers. IBM Journal of Research and Development *3*, 210-229.

Savoian, M.S., and Rieder, C.L. (2002). Mitosis in primary cultures of Drosophila melanogaster larval neuroblasts. J Cell Sci *115*, 3061-3072.

Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., Preibisch, S., Rueden, C., Saalfeld, S., Schmid, B., *et al.* (2012). Fiji: an open-source platform for biological-image analysis. Nat Methods *9*, 676-682.

Schmitz, C., Eastwood, B.S., Tappan, S.J., Glaser, J.R., Peterson, D.A., and Hof, P.R. (2014). Current automated 3D cell detection methods are not a suitable replacement for manual stereologic cell counting. Front Neuroanat *8*, 27.

Shuvaev, S.A., Lazutkin, A.A., Kedrov, A.V., Anokhin, K.V., Enikolopov, G.N., and Koulakov, A.A. (2017). DALMATIAN: An Algorithm for Automatic Cell Detection and Counting in 3D. Front Neuroanat *11*, 117.

Sidman, R.L., and Rakic, P. (1973). Neuronal migration, with special reference to developing human brain: a review. Brain Res *62*, 1-35.

Simian, M., and Bissell, M.J. (2017). Organoids: A historical perspective of thinking in three dimensions. In J Cell Biol, pp. 31-40.

Simon, C.S., Downes, D.J., Gosden, M.E., Telenius, J., Higgs, D.R., Hughes, J.R., Costello, I., Bikoff, E.K., and Robertson, E.J. (2017). Functional characterisation of cis-regulatory elements governing dynamic Eomes expression in the early mouse embryo. Development *144*, 1249-1260.

Skeath, J.B., and Thor, S. (2003). Genetic control of Drosophila nerve cord development. Curr Opin Neurobiol *13*, 8-15.

Sommer, C., and Gerlich, D.W. (2013). Machine learning in cell biology - teaching computers to recognize phenotypes. J Cell Sci *126*, 5529-5539.

Sommer, C., Straehle, C., Koethe, U., and Hamprecht, F.A. (2011). ILASTIK: Interactive Learning And Segmentation Toolkit. 2011 8th Ieee International Symposium on Biomedical Imaging: from Nano to Macro, 230-233.

Soulez, F., Denis, L., Tourneur, Y., and Thiébaut, É. (2012). Blind deconvolution of 3D data in wide field fluorescence microscopy. In Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on (IEEE), pp. 1735-1738.

Stegmaier, J., Amat, F., Lemon, W.C., McDole, K., Wan, Y., Teodoro, G., Mikut, R., and Keller, P.J. (2016). Real-Time Three-Dimensional Cell Segmentation in Large-Scale Microscopy Data of Developing Embryos. Dev Cell *36*, 225-240.

Stephens, D.J. (2003). Light Microscopy Techniques for Live Cell Imaging. Science *300*, 82-86.

Stephenson, R., and Metcalfe, N.H. (2013). Drosophila melanogaster: a fly through its history and current use. J R Coll Physicians Edinb *43*, 70-75.

Stiles, J., and Jernigan, T.L. (2010). The Basics of Brain Development. In Neuropsychol Rev, pp. 327-348.

Stoter, M., Niederlein, A., Barsacchi, R., Meyenhofer, F., Brandl, H., and Bickle, M. (2013). CellProfiler and KNIME: open source tools for high content screening. Methods Mol Biol *986*, 105-122.

Sullivan, D.P., and Lundberg, E. (2018). Seeing More: A Future of Augmented Microscopy. Cell *173*, 546-548.

Sundararajan, S., Wakamiya, M., Behringer, R.R., and Rivera-Pérez, J.A. (2012). A fast and sensitive alternative for β-galactosidase detection in mouse embryos. Development *139*, 4484-4490.

Svoboda, D., Homola, O., and Stejskal, S. (2011). Generation of 3D digital phantoms of colon tissue. In International Conference Image Analysis and Recognition (Springer), pp. 31-39.

Svoboda, D., Kozubek, M., and Stejskal, S. (2009). Generation of digital phantoms of cell nuclei and simulation of image formation in 3D image cytometry. Cytometry Part A: The Journal of the International Society for Advancement of Cytometry *75*, 494-509.

Syed, M.H., Mark, B., and Doe, C.Q. (2017). Steroid hormone induction of temporal gene expression in Drosophila brain neuroblasts generates neuronal and glial diversity. Elife *6*.

Tinevez, J.Y., Perry, N., Schindelin, J., Hoopes, G.M., Reynolds, G.D., Laplantine, E., Bednarek, S.Y., Shorte, S.L., and Eliceiri, K.W. (2017). TrackMate: An open and extensible platform for single-particle tracking. Methods *115*, 80-90.

Tola, E., Lepetit, V., and Fua, P. (2008). DAISY: A Fast Local Descriptor for Dense Matching. CVLAB, 1-8.

Toyoshima, Y., Tokunaga, T., Hirose, O., Kanamori, M., Teramoto, T., Jang, M.S., Kuge, S., Ishihara, T., Yoshida, R., and Iino, Y. (2016). Accurate Automatic Detection of Densely Distributed Cell Nuclei in 3D Space. PLoS Comput Biol *12*, e1004970.

Trcek, T., Lionnet, T., Shroff, H., and Lehmann, R. (2017). mRNA quantification using single-molecule FISH in Drosophila embryos. Nat Protoc *12*, 1326-1348.

Tsuji, T., Hasegawa, E., and Isshiki, T. (2008). Neuroblast entry into quiescence is regulated intrinsically by the combined action of spatial Hox proteins and temporal identity factors. Development *135*, 3859-3869.

Van der Walt, S., Schönberger, J.L., Nunez-Iglesias, J., Boulogne, F., Warner, J.D., Yager, N., Gouillart, E., and Yu, T. (2014). scikit-image: image processing in Python. PeerJ *2*, e453.

Waithe, D., Hailstone, M., Lalwani, M.K., Parton, R., Yang, L., Patient, R., Eggeling, C., and Davis, I. (2016). 3-D density kernel estimation for counting in microscopy image volumes using 3-D image filters and random decision trees. In European Conference on Computer Vision (Springer), pp. 244-255.

Waithe, D., Rennert, P., Brostow, G., and Piper, M.D. (2015). QuantiFly: Robust Trainable Software for Automated Drosophila Egg Counting. Plos One *10*, e0127659.

Walsh, A.J., Cook, R.S., and Skala, M.C. (2017). Functional Optical Imaging of Primary Human Tumor Organoids: Development of a Personalized Drug Screen. J Nucl Med *58*, 1367-1372.

Walt, S.v.d., Colbert, S.C., and Varoquaux, G. (2011). The NumPy array: a structure for efficient numerical computation. Computing in Science *13*, 22-30.

Weissman, I.L., Anderson, D.J., and Gage, F. (2001). Stem and progenitor cells: origins, phenotypes, lineage commitments, and transdifferentiations. Annu Rev Cell Dev Biol *17*, 387-403.

Wernike, D., Chen, Y., Mastronardi, K., Makil, N., and Piekny, A. (2016). Mechanical forces drive neuroblast morphogenesis and are required for epidermal closure. Dev Biol *412*, 261-277.

Willoughby, L.F., Schlosser, T., Manning, S.A., Parisot, J.P., Street, I.P., Richardson, H.E., Humbert, P.O., and Brumby, A.M. (2013). An in vivo large-scale chemical screening platform using Drosophila for anti-cancer drug discovery. Dis Model Mech *6*, 521-529.

Wilson, C., Leiblich, A., Goberdhan, D., and Hamdy, F. (2017). The Drosophila accessory gland as a model for prostate cancer and other pathologies. In Current topics in developmental biology (Elsevier), pp. 339-375.

Wodarz, A., and Gonzalez, C. (2006). Connecting cancer to the asymmetric division of stem cells. Cell *124*, 1121-1123.

Wolff, C., Tinevez, J.Y., Pietzsch, T., Stamataki, E., Harich, B., Guignard, L., Preibisch, S., Shorte, S., Keller, P.J., Tomancak, P.*, et al.* (2018). Multi-view light-sheet imaging and tracking with the MaMuT software reveals the cell lineage of a direct developing arthropod limb. Elife *7*.

Wu, C.F., Suzuki, N., and Poo, M.M. (1983). Dissociated neurons from normal and mutant Drosophila larval central nervous system in cell culture. J Neurosci *3*, 1888-1899.

Wu, J.S., and Luo, L. (2006). A protocol for dissecting Drosophila melanogaster brains for live imaging or immunostaining. Nat Protoc *1*, 2110-2115.

Xing, F., and Yang, L. (2016). Robust Nucleus/Cell Detection and Segmentation in Digital Pathology and Microscopy Images: A Comprehensive Review. IEEE Rev Biomed Eng *9*, 234-263.

Yang, C.P., Samuels, T.J., Huang, Y., Yang, L., Ish-Horowicz, D., Davis, I., and Lee, T. (2017). Imp and Syp RNA-binding proteins govern decommissioning of Drosophila neural stem cells. In Development, pp. 3454-3464.

Yu, H.H., Kao, C.F., He, Y., Ding, P., Kao, J.C., and Lee, T. (2010). A complete developmental sequence of a Drosophila neuronal lineage as revealed by twin-spot MARCM. PLoS Biol *8*.

Zhang, H., Fritts, J.E., and Goldman, S.A. (2008). Image segmentation evaluation: A survey of unsupervised methods. Computer Vision and Image Understanding *110*, 260-280.

Zhang, Y., Fuger, P., Hannan, S.B., Kern, J.V., Lasky, B., and Rasse, T.M. (2010). In vivo imaging of intact Drosophila larvae at sub-cellular resolution. J Vis Exp.

Zuiderveld, K. (1994). Contrast limited adaptive histogram equalization. In Graphics gems IV (Academic Press Professional, Inc.), pp. 474-485.

APPENDIX

A) User Manual: CytoCensus Quickstart Guide

CytoCensus is intended to allow biologists to identify and count cells and objects of interest in complex 3D microscopy images, without detailed knowledge of image analysis. Specifically, it is targeted at users with large numbers (>10) of microscopy images or large time-lapse datasets, and asks the user to do some very simple annotations (training the program) to identify the cells of interest. Test data sets can be found at: https://github.com/hailstonem/CytoCensus

What can I do with CytoCensus?

- **Count cells/objects in 3D** – if they are approximately round
- **Find cells** of a particular size
- **Identify cell/object centre**s i.e. XYZ coordinates
- **Determine counts and coordinates over time-lapse** image series
- Determine object counts in a **user defined regions of interest (ROI)** over time
- Compare the relative numbers of different object classes (or subclasses)

With a little FIJI/ImageJ knowledge, CytoCensus outputs can be used to:

- **Determine average intensity** across a class of cells
- **Track cells** (using CytoCensus probability maps as input to the ImageJ TrackMate plugin to track previously untrackable cell classes)

What can't I do with CytoCensus?

- **Identify things than you cannot recognise.** If you cannot find it and train on it, you cannot use CytoCensus.
- **Find objects of extremely variable size.** CytoCensus allows some variation in size, but within limits - but if you have a large set of cells and a small set, you can make a separate model to recognise each but not both together.
- **Trace long, thin, or strangely shaped cells, such as neurons.** However, in you can use CytoCensus to identify roughly round features, such as cell bodies or nuclei, and pass those co-ordinates to other analysis tools
- **Measure cell areas\volumes directly.** Unlike segmentation tools such as Ilastik, CytoCensus determines cell centres, not their edges (boundaries). To measure cell areas/volumes it is possible to use the XYZ centres output from CytoCensus as the starting point in an image analysis pipeline.
- **Measure total fluorescence of cells directly.** As CytoCensus does not identify boundaries, it cannot determine the extent of the cell, therefore, total fluorescence cannot be measured directly. However, as described for cell volume above, CytoCensus outputs may be used as the starting point for an image analysis pipeline.

What kind of data should I use?

CytoCensus is designed to work with multiple 3D image files, with multiple z-slices, multiple channels, and/or multiple timepoints. CytoCensus requires TIFF files. If your data is in another format, convert it using the excellent FIJI/bioformats plugin found at https://imagej.net/Fiji/Downloads (Linkert *et al.*, 2010).

HOW DO I USE CytoCensus?

CytoCensus is best used on a computer with 16GB or more RAM with a mouse. CytoCensus works with a PC, Mac or Linux platform.

SETUP

1. Download the latest version of CytoCensus from https://github.com/hailstonem/CytoCensus/releases/
2. Run 'CytoCensus_Train_0_1_xxx'

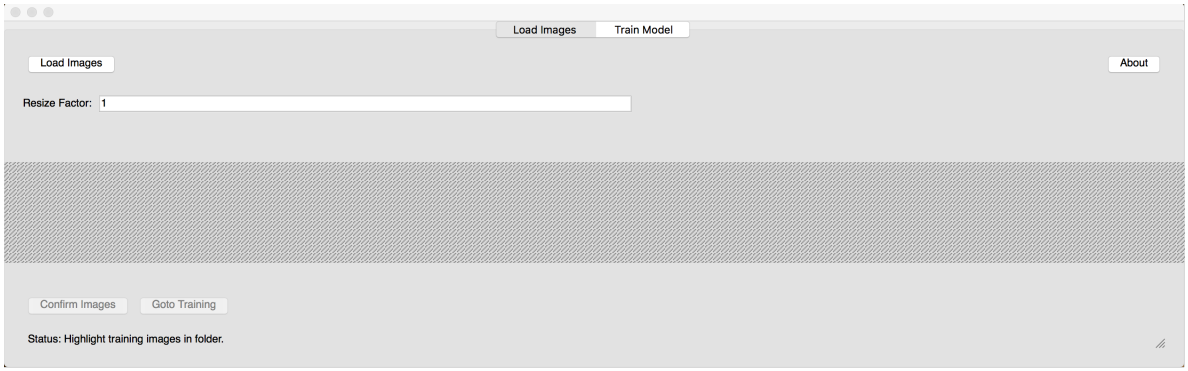

3. **Load images** (TIF/TIFF file support only) - 3D images should be stacks or hyperstacks: https://imagej.nih.gov/ij/docs/guide/146-8.html
    a. If your images are in another format, use ImageJ/FIJI to convert them https://imagej.net/Fiji/Downloads
    b. You can choose multiple images files but at this point it is unnecessary to load your whole image dataset, simply load enough to capture <<represent>> the variation in your images
4. If your images are large (>512x512pixels and/or z>50), and the objects you want to look at are also large (>20 pixels), choose a '**Resize factor**' (e.g. 2,4,5). This will speed up loading and processing, but setting this too high will reduce your image quality.
5. Set '**Sampling**': 5 is the default and generally works. Decreasing this to 2 (minimum = 1) can provide better accuracy with rare and small objects (<5pixels). 'Sampling' sets how many pixels are used in calculations to determine the probability density values.

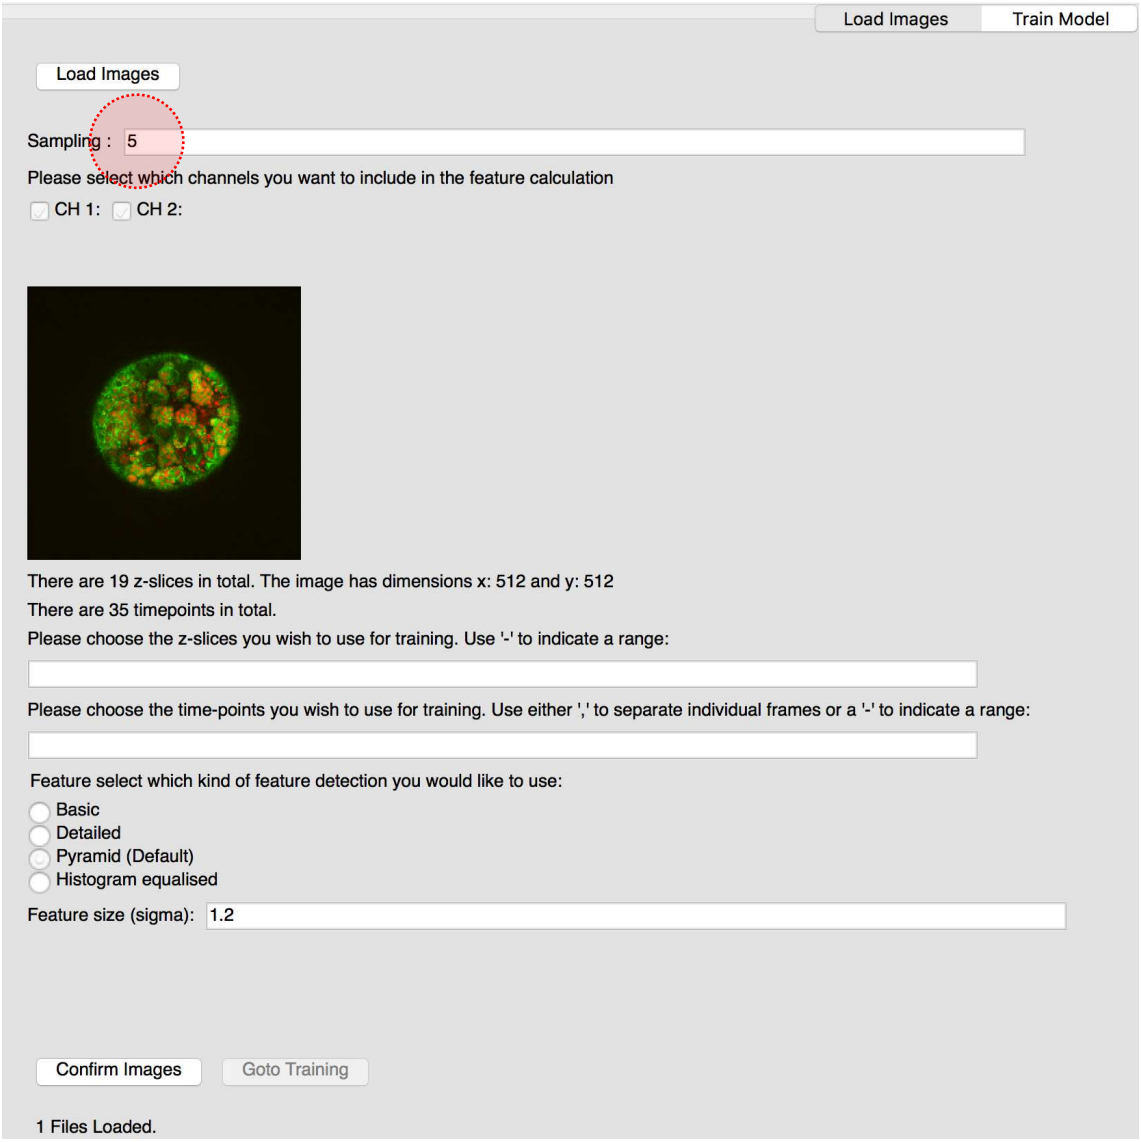

6. '**Feature set'**. Leave this at 'Pyramid' for most cases. Basic is faster, but requires you to set the 'feature size' parameter (see later). 'Histogram equalised' is slower, but can help if you image contrast varies a lot.
7. '**Confirm Images'** and '**Goto Training'**.

TRAINING

8. **Navigate** through your images using up down keys (z), left right keys (t), < > keys (files) or the corresponding buttons on the GUI.

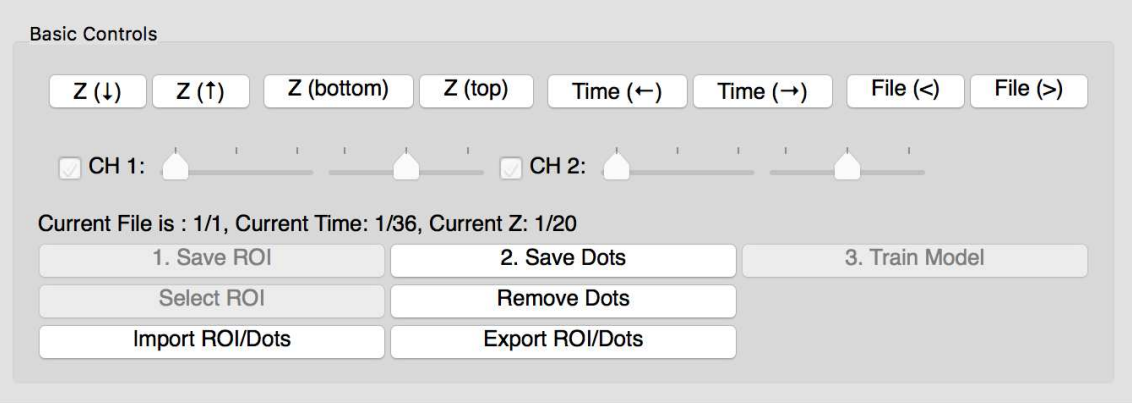

9. Find a z-slice with an example of your cell or object of interest. Adjust contrast using the contrast sliders if necessary.
10. **Right-click and drag to make an ROI** containing a number of objects of interest (and some background).

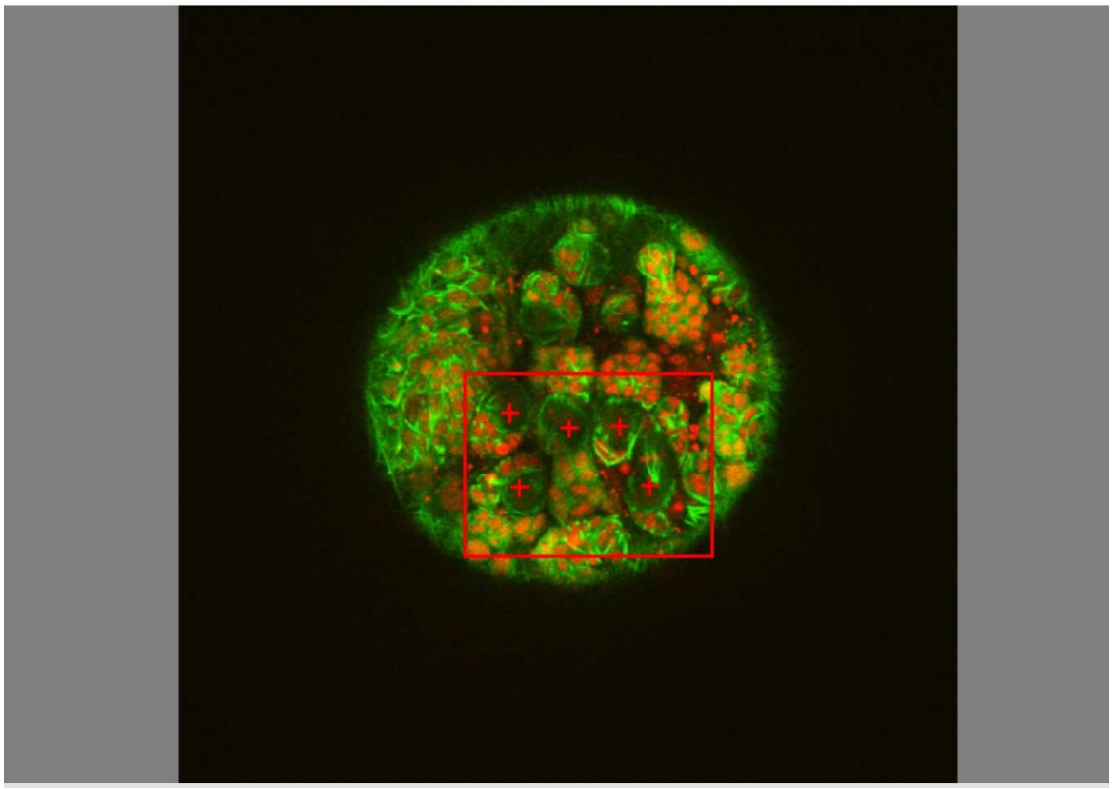

11. When happy with your ROI, **left-click your objects** of interest in it. If you're not happy simply right-click and drag to redraw a new ROI. Boxes can't be moved or resized. (Don't change image/z or you'll lose your ROI.)
12. **Save** your ROI.

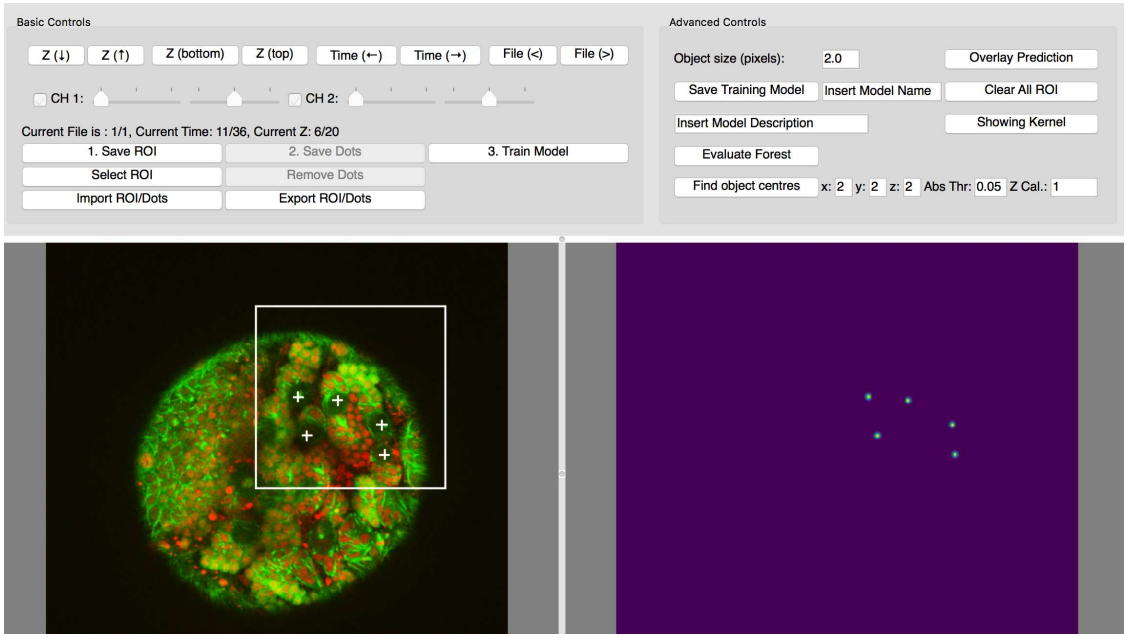

13. You should now see some points in the right hand plot. These should be roughly the same size as your object of interest. If they're much smaller, or barely visible, increase the '**Object Size'** parameter until the size roughly matches that of your objects. Do not worry too much if this isn't perfect, it can be changed later.

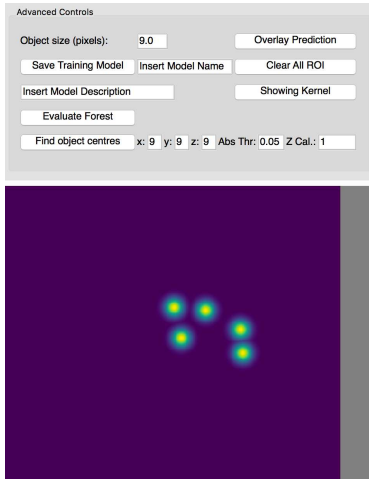

14. Add more training if desired (8-12), ideally in a different z/t. Note: adding multiple annotations in Z planes rather than timepoints is more efficient for the program but that results will only be displayed for the current Z-stack (all your training is used!).
15. When happy with your training '**Train Model'.**

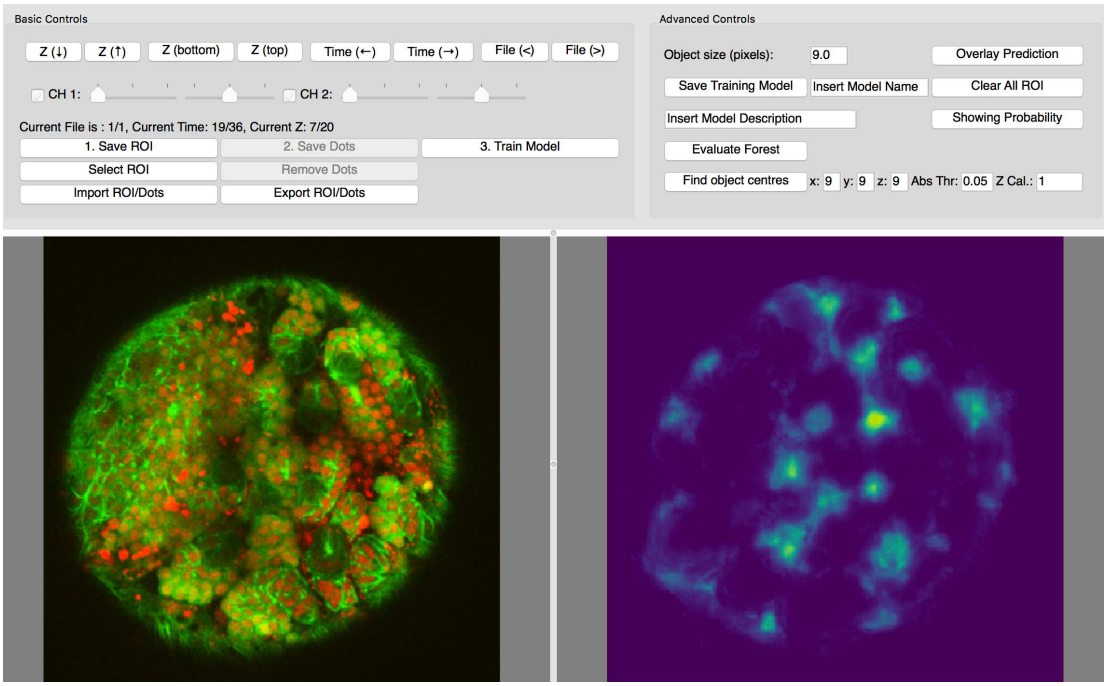

16. Use the **probability map** on the right to evaluate the training in the current z stack. To examine the result in a different z-stack, use the '**Evaluate Forest'** button. Use the '**Overlay Prediction'** button to help. In an ideal map there should be clear discrete regions of green/yellow/red corresponding to your objects of interest, and faint blue or very little elsewhere. Add more training as described in 8-14 to improve your training. Note that you can add more points, but also that you can make new ROIs with no points in them if there is some background that is causing an issue.

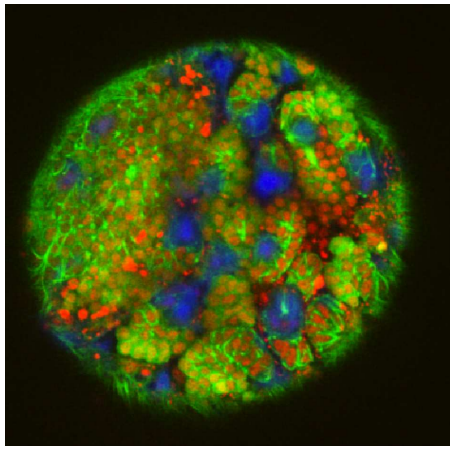

17. When happy with the **probability map,** click '**Find Object centres**'. This will attempt to find the highest intensity points of interest in 3D! This is the hardest part to get right, and can be quite variable depending on how consistent your objects are. There are 4 parameters that are important for this. The '**threshold'** parameter is the most important. Typically a value between 0 and 20 is best. The **x and y sizes** should be a good estimate based on the size of objects chosen earlier, but the **z size** can vary a lot. If the **z calibration** on the right is accurate (this is the difference between your x/y resolution and your z resolution) then matching it to the object size make sense. If you have 1um pixel spacing in x/y and take 5um steps, then z calibration would be 5. If your object size is 10, then, a reasonable value for x,y,z sizes is 10, but you may need to try changing the z size between 5 and 15 to see what gives the best result.

18. When you're happy with this, **choose a name** for your "Training Model" and the **Save Training Model.**

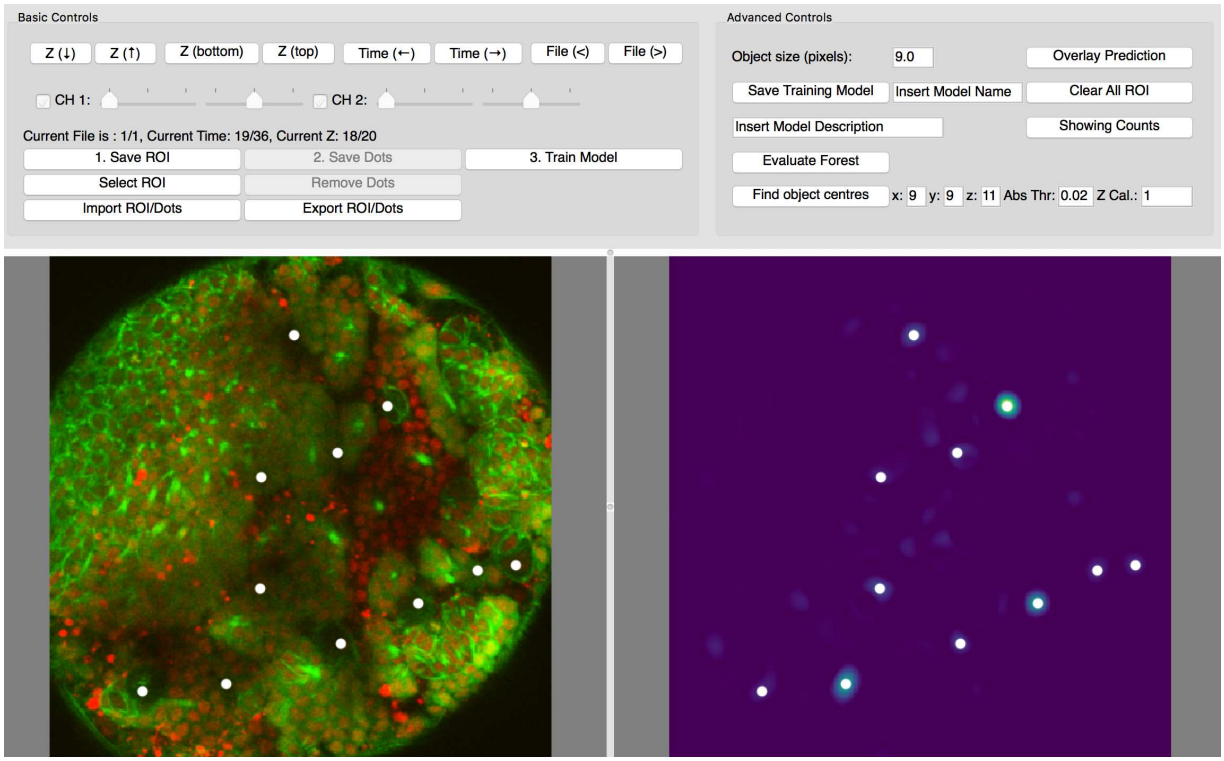

BATCH MODE

After training by the user, CytoCensus employs a collection of filters and scores and combines them to find features in the image that identify the user-defined cell centres. In this way, a series of transformations of the image data (referred to as the "**trained model**") are learned. The trained model is subsequently applied pixel by pixel to new data sets, outputting an estimated probability map of the cells of interest with their predicted cell centre co-ordinates.

19. The hard part is now done. Now open 'CytoCensus_Evaluate_0_1_xxx'
20. Load in your images as before (**Add images**). This time load the whole image dataset (i.e. leave the options for Z and T blank).

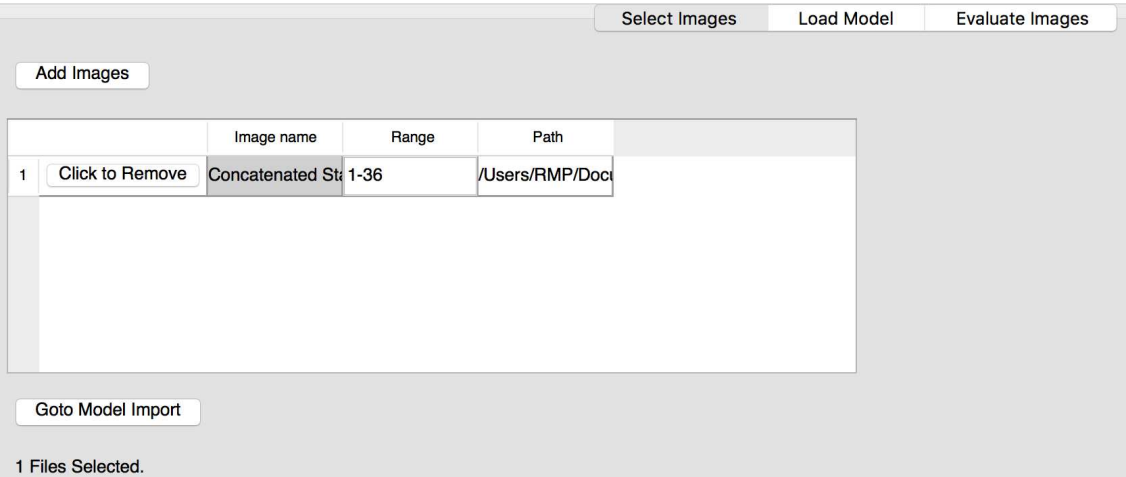

21. **'Goto Model Import**'. Note, take care in applying models across data sets, particularly, for example WT vs mutant data sets. If there is a significant difference in data quality of in the object (cell type) of interest between datasets, then retraining is required. However, if may be of value to apply training from a WT case to a mutant - if the cells of interest are not detected then there may be a significant difference from WT - i.e. WT cells are not present in the mutant.

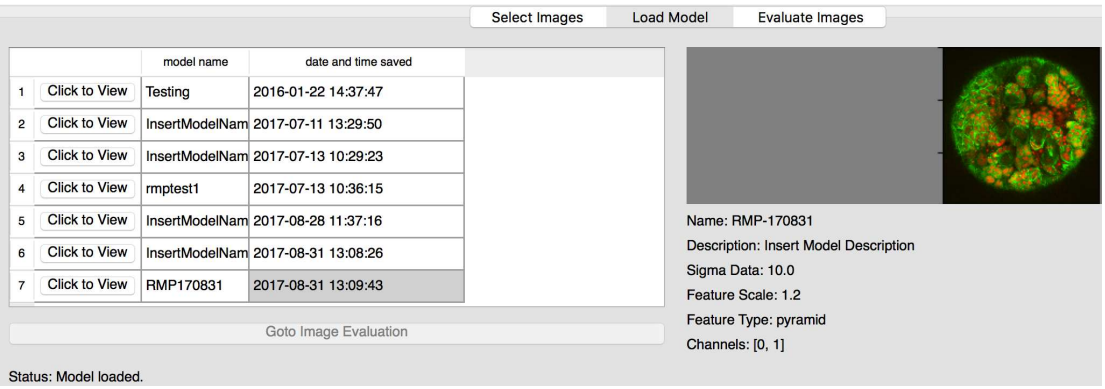

22. Click on your **saved Model**
23. Goto **'Image Evaluation'**

24. **Evaluate Images.** This is the most time-consuming step (for us, it takes a little longer than 5 minutes per 50 z-planes, so can be a couple of hours for a 150 timepoint movie), but this will depend on the computer you're running it on.

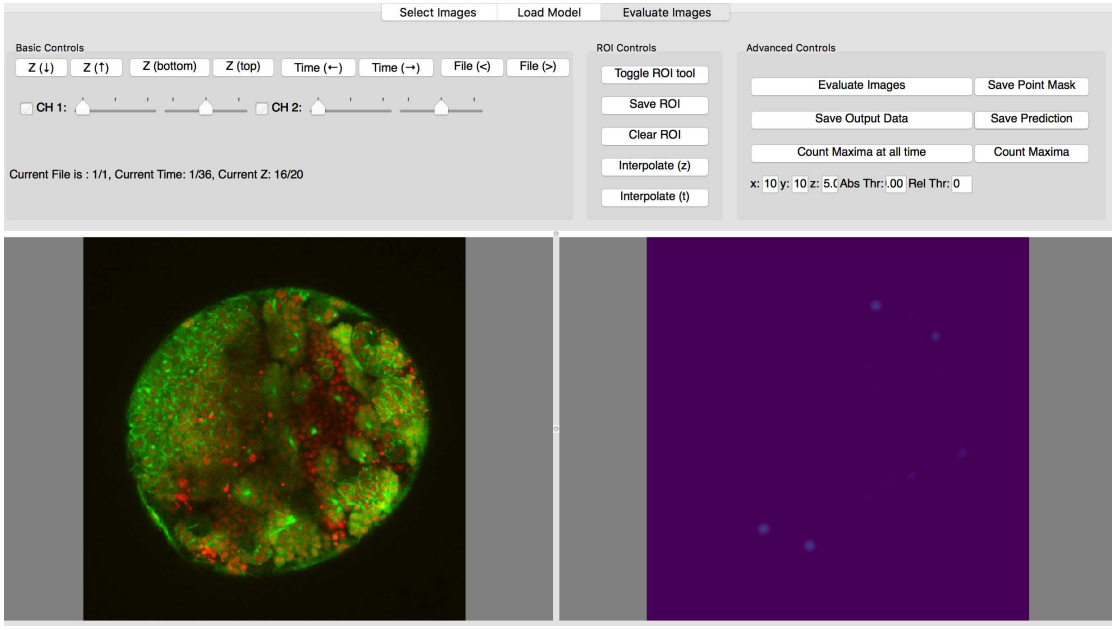

25. When evaluation is complete, you can export the probability maps as tiff files **'Save Prediction'.** Note: this will this be a **hyperstack** if the **input data** is a hyperstack.
26. If you want the counts and cell centre coordinates use '**Count Maxima at all time**', (this make take a few minutes), and then use **'Save Output Data'** to save the results as excel files. You get one file per image file, with the x,y,z coordinate outputs, and one additional file with total counts per image file.
27. If you want to restrict your analysis to a subregion of your image, use the **ROI tool**.

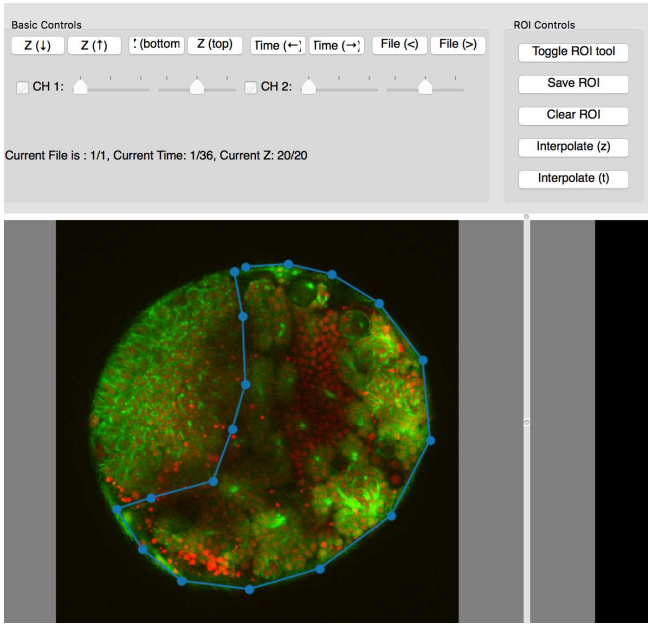

28. First **Toggle ROI tool**
29. Next click a series of points in the current z plane.
30. **Save ROI**
31. Go to another Z plane and click another series of points.

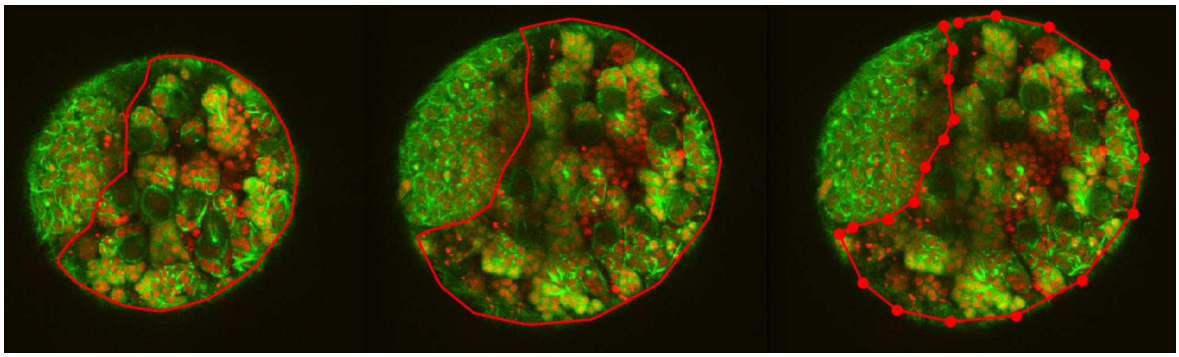

32. Click **Save ROI**
33. Click **Interpolate (z),** to automatically estimate a 3D region of interest between the defined planes
34. Check you're happy with this ROI. If you like you can edit the ROI by adding additional points in more Z planes.
35. If you have images with a T dimension, you can also interpolate between ROIs in T
36. Repeat for additional files, if necessary.
37. use '**Count Maxima at all time**' to find the counts and centres in these subregions

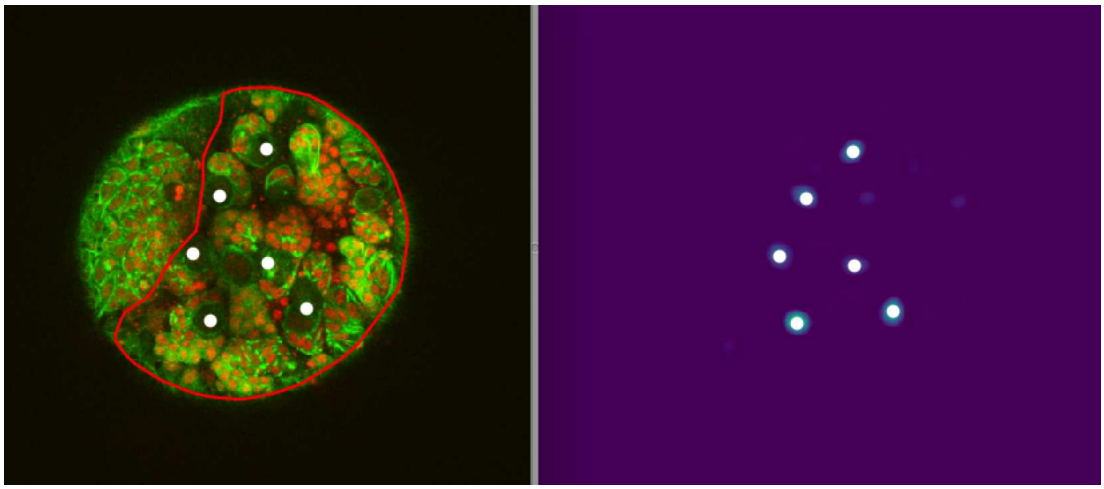

38. Use '**Save Output Data**' to save the results