UNIVERSITAT POLITÈCNICA
DE CATALUNYA
BARCELONATECH

PROGRAMA DE DOCTORAT EN ENGINYERIA BIOMÈDICA
DEPARTAMENT D'ENGINYERIA DE SISTEMES, AUTOMÀTICA I INFORMÀTICA INDUSTRIAL
CENTRE DE RECERCA EN ENGINYERIA BIOMÈDICA

Multiscale Image Analysis of Calcium Dynamics in Cardiac Myocytes

Alexander Vallmitjana Lees

June 2017

Advisor

Raúl Benítez Iglesias

Abstract

CARDIAC myocytes constitute a unique physiological system. They are the muscle cells that build up heart tissue and provide the force to pump blood by synchronously contracting at every beat. This contraction is regulated by calcium concentration (among other ions) which exhibits a very complex behaviour, rich in dynamical states at the molecular, cellular and tissue levels. Details of such dynamical patterns are closely related to the mechanisms responsible for cardiac function and also cardiac disease, which is the first cause of death in the modern world. The emerging field of translational cardiology focuses on the study of how such mechanisms connect and influence each other across spatial and temporal scales finally yielding to a certain clinical condition.

In order to study such patterns, we benefit from the recent and very important advances in the field of experimental cell physiology. In particular, fluorescence microscopy allows us to observe the distribution of calcium in the cell with a spatial resolution below the micron and a frame rate around the millisecond, thus providing a very accurate monitoring of calcium fluxes in the cell.

This thesis is the result of over five years' work on biological signal and digital image processing of cardiac cells. During this period of time the aim has been to develop computational techniques for extracting quantitative data of physiological relevance from microscopy images at different scales. The two main subjects covered in the thesis are image segmentation and classification methods applied to fluorescence microscopy imaging of cardiac myocytes. These methods are applied to a variety of problems involving different space and time scales such as the localisation of molecular receptors, the detection and characterisation of spontaneous calcium-release events and the propagation of calcium waves across a culture of cardiac cells.

The experimental images and data have been provided by four internationally renowned collaborators in the field. It is thanks to them and their teams that this thesis has been possible. They are Dr. Leif Hove-Madsen from the Institut de Ciències Cardiovasculars de Catalunya in Barcelona, Prof. S. R. Wayne Chen from the Department of Physiology and Pharmacology in the Libin Cardiovascular Institute of Alberta, University of Calgary, Dr. Peter P. Jones from the Department of Physiology in the University of Otago, and Prof. Glen Tibbits from the Department of Biomedical Physiology & Kinesiology at the Simon Fraser University in Vancouver.

The work belongs to the biomedical engineering discipline, focusing on the engineering perspective by applying physics and mathematics to solve biomedical problems. Specifically, we frame our contributions in the field of computational translational cardiology, attempting to connect molecular mechanisms in cardiac cells up to cardiac disease by developing signal and image-processing methods and machine-learning methods that are scalable through the different scales. This computational approach allows for a quantitative, robust and reproducible analysis of the experimental data and allows us to obtain results that otherwise would not be possible by means of traditional manual methods.

The results of the thesis provide specific insight into different cell mechanisms that have a non-negligible impact at the clinical level. In particular, we gain a deeper knowledge of cell mechanisms related to cardiac arrhythmia, fibrillation phenomena, the emergence of alternans and anomalies in calcium handling due to cell ageing.

Index

List of Figures

Microscopy images shown in the thesis are all courtesy of our collaborators. Diagrams and graphs are either original or have appeared in one of the publications associated to this work.

Four figures are taken from publications by other authors and this is specified in the figure caption. In the four cases the authors are aware that the figures appear in this thesis and have given their permission.

List of Tables

1 Introduction

1.1 Motivation

During the past three decades Mankind has entered what in the future is going to be known as the information revolution or age of information. Most probably it will be grouped as part of a longer, several-hundred-year-old, technological revolution that will include the industrial revolution [1, 2]. This information age was born to internet with the digital revolution and is characterised by the ability to broadcast and store great amounts of data. It is estimated the world's total storage capacity in 2014 was of 5ZB [3] and the total annual traffic of 2016 has surpassed the ZB [4], a zettabyte being 10^{21}bytes. If we were to store this amount of information as text, without compression each character taking up one byte, and we were to print it out using the format of this thesis, without figures each page containing about 3000 characters and considering a paper thickness of 0.1mm... our *book* would have a thickness of 10^3AU (astronomical units), that is a book with the thickness of one thousand times the mean distance from Earth to the Sun.

These overwhelming amounts of data cover all kinds of information and nonsense, of which a very small part is dedicated to science, and of this very small part a great deal will either never be applied for it is either too specialised and condition dependent, or simply false [5, 6]. Nevertheless, progress in science does require a whole load of parallel minute contributions (and failures) for a single breakthrough and for its author to go down in history and therefore all data, good or bad, needs to be processed.

Specifically in the field of biology, the onset of the information age has forced image processing and analysis to become a key step. Thanks to the combined progress of microscopy and computer engineering current researchers are producing huge amounts of data, far more than they can store and far, far more than they can analyse by means of the manual

and time-consuming traditional techniques [7]. Techniques such as segmenting, labelling or counting objects in an image and measuring their morphological features have historically been human dependent and subjective whereas now, with the aid of software tools, they are becoming objective and more reproducible and results that were previously limited to qualitative observation are now gradually becoming quantitative [8]. Hence the need for bridge disciplines: disciplines that are halfway between the classical science pillars, like biomedical engineering, that cover the gap in this case between researchers applying bio-imaging and the advanced mathematical and computational tools available [9, 10].

Each research group has overcome their own problems in their own ways and this has led to a gradually increasing proliferation in methods and software tools, both commercial and of free access, covering a wide range of utilities like image viewing, manipulation, annotation, analysis, classification and storage [11]. There are currently hundreds of them available but each individual tool is hardly used in the sense that they are highly tailored to resolve specific problems and so only applicable under very particular conditions and are therefore very hard to tune and adapt to new experiments [12, 13, 9]. Besides, it is very common that when attempting to widen the spectrum of applicability of a tool one tends to compromise its performance and accessibility and there is a general lack of usability [12, 14].

1.2 Framework

The field of translational cardiology is no different from the rest of the biomedical world. It involves experimental data that may come in very different formats and belonging to different physical scales, hence the thesis title with the term multiscale. A scale integration is required that connects molecular, sub-cellular, cellular and multicellular cultures. In this sense, this thesis is an attempt to connect the molecular scale to the tissue scale by developing a set of tools for signal and image processing.

Some of the specific problems that overall make processing experimental data from translational cardiology so hard and complex are listed next:

- Many different data types: electrophysiological recordings, line scan images, frame scan images, 3D stacks of images and image sequences.

- Necessity to be able to compare amongst sets of data that may have been taken under very different conditions such as different sampling rates or pixel physical sizes. Also to be able to deal with all the available digital file formats such as images in grayscale versus colour-indexed versus RGB or image sequences versus multipage images versus video files.

- Multiple temporal and spatial scales showing a great variety of spatial and temporal dynamics and patterns. Calcium events can range from milliseconds (sparks) to seconds (calcium waves) and show different morphological and dynamical characteristics.

- The characteristics that define each type of event also changes depending on the species that is under study, even amongst close mammals such as mouse and rat.

- Noise is always present and is both of biological and experimental origin. It is a central feature that heavily restricts the use of methods designed for broad use and often makes it necessary to design specific algorithms or to adapt conventional methods to each dataset.

- Large amounts of data per experiment, especially in image sequences, which require both huge hard drive capacity and processing power.

- Great inter and intra-experiment variability requires the processing of huge amounts of experiments to test the developed tools.

This thesis provides a best practice guide towards dealing with this kind of data with an automatic approach that allows both to quantify from a neutral point of view and also to reproduce the results obtained. It specifically provides state-of-the-art solutions to three main problems and many satellite problems that are also encountered while dealing with calcium imaging. These are:

- Detection and characterisation of intracellular molecular structures.

- Detection and characterisation of intracellular calcium events.

- Characterisation of the calcium dynamics of single cells and cellular cultures.

These three main problems are covered individually in the following chapter while reviewing the current state of the art and are described in detail in the subsequent chapters. In this chapter, a not-so-brief introduction to basic cardiac physiology, fluorescence microscopy and the main techniques involved in digital signal and image processing is required due to the fact that we are dealing with a broad discipline that is halfway between fundamental branches of science that historically have been independent.

1.3 Cardiac physiology

Cells are the basic biological unit of all known lifeforms excluding viruses. The name cell comes from the fact that they are closed compartments that form part of a whole. Some organisms are unicellular, like bacteria and protozoa, while all plants and animals are multicellular. As an example, a human body is formed of some 10^{14} cells. Cells have a membrane that is its boundary and are full of water acting as the solvent to all chemical compounds needed for life. Inside the cell is also a complex network of organelles, each having its own function, analogous to the organs in the human body.

Myocytes are the cells that form muscular tissue. The main feature that characterises this kind of cells is their ability to contract, that is to reduce their size. When all cells in a particular muscle tissue contract simultaneously the whole muscle reduces its total length and produces a tension force in the direction of contraction. Cardiac myocytes are the cells that form cardiac tissue and their contraction at each heart beat reduces the total volume of the heart cavities which allows the continuous pumping of blood through the circulatory system.

Cardiac myocytes are elongated cells attached to one another in the longitudinal direction forming the cardiac muscular fibres. Each of them is around 100μm long and around 25μm wide and a human heart consists of some 10^9 cardiac myocytes all contracting at each heart beat, pumping around 7000 litres of blood a day and producing a total of around $3{\cdot}10^9$ beats in an average human lifetime [15]. Vertebrate animals have hearts composed of between two to four chambers, mammals having two atria and two ventricles. Atria are smaller and with each contraction they collect blood from the system and serve it into the ventricles, which in turn during the same beat expel the blood from the heart sending it out again into the circulatory system. The heart is divided into two halves, the right atria serving to the right ventricle which pumps to the lungs for oxygenation, this blood then returns oxygenated back to the heart via the left atria which serves to the left ventricle which pumps into the rest of the system via the aorta.

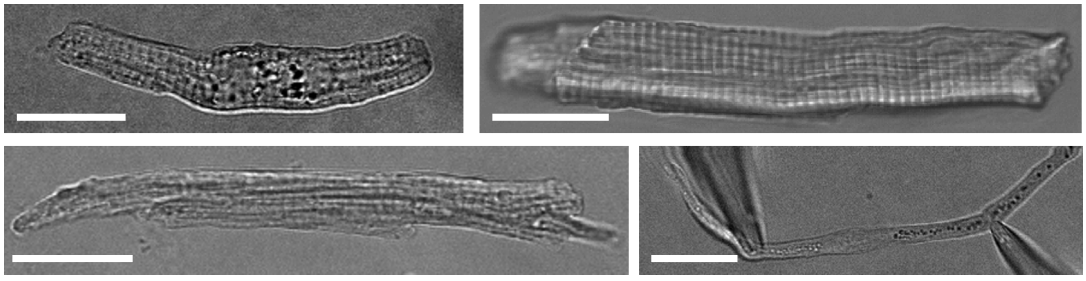

Fig. 1 *Atrial myocytes. Top row, left to right, human and pig, bottom row, left to right, mouse and trout (held by two pipettes). All scales are 20μm long.*

Contraction of myocytes is dictated via electrical stimulus and the process converting this electrical stimulus into mechanical force is called excitation-contraction coupling [16–18]. The

charge carriers at the cellular level are ions of different chemical elements, the most relevant being Na^+,K^+,Ca^{2+},Cl^-. In the case of cardiac myocytes the ion that is mainly responsible for regulation of contraction is Ca^{2+} [19].

Each cell is limited by a membrane that isolates it from the exterior. This membrane consists of by two layers of phospholipids, a kind of lipid (fat) molecule, and together they form what is known as a lipid bilayer. The main property of phospholipids is that one end of the molecule is hydrophilic (attracted to water) while the other end is hydrophobic (repelled by water) and therefore when in water the hydrophobic ends are forced to aggregate exposing the hydrophilic ends on each side of the membrane. This is the reason why the cell membrane is impermeable to water and any molecules and free ions that may be dissolved in water thus allowing it to hold considerable charge gradients.

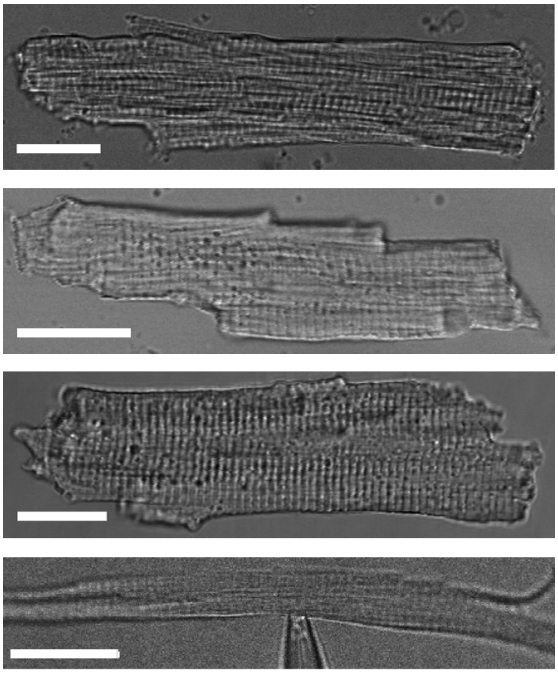

Fig. 2 *Ventricular myocytes. From top to bottom: mouse, rat, rabbit and zebrafish (held by a pipette). All scales are 20μm long.*

For the particular case of Ca^{2+}, the concentration in extracellular medium is around $1mM$ whereas the concentration in the cytosol (intracellular fluid) of a resting cardiac myocyte is around $0{,}1\mu M$, a factor of 10^4 between each other [20]. This great difference in concentration allows the cell to very rapidly increase its inner concentration by opening the membrane up and letting Ca^{2+} fall in.

This ability of the cell's membrane to open up is due to transmembrane proteins. These are molecules that are embedded in the membrane, going straight through it from one side to the other, and acting as doorways for the cell. There are many different types of transmembrane proteins depending on wether they are selective to particular ions, if they are activated by the presence of a certain molecule or by an electrical potential difference, or if they allow only the ion flow in the direction of gradient or if they are capable of pumping ions counter-gradient. The three main transmembrane proteins for Ca^{2+} transport are the L-type calcium channel (voltage regulated opening allowing ions to pass in direction of gradient), the Ca^{2+} ATPase (which pumps Ca^{2+} out of the cell consuming ATP), and the sodium-calcium exchanger (which removes one Ca^{2+} from the cell while letting three Na^+ in).

Myocytes have also a skeleton-like structure of proteins that holds them together and acts as the contractile machinery. This structure can be seen as a grid of proteins attached to one

another with the ability to slide through one another reducing the effective lattice of the grid. This lattice grid unit is around $1\,\%$ of the length of the myocyte and the main proteins that form it are myosin, actin and troponin. The latter has a series of Ca^{2+}-binding sites so that in the presence of Ca^{2+}, the sites are filled and the proteins undergo their conformational change. It is a sort of double direction rowing boat, myosin being the boat with oars holding it to actin on either end of it so that the rowing motion brings together actin at either end reducing the distance between them and therefore the total length of the myocyte. This all seems rather confusing from an engineering point of view, but the bottom line, however, is that the presence of Ca^{2+} in the cell triggers its contraction.

To regulate this presence of Ca^{2+}, myocytes have a much more clever way than just letting it in through the membrane. The process is called *calcium-induced calcium release* [21, 22], and the name comes from the fact that myocytes have a Ca^{2+} reservoir inside them that contributes largely to raising the concentration of Ca^{2+} far above the concentration they would reach by just letting Ca^{2+} from the outside. This reservoir is the sarcoplasmatic reticulum, a closed structure within the cell, with its own electrically insulating membrane, and with its own transmembrane proteins to allow ion flux through it. This reticulum releases large amounts of Ca^{2+} into the intracellular fluid and this release is triggered by the presence of Ca^{2+} coming from ouside of the cell. The two main transmembrane proteins that allow this calcium flux from the sarcoplasmatic reticulum are the ryanodine receptor (RyR2) for the release flux and the SERCA pump for the subsequent regathering back into the reticulum.

To summarise, here is a step list of the cycle each cardiac myocyte undergoes with each heart beat:

1. In the resting state the myocyte keeps an electrical membrane potential of around -90mV. The cell is polarised, with an excess of negative charge.

2. A variation in the extracellular charge (due to an electric current coming from a neighbouring cell) triggers a depolarisation of the cell, allowing great amount of positive charge through the membrane, mainly Na^{+} and K^{+}.

3. Membrane potential changes sign and reaches a value of up to +50mV. At this point Ca^{2+} starts to enter the cell while K^{+} are ejected, fixing the positive voltage value (this phase is called the plateau of the action potential), and this rise in Ca^{2+} concentration starts the calcium-induced calcium release process described above, by which the release of Ca^{2+} from the sarcoplasmatic reticulum grants a sudden rise in inner Ca^{2+} concentration which initiates contraction of the myocyte.

4. Finally, Ca^{2+} is collected back into the sarcoplasmatic reticulum and ejected out of the cell to return to the negative potential resting state. To avoid a cell initiating the cycle again with its own charge release, the transmembrane proteins that allow the entrance of charge are shut

down for a short interval of time until the cycle is complete. This is known as the refractory period.

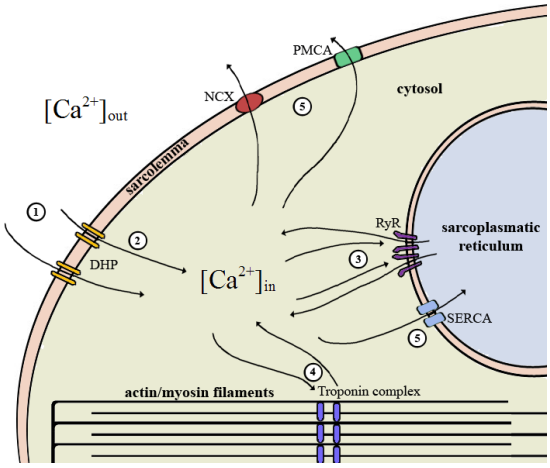

Fig. 3 *Calcium-induced calcium release diagram of the heart beat cycle at the cellular level: Extracellular calcium ions (1) enter the cell (2) trigger the calcium release from the sarcoplasmatic reticulum (3) which all together activates the contractile machinery (4) and finally calcium levels are recovered (5).*

1.4 Fluorescence microscopy

1.4.1 Optics introduction

Refraction is a property of light that defines the fact that it changes the direction of propagation whenever it crosses a change of transmission medium. Typically the two media are air to glass or vice versa, but for any given two media this direction change is described by Snell's law:

$$n_1 sin(\theta_1) = n_2 sin(\theta_2)$$

where n_1 and n_1 are the refractive indices of the two transmission mediums and θ_1 and θ_1 are the angles of incidence and refraction perpendicular to the boundary surface between the two media.

A flat piece of glass will make light change direction when entering and then rectify when leaving to recover its original direction. If the two surfaces of the glass are not parallel (i.e. the thickness is not homogeneous) the light will effectively change its direction. The simplest lens is a piece of transparent material, polished to a spherical curvature on either side and with thickness decreasing from the centre, with the result that light entering perpendicular to the lens plane is focused to a single point in space regardless of the position of incidence on the plane.

A microscope is an optical device designed to produce an image of an object that is larger than the object itself and therefore achieving magnification of the object. The magnification

power is defined as the ratio between the angle under which we would see the object without the microscope (at the closest possible distance the eye can focus it, standardised to 25mm) and the angle under which we actually see the object through the microscope. This can be achieved using a single convergent lens but to a much greater degree by combining a set of lenses. Figure 4 shows an example of an object being viewed under a single convergent lens and also exemplifies the simple rules required for understanding geometrical optics (approximation for thin lenses and small angles). These rules are the following:

- The image of a point is in focus only where rays coming from the point converge to form another point.
- Light entering a lens through its centre will not change its path.
- Light entering a lens parallel to the optical axis will leave the lens passing through the lens focus (the optical axis is the line perpendicular to all lenses in the system that goes through the centre of them).
- Light entering a lens passing through its focus will leave the lens parallel to the optical axis.

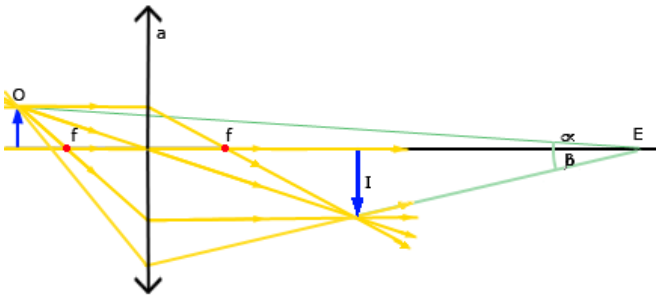

Fig. 4 *Single convergent lens symbolised by the vertical line with arrows (a) and foci symbolised by red dots (f). The light leaving object (O) converges to form the object image (I). An observer (E) will see the object magnified (and inverted in this case) since the angular size of the image (β) is larger than the angular size of the object (α).*

The optical microscope (figure 5) consists of two sets of lenses (objective and eye piece) arranged one behind the other along the optical axis in a manner that the angular separation of light rays coming from two different parts of an object is even wider than if seen through a single lens.

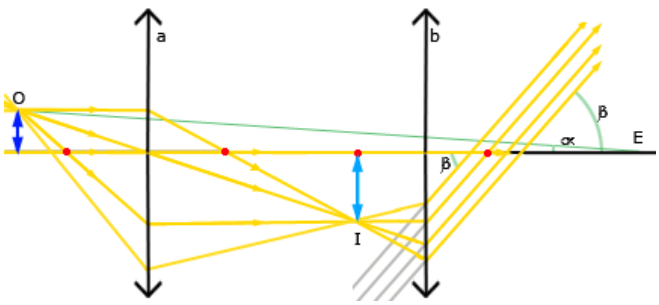

Fig. 5 *Diagram for an optical microscope. The object (O) is viewed from the eye (E) inverted and at infinity. The magnification is the ratio $tan(\beta)/tan(\alpha)$. The red dots are the focal points for each of the two lenses, objective lens (a) and eyepiece lens (b). The pale blue arrow is the image of the object through the objective lens, which the eyepiece lens sends to infinity.*

If we need to record the image through the microscope we have to replace our eye by a camera. In order to view the object in-focus we have to set the camera lens position and detector position so that the points of the object are focused as points on the detector (not as blurred blobs). Figure 6 depicts the situation where we have added the camera to the system and also exemplifies how the image of the object becomes blurred if we change the detector plane position. In the same way, objects lying closer or farther away from the microscope will be in focus on a different plane and for this reason the microscope-camera system can only focus objects that are lying on a plane at a particular distance from the system.

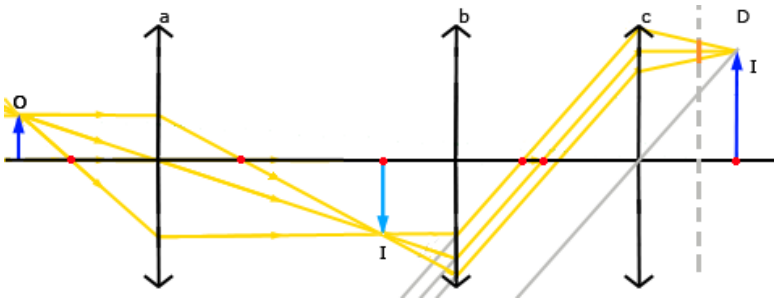

Fig. 6 *Extended version of figure 5 with a camera/eye attached to the microscope. The camera/eye lens (c) produces the final image (I) where the detector chip/retina lies (D). The grey dashed line is a hypothetical detector where the tip of the object is seen as an extended blob (orange) instead of a point. In this hypothetical plane the object is seen out-of-focus.*

The aperture of the system determines the amount of light the system can collect. It is limited by the size of the physical tube holding the microscope and the size of the lenses, and it can also be limited by the use of diaphragms. These are simply opaque elements with an opening in the centre with the single purpose to limit the cone of light that enters the system and therefore limiting the brightness of the image itself. One would expect this to be an undesired result (we want our system to produce bright images in order to detect very dim objects) but it has two clear benefits: On one hand it reduces the effect of optical aberrations [23] and on the other, and more importantly, by reducing the cone of light it also reduces the size of out-of-focus points (see figure 7).

The detector will always be composed of a series of unit detectors, in the case of the digital camera they are the MOS capacitors (producing the pixels in the image), in classical photography it is the size of the crystals sensitive to light (film grain) and for the human eye it is the size of the light sensitive neurones (rods and cones). If our aperture is small enough so that the cone of light coming from an out-of-focus object is smaller than the unit detector then this object will actually be viewed in-focus. This combined effect of both aperture and size of unit detector allows the system to record focused images of objects that are in a range of distances to the microscope, not only at one particular distance. This range of distances is what is known as depth of field.

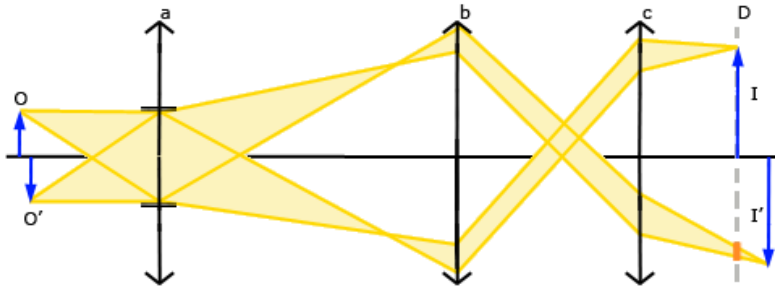

Fig. 7 *Simplified version of the system in figure 6 with a second object closer to the microscope (O') and its image (I'). Also a diaphragm has been added to limit the amount of light and instead of single light rays we see the whole light cone. The image of the new object does not fall on the detector and therefore every point of the new object is recorded as an extended blob (orange). Note that the size of the aperture limits the size of this extended blob, the smaller the aperture is the smaller the extended blob. If we reduce the aperture enough so that the size of the blob is smaller than the unit detector on the recording plane (size of the pixel for a digital recorder) we will obtain focused images of the two objects even if they are not on the same plane. The aperture together with the size of the unit detector controls the depth of field of the system.*

The first microscopes would illuminate the sample by placing a source of light so that some of the light reflected by the sample in all directions would be caught by the microscope to form the image. This system fails if the sample is transparent since light goes straight through and none is reflected, and that is the exact case when the sample is a cell. For this reason the first cells to be viewed through a microscope were dyed using colorants that would give contrast and therefore reveal the basic morphology. The next step was to put the source of light along the optical axis before the sample and to use a lens (condenser lens) to focalise it on the sample. This allows a better control of the amount of light that is being received by the microscope and as the image is now produced by transmission of light through the sample the images are called transmission images or simply *trans*. These types of setup are referred to as bright-field microscopy due to the fact that the resulting images are bright and the light absorbing sample is darker.

Another early improvement was the scanning optical microscope. This technique consists in limiting the illumination so that it hits a single point of the sample and recording the image of that point in the detector. Then the image is built by moving the sample and detector along the plane perpendicular to the optical axis (or moving the source of light), to form the image point by point. The sample is being scanned. This again was an improvement on the control of the amount of light being used, allowing images of very low or very high contrast specimens, as the image is formed point by point, the illumination conditions or microscope parameters can be changed depending on the area that is being scanned. The scanning microscope allowed a further improvement though, which is the confocal microscope.

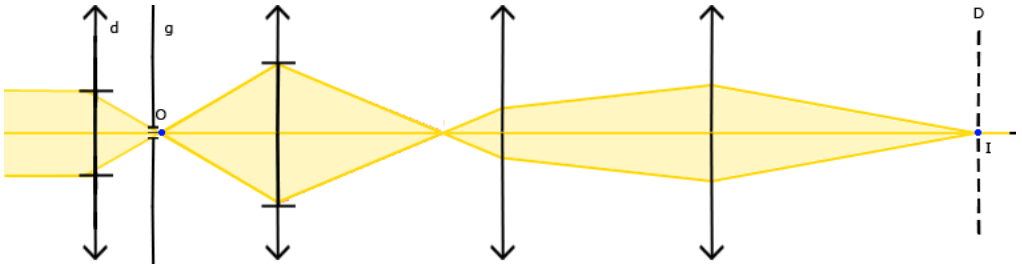

Fig. 8 *Scanning optical microscope. Condenser lens (d) is focusing light coming from our source onto the plane where our specimen is (O), and the diaphragm (g) is blocking the light so that only points along the optical axis are illuminated*

1.4.2 The confocal microscope

The confocal microscope is a scanning microscope with the addition that it is able to greatly reduce the amount of light coming from out-of-focus regions and therefore achieves much sharper images. It was first developed during the 1950s but together with the improvements in fluorescence techniques, since the 1980s it has become the king of microscopes in the study of biological sciences. The principle behind it is the pin-hole, a very small diaphragm, placed in the plane where the microscope objective lens focuses the image of the sample. Figure 9 exemplifies this situation and demonstrates how light coming from points lying out of the focal plane is blocked by the pin-hole. It has a clear disadvantage: we are greatly reducing the amount of light coming from the object, so we need a very sensitive detector (or long exposures) and images become noisier, we are adding diffraction (described further on in section 1.4.5), but on the other hand, we also greatly reduce the blur due to out-of-focus light and, the main benefit of this feature, we can now take 3D images of objects. The microscope scans a whole plane of the sample taking a 2D image, and as the pin-hole ensures the image contains none or very little information from other planes, we can then move the sample away from the microscope and scan another image off a parallel plane to the first. By continuing to do this we obtain a series of stacks that reconstruct the sample.

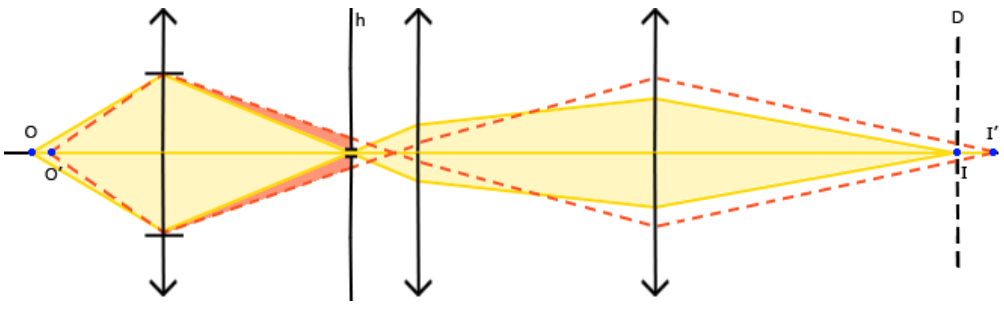

Fig. 9 *Confocal microscope principle: The pin-hole (h) is letting through the whole cone of light for the relevant object (O). Other objects lying before or after (O') have a fraction of the cone blocked at the pinhole, the further they are from the relevant object, the higher the fraction of light blocked.*

1.4.3 Line scan versus frame scanning

If one needs to record the time evolution of a sample, this scanning technique can be very limiting since the time it takes to record the whole scene limits the frame rate of the final sequence. This is the reason why line scans were first introduced. The technique consists in scanning a single line in the sample at each frame and the final result is, instead of a series of 2D images forming a film, a set of 1D lines that can be arranged to form a 2D image with one dimension being space and the other being time. This technique allows a greater temporal resolution at the cost of losing a spatial dimension (see figure 10 for an example).

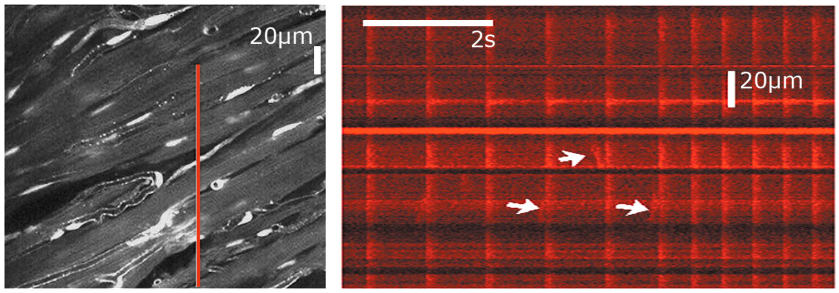

Fig. 10 *Line scan example. Left panel shows a bright field image of a region of cardiac tissue with several visible cells. The vertical red line represents the position where the microscope is recording to obtain the line scan, in this case crossing seven cells. The right panel shows the line scan where one can observe the temporal evolution (time going left to right). Many calcium transients can be seen, approximately one every second, and the arrows indicate calcium waves travelling within single cells. Image taken from work by Rubart et al. [24]*

1.4.4 Fluorescence

Fluorescence is a natural phenomenon by which an object will absorb light of a certain energy and re-emit it in a lower energy. The habitual mental concept of a spontaneous emission of light is erroneous and is due to the fact that an object apparently emitting light is actually being illuminated by higher energy light that the human eye cannot perceive (such as ultraviolet). To understand how this absorption and emission works a little knowledge on the interaction of electromagnetic radiation and matter is required.

Electromagnetic radiation is an oscillatory perturbation in the electric and magnetic field that propagates through space and is responsible for the interactions of charged matter. The frequency of this oscillation determines the energy of the radiation, high energy radiation oscillating faster than low energy, and this energy determines the possible interactions with matter. What is commonly known as light is a small range of the whole spectrum of energies electromagnetic radiation can cover, it is the most energetic part of the range that the Earth's atmosphere does not filter out and, for this same reason, the range the human eye can perceive.

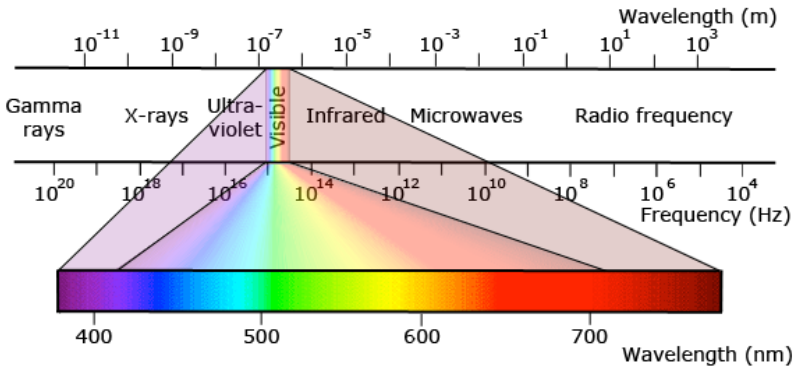

Fig. 11 *Electromagnetic spectrum with the associated orders of magnitude in frequency and wavelength. Visible light is the expanded region where wavelength values associated to each colour are seen. White colour is a combination of all wavelengths together, black is absence of light.*

Visible matter is composed of atoms. They are the chemical units of matter and their sizes are measured in picometers (10^{-12}m). Atoms at the same time are composed of three main particles: neutrons, protons and electrons. The first two provide mass, they lie in the centre of the atom forming its nucleus, and are bound together by a very strong force; the nuclear force. The second two provide electrical charge, of opposing sign, and this charge bonds electrons around the nucleus into specific possible states known as shells. These shells (or energetic states) can hold a specific number of electrons and depending on this number it will allow certain atoms to produce certain types of bond with other atoms to form molecules. The transition of an electron from one shell to the other is possible by exciting the electron with electromagnetic radiation, and vice versa, an excited electron can decay to a lower energetic state by emitting electromagnetic radiation.

Substances that are fluorescent (fluorophores) have atoms with electrons in a state that can be excited to a higher state and then decay to the ground state by stopping in at least one intermediate state. In this manner the electrons are absorbing light of a certain energy when being excited, to re-emit it in at least two different lower energies.

Some of the parameters that characterise a fluorophore are: The peak wavelengths of both the absorption and emission spectra (i.e. the colour of the light used to illuminate and the colour of the light it will then re-emit), the quantum yield (the ratio between the intensities of the absorption and emission light), the extinction coefficient (the relation between the intensity of light absorbed and the concentration of the fluorophore) and the lifetime (mean time the excited state will take to decay to the ground state).

The decay paths the electrons take after being excited determine the effectiveness of the fluorophore. If the decay is slow, generally by going through close to stable states, the fluorophore is less useful since the light emission at a particular time cannot be directly related to the simultaneous concentration of the fluorophore. If the decay is fast the fluorophore is more reliable since the light emitted is proportional to a quasi-simultaneous concentration of the fluorophore. This decay being a quantum process, we cannot ensure a certain decay path will take place, all we can do is to select those fluorophores that have empirically proven

certain decays. Moreover, the fluorophore will tend to change its chemical properties under continuous exposure to light, depending on temperature, the concentration of the fluorophore and obviously the chemical composition of the fluorophore itself and, in practice, it will deteriorate in time. All these combined effects produce what is known as photobleaching.

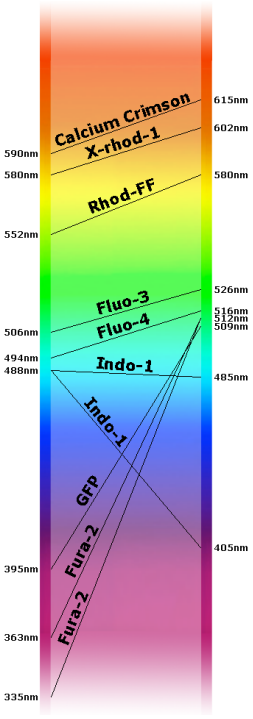

Fig. 12 *Some calcium indicators with their excitation and emission wavelengths. Those that have either two excitation or emission wavelengths are ratiometric. Wavelength values are taken from [25].*

Fluorescence microscopy is the technique by means of which one uses a fluorophore that has been designed to attach to a target biological structure to reveal its position and allow its viewing under the microscope. Non-specific staining is a general term used to refer to the fact that on some occasions a fluorescent signal will be emitted from an undesired location in an experiment. An example of this can be seen in figure 29, subfigure *a*, where the fluorophore has attached to cellular structures other than the targeted RyR2. Fluorophores can also be designed to become fluorescent only when bound to a particular chemical or molecular structure, hence they can be used to measure concentration of the chemical.

In fluorescence microscopy, one loads the specimen with a fluorophore (or genetically encodes it in), then illuminates with light of the appropriate energy for the absorption and records the emission using a filter to block light of all other wavelengths. The first fluorophores were derivatives of the green fluorescent protein (GFP), first extracted from a kind of jellyfish in the 1960s [26] but not applied in life sciences until the 1990s [27, 28]. For the particular case of measuring Ca^{2+} concentration there are many available fluorophores (fig. 12) and one of the most common of them is Fluo-4. The excitation wavelength peak of Fluo-4 is 494nm and the emission peak is 516nm. Using this fluorophore we would illuminate the sample with laser of a wavelength around 494nm, record the emission with a filter around 516nm and the intensity of light recorded would be directly proportional to the intracellular Ca^{2+} concentration. Ratiometric fluorophores either require two excitation wavelengths or emit in two different wavelengths and the concentration is measured as a ratio of the intensities.

1.4.5 The diffraction limit

Due to the wave nature of light there are some considerations to be made when imaging objects the size of which is close to the wavelength of the light we are using to view them. Diffraction is a term used to describe some of the effects of this wave nature and it can be understood by considering Huygens' principle which states that 'Every point in a flat wave-front can be considered as a source of a spherical front' (see figure 13).

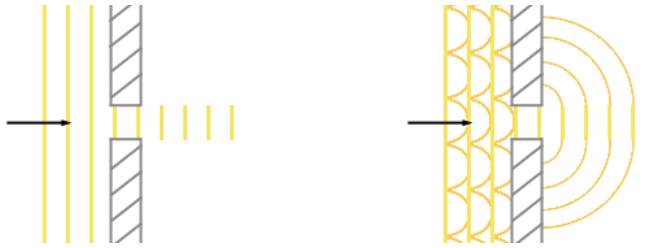

(a) *Wrong interpretation of a flat wave-front encountering a slit.*

(b) *The same situation with Huygens principle being applied.*

Fig. 13 *Left panel depicts the naive idea of light encountering a slit. Right panel depicts Huygens principle where we imagine every point in the wave-front as a source of a spherical front, so that after the slit light can reach regions we would not expect.*

The most important implication of the principle is that light is capable of changing direction after encountering an obstacle, therefore reaching places one would not expect from a naive point of view. If an object is illuminated by a light source one would expect that in the space behind the object, where the shadow is cast, absolutely no light would reach. Applying Huygens' principle, someone standing in the object's shadow sees light coming from the edges of the object.

Interference is another property of light also due to its wave nature. Two identical waves that meet in a point in space can produce very different results depending on the phase of the waves. If at the point where they meet both waves are at maximum amplitude the resulting wave will be of double amplitude whereas if one is at maximum amplitude and the other is at minimum amplitude the two waves will cancel out. The two waves are interfering with each other.

Taking both phenomena into account we discover that light going through a pinhole will propagate in all directions from the pinhole and furthermore it will interfere with itself since every point in the pinhole can be considered a different source of light. If we set a screen behind the pinhole on which to project the light that goes through the pinhole we will not see an exact image of the hole, rather we will see the diffraction pattern of the hole, with dark regions due to interference. This diffraction pattern depends on the size and shape of the hole, the wavelength of the light used and the medium the light is travelling through.

This is the reason why we define the point spread function (PSF) of an optical setup. The PSF is the image of a point source of light through the optical setup, and it is not a point, it is instead a pattern due to diffraction and interference, that depends on the optical setup and the light itself. Optical aberrations will also contribute to the PSF of the system, see Hecht & Zajac, chapter5 [23] for more information. Any light-emitting object with a certain extension has to be understood now as a series of points emitting light that cover the extension of the object. The image of the object through the optical system will be the sum of the PSFs of the

system at every location of the object. This procedure of summing a pattern by locating it at every position is known as convolution. If the objects being viewed are much bigger than the PSF of the system then the image is hardly affected by it but when viewing objects of sizes close to the PSF the image is going to be altered by the pattern of the PSF. In conventional photography the mental concept of an image is that of a projection copy of the scene we are photographing. From now on we need to set aside this idea and think of an image as the convolution of the scene by the PSF of the optical system.

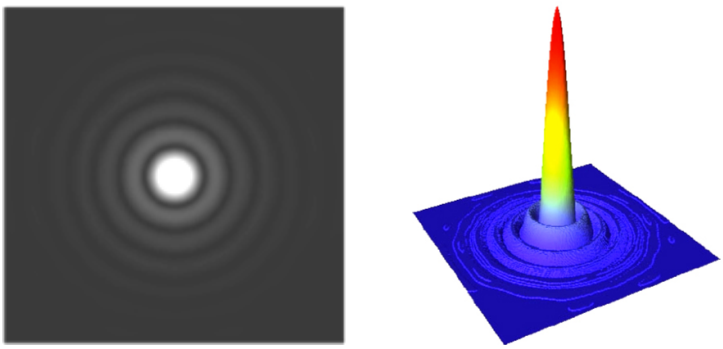

(a) *Image of a circular pin-hole.* **(b)** *Intensity distribution.*

Fig. 14 *The image of a circular pin-hole is the Airy pattern. Left panel shows a real image obtained on an optical bench and right panel depicts a simulation of the intensity distribution. The diameter of the central disk is the diffraction limit of the optical setup. Images courtesy of Optics and Photonics Unit, Department of Applied Phisics, Universitat de Barcelona.*

For the case of an optical setup with a circular aperture the PSF has the morphology of the Airy pattern (see figure 14). This pattern has a very bright, central, circular zone and a series of concentric bright rings decreasing in intensity the further away they are from the centre. The radial distribution of this pattern is a *sinc* function; $sinc(x) = \frac{sin(x)}{x}$. The central circle is known as the Airy disk and its size is our main handicap in microscopy imaging. The diameter of this central bright disk is what is known as the diffraction limit of the optical system since it is the smallest size resolvable by the optical system, typically in fluorescence micoscopy around the quarter of micron [29]. In other words, if we have two point sources of light that are closer to each other than the diffraction limit, their images will be superimposed and we shall not be able to tell they were originally two separate sources (see figure 15). Another less-restrictive definition of this minimal resolvable distance is by using the radius of the Airy disk instead of the diameter and it is known as the Rayleigh limit [23].

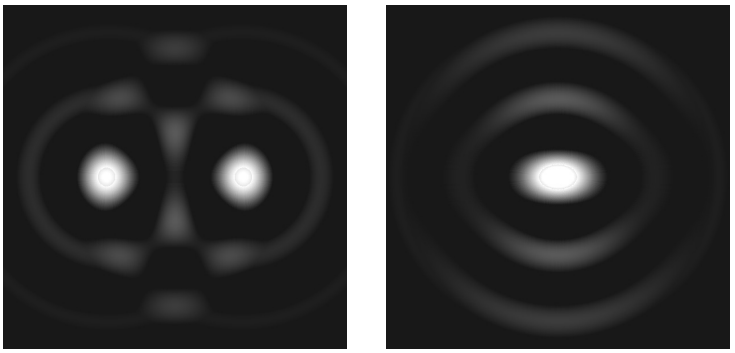

(a) *Separation* $\sim 4R$. **(b)** *Separation* $\sim R$.

Fig. 15 *The diffraction limit. Two points emitting light appear to form an image with the two corresponding Airy disks. In the left panel the separation of the sources is around four times the radius and so the disks are separated. In the right panel the images appear superimposed because the separation is smaller than the disk's diameter, and so the objects start to be irresolvable. We cannot tell the number of point sources under the pattern. Also interference can be observed in the outer rings of the left panel.*

1.5 Digital data processing

This section covers a series of general key concepts in the field of digital data processing, both for signals and images, that are necessary before entering the actual thesis body.

1.5.1 The normalisation problem

Normalisation is a key point in the preprocessing stage of a processing pipeline because it allows certain image transforms to be applied later on in the pipeline. When batch processing images from which data is going to be extracted one needs to ensure is that all images have been taken under the same conditions so that the brightness intensity of objects can be compared between all images. In general this is not necessarily possible. We may have to compare images taken with different optical setups or under different ambient conditions and therefore images need to have their intensity normalised.

Normalisation is the process by which one establishes a reference minimum and maximum value (generally zero and unity) in order to measure a magnitude regardless of the units it is measured in. The main purpose for normalisation is to fix the range of values a magnitude can take to an easily comprehensible and normal range for our perception.

For example, someone may say *The mean distance from Earth to the Sun is 149597870700m* [30]. This is not much of an intuitive sentence since on one hand we are not accustomed to thinking of numbers that big and on the other hand the meter is not an appropriate unit for measuring something that is millions of millions of times bigger. We might just as well say *The distance from Earth to the Sun is huge* and to most of us it will be just as informative. For this

reason, when measuring distances in the Solar System, astronomers define the astronomical unit (AU) as the distance from Earth to the Sun so the sentence becomes *The mean distance from the Earth to the Sun is 1AU*. Now any other distance can be referred to this astronomical unit and it will easily enter our mind without confusing us. For instance we can now say *the mean distance from Venus to the Sun is 0.72AU*. We have normalised the distances to the astronomical unit.

In the particular case of digital image processing the magnitude that needs to be normalised is the brightness or intensity values of the pixels in the image and the normalisation process is usually referred to as histogram stretching or level equalisation.

From a mathematical perspective we would say:

$$I_n = k\frac{(I_o - min(I_o))}{(max(I_o) - min(I_o))}$$

where I_o is the original image, I_n is the normalised image and k is the normalisation constant. We are making the darker pixel in the image appear black and the brighter pixel in the image appear white.

The dynamic range of a digital image is the range of possible values the image takes. A pure black and white image (no grey, just two colours, black and white) has a dynamic range of two values. This hypothetical image is very poor in terms of human perception but it has the advantage that it fills very little memory when recorded since we can use only one bit of memory for each pixel. It is generally accepted that a minimum of 8 bits (one byte) is required so that the intensity transition becomes smooth for human perception [31]. In this case a greyscale image (what is known as a black and white image) has 256 possible tones of grey and the dynamic range is set in the interval ranging from 0 to 255 (255 being the normalisation constant). The darkest pixel in this image, pure black, should take the value zero, the brightest pixel, pure white, should take the value 255, and any value in between the two will be a grey pixel. The transition of brightness is said to be smooth to human perception due to the fact that two pixels differing one unit are indistinguishable to the human eye, meaning that a pixel with value 254 will be as equally white to our eye as a pixel with value 255. The higher the dynamic range in the image the higher the precision we will have when measuring brightness but at the same time the higher the memory usage will be when storing the image. When recording an image it might happen that we do not use up the whole of the dynamic range. This can be due to many reasons, intended and unintended, but in terms of resource and memory usage the image is badly exposed. We may allow 256 possible values for each pixel but if in the image fewer than 256 values are present we are reserving too much memory for each pixel.

As an example, imagine two images of the same field of view taken with the same exact experimental conditions and optical setup except that the second is using a smaller aperture

and therefore is darker. In this case the two images will not be using up the same dynamic range, the second image with the darker light conditions will be using a smaller interval of the dynamic range (and closer to zero) and so an object appearing in this second image will have lower intensity values compared to its intensity in the first image. In order to be able to compare the two images, we will need to normalise both images to the whole of the dynamic range, so that in both images the darkest pixel takes the lowest possible value and the brightest pixel takes the highest possible value. Once the two images have been normalised they should appear very similar to the eye, although the grey transition of the first will be much smoother since it was originally using a higher fraction of the dynamic range. Nevertheless, once the two have been normalised, an object in both images will appear with similar intensity values in the two images.

Now imagine we take a third picture but, in this case, in the field of view we have a new object that is brighter in intensity than anything else in the image. If we normalise this image, the brighter object will be setting the maximum value and in this case the image will not be comparable to the other two because all the objects that in the other two images were equally bright (after normalisation) will now appear darker in this third image. This is when the normalisation problem arises. An example is shown in figure 16 with a lamp post playing the role of the bright object.

Under the assumption that in a set of images they were all taken with the same exact optical setup and experimental conditions, the images should not be normalised, or at least they should not be normalised to their own range, instead they should be normalised to the overall dynamic range of the whole set, so that the intensity of objects remains unchanged independently of whatever is in the field of view. On the other hand when this assumption cannot be made, images have to be normalised to their own dynamic range but, in this case, one needs to either ensure the brighter and darker objects in all images are the same or at least be able to tell when brightness has shifted. If this is not the case a solution is to reference the normalisation to a known object that appears in all images but that of course will require further processing for the object will need to be found in each image.

In summary, each set of images will require its own treatment depending on what is being pictured and how the images were taken (even on who took them), and normalisation is an important step that needs to be treated carefully because it can easily distort image content.

Normalising fluorescence signals

All that has been said for conventional photography is equally applied to fluorescence imaging. We still have to consider image content when deciding on how to normalise an image and especially a set of images, and we also still have to consider how the images were taken; the image intensity values are still dependant on all the variables already mentioned

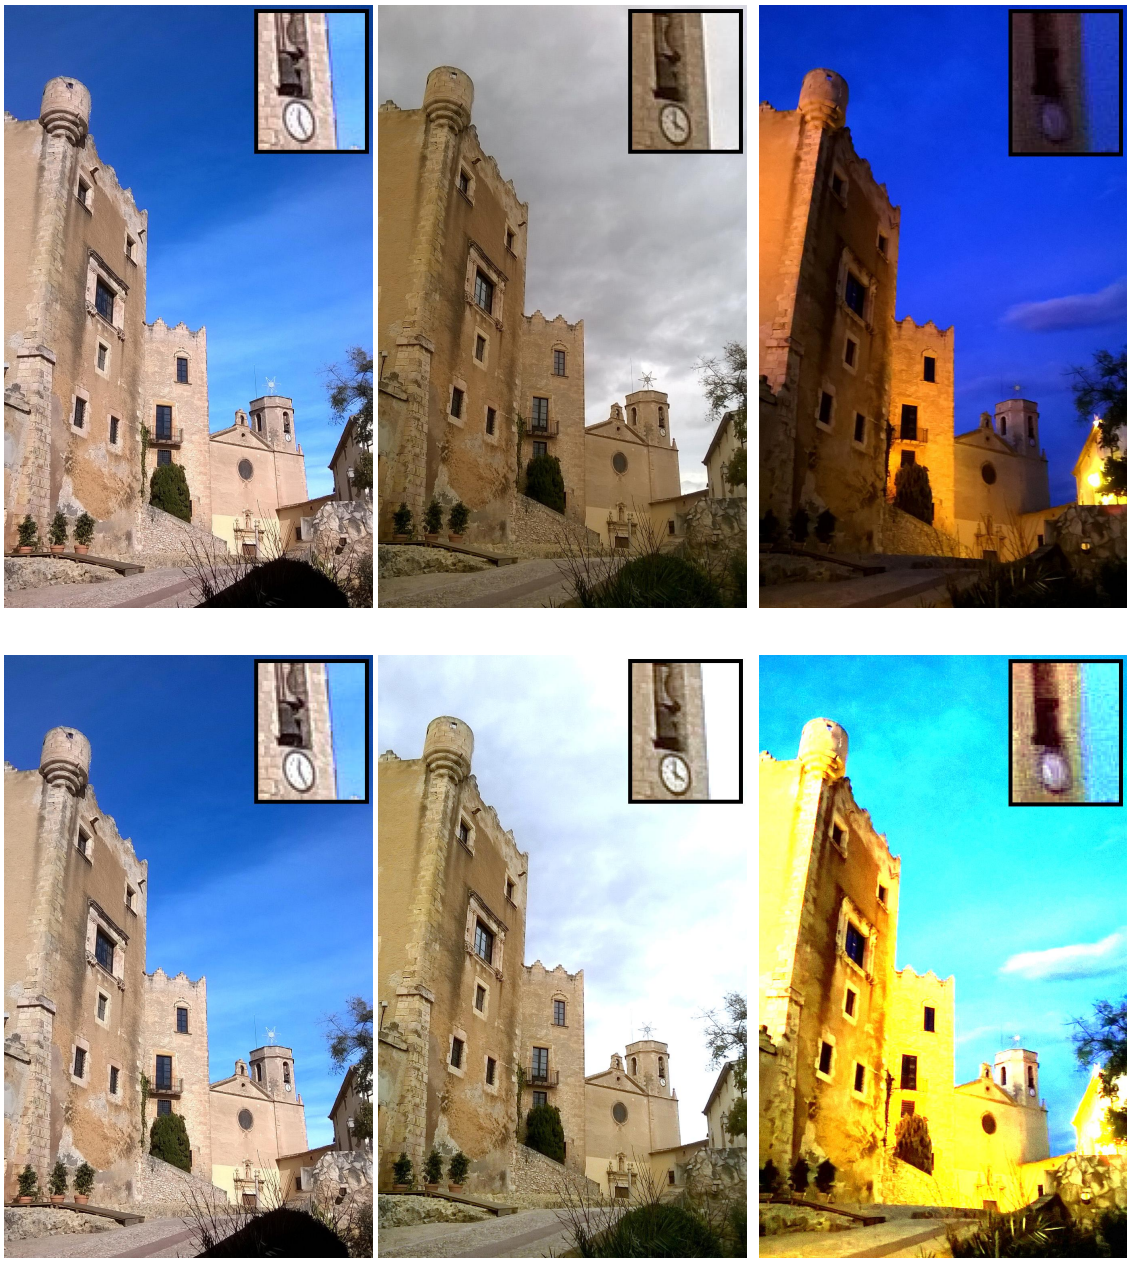

Fig. 16 *The normalisation problem. The top row shows three images of Altafulla castle taken in different light conditions. All three have been normalised to stretch the dynamic range to the maximum [0,255] and a region of interest is shown with the detail of the clock and bell in the church steeple (rightmost building).*
The bottom row shows a renormalisation using the brighter pixels belonging to the clock for the maximum value and the darker pixels belonging to the bell shadow for the minimum value. Regarding the whole field of view, the first image hardly changes, but the second and especially the third have saturated regions. On the other hand, the steeple wall in the region of interest now has similar colour tones in all three images.

(exposition, aperture, detector sensitivity). There is though another variable to take into account: the emitted light depends on fluorescence, both on the fluorophore concentration and the excitation light intensity.

This produces additional experimental variability and makes comparison of images harder. The way this has historically been overcome is by normalising the fluorescence to the basal level of fluorescence in the image or sequence. That is to find the darker areas of the image and normalise to the mean fluorescence in that area F_o. Authors use two variants of this normalised fluorescence value, one is absolute normalisation F_a and the other is relative normalisation F_r:

$$F_a = \frac{F}{F_o} \qquad or \qquad F_r = \frac{F - F_o}{F_o}$$

Both are incomplete normalisation procedures in the sense that they consider only the lower values and not the higher. This is done to remove the fluorescence offset due to concentration and or intensity of the light, but the normalisation problem is still present when attempting to compare different images or different sets of images.

Moreover there is a new problem that has to be taken into account and that is that when basal fluorescence is very low, the normalised fluorescence becomes very large due to division by a close to zero value (see figure 17).

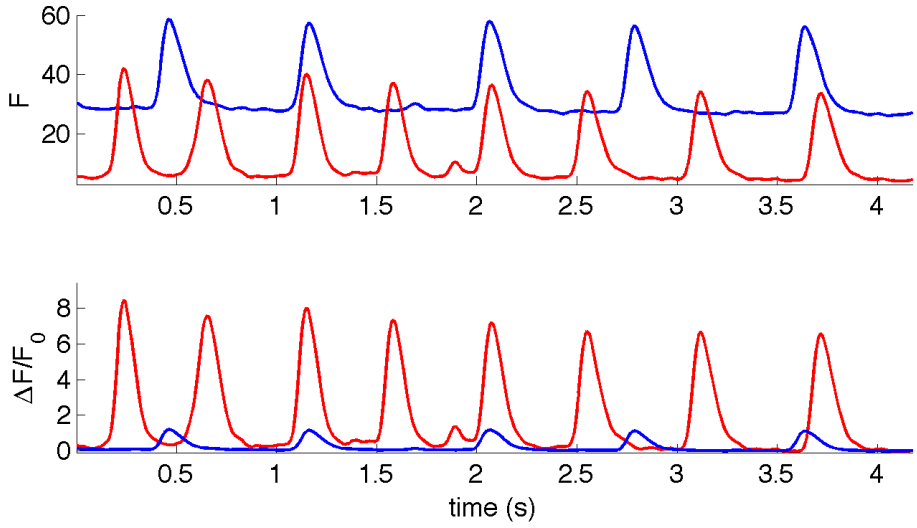

Fig. 17 *Normalising fluorescence signals may lead to artifacts in amplitude. The top panel shows two sample signals taken from different places in an image with one with a higher baseline value (blue). The bottom panel shows the same two signals normalised using $(F - F_o)/F_o$. As the red signal has a close to zero baseline in the top panel, the normalised version has a much larger amplitude than the blue signal.*

1.5.2 Baseline estimation

Normalisation of fluorescence images requires an accurate estimation of the basal fluorescence. The most simple approach is to take the lower values in the signal, some low percentile threshold, for example 0.05, and suppose that that is the baseline. This is a very unreliable method because:

- Noise oscillations may be well above the percentile threshold used, making our baseline estimation really be the lower peaks of the noise.
- Morphology of the signal changes the distribution of the values, so if the signal contains lots of activity the threshold may include signal.
- Most importantly it does not take into account variations in baseline due to fluorophore wash-out or degradation.

The first two can be addressed by either first detecting the activity in the signal and using the rest for baseline estimation or alternatively if the signal belongs to an experiment with a preset resting state (see section 5.1 and figure 58 for an example stimulation protocol) then one can use the mean fluorescence in the window of the experiment without activity as the basal fluorescence.

The last item requires an additional step. One needs to not only know the activity that is present in the signal to exclude the events from baseline estimation, but also to consider the variability in the baseline. If one knows the nature of the variability, the most widely used method is to fit a curve to the values belonging to baseline (see figure 18).

The proposed method in this thesis is to use a sliding window and produce a fluctuating local baseline estimation that adapts to the signal (this is described in detail in section 4.3.1).

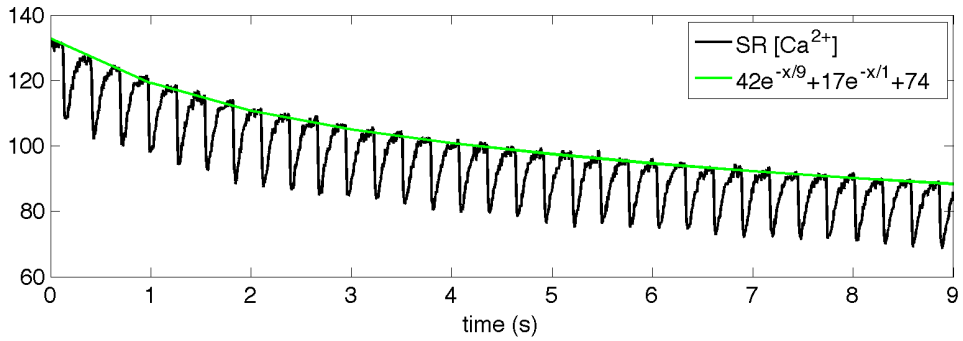

Fig. 18 *Example signal with decaying baseline. The particular signal corresponds to the calcium concentration inside the sarcoplasmatic reticulum, which at each beat is ejecting calcium into the cytosol and so the inner concentration decreases from a basal fluorescence which in turn is also decreasing. For this reason basal fluorescence corresponds to the upper part in each peak. A double exponential is fit to model the loss of fluorescence in time (y axis is in image intensity units over 255).*

1.5.3 Noise estimation

Noise is the part of information that is irrelevant and will tend to hide the relevant part. In image processing it will usually be due to the electronics in the digital detector, either leakage currents produced by long exposures or heat or faulty capacitors in the array, but it can also be due to the experimental conditions such as the fluorophore concentration or binding specificity. Quantifying the noise in a signal or an image is useful to establish an adequate hardness in the noise removal step, especially when dealing with sets of data with different origins, where maybe some will be noisier than others, and we will need to apply filters of different intensity.

Common measures of noise are standard deviation, coefficient of variation and mean absolute deviation, but in general, any statistical measure of variance can be used with the only consideration is that it must only measure noise and not signal. This is easy to say but very hard to achieve. Say for example we take the standard deviation of all the pixel values of two image sequences. The first contains lots of activity, e.g. moving objects, bright objects appearing or disappearing in the field of view, changes of plane, etc. All of this activity will contribute to the calculated standard deviation. The second image sequence is a near to static plane of a very smooth field of view, for example the sky, or a snowy landscape, but in this case, the image sequence is taken under very poor light conditions which produces high variability in the pixels. These two film examples may well produce very similar pixel standard deviation when in fact only one of them has the noise contributing to the variance measure.

A solution to this is to measure the variability using a sliding window that iteratively covers the image sequence. If the window is small enough so as not to capture the real activity in the film, then the noise estimation will not include signal.

It is very common though in calcium imaging that the spatio-temporal resolution is close to the dimensions of the objects/events we are trying to observe and so the sliding window will be too small to achieve a proper statistical measure. The best solution to this is to exclude the objects/events from the noise estimation and take the statistical measure from the background pixels. In the image sequence example of the previous paragraph, if we are capable of first detect all moving objects, then we can estimate he noise by taking the standard deviation of the pixels that belong only to background. This may seem a bit awkward since the noise estimation step will usually be part of a preprocessing phase to ease the posterior object detection, but it is not that strange. One can start the analysis by applying a simple and rough detection step, ideally with more false alarms than omissions, to exclude all areas with signal and then use the noise estimation of the background in the preprocessing phase to clean the data and normalise it to finally ensure a finer and more robust detection step at the end.

1.5.4 Image Segmentation

Image segmentation is a general term that refers to the process of grouping the pixels in an image into separate regions. It can be seen as identifying objects or as labelling areas. It is one of the main blocks in image processing if not *the* most important one and in this thesis it is widely used. The following chapter goes deeper into the segmentation methods applied to calcium imaging but it is mentioned here as an introduction and the two most important methods are briefly described.

(a) *Example image of the chapel outside Tarrés, Lleida.* **(b)** *Ask a computer to identify three regions.* **(c)** *Ask a human to identify three regions.*

Fig. 19 *Image segmentation is a subjective task. A human will take into account its knowledge of the image content whereas the computer needs to be programmed. The segmentation algorithm used in the middle image is a region clustering according to similar colour (k-means). Of the three objects covering the chapel only the middle one is a cyprus tree, the other two are shadows.*

Histogram thresholding

Histogram segmentation is a simple and effective technique consisting in using the hills and valleys in the histogram of an image to cluster the pixels in separate groups. If we rather just want to separate the image into two regions (background and foreground) it can be called histogram thresholding because we use the histogram to find a threshold value to binarise our image. This method, or variations of it, can be applied when content and background pixels have distinctive enough values and at the same time the image has enough pixels to allow a smooth histogram. In calcium imaging of isolated cells one will always have a near-black background versus a bright foreground, which is the cell, and therefore this is a good technique to obtain the cell mask for further processing. When the experiment consists of a sequence of film frames or a series of stacks, obtaining the sum (or average) of the sequence (or series) can sometimes give the pixel value richness that a single image may not have.

In the particular case of histogram segmentation of a cell in a fluorescence image we can encounter several possibilities. In general there will be two *hills* in the image histogram; the bulk values corresponding to pixels belonging to the background and to the right of that (brighter) the values of pixels corresponding to the cell. The background pixels can be around

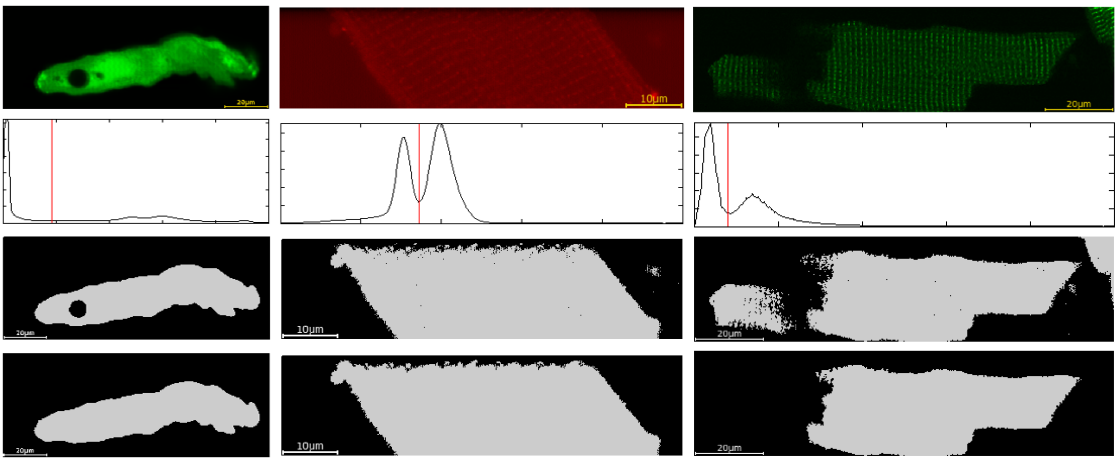

(a) (b) (c)

Fig. 20 *Cell mask evaluation in three different cases. In each of the three, the top panel shows the original image, second panel shows the image histogram and threshold determination in red, next is the binarised version of the image and last is the final mask after filling holes and deleting islands. The first and last exemplify presence of wholes, the second and third a background above zero and presence of islands.*

zero so that the shape of the histogram is only half a hill and in some cases the image can present other bright areas contributing to a third or fourth *hill* in the histogram. Therefore we can say that the first local minimum in the histogram corresponds to a good separation between background and foreground pixels. To further reinforce this in cases with a low number of pixels, one should choose as a threshold the first value (going from left to right) that is a local minimum in the histogram and that is also below a certain percentage of the maximum value to the left of it. Some examples of this are shown in figure 20, in the second row of panels.

Two further steps are generally required after thresholding: the first is to delete any isolated areas in the mask and the second is to fill in the holes in the mask. In the first case, a simple procedure is to find the unconnected regions belonging to the mask and delete all but the one with the largest area. In the second case, a similar procedure is again to find unconnected regions but now belonging to the background and all of those that are not in contact with the image borders can be set to cell mask. Examples of this are shown in figure 20.

Watershed segmentation

There is an infinite family of watershed-based algorithms but they all have in common a step called the watershed transform [32, 33]. The name given to these algorithms is connected to the geographical idea of a watershed (also catchment or drainage basin), that is a region of land such that all the rain water that falls upon it will eventually converge to a single point where it may or may not exit the region. If one thinks of an image as a section of terrain,

with the brighter regions being the mountains and the darker regions being the valleys, the watershed transform can be seen as if we were to pour water on the terrain and draw lines separating regions that send the water to different points. Typically the separation borders will be high elevation points along the mountain peaks and ridges. We can then use further information to limit the growth of these regions, using for example the area covered or the depth of the region. The reason why the family is infinite is because we can choose an infinite combination of parameters to modify or to limit the watersheds.

These algorithms have many applications, and in the case of calcium imaging, by adding a few constraints , they are very powerful for segmentation of excitable cells in a culture. This is expanded in the following chapter and in section 6.1.

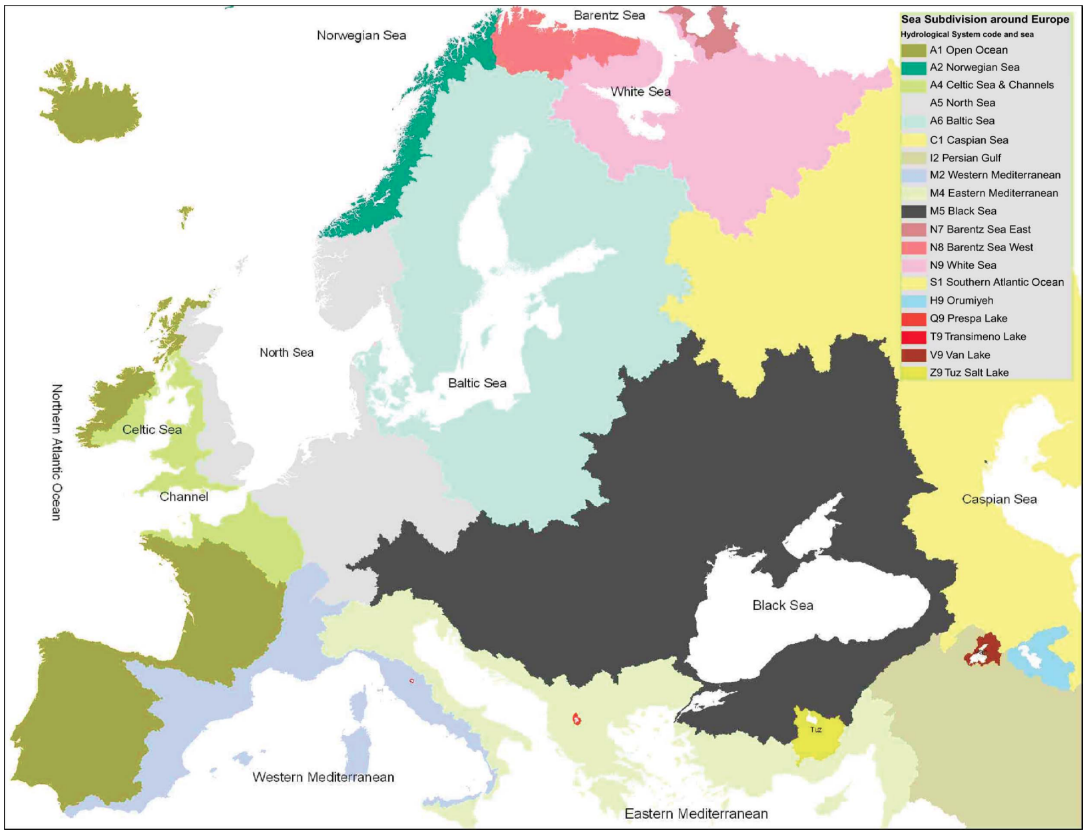

Fig. 21 *European seas and the landmasses draining into each of them. From De Jager & Vogt [34]. Further subdivisions of these watersheds are seen in figure 26.*

1.5.5 Integral transforms in image processing

The Fourier transform

The Fourier transform is an integral transform with many applications in a wide range of fields. It is named after Jean-Baptiste Joseph Fourier who first introduced the idea in his work

on heat flow in 1822. The basic concept behind it is that any function can be decomposed into a sum of periodic functions, meaning that we can express the original function into another function that will now depend on the periods of the functions it is composed of (instead of the original variable it depended on). The typical example is a time function that is transformed into a frequency function, but in general the original variable can be any magnitude, for example in image processing it is a space function that is transformed into a spatial-frequency function.

The definition of the transform is as follows:

$$F(\nu) = \int_{-\infty}^{\infty} f(t) e^{-2\pi i t \nu} \, dt$$

where $f(t)$ is the original function of time (t), $F(\nu)$ is its transform function of frequency (ν) and the term $e^{-2\pi i t \nu}$ is the complex number $cos(-2\pi t\nu) + i\, sin(-2\pi t\nu)$ with i being the complex unit $\sqrt{-1}$. For a particular frequency, we are multiplying the the original function point by point (from minus infinity to plus infinity) by a periodic function and adding up all the terms so, if the original function has some kind of periodicity with the frequency of the periodic function, the maxima and minima will match up and the result of the integral will be a large number. On the other hand, when the frequency is not present in the original function, the result of the integral will be much lower. After covering all possible frequencies, we obtain a function that, at every frequency, will tell us how much that frequency is present in the original function, regardless of the position (temporal position usually). A good example in musical terms is to take the Fourier transform of the signal of a chord, say C major, composed of three separate notes played at the same time (do, mi, sol). The transform will produce three peaks at the frequencies corresponding to these three notes and also at the harmonics of these notes (signals with frequencies that are multiples of each other fit perfectly into one another). The inverse Fourier transform is simply the integral transform that will do the inverse process: given the function of frequency it will recover the original time function.

To put it simply, a function that has very rapid changes will have high frequencies present, whereas a function that is smooth and does not oscillate much, will have low frequencies present. Extrapolating this concept to images, if our image has many details, and so the values change drastically from one pixel to the next, the transform will give us high values at high frequencies. On the other hand, if the image is smooth and does not have great contrast changes from one pixel to the next, the transform will return us high values at low frequencies.

The calculation of this transform may seem tedious, having to integrate the whole range $[-\infty, \infty]$ for every frequency, but it is in fact computationally very fast thanks to an algorithm

already developed by Gauss around 1805 and later re-formulated and published in 1965 by Cooley and Tukey [35]. This algorithm, known as the Fast Fourier Transform (FFT), uses a divide and conquer approach which drastically reduces the computational load and, for this same reason, gives new applications to the transform.

Extending this transform to a two-dimensional discrete function in the spatial domain, that is an image, we obtain:

$$F(\mu,\nu) = \sum_m \sum_n f(m,n) e^{-2\pi i(m\mu+n\nu)}$$

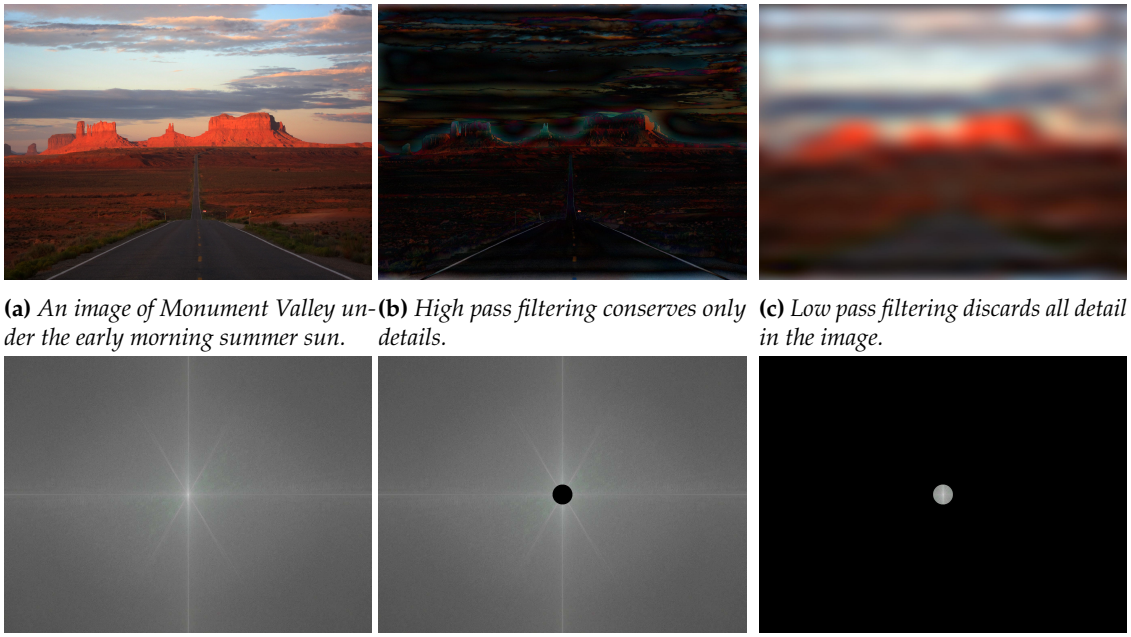

(a) *An image of Monument Valley under the early morning summer sun.*

(b) *High pass filtering conserves only details.*

(c) *Low pass filtering discards all detail in the image.*

(d) *Natural logarithm of the FFT of original image.*

(e) *FFT without the excluded central zone.*

(f) *FFT with only the central low frequency zone.*

Fig. 22 *Example of FFT used for image filtering. Starting from the original image (a), we transform it to its frequency domain (d) and then we can either crop low frequencies (e) to anti-transform and obtain the image (b) or we can crop high frequencies (f) and obtain image (c).*

One of the best known applications in image processing is regarding convolution. As described in section 1.4.5, convolution is the process of multiplying one object by another at every possible position. Thanks to the convolution theorem and the FFT algorithm, convolution is nowadays performed in a fraction of a second. Read more about this in [36].

Another very common application of the FFT in image processing is filtering. As the transform separates the information in our image into frequencies, we can crop an undesired part of the transform and then do the inverse transform to recover the original image without the undesired frequencies. For example one can carry out a low-pass filtering to reduce noise in the image, or a high-pass filter to enhance details, or even a band-pass filter to keep only certain frequencies. See an example of this in figure 22.

Wavelet transform

The wavelet transform is another very useful transform with again many applications in a wide range of fields like data compression, pattern recognition or signal denoising. Its mathematical definition is as follows:

$$W(r,s) = \int_{-\infty}^{\infty} f(t)\,\psi\left(\frac{t-r}{s}\right)\,dt$$

where $f(t)$ is the original function of time (t), $W(r,s)$ is the resulting transform function of both time shift (r) and scale (s), and ψ is the particular mother wavelet function used for the transform. This mother wavelet (or wavelet family) is a function that complies with a series of mathematical conditions but is basically a template of a certain shape that we are using to convolve the original function with. Just like in conventional convolution, if the original function has a similar shape at a certain time, the integral will produce high values for the time shift corresponding to that time, but the power of the method lies in the fact that not only can we test our template at every possible time, but also for every possible scale of the template. For this reason the wavelet transform becomes a very powerful tool for signal morphology inspection.

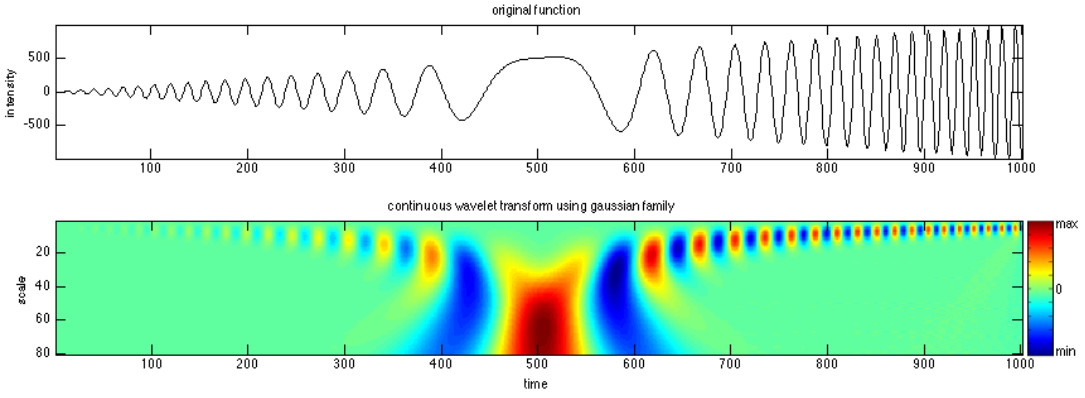

Fig. 23 *Top panel: example signal with increasing intensity in time, decreasing frequency in the first half and increasing frequency in the second half. Bottom panel: Continuous wavelet transform of the example signal using gaussian family. As can be seen the transform is a 2D function that effectively obtains high values (red) at the times where the original signal had a 'gaussian shape', and also identifies the scale of the gaussian shape in the vertical axis (small at the top, large at the bottom). Note how the term scale has a more temporal connotation than in the sense of size, meaning that high scales correspond to long objects more than big objects.*

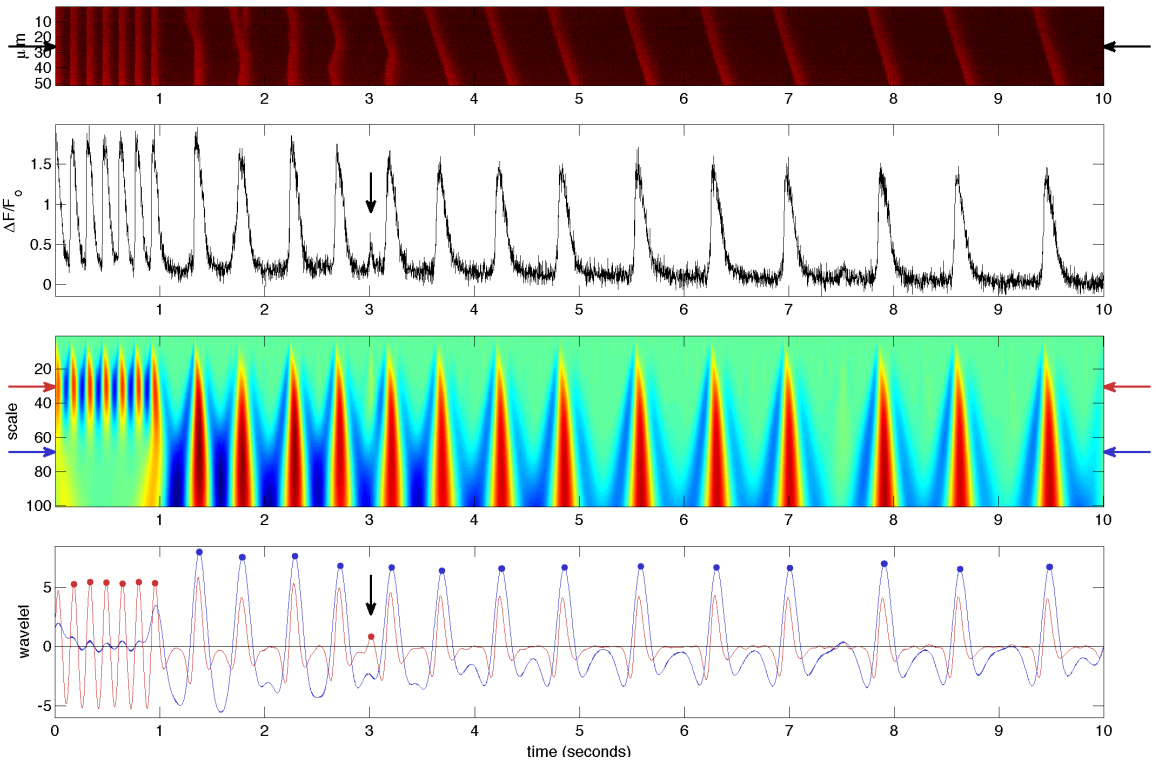

Fig. 24 *Example application of the wavelet transform in detecting events of two scales.*
Top panel: line scan image with some calcium activity. At start a series of transients, followed by a series of calcium waves. The black arrows mark the position from which a profile is extracted for the following panel.
Second panel: A calcium profile extracted from the line scan reveals the presence of a spark (arrow) which was hardly visible in the original image. The profile shown is the mean of a 17pixel central region in the line scan (out of a total of 512) to reduce the amount of noise.
Third panel: Continuos wavelet transform of the previous signal using gaus2 family. Two scales of interest have been marked with the red and blue arrows.
Last panel: The two profiles extracted from the wavelet transform where the arrows are marked in the previous panel. A simple maxima detection function is applied to locate the events (marked with dots) and each maximum is assigned to one of the two scales if the wavelet transform is higher for that scale than it is for the other. Note how the smaller scale (red) enhances the spark and note also how the initial transients are assigned to the small scale due to their short duration.

1.5.6 Local extremal detection

This is a common problem found in many fields and therefore it is covered here for it will be used in several sections further on. Starting in one dimension, we would have a signal of which all the peaks and valleys have to be found. The mathematical approach is to find the places where the derivative of the signal is zero. As our signal is discretised we are likely not to have a zero derivative sample, rather we will have a positive derivative sample followed by a negative derivative sample (or vice-versa), so the computational translation of the mathematical approach is to find between which samples there is a zero crossing of the derivative. The easiest way to do this is to compute the derivative of the signal by subtracting each value from its susequent neighbour and then to find extrema by multiplying each value

in the derivative by its following neighbour. At this point all negative values of the resulting vector are local extrema, those with positive second derivative are valleys, those with negative second derivative are peaks. Here is an example in matlab code:

```
% the original signal is in vector S with length L
D=S(2:L)-S(1:L-1); % first derivative
P=D(1:L-1).*D(2:L); % product of each sample with its following neighbour
E=find(P<0); % vector containing sample indexes for local extrema
```

This simple code is the one used throughout the thesis whenever peak detection of a signal is mentioned. When moving on to higher dimensional data the problem becomes a little harder because null derivatives do not necessarily imply local extrema (in 2D a saddle point is a maxima in one direction and a minima in another and therefore it is not a local extremum). A method that is applicable all the same to any dimensional data and that is used throughout the thesis whenever detection of image local maxima is required is the following: loop through all samples and compare them to the n neighbours (two for signals, in images four or eight depending on wether diagonals are excluded or not). If any of the neighbours has a higher value we set the current sample to zero (or the minima in the set). This leaves us with all samples set to zero except for local maxima. If we are searching for minima the search is performed all the same but whenever a neighbour has a lower value we set the sample to the maximum in the set.

The two methods effectively find all samples that are local extrema, but this does include any noisy bumps the signal or image might have. The first solution one would come up with is to filter the original signal by some low pass filter but this is not recommended for it affects the signal morphology. The usual approach is to detect all extremal points and filter them out afterwards using some custom rule appropriate for the occasion. This could be an absolute or relative height rule (keep only extrema that have a height/depth above a certain threshold) or it could be a width rule (keep only extrema that are at least a certain distance away from other neighbouring extrema), or even a shape rule (keep only those extrema with surrounding samples satisfying certain morphology criteria).

1.6 Thesis outline

After this long but necessary introduction we are now ready to enter the subject of the thesis. The next chapter provides another sort of introduction but in this case to the main methods used in calcium imaging. It is a review on the current state of the art and therefore sets the framework to what has been carried out during the development of the thesis.

The following four chapters contain the body of the work and detail the computational methods developed aiming to solve the typical problems an experimental physiologist encounters while studying calcium dynamics in cardiac cells. This includes detecting objects in signals and images, establishing appropriate features to characterise them, using machine-learning techniques to classify them and extracting statistical information in order to prove or refute hypothesis of either medical or biological implications. The chapters are organised according to the physical scale of the problem to be treated, starting from the cell organelle scale through the cellular scale up to the tissue scale.

The link section that connects the chapters dedicated to one scale to the next should intuitively be found in between the two but instead it is found at the end of the second. For example the section that links the molecular scale with the intracellular scale is found after the chapter describing the latter, so that first you read the chapter dedicated to the molecular scale, next you read the chapter dedicated to the intracellular scale and at the end of this one you encounter the section that connects the two.

At the end of each of the four chapters that constitute the body of the thesis there is a small section with the results of the methods presented in the each chapters, and a summary of the datasets that were used for developing the methods.

The last chapter contains the conclusions to the work as an itemisation of the contributions of the thesis both from a computational and also from a more general biomedical point of view. Also a summary of the most important publications is provided.

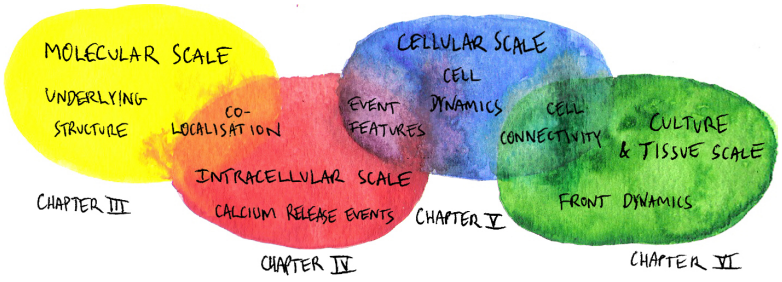

Fig. 25 *Thesis outline. The body of the thesis is composed of four chapters each dedicated to a spatial scale.*

2

State of the Art

This chapter is an up-to-date review of the research in the three main themes covered by this thesis. They are:

- Detection and characterisation of intracellular molecular structures.
- Detection and characterisation of intracellular calcium events.
- Characterisation of the calcium dynamics of single cells and cellular cultures.

All three contain the word characterisation. In the first two it refers to measuring the morphological features of the objects detected. In the third case it is considered in a broader sense and consists in defining a closed set of dynamic regimes and classifying amongst them. Most importantly though, all three of them are based on image segmentation. The first two contain the detection part, in the first case static structures, in the second, time-evolving events and the third theme also contains a segmentation part because it requires detecting the cells in the culture as the first step.

Image segmentation is a very common problem in image processing, if not the most common, and in many cases a crucial step. This is because it is usually one of the first blocks in a processing pipeline and the following analysis heavily depends on it. It consists in separating the relevant or foreground pixels in a digital image from the irrelevant or background pixels according to specified relevancy conditions. There are many studies that treat the general problem of image segmentation together with methods for evaluating its performance and they all agree on the fact that segmentation is a very challenging problem and an open one [37–42]. The main reason for this is that it is, as mentioned, a process that depends on predefined conditions and therefore it is essentially subjective. On top of that there is the fact that raw images can have very different properties and content (see

the normalization problem in section 1.5.1) which overall makes it very hard to choose a particular method or technique. This is the reason why there are so many studies that either focus on a particular type of image and compare different segmentation algorithms [43–46] or, the great majority including this thesis, that attempt to develop new algorithms by applying or combining or modifying existing techniques to their particular field of research.

Our case is restricted to fluorescence images and particularly blob segmentation (sub-resolution objects), spark segmentation and cell segmentation. The following sections cover the methods applied in each of these cases. There is one general tip for image segmentation in fluorescence imaging (and probably any other kinds of imaging) that is worth mentioning for it is used several times later on and is regularly stated in the literature. It is that if one wants to guarantee no omissions, it is always better to be generous and over-segmentate by having more false alarms and then to use a subsequent filtering step where one can discard the false alarms [47, 48].

2.1 Image Processing Tools

Some of the best known general-purpose tools for bio-image processing and analysis in fluorescence microscopy are ImageJ [49], CellProfiller [50, 51], BioImageXD [52], Icy [53], amongst many others. These are high-level examples which do not require great computer science skills to use, but there are also sets of libraries that can be called from code, and in fact are used by the already mentioned, containing the basic image processing functions. Examples are ITK [54] and openCV [55]. Subsets of these are continuously being reviewed, expanded and tested for specific applications and under specific conditions [48, 56] because, as already mentioned, the variability in image conditions makes image processing a very artisanal subject where each problem has many possible solutions: many acceptable, none perfect. For this reason there is still no unified approach or gold-standard and a great amount of literature available in the field. The specific applications covered in this thesis are discussed in the following sections.

2.2 Blob Segmentation

Blob or particle segmentation refers to the detection of isolated objects that are smaller than the diffraction limit of the optical setup. Isolated because the distance from one to the other has to be larger than the diffraction limit and so they can be resolved as independent PSFs (see section 1.4.5 for details). This is a common problem, on the one hand for straight localization of molecular structures (and also applied in astronomy images for star detection [57]), and on the other for microscope PSF estimation purposes in which luminous beads of known size are imaged [58].

There have been many developments in recent years towards overcoming the diffraction limit. Superresolution methods such as the early STED [59, 60] or the more recent STORM [61] or PALM [62] allow resolving light sources that are closer to each-other than the system diffraction limit by means of special fluorescence technology and taking thousands of images. The basic idea is that by using fluorophores that have been engineered to be switchable and with a low probability of activation, much lower than the time between images, then in a single image of the sample the chances are that we record light from a small subset of emitting molecules. By taking many images of the same cluster of molecules we can observe slight shifts in the PSF and therefore infer the individual locations of emission and reconstruct a superresolution image of the structure [63]. The sample needs to be treated during the experiment to remove the subset of light emitting probes that were switched on and allow another subset for subsequent images .These procedures have proven very successful and produce impressive images with very high detail, some ten times higher resolution than was thought possible by means of visible light (in nanometers from the hundreds down to the dozens [64]). However, these techniques are expensive and time consuming, and are unnecessary when the required precision in localisation is not below the few hundreds of a micron.

In conventional optical microscopy all sub-diffraction limit objects appear as blobs. The best case scenario is that all blobs are of homogeneous intensity and morphology against a background that also has homogeneous intensity. Then, more or less any approach will work, and possibly the most widely used is simple thresholding [65]. The image is binarised according to a single intensity value that separates blobs from background. There are many ways to choose a value the best known being Otsu's method [66] in which the threshold is the value that minimises the variance in the two classes. Example applications are by Allalou et al. [14] where Otsu is used as a first step and then is combined with a watershed method and Ghaye et al. [67] that use a local thresholding method thus overcoming intensity variability throughout the images. Similarly histogram based methods will attempt to find the threshold based on the morphology of the histogram, by selecting a local minimum that separates the pixel values into distinct populations (see section 1.5.4).

Other common methods are those based on image morphological operations. One of the most famous of these is the top-hat filter. It consists in applying an opening filter (erosion followed by dilation) and then subtracting the result image from the original image. The erosion suppresses any areas smaller than the structuring element (presumably the blobs), the following dilation recovers the original image without the blobs which, when subtracted from the original image, only the blobs survive. The method performs very well when objects have very clear borders or alternatively are very bright because the image can be thresholded and a binarised version used for segmenting. Examples of applications can be found in [68, 69]. The other very widely used morphological operation is the h-dome, also named h-maxima

transform. This image transformation deletes all maxima that are below a certain height h while lowering all the other areas to the same height. Its main feature is that the maxima are all then set to a common height, which can be both an advantage or a disadvantage. Another feature that is an advantage over the top-hat transform is that it is independent of object size and shape. Its main counterpart is that it is heavily dependent on the height h which requires a very precise estimation and so the method is usually combined with other methods to maximise performance [70–72].

Another family of methods are those that consist in some sort of filtering in order to reduce noise and to enhance the target blobs and subsequently apply any of the previously mentioned to segment them. Examples are techniques that model the blob by a Gaussian function, including the one proposed in this thesis, and apply a filter accordingly, or others based on wavelet decomposition in which the image is filtered by choosing appropriate coefficients that remove noise and large structures [73–75]. The most common of all blob detectors though is the Laplacian of a Gaussian filter, also called Mexican hat, which as the name suggests has a bell-like central shape that decays rapidly. An example is by Sage et al. [76], who use the Mexican hat filter in a particle tracker algorithm and also demonstrate that it is the optimal detector for Gaussian-like blobs under certain noise conditions.

For a review and comparison between several blob segmentation methods see Smal et al. [77] and Ruusuvuori et al. [78].

2.3 Event Segmentation

Calcium events is a term used to refer to a localised increase in calcium concentration, both spatially and temporally, within an excitable cell. There are several types of events classified depending on their spatial and temporal scale (see table 5 for a summary), but of these, the most studied is by far the calcium spark. Since its first description in the 1990's [79] there have been many attempts at producing a robust spark detector and we will now briefly review them.

One of the most widely used methods was developed by Cheng et al. [80], which is a thresholding method based on an estimation of background noise and setting the threshold to a factor times the background standard deviation. The factor is user-inputted and therefore depends on the image type and conditions. This method is surprisingly still in use [81–84] mainly for the simplicity of its implementation and the comprehensiveness of its performance.

Other thresholding based methods that are in fact variations of the previous one are for example cases in which Cheng's method is combined with other methods for performance improvement with an automatically estimated threshold [85], or with an improved background noise estimation [86], or with relaxed threshold, relying on precise morphology post

filtering [87–89] or even with user-supervising to reject or accept candidates [90]. Others have taken into account the temporal nature of sparks, for example by thresholding the difference of consecutive frames [91] or by adapting the threshold considering the previous frames [92].

Just as in the case of blob segmentation, there is a family of methods that are based on some sort of template matching which altogether smoothens out the images and enhances locations of sparks. In all of these cases a posterior method for determining the exact spark location is required and it generally involves some sort of thresholding. As an example, Kong et al. [93] detect sparks in line scans using a hand-made spark template which is cross-correlated through the line scan and local maxima are recorded. Statistical tests are performed to assess the correlation and in this case the threshold is the significance applied to the p-value.

Within the template-correlating family included are the methods based on the wavelet transform. As described in section 1.5.5, wavelet transform is very powerful for it allows detection of objects that have a particular shape but do not have a fixed size. It is a kind of correlation by a template where the template is modified through a range of scales. Again the method requires some sort of thresholding after the transform has been applied in order to detect the exact location of the sparks. Examples of methods that use wavelet transform can be found in [94–97]

2.4 Cell Segmentation

Cell segmentation is a very broad term due to the vast variety of microscopy imaging techniques. As mentioned in the previous chapter for the general case of biomedical data processing, when developing a tool for cell segmentation one needs to find a balance between producing a solution with a wide range of applicability and at the same time not losing performance in the individual cases. On the other hand, it is also very common to find papers that propose a particular solution to a problem that is so particular that it will hardly be useful elsewhere or are simply using a very small set of images lacking statistical significance [98]. A curious example of a highly tailored solution to specific image conditions is by Thompson et al. [99].

As in the previous cases, the simplest approach is image thresholding which are fast methods but highly dependent image conditions and therefore only applicable under very specific situations or as a step in a larger segmentation pipeline [100]. The renowned Otsu method dating from the late seventies is still applied in some cases, generally as part of a combination of methods, for example in the work by Amanfu et al. [101] which is discussed in the following section. An example of a thresholding method using commercial software that demonstrates the limited range of application is by Walsh et al. [102].

Generally speaking, in cell segmentation, the approaches that produce the best results are contour based, in which search for local maxima in brightness to detect the border

between cell membrane and background [103]. They are the best results in the sense that the pixels belonging to the cell will be more precisely identified from the background, not necessarily from other cells [104]. These contour based methods are only applicable when there is a clear border between cell and background [105] and therefore require certain types of imaging (e.g. phase contrast, bright field), and single cell images (as opposed to cell cultures where clumping may occur). Examples of contour based methods are segmentation via edge detectors [106, 107], contour shape and morphology-based [108, 109], active contour models or snakes [110–112] or the subgroup of these, level-set based [113–115] and combinations of contour-based with others like local thresholding [116], or with watershed [117].

In calcium imaging cell borders are hardly visible and cell segmentation usually becomes cytoplasm segmentation. This is due to the combined facts that light is originated in the fluorophore that is in the cell cytoplasm and also that due to the aqueous nature of the cells, being laid on a slide they do not have constant thickness. Instead they rather rest in a manner having a thick centre and spreading out with shallow thickness towards the edges, in a mountain-like fashion, or in Dr Hove's words "like fried eggs in a pan". This makes cells in fluorescence images appear with high intensity in the central region and gradually dimming towards the cell membrane, where it fades out to the background without a clear border. Despite this there have been some attempts to use contour-based methods like the work by Srinivasa et al. [118] and Bergeest et al. [119] where a combination of methods based on active contours and minimising certain energy functionals are applied to hard problems in fluorescence imaging to produce remarkable results. Nevertheless, in general when dealing with lack of borders, the other great family of cell segmentation methods arises, those relying on smooth transition from the pixels belonging to the cell to the background. Examples from this family are statistical- or clustering-based methods [120–122], region-growing-based methods [123, 124] or, most importantly the watershed-based family [125], which are currently the preferred methods for their balance of performance and computation time [126, 118]. The watershed methods are discussed at the end of this section.

All of the mentioned have both semi-automatic versions, requiring user-inputted data, and pure automatic versions. The user-input data usually comes in the form of either an initial point from which the clustering, region growing or watershed will start [127–129], or an initial line or contour from which the edge detection or active contour method will start [130, 131]. These semi-automatic versions tend to perform better due to the clear advantage of the human inputted data [46, 48] and the fully automatic tend to rely on a post-processing step to remove, merge or join errors [132]. Some authors have successfully attempted to produce a seed detector algorithm to produce hybrid methods [133, 134].

Another family of methods that has gained recent popularity are those that apply graph theory to image segmentation [135]. Graph theory is a branch of maths that formalises the study of sets of objects that have pairwise relations (see Rubinov et al. for an introduction

[136]). In this sense an image can be treated as a regular lattice of points (the pixels) with connections to their neighbours. This allows distances to be established between the pixels and then the graph-partitioning methods can be applied to cluster them in regions, thus segmenting it [137]. Another graph-based interpretation of an image that is very close to the idea of watersheds [138, 139], builds a tree of nodes in which the root node represents the whole image and the leaf nodes are each individual pixel. The nodes at the branching points in the tree represent areas of similar pixel values according to some established distance metric, and again, graph-partitioning methods can be applied in this case to establish a pruning rule of the tree [140]. As an example, using figures 21 and 26, the North Sea catchment basin (figure 21) would be represented by a node from which several branches would grow, one of them being the river Thames basin from which many branches would grow, one for the basin of each tributary (figure 26). This modelling produces impressive results in real-world images because it can easily handle the complexity found of shapes, borders, sizes and brightness, but it is also computationally expensive and altogether is hardly applied in cell segmentation [141].

As explained in the previous chapter, section 1.5.4, watershed-based methods understand the image as if it were a surface and then use a rain catchment basin idea to separate it in regions. There are two interpretations of the watersheds and both can share a common result [142]. The first is the rain model in which water falling on each point will flow down the surface and end up at a certain point, either inside or at the surface limits; the whole area sending the water to that point is the watershed. This is the interpretation for the method discussed in this thesis and some examples of similar methods found in the literature are mentioned in the following paragraphs. The other interpretation is the flooding model in which the surface is flooded from sources placed at particular points so that the areas that get filled up from these points are the watersheds and when two watersheds meet, a dam is setup to separate them. The flooding points, also called seeds, can be set as regional minima [143], user-inputted [144] or randomly placed [133]. The two interpretations give the same results when enough seeds are taken to cover the whole image

There are also two interpretations of the term watershed itself, again analogous to one another. The first is that the actual regions are the watersheds [145] and the other is that the border or dam separating the regions is the watershed [146, 147]. Obviously, these methods can be used both to detect local brightness minima and also local brightness maxima by inverting the image. A not so obvious but very common approach is to use watershed-based techniques on the derivative of the image, or the norm of the gradient of the image, that is to segment the regions of steep brightness changes, converting the method into a border detection technique.

Watersheds are the most widely-used methods for cell segmentation in calcium imaging, because, by definition, they are designed to locate local maxima/minima intensity regions

which are the standard images (fried egg model mentioned earlier). Also, the watershed method is immune to shape and size variations [148] and it can easily handle images of high density cultures with clustered cells [149, 150]. The main drawback of the watershed method is that it is sensitive to noise and will tend to over-segmentate if no constraints are established for the regions [151]. This sets the wide variety of methods found in the literature and also allows some authors to claim it gives bad results by applying no or poor constraints to the technique.

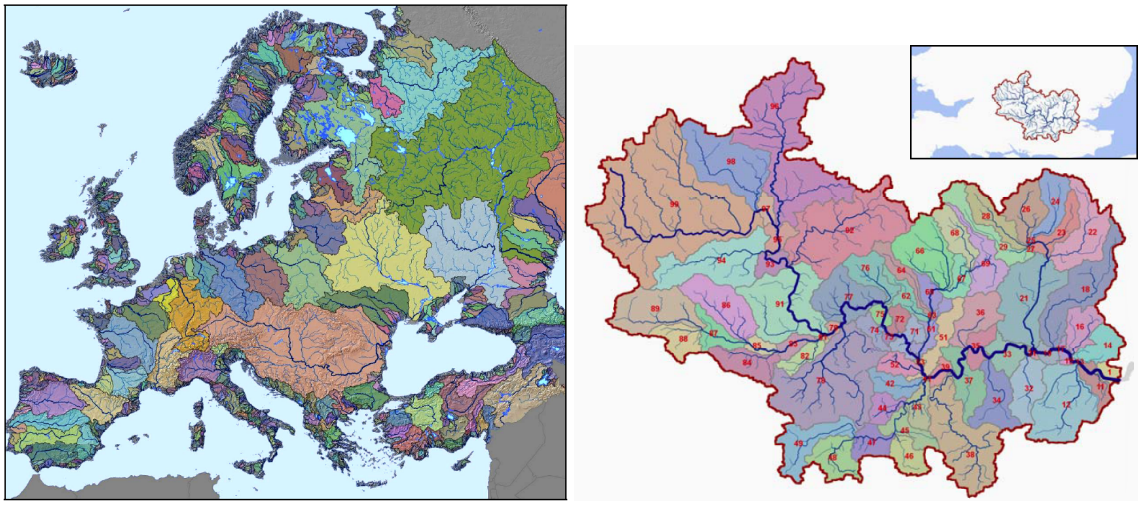

(a) *Catchment basins of the major rivers in Europe.* **(b)** *Sub-basins of the River Thames (UK).*

Fig. 26 *Watershed segmentation requires a stopping rule of some kind to avoid over-segmentation. Both figures from Vogt et al [152]. These are further subdivisions of the watersheds shown in figure 21.*

There are many examples of watershed based approaches to cell segmentation, the following are worth mentioning because of their similarity to the method applied here: Lindblad et al. [153] overcome the over-segmentation by previously detecting cell nuclei (using another fluorescent marker) and seeding the watersheds from the nuclei, thus allowing a single cell (watershed) from each nuclei. A similar approach is used by Pinidiyaarachchi et al. where they use a sequence of images to track cells by using the segmentation in one frame as the seed for the next [143]. Arteta and others [47, 154, 155] have applied a well known method in which the watersheds are selected depending on the speed at which they grow while the surface is filled, i.e. sudden changes in area can indicate two regions have been merged thus solving the over-segmentation problem. Wahlby et al. [156] propose a method in which they seed from local maxima (inverted image model) and allow for regions to grow from this maxima on a step-by-step basis in a range of intensities. The range is gradually lowered thus finding cell borders when two regions meet.

Finally, for a some articles where several methods are reviewed and comparison measures are proposed are [43, 44, 120, 157] and a nice recap on cell segmentation in the last 50 years is the work by Meijering [100].

3

Molecular Scale

In this chapter we cover a collection of methods for detection and measurement of structures at the intracellular scale. These are generally static cell organelles and the methods described can be applied to the analysis of any sub-diffraction limit object. They are molecular size structures in the size order of magnitude of the nanometer, tagged by some fluorophore and under the microscope they will be imaged as shapes close to the system PSF (see section 1.4.5). The chapter focuses mainly on ryanodine receptor 2 (RyR2) detection because it is currently a research focus in calcium imaging, but it could simply be called *blob detection*.

3.1 Ryanodine Receptor 2

As mentioned in section 1.3 the RyR2 is a transmembrane protein acting as a calcium channel in the sarcoplasmatic reticulum's (SR) membrane and its role in the calcium induced calcium release process is to mediate the massive release of calcium ions into the cytosol. Its name is due to ryanodine, an alkaloid found in a Caribbean plant *ryania speciosa*, which has nothing to do with cardiac function other than its toxicity to mammals for it blocks the RyR, and thus the historical name for the protein [158, 159].

RyR2 is the isoform of the protein found in cardiac cells (RyR1 are found in skeletal muscle cells and RyR3 in neurones) [160]. They are grouped in clusters, each containing a number ranging from a few dozen [161] to a few hundred [162] and the clusters are distributed all over the SR which at the same time is evenly distributed throughout the myocyte.

It is well known that calcium release through the RyR2 plays a crucial role in the regulation of intracellular calcium and cardiac contraction [164, 165] and that some heart diseases are linked to mutations in the RyR2 [166]. However little is known about the 3D distribution throughout the cell and if they are the only agents responsible for calcium release from the

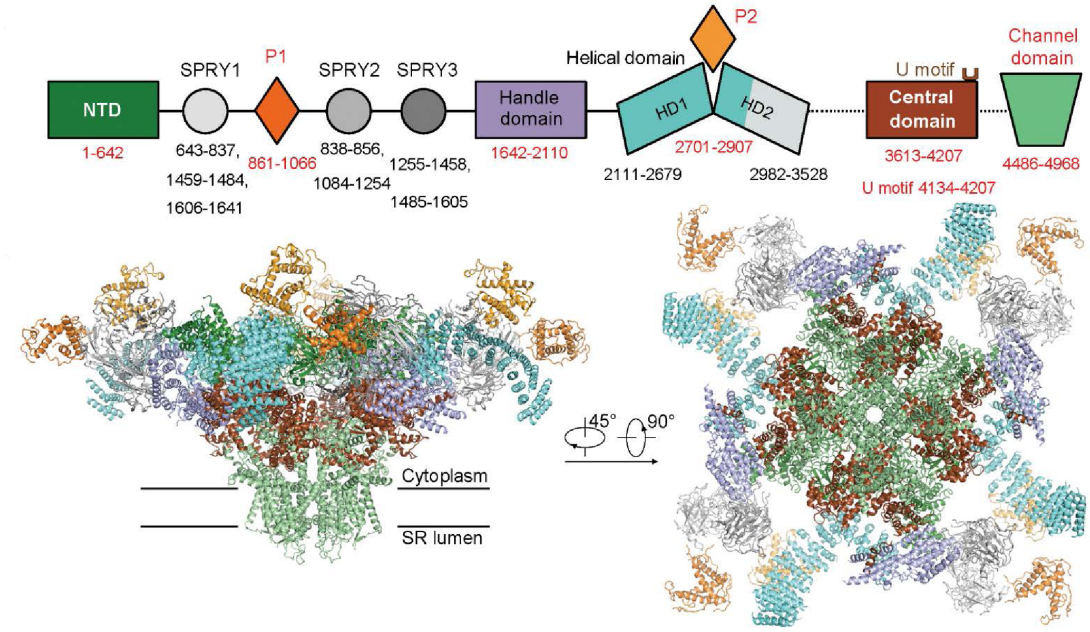

Fig. 27 *Ryanodine receptor structure in closed state obtained from cryogenic electron microscopy by Peng et al. [163]. Top panel shows the different domains identified in the receptor and in the bottom panel we see a side view (left) and top view (right) of the receptor using the same colour code for each domain. The sarcoplasmatic reticulum membrane position is indicated in the left view with the two black lines, on top would be the cytoplasm and below the sarcoplasmatic reticulum interior.*

sarcoplasmatic reticulum. Therefore this molecule is currently a common focus of study in many research groups around the world and the development of computational techniques for the detection of RyR2 clusters from fluorescence imaging is currently highly relevant.

3.2 RyR2 localisation

RyR2 can be tagged with fluorophores in order to be viewed through the microscope. There are several different techniques for doing this but in any case, as they are objects some ten times smaller ($\sim 30nm$ [167]) than the diffraction limit of optical microscopes, they appear as blobs a little larger than the size of the system's PSF (see figure 29).

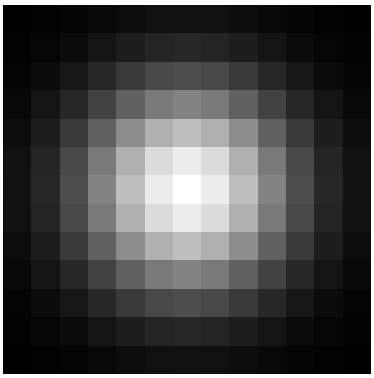

Fig. 28 *Example template used for filtering RyR2 images. It is a 2D Gaussian with its mean subtracted, so that the outer values are negative, the central values are positive and the sum of all values is zero.*

A simple and effective way to localise these structures in an image is to perform a search for local brightness intensity maxima. The problem is that very often noisy images will have local maxima that are purely noise peaks, so we need to free the image of these noise peaks and so a filter has to be applied. A good approach is to use a Gaussian filter of the same size as the PSF of the system (figure 28). The reason is two-fold; the 2D Gaussian is a good approximation to the PSF [77, 168] so it will nicely adjust to the image wherever clusters are located and, in this manner, RyR clusters will be enhanced while at the same time the whole image will be smoothed out reducing the number of false detections when localising the clusters.

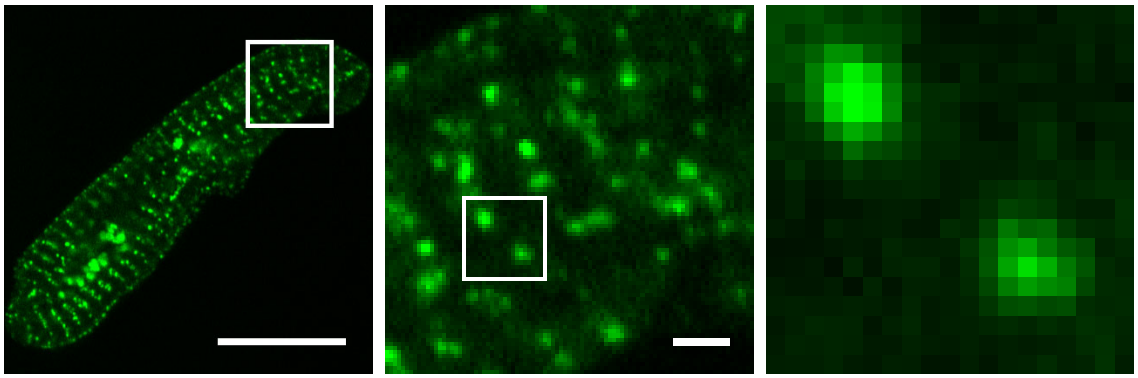

(a) *Scale length 20μm.* **(b)** *Scale length 2μm.* **(c)** *Pixel size 0.135μm.*

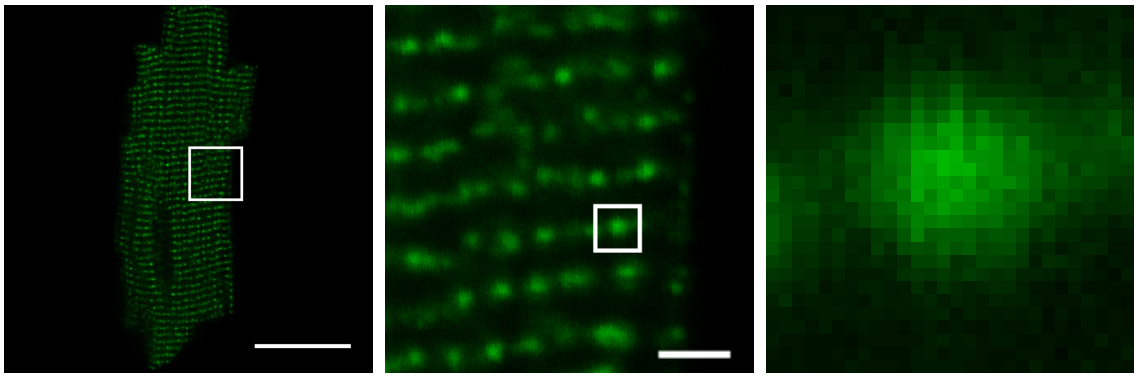

(d) *Scale length 20μm.* **(e)** *Scale length 2μm.* **(f)** *Pixel size 0.04μm.*

Fig. 29 *Two example images of RyR2 clusters in different conditions. Top row: Human atrial myocyte labelled by means of immunofluorescence. Bottom row: Mouse ventricular myocyte from a line of mice that have been genetically modified to express GFP at RyR2s. In both cases, brighter objects correspond to clusters that fall on the focal plane and fainter objects correspond to out-of-focus clusters. In the first image, the bigger bright regions correspond to non-specific staining, possibly vacuoles or other cellular organelles that are not the real target of the study. In the second case the pixel size is unnecessarily small: the cluster is a sub diffraction-limit object an so the blob we see is the microscope PSF and it will not get any better by improving the resolution of the detector.*

Cluster detection then becomes easy by finding those pixels with zero derivative as explained in section 1.5.6 although it is very likely that some kind of post filtering will be needed to discard local maxima with too low intensity values (possibly out of focus clusters and/or surviving noise peaks) or objects that do not have the expected shape of the PSF.

Another approach is to use a watershed based method and impose some sort of limitation to the detection by applying known properties of the objects to be detected such as shape or size, but to do this we first need to cover how to measure these properties.

3.3 Shape and size measurements

Due to the fact that these sub-resolution objects do not have a clear border, as they rather fade into the background, a definition of how size is going to be measured is required. This definition can either be relative or absolute in terms of image intensity: We can say an object border corresponds to those pixels with intensity below a certain threshold, this threshold being a fraction of the brightest pixel of the object itself, or instead a fraction of the brightest pixel in the whole image.

Using one definition or another will depend on the image content. If for example we know all objects are more or less the same size but the fact that some are farther away than others makes them appear smaller in the image, the first option will be more appropriate. On the other hand, if we know all objects are more or less on the same plane, and that those we see smaller are truly smaller, then the second option is a better approach.

Other definitions can be made using the derivative of the object's cross section. For example a common criteria is to search for the inflection point in the decay (i.e. where the slope is maximum). Once a size definition is set, it is best practice to fit the object to a known function and measure the size using this fit instead of measuring from the raw pixel values in order to avoid errors due to image noise. For the particular case of RyR2 cluster images, a Gaussian fit is a straightforward approach.

As RyR2 cluster images have radial symmetry, in principle their radius can be measured in any direction and should produce the same result. For this reason, a sensible approach is to do so in a series of cross sections and give the radius result as a mean of the values measured in every direction. One last consideration though is to take each cross-section and crop from the first minima on either side of the local maxima so that any possible neighbouring clusters do not affect the size measurement (see figure 30 magenta cross-section). Note also that in each cross-section, pixel size depends on the obliquity of the cross-section.

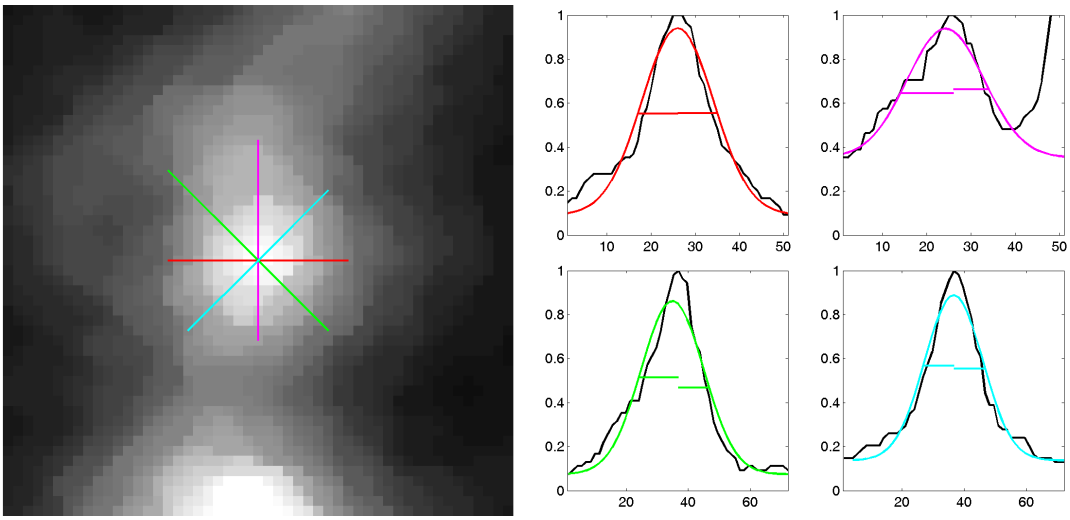

Fig. 30 *Cluster size measurement using four cross-sections. In each of them a Gaussian function is fitted to the segment of the cross-section in between the two minima before and after the maximum. As an example, the top right panel shows a case where the final uprise was excluded in the fit. Note that the radii are measured from the true centre of the object to the intersection with the fit, so in this case we are actually measuring eight radii. Horizontal axis in the plots are measured in pixels.*

3.4 Z-line distance estimation

Z-lines are the name given to the observed dark lines between sarcomeres when viewing a myocyte under the microscope. They correspond to the place where actin molecules are bound together in the contractile machinery of a myocyte, but they are merely a name given to a visual perception, they do not define a structural component of a myocyte. Nevertheless the distance between Z-lines in a myocyte does give information about the grade of contraction of a myocyte and so the measurement of this distance is relevant when studying myocytes.

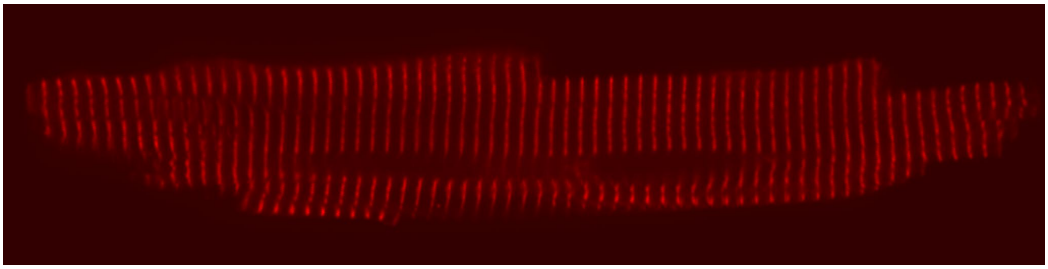

Fig. 31 *Mouse ventricular myocyte with actin labelling. The centre of each band corresponds to the position of the Z-line. RyR2 clusters and actin filaments are arranged with the same lattice constant, known as Z-line distance.*

The reticular distribution of RyR2 clusters corresponds to the periodic distribution of actin myosin filaments and for this reason the distance between bands of RyR2 clusters corresponds to the Z-line distance.

One would think that the best way to tackle this problem is to connect each detected cluster with its nearest neighbour in order to to segment the lines clusters are distributed along, fit a

straight line and then measure the distance between these lines in the perpendicular direction. This method is tedious and is easily dependent on errors due to regions where the confocal plane changes from one plane of clusters to the next or regions close to the cell membrane where the mesh seems to disappear. A statistical approach to this measurement can be used thanks to the fast algorithms available for computing all pairwise distances bewteen a set of objects [169], in this case between RyR2 clusters. If an image has N detected clusters, the number of distances to measure is $(N^2 - N)/2$. If one then takes the histogram of the values measured from the periodic distribution of clusters, this histogram will produce a series of maxima matching the multiples of the Z-line distance. By autocorrelating this histogram one can measure the mean of the most present distance in the image and therefore a mean Z-line distance is obtained by taking into account all clusters (see figure 32).

This procedure can equally be applied to an image with labelling on any of the proteins that form the contractile machinery of a myocyte. As an example figure 31 shows actin labelling corresponding to the Z-lines.

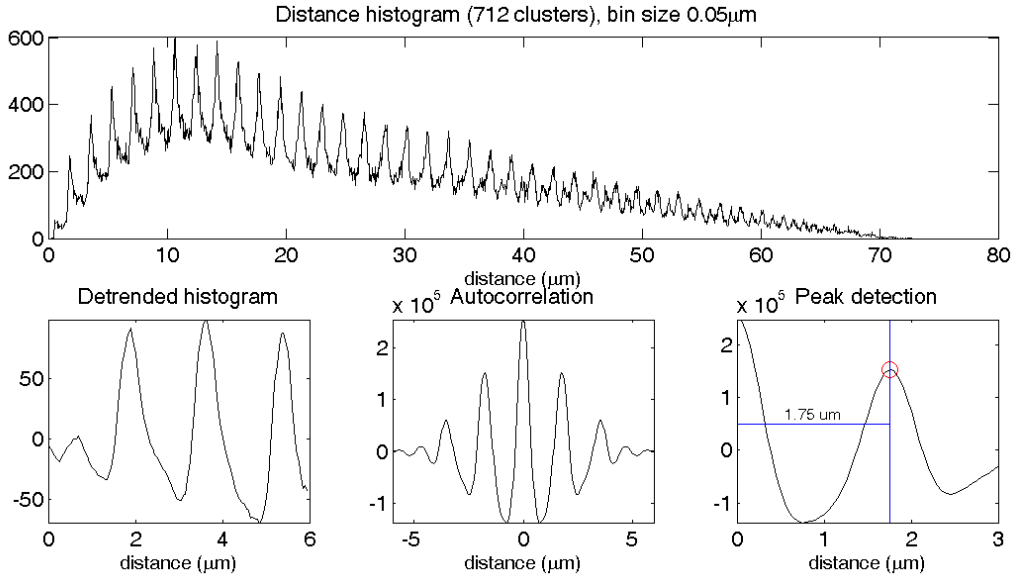

Fig. 32 *Z-line estimation after RyR cluster detection. Top panel shows the histogram of all distances present between clusters, in this case a total of 712 clusters yield to ∼250k distances. As can be seen there is a periodic pattern corresponding to multiples of the separation between cluster rows. To measure the distance corresponding to this periodicity, we first cut and de-trend the initial region of the histogram (bottom row, left-most panel), we then autocorrelate this signal (bottom row, middle panel), and we finally measure the distance to the first peak (bottom row, rightmost panel).*

3.5 Example applications

The set of sub-resolution object detection techniques described above have successfully been applied in a series of experimental conditions and the results have been manually

verified in a subset of cases. As one may imagine, the best results are obtained with the higher signal to noise ratio, higher resolution and also when no non-specific fluorescent signal is present. The following subsections are some examples of applications.

3.5.1 Straightforward detection

The first example is from a set of experiments of transgenic mice expressing GFP at RyR2 clusters where detection was used for exploration of spatial distribution, size and distance measurement [170, 171]. These images were processed using some of the techniques described in the previous sections; image pre-filtering for noise removal, convolution by a Gaussian filter (0.5 divided by image resolution of 0.14μm yielding around 35 pixel filter similar to figure 28), image binarisation with a threshold depending on image intensity range, connected component filter, and final removal of objects out of expected size interval. Post-processing included measuring nearest neighbour distances, size measuring (using eight cross sections as described in section 3.3) , and Z-line distance estimation (using method described in 31). See figure 33 for examples of these experiments.

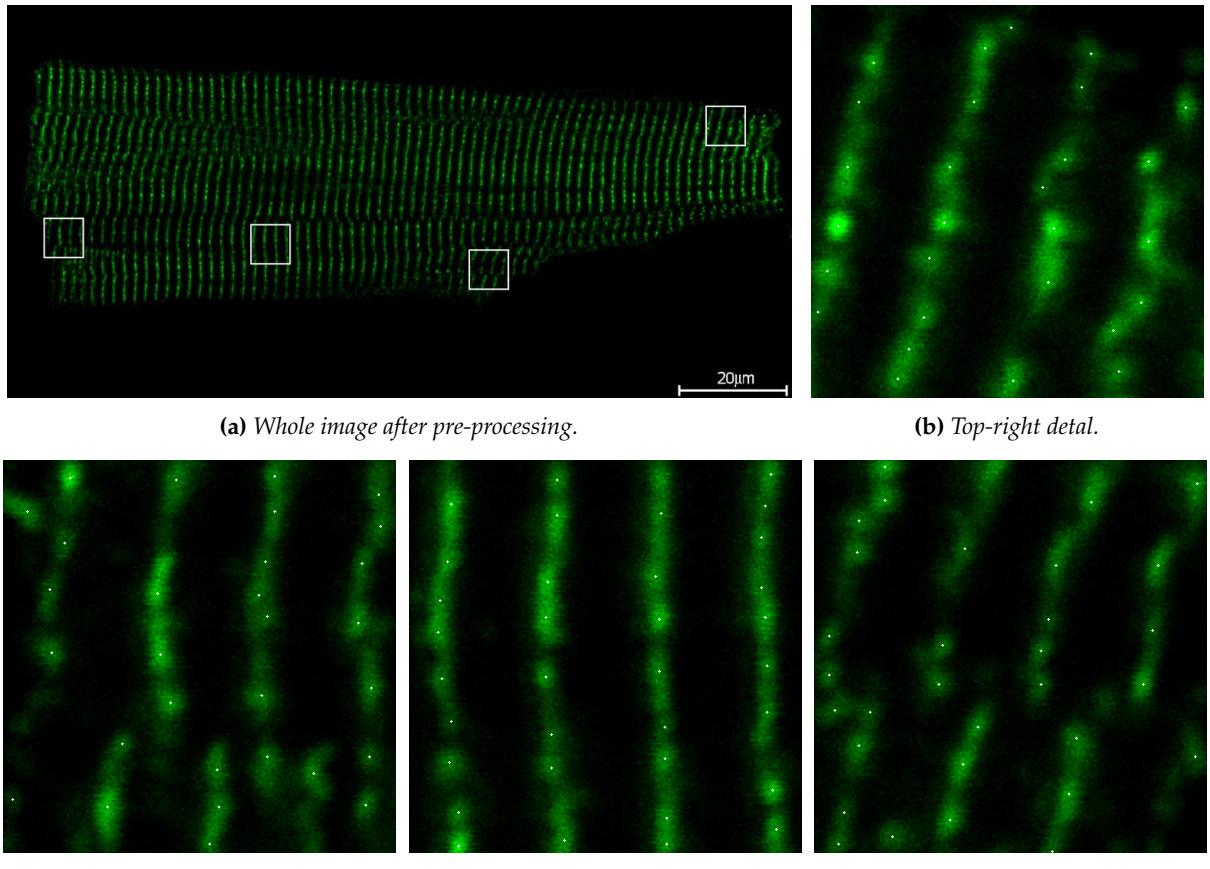

(a) *Whole image after pre-processing.* **(b)** *Top-right detal.*

(c) *Bottom-left detail.* **(d)** *Bottom-center detail.* **(e)** *Bottom-right detail*

Fig. 33 *Mouse ventricular myocyte expressing GFP at RyR2 clusters. After the image has been processed for detection of clusters, a white dot has been positioned on detected sites. Pixel size is of 0.14μm.*

3.5.2 Phosphorylated versus non-phosphorylated clusters

Another application of the methods described was applied to a set of experiments in which the labelling technique used was immunofluorescence. It has been proven that phosphorilation of RyR2 will produce a conformational change and promote its opening [172] and for this reason in this study a fluorophore was used to tag RyR2 clusters, while another fluorophore was used to tag those clusters that were phosphorylated (see figure 34). This work was presented in [173] and also used in [174]. Detecting the clusters in each of the images allowed computing of ratios of phosphorylated versus non-phosphorylated clusters under different pharmacological conditions. Also these phosphorylation ratios were measured in separate concentric regions in the cell in order to measure how the distance to membrane would affect phosphorylation under the different pharmacological conditions (see figure 35).

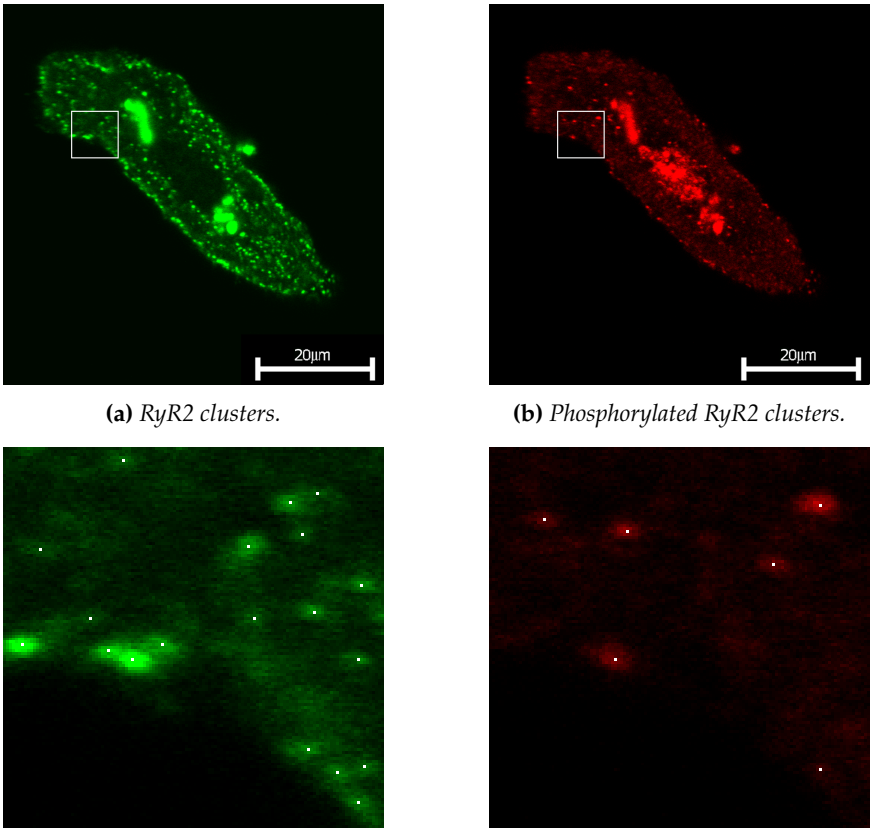

(a) *RyR2 clusters.* **(b)** *Phosphorylated RyR2 clusters.*

(c) *RyR2 clusters detail.* **(d)** *Phosphorylated RyR2 detail.*

Fig. 34 *Sub-diffraction limit blob detection applied to Phosphorylated RyRs versus non-phosphorylated. Upper row shows original images (after preprocessing), lower row shows the same detail with white pixels on detected blobs.*

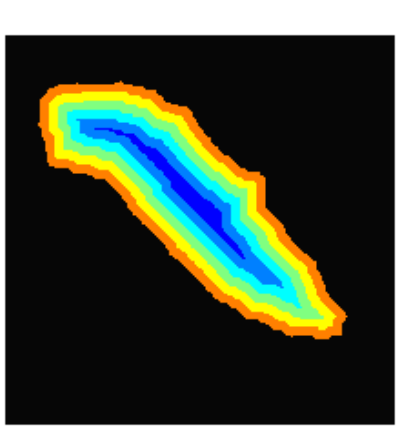

(a) *Cell regions.*

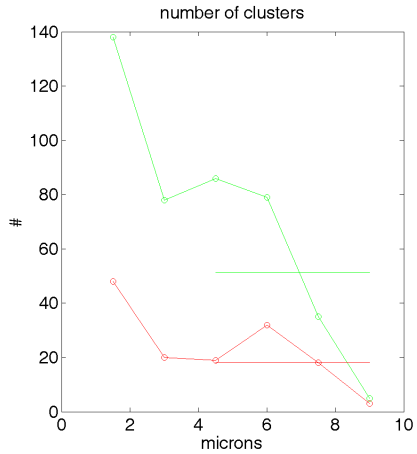

(b) *Number of clusters.*

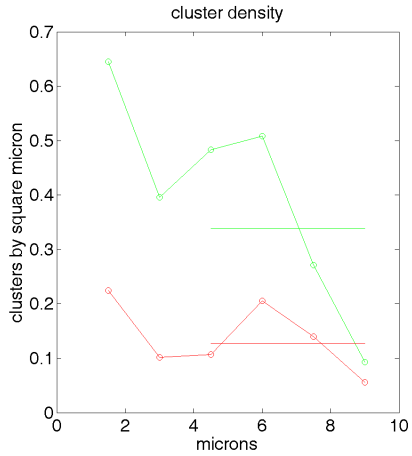

(c) *Cluster density.*

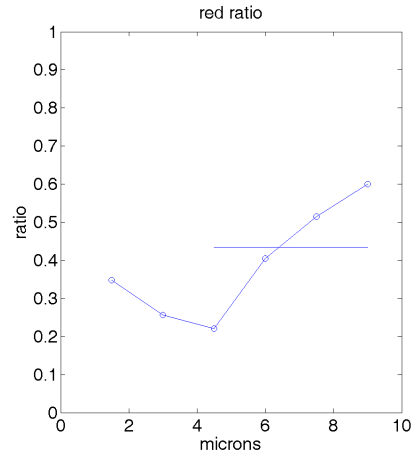

(d) *Phosphorlated ratio.*

Fig. 35 *Peeling example results for the same cell as in figure 34. First panel (a) shows concentric regions of width 1.5μm, the following panels show number of clusters in each region (b), cluster density in each region (c) and ratio of phosphorylated RyR2 over the total number of RyR2 clusters (d). In these graphs the horizontal line corresponds to the mean of the inner regions (third and higher).*

3.5.3 Co-localisation of RyR2 and membrane calcium channels

This set of experiments is very similar to the previous example regarding image conditions and image processing but in this case one channel contained the RyR2 cluster fluorescence whilst the other was of L-type calcium channels. See figure 36.

3.5.4 Line scan example

This example corresponds to a set of experiments designed to co-localise calcium sparks (section 4.1) and RyR2 clusters in line scans (section 1.4.3) used in [170, 175].

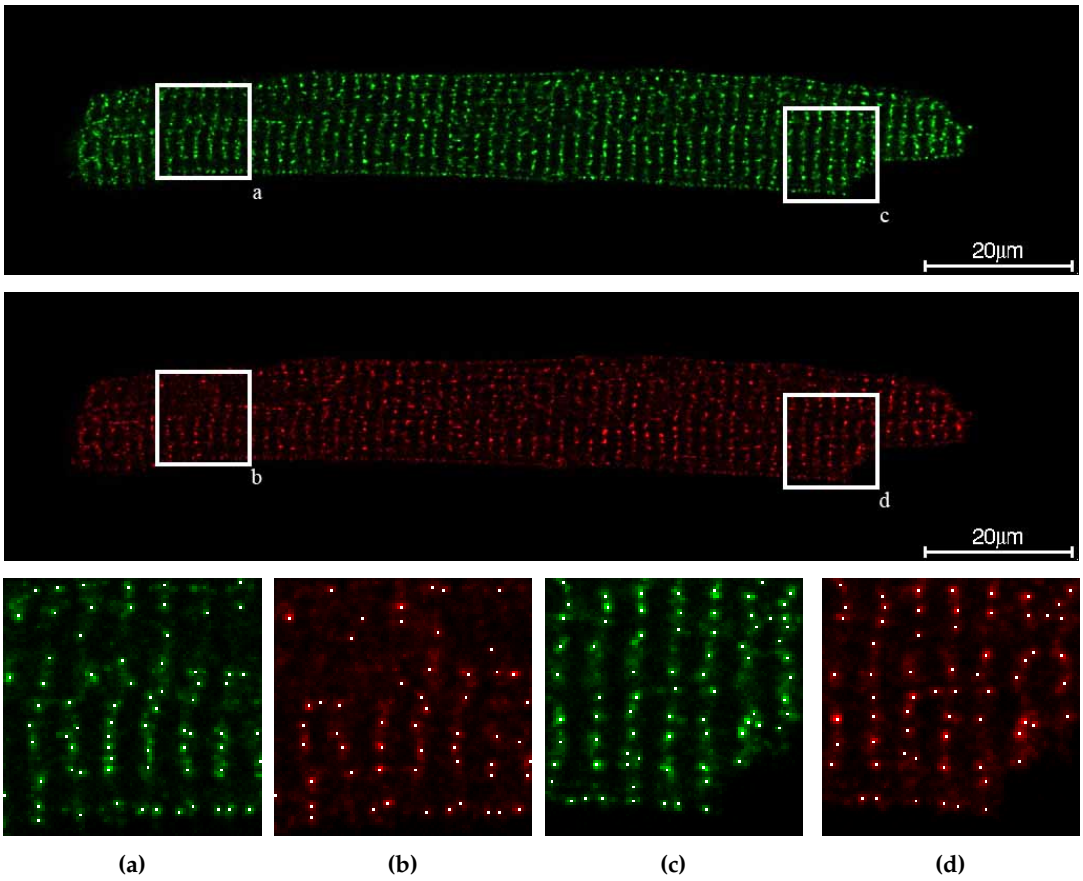

Fig. 36 *Mouse ventricular myocytes. Green fluorescence reveals RyR2 cluster locations and red fluorescence reveals calcium channels ($Ca_v1.2$). Detailed crops have been marked with a white pixel in detection sites.*

In this set of experiments the image temporal pixel size was of 1.8ms and the spatial pixel size ranged from 0.04μm to 0.06μm. These values are relevant for the filtering parameters mentioned further on. As RyR2 clusters are static objects, they appear as bands following the time direction, with only a slight shortening of the distance between them whenever the cell is contracting. For this reason the signal was further enhanced by summing the spatial signal in the temporal direction in windows of 100ms (a time long enough so as to ensure noise reduction and short enough so as not to include contraction artifacts). The resulting spatial fluorescence signal from each time sample was then run through a continuous wavelet transform with the scales of the estimated sizes of the clusters (sum of scales 4 to 10) and then thresholded to values of the transform above zero. The resulting binarised image was run through a closing filter to fill any gaps in the logical true regions and finally a *skel* filter was applied to obtain the central line for each of the bands. This process can be seen in figure 37.

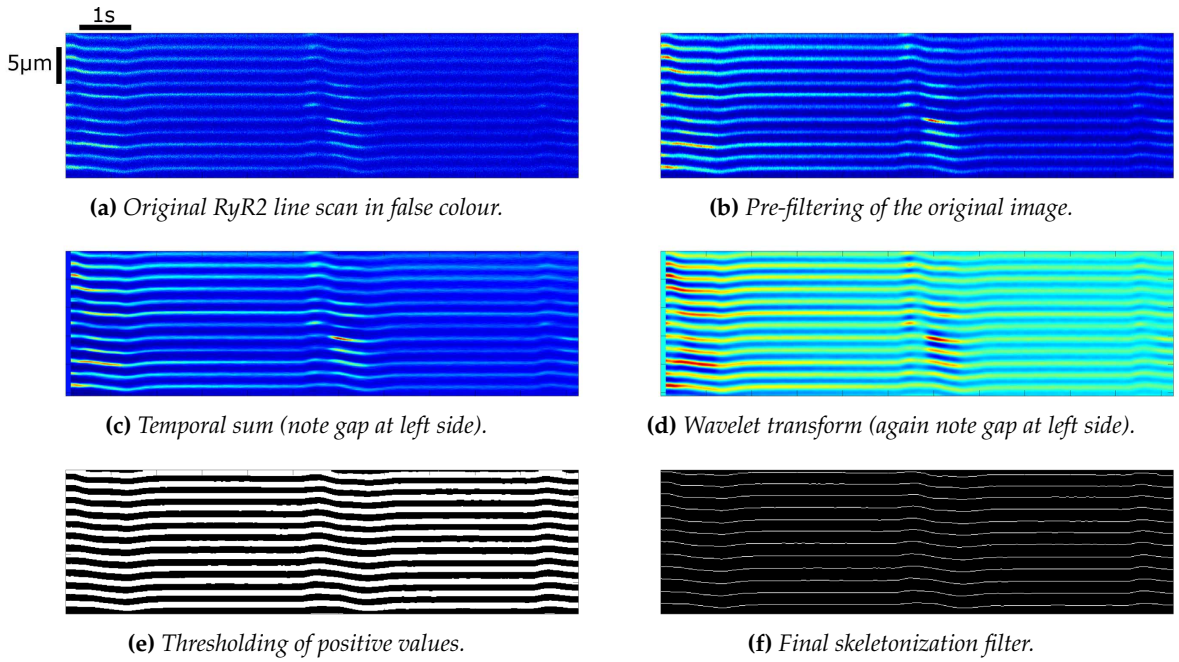

(a) *Original RyR2 line scan in false colour.*

(b) *Pre-filtering of the original image.*

(c) *Temporal sum (note gap at left side).*

(d) *Wavelet transform (again note gap at left side).*

(e) *Thresholding of positive values.*

(f) *Final skeletonization filter.*

Fig. 37 *RyR2 clusters detected in line scans. Note that the horizontal dimension is time, and that it has been significantly compressed to show the two contractions present during the recording. We are seeing 20µm in the vertical direction by 10s in the horizontal direction. The crucial step in the process corresponds to panel (d) where the image is composed of concatenated vertical lines resulting of wavelet filtering using gauss2 family for a range of appropriate scales. Note that the gap to the left of the image corresponding to the size of the temporal window used for the sum in step (c), that is then filled up at step (e) by extending the values of the first line different to zero.*

3.6 Results

The following sections present some of the results directly obtained from the methods in this chapter. Whenever a value is reported accompanied by a plus/minus sign (±) and second value behind it, this second value always corresponds to the standard deviation (not standard error as commonly used in biology).

3.6.1 RyR2 detection in line scan

The RyR2 cluster detection method for line scans described in section 3.5.4 was used on several image sets in order to report cluster distribution and its relation to spark occurrence sites [170]. For the study, a total of 68 line scans were used from which 1391 RyR2 clusters were detected.

The method proved to be near perfect, successfully detecting clusters in challenging conditions such as while the cell is contracting or in situations where the clusters are very faint. Figure 38 shows examples of dim clusters due to image conditions, out-of-focus clusters

Image Set	# of cells	scanned distance (μm)	cluster density (#/μm)	nearest neighbour(μm)	# of clusters
Jun2015	11	384	0.58 ± 0.04	1.77 ± 0.13	223
Aug2015	23	589	0.59 ± 0.05	1.76 ± 0.12	342
Sep2015	34	1408	0.59 ± 0.03	1.72 ± 0.09	826

Table 1 *Summary of the cluster detection in line scan datasets. Note that the nearest neighbour distance is in fact closer to the z-line distance because the cells were scanned in the contraction direction.*

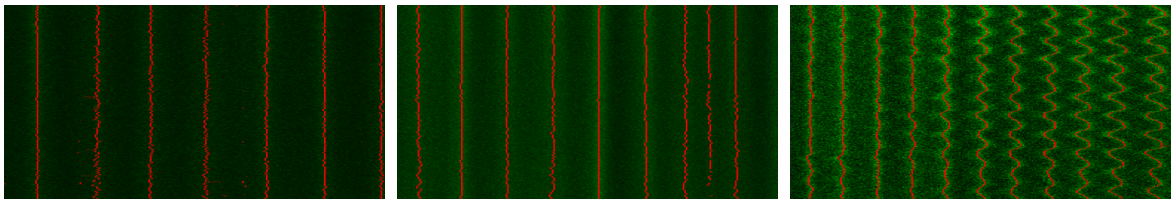

(a) *Faint clusters.* **(b)** *Out-of-focus clusters.* **(c)** *Clusters under contraction.*

Fig. 38 *RyR2 detection in line scans showing examples of challenging situations. The original image is the green channel line scan and a red line has been pained on the detected clusters. In all cases spatial direction is horizontal and temporal direction is vertical.*

due to structural reasons like the sarcoplasmatic reticulum folds or transition areas from one level of clusters to the next and a contraction situation.

The scans were taken in the contraction direction so the nearest neighbour distance in fact corresponds to the z-line distance. For each cluster detected, the distance to the next was measured, so that each scan i with n_i clusters contributed with $n_i - 1$ distances. The final z-line distance obtained was 1.73 ± 0.11 (μm), N =1323.

3.6.2 RyR2 detection in frame scan

The frame scan RyR2 cluster detection methods described in this thesis have been published as independent methods and also have allowed publications towards further understanding of both structure and function of cardiac cells from a physiological point of view.

In order to develop the RyR2 cluster detection, two sets of images were used; one of inmuno-fluorescence human atrial myocytes taken from patients under heart surgery, and another of mice ventricuar myocytes genetically engineered to express green fluorescent protein at RyR2 clusters.

The first is summarised in table 2, and is the set that was initially used for developing the cluster detector. The images in immunofluorescence are by far noisier and therefore are a harder set. The set contains images taken from 431 different cells under different drug conditions, some of them imaged in a single frame and others taken as a series of stacks along the optical axis to obtain a 3D structure approximation. The experiments were part of a study on the effect of phosphorilation in the opening of RyR2s and for this reason all images are

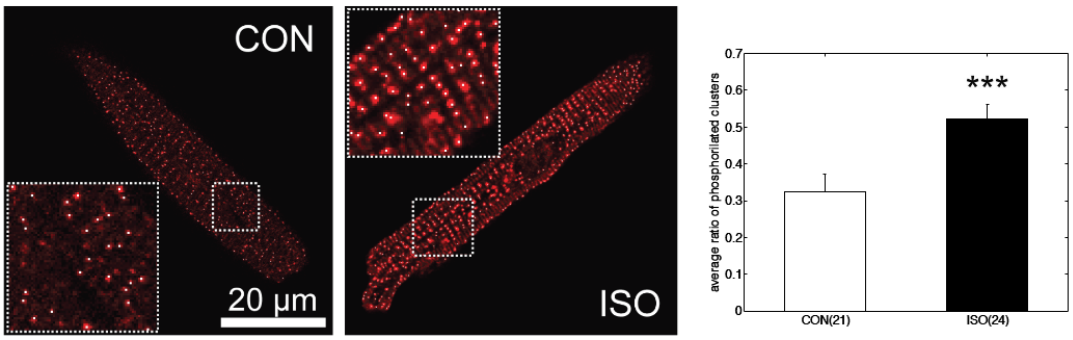

Fig. 39 *Phosphorilation dataset results. The two images shown correspond to the phosphorilated cluster tag images, both of which are compared to a RyR2 cluster image. The reported ratio is the number of clusters that appear in the phosphorilated image over the number of clusters that appear in the RyR2 cluster image.*

duplicated, one corresponds to the raw RyR2 tagging and the other to the phosphorilated RyR2 clusters (ratiometric approach described in section 3.5.2). This leads to a huge dataset with over 10k images containing near 2M clusters.

Dataset characteristics	
Number of patients	44
Number of cells	431
Number of treatments	6
Number of image stacks	5064
Number of images	10854
Number of detected clusters	1972366

Table 2 *The human dataset used for developing the RyR2 cluster detector.*

The dataset included cells under six different treatments, one being the control and the other five being drug conditions. Of the five drug conditions only one was used for the study results, isoprotheranol (ISO), to compare the ratio of phophorilated clusters under this drug treatment versus the control set. Of the total 45 patients, 9 where used in the study, from which 21 untreated control cells where used and 24 cells under ISO treatment. These two subsets produced mean phosphorilated ratios of 0.32 ± 0.03 (CON) versus 0.52 ± 0.06 (ISO), with a t-test p-value of 0.002 (figure 39)[173], the p-value in this case being the probability of the two distributions having equal means and producing these results. This result was used as a positive control towards proving the detection method since it was a known fact that ISO would promote cluster phosphorilation.

The second dataset was the one used in the co-localisation study of sparks and RyR2 clusters (table 3). In this case a total of 173 confocal images of mice ventricular myocytes where used, in which over quarter of a million RyR2 clusters where detected (263402). This study presented cluster size measurements, cluster density and nearest neighbour distance, together with z-line distance shown in the following section.

The conditions of this set where slightly different to the previous one, mainly because

there was no non-specific staining, which made the detection simpler and therefore more robust. Of the whole dataset used for developing the detector, in the final publication only a small fraction was used because many images where discarded for several reasons. Some of the reasons are image saturation (image dynamical range is not optimal), out-of-focus images (object sizes cannot be properly estimated), cropped cells or more than one cell in a single image (producing false cell area estimation and therefore offsetting the cluster density). Nevertheless the final numbers presented [170] coincide with the ones that appear here in table 3. Figure 40 shows the distributions for nearest neighbour, cluster size and cluster intensity for one of the image subsets.

Image set	# of cells	cluster radius (nm)	cluster density (#/μm^2)	nearest neighbour (nm)	# of clusters
Live cells	59	49.5 ± 5.3	0.719 ± 0.093	763 ± 37	94029
Fixed cells	58	53.5 ± 6.2	0.699 ± 0.105	785 ± 34	85201
March set	28	55.7 ± 5.7	0.725 ± 0.062	787 ± 32	38715
July set	28	62.5 ± 5.2	0.316 ± 0.064	762 ± 29	45457

Table 3 *Summary of the cluster detection datasets with the measured parameters. Note that reported cluster radius are sub diffraction limit values, they have been estimated using the measured PSF of the optical system (see section 3.6.4).*

3.6.3 Z-line distance estimation

Z-line distance estimation by means of the procedure described in section 3.4 was applied to the same datasets as in the previous section, but in this case adding an extra subset of

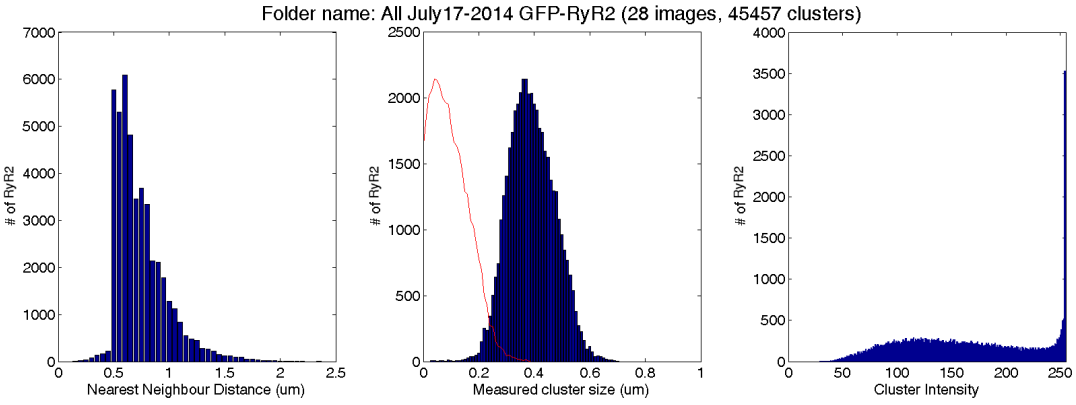

Fig. 40 *RyR cluster measurements for the July image set. Left panel shows nearest neighbour distribution, the shape of which is highly conditioned by the detection method (section 3.2). Central panel shows measured size together with the estimated size by shifting the distribution using the known optical system PSF (section 3.6.4). The third panel shows cluster intensity in order to demonstrate that some of the images have clusters that are saturated, i.e. they are cut flat at the top and therefore are really brighter than they appear. This is usually human error and due to over exposition or too much gain in the detector amplifier; we erroneously tend to want to see the dimmer clusters at the cost of capping the brighter ones when the information is already in the image, just that the naked eye cannot see it.*

experiments. In this case a total of 218 cells where analysed and the mean z-line distance obtained was of 1.74 $\pm$ 0.20 μm.

The extra subset corresponds to a set of images that had had a manual measurement of the z-line distance and therefore was included as a validation set. The manual measurements had reported a z-line distance of 1.78 $\pm$ 0.16 μm (N=28), whereas the automatic determination for this subset produced a z-line distance of 1.70 $\pm$ 0.21 μm. Both the automatic and manual values overlap each-other in terms of the standard deviation which is a measure of the goodness of the method. As always though, manual validation is in general dangerous because humans will tend to be prone to bias and fatigue when manually segmenting or measuring. In this case for example, the 28 images came from a larger set of 31 but three of them had been already excluded because they were declared as presenting *contracted regions*. Of the final 28, another 13 were included although they were tagged as *slightly contracted* and it was thought that they would contribute to underestimate the z-line distance.

Image set	# of cells	Estimated z-line distance (μm)
Live cells	59	1.75 $\pm$ 0.19
Fixed cells	58	1.73 $\pm$ 0.21
March set	28	1.66$\pm$ 0.14
May set	45	1.70 $\pm$ 0.21
July set	28	1.84 $\pm$ 0.17
Totals	218	1.74 $\pm$ 0.20

Table 4 *Summary of the z-line estimation datasets with the automatic measurements. The May dataset is the one that was used for manual measurement.*

3.6.4 PSF estimation

Measuring of the point spread function (PSF) of an optical system is useful if one wants to estimate the size of objects smaller than the actual PSF. In this case we were attempting to perform a rough estimation of the RyR2 cluster sizes, aiming at later developing an image deconvolution method to improve resolution of the system. As explained in section 1.4.5, an image of a sample is actually the convolution of every light source of the sample by the system's PSF. Big objects (compared to the PSF) hardly appear modified but small objects will be viewed with a slight size increase due to the PSF and a single point source is viewed as the PSF itself.

In order to measure the microscope PSF, a light emitting source of a known size that is as small as possible needs to be imaged. For that, artificial fluorescent beads with a diameter of 0.1μm were bought and imaged under the microscope with the same setup that had been used for viewing the cardiac myocytes (the beads commercial name is TetraSpec™microspheres, catalog number T7279 at TermoFisher) .

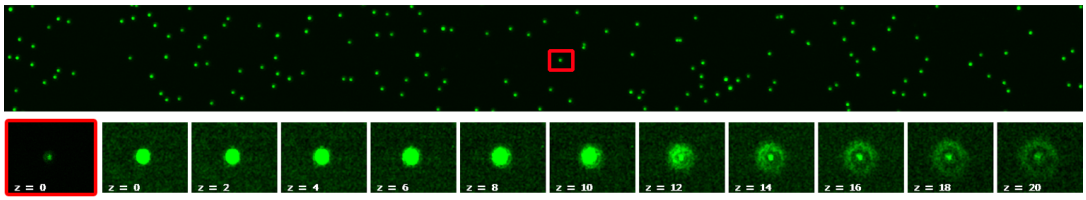

Fig. 41 *PSF estimation. Top panel shows a crop of a high resolution image of fluorescent beads of 0.1μm diameter. The image has been saturated in order to be able to actually see the beads. There is a region of interest marked of an example bead which appears magnified at the bottom left (unsaturated). Next to it a sequence of stacks (again saturated in order to perceive the diffraction rings), moving away from the focal plane where the original image was taken. As can be seen the PSF grows as we move away from the focal plane forming conical shape, sort of a sand glass. Step size between stacks is* $0{,}1\mu m$*, as the figure shows only even numbered stacks, the distance from one to the next is* $0{,}2\mu m$*.*

If one measures a bead's diameter in an image, it appears bigger than the known size due to the optical system's PSF. The excess radius with respect to the known $0.05\mu m$ is due to the PSF's radius. This is a very simple first approach but it allows to estimate the measured size of any other object by subtracting the PSF's estimated radius.

The approach is applicable only for objects that are lying on the same plane as the plane where the PSF estimation was made, so if we measure the bead's radius on the focal plane, we can only ensure an estimation for objects on the focal plane. Moreover as it is a very rough method, it is really only valid when coping with great numbers; subtracting the PSF radius to the mean radius of the objects makes sense, but subtracting to individual values will easily produce impossible results like negative sizes (see figure 40, central panel).

For the optical setup used for viewing the RyR2 clusters, a total of 628 beads where measured using the method described in section 3.3 to produce a bead radius of 0.314 ± 0.078 μm and thus a PSF estimation radius of $268nm$. This allowed estimating the real size of the clusters in the image sets that had been used for characterising RyR2 distribution, producing mean cluster radius of 55.7 ± 5.6 nm.

4

Sub-cellular scale

4.1 Types of event

During cardiac contraction each myocyte undergoes the process known as *calcium-induced calcium release* described in section 1.3. It is a complex process in which many factors intervene and with a variety of possible outcomes that can occur depending on specific conditions. A heuristic classification of calcium release events has been established based on visual inspection of cardiac myocytes in fluorescent microscopy of the spatio-temporal scale and the ability of the event to depolarise neighbouring cells (see table 5).

Event Name	Spatial Scale	Temporal Scale	Description
Spark	$\sim 1\mu$m	<100ms	Localised and generally isolated calcium release event from the sarcoplasmatic reticulum.
Mini-wave	$\sim 10\mu$m	>100ms	Propagating calcium release event that covers a small fraction of the cell size and is not capable of depolarising neighbouring cells.
Wave	$\sim 100\mu$m	~1s	Major slow-propagating calcium release event covering a large fraction of the cell volume and is capable of depolarising neighbouring cells.
Transient	*whole cell*	~1s	Synchronised calcium concentration increase during cell contraction that depolarises neighbouring cells.

Table 5 *Summary of intracellular calcium events [176].*

Cell transients are the normal cell function in which the intracellular calcium rises and triggers cell contraction. The release from the sarcoplasmatic reticulum is homogeneous and synchronous so in a fraction of a second the intracellular calcium of the whole cell rises to maximum value and then decays to prepare for the next beat. Waves are said to be slow-propagating events meaning that they are slower than a regular transient that is practically instant. This slow-travelling speed allows for a neighbouring cell to complete its refractory

period meaning that it can be stimulated again by the wave which can produce a variety of undesired propagation faults that on the large scale show up as cardiac malfunction.

Calcium sparks are the smallest and the only type of event that can only be spontaneous whereas all other events can be both spontaneous and induced by other types of events [79, 177]. It is believed that sparks are a healthy regulation mechanism that allows for a flexible heart function and that a lack of sparks implies a stiff dynamics and would induce heart pathologies [178]. On the other hand, waves are seen as unhealthy events because they are capable of altering the normal front propagation patterns [179].

4.2 Event segmentation

Image segmentation is the process of separating the relevant or foreground pixels of an image from the background pixels. There are many algorithms that can be applied as already discussed in section 2.3 depending on the conditions of the image and the object that has to be segmented.

In the particular case of calcium dynamics, the events need to be segmented from a space-time point of view, in the case of line scans it is therefore 2D and in the case of image sequences it is a 3D segmenting problem. The following two sections cover these cases.

4.2.1 Line scan imaging case

As already described in section 1.4.3, line scans are images composed of a sequence of single pixel lines taken at consecutive frames. In such images, calcium transients will appear as brighter areas occupying the whole spatial dimension and a short temporal size. Traveling waves and mini-waves will appear as tilted brighter areas (the travelling speed can be measured from the tilt) and sparks will appear as localised rises in fluorescence (see figure 42 for some examples with image pre-processing and figure 43, panels d and e, for examples of detection and classification of events). The following subsections describe the image processing techniques applicable to line scans in calcium imaging. These techniques have been successfully applied in several publications in the field of biological science [170, 180].

Line scan pre-processing

The scanned line will usually be taken from within the cell but in some cases it can include a membrane section and regions outside the cell. Therefore cell boundaries need to be detected and this can be done by searching for regions presenting a large spatial derivative of the time integral of the fluorescence. In cases where cell contraction results in a movement of

the cell boundaries, cell limits can be defined by the minimum width of the cell during the experiment thus avoiding the use of fluorescence data from regions outside the cell.

Images should be normalised using a time-dependent basal fluorescence $b_i(t)$ at each spatial location c inside the cell. The reason for using a space-dependent baseline is to remove imaging artifacts of brighter regions, and the reason for using a time-dependent baseline is twofold: first, to account correctly for consecutive events between which fluorescence does not return to the basal level. Second, to correct for temporal drifts in the basal fluorescence due to experimental factors such as photobleaching. If $s_i(t)$ is the time-dependent fluorescence signal at cell pixel i, at each time t, the baseline $b_i(t)$ is estimated as the first decile (meaning lower 10 %) of $s_i(t)$ computed in a time window centred at sample t. In order to avoid overestimation of the basal fluorescence due to the occurrence of large events within the temporal window, the window size is suitably chosen as the maximum number of consecutive time samples in which the fluorescence exceeds the mean fluorescence of the experiment (section 4.3.1 details this process of baseline estimation). Normalised line scan at each pixel i can now be defined as $z_i(t) = s_i(t)/b_i(t)$, so that the normalised fluorescence of release events is measured relative to its local baseline in both space and time. Figure 42 shows some examples of line scan pre-processing.

In order to ensure a common global contrast when comparing sets of images, the contrast of each image should be adjusted so that pixel values range similarly, for example from 1 to the average maximum fluorescence of all the images in the study.

Actual detection

For each pixel location i in the line, a wavelet-based detection applied to the normalised time-dependent fluorescence signal $z_i(t)$ will robustly enhance Ca^{2+} release events. The continuous wavelet transform (section 1.5.5) measures the similarity between the local shape of the signal $z_i(t)$ and the shape of a reference template function (the wavelet). In particular for calcium release events, a good approach is the use of the bell-shaped Gaussian wavelet function with a duration in the interval [50,100]ms, which allows the localisation of calcium release events with a duration ranging from 20-40ms (sparks) up to 300ms (mini-waves). Events with longer durations such as waves or transients can be easily detected using an amplitude threshold given the fact that they are the brightest events in the line scan.

As already mentioned, it is always better to have more false alarms and fewer omissions (by using relaxed constraints in detection such as a wide wavelet scale range or a low threshold) and apply a filtering step after detection. This subsequent filtering step can be performed by removing events that do not present a statistically significant increase in fluorescence with respect to the local activity in the surrounding region. Specifically, a two-sample-Student's t-test with a low significance level *p=0.01* can be used to compare the

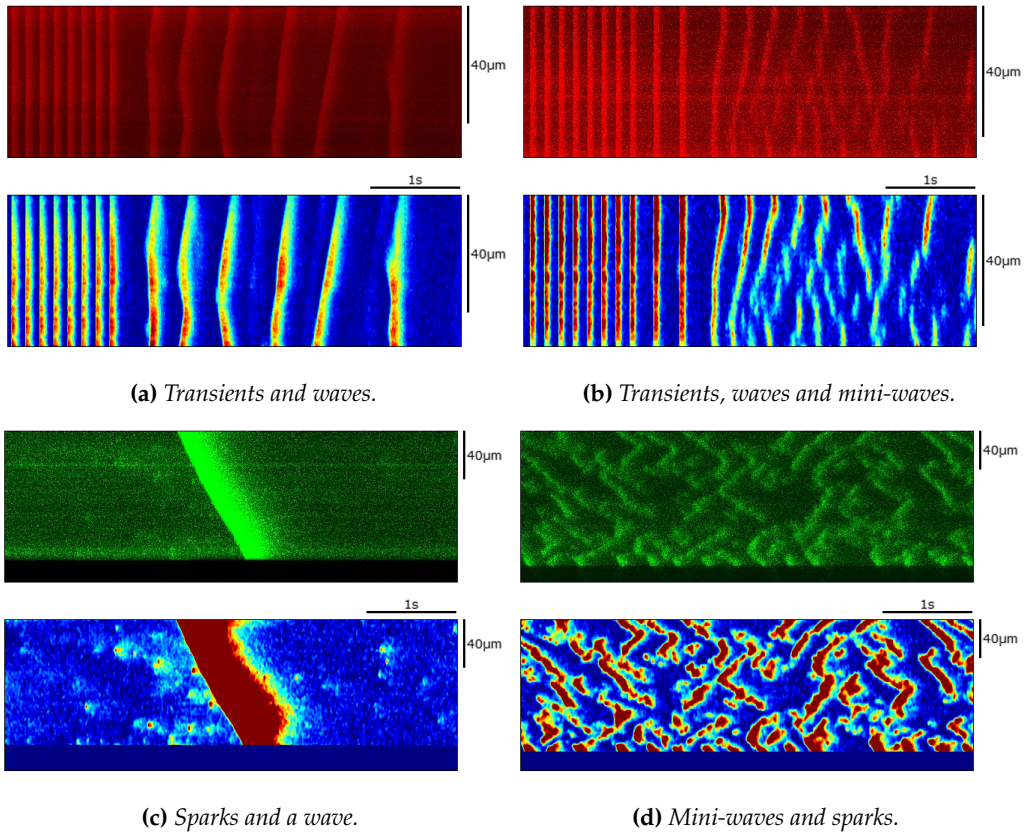

(a) *Transients and waves.* **(b)** *Transients, waves and mini-waves.*

(c) *Sparks and a wave.* **(d)** *Mini-waves and sparks.*

Fig. 42 *Four examples of the line scan pre-processing. In each case the top panel shows the original line scan and the bottom panel shows the image after preprocessing (it has been converted to a single channel image and hence the false colour). We can see examples of brighter horizontal lines removed by a localised baseline normalisation, examples of image cropping due to line scan covering the membrane area, and general event enhancement thanks to the normalisation step.*

distribution of pixels belonging to an event and the distribution of pixels in a surrounding region with equal pixel area.

In order to avoid erroneous grouping of close events in a single detected region an event separation phase should be applied. As an example, events can be skeletonised and divided using a probabilistic clustering method based on orientation, occurrence time and spatial overlap (skeletonisation is a process where a region is eroded until only a line is left [31]). An example of such a method is provided in section 4.4. Indeed, these features allow the distinguishing of cases in which calcium waves with different origins converge or calcium waves with a common origin that propagate in opposite directions. Figure 43 (panels a to c) shows an example of wavelet detection and skeletonisation allowing the event separation.

4.2.2 Frame scan image sequence case

A frame scan is a traditional 2D image, frame as a contraposition to line scans (section 1.4.3). In this case, for studying temporal dynamics we need a temporal sequence of frame

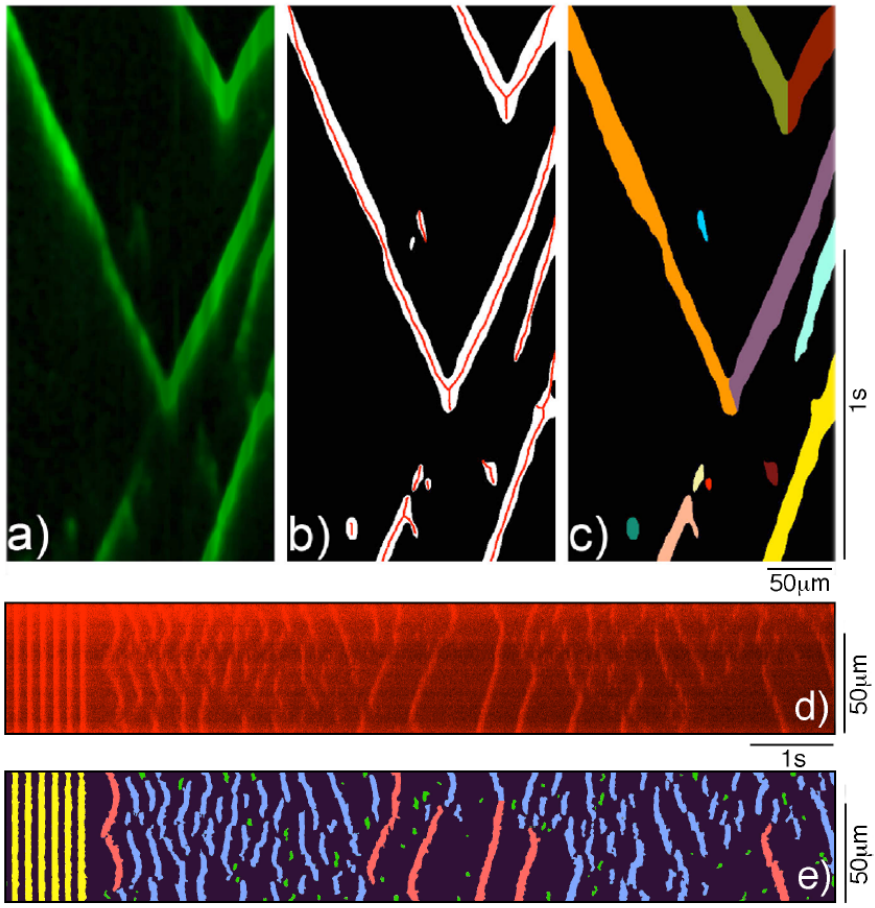

Fig. 43 *Event segmentation in line scans. The top three panels show a region in a line scan (a), event detection in central panel (b) with the wavelet detection (white) and skeletonisation of detected events (red), and final event separation (c) with each event in a different colour (temporal direction is vertical and spatial direction is horizontal). Panel (d) shows an example of line scan with a variety of calcium events and panel (e) shows the detected events and colour-coded classification. In this case temporal direction is horizontal and space direction is vertical. Cell transients correspond to the yellow objects, waves are the red objects, mini-waves blue and sparks in green.*

scans, in everyday terms a film. The methods described in the following sections have been used as a tool in several biomedical publications for linking spark spatial distribution to atrial fibrillation [181, 182], proving causes to changes in spark distribution, morphology and kinetics [183].

Image sequence preprocessing

As always, images will require some sort of pre-processing depending on their original conditions. The two main steps are noise estimation and removal and image normalisation. The first is fairly straightforward (see section 1.5.3) but, as explained in section 1.5.1, there are some considerations to be taken into account when normalising the image sets, especially when the analysis is going to be performed in several experiments. If this is the case, it is very possible that each experiment will have a slightly different fluorophore load, and also that some experiments may have non-specific staining that can shift the normalising factor to a very high value compared with experiments without non-specific staining.

The best approach is to first extract the cell mask as described in 2.4, then use this mask to

measure the time signal of pixels belonging to the cell, next apply detection methods, and finally use pixels belonging to the cell, but not belonging to detected events, to establish basal fluorescence. This method should really be called post-processing normalisation but it is equally valid. Basal fluorescence f_0 can then be the mean of the values below a certain pre-established percentile of the pixels that do not belong to an event. Once f_0 has been measured we can normalise the whole sequence by either applying $F_n = F/f_0$ or $F_n = \Delta F/f_0$ where ΔF is fluorescence increment from baseline $\Delta F = F - F_0$.

Detecting events

There are two basic ways to approach sequences of calcium imaging: one is a two-step method, in which we analyse each temporal frame independently for a later merge of information in the temporal direction, and the other is a straightforward 3D method, in which we take the sequence as a 3D space x-y-t where we have cubic pixels (voxels), with one direction being temporal, and treat the whole volume for the analysis. Each approach has its own pros and cons; the first is easier to work with since it allows performance during development because viewing 2D images is much more simple (when viewing a volume the inner voxels are hidden by the outer) and it is also cheaper in terms of computation cost. On the other hand it is not as neat because it is a two step approach and one has to establish a criterion for the later merging in the temporal direction and, above all, the approach gives priority to either spatial or temporal characteristics depending on the order in which the problem was tackled. In general, good advice is to start with the combined 1D and 2D approach for data inspection but when targeting a definitive analysis go for the 3D view.

Here we cover segmentation of non-propagating events such as sparks, the next section covers segmentation of propagating events. As in general the image sequence will have all kinds of events, a good approach for the 1D temporal direction, as always, is to use wavelet transform to enhance objects with the desired temporal support. We take the time signal for every pixel in the field of view and compute its transform with either a fixed single scale (all objects have similar temporal length) or by taking a range of scales and summing or averaging the result (objects cover a range of temporal lengths). Treatment of individual sequence frames for the case of the 2D spatial analysis will generally require some sort of smoothing as in the case of RyR detection in the previous chapter. The spatial filtering can be of any kind, we have to choose an appropriate filter considering the the size of the objects we are targeting, the resolution of the images and the remaining noise after pre-processing. A typical case for calcium sparks would be convolution by a Gaussian filter with a diameter in the range of the micron. If our approach is 3D we can equally filter by a 3D smoothing template.

The actual segmentation in the combined 1D temporal and 2D spatial and x-y-t 3D cases can be performed with any method depending on the resulting data sets. The simplest

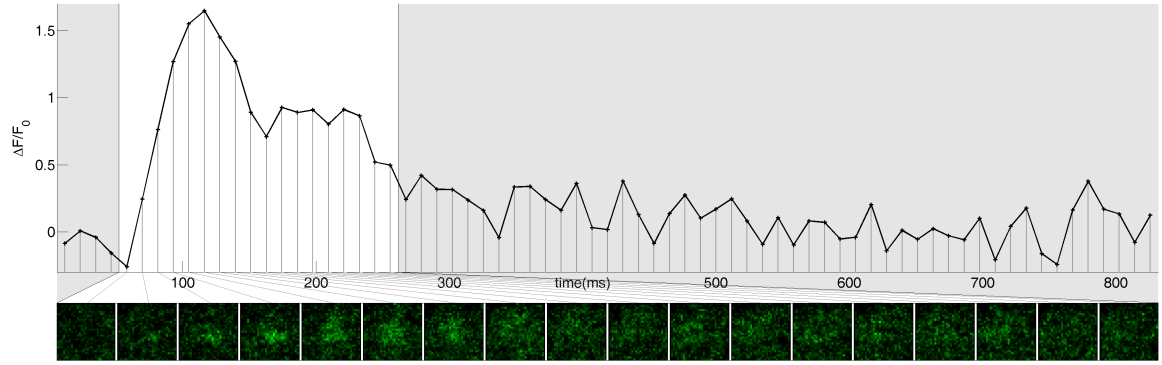

Fig. 44 *A sample trace of a calcium spark from an experiment taken at a spatial resolution of* $0{,}28\mu m/pix$ *and temporal resolution of* $11{,}63ms/frame$*. Each time sample is the mean fluorescence of a* $4 \times 4\mu m$ *region around the centre of the calcium release. Beneath the trace is a cropped image sequence of* $8 \times 8\mu m$ *corresponding to the 18 frames around the maximum (white area in top panel).*

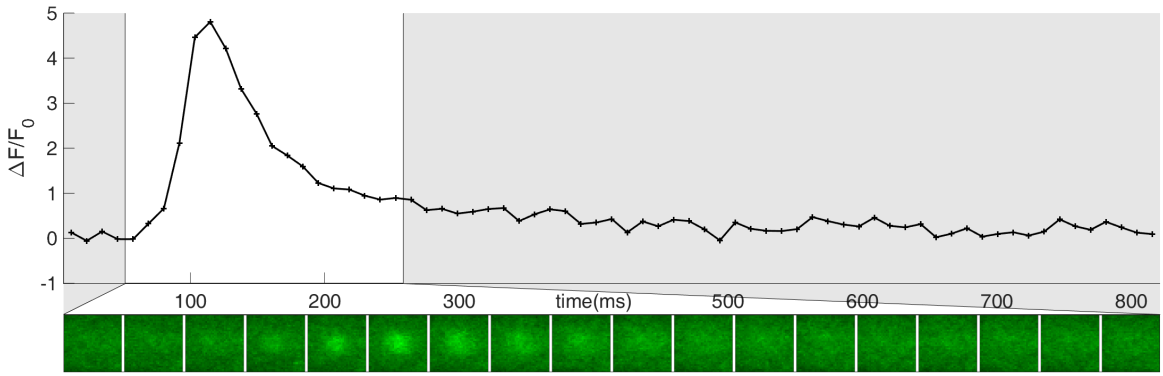

Fig. 45 *Mean spark trace with image sequence. In this case, instead of a single spark we show the mean calcium spark by averaging the traces of all sparks (13) and also the the regions of interest around them are combined into a single image sequence. This reveals both a much clearer profile and calcium diffusion pattern. Spatial resolution is of* $0{,}23\mu m/pix$ *and temporal resolution of* $11{,}63ms/frame$*.*

approach is the straightforward thresholding but in general an almost infallible method is a watershed-based algorithm 1.5.4. As an example, the method applied in [175, 171] was to first filter using wavelet transform in the temporal direction obtaining quite a smoothed 3D volume, with which a watershed transform was used to search for local maxima in the volume. These were treated as candidates for sparks but were further filtered after measuring morphological parameters. This takes us to the rule of thumb in processing calcium imaging: Always use a permissive segmentation step with false alarms and no omissions because objects can always be post-filtered using morphological features.

4.3 Event features

In calcium digital imaging, calcium events can be measured by taking the mean fluorescence value in a region around the event and plotting this fluorescence over time. They all

present common characteristics, such as a sudden rise followed by a slower decay. For this reason a series of signal features have long been established by the community and are used to define and classify an event. Furthermore, measuring these features is crucial to be able to link cellular level conditions to cardiac pathologies. Some of the most relevant of these features are summarised in table 6.

Feature	Abbreviation	Units	Description
Absolute amplitude	AMP	F	Intensity of maximum measured from zero.
Relative amplitude	amp	F	Intensity of maximum measured from local baseline.
Baseline	BL	F	Local baseline intensity before event start.
Time to peak	$t2p$	ms	Time lapse from baseline intensity to maxima.
Rate of rise	RoR	F/ms	Mean intensity variation over time during event rise.
Rate of decay	RoD	F/ms	Mean intensity variation over time during event decay.
Decay constant	tau	ms	Denominator parameter in an exponential fit to the decay of the event.
Decay time	TX	ms	Time it takes for the intensity to reach X percent of that at the maxima (typical features are $T10$, $T50$, $T90$) .
Full duration at half maximum	FDHM	ms	Temporal width of the event measured at half the intensity of the maximum.
Full width at half maximum	$FWHM$	μm	Spatial width of the event measured at half the intensity of the maximum.
Distance to membrane	$d2m$	μm	Distance from the centre of the event to the cell membrane measured on the focal plane (makes sense only for small events like sparks).
Speed	v	$\frac{\mu m}{ms} \equiv \frac{mm}{s}$	Scalar quantity measuring the mean displacement over time.
Velocity	$\vec{v}$	$\frac{\mu m}{ms} \equiv \frac{mm}{s}$	Vector quantity measuring the speed in each of the N predefined axes in an N-dimension space.

Table 6 *Summary of typical event features. F stands for fluorescence, which can be raw image intensity values, or more frequently will be normalised to basal fluorescence: either F/Fo or $\Delta F/Fo$. ΔF stands for fluorescence increment; ΔF=F-Fo.*

These features or a subset of them can be used for classification of events into the set of existing predefined event types (table 5) or blindly classified into new types depending on the experiment. In general though, as the detection technique will already be different depending on the type of event, it is not very common to use a posterior classifier.

The size one chooses for the region around the event is very important towards the resulting time signal. The bigger the region chosen the smoother the signal will be but, at the same time, the smaller the intensity of the event because we will be averaging with surrounding lower fluorescence values. A good procedure is to choose a fixed radius that is slightly smaller than the average size of the events, ensuring we are not diluting the signal with surrounding background values, but big enough for the signals to be clean and allow measurement of signal properties. For example if we are measuring sparks, a diameter of one or two microns will allow the inclusion of the main area of the spark while at the same time

averaging enough pixels to obtain a reasonable signal-to-noise ratio. This of course is totally dependent on the time-space resolution of the experiment and therefore every case will be different and no absolute statements can be declared.

In some cases, measuring these features is fairly straightforward, in others the process can be a little problematic. The following sections are dedicated to each of the signal features where some of the typical problems one can encounter are detailed, together with possible solutions to overcome them.

4.3.1 Baseline

Baseline or local baseline, is defined as the fluorescence value *just before an event*. This is a somewhat ambiguous definition and for this reason some considerations need to be made when attempting to measure an event's baseline.

The first consideration is that fluorescence decays with usage during an experiment, meaning that the overall intensity of the signal will gradually be dimmer and so the local baseline before a particular event will be different depending on when the event occurred in an experiment.

Also, the baseline signal for a particular time signal should be obtained prior to event segmentation for two main reasons. For a start if we try to obtain the baseline value from a cropped section of a time signal, we are subject to possible errors in the event segmentation. A common case is that of events that occur *one on top of the other*, for example a series of sparks that appear during a calcium wave. If we consider only the time signal of the spark, we will most probably overestimate local baseline because of the high fluorescence values produced by the wave.

The other reason is that the only way to estimate the baseline of a signal is to establish a sliding window from which to take the lower values, and the size of the sliding window will depend on the events that take place in the signal.

The procedure is simple, we choose a window of a certain temporal length and for a particular position of the window we measure the baseline value. We then slide the window one sample along the signal and repeat the operation to obtain the next value for the baseline signal. The way in which we take the baseline value in the window can be chosen from a range of possibilities, but the overall structure of the method is the same. One simple and computationally effective approach is to take the minimum value directly, but this method is not recommended because it is subject to noise and can easily produce jagged signals of the size of the window.The best approach is to take a particular low percentile of the window or, similarly, the mean of the values below a particular fraction of the range of values in the signal.

This method therefore requires choosing two parameters. The first is the fraction threshold or percentile. One should choose a low enough value but one that makes sense with the length of the window in samples, so that we are not simply forcing the minimum value in the window. For example, if the window measures 10 samples, taking the lower 1 % makes no sense because we will simply be taking the lower of the 10 samples.

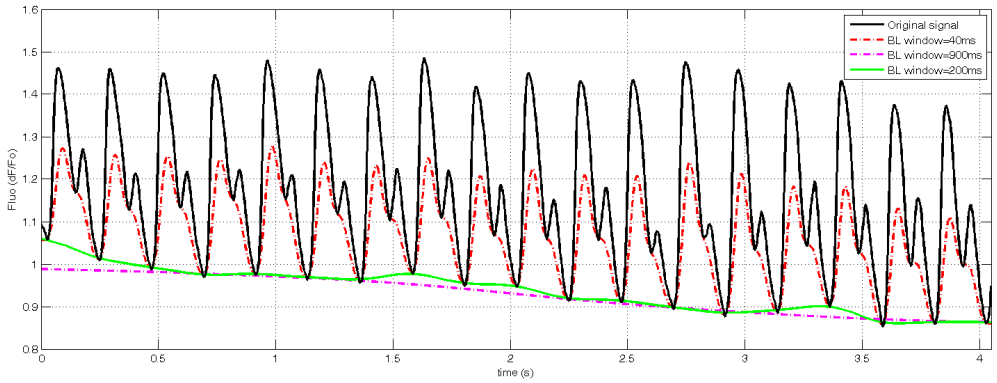

Fig. 46 *Baseline method applied to a particular signal of a cell producing alternating calcium transients. The method applied is by taking the mean of the lower 6 % of the sliding window. The black plot is the original signal and the green plot is a baseline signal taken with a reasonable window size. The other two are respectively a clearly too small window (red, the baseline follows the signal up the transients), and a too big a window (magenta, during the first and last half second the baseline is static).*

The other parameter is the length of the window and it is not so simple to choose. The problem of the method is that the resulting baseline signal is shorter than the original signal by exactly the length of the sliding window in samples (minus one). These empty samples can be either at the start or the end of the signal, or even half on each side depending on how we applied the algorithm, and they have to be filled typically by extending the first/last value to the extension of the original signal. Obviously, we do not want to have to *invent* too many values, so we want our sliding window to be the shorter the better. On the other hand, the sliding window cannot be too short, otherwise our baseline signal will simply become a smoothed version of the original signal. We cannot accept a sliding window that is smaller than the bigger events taking place in the signal because otherwise our baseline will follow the events as they go up in intensity (see figure 46 for clarity). A good option is to take a window of the duration of the largest possible event, which, in our case has to be calcium waves that are in the order of magnitude of the second. Nevertheless, if we know for sure that our signal contains only sparks, we can then reduce the size of the window down to a tenth of a second.

Global baseline corresponds to a background fluorescence value that is measured in the first frames of a sequence and is therefore higher than the local baseline in events throughout the sequence.

4.3.2 Amplitude

Amplitude is the peak value of the intensity of an event measured from either zero intensity value (absolute amplitude) or from global or local baseline (relative amplitude). The only consideration to take into account is that if the signal is noisy one cannot simply take the maximum value in the event for it will really be measuring the peak noise. The first and most simple option is to ensure a pre-filtering step to clean the signal in order to assign the maximum value to the amplitude but this can deform the original time signal, especially in cases when it is not symmetric, and will usually produce an under-estimated value of the amplitude.

Another option is to fit a known function to the whole event and take the maximum of the fit. This second option is not always possible: when we are taking the mean fluorescence values of a whole cell or a large area, the time signal is a rounded function somewhere in between a Gaussian and a rounded exponential (also named *roex*; the subtraction of decaying exponentials, the first with a slightly higher decay time). Usually though, the time signal belongs to a small region of interest and, especially for sparks, will consist of a sudden uprise followed by a long decay.

Therefore, the third option and most usual approach is to divide the signal into two segments and fit a straight line to the first half and an exponential decay to the second half. The time a which the maximum intensity occurs can then be defined as the time at which the these two fits intersect. Choosing the point at which to divide the signal is not so crucial now so one can filter the signal to smooth out the values simply by taking the maximum as the breaking point.

4.3.3 Time to peak

Time to peak is the time it takes from the initial upstroke to reach the maximum of the signal.

The only complication here is to properly define where this upstroke starts because it depends on the segmentation step. If the temporal borders of the event are either exact or they crop some of the event off, then one can set the start of the event as the start of the upstroke, but this is not usually the case and an initial baseline segment has to be excluded from the signal. A robust method to find the exact sample where the the upstroke starts is depicted in figure 47. It consists in fitting a constant value followed by a straight line to the part of interest of the signal (from start to maxima) and taking the sum of distances from the fit to the real signal. The position i that separates the constant fit and the linear fit for which the sum of distances is minimum is the best fit and so i can be set as the sample where the upstroke starts. For a given sample i the constant value is directly the mean of the signal from the first

sample to i, the rest of the signal until the sample where the maximum value is fitted to a linear polynomial.

4.3.4 Rate of rise

Rate of rise is the rate at which the fluorescence grows during the uprise. It is used only for rapid increases in calcium release events in which the upstroke is near lineal and so the rate of rise is a constant given by the slope of a straight line fit to the upstroke.

A direct way to obtain this feature is to simply divide the relative amplitude by the time to peak but a better approach is to fit the upstroke to a straight line. As described in the previous section, by fitting a constant term followed by a straight line to the samples from the start of the signal to the maximum value, we can obtain the exact sample where the upstroke starts. The coefficient of the degree-one term of the linear fit is taken as the rate of rise.

4.3.5 Decay constant

The intensity decay in calcium events is close to exponential. In some cases it has been modelled by a double or even triple exponential decay justified by the several processes involved (diffusion of calcium in the cytosol, regathering into the sarcoplasmatic reticulum and ejection through the cell membrane out of the myocyte), but the most common is a fit to a simple exponential function of time t: $I(t) = Ae^{-\frac{t-t_0}{\tau}}$, where e is the natural exponential base (the transcendental number $e = 2{,}7183..$), A the normalisation factor, t_0 the temporal shift and, most importantly, τ (tau) the decay constant.

This decay constant is a parameter expressed in the temporal units we are sequencing our images (generally ms) and it indicates the rate at which the signal intensity decays regardless of its current intensity value. It is defined as the time it takes for the intensity to decay by a factor of e^{-1} ($\approx 0{,}368$), that is for the intensity to reduce to around 36.8 % of its value. It is important to see that this decay gives no measure of the temporal duration neither of the amplitude, it is only a measure of the rapidness of the decay.

Given a temporal signal of an event, with the time of the maximum already found (see sections 4.3.2 and 4.3.3), one can crop the decay so that the maximum corresponds to the initial sample, and divide the decay into the value at the maxima so it is normalised to unity. This is equivalent to setting the temporal shift to zero ($t_0 = 0$) and the normalisation factor to unity ($A = 1$), thus allowing a computationally faster fit to a single parameter function: $I(t) = e^{-t/\tau}$.

The one consideration to take into account when fitting the decay is that it is possible for a new event to occur before the previous one has decayed to rest state. In this case the only solution is to exclude the new event from the decay so the fit is not altered.

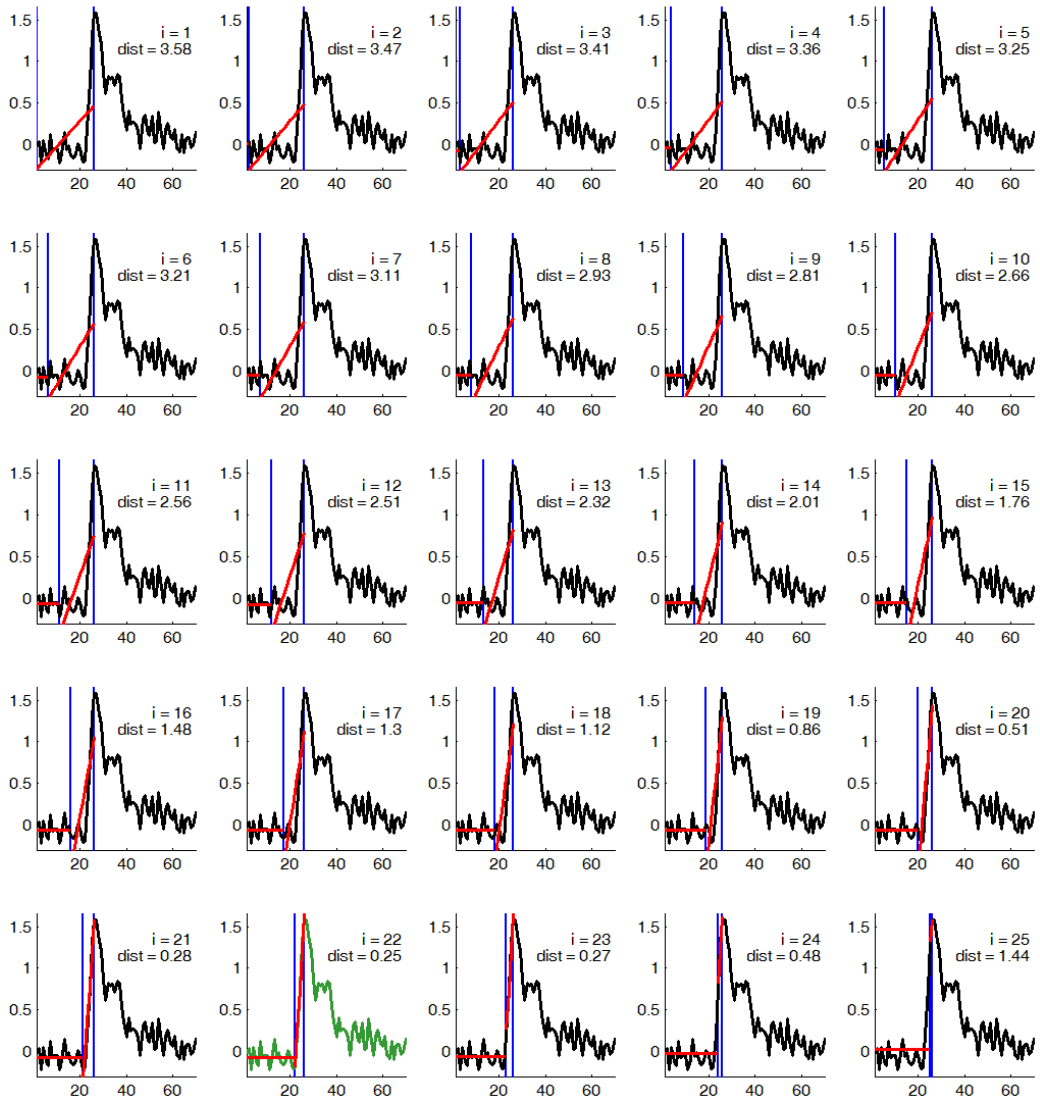

Fig. 47 *Finding the upstroke start of a spark time signal. We scan every possible sample i (in this case from $i = 1$ to $i = 25$; the sample of the maximum is $i = 26$) and in every case we fit the signal with a constant value from the start to i and a straight line from i to the maxima. In each case we take the sum of squared differences between the fits and the signal to obtain a mean distance value. The position that produces a minimal distance is set as the start of the upstroke (in this case $i = 22$ with the signal plotted in green). The two blue lines mark the sample separating the two fits and the sample of the maximum.*

4.3.6 Decay time

Decay time is another feature to describe the decay and it is a little more intuitive and less mathematical than the decay constant because it is given in *normal* numbers and not an irrational constant. It is defined as the time it takes for the signal to be reduced by a certain percentage and it is expressed as TX where X is the particular percentage. For example $T10$ would be the time it takes for the signal to decay from the peak value to exactly 90 % of the peak value. The most commonly used are $T10$, $T25$, $T50$, $T75$ and $T90$. The decay constant

from the previous section expressed in this way would be $T63{,}2$ because it was defined as the time it takes for the intensity to drop to 36.8 %.

Unlike the decay constant in the previous section, these parameters do not require a fit to a particular function and they can be measured straight from the raw signal, but it is likewise recommended to previously fit a function to the decay and measure the parameters from the decay so that the results are subject to noise.

4.3.7 Rate of decay

This is the mean intensity variation in the decay phase of the event, as it is a rate, we are making a linear approximation. This is an uncommon simplification applicable only when the decay is very fast or the temporal resolution is too low or events are so close together that the calcium cannot decay to resting state. In any case measuring it is a mirror version of rate of rise described in section 4.3.4 together with section 4.3.3 and figure 47.

4.3.8 Full duration at half maximum

This is simply the temporal width of an event at half its intensity height. Assuming the uprise of the event is linear, this feature can be defined as the decay time to 50 % plus half of the time to peak: $FDHM = T50 + t2p/2$.

4.3.9 Full width at half maximum

Of all the event features described up to this point this is the first that is of spatial nature instead of temporal. It is defined as the width of the event measured at half its intensity.

Several techniques for measuring this feature are described in the previous chapter in section 3.3 and an example is provided in figure 30 where objects are fitted to a Gaussian function across several cross-sections and the width is measured from the respective fits. Measuring the size of the spark is very much the same: we take the image where the spark is at its maximum (or the mean image of a window around the maximum if the images are noisy) and proceed to fit a Gaussian function in a series of cross-sections. In each of these cross-sections we take the width at half the height of the fit. See figure 48.

4.3.10 Distance to membrane

This is not so much a feature but a characteristic that only makes sense for localised small events like sparks and is useful to analyse the spatial distribution of the events. In general we

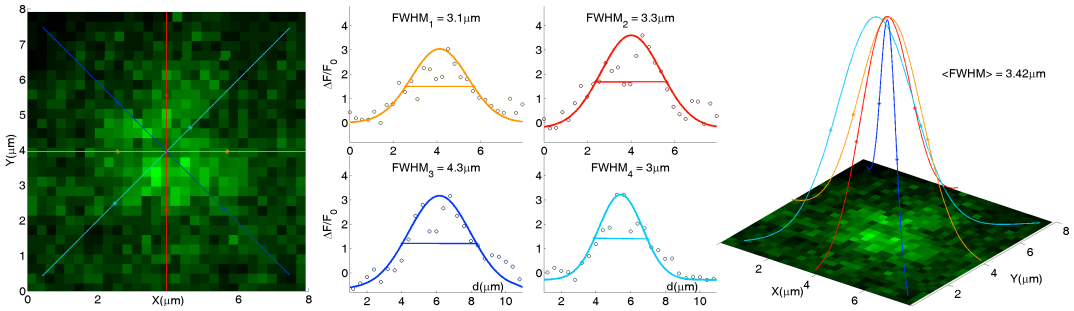

Fig. 48 *Measuring full spark width at half its maximum. In this case the left panel shows the mean image of five frames around the maximum ($\sim 60ms$) and four cross-sections at $45°$ from each other, each taken as the mean of a $3pix$ wide line ($\sim 0{,}8\mu m$). The central panel shows each of these cross-sections with the Gaussian fit and the diameter measured from the fit. The last panel shows a 3D representation of the fits on the image where the fits have been normalised in height for the sake of aesthetics. This example is the same spark as in figure 44.*

will have a single image so we can only measure the shortest distance to the membrane in the focal plane or plane of the image. The best approach is usually first to establish the pixels that belong to the cell as described in section 2.4 obtaining a cell mask, and from this mask obtain a contour of it by, for example, subtracting its one-pixel eroded version of itself. Next one measures all distances from the event location to the pixels in the contour and takes the smallest of them.

4.3.11 Speed and velocity

Speed is a scalar magnitude that gives the rate of change in position, whereas velocity is the vectorial equivalent and the norm of which is speed. In both cases, this magnitude is the derivative of position and so the straightforward approximation is to define speed at a particular frame $s(f)$ as the difference of position $s(f) = p(f) - p(f-1)$. There are many other definitions for a discrete estimation of derivative, like using the next frame instead of the previous $s(f) = p(f+1) - p(f)$, or maybe defining speed as a shifted vector with measureents in between samples (at pixel vertices in 2D) $s(f - \frac{1}{2}) = p(f) - p(f-1)$, or also using higher order approximations $s(i) = p(f+1) - 2p(f) + p(f-1)$. Nonetheless, the derivative will be a signal with a length at the most one sample shorter and this has to be taken into consideration; for example one cannot expect to give speed at the first sample $s(1)$ if it is being calculated using the most common backward definition $s(f) = p(f) - p(f-1)$ because position at frame zero $p(0)$ is not defined.

4.4 Wave tracking

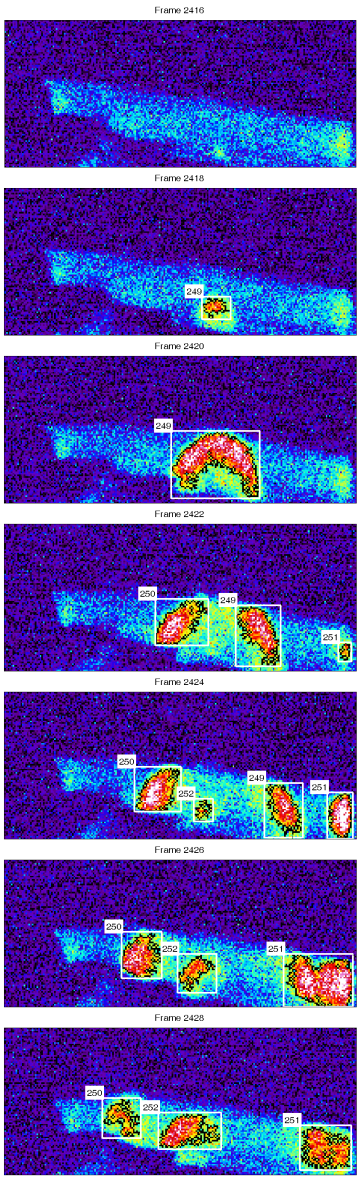

Fig. 49 *Video tracking of calcium waves. Sample frames are shown in which waves are originating at the centre of the cell and travelling in opposing directions. Images are displayed in false colour.*

Large propagating objects such as waves can also be approached as described in the previous sections; temporal detection based on object duration, spatial detection based on object size, with these two methods applied in parallel or consecutively. Alternatively detection can be performed in the 3D volume as x-y-t blobs, if the resolution allows it, by simply segmenting objects in the volume.

Another option though is to apply a motion-tracking algorithm. This kind of approach is especially useful when the image sequence contains many propagating events at a time, because a raw 3D segmentation will tend to join them if the events either collide or start from the same point.

There are two main types of tracking algorithm; the first being the so-called detect-before-track (DbT) and the other being the track-before-detect (TbD). The first group uses some segmentation method to separate objects from background and then applies a tracking method to assign each object in one frame to an object in the next frame (i.e. deciding which object is which in two different frames). The second group, the TbD family, uses some blind method to follow moving objects without clearly resolving the limits of the object and after estimating the new position, at each frame one has to apply a segmentation method to clearly define the object. Each method has its pros and cons and is useful in different conditions. For example DbT is useful when object movement is such that their position does not overlap from one frame to the next (typically small fast moving objects related to the frame rate), when the scene has many objects that can change size throughout the sequence or when objects tend to leave the field of view or new ones appear. On the other hand, TbD is very powerful when we have few objects that do not change much in size and move slowly enough so that their position is similar from one frame to the next and therefore near continuous.

An example of an application case is detailed next with a sequence of images of mouse ventricular myocytes with a heavy calcium load that triggered many spontaneous waves.

For this particular problem a DbT method was applied and a custom probabilistic tracking algorithm was developed [184] and further used in [185].

Knowing the image sequence contains waves and that these are the brighter events in the sequence, the detection step is straightforward; simple thresholding at a fraction of the value distribution in the sequence provides a fast and easy way to separate foreground objects from the background. At each frame a series of features are measured for each surviving region and these features are used to establish the probability of a region i in frame f being the same event as region j in frame $f+1$. A general expression for this probability function in terms of a weighted normalised discrepancy is:

$$P_{ij}^{f} = 1 - \sum_{l=1}^{N_f} \omega_l \, \hat{d}_l$$

where N_f stands for the number of extracted features, ω_l is the weight assigned to feature f_l and $\hat{d}_l$ is the normalised feature discrepancy $d_l = f_l(i,f) - f_l(j,f+1)$. This discrepancy is simply the difference in values the feature takes from one frame to the next, and it is normalised to a value that can be the maximum value possible for the difference: $\hat{d}_l = d_l / L_l$. For example if our feature is distance, the normalisation value L can be the maximum distance possible in the image, i.e. the image diagonal, so that any difference in position between objects is a value in the interval [0,1].

The particular probability function used in the study was built using four features: position norm, the two velocity components and area, and so the its complete expression was:

$$P_{ij}^{f} = 1 - \omega_1 \frac{\|\vec{r}_i - \vec{r}_j\|}{\sqrt{{L_x}^2 + {L_y}^2}} - \omega_2 \frac{|v_{ix} - v_{jx}|}{L_x} - \omega_3 \frac{|v_{iy} - v_{jy}|}{L_y} - \omega_2 \frac{|a_i - a_j|}{L_x L_y}$$

where i and j represent two detected regions belonging to consecutive frames, L_x and L_y are the spatial dimensions of the frames, $\vec{r}_i$ and $\vec{r}_j$ are the position vectors of the centre of each region, v_{ix}, v_{jy}, v_{ix} and v_{jy} are the velocity vector components of each region, and a_i and a_j are the areas of each region.

This probability function returns a maximum value of unity when we are comparing two regions that have exactly the same area, position and velocity, and will return a minimum value of zero when the difference of the features is the highest possible (determined by the frame dimensions). It is key to say that the velocity parameter is measured from frame to frame, so when comparing two particular regions, the velocity will be set assuming that they are the same region. That is to say that when comparing regions i in frame f and j in frame $f+1$, the velocity of j will be set as the difference $\vec{v_j}^{f+1} = \vec{r}_j^{f+1} - \vec{r}_i^{f}$ but then when comparing another region, say k in frame f with the same j in frame $f+1$, the velocity of j

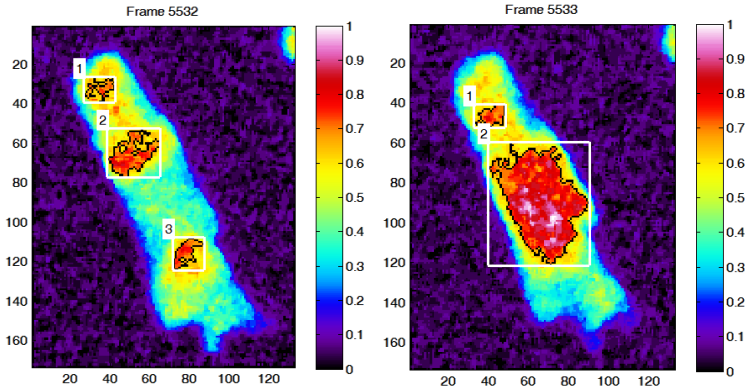

(a) *Frame A.* **(b)** *Frame B.*

	B1	B2
A1	0.95	0.82
A2	0.94	0.93
A3	0.82	0.92

(c) *Crossed probabilities.*

Fig. 50 *Two consecutive frames, the first with three detected regions and the second with two detected regions together with the associated probabilities. We see that region A1 is more likely to be region B1 (P=95 %) and that region A2 is also more likely to be region B1 (P=94 %). The algorithm would start with the highest value in the table (95 %) and assign A1 to B1. Then it would ignore all other values in that row and column and continue with the highest value remaining (93 %), therefore assigning A2 to B2. A3 would be left with no assignation in this case. Frame rate is* $80ms$ *and thresholding value is* $0{,}7$.

will be set as the difference $\vec{v_j}^{f+1} = \vec{r}_j^{f+1} - \vec{r}_k^{f}$. Figure 50 shows an example of two frames with a few regions and the table shows the crossed probabilities measured for these regions.

Now for a given region in a given frame, it is suffice to choose the region of highest probability from the following frame, discarding all other combinations. The assignation should be made starting from the higher probability values and going down from there (see figure 50). However a certain probability minimum has to be required to associate two regions (depending on the experimental setup), otherwise it could happen for example that a calcium wave disappearing out of the field of view at one end of the image could be associated with a new wave appearing from the opposite end of the image in the following frame if no other events are present at that time.

4.5 Connecting the scales

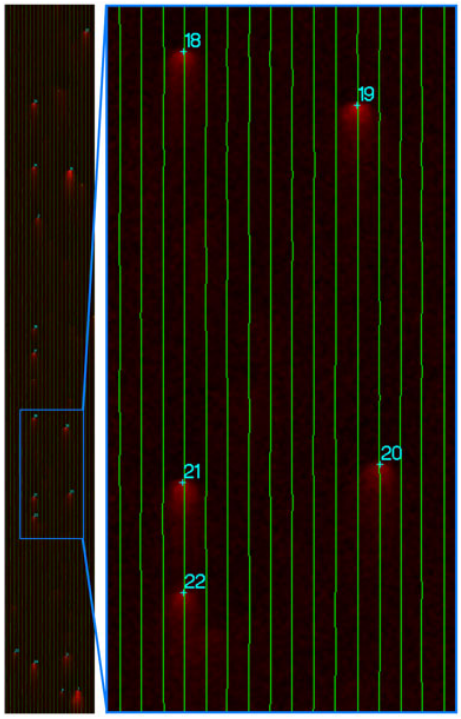

Fig. 51 *RyR2 cluster and spark detection. Left is an example line scan with calcium concentration in red. Right hand panel is a detailed area of the line scan. Vertical lines are the RyR2 cluster position detected and in cyan the detected sparks.*

As explained in the previous chapter, ryanodyne receptor 2 (RyR2) are molecular channels that release calcium ions from the sarcoplasmatic reticulum into the cytosol in order to initiate cell contraction. As explained in this chapter, sparks are localised calcium release events that may or may not trigger further events such as waves. Co-localisation refers to the act of testing if the locations of both RyR2 and sparks correspond to the same positions. This was part of a study in which the aim was to prove that calcium sparks were only produced in RyR2 clusters [170], and so there were no other ion channels responsible for calcium release. The analysis was performed both in line scans and frame scans.

For this, first RyR2 cluster position was determined (detailed in the previous chapter, section 3.5.4) and then sparks were detected using the wavelet procedure detailed in the this chapter (section 4.2) . Sparks were also filtered using spark morphology features described in this chapter and with the final set of sparks the distance to the nearest cluster was then measured.

The results of this study can be seen in figures 51 and 52, where it can be seen that sparks form almost exclusively in RyR2 locations. The few exceptions found in the study seemed to be due to out-of-focus sparks in areas where the confocal plane was in between RyR2 clusters. The frame scan version of these results is in figure 53.

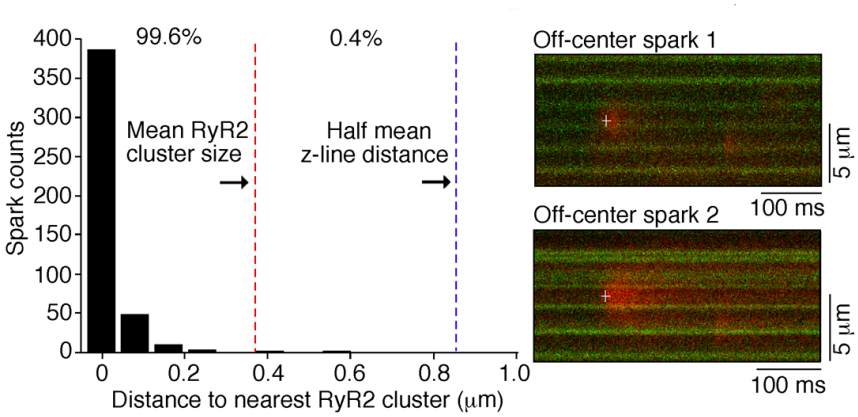

Fig. 52 *Results of the co-localisation study in line scans [170]. The only examples of sparks that do not fall in the apparent RyR2 cluster size seem to be either very close to the cluster or in an area where the confocal plane is between clusters and therefore the spark is most probably out of focus.*

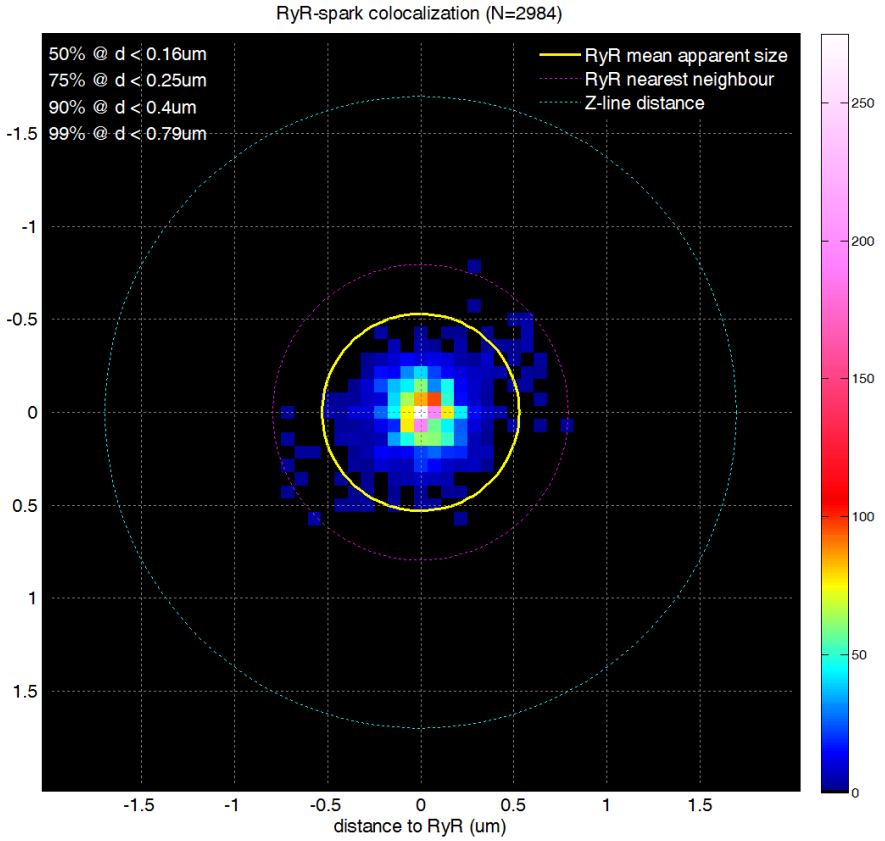

Fig. 53 *A 2D histogram of the spark occurrences around the closest RyR2 cluster. The total spark count is $N = 2984$ of which over half fall in the central 9 pixels (spatial resolution is $0{,}07 \mu m/pix$). As can be seen the outer sparks fall in approximately one nearest cluster neighbour distance, suggesting that possibly these sparks were released from an undetected cluster and were therefore wrongly associated to a neighbouring cluster.*

4.6 Results

As in the previous chapter, here some of the results directly obtained from the methods in this chapter are presented and a summary of the datasets form which they were measured is also provided. Whenever a value is reported accompanied by a plus/minus sign ($\pm$) and second value behind it, this second value always corresponds to the standard deviation (not standard error as commonly used in biology).

4.6.1 Event detection in line scan

Section 4.2.1 describes a method to segment events in line scans, and classify them by means of their morphological features. This method was developed specifically for a study that related certain drugs to the increase in mini-waves versus decrease in waves [180] and successfully quantified the mentioned results.

Total detected event numbers	**285 linescans**
Sparks	24563
Mini-Waves	14245
Waves	1065
Stimulated Transients	232
Spontanous Transients	167
Other Events	708
Total	40980

Table 7 *Events detected in line scans.*

The whole dataset with which the method was tested comprised a total of 285 line scan images, all together summing a scan distance of $\sim$33mm and a scan time of $\sim$48min. These images where analysed using the described methods, to detect over 40k events in total. Table 7 gives the exact number of events detected of each type, described in section 4.1. The term *other events* refers to discarded events due some features not being measured, usually because the events are either too close to image border or are overlapping some other events.

The line scan dataset was divided into 6 different subsets depending on the drug treatment that had been applied to the cells. Of the subsets, 4 of them where actually divided into a *before* and *after* the drug set, so they actually contributed with 8 sets. The other two were divided into 3 due to the study purpose; two independent phenotypes plus the expression of the two at the same time (further explained in section 7.2). This produced a total of 14 subsets, each of which was analysed independently and the properties of all detected events where measured.

The relevant events for the study were actually only sparks, mini-waves and waves, and this is the reason why these are the ones reported in the following tables. The first, table

	#events/scan			#events/s	#events/mm/s		FluoContribtion		
DataSet	waves	m-wav	spks	waves	m-wav	spks	waves	m-wav	spks
RC	8.2±0.4	3.4±0.5	10.6±2.2	0.9±0.0	6.7±1.0	22.6±5.5	81 %	18 %	1 %
PLN	1.3±0.4	56.7±4.0	88.0±9.9	0.1±0.0	61.6±4.3	95.2±10.5	6 %	70 %	24 %
PLNRC	2.3±0.3	41.1±2.3	52.9±3.7	0.2±0.0	61.6±2.4	85.1±6.5	10 %	77 %	14 %
ISO RC	5.6±0.8	5.7±2.0	16.8±2.9	0.3±0.0	6.0±2.3	15.4±2.7	87 %	11 %	2 %
ISO PLN	2.3±0.6	29.7±12.8	64.5±18.5	0.1±0.0	15.0±6.2	31.6±8.8	22 %	69 %	10 %
ISO PLNRC	2.5±0.6	24.7±4.8	113.2±19.5	0.1±0.0	24.7±4.8	113.2±20.0	25 %	53 %	22 %
befCaff+6mMCa	6.4±1.0	53.1±9.1	69.9±9.4	0.6±0.1	41.8±6.7	55.3±7.1	41 %	54 %	5 %
0.5mMCaff	2.9±0.9	78.6±9.2	221.5±13.5	0.3±0.1	48.7±6.1	136.9±8.6	17 %	67 %	16 %
KRH+6mMCa	4.6±0.9	56.0±7.6	59.1±10.7	0.4±0.1	39.2±4.4	40.5±6.2	27 %	68 %	5 %
LiCl+6mMCa	7.6±1.2	91.4±12.3	113.8±11.2	0.7±0.1	61.7±6.9	79.1±8.7	29 %	65 %	6 %
befBayK+6mMCa	5.3±1.2	70.8±11.4	81.2±7.0	0.5±0.1	51.2±8.3	58.3±5.8	30 %	65 %	6 %
1μmBayK	4.6±0.8	104.1±13.8	131.1±15.7	0.4±0.1	66.4±8.1	84.3±9.9	19 %	73 %	8 %
befBHQ	2.0±0.4	106.3±9.3	166.1±19.0	0.2±0.0	63.8±6.4	101.0±15.5	7 %	82 %	11 %
aftBHQ	2.5±0.3	3.8±1.6	21.9±7.2	0.2±0.0	2.7±1.8	12.3±4.3	97 %	2 %	1 %

Table 8 *General event properties of the line scan datasets. For each dataset, three main properties are presented for each of the three relevant types of event . Of the three properties, the normalised number of events (central columns) was performed per unit time and area in the case of sparks and mini-waves, but only per unit time in the case of waves. This is because the line scans only scanned a small section of the cell and so the waves always would cover the whole scan distance, therefore it made no sense to normalise by spatial dimension because the real size was actually unknown.*

8, gives the general event numbers (event per scan, event per unit area and fluorescence contribution), where as the second, table 9, gives the three most relevant features (amplitude, duration and rate of rise).

	Mean Amplitude			Mean FDHM			Mean RoR		
DataSet	waves	m-wav	spks	waves	m-wav	spks	waves	m-wav	spks
RC	0.9±0.0	0.8±0.0	0.3±0.0	118.6±1.6	114.1±2.4	37.4±2.2	20.4±1.3	19.0±2.0	8.0±1.0
PLN	0.6±0.1	0.6±0.0	0.4±0.0	68.8±2.2	59.9±1.1	44.1±0.8	14.7±1.8	14.8±1.0	11.5±0.8
PLNRC	0.7±0.0	0.6±0.0	0.3±0.0	69.8±1.4	65.1±0.9	43.5±0.4	14.8±0.8	13.0±0.7	9.1±0.6
ISO RC	2.8±0.2	2.6±0.2	0.6±0.0	155.4±12.9	144.4±16.9	35.3±0.7	64.5±11.5	61.1±12.0	27.9±3.4
ISO PLN	3.4±0.2	2.3±0.2	0.9±0.1	155.5±21.0	94.7±3.3	41.4±1.5	266.1±80.8	39.4±3.9	38.3±3.7
ISO PLNRC	3.3±0.3	1.9±0.1	0.9±0.1	185.3±49.8	99.7±4.1	44.7±1.2	145.4±29.5	27.3±2.5	27.1±1.8
befCaff+6mMCa	4.4±0.2	2.8±0.1	1.3±0.1	131.8±9.9	92.7±2.8	44.6±1.0	74.3±5.6	44.0±2.8	41.6±2.5
aft0.5mMCaff	2.2±0.2	1.3±0.1	0.7±0.0	180.9±7.7	111.1±2.1	49.9±1.0	25.5±4.0	12.3±0.9	14.4±0.8
KRH+6mMCa	3.8±0.2	2.6±0.1	1.1±0.1	95.6±2.7	79.5±1.3	43.0±0.9	68.0±6.2	48.5±3.4	39.9±3.5
LiCl+6mMCa	2.6±0.3	1.9±0.2	0.7±0.0	92.4±2.0	77.0±1.4	41.1±1.3	53.5±5.8	35.0±3.4	26.6±2.5
befBayK+6mMCa	3.4±0.2	2.4±0.1	1.0±0.0	122.2±8.8	89.2±2.5	45.6±0.9	56.3±4.7	37.4±2.7	33.0±2.1
1μMBayK	3.0±0.2	2.0±0.1	1.0±0.0	108.9±4.2	85.2±1.7	46.2±0.7	47.7±4.5	32.0±2.2	28.3±1.7

Table 9 *Morphological features of the events in the line scan datasets (continuation of previous table).*

4.6.2 Transient parameter measurements in line scan

This study was a follow-up on the previous, in this case though the focus type of events where whole cell transients produced by external stimulation of the cells. Again, line scan images where taken and the task was to detect the transients and extract the relevant parameters. As the transients were stimulated, the pacing frequency was known so it was relatively simple to detect the transient peaks. Simply by correlating the time signal with a train of deltas separated by the appropriate distance would give a lag between the signals from which the position of the transients was determined. Then the parameters of each transient were measured by following the methods described in section 4.3.

Set ID	# of linescans	# of transients
EQwt#9 GCaMP6f-T	46	426
EQwt#11 GCaMP6f-T	68	625
EQwt GCaMP6f-T	66	603
EQHet#10 GCaMP6f-T 500uMCaff	32	294
EQHet#10 GCaMP6f-T	36	330
EQHet#10 GCaMP6f-T 1mMCaff	38	349
EQHet#10 GCaMP6f-T 100nMIso	48	440
EQHet GCaMP6f-T	106	966
EQHet 500uMCaff	28	260
EQHet 100nMIso	28	256

Table 10 *Transient parameter measurement line scan dataset.*

Set ID	amp($\Delta F/F_o$)	tau(ms)	ror($\Delta F/s$)	t2p(ms)	fdhm(ms)	T10(ms)	T50(ms)	T90(ms)
EQwt#9 GCaMP6f-T	0.82±0.37	95±16	13±7	105±26	184±31	110±18	176±27	263±40
EQwt#11 GCaMP6f-T	0.95±0.38	101±16	13±7	112±26	188±31	111±19	181±28	274±41
EQwt GCaMP6f-T	0.63±0.36	98±20	9±6	110±24	179±30	105±18	173±29	262±45
EQHet#10 GCaMP6f-T 500uMCaff	0.84±0.40	104±20	12±7	119±25	205±30	121±21	192±33	287±49
EQHet#10 GCaMP6f-T	0.89±0.51	103±26	13±9	108±19	188±28	114±24	184±41	278±63
EQHet#10 GCaMP6f-T1 mMCaff	0.87±0.37	135±297	11±6	137±28	236±44	139±27	221±37	327±53
EQHet#10 GCaMP6f-T 100nMIso	1.63±0.51	87±14	29±12	103±19	190±24	113±16	173±24	253±35
EQHet GCaMP6f-T	0.65±0.34	114±29	10±6	118±38	215±46	132±32	210±47	312±67
EQHet 500uMCaff	0.73±0.36	100±19	12±7	98±27	185±39	115±25	184±37	274±53
EQHet 100nMIso	1.20±0.41	85±14	26±10	86±16	174±19	110±16	169±25	247±37

Table 11 *Morphological features of the transients.*

This section is mentioned as further proof that the methods described in the previous chapters have been extensively used in many sets and that they have proven to be robust. The dataset allowed measuring the parameters of over 4.5k transients, which are listed in figure 11. Of them, the time of decay to half the amplitude (T50) is shown as an example of statistical comparison of the sets in figure 54.

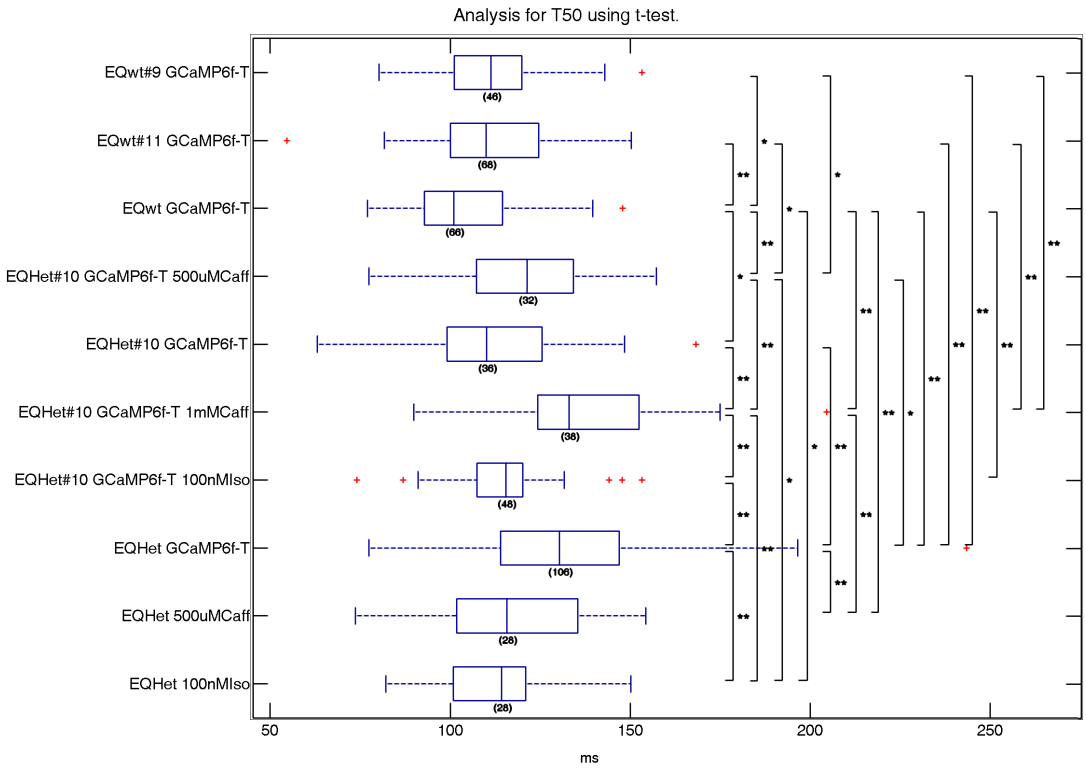

Fig. 54 *Example parameter statistical comparison. In this case T50 is compared in the different datasets (time it takes for the signal to decay to 50 % of its initial amplitude). A t-test was performed to measure the probability of the given distributions coming from distributions of equal mean. Sets linked with one star represent a p-value under 1 %, two stars represent p-values under 0.1 %.*

4.6.3 Spark detection in frame scan

The frame scan spark detector was developed using a manually validated set of experiments in which cells were filmed while presenting spontaneous calcium release events. The manual validation was performed by an expert member of Dr Hove-Madsen's lab, and the validation consisted in manual labelling of the pixel coordinates and temporal frame of each spark within each experiment. A total of eight experiments where used with an accumulated spark count of 621. The experiments were chosen to cover two different recording set-ups (traditional photomultiplier versus hybrid detector using also avalanche photo diode), had been taken in two different seasons of the year (implying different temperature conditions) and had been performed by two different technicians (some parameters are subjective and produce discrepancies amongst experiments; for example the amplification applied to the recording device which can easily saturate the images).

It is necessary to say that this manually validated ground truth was far from perfect due to human error, mainly because the image sequences are very long and the sparks can be very

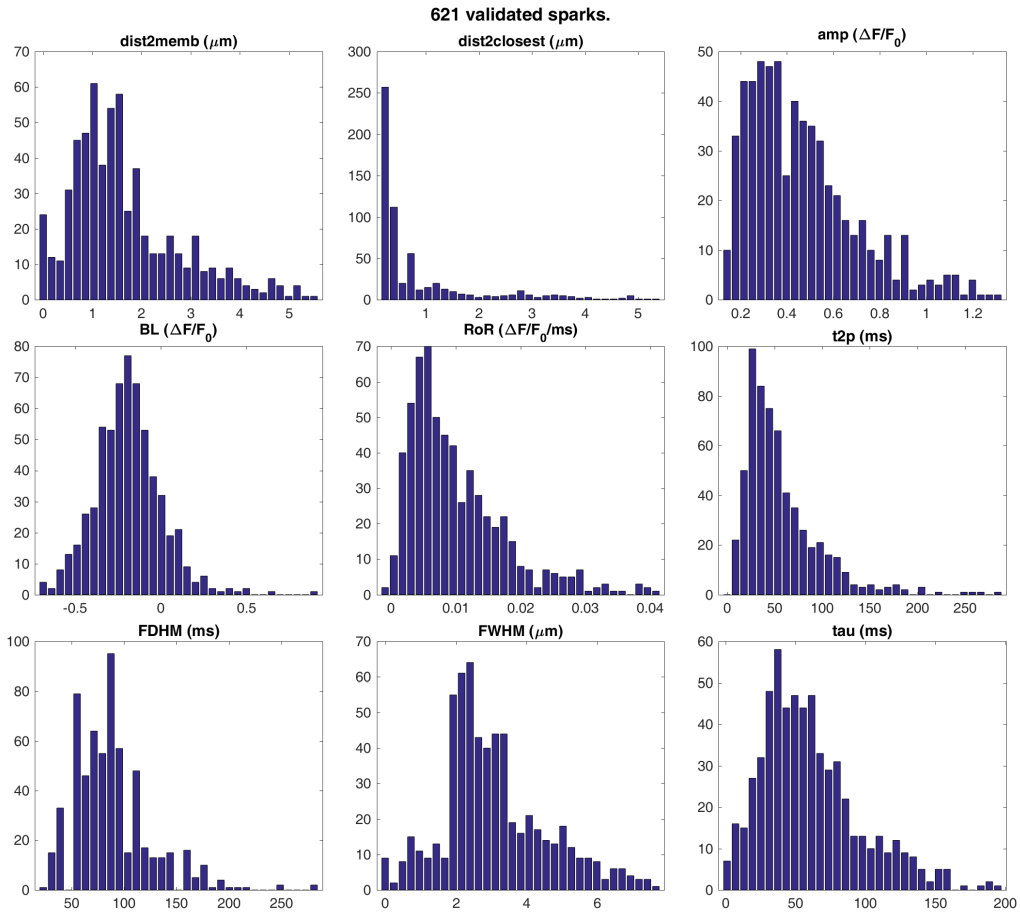

Fig. 55 *Spark features of the validation dataset.*

faint. Altogether, a posterior analysis has given an estimation of the ground truth data set error: less than 1 % of false positives, but between 5 and 25 % of omissions depending on the experiment.

The set of validated experiments (table 12) allowed to tune some of the detector parameters by repeatedly processing the data-set. All parameters where fixed to a particular value except for one that was used as a test parameter by scanning across different possible values and recording the performance in each case. This was done for a series of parameters in order to obtain the combination of filtering parameters that best would fit the validation results. Figure 56 shows an example of the sensitivity and specificity versus two of the filtering parameters and figure 57 presents the receiver operating characteristic (ROC, [186]) analysis obtained. These plots show sensitivity versus one minus specificity for the selected combinations of the parameters (the intersections of the grid in figure 56, top panel). Of all the points obtained, the one that presents a higher *are under the curve* (AUC) is selected as the best performance (AUC in the case of a single point is the area of the quadrilateral under the lines that go from the origin (0,0) to the selected point and then up to the opposing edge at (1,1). ROC analysis

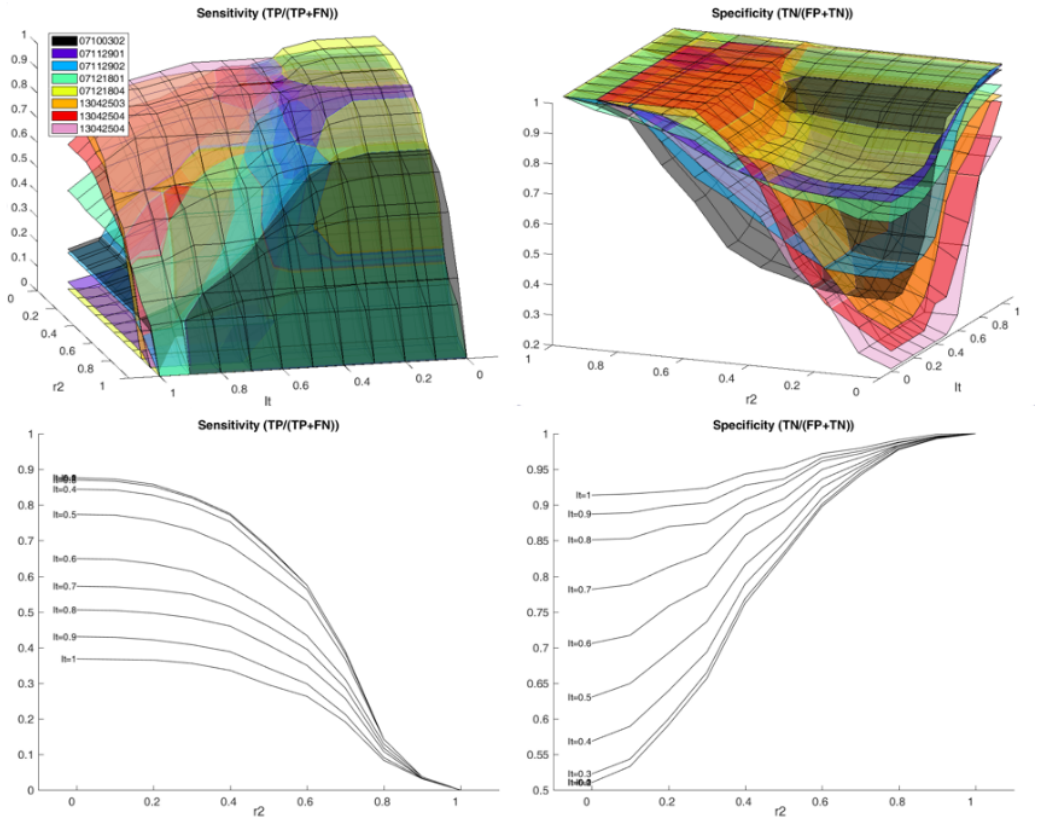

Fig. 56 *Parameter sweeping. The algorithm is tested with the eight experiments forming the dataset; top two panels show a surface for each experiment of sensitivity and specificity, both as a function of intensity threshold (I_t) and of the goodness of fit of the exponential fit to the spark decay (r^2). The mean value of the eight experiments is seen in the bottom panels as a function of r^2 for different I_t values.*

of the eight experiments produces a mean AUC of 0.798 ± 0.047.

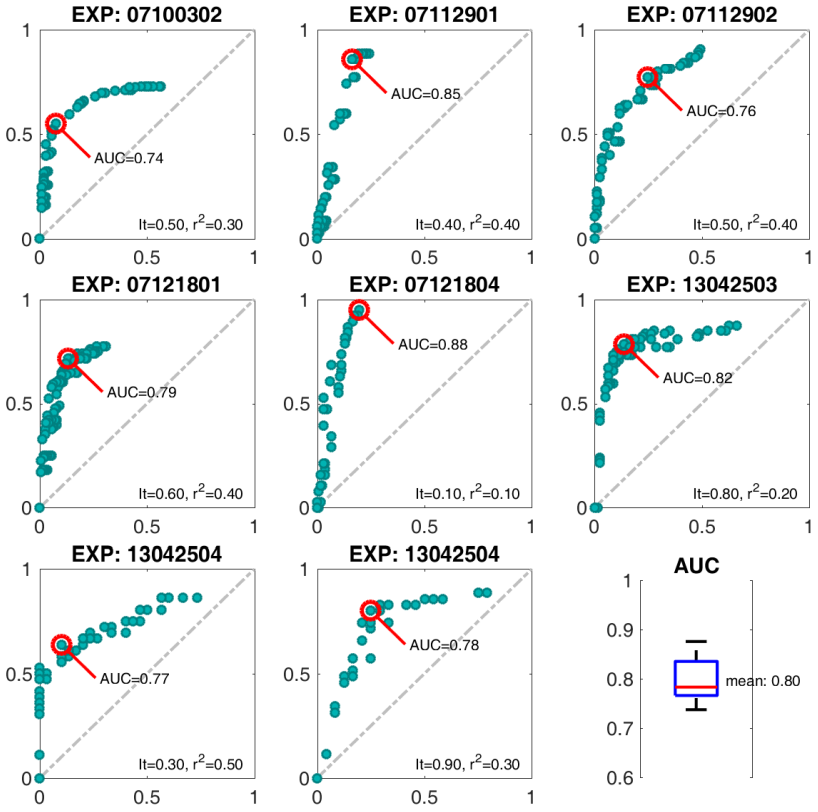

Fig. 57 *Parameter tuning example. Each experiment is run under all possible combinations of the parameters (in this case r^2 and I_t). Each of these combinations gives a point in the sensitivity versus specificity (1-specificity). The one that gives a higher area under he curve (AUC) is chosen as the best combination for the set. It is interesting to see how setting the same parameters changes completely the ROC analysis depending on the experiment. As opposing examples, the top-left panel has a clear saturation at sensitivity around 75 %, whereas the central panel does not reach saturation.*

Experiment Id	Cell area (μm^2)	Duration (s)	# of frames	# of Sparks
07100302	1654	31.43	2125	235
07112901	1309	52.34	4547	35
07112902	1001	31.41	2729	75
07121801	1583	31.35	2729	88
07121804	728	41.90	3640	38
13042503	632	30.35	2610	79
13042504a	738	30.34	2610	36
13042504b	719	30.34	2610	35
Totals	8365	279.46	23600	621

Table 12 *Summary of spark detector validation dataset. All experiments were recorded at a spatial resolution of $0.24 \mu m/pix$)*

4.6.4 Spark-RyR2 co-localisation

This study was a combination of methods for RyR2 detection and spark detection. It was performed both in line scans and frame scans, and the results are already shown in section 4.5. Nevertheless here are a couple of tables with the resulting distribution of distances for the two sets.

sparks lying	at	less than
50 %		$0.10\mu m$
75 %		$0.24\mu m$
90 %		$0.48\mu m$
99 %		$0.88\mu m$

sparks lying	at	less than
50 %		$0.16\mu m$
75 %		$0.25\mu m$
90 %		$0.40\mu m$
99 %		$0.79\mu m$

Table 13 *Co-localisation results, left for line scans (N=1008sparks) and right for frame scans (N=2984).*

These distributions cannot be shown as histograms because of the digitalisation of space, both in the line scans and in the frame scans. In the line scans, the images were taken in four different pixel sizes [0.03,0.04,0.05,0.06]μm. This uneven discretisation means not only that there are only certain possible values of distance measured from a spark to a RyR2 cluster, but also that certain distances cannot happen in certain sets. A histogram with a small distance would present wild peaks due to the number of images with each resolution. The only solution to this would be to use a binning size of the least common multiple of the resolutions ($0.6\mu m$), but that would produce a histogram of just 3 columns from one z-line to the next, with all sparks falling in the first bin.

The frame scans were taken all with the same resolution, but again, being the plane discretised into pixels, only certain distances are possible: with r being the pixel size, objects can only be in the same pixel $d = 0$, one pixel apart $d = r$, one pixel apart in each direction $d = \sqrt{2}r$, two pixels apart $d = 2r$, one pixel apart in one direction and two in the other $d = \sqrt{5}r$, two pixels in each direction $d = 2\sqrt{2}r$ and so on. This would populate the histogram according to the square lattice mesh, rather than giving us real information of the distance distribution, and we would only get a smooth histogram with a bin size that is far to big to capture any information. These are the reasons why the distributions are presented as percentiles in the tables.

5

Cellular scale

5.1 Cell dynamical regimes

Cardiac myocytes normal function is to respond to a neighbouring cell action potential with a calcium transient triggering its own contraction and at the same time triggering the next cell to start the same process. Therefore, a typical stress experiment for a single cell (or a group of cells) is to submit them to an electrical stimulation protocol in which the cell is electrically paced at increasing pacing frequency. The experiment would consist of different sections during which first the pacing frequency would be low for a some time and then we would have periods with pacing at increasing frequency. Typically a section at the start would be included with no pacing to see if spontaneous activity is present and a final section at the end with the original low frequency pacing to see if the cell returns to its original state. See figure 58 for an example.

The calcium concentration is measured as the average fluorescence value of the pixels within a cell at each frame normalised to the cell baseline. This produces a time-dependent signal $s(t) = (f(t) - f_0)/f_0$, where $f(t)$ is the average fluorescence of the cell in each frame t and f_0 the average baseline value (see section 4.3.1 on baseline estimation). In normal conditions this signal consists of a series of transients, each corresponding to an action potential of the cell as a result of external electrical stimulation. However cell response is not always as such.

In order to study cell dynamics of human myocytes, our collaborator Dr Leif Hove-Madsen established a total of six groups into which the dynamics of a cell is classified according to its response to external electrical stimuli at a constant frequency with the aim of developing an automatic classification system. These six dynamical regimes are not unique and even they can slightly overlap each-other, but they are useful to characterise the cell function and

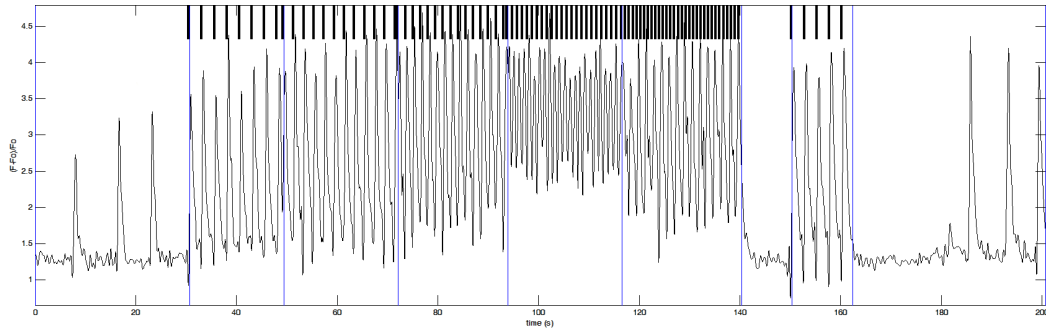

Fig. 58 *An example of a calcium concentration plot of a cardiac myocyte during a stimulation protocol. The pacing times are marked as black ticks at the top of the plot and the blue lines separate periods of different pacing frequency. As can be seen the experiment consists of an initial section lasting 30s without pacing during which the cell produces three spontaneous transients. The following five periods are of increasing frequency: pacing every 2.5s, 2s, 1.5s, 1s, 0.75s. After that there is a 10s period without any pacing followed by another 10s period with pacing at the original rate of 2.5s and finally 40s without any pacing again during which the cell produces another three spontaneous transients.*

some of the behaviours that can lead to pathologies at the tissue level. They are subjective classification set from an experimental point of view of the typical responses one can observe. The six possible regimes are described in the following table:

Regime Name	Description
Regular	The cell responds with a normal transient after each stimuli.
Alternans	The cell responds with a lower intensity every other stimulus.
Phase-lock	Severe case of alternans in which every other stimulus there is no cell response.
Wave	The cell initiates a slow propagating event that lasts over several stimuli.
Irregular	The cell responds in an irregular fashion that does not correlate to the stimulus.
Inactive	The cell does not respond to the stimulus.
Sub-classes of Alternans	
Spatially Discordant	The cell has two distinct spatial regions which activate separately contributing to an alternant behaviour as a whole.
Spatially Concordant	The alternans is not due to spatial differences, rather the whole cell has an alternating pulsation.

Table 14 *Possible behaviour of an excitable cell under periodic electrical stimuli.*

It is relevant to say that the six cases do not exclude each other, for example a small wave and an irregular response are indistinguishable, and that the border between some of them can be a fuzzy region, for example between alternans and phase-lock. Therefore the classification is rather a subjective matter and this is why for the automatic classification a machine learning method was chosen as the best option so that the method would be capable of adapting to the criterion of the person using it. Regarding the two types of alternans, since it is a spatial difference, it cannot be distinguished only with the time signal and so the problem is tackled from another point of view in the next section.

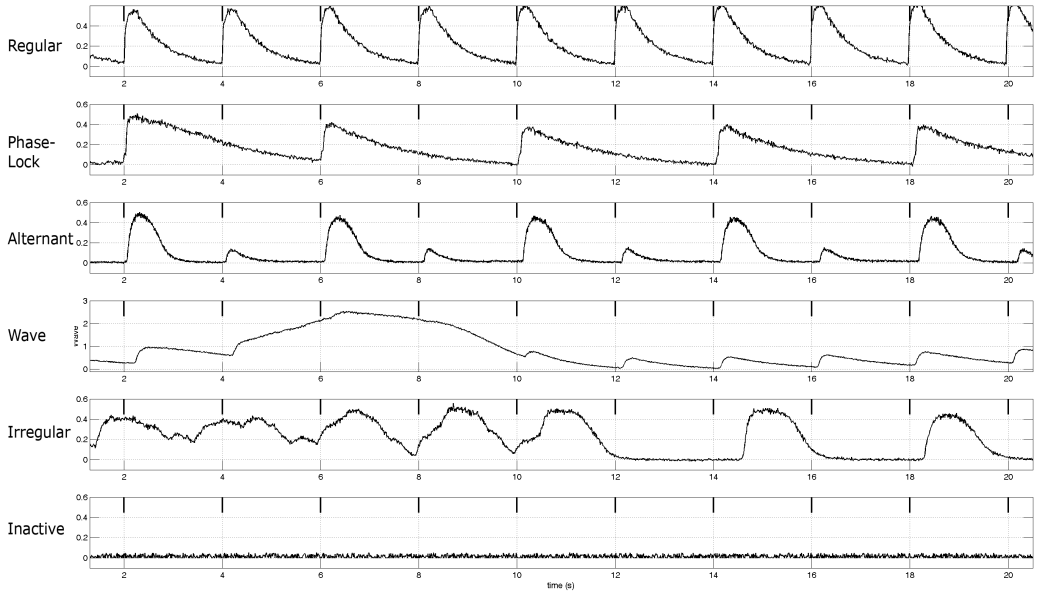

Fig. 59 *Example signal for the six different dynamical regimes for cells stimulated every 2s (black ticks mark stimuli).*

5.2 The training set

In order to prepare a classifier we selected a total of 569 signals from different experiments and three separate experts manually classified them into one of the six regimes. The cases in which none of the three coincided with the verdict were removed leaving a final training set of 512 signals with their class set. Because some of the regimes are not very common the set turned out to be un-balanced: 180 cases of regulars the most common versus only 36 cases of alternans the least common (see table 18 at the end of the chapter for a summary of the training set). This un-balanced number of elements in each of the classes can be a handicap for some classifiers, and for this reason the set is currently being rebuilt with the aim of having 200 instances of each of the 6 classes to build a final training set of 1200 instances.

Choosing a correct set of features to best describe the signals is an open problem and is also a critical step because it can limit the performance of the classifier. A set of features was selected such that at least one feature would best describe one of each of the six possible classes and some extra features were aded for robustness. The features were chosen so that they would emulate the mental process one undergoes when tackling the problem as a human classifier. The feature set is detailed in table 15.

5.3 The classifier

The next step is to choose the appropriate machine learning classifier which is a somewhat tedious task if one has to implement the whole collection of classifiers available in the market.

Chosen feature descriptor	Target Regime
The coefficient of variation([1]) of the time between detected peaks in the signal.	Regular
The mean time between detected peaks in the signal over the pacing period.	Phase-Lock
The number of peaks that alternate (more than 10 % in intensity) over the total number of peaks([2]).	Alternans
The intensity of the highest peak in the continuous wavelet transform([3]).	Waves
The amplitude of the signal (maximum value minus minimum value).	Inactive
The mean of the signal.	Extra
The standard deviation of the signal.	Extra
The standard deviation of the baseline of the signal (measured as described in 4.3.1).	Extra

Table 15 *Selected features used to classify a signal together with the dynamical regime they are targeting. Peak detection was performed by searching for 2D maxima in the continuous wavelet transform with wavelet family gaus2 and scales ranging from corresponding to 0.2s (minimal length of an event) to the length of the signal.*
([1]) *Coefficient of variation is defined as the standard deviation over the mean.*
([2]) *Over the total number of peaks minus two since at least three peaks are needed to establish if alternans is present.*
([3]) *At scales corresponding to more than twice the pacing period (feature is set to zero if no peaks are present).*

The problem is highly non-linear with possible hidden feature interactions and presence of outliers and it was first tackled using a hand-made decision tree [187]. This early stage classifier had only four classes but already implemented the alternans discrimination using PCA (described in 5.4.1). Shortly afterwards, the performance of the classifier was greatly improved with the implementation of the random forest algorithm [188] which seems to outperform many of the newer classification methods. This method has proven to give good results and it has been used continuously for a couple of years and given rise to several publications. After testing several methods, it seems that simple bagging could give better results than the random forest, although this may be due to the fact that the set is unbalanced and the random forest tends to prioritise classes with a low representation in the training set. For this reason a new exploratory analysis will be carried out when the new training set is ready.

Table 16 shows the accuracy of a collection of classifiers applied to the training set. Luckily, new versions of MATLAB have a classification tool (called *classification learner*) that applies a fixed set of classifiers to your data thus allowing a quick approach to the best family of classifiers. The tool rapidly showed that ensemble methods using decision trees were the best option for this problem and so the table was completed with a collection of variations in the algorithms. All algorithms in the table were tested by applying a cross validation of 5 folds, that is dividing the training set into 5 random subsets, training the classifier using 4 of the subsets and then testing the classifier with the remaining subset. The process is repeated 4 more times with the other possible combinations to finally produce an accuracy value as the total ratio correctly classified in the 5 tests. Because the training set is unbalanced and the testing method is random there is a certain drift in the accuracy values every time one repeats the test ($\sim 1\,\%$) so to be strict the values have been rounded to the closest integer percentage.

The table is divided into families. First the Naive Bayes stands on its own with a poor

Classifier Name	Accuracy	Parameter Details
Naive Bayes Classifier	54 %	-
Linear Discriminant	61 %	Regularisation: Diagonal covariance.
Quadratic Discriminant	49 %	Regularisation: Diagonal covariance.
Linear SVM	66 %	Kernel scale: Automatic.
Quadratic SVM	64 %	Kernel scale: Automatic.
Cubic SVM	62 %	Kernel scale: Automatic.
Fine Gaussian SVM	56 %	Kernel scale: 0.75.
Medium Gaussian SVM	65 %	Kernel scale: 3.
Coarse Gaussian SVM	57 %	Kernel scale: 12.
Fine KNN	56 %	Neighbours: 1, Distance metric: Euclidean, Distance weight: Equal.
Medium KNN	63 %	Neighbours: 10, Distance: Euclidean, Distance weight: Equal.
Coarse KNN	53 %	Neighbours: 100, Distance: Euclidean, Distance weight: Equal.
Cosine KNN	61 %	Neighbours: 10, Distance: Cosine, Distance weight: Equal.
Cubic KNN	61 %	Neighbours: 10, Distance: Minkowski (cubic), Distance weight: Equal.
Wighted KNN	64 %	Neighbours: 10, Distance: Euclidean, Distance weight: Squared inverse.
Simple CT	61 %	Number of splits: 4, Criterion: Gini's diversity index.
Medium CT	65 %	Number of splits: 20, Criterion: Gini's diversity index.
Complex CT	60 %	Number of splits: 100, Criterion: Gini's diversity index.
Wild C4.5	61 %	Number of splits: Unlimited, Criterion: Information gain, Pruning: No.
Complex C4.5	63 %	Number of splits: Unlimited, Criterion: Information gain.
Subspace Discriminant	59 %	Number of learners: 30, Subspace dimension: 5.
Subspace KNN	61 %	Number of learners: 30, Subspace dimension: 5.
AdaBoosted Trees	66 %	Number of splits: 20, Trees: 30, Learning rate: 0.1.
RUSBoosted Trees	66 %	Number of splits: 20, Trees: 30, Learning rate: 0.1.
Small RF	64 %	Trees: 30, Criterion: Information gain.
Small Wild RF	62 %	Trees: 30, Criterion: Information gain, Pruning: No.
Medium RF	66 %	Trees: 100, Criterion: Information gain.
Medium Wild RF	67 %	Trees: 100, Criterion: Information gain, Pruning: No.
Large RF	65 %	Trees: 1000, Criterion: Information gain.
Large Wild RF	67 %	Trees: 1000, Criterion: Information gain, Pruning: No.
Small Bagging CT	67 %	Trees: 30, Criterion: Gini's diversity index.
Small Bagging C4.5	68 %	Trees: 30, Criterion: Information gain.
Small Bagging Wild C4.5	68 %	Trees: 30, Criterion: Information gain, Pruning: No.
Medium Bagging C4.5	69 %	Trees: 100, Criterion: Information gain.
Medium Bagging Wild C4.5	69 %	Trees: 100, Criterion: Information gain, Pruning: No.
Large Bagging C4.5	67 %	Trees: 1000, Criterion: Information gain.
Large Bagging Wild C4.5	69 %	Trees: 1000, Criterion: Information gain, Pruning: No.

Table 16 *Comparison of different classifiers where clearly the winners are ensemble methods using decision trees. The training set consists of 512 instances and 8 features with a possible outcome of 6 classes (notice that a random classifier would produce an accuracy of 16.7 %). In each case the training set was cross-validated using 5 random subsets with no overlapping and for this reason there is a dispersion of ~ 1 % in the accuracy and the values have been rounded to the integer.*
Acronym explanation: SVM is support vector machine, KNN is k-nearest neighbours, CT is a classification tree (as in CART), C4.5 is a particular algorithm for CTs and RF is random forest.
Regarding classifier parameters, some have been omitted in the details column for redundancy and are mentioned here: all SVM and KNN were applied setting standardise data to 'true' and all SVM were applied using box constraint level of 1 and multi-class method of 'one-vs-one'. The term 'wild' in decision trees is used to refer to trees without pruning, in all other decision trees the pruning threshold is set to 3 %. Finally, bagging and random forest methods were applied using an outbag of 55 %.

accuracy as expected being a non-independent set of features. The next are the discriminant analysis where the linear performs surprisingly well again considering the highly probable correlation in the features and that they do not necessarily come from normal distributions. The following groups are the support vector machines, the nearest neighbour methods and the single decision trees, and the final group consists of the ensemble methods where the bagging and random forests have been thoroughly explored because they have proved to be the best.

Just to make it clear as there is a little confusion in the literature on these two terms: Bagging is the act of building a series of learners (decision trees in this case) each of which is trained (grown in this case) using a subset (generally random) of the training set. The term random forest is a kind of bagging in which the trees are also built including a certain randomness; at every branch a feature is chosen at random and the value that best separates the remaining classes is chosen (using an established criterion). In the case of the C4.5 trees the feature is not chosen at random, rather the chosen feature is the one that will best separate the remaining classes once the best separating value is found.

The trees can be grown with a limited number of branching points or instead they can be grown until only elements of one class are left on that branch (called a leaf). If they are allowed to grow to this point it is highly probable that the tree is over-fitting the data. In principle the ensemble methods using several trees will overcome this problem by averaging the results, but for the sake of testing a pruning method was applied anyway. Once the tree is grown without limits, the training set used to grow that tree is run through the tree to see how many instances fall down each branch. The pruning method cuts any branch that contains less than a certain percentage of the training set. As can be seen, small forests/bagging ensembles produce slightly better results when the trees have been pruned and on the other hand when the forest/ensemble is big the results are slightly better if the trees are wild (no pruning).

It seems that bagging of C4.5 trees outperforms the random forest, but as mentioned earlier we assume this is an effect of the fact that the set is un-balanced and that may not be the case once the training set is rebuilt.

Regarding the apparent low performance of all classifiers, this is due partly to the fact that it is a subjective problem, with overlapping classes and with a high variability in the data. The relatively small feature set is surely limiting the performance, and also the fact that the signals come from ver different sources (atria and ventricular, isolated single human cells imaged at high resolution, cultured cells imaged at low resolution). As already mentioned we are currently building a new training set, which hopefully will overcome all these imitations and raise the classifier performance.

5.4 Alternans analysis

Alternans in cardiac myocytes has proved to be linked to cardiac pathologies such as ventricular fibrillation and sudden death but the hidden mechanisms that trigger this behaviour are not yet understood nor are the exact consequences it implies. For this reason its study is a hot topic and in fact has been so for the last 10 years or so.

5.4.1 SCA versus SDA using PCA

As mentioned in table 14 there are two types of alternans regarding their spatial origin, namely *spatially concordant alternans* (SCA) and *spatially discordant alternans* (SDA). This section provides a short description on how *principal component analysis* (PCA)[189] can be used to distinguish between these two cases. The method is was presented in [187] and further applied in [190].

PCA is a technique that is usually used to reduce the dimensionality of a dataset. Given a set of observations in an N-dimensional space, the method searches for the linear combination of the N dimensions that gives the direction of maximum variability in the observation set. It will then search for a new direction that is orthogonal to the previous and which also contains the maximum variability and so on until a set of N vectors is obtained. These N vectors provide a new orthogonal basis for the N-dimensional space. Now if we apply a change of basis of the observations to this new set of vectors we are simply projecting the observations to each of the vectors to find their components in each of the N new directions. Because of the fact that this new basis is sorted (by the order in which we obtained it) we can keep only a fraction of the components and ignore the rest to reduce the dimensionality of the set. The number of components that are kept depends on the focus of the problem, it can either be a fixed value, for example we can set it to two if we want to do a graphical representation of the data, or we can set it to the number such that the ratio of reconstruction of the original data is above a certain threshold.

Going back to our problem, if we imagine each pixel in an image of the sequence as an observation and the values it takes over time as each of its components in an N-dimensional space (where N is the number of frames in the film), we can apply PCA to obtain an image for each of the N directions of variability. In each of these images each pixel will show its component in that direction but the significance of these principal images will depend on the contents of the original sequence, that is it will *"show whatever are the most exciting things happening in the film"*. Examples of this are shown in figure 60 where we see the first component capturing extremal positions in a contracting cell, a similar case of a moving cell, a case in which simply the presence versus lack of activity is captured and an example where the discordant activity is captured.

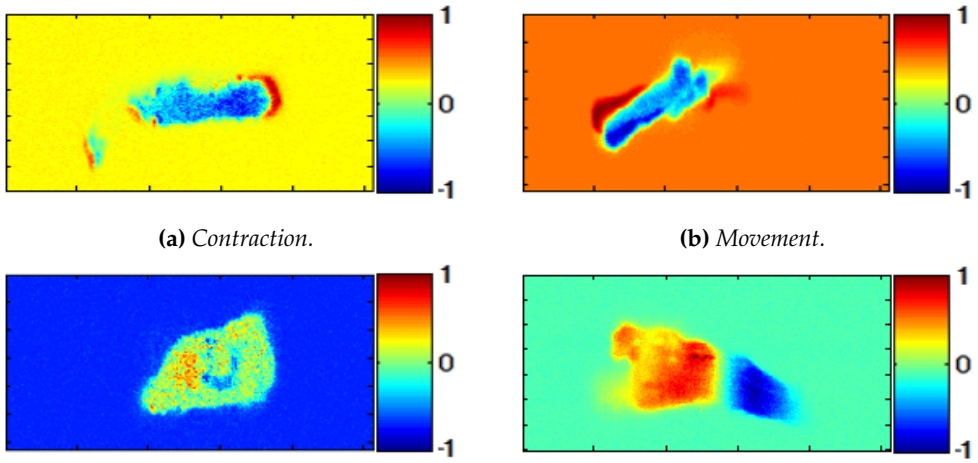

(a) *Contraction.* **(b)** *Movement.*

(c) *Activity versus background.* **(d)** *Spatially discordant alternans.*

Fig. 60 *First component of four different sequences of images. Each pixel displays its component in the direction of maximum variability of the data, with the values normalised to the maximum range. As can be seen, a variety of situations can show up.*

Assuming we have applied a previous classifier that already has established that the signal of the mean fluorescence of a cell is alternating we can test for SDA. We will have to take not only the first but up to the third or fourth component (considering PCA could have first captured movement, contraction and presence of activity). An option is to not include the pixels of the background in the analysis but this can be dangerous if the cell is moving throughout the image sequence. A good practice is to apply a k-means clustering to the resulting components with $k = 3$ allowing for empty clusters. The cluster to which the borders of the image belong is discarded as being considered background. If two clusters are still remaining and the pixels belonging to each are connected in a single region and at the same time are enough to cover a minimal area of the cell, we can assume we have found the alternating areas.

By obtaining the mean fluorescence signal for the surviving regions we can test if they are actually contributing to the total signal as a counter-phase system (example in figure 61). The dynamical regime classifier should classify these signals as either alternant or phase-locking and their lag at maximum cross-correlation should be of exactly the pacing time. If not then we might have captured movement or contraction and so we would need to search further in higher components or else the alternans is spatially concordant.

5.4.2 Characterising alternans

Given a signal presenting alernans there are several ways to characterise it. The severity of the alternans and its duration are two independent matters. This problem was addressed

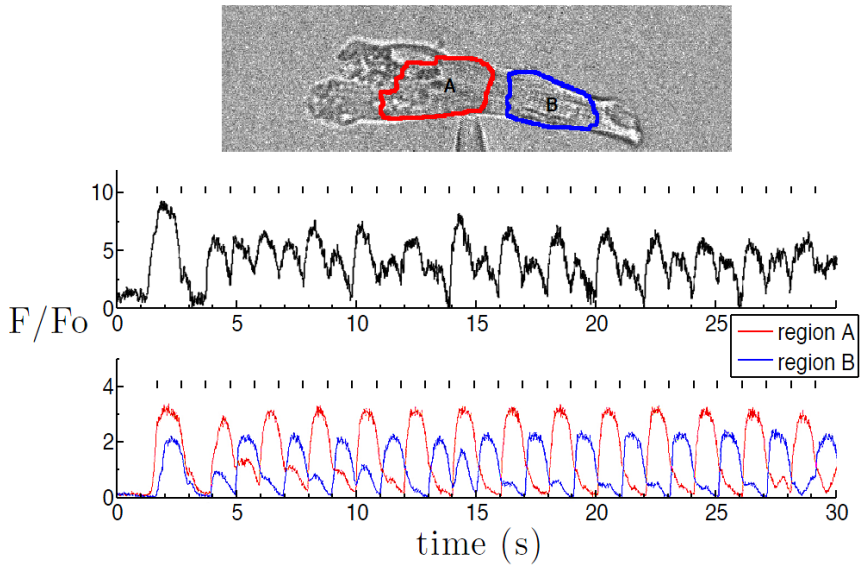

Fig. 61 *The top panel shows two distinct regions found at opposing ends of the first principal component. The middle panel shows the mean fluorescence of the whole cell and the bottom panel shows the mean fluorescence of each of the regions. This example is the same cell as in figure60(d).*

in [191] where alternans was linked to RyR2 function and ventricular tachycardia. It is briefly described in this section.

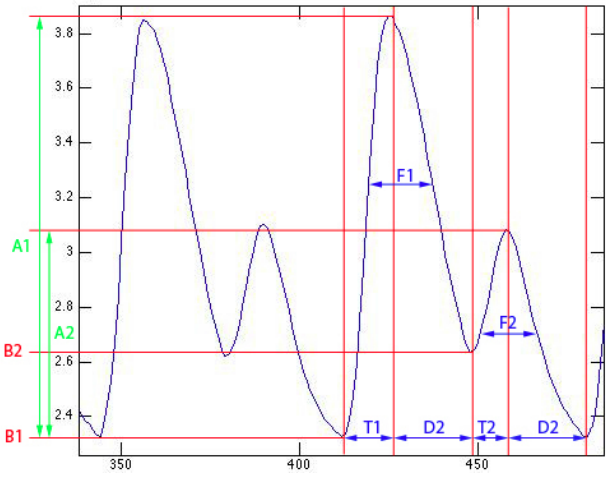

Fig. 62 *Features measured in alternating transients. After classifying peaks into wither big or small, the features are measured and kept in a different pool. A is for amplitude, B is for baseline, T is for time to peak and D is for duration.*

Having already measured the average fluorescence signal of a cell, as always, continuous wavelet transform is a reliable way to clean of noise and at the same time enhance peaks in the signal. Zero-crossing of the derivative can be used to detect peaks and valleys in the signal (see section 1.5.6). Knowing the pacing frequency we can cross-correlate the signal with a train of deltas to separate stimulated from spontaneous peaks. Amplitude for each peak can be measured as the difference in height between the peak and the previous valley (relative amplitude). For a particular couple of consecutive stimulated peaks, alternans ratio can be defined as the difference in amplitude between the two over the amplitude of the largest of the two (taking the absolute value).

The criterion we used for establishing the presence of alternans in one cell was that at least six consecutive peaks should all present an alternans ratio above 0.05. The mean value of the alternans ratio for every couple of peaks provided the alternans ratio of the whole signal, and the alternans duration was defined as the cumulative elapsed time of alternans periods over the total duration of the line-scan.

5.5 Cell contraction

Somewhat to my displeasure, measuring cell contraction has stayed slightly at the edge of the scope of this thesis when it is one of the most important aspects of cardiac myocytes. The main reason behind this is that it is not a fluorescence imaging problem (also the fact that time is finite). Nevertheless, attempts have been made by applying image processing techniques to the visible channel image of the cell because it is computationally speaking cheaper than other tecniques [192, 193].

Characterising cell contraction requires a method that will produce a vector field of displacement at each position of the cell. A very well-known set of methods that can produce this displacement field are algorithms that attempt to measure optical flow [194, 195]. For big images and high frame rates, these methods rapidly become slow whereas image correlation based methods can be tuned easily by reducing the window of correlation. Image correlation is an extension of the regular signal correlation, in which an image template is multiplied at each position of a given image to obtain a sort of match score in the image. It can be used as a tracking technique; one uses as a template a small region of an image in one frame of a sequence and then correlates it to a bigger region around it in the following frame. If the two are exactly the same a maximal value is obtained in the centre of the correlation. If there is a displacement of the bright objects the maxima appears away from the centre. This provides a mean displacement vector (proportional to the intensity) of the objects in the image. The procedure is applied frame after frame to obtain a vector field of displacement in each frame. The technique is simple to apply but has the drawback that it is a polarised method, in the sense that it will match objects that have high intensity and therefore, for example, black objects in a grayscale image that have displaced will show a maximal correlation in the previous position and not the current.

The cell was first segmented and then divided into regions. For each region the image correlation method was applied throughout the film to obtain a position signal as a function of time with respect to resting state (figure 63). The position signal of the regions corresponding to one end of the cell plus the position signal of the regions at the opposite end of the cell (they have opposing signs) produced a total length variation signal which measures contraction in an absolute sense, and it was then divided into the length of the cell to give a relative contraction measure. Figure 64 shows the contraction signal in red together with the calcium fluorescence signal (of an alternating cell), where one can see how when calcium concentration raises, shortly after contraction starts. Also it is seen that a small increase in calcium is not enough to trigger contraction.

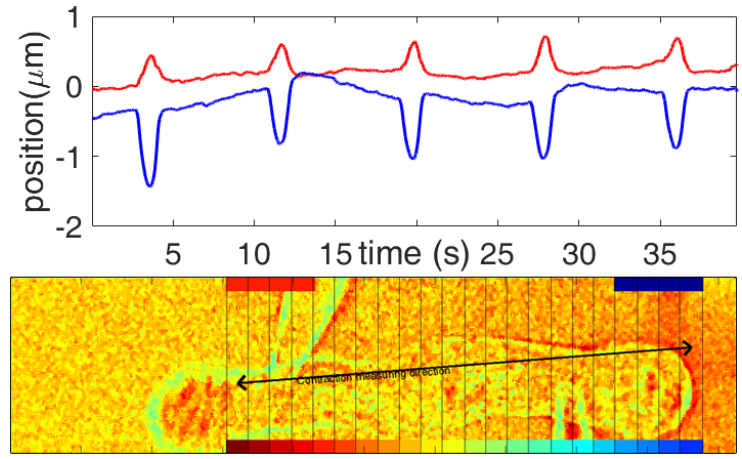

Fig. 63 *The mean displacement as a function of time of the opposing regions of the cell (the plots correspond to the mean of the values measured in the four leftmost and rightmost columns). The the cell is slightly out of focus in the left area and that is why that zone was excluded from the cell segmentation.*

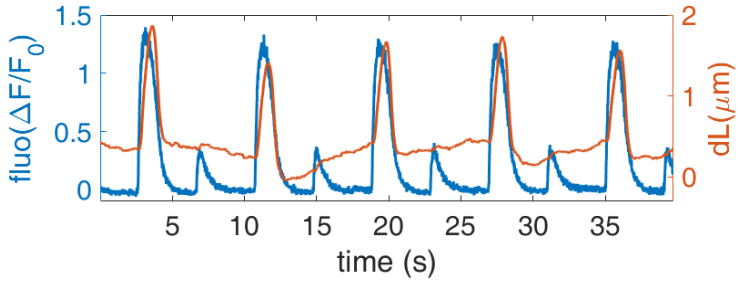

Fig. 64 *The sum of the two signals in the previous figure appears here in red together with the cell fluorescence in blue. The cell is alternating and small responses do not trigger contraction.*

5.6 Connecting the scales

In the previous chapter we have seen a description of the typical calcium events that appear in calcium imaging together with a series of methods to detect them. Also the main features that characterise these events have been explained and methods on how to measure these features are detailed. Now, in this chapter, when imaging whole cells developing transients as a whole, this set of features can still be applied to describe the whole cell calcium signal that are obtained.

Indeed, these features constitute the basis of the features that have then been extracted to build the cell response classifier, which has allowed to characterise the response of a cell to electrical stimulation. This characterisation of the cell behaviour under stimulation can be extended to the cell response is due to neighbouring cells activation instead of artificial stimulation of an isolated cell. Therefore it will still be applicable in the following chapter when analysing cultures of cells and the propagation of a calcium front across the culture. Applying this dynamical regime classifier to each of the cells in a culture will allow to describe the propagation patterns from the point of view of an individual cell.

5.7 Results

5.7.1 Alternans analysis

In this study the aim was to develop a way to characterise alternans in intact heart line scans while under electrical stimulation. The line scans taken across the tissue, covering several cells in the direction in which the cells are shorter (see figure 10 for an example). During stimulation the cells would respond to each stimulus by producing a transient which in the line scan appeared as a sudden rise in calcium concentration. The developed algorithm had to extract the mean fluorescence signal of each cell, detect each of the transients, use the temporal location of the transients to determine the exact pacing frequency (an estimate was known but with a precision around 10 times the sampling time) and measure the properties of each transient (section 5.4.2).

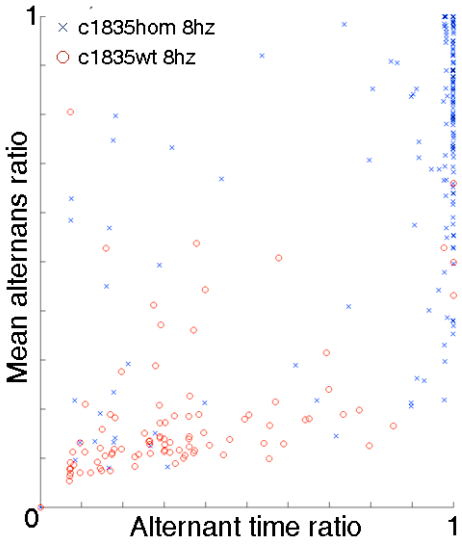

Fig. 65 *Characterising alternans. Scatter plot of the two alternans measures used in [191]. In this particular example, a total of 328 cells are depicted, each having been paced at 8Hz during 10s. Half of the cells (blue crosses) had a mutation that suppressed the function of RyR2 which clearly triggered the appearance of alternans.*

The experiments were performed under different drug conditions and many stimulation frequencies, which in total provided us 653 line scan images. In these images, over half a million transients were detected and classified either into spontaneous, non-alternating or if alternating into big or small. These numbers can be seen in table 17. All of these transients had their properties measured but this is not shown here.

Characterisation of alternans was performed using two parameters. The first the alternans ratio, that is the mean difference of intensities between consecutive peaks normalised to the larger of the two (equal to 0 if no alternans is present and equal to 1 in the extreme case of phase-locking). A cell was considered to present alternans if at leas six consecutive peaks mantained a ratio above 5 %.

The other parameter used was duration of alternating periods over the duration of the experiment. These two parameters proved to adequately quantify the presence of alternans and the results of these methods were used in a study to prove the relation between RyR2 function and the appearance of alternans at the cellular level and problems at the organ level [191]. An example of a scatter plot of these parameters can be seen in figure 65.

Set ID	freqs	# of linescans	cells	% Cells alternating	# of transients	% Transients alternating	% Transients spont	Total scan dist(μm)	Total scan time(s)
EQWT 1.8mMCa	7	91	857	56.5	78944	50.7	0.0	17900	949
EQHET 1.8mMCa	7	103	988	83.8	84531	86.7	0.6	20190	1074
Bef 0.5mMCaff	8	14	121	89.3	10933	98.3	0.0	2380	136
Aft 0.5mMCaff	5	14	132	65.9	11268	57.1	6.4	2660	147
Bef 1mM Caff	8	66	548	84.7	47614	88.7	0.0	11040	624
Aft 1mMCaff	7	45	371	46.9	32010	46.6	2.3	7790	456
Bef 3uMCarv	6	15	133	75.2	9172	79.5	0.0	2490	136
Aft 3uMCarv	4	10	109	69.7	7538	31.2	25.3	1910	104
Bef 1uMEGTAAM	7	42	475	89.3	40317	88.2	1.8	8360	437
Aft 1uMEGTAAM	6	72	768	87.6	62843	86.3	2.0	13600	755
Bef 2uM EGTAAM	7	45	522	87.7	44163	89.3	0.9	8880	469
Aft 2uMEGTAAM	6	75	795	83.4	67221	84.2	0.2	13950	786
Bef 4uMEGTAAM	6	24	270	90.7	23874	95.2	0.0	4600	251
Aft 4uMEGTAAM	6	37	388	88.4	34461	71.4	2.3	6460	389

Table 17 *The alternans dataset. The dataset is divided into 14 conditions, each of them has several stimulation frequencies (*freqs*). Other abbreviations in the table are* spont*, standing for spontaneous, and* dist*, standing for distance.*

5.7.2 Cell response

This section needs not to be extended since the procedure of building the training set, choosing the classifier and testing the performance is already described in sections 5.2 and 5.3. Nevertheless, here is a summary of the original 512 signals chosen to form the training set (table 18). Also, another table shows a summary of the properties of the random forest that was built to solve the classification problem (table 19). Regarding the characterisation of decision trees, the number of branches refers to the total number of branching points, and so it gives an idea of the complexity of the tree. The tree height is, of all the paths from the root of the tree to all the leafs, the path that has the most branching points, this number of branching points determines the height. It is the maximum number of branching points one needs to go trough to get to a leaf, therefore providing also information on the computational load of the decision tree.

Dynamical regime	Number	Percentage
Regular	152	29.7 %
Phase-lock	46	9.0 %
Alternant	37	7.2 %
Wave	84	16.4 %
Irregular	135	26.4 %
Inactive	58	11.3 %
Total	512	100 %

Table 18 *Summary of the classifier training set. The 512 signals are the signals that, after the manual classification, at least two of the three experts coincided in the veredict.*

Random Forest	characteristics
Total number of trees	1000
Average # of branches	148 ± 10
Average tree height	15 ± 1
Largest tree:	181 branches
	24 height
Smallest tree:	125 branches
	12 height

Table 19 *Random forest description. As can be seen the trees were not pruned (large heights and number of branches). Clearly a single tree will over fit the data, but having 1000 trees allows to overcome this.*

6

Culture and tissue scale

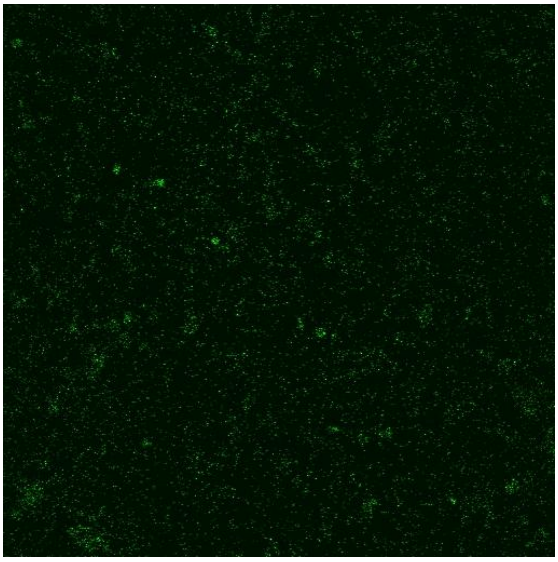

Fig. 66 *Sequence of fluorescence images of a culture under electrical stimuli every 2.5s. Planar fronts can be seen travelling with a slight anti-clockwise drift after which a couple of clear spiral fronts can be seen. Towards the end flat fronts are seen again but with a clockwise drift. The field of view is approximately 1mm by 1mm with around 550 cells, and the total film duration is around 30s.*

Cardiac tissue is a complex network of myocytes connected to each other that allows the propagation of the action potential that triggers the contraction of each individual cell. The study of patterns of the electrical propagation of this action potential in cardiac tissue is very important to understand how pathologies at cellular level can affect myocardium function. Studies can be carried out using ex-vivo intact heart (removing a whole heart and keeping it pumping alive while performing calcium imaging) or for simplicity of equipment and to avoid artifacts due to movement among others, tissue can be emulated by means of a cell culture (growing cells on a flat surface). Cultures also allow researchers to manipulate the environment for instance by injecting drugs into individual cells or adding mutated cells to regions of the culture or simply by erasing or burning an area to affect electric conductivity.

The stimulation protocols described in the previous chapter (section 5.1, figure 58) in which a single cell is being electrically stimulated to test its response and correct function can now be equally applied to a network of cells. A pair of electrodes are placed at one end and an electrical pulse triggers an action potential that should be transmitted throughout the cells, if the cells have been loaded with the appropriate fluorophore we will be able to record brightness proportional to calcium concentration.

The analysis of calcium imaging at the tissue/organ scale requires segmentation of individual cells, measuring of propagation of calcium fronts and characterising this propagation in order to link cellular anomalies to tissue propagation dynamics. The following sections describe these steps.

6.1 Culture segmentation

When imaging fluorescence from an excitable cell under stimulation or exhibiting spontaneous activity, the light originates inside of the cell and the transition to lack of light outside the cell towards the edge of the cell is rather gentle or smooth. This is because the cells are sitting on a flat surface (the culture dish) and tend to be thicker around the centre where the nucleus or the bulk of the contractile machinery may be. As the focal-plane will always be slightly above the culture dish, each cell has a larger portion within the focal-plane around the centre than around the edges and therefore we register more light from the centre than from the edges.

Moreover, each individual image in a sequence tends to be very noisy and it is not until one sees the whole sequence at fast frame rate that one can start to notice the presence of individual cells. See figure .

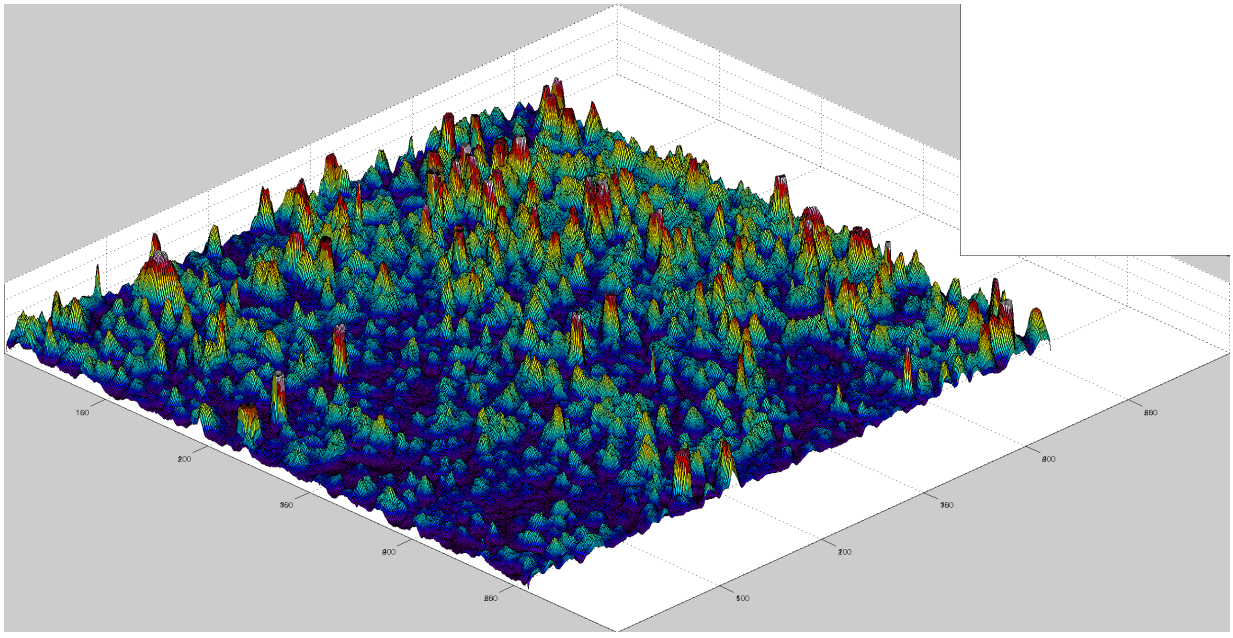

Fig. 67 *Animation of the segmentation watershed method of the variability image.*

The usual approach when a single image does not contain enough information to reveal sharp edges is to take the mean of several images. As these are excitable cells which throughout the sequence have an oscillatory behaviour, a much better approach is to use this

variance as an indicator of the positions where a cell is (at least an active cell). The approach was originally developed for cardiac myocytes [187] but can be applied to any excitable cell. That produce an image where at each pixel we depict the variance or standard deviation of the same pixel throughout the image sequence. This provides a clear image where individual cells are revealed. We call this image the variability image because it enhances the areas that take a larger distribution of values in the image sequence (see figure 68a for an example).

This variability image is the input to the segmentation method. It is a fairly smooth image, with no clear edges and all the cells are more or less the same size. Furthermore each cell has a single local brightness maximum around the centre and then decreases in brightness towards the edges. This clearly calls for a watershed method as mentioned in section 1.5.4, where each cell will be a separate watershed (if the image is inverted, the cells are local brightness minima which if we were to pour water on would form independent puddles). We do not want to constrain the watersheds by absolute or local intensity because some cells have greater variability than others and therefore appear brighter, but we do want to limit the size of the watersheds to avoid several cells being joined and demand for a minimal size to avoid spots of noise being tagged as an independent cell.

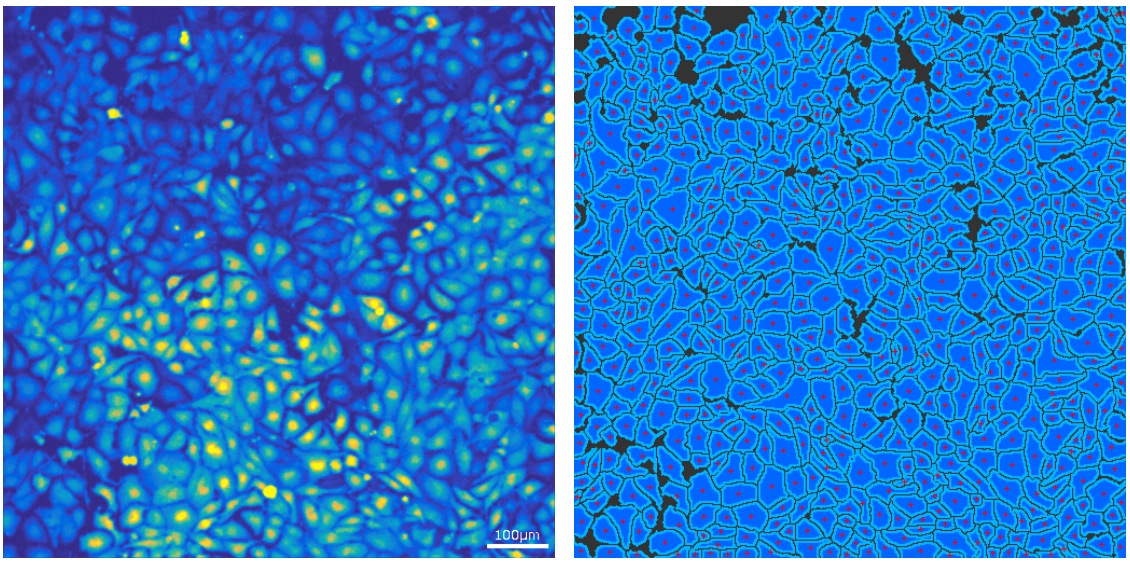

(a) *Variablilty image.* **(b)** *Algorithm result.*

Fig. 68 *An image sequence of a culture is preprocessed into a single image (a) and the cell segmentation algorithm is applied (b) to detect a total of 651 cells in this case. For clarity each cell has had its centre marked with a red cross and its limits marked with a brighter shade of blue.*

An algorithm was developed to accomplish this using a minimal and maximal size as constraints to the watershed. This segmentation method was developed reversed with respect to the usual watershed approach: we can imagine the variability image as a section of terrain where each cell is a hill and then we cover the whole terrain with water making each cell an underwater hill. We then lower the sea level and at each step we measure the sizes of the

islands that appear. The islands will tend to grow in size and merge with each other as the sea level gets lower and lower. The method records the limits of an island only when it is between the expected sizes. Figure 67 depicts this process.

Apart from the minimal and maximal sizes of the desired objects the other parameter that defines the precision of the method is the step size, that is the amount we lower the water level at each iteration. This parameter greatly affects the speed of the algorithm, a big step height vastly reduces the computation time at the cost of producing worse segmentation results because from one step to another an object can go from being smaller than the minimal size to being bigger than the maximal size. This is why not being too restrictive with the sizes allows for a bigger step size and faster computation speed. A balanced value to which it has been fixed is the range of values in the image divided by 100, meaning that we do 100 steps. It has successfully been used like this in [190, 196–198]. As an example, figure 68b shows the results of the method being applied with 1000 steps and producing 651 distinct cells. The method applied to the same image but with 100 steps produced 650 distinct cells, so we gained a ridiculous 0.15 % accuracy at the ludicrous computational cost of ∼10 fold.

6.2 Front propagation

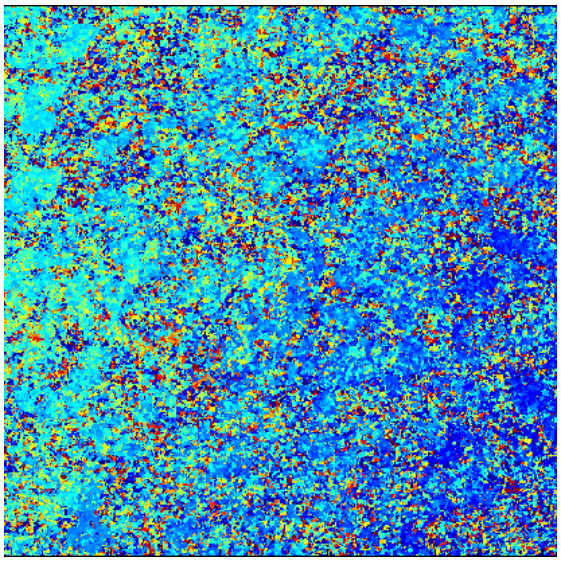

Fig. 69 *Time of front arrival at every pixel of the field of view. The image is very noisy, especially in the regions inbetween cells and for this reason it is far better to only use information from pixels that belong to a cell than to simply smooth out this image by means of any sort of filter.*

Measuring front propagation may seem a simple problem with a straight-forward solution but this is not as so. The process is, given a sequence of images containing several fronts, first to cut the sequence into groups of frames containing a single front, then to estimate the time of arrival of the front at each pixel, with which it should be possibe to construct a propagation map where the colour codes the time of arrival of the front. This colour-coded map is what is known as an isochronal map.

There are several problems to overcome before obtaining the isochronal maps though. Dividing the sequence into frames belonging to each front can be very hard (or impossible) if the fronts overlap eachother or if two or more of them appear at the same time in the field of view. Besides, taking the time fluorescence of a single pixel can produce noisy signals which can then hinder the estimation of time of arrival of a front (see figure Figure 69).

If the fronts are spaced out in time the simplest approach is to take the mean fluorescence value of the whole field at each frame and use the resulting time dependant signal to establish the presence of a front by simply finding local minima in the signal (see section 1.5.6). If the activity is high, turbulent-like so the fronts occur simultaneously and therefore cannot be separated in the time signal, a possible solution is to use a tracking technique as described in section 4.4. The third option is not to cut the image sequence into distinct stretches until the times of arrival have been found and then attempt to relate them to one another.

Regarding the noisy signals, as always the first solution is to take the mean value of a region around every pixel to produce the time signal but an approach that is a little more tailor-made to the problem is to consider the previously segmented cells in the culture. As mentioned in the previous section, each cell will tend to emit more light from its central regions and therefore the areas close to cell edges (or where no cell is present) will produce noisy and almost random signals.

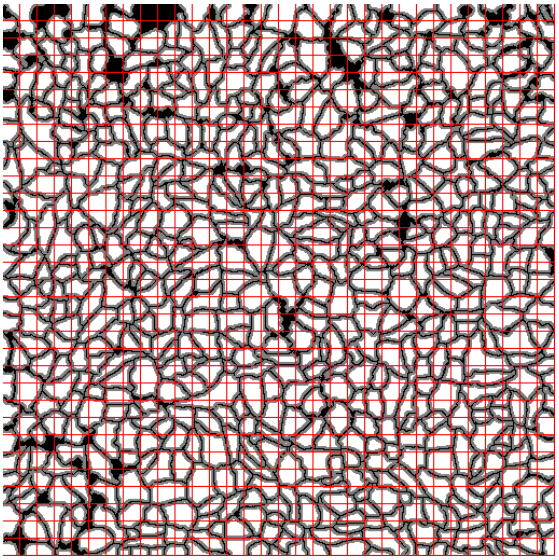

Fig. 70 *Example of the same culture as in figure 68 where the cells have been eroded by 5μm (white areas), the mean distance to nearest neighbour is of 29μm and so in the field of 912×912μm has been divided into a 32×32 mesh. The time signal of each square will now be at each sample the mean of the pixels belonging to a white area in the square.*

Taking this into account, a method was developed in which the temporal signal from which the time of arrival was measured would use only fluorescence values belonging to areas close to the centres of the cells. This method was published as part of a more extensive culture dynamics work [199] and also used as a tool in [200, 198]. In this method, the field of view was divided into a square mesh with a lattice constant of the mean nearest neighbour distance of all cells. Then the temporal signal was constructed based on this mesh; for a particular square of the mesh, at each frame the fluorescence of the square was taken as the mean value of only the pixels belonging to the central regions of a cell. The central region can easily be found by applying an erode filter to the segmented cell image (using a template of radius around 5μm would produce an image with only the white areas in figure 70). This process allowed us to obtain a smooth time-dependent signals where the arrival of the front was made obvious by the presence of a sudden upstroke in the signal followed by a slow decay. Detecting this upstroke was then very straightforward by searching for the highest maximal point in the derivative of the signal. For each front an image was constructed where the pixels belonging to each square of the mesh were assigned the number of the frame of arrival of the front at that square. Finally a smoothing step was applied by using a low-pass

Fourier transform as described in section 1.5.5 and figure 22. The resulting image was then rounded to the closer integer (the procedure is applied on a frame basis) and a final step of eroding each region by one pixel to daw lines in between regions. Multiplying the whole image by the frame rate in milliseconds produced some nice isochronal maps as shown in figure 71.

6.3 Front dynamics

The purpose of this section is to characterise propagation dynamics and at the same time to classify a front amongst a set of predefined cases. These cases or situations were set with the aim of establishing a connection between the inividual cellular calcium dynamics (discussed in section 5.1 and summarized in table 59 and figure 14) and the properties of the propagating fronts. For this reason the situations are rather general, only four different cases have been established, but one could always subdivide into further subgroups. These four basic situations are regular flat fronts, flat reentry fronts, spiral fronts and a fourth left-over group for anything that does not fit in the previous three. The spiral case is probably the most interesting since this kind of dynamics is the cause of pathologies such as cardiac fibrillation. The origin of this is a region in the network that is blocking the transmission of a front which causes the front to bend around this area. This can then cause a permanent activity around the area and the breaking of future fronts arriving at the region. The reentry case is usually when a spiral has taken place out of the field of view and we see a front that, instead of coming from the electrodes, is coming from an unexpected direction. This tells us there is some connectivity problem outside of the field of view, so the reentry can be seen as a sign of unhealthy culture but at the same time as a sign of a healthy field of view (in the part of the culture we are imaging the propagation is flat). The four situations are summarised in the following table (table 20) and an example of each is shown in figure 71.

Front Name	Description
Regular	A flat front that is originated at the field electrodes.
Reentry	A flat front that is originated elsewhere.
Spiral	A front with constant angular velocity different to zero.
Other	Either non-constant angular velocity or a combination of various fronts.

Table 20 *Summary of possible front dynamics. The second case is relevant because as the field of view is always smaller than the whole culture and for this reason a reentry usually indicates a spiral or a spontaneous front has occurred outside the field-of-view. The last case is more of an escape case for the situations in which the fronts have either not been properly separated or a strange unexpected situation appears.*

The main feature that characterises these situations is angular velocity, allowing us to distinguish the spiral case, and the other feature would be linear velocity, its direction

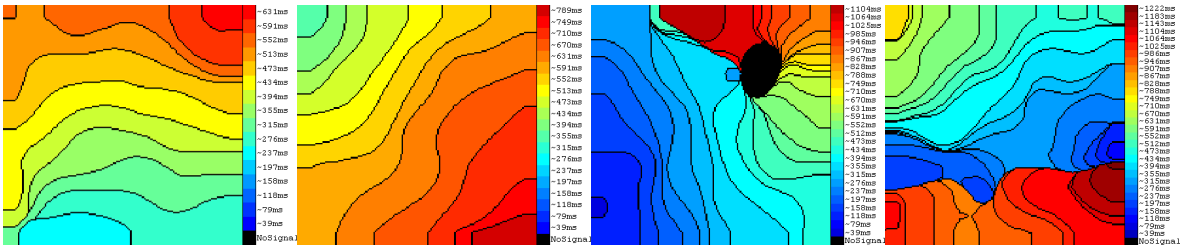

(a) *Regular flat front.* **(b)** *Reentry.* **(c)** *Spiral front.* **(d)** *Other situation.*

Fig. 71 *Sample isochronal maps displaying different propagation dynamics. Each successive colour indicates the front arrived at a particular frame and so that is associated with a time of arrival.*

distinguishing between reentry and regular flat. We will now focus on how to measure these magnitudes.

6.3.1 Linear velocity

Remembering the way in which the isochronal maps were constructed, we have a frame-wise approach, meaning that each of the distinct coloured bands in the map indicate that the front arrived during the recording of that frame.

Consider a particular coloured band in the map. All the pixels in the band have in common that the front arrived at the same frame in the sequence, and the region is limited by two lines that are the borders with the previous and next region. We will call these two border lines isochrones; one is the previous isochrone and the other is the next isochrone.

Now, taking the previous isochrone, we can draw a perpendicular straight line from it which will then intersect the next isochrone at one point. This line going from one isochrone to the next (perpendicular to the first) we will define as a velocity vector. If we draw many of these starting from all along the previous isochrone we will obtain a collection of velocity vectors all associated with the region between the two isochrones. The mean of all these vectors can be set as the velocity vector for that frame.

The modulus of this velocity vector has been constructed as the mean width of the region in the map; it is the distance covered by the front during one frame, so dividing it by the time between frames we obtain proper speed units.

There are some considerations to be taken into account. One is that when drawing the perpendicular lines to the isochrone we need to fix a distance from one to the next (the spatial resolution of the velocity measure). We can do this pixel-wise at high computational cost or relax it a little. Another consideration is that sometimes the velocity vector will not be measurable, it may not intersect the next isochrone either because it hits the limits of the field of view or the next region has zero width, or this region has zero width. All these cases should be ignored otherwise we might obtain strange things like infinite velocity vectors. The

last consideration is that to be able to measure the velocity vector we need the two borders of a region in the map, so if the front takes N frames to cover the field of view, we are only able to measure the velocity for $N-2$ fronts, that is from the second to the last but one.

6.3.2 Angular velocity

Having previously measured the velocity vector at each frame, obtaining the angle of propagation is now very straightforward. We only need to set a zero angle direction, in our case having the electrodes beneath the field-of-view, we set the vertical axis as the zero direction and the angles were measured anticlockwise, so a front travelling normally towards the top would have an angle of $0°$ and a front travelling towards the left would have a direction of $90°$. Each velocity vector has its angle measured from this direction to obtain an angle of propagation. As we only need the angle and not the modulus of the velocity vector, we can go back and measure it for the first and last isochrone in the map, to obtain an angle of propagation for each of the isochrones in the isochronal map. This means that if the field of view was covered during N frames we would now be able to measure the angle of propagation for the $N-1$ isochrones present in the map.

The angular velocity is simply the derivative of the sequence of measured angles so by subtracting the sequence from one another and dividing by the time between frames we are left with a sequence of N-2 angular velocities.

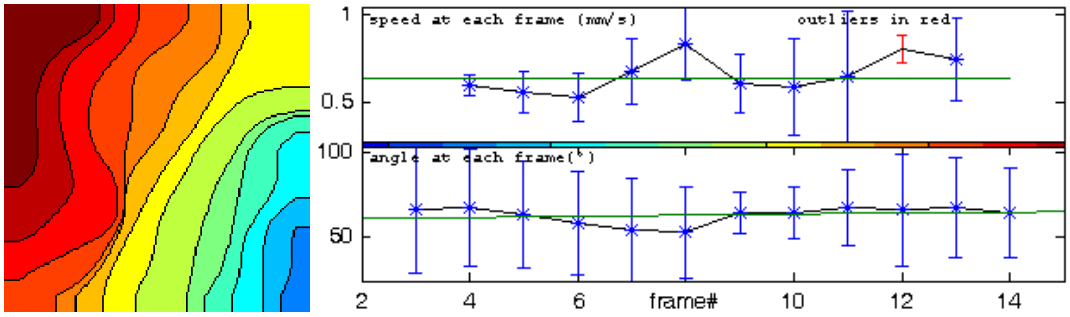

Fig. 72 *Linear and angular velocities measured for a planar front.*
Speed is the mean of the top graph $s = 0{,}6 \pm 0{,}1(mm/s)$, angle of propagation is the mean of the bottom graph $\alpha = 63 \pm 5(°)$ and angular velocity is the mean slope of the bottom graph $\omega = 8 \pm 25(°/s)$. This dispersion including zero in angular velocity is proof that the front is planar.

6.3.3 Classifying the dynamics

The procedure here should be equal to that in the previous chapter when attempting to classify cell dynamics (section 5.1); constructing a training set, manually labelling it, establishing a collection of features, testing several classifiers and choosing one from the best. Unfortunately, as we are still pending rebuilding the cell dynamics training set we have not

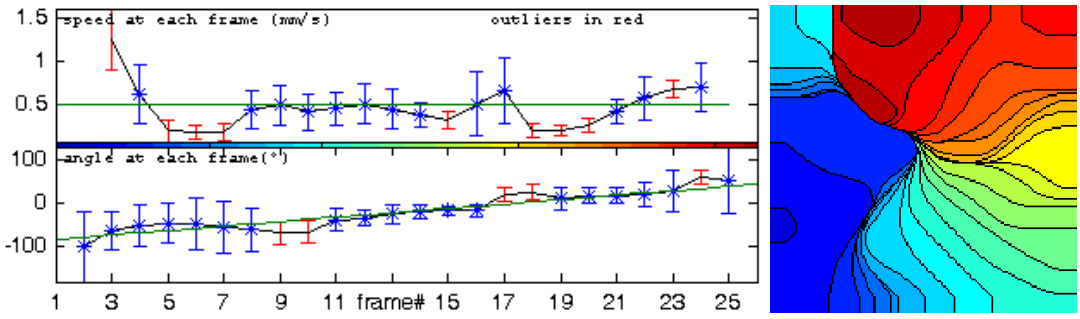

Fig. 73 *Linear and angular velocities measured for a spiral front. In this case $s = 0{,}5 \pm 0{,}2(mm/s)$, $\alpha = -23 \pm 39(^\circ)$ and $\omega = 129 \pm 18(^\circ/s)$. Now, the enormous dispersion in α prooves the front is not planar.*

managed to build and label the culture propagation training set. For this reason, a temporary solution was to build a decision tree with manually chosen threshold values.

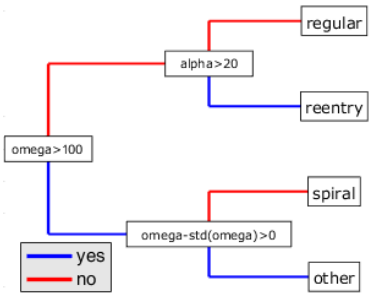

Fig. 74 *Front classification decision tree.*

The custom threshold values were: angular speeds below $100^\circ/s$ implied a flat front, (flattish that is, as opposed to spiral). Then for the flat fronts, mean angles above 20° were assigned as reentries and below that as regular fronts. For the case of angular speed above $100^\circ/s$ if the standard deviation was lower than the actual value of the angular speed the front was classified as spiral, otherwise it was classified as the fourth case.

6.4 Connecting the scales

The first approach to connect the cellular scale with the tissue/culture scale is to see if the number of cells behaving in a particular dynamical regime grow or decrease when pacing frequency gets higher. A statistical test that allows this is the McNemar test [201] for significance of changes, which allows testing hypotheses such as: *When the pacing frequency goes from 1Hz to 1.3Hz the number of irregular cell responses is significantly larger*. This test can be applied between every two pacing sections in the protocol and for every dynamical regime.

The connection of the cell dynamics with front dynamics depends more on cell size than cell number. Rather than cell size it is the sum of sizes of cells behaving in a particular way that contributes to the propagation patterns. This is why a good measure is the area ratio covered by cells behaving in a particular way over the total area.

As previously mentioned, the reentry fronts are a sign that the culture is not propagating properly outside the field of view, but, in terms of propagation within the field-of-view the behaviour is flat and therefore the connectivity is healthy. This is why a good measure is the

ratio of flat fronts to non-flat fronts present at a particular pacing frequency that characterises the connectivity (figure 75b).

We can then connect these two magnitudes to the pacing frequency (figure 75a and 75a), but above all, we can statistically compare the distribution of one of these magnitudes fixing the other and therefore establishing a direct connection from the individual cell response to the propagation of the calcium signal through the culture (figure 75c). This work was published [196] and figure 75 is an extract of the work. The experimental data was taken from 35 cultures of the HL-1 cell line, with between 500 and 1000 cells each field-of-view, and all paced with the same protocol consisting of 5 periods of increasing stimulation rates and two resting periods (one before and one after the protocol). The statistical comparison of the distribution of area ratios covered by irregular cells for the cases when all fronts were flat versus cases in which at least one front was not flat was carried out with the Wilcoxon rank-sum test and a p value under 10^{-3} was obtained.

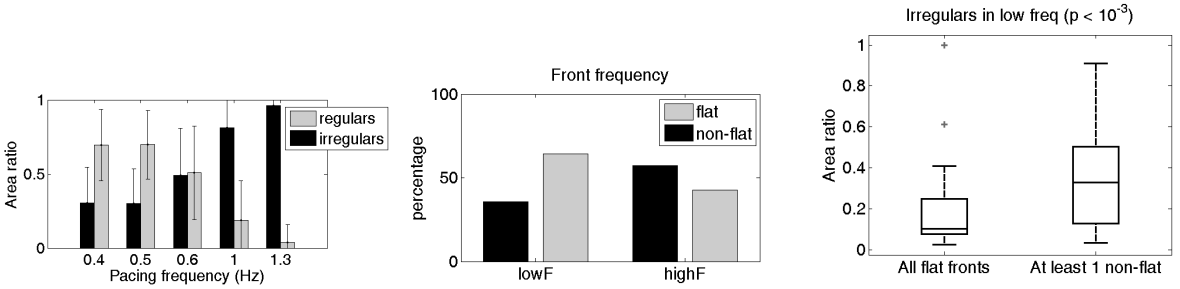

(a) *Mean area occupied by a cells of each (complementary) class is close to linear with frequency change (slope* $0{,}77 \pm 0{,}11 Hz^{-1}$*).*

(b) *Mean values of the percentage of fronts present in each (complementary) class. Wilcoxon rank-sum test produced* $p < 0{,}001$*.*

(c) *Irregular cells area distribution in two front dynamics situations. Wilcoxon rank-sum test produced* $p < 0{,}001$*.*

Fig. 75 *Results of a particular culture study. The last panel is the most important result where we can see that when non-flat fronts start to appear, the area occupied by irregular responding cells is greater.*

All the methods described in this chapter together with the single cell response classifier from the previous chapter constitute an automatic processing pipeline that analyses a sequence of fluorescence images from an excitable cell culture and provides numeric quantification of the spatio-temporal dynamics that are present, both at cellular scale and the culture scale. This pipeline is the body of a patent published in January 2016 [197].

The methods allow for comparison of control versus test cells, wether in a group of cells a mutation has been introduced or that a drug has been applied. If these cells in the culture have been marked by means of a different dye (other than the calcium indicator) a segmentation technique can be applied separately. Then a group of equal size of controls can be selected for comparison. A good way to select the controls is to search for cells of similar size and similar position. For each dynamical regime, a contingency table can be built where one counts the number of cells presenting that dynamical regime versus the number of cells not presenting the regime, and that can be done for the control cells and for the test cells. Each

of the contingency tables can be statistically tested using Fisher exact test [202] against the hypothesis that the data belongs to the same distribution.

6.5 Results

The culture analysis was performed using a total of 47 different culture experiments. These experiments were part of different projects, and had different types of cells under different drug conditions, but they all served to develop, test and improve the different parts involved such as the cell response classifier, the culture cell segmentation or the front analysis.

Culture characteristics	
Total number of cultures	47
Number of sections in each	7
Total number of cells	25988
Mean number of cells	553 ± 65
Total number of signals classified	181916

Table 21 *Summary of the dataset used for developing the culture analysis programs.*

As described in section 5.1 the experiments are divided into different sections according to the stimulation frequency applied. With a total of 47 experiments each with 7 sections and around 550 cells each, this produced a total of nearly 182k cell response signals which were further classified into one of the six cell dynamical regimes established (also section 5.1). Table 21 summarises this dataset.

Section 5.2 describes how a subset of these signals where selected to build the training set for the classifier. Here the total 182k sample signals where classified (table 22) and these numbers are proof that the method has been extensively tested.

Being under development during five years, several parts of it have been published as individual methods: the cell response classifier [187], the cell segmentation algorithm [199] or the front propagation analysis,[196], and the whole method is under a spanish patent [197]. Several studies have used this culture analysis method to in different conditions [200, 190] they are discussed in the conclusions chapter, section 7.2.

Dynamical regime	Number	Percentage
Regular	48026	26.4 %
Phase-lock	9776	5.4 %
Alternant	2887	1.6 %
Wave	48819	26.8 %
Irregular	64351	35.4 %
Inactive	8057	4.4 %

Table 22 *Summary of the signal classes used for developing the cell response classification and culture analysis programs.*

7 Conclusions

This thesis is a collection of methods for image processing in calcium imaging using fluorescence microscopy, ranging from the diffraction limit scale to the culture scale. The development of these methods has allowed significant contributions to science in the understanding of cardiac function at the cellular level and its relation to cardiac pathologies such as arrhythmia and fibrillation. The last section of this chapter includes a brief summary of the most important publications that have been carried out during the thesis.

7.1 Summary of Contributions

The specific contributions in the field of calcium fluorescence microscopy image processing are listed as individual methods in the following section. After that another itemisation lists the findings that the thesis has allowed from a more general physiology and medical point of view.

7.1.1 Computational tools

Organelle scale

- A fast, reliable and fully automatic method for localisation of sub-diffraction limit structures in frame scans together with the measurement of their sizes [173, 174, 203, 204].
- A simplification of the method that is applicable to line scan images [170].
- A mean Z-line distance estimation measurement method based on the position of detected RyR2 clusters [175].

Intracellular scale

- A very sensitive calcium spark detector for frame scans that measures morphological properties of the detected objects to further refine the detection based on a set of tuneable parameters [205, 206, 181–183].

- In particular an uprise position detector is proposed for sparks and transients, and in general for any signal with a rapid upstroke [171].

- A calcium wave tracking method based on a probabilistic approach using information of the position, speed and size of the wave [184, 207].

- A general calcium event detector for line scans that is capable of detecting and classifying sparks, mini-waves and waves [180, 208, 207].

Cellular scale

- Definitions of the possible dynamical regimes an excitable cell may enter during its stimulation [185, 187].

- A proposal for a set of features to describe these dynamical regimes together with a classification method [209, 185].

- An alternans detector and quantification method for line scans and frame scans capable of distinguishing spatially concordant from spatially discordant alternans in the case of frame scans [191, 190].

Culture scale

- A variation of the well known watershed transform focused on segmentation of cardiac myocytes in a culture during a stimulation protocol [190, 196, 198].

- A method to determine the propagation of fronts during this type of experiments by producing isochronal maps and measuring the linear and angular velocity of the action potential throughout the culture [196, 197, 200, 199].

7.1.2 Physiological findings

The contributions to knowledge in cardiac physiology that this thesis has allowed can be divided into four main subjects of study. The next four sections correspond to these four subjects and the contributions are itemised under each section.

Arrhythmia

- Identification of a mutation responsible of increased spontaneous calcium release and ultimately of increased risk of atrial fibrillation [206].
- Promotion of calcium intake by the SERCA pump into the sarcoplasmatic reticulum breaks down calcium waves into separate smaller waves which suppresses ventricular tachycardia [180].
- LDL concentration affects sarcoplasmatic reticulum load and calcium transient amplitude and regularity where as HDL does not. Higher LDL levels also negatively affect action potential propagation across cardiac myocite cultures [200].

Alternans

- RyR2 cluster function is closely linked to the appearence of alternans at the cellular level. Suppression of its function promotes cardiac alternans at the heart level [191].
- A2A receptor activation also promotes alternans at the cellular level [190].

Ageing

- Low calcium release by the sarcoplasmatic reticulum and long calcium transients are linked to a pathologically slow heart rhythm related to ageing [209].
- Low calcium concentration in the sarcoplasmatic reticulum and low membrane calcium current is also related to ageing [185].

Structure & Function

- RyR2 clusters are the only calcium release receptors from the sarcoplasmatic reticulum [170].

7.2 Publication Review

This section is a review of the journal publications that have been carried out during the development of this thesis and is therefore a proof of the thesis impact.

The 4q25 variant rs13143308T links risk of atrial fibrillation to defective pitx2c and calcium homeostasis. Pending publication **[U1].**

In this paper, a particular single nucleotide mutation in chromosome four is shown to be more frequent in atrial fibrillation patients and is linked to spontaneous release of calcium from the sarcoplasmatic reticulum.

The frame scan spark detector presented in this thesis was used to quantitatively measure the increase in spontaneous calcium release. Sparks were located, morphological features were measured and also a clustering method of the spark sites was applied. This allowed to report both a significantly higher spark frequency and higher number of release sites in patients with the mutation.

Suppression of Ryanodine Receptor Function Prolongs Ca2+ Release Refractoriness and Promotes Cardiac Alternans in Intact Hearts. Biochemical Journal, October 2016 **[J1].**

This publication reports a longer refractory period in calcium release from the sarcoplasmatic reticulum and an increase in alternating behaviour in cardiac myocytes of mice that have been genetically engineered to present a mutation that supresses function of the RyR2. It also demonstrates that drugs affecting RyR2 activity (caffeine suppressing and carvedilol promoting it) have the same effect on myocyte alternans, therefore linking RyR2 to alternans at cellular and whole heart level and arrhythmia.

To demonstrate the increase and decrease of alternans depending on the drugs applied, the alternans identification and quantification method presented in this thesis was used together with the transient morphology characterisation methods. It showed statistically significant changes in the severity of alternans and in the duration throughout the experiments and also showing the dependency of pacing frequency with appearance of alternans.

Cardiac electrical defects in progeroid mice and Hutchinson-Gilford progeria syndrome patients with nuclear lamina alterations. PNAS, October 2016 **[J2].**

Abnormal calcium dynamics is reported in cardiac myocytes of mice with a particular mutation and human patients with a rare disease. In both cases, alterations in an essential component of the cell nucleus leads to symptoms related to premature ageing. In the paper it

is shown that the cardiac rhythm of the mice presents a slower rhythm and this is linked at the cellular level to longer calcium transients and lower sarcoplasmatic reticulum capacity and release. These cellular level abnormalities are proposed to be the cause of the premature death in the patients due to myocardial infarction.

Single cell response to increasing pacing was classified with the discussed methodology in this thesis permitting differences to be established in both in the calcium dynamics and in the transient kinetics.

Prevention of adenosine A2A receptor activation diminishes beat-to-beat alternation in human atrial myocytes. Basic Research in Cardiology, January 2016 **[J3].**

This article reports a relation between the activation of a transmembrane protein and irregular response of human atrial myocytes under high frequency pacing. It also shows similar results in mouse myocyte cultures and in porcine atria.

The contribution here was analysing the cultures, in which the cell segmentation algorithm was used to automatically measure the calcium signal of all cells in the cultures. The cell response classifier was further applied showing that the fraction of cells that were responding irregularly increased significantly when the receptor was activated therefore proposing its inhibition as a means of preventing atrial fibrillation.

Distribution and Function of Cardiac Ryanodine Receptor Clusters in Live Ventricular Myocytes. Journal of Biological Chemistry, August 2015 **[J4].**

This paper shows the spatial distribution of RyR2 clusters in mouse ventricular myocytes and links them to spark origination sites. Here the contribution was both in line scans and frame scans. In both cases the RyR2 cluster detectors proposed in the thesis were used and in both cases also the spark detection methods were used. Once sparks and RyR2 clusters were located in space and time directions, the distance from each spark to the closest RyR2 was measured, which proved that sparks originate exclusively from RyR2 clusters.

Ageing is associated with deterioration of calcium homeostasis in isolated human right atrial myocytes. Cardiovascular Research, April 2015 **[J5].**

This publication is a study on how age effects calcium dynamics at the cellular level. It describes a general depression in function including lower membrane calcium current and lower intacellular and sarcoplasmatic reticulum calcium concentration amongst others. Human myocytes were analysed using the image processing techniques presented in this thesis, such as the single cell segmentation, signal extraction from concentric layers, and

morphology and kinetics measurement. Furthermore the signals were analysed and classified to automatically quantify cell response as a function of age and the frame scan spark detector was applied to report no changes in spontaneous activity.

Phospholamban knockout breaks arrhythmogenic Ca2+ waves and suppresses catecholaminergic polymorphic ventricular tachycardia in mice. Circulation Research, August 2013 **[J6].**

This article presents an analysis on the effects of a particular protein, phospholamban (PLN), in calcium dynamics and relates these effects to ventricular tachycardia. The protein is a known inhibitor of the SERCA pump that collects calcium into the sarcoplasmatic reticulum, and here it is shown that by inhibiting PLN, therefore promoting calcium intake by SERCA, tachycardia inducing calcium waves are broken down into smaller units which suppresses the tachycardia. Mice were genetically engineered to present lack of PLN (knockout) and these were crossbred with another line of mice that have a mutation making them prone to ventricular tachycardia producing a third set of mice presenting both features.

For this particular study the line scan general calcium event detector and classifier was developed and it allowed us to automatically produce event frequency distribution for the three sets of mice studied, thus quantifying the absence of waves. This batch processing also allowed the testing of the counter-effects on a series of drugs attempting to recover the calcium waves.

Low Density Lipoproteins Promote Unstable Calcium Handling Accompanied by Reduced SERCA2 and Connexin-40 Expression in Cardiomyocytes. PLoS ONE, March 2013 **[J7].**

In this paper the negative effect of a type of cholesterol (LDL) in the accommodation of cardiac cells to increasing frequencies is shown and linked to changes at the cellular level. It is shown that transient amplitude is reduced, sarcoplasmatic reticulum load is reduced and overall the calcium dynamics becomes irregular. For the study cultures of mouse myocytes were used and signal propagation speed across the culture was also reported to decrease with LDL concentration.

The analysis of the cultures was performed using the methods described in this thesis, such as the cell segmentation, the dynamical regime classifier and the front propagation quantification technique. This showed smaller calcium transients under LDL and that as pacing frequency increased, the number of cells presenting irregular response or alternans was significantly greater at higher LDL concentration.

7.3 Further work

There is still plenty of work to do in the field. In my opinion, the one thing that is needed and that I regret not having achieved during this thesis is a unification of the calcium event detector. This would be a sort of general purpose tool for x-y-t calcium imaging that identifies objects without taking into account any parameters or a pre-established event classification. More like what a naive infant would do: simply segment relevant concentration increases and maybe after that attempt to perform some kind of unsupervised clustering that is ignorant of current trends in event classification.

The RyR2 cluster distribution inside cardiac myocytes is still not fully understood and for that stacks of images need to be analysed (x-y-z). De-convolution type techniques need to be developed to overcome the greater diffraction limit in the optical axis direction to be able to locate the clusters properly in the 3D space. Further work also needs to be done in connecting the calcium release from these clusters with their spatial distribution by, for a start, performing proper local identification of z-lines and by measuring diffusion of sparks and relating it to such z-lines.

Regarding the classification of calcium dynamics it is possible that neural network based methods will achieve better results and that is something else that is yet to be done. The current classifier is certainly limited by the set of features, so might be a good idea to broaden it up or maybe even not take features and rather use the signal samples as features in the neural network (by adjusting the signal to a particular number of samples).

Another idea close to the afore-mentioned is to develop a convolutional neural network that can take the culture image sequences and maybe its segmentation as input and establish either cell response classification or culture response classification.

List of Publications

This is a list of articles that either publish some of the sections in this thesis or, rather more frequently, use the techniques described to show findings in the field of biomedical science. They are divided into sections according to the type of publication and within each section they are presented in descending order by date.

Journal Articles:

[J1] X. Zhong, B. Sun, A. Vallmitjana, W. Gou, R. Wang, A. Guo, H. Duff, A. Gillis, L. S. Song, L. Hove-Madsen, R. Benitez, and S. R. W. Chen. Suppression of Ryanodine Receptor Function Prolongs Ca2+ Release Refractoriness and Promotes Cardiac Alternans in Intact Hearts. *BIOCHEMICAL JOURNAL*, 473(21):3951-3964, OCT 2016.

[J2] J. Rivera-Torres, A. Llach, G. Guzmán-Martínez, C. Calvo, R. Caballero, C. González-Gómez, L. J. Jiménez-Borreguero, J. A. Guadix, F. G. Osorio, C. Lopez-Otin, A. Herraiz-Martínez, N. Cabello, A. Vallmitjana, R. Benitez, L. B. Gordon, J. Jalife, J. M. Pérez-Pomares, J. Tamargo, E. Delpon, L. Hove-Madsen, D. Filgueiras-Rama, and V. Andres. Cardiac electrical defects in progeroid mice and Hutchinson-Gilford progeria syndrome patients with nuclear lamina alterations. *PROCEEDINGS OF THE NATIONAL ACADEMY OF SCIENCES*, 113(46):E7250-E7259, OCT 2016.

[J3] C. E. Molina, A. Llach, A. Herraiz-Martinez, C. Tarifa, M. Barriga, R. F. Wiegerinck, J. Fernandes, N. Cabello, A. Vallmitjana, R. Benítez, J. Montiel, J. Cinca, and L. Hove-Madsen. Prevention of adenosine A2A receptor activation diminishes beat-to-beat alternation in human atrial myocytes. *BASIC RESEARCH IN CARDIOLOGY*, 111(1):5, JAN 2016.

[J4] F. Hiess, A. Vallmitjana, R. Wang, H. Cheng, H. E. D. J. ter Keurs, J. Chen, L. Hove-Madsen, R. Benítez, and S. R. W. Chen. Distribution and Function of Cardiac Ryanodine Receptor Clusters in Live Ventricular Myocytes. *JOURNAL OF BIOLOGICAL CHEMISTRY*, 290(33):20477–20487, AUG 2015.

[J5] A. Herraiz-Martínez, J. Alvarez-Garcia, A. Llach, C. E. Molina, J. Fernandes, A. Ferrero-Gregori, C. Rodríguez, A. Vallmitjana, R. Benítez, R. Padró, J. Martínez-González, J. Cinca, and L. Hove-Madsen. Ageing is associated with deterioration of calcium homeostasis in isolated human right atrialmyocytes. *CARDIOVASCULAR RESEARCH*, 106(1):76–86, APR 2015.

[J6] Y. Bai, P. P. Jones, J. Guo, X. Zhong, R. B. Clark, Q. Zhou, R. Wang, A. Vallmitjana, R. Benítez, L. Hove-Madsen, L. Semeniuk, A. Guo, L. S. Song, H. J. Duff, and S. R. W. Chen. Phospholamban knockout breaks arrhythmogenic Ca2+ waves and suppresses catecholaminergic polymorphic ventricular tachycardia in mice. . *CIRCULATION RESEARCH*, 113(5):517–26, AUG 2013.

[J7] M. Barriga, R. Cal, N. Cabello, A. Llach, A. Vallmitjana, R. Benítez, L. Badimon, J. Cinca, V. Llorente-Cortes, and L. Hove-Madsen. Low Density Lipoproteins Promote Unstable Calcium Handling Accompanied by Reduced SERCA2 and Connexin-40 Expression in Cardiomyocytes. *PLoS ONE*, 8(3):e58128, MAR 13 2013.

Article Pending Publicaton:

[U1] A. Herraiz-Martínez, A. Llach, C. Tarifa, E. Lozano-Velasco, S. A. Serra, J. Gandía, A. Vallmitjana, E. Vázquez Ruiz de Castroviejo, R. Benítez, A. Aranega, C. Muñoz-Guijosa, D. Franco, J. Cinca, , and L. Hove-Madsen. The 4q25 variant rs13143308T links risk of atrial fibrillation to defective pitx2c and calcium homeostasis. , 2017.

Patent:

[Pt1] A. Vallmitjana, L. Hove-Madsen, and R. Benítez. Método implementado por ordenador para caracterización dinámica de celulas en cultivos celulares y programas informáticos para llevar a cabo el método. Spanish patent number P201330974, sent 28 JUN 2013, published 14 JAN 2016.

Conference Presentations:

[C1] C. Tarifa, A. Herraiz, A. Llach, N. Cabello, C. E. Molina, A. Vallmitjana, J. Montiel, R. Benítez, J. Cinca, and L. Hove-Madsen. A higher frequency and proximity of calcium sparks to the cell membrane may potentiate arrhythmogenic afterdepolarizations in atrial myocytes from patients with atrial fibrillation. *European heart journal*, 35:183–183, SEP 2014. ESC Congress; AUG 30-SEP 03, 2014; Barcelona, SPAIN.

[C2] A. Vallmitjana, P. E. Molina, A. Herraiz, J. Montiel, J. Cinca, R. Benítez, and L. Hove-Madsen. Detection, quantification and visualization of ryanodine receptor phosphorylation in human atrial myocytes using a novel ratiometric immunofluorescent analysis. *Cardiovascular Research*, JUL 4 2014. 3rd Congress of the ESC-Council-on-Basic-Cardiovascular-Science on Frontiers in Cardiovascular Biology; JUL 04-06, 2014; Barcelona, SPAIN.

[C3] A. Llach, J. Rivera, A. Herraiz, N. Cabello, C. Gonzalez, A. Vallmitjana, R. Benítez, V. Andres, and L. Hove-Madsen. Abnormal calcium handling in the progeric mouse model Zmpste24-/- is linked to reduced sarcoplasmic reticulum calcium loading and depression of SERCA2 and calsequestrin-2 expression. *European heart journal*, 34:7–7, Aug 2013. Congress of the European-Society-of-Cardiology, AUG 31-SEP 04, 2013, Amsterdam, NETHERLANDS.

[C4] A. Vallmitjana, M. Barriga, L. Hove-Madsen, and R. Benítez. Multilevel analysis of calcium dynamics in stimulated cultures of cardiomyocytes. *2013 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 6514–6517, July 2013. 2013 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Osaka, JAPAN.

[C5] M. Barriga, E. Alvarez-Lacalle, A. Llach, A. Vallmitjana, C. E. Munoz, J. Cinca, R. Benítez, and L. Hove-Madsen. Effect of the resting membrane potential on the distribution, dimensions, and frequency of calcium sparks in human atrial myocytes. *European heart journal*, 31:895–895, SEP 2010. ESC Congress; AUG 28-SEP 01, 2010; Stockholm, SWEDEN.

[C6] M. Barriga, E. Alvarez-Lacalle, A. Llach, A. Vallmitjana, C. Munoz, J. Cinca, R. Benítez, and L. Hove-Madsen. Efecto del potencial de membrana sobre las dimensiones, la cinetica y la frecuencia de la liberation local de calcio en miocitos auriculares humanos. *Revista Espanola de Cardiología*, page 95, OCT 2010. Congreso de la Sociedad Espanola de Cardiología, Valencia, SPAIN.

[C7] A. Vallmitjana, M. Barriga, Z. Nenadic, A. Llach, E. Alvarez-Lacalle, L. Hove-Madsen, and R. Benítez. Identification of intracellular calcium dynamics in stimulated cardiomyocytes. *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 68–71, SEP 2010. 32nd Annual International Conference of the IEEE Engineering-in-Medicine-and-Biology-Society; AUG 30-SEP 04, 2010; Buenos Aires, ARGENTINA.

Conference Posters:

[P1] A. Vallmitjana, C. Tarifa, L. Hove-Madsen, and R. Benítez. Detection and characterization of spontaneous calcium release events in cardiac myocytes. JUN 2017. Single-Cell Biophysics: Measurement, Modulation and Modeling, Taipei, TAIWAN.

[P2] V. Jiménez-Sábado, A. Herraiz, C. Nolla-Colomer, C. Tarifa, A. Vallmitjana, R. Benítez, J. Cinca, and L. Hove-Madsen. Effects of atrial fibrillation on the distribution and phosphorylation of ryanodine receptor clusters at the sites ser2808 and ser2814 in human atrial myocytes. JUN 2017. Cardiac Arrhythmia Copenhagen Meeting, Copenhagen, DENMARK.

[P3] C. Tarifa, A. Herraiz-Martinez, A. Vallmitjana, S.A. Serra, D. Franco, R. Benítez, and L. Hove-Madsen. Atrial specific pitx2 insufficiency increases the frequency of calcium sparks, waves and after-depolarisations in mouse atrial myocytes. *Biophysical journal*, 112(3):400A, FEB 14 2017. 61st Annual Meeting of the Biophysical-Society; FEB 11-15, 2017; New Orleans, LA, USA.

[P4] X. Marimon, A. Vallmitjana, P. Jones, and R. Benítez. A New User-Friendly Tool for Localizing Spots in Fluorescence Microscopy Images. *2015 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, AUG 29 2015. 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Milan, ITALY.

[P5] A. Vallmitjana, C. Nolla, L. Hove-Madsen, and R. Benítez. Spatial localization of ryanodine receptors in human cardiac cells. *2015 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 6297–300, AUG 28 2015. 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Milan, ITALY.

[P6] A. Vallmitjana, C. Nolla, A. Herraiz, L. Hove-Madsen, and R. Benítez. Multiscale Detection of Intracellular Protein Structures from Microscope Fluorescence Images. APR 17 2015. ISBI - IEEE International Symposium on Biomedical Imaging, New York, NY, USA.

[P7] A. Vallmitjana, F. Hiess, S. R. W. Chen, L. Hove-Madsen, and R. Benítez. Simultaneous Detection and Co-localization of Calcium Sparks and Ryanodine Receptor Clusters in Cardiac Myocytes. *Biophysical journal*, 108(2):262A–262A, JAN 27 2015. 59th Annual Meeting of the Biophysical-Society; FEB 7-11, 2015; Baltimore, MD, USA.

[P8] F. Hiess, R. Wang, A. Vallmitjana, D. R. L. Scriven, L. Hove-Madsen, R. Benítez, E. D. W. Moore, and S. R. W. Chen. Superresolution Imaging of RYR2 Clusters in GFP RYR2 Knock in Mouse Cardiomyocytes. *Biophysical journal*, 106(2):398A–398A, JAN 28 2014. 58th Annual Meeting of the Biophysical-Society; FEB 15-19, 2014; San Francisco, CA, USA.

[P9] R. Benítez, A. Vallmitjana, and L. Hove-Madsen. A Method to Characterize Calcium Activity in Stimulated Cultures of Cardiac Myocytes. *Biophysical journal*, 104(2):607A–607A, JAN 29 2013. 57th Annual Meeting of the Biophysical-Society; FEB 02-06, 2013; Philadelphia, PA, USA.

[P10] P. P. Jones, Y. Bai, A. Vallmitjana, R. Wang, R. Clark, Q. Zhou, R. Benítez, L. Hove-Madsen, and S. Chen. Disrupting Ca2+ waves, but not Ca2+ leak, protects against cardiac arrhythmia. *Heart, Lung and Circulation*, 21, 2012. 60th Annual scientific meeting of the Cardiac Society of Australia and New Zealand (CSANZ); AUG 16-19, 2012; Brisbane, QLD, AUSTRALIA.

[P11] A. Vallmitjana, E. Alvarez-LAcalle, L. Hove-Madsen, and R. Benítez. Automatic method for the detection and characterization of calcium activity in cardiac cells. MAY 5 2012. ISBI - IEEE International Symposium on Biomedical Imaging, Barcelona, SPAIN.

[P12] M. Barriga, C. Molina, A. Llach, N. Cabello, A. Vallmitjana, R. Benítez, J. Padro, J. Cinca, and L. Hove-Madsen. Adenosine A2A receptor activation decreases beat-to-beat stability of intracellular calcium transients and their propagation in atrial myocytes, FEB 25-29 2012. 56th Annual Meeting of the Biophysical Society, San Diego, CA, USA.

[P13] R. Benítez, A. Vallmitjana, L. Hove-Madsen, and P. P. Jones. Detection and classification of spontaneous calcium release events in cardiac myocytes, FEB 25-29 2012. 56th Annual Meeting of the Biophysical Society, San Diego, CA, USA.

[P14] M. Barriga, A. Llach, E. Lacalle, N. Cabello, B. Ballester, A. Vallmitjana, C. Munoz, R. Benítez, J. Cinca, and L. Hove-Madsen. Effects of atrial fibrillation on the distribution, dimensions, and frequency of calcium sparks in human atrial myocytes. *European heart journal*, 32:468–468, AUG 2011. ESC Congress; AUG 28-31, 2011; Paris, FRANCE.

References

[1] K. T. Mehta and V. Shah, "Information revolution: Impact of technology on global workforce," *Journal of Internation Information Management*, vol. 6, no. 2, pp. 85–94, 1997.

[2] M. Humbert, "Technology and Workforce : Comparison between the Information Revolution and the Industrial Revolution," *University of California, Berkeley*, no. December, 2007.

[3] M. R. Gillings, M. Hilbert, and D. J. Kemp, "Information in the Biosphere: Biological and Digital Worlds," *Trends in Ecology and Evolution*, vol. 31, no. 3, pp. 180–189, 2016.

[4] Cisco®, "The Zettabyte Era—Trends and Analysis," *Cisco VNI Forecast and Methodology, 2015-2020*, 2016.

[5] J. P. A. Ioannidis, "Why Most Published Research Findings Are False," *PLOS Medicine*, vol. 2, no. 8, pp. 696–701, 2005.

[6] Editorial, "Fix the PhD.," *Nature*, vol. 472, no. 7343, pp. 259–260, 2011.

[7] C. Kervrann, C. Óscar, S. Sorzano, S. T. Acton, and M. Unser, "A Guided Tour of Selected Image Processing and Analysis Methods for Fluorescence and Electron Microscopy," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 1, pp. 6–30, 2016.

[8] A. E. Carpenter, "Software opens the door to quantitative imaging," *Nature Methods*, vol. 4, no. 2, pp. 120–121, 2007.

[9] A. E. Carpenter, "Extracting Biomedically Important Information From Large , Automated Imaging Experiments," *IEEE International Symposium on Biomedical Imaging*, vol. 2011, pp. 1723–1726, 2011.

[10] M. Watabe, S. N. V. Arjunan, S. Fukushima, and K. Iwamoto, "A Computational Framework for Bioimaging Simulation," *Plos One*, pp. 1–19, 2015.

[11] K. W. Eliceiri, M. R. Berthold, I. G. Goldberg, L. Ibáñez, B. S. Manjunath, M. E. Martone, R. F. Murphy, H. Peng, A. L. Plant, B. Roysam, N. Stuurmann, J. R. Swedlow, P. Tomancak, and A. E. Carpenter, "Biological imaging software tools," *Nature*, vol. 9, no. 7, 2012.

[12] A. E. Carpenter, L. Kamentsky, and K. W. Eliceiri, "A call for bioimaging software usability," *Nature Methods*, vol. 9, pp. 660–670, 2012.

[13] L. Shamir, J. D. Delaney, N. Orlov, D. M. Eckley, and I. G. Goldberg, "Pattern Recognition Software and Techniques for Biological Image Analysis," *Plos Computational Biology*, vol. 6, no. 11, 2010.

[14] A. Allalou and C. Wählby, "BlobFinder, a tool for fluorescence microscopy image cytometry," *Computer Methods and Programs in Biomedicine*, vol. 4, pp. 58–65, 2009.

[15] N. J. Severs, "The cardiac muscle cell," *Bioessays*, vol. 22, no. 2, pp. 188–199, 2000.

[16] A. Sandow, "Excitation-contraction coupling in muscular response.," *The Yale journal of biology and medicine*, vol. 25, pp. 176–201, dec 1952.

[17] M. D. Stern, "Theory of excitation-contraction coupling in cardiac muscle.," *Biophysical journal*, vol. 63, pp. 497–517, aug 1992.

[18] D. M. Bers, "Cardiac Excitation Contraction Coupling," *Nature*, vol. 415, no. 6868, pp. 198–205, 2002.

[19] D. M. Bers, "Calcium and cardiac rhythms: physiological and pathophysiological.," *Circulation research*, vol. 90, pp. 14–7, jan 2002.

[20] A. Peskoff and G. A. Langer, "Calcium Concentration and Movement in the Ventricular Cardiac Cell during an Excitation-Contraction Cycle," *Biophysical journal*, vol. 74, pp. 153–174, 1998.

[21] M. Endo, "Calcium Release from the Sarcoplasmic Reticulum," *Physiological Reviews*, vol. 57, no. 1, pp. 71–108, 1977.

[22] A. Fabiato, "Calcium-Induced Release of Calcium from the Cardiac Sarcoplasmic-Reticulum," *American Journal of Physiology*, vol. 245, no. 1, pp. C1–C14, 1983.

[23] Hecht Eugene and Zajac Alfred, *Optics*. Addison-Wesley series in physics, 1974.

[24] M. Rubart, E. Wang, K. W. Dunn, and L. J. Field, "Two-photon molecular excitation imaging of Ca2+ transients in Langendorff-perfused mouse hearts.," *American journal of physiology. Cell physiology*, vol. 284, pp. C1654–68, jun 2003.

[25] R. M. Paredes, J. C. Etzler, L. T. Watts, W. Zheng, and J. D. Lechleiter, "Chemical calcium indicators," *Methods*, 2008.

[26] O. Shimomura, F. H. Johnson, and Y. Saiga, "Extraction, Purification and Properties of Aequorin, a Bioluminescent Protein from the Luminous Hydromedusan, Aequorea," *Journal of Cellular and Comparative Physiology*, vol. 59, no. 3, pp. 223–239, 1962.

[27] M. Chalfie, Y. Tu, G. Euskirchen, W. W. Ward, and D. C. Prashert, "Green Fluorescent Protein as a Marker for Gene Expression," *Science*, vol. 5148, pp. 802–805, 1994.

[28] J. Marshall, R. Molloy, G. W. Moss, J. R. Howe, and T. E. Hughes, "The jellyfish green fluorescent protein: A new tool for studying ion channel expression and function," *Neuron*, vol. 14, pp. 211–215, feb 1995.

[29] M. Bates, B. Huang, G. T. Dempsey, and X. Zhuang, "Multicolor super-resolution imaging with photo-switchable fluorescent probes.," *Science*, vol. 317, pp. 1749–53, sep 2007.

[30] E. V. Pitjeva and E. M. Standish, "Proposals for the masses of the three largest asteroids , the Moon-Earth mass ratio and the Astronomical Unit," *Celestial Mechanics and Dynamical Astronomy*, vol. 103, pp. 365–372, 2009.

[31] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*. Prentice-Hall, 2006.

[32] J. B. T. M. Roerdink and A. Meijster, "The Watershed Transform : Definitions , Algorithms and Parallelization Strategies," *Fundamenta Informaticae*, vol. 41, no. 1, pp. 187–228, 2001.

[33] F. Meyer, "Topographic distance and watershed lines," *Mathematical Morphology and its Applications to Signal Processing*, vol. 38, no. 1, pp. 113–125, 1994.

[34] A. L. De Jager and J. V. Vogt, "Development and demonstration of a structured hydrological feature coding system for Europe Development and demonstration of a structured hydrological feature coding system for Europe," *Hydrological Sciences Journal*, vol. 55, no. 5, pp. 661–675, 2010.

[35] J. W. Cooley and J. W. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series," *Mathematics of Computation*, vol. 19, pp. 297–301, 1965.

[36] R. Bracewell, *The Fourier Transform and Its Applications*. McGraw Hill, 1999.

[37] V. Ljosa, A. E. Carpenter, S. Floyd, M. Pacold, R. Thouis, and A. E. Carpenter, "Introduction to the Quantitative Analysis of Two- Dimensional Fluorescence Microscopy Images for Cell- Based Screening," *PLOS Computational Biology*, vol. 5, no. 12, pp. 1–10, 2009.

[38] K. Mcguinness, *Image Segmentation , Evaluation , and Applications*. PhD thesis, Dublin City University, 2009.

[39] L. Yang, *Robust Segmentation and Object Classification in Natural and Medical Images*. PhD thesis, The State University of New Jersey, 2009.

[40] H. Zhang, J. E. Fritts, and S. A. Goldman, "Image segmentation evaluation : A survey of unsupervised methods," *Computer Vision and Image Understanding*, vol. 110, pp. 260–280, 2008.

[41] M. Paulinas and A. Ušinskas, "A Survey of Genetic Algorithms Applications for Image Enhancement and Segmentation," *Information Technology and Control*, vol. 36, no. 3, pp. 278–284, 2007.

[42] C. Carson, S. Belongie, H. Greenspan, and J. Malik, "Blobworld : Image Segmentation Using Expectation-Maximization and Its Application to Image Querying," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 8, pp. 1026–1038, 2002.

[43] C. Hagwood, J. Bernal, M. Halter, and J. Elliott, "Evaluation of Segmentation Algorithms on Cell Populations Using CDF Curves," *IEEE Transactions on Medical Imaging*, vol. 31, pp. 380–390, feb 2012.

[44] A. A. Dima, J. T. Elliott, J. J. Filliben, M. Halter, A. Peskin, J. Bernal, M. Kociolek, M. C. Brady, H. C. Tang, and A. L. Plant, "Comparison of Segmentation Algorithms For Fluorescence Microscopy Images of Cells," *Cytometry Part A*, vol. 79A, no. 7, pp. 545–559, 2011.

[45] P. Kramer, F. Boto, D. Wald, F. Bessy, C. Paloc, C. Callol, A. Letamendia, I. Ibarba, O. Holgado, and J. M. Virto, "Comparison of Segmentation Algorithms for the Zebrafish Heart in Fluorescent Microscopy Images," *Lecture Notes in Computer Science*, vol. 5876, pp. 1041–1050, 2009.

[46] K. Mcguinness and N. E. O. Connor, "A Comparative Evaluation of Interactive Segmentation Algorithms," *Pattern Recognition*, vol. 43, no. 2, pp. 434–444, 2009.

[47] P. D. Kostelec, L. M. Carlin, and B. Glocker, "Learning to detect and track cells for quantitative analysis of time-lapse microscopic image sequences," *ISBI - IEEE International Symposium on Biomedical Imaging*, pp. 4–7, 2015.

[48] Q. Wang, J. Niemi, C.-m. Tan, L. You, and M. West, "Image segmentation and dynamic lineage analysis in single-cell fluorescence microscopy," *Cytometry Part A*, vol. 77, no. 1, pp. 101–110, 2009.

[49] M. D. Abràmoff, P. J. Magalhães, and S. J. Ram, "Image Processing with ImageJ," *Biophotonics International*, 2004.

[50] M. S. Vokes and A. E. Carpenter, "CellProfiler: Open-Source Software to Automatically Quantify Images," *Microscopy Today*, pp. 38–39, 2008.

[51] L. Kamentsky, T. R. Jones, A. Fraser, M.-a. Bray, D. J. Logan, K. L. Madden, V. Ljosa, C. Rueden, K. W. Eliceiri, and A. E. Carpenter, "Improved structure , function and compatibility for CellProfiler : modular high-throughput image analysis software," *Bioinformatics*, vol. 27, no. 8, pp. 1179–1180, 2011.

[52] P. Kankaanpää, L. Paavolainen, S. Tiitta, M. Karjalainen, J. Päivärinne, J. Nieminen, V. Marjomäki, J. Heino, and D. J. White, "BioImageXD: an open, general-purpose and high-throughput image-processing platform," *Nature methods*, vol. 9, no. 7, pp. 683–689, 2012.

[53] F. de Chaumont, S. Dallongeville, N. Chenouard, N. Hervé, S. Pop, T. Provoost, V. Meas-Yedid, P. Pankajakshan, T. Lecomte, Y. L. Montagner, T. Lagache, A. Dufour, and J.-C. Olivo-Marin, "Icy: an open bioimage informatics platform for extended reproducible research," *Nature Methods*, vol. 9, pp. 690–696, 2012.

[54] M. McCormick, X. Liu, J. Jomier, C. Marion, and L. Ibanez, "ITK: enabling reproducible research and open science.," *Frontiers in Neuroinformatics*, vol. 8, no. February, p. 13, 2014.

[55] K. Pulli, A. Baksheev, K. Kornyakov, and V. Eruhimov, "Real-time computer vision with OpenCV," *Communications of the Association for Computing Machinery*, vol. 55, no. 6, p. 61, 2012.

[56] V. Wiesmann, D. Franz, C. Held, C. Münzenmayer, R. Palmisano, and T. Wittenberg, "Review of free software tools for image analysis of fluorescence cell micrographs," *Journal of Microscopy2*, vol. 257, no. 1, pp. 39–53, 2014.

[57] H. Chang, Q. Yang, and B. Parvin, "Segmentation of heterogeneous blob objects through voting and level set formulation," *Pattern Recognition Letters*, vol. 28, pp. 1781–1787, jan 2007.

[58] R. W. Cole, T. Jinadasa, and C. M. Brown, "Measuring and interpreting point spread functions to determine confocal microscope resolution and ensure quality control," *Nature Protocols*, vol. 6, no. 12, pp. 1929–1941, 2011.

[59] T. a. Klar, E. Engel, and S. W. Hell, "Breaking Abbe's diffraction resolution limit in fluorescence microscopy with stimulated emission depletion beams of various shapes.," *Physical Review*, vol. 64, no. 6 Pt 2, p. 066613, 2001.

[60] T. A. Klar, S. Jakobs, M. Dyba, A. Egner, and S. W. Hell, "Fluorescence microscopy with diffraction resolution barrier broken by stimulated emission," *PNAS*, vol. 97, no. 15, 2000.

[61] M. J. Rust, M. Bates, and X. Zhuang, "Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM) Michael," *Nature Methods*, vol. 3, no. 10, pp. 793–796, 2006.

[62] E. Betzig, G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. Olenych, J. S. Bonifacino, M. W. Davidson, J. Lippincott-Schwartz, and H. F. Hess, "Imaging intracellular fluorescent proteins at nanometer resolution.," *Science*, vol. 313, pp. 1642–5, sep 2006.

[63] S. van de Linde, S. Aufmkolk, C. Franke, T. Holm, T. Klein, A. Löschberger, S. Proppert, S. Wolter, and M. Sauer, "Investigating cellular structures at the nanoscale with organic fluorophores.," *Chemistry & biology*, vol. 20, pp. 8–18, jan 2013.

[64] G. Patterson, M. Davidson, S. Manley, and J. Lippincott-Schwartz, "Superresolution imaging using single-molecule localization.," *Annual review of physical chemistry*, vol. 61, pp. 345–67, jan 2010.

[65] M. Sezgin, "Survey over image thresholding techniques and quantitative performance evaluation," *Journal of Electronic Imaging*, vol. 13, no. 1, pp. 146–165, 2004.

[66] N. Otsu, "A Threshold Selection Method from Gray-Level Histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.

[67] J. Ghaye, M. A. Kamat, L. Corbino-Giunta, P. Silacci, G. Vergeres, G. De Micheli, and S. Carrara, "Image thresholding techniques for localization of sub-resolution fluorescent biomarkers," *Cytometry Part A*, vol. 83, no. 11, pp. 1001–1016, 2013.

[68] E. J. Breen, G. H. Joss, and K. L.Williams, "Locating objects of interest within biological images: The top hat box filter," *Journal of Computer Assisted Microscopy*, vol. 3, no. 2, pp. 97–102, 1991.

[69] Y. Kimori, N. Baba, and N. Morone, "Extended morphological processing: a practical method for automatic spot detection of biological markers from microscopic images.," *BMC bioinformatics*, vol. 11, p. 373, 2010.

[70] I. Smal, M. Loog, W. Niessen, and E. Meijering, "Quantitative Comparison of Spot Detection Methods in Live-Cell Fluorescence Microscopy Imaging," *IEEE Transactions on Medical Imaging*, no. c, pp. 2–5, 2009.

[71] I. Smal, K. Draegestein, N. Galjart, W. Niessen, E. Meijering, and S. Member, "Particle Filtering for Multiple Object Tracking in Dynamic Fluorescence Microscopy Images : Application to Microtubule Growth Analysis," *IEEE Transactions on Medical Imaging*, vol. 27, no. 6, pp. 789–804, 2008.

[72] S. H. Rezatofighi, R. Hartley, and W. E. Hughes, "A new approach for spot detection in total internal reflection fluorescence microscopy," *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 860–863, 2012.

[73] M. Dimiccoli, J.-P. Jacob, and L. Moisan, "Particle detection and tracking by a-contrario approach: application to fluorescence time-lapse imaging," *Machine Vision and Applications*, vol. 27, no. 511, 2016.

[74] N. Chenouard, S. Vernhettes, I. Bloch, and J. C. Olivo-Marin, "Morphological source separation for particle tracking in complex biological environments," *IEEE International Conference on Pattern Recognition*, pp. 1–4, 2009.

[75] Z. B, F. J, S. J, and O.-M. J, "Multiscale variance-stabilizing transform for mixed-Poisson-Gaussian processes and its applica- tions in bioimaging," *IEEE International Conference on Image Processing*, vol. 6, 2007.

[76] D. Sage, F. R. Neumann, F. Hediger, S. M. Gasser, and M. Unser, "Automatic tracking of individual fluorescence particles: Application to the study of chromosome dynamics," *IEEE Transactions on Image Processing*, vol. 14, no. 9, pp. 1372–1383, 2005.

[77] I. Smal, M. Loog, W. Niessen, and E. Meijering, "Quantitative Comparison OF Spot Detection Methods in Live-Cell Fluorescence Microscopy Imaging," *IEEE Transactions on Medical Imaging*, vol. 29, no. 2, pp. 282–301, 2010.

[78] P. Ruusuvuori, T. Aijö, S. Chowdhury, C. Garmendia-Torres, J. Selinummi, M. Birbaumer, A. M. Dudley, L. Pelkmans, and O. Yli-Harja, "Evaluation of methods for detection of fluorescence labeled subcellular objects in microscope images.," *BMC bioinformatics*, vol. 11, p. 248, jan 2010.

[79] H. Cheng, W. J. Lederer, and M. B. Cannell, "Calcium Sparks - Elementary Events Underlying Excitation-Contraction Coupling in Heart-Muscle," *Science*, vol. 262, no. 5134, pp. 740–744, 1993.

[80] H. Cheng, L.-s. S. Song, N. Shirokova, A. Gonzalez, E. G. Lakatta, E. Rios, and M. D. Stern, "Amplitude distribution of calcium sparks in confocal images: theory and studies with an automatic detection method.," *Biophysical Journal*, vol. 76, no. February, pp. 606–617, 1999.

[81] G. Q. Zhang, H. Wei, J. Lu, P. Wong, and W. Shim, "Identification and Characterization of Calcium Sparks in Cardiomyocytes Derived from Human Induced Pluripotent Stem Cells," *PLOS One*, vol. 8, no. 2, 2013.

[82] N. Weisleder, J. Zhou, and J. Ma, "Detection of calcium sparks in intact and permeabilized skeletal muscle fibers," *Methods in Molecular Biology*, vol. 798, no. 11, pp. 395–410, 2012.

[83] L. Zsolt Szabó, B. Dienes, and L. Csernoch, "An image analyser program written in MATLAB to detect and analyse calcium sparks," *Proceedings of the 7th International Conference on Applied Informatics*, vol. 1, pp. 155–162, 2007.

[84] X. Wang, N. Weisleder, C. Collet, J. Zhou, Y. Chu, Y. Hirata, X. Zhao, Z. Pan, M. Brotto, H. Cheng, and J. Ma, "Uncontrolled calcium sparks act as a dystrophic signal for mammalian skeletal muscle," *Nature Cell Biology*, vol. 7, no. 5, pp. 525–530, 2005.

[85] M.-A. Bray, N. A. Geisse, and K. K. Parker, "Multidimensional detection and analysis of Ca2+ sparks in cardiac myocytes," *Biophysical journal*, vol. 92, pp. 4433–4443, jun 2007.

[86] E. Picht, A. V. Zima, L. A. Blatter, and D. M. Bers, "SparkMaster : automated calcium spark analysis with ImageJ," *American Journal of Physiology*, vol. 60153, pp. 1073–1081, 2007.

[87] E. M. Steele and D. S. Steele, "Automated detection and analysis of Ca2+ sparks in x-y image stacks using a thresholding algorithm implemented within the open-source image analysis platform ImageJ," *Biophysical Journal*, vol. 106, no. 3, pp. 566–576, 2014.

[88] J. Zhou, G. Brum, A. Gonzalez, B. S. Launikonis, M. D. Stern, and E. Rios, "Ca2+ sparks and embers of mammalian muscle. Properties of the sources.," *The Journal of General Physiology*, vol. 122, no. 1, pp. 95–114, 2003.

[89] A. González, W. G. Kirsch, N. Shirokova, G. Pizarro, M. D. Stern, and E. Ríos, "The spark and its ember: separately gated local components of Ca(2+) release in skeletal muscle.," *The Journal of general physiology*, vol. 115, no. 2, pp. 139–158, 2000.

[90] S. Sebille, A. Cantereau, C. Vandebrouck, H. Balghi, B. Constantin, G. Raymond, and C. Cognard, "Calcium sparks in muscle cells: interactive procedures for automatic detection and measurements on line-scan confocal images series," *Computer Methods and Programs in Biomedicine*, vol. 77, no. 1, pp. 57–70, 2005.

[91] T. Bányász, Y. Chen-Izu, C. W. Balke, and L. T. Izu, "A new approach to the detection and statistical classification of Ca2+ sparks.," *Biophysical journal*, vol. 92, no. 12, pp. 4458–4465, 2007.

[92] K. L. Ellefsen, B. Settle, I. Parker, and I. F. Smith, "An algorithm for automated detection, localization and measurement of local calcium signals from camera-based imaging," *Cell Calcium*, vol. 56, no. 3, pp. 147–156, 2014.

[93] C. H. T. Kong, C. Soeller, and M. B. Cannell, "Increasing sensitivity of Ca2+ spark detection in noisy images by application of a matched-filter object detection algorithm.," *Biophysical journal*, vol. 95, no. 12, pp. 6016–6024, 2008.

[94] F. von Wegner, M. Both, and R. H. A. Fink, "Automated detection of elementary calcium release events using the a trous wavelet transform.," *Biophysical journal*, vol. 90, no. 6, pp. 2151–2163, 2006.

[95] F. von Wegner, M. Both, R. H. A. Fink, and O. Friedrich, "3D high-speed multifocal multiphoton microscopy in muscle and wavelet analysis of fast elementary calcium release events," *IEEE International Symposium on Biomedical Imaging*, 2006.

[96] F. von Wegner, M. Both, R. H. A. Fink, and O. Friedrich, "Fast XYT imaging of elementary calcium release events in muscle with multifocal multiphoton microscopy and wavelet denoising and detection," *IEEE Transactions on Medical Imaging*, vol. 26, no. 7, pp. 925–934, 2007.

[97] L. Z. Szabó, J. Vincze, L. Csernoch, and P. Szentesi, "Improved spark and ember detection using stationary wavelet transforms Author ' s Accepted Manuscript," *Journal of Theoretical Biology*, vol. 264, no. 4, pp. 1279–1292, 2010.

[98] F. Ge and S. Wang, "New benchmark for image segmentation evaluation," *Journal of Electronic Imaging*, vol. 16, no. 3, pp. 1–16, 2007.

[99] G. C. Thompson, T. A. Ireland, X. C. Larkin, J. Arnold, and R. M. D. Holsinger, "A Novel Segmentation-Based Algorithm for the Quantification of Magnified Cells," *Journal of Cellular Biochemistry*, vol. 1854, no. May, pp. 1849–1854, 2014.

[100] E. Meijering, "Cell Segmentation: 50 Years Down the Road," *IEEE Signal Processing Magazine*, vol. 29, no. 5, pp. 140–145, 2012.

[101] R. K. Amanfu, J. B. Muller, and J. J. Saucerman, "Automated Image Analysis of Cardiac Myocyte Ca2+ Dynamics," *33rd Annual International IEEE EMBS Conference*, p. 4664, 2011.

[102] A. J. Walsh and M. C. Skala, "An automated image processing routine for segmentation of cell cytoplasms in high-resolution autofluorescence images," *Proceedings of SPIE - The International Society for Optical Engineering*, p. 8948, 2014.

[103] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik, "Contour Detection and Hierarchical Image Segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 5, pp. 898–916, 2011.

[104] J. Ali, R. Khan, N. Ahmad, and I. Maqsood, "Random forests and decision trees," *IJCSI International Journal of Computer Science Issues*, vol. 9, no. 5, pp. 272–278, 2012.

[105] B. Selig, *Image Segmentation using Snakes and Stochastic Watershed With Applications to Microscopy Images of Biological Tissue*. PhD thesis, University of Uppsala, 2015.

[106] J. Canny, "A Computational Approach to Edge Detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 679 – 698, 1986.

[107] T. Lindeberg and M.-X. Li, "Segmentation and Classification of Edges Using Minimum Description Length Approximation and Complementary Junction Cues," *Computer Vision and Image Understanding*, vol. 67, no. 1, pp. 88–98, 1997.

[108] W. N. Goncalves and O. Martinez Bruno, "Automatic system for counting cells with elliptical shape," *Learning and NonLinear Models*, vol. 9, no. 1, 2012.

[109] J. A. M. Belien, H. A. H. M. van Ginkel, P. Tekola, L. S. Ploeger, N. M. Poulin, J. P. A. Baak, and P. J. V. Diest, "Confocal DNA Cytometry : A Contour-Based Segmentation Algorithm for Automated Three-Dimensional Image Segmentation," *Cytometry*, vol. 49, no. 1, pp. 12–21, 2002.

[110] S. T. Acton, C. Yang, J. a. Hossack, and B. R. Wamhoff, "Poisson inverse gradient approach to vascular myocyte detection and segmentation," *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pp. 1298–1301, jun 2009.

[111] T. M. Lehmann, J. Bredno, and K. Spitzer, "On the design of active contours for medical image segmentation - A scheme for classification and construction," *Methods of information in medicine*, vol. 42, no. 1, pp. 89–98, 2003.

[112] M. Kass, A. Witkin, and D. Terxopoulos, "Snakes: Active Contour Models," *International Journal of Computer Vision*, pp. 321–331, 1988.

[113] O. Dzyubachyk, W. A. V. Cappellen, J. Essers, W. J. Niessen, S. Member, E. Meijering, and S. Member, "Advanced Level-Set-Based Cell Tracking in Time-Lapse Fluorescence Microscopy," *IEEE Transactions on Medical Imaging*, vol. 29, no. 3, pp. 852–867, 2010.

[114] O. Dzyubachyk, W. Niessen, and E. Meijering, "Advanced level-set based multiple-cell segmentation and tracking in time-lapse fluorescence microscopy images," *5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2008.

[115] R. Malladi and J. A. Sethlan, "Image processing via level set curvature flow," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 92, no. July, pp. 7046–7050, 1995.

[116] B. Nielsen, F. Albregtsen, and H. E. Danielsen, "Automatic segmentation of cell nuclei in Feulgen-stained histological sections of prostate cancer and quantitative evaluation of segmentation results.," *Cytometry. Part A : the journal of the International Society for Analytical Cytology*, vol. 81, pp. 588–601, jul 2012.

[117] S. Tse, L. Bradbury, J. W. L. Wan, H. Djambazian, R. Sladek, T. Hudson, and D. R. Cheriton, "A Combined Watershed and Level Set Method for Segmentation of Brightfield Cell Images," *Proceedings of SPIE - The International Society for Optical Engineering*, vol. 7259, 2009.

[118] G. Srinivasa, M. C. Fickus, Y. Guo, and A. D. Linstedt, "Active Mask Segmentation of Fluorescence Microscope Images," *IEEE Transactions on Image Proceessing*, vol. 18, no. 8, pp. 1817–1829, 2009.

[119] J.-P. Bergeest and K. Rohr, "Fast globally optimal segmentation of cells in fluorescence microscopy images.," *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention*, vol. 14, pp. 645–52, jan 2011.

[120] X. Du and S. Dua, "Segmentation of fluorescence microscopy cell images using unsupervised mining," *The Open Medical Informatics Journal,*, 2010.

[121] Y.-y. Chuang, B. Curless, D. H. Salesin, and R. Szeliski, "A Bayesian Approach to Digital Matting," *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, no. 2, pp. 264–271, 2001.

[122] L. Lucchese and S. K. Mitra, "Color Image Segmentation : A State-of-the-Art Survey," *Proceedings of the Indian National Science Academy*, pp. 207–221, 2001.

[123] D. Fenistein, B. Lenseigne, T. Christophe, P. Brodin, and A. Genovesio, "A fast, fully automated cell segmentation algorithm for high-throughput and high-content screening," *Cytometry Part A*, vol. 73A, pp. 958–964, oct 2008.

[124] R. Nock and F. Nielsen, "Statistical region merging," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 11, pp. 1452 – 1458, 2004.

[125] F. Meyer, "The watershed concept and its use in segmentation : a brief history," *Publications from the Center of Mathematical Morphology*, pp. 1–11, 2012.

[126] C. Held, R. Palmisano, L. Haeberle, M. Hensel, T. Wittenberg, C. Held, R. Palmisano, and L. Ha, "Comparison of Parameter-Adapted Segmentation Methods for Fluorescence Micrographs," *Cytometry Part A*, vol. 79A, no. 11, pp. 933–945, 2011.

[127] J. Sigut, F. Fumero, O. Nuñez, and M. Sigut, "Automatic marker generation for watershed segmentation of natural images," *IEEE Electronics Letters*, vol. 50, no. 18, pp. 1281 – 1283, 2014.

[128] J. Ning, L. Zhang, D. Zhang, and C. Wu, "Interactive image segmentation by maximal similarity based region merging," *Pattern Recognition*, vol. 43, no. 2, pp. 445–456, 2010.

[129] L. C. Wong, B. Lu, K. W. Tan, and M. Fivaz, "Fully-automated image processing software to analyze calcium traces in populations of single cells," *Cell calcium*, vol. 48, no. 5, pp. 270–274, 2010.

[130] K.-h. Tan and N. Ahuja, "Selecting Objects With Freehand Sketches," *Proceedings of CVPR*, 2001.

[131] Y. Li, J. Sun, C.-k. Tang, and H.-y. Shum, "Lazy Snapping," *Proceedings of ACM SIGGRAPH 2004*, vol. 23, no. 3, pp. 303–308, 2004.

[132] D. Yu, T. D. Pham, and X. Zhou, "Analysis and recognition of touching cell images based on morphological structures," *Computers in biology and medicine*, vol. 39, no. 1, pp. 27–39, 2009.

[133] B. Selig and C. L. Luengo, "Stochastic watershed—an analysis," *Proceedings of Swedish Society for Automated Image Analysis*, pp. 82–85, 2012.

[134] Z. Xie, G. Chen, W. Lin, K. Yang, J. Chen, and R. Chen, "A new center location algorithm used in cell segmentation," *IEEE Congress on Image and Signal Processing*, p. 614, 2008.

[135] O. Lézoray and L. Grady, *Image Processing and Analysis with Graphs: Theory and Practice*. CRC Press, 2012.

[136] M. Rubinov and O. Sporns, "Complex network measures of brain connectivity: uses and interpretations.," *NeuroImage*, vol. 52, no. 3, pp. 1059–69, 2010.

[137] L. Bradbury and J. W. Wan, "A spectral k-means approach to bright-field cell image segmentation," *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*, 2010.

[138] F. Meyer, "Watersheds on weighted graphs," *Pattern Recognition Letters*, vol. 47, pp. 72–79, 2014.

[139] C. Allene, J.-y. Audibert, M. Couprie, and R. Keriven, "Some links between extremum spanning forests , watersheds and min-cuts," *Image and Vision Computing*, pp. 1–27, 2009.

[140] J. Cousty, G. Bertrand, and L. Najman, "Watershed cuts," *Proceedings of the 8th International Symposium on Mathematical Morphology*, pp. 301–312, 2007.

[141] Y. Guo, X. Xu, Y. Wang, Z. Yang, Y. Wang, and S. Xia, "A computational approach to detect and segment cytoplasm in muscle fiber images.," *Microscopy research and technique*, vol. 78, pp. 508–18, jun 2015.

[142] F. Meyer, "Watersheds on edge or node weighted graphs," *Publications from the Center of Mathematical Morphology*, 2013.

[143] A. Pinidiyaarachchi and C. Wahlby, "Seeded watersheds for combined segmentation and tracking of cells," *Lecture Notes in Computer Science*, 2005.

[144] E. Hodneland, T. Kögel, D. M. Frei, H.-h. Gerdes, and A. Lundervold, "CellSegm - a MATLAB toolbox for high-throughput 3D cell segmentation," *Source Code for Biology and Medicine*, vol. 8, no. 1, pp. 1–10, 2013.

[145] J. Cousty, G. Bertrand, L. Najman, and M. Couprie, "Watershed Cuts : Minimum Spanning Forests and the Drop of Water Principle," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 8, pp. 1362–74, 2009.

[146] F. Meyer, "The steepest watershed : from graphs to images," *Publications from the Center of Mathematical Morphology*, 2012.

[147] L. Najman and M. Couprie, "Watershed algorithms and contrast preservation," *Discrete Geometry for Computer Imagery*, vol. 2886, pp. 62–71, 2003.

[148] E. Meijering, O. Dzyubachyk, I. Smal, and W. A. van Cappellen, "Tracking in cell and developmental biology," *Seminars in Cell & Developmental Biology*, vol. 20, no. 8, pp. 894–902, 2009.

[149] L. P. Coelho, A. Shariff, and R. F. Murphy, "NUCLEAR SEGMENTATION IN MICROSCOPE CELL IMAGES : A HAND-SEGMENTED DATASET AND COMPARISON OF ALGORITHMS," *IEEE International Symposium on Biomedical Imaging*, pp. 518–521, 2009.

[150] N. Malpica, C. O. D. Solo, J. Jose, I. Vallcorba, M. Garcı, and F. Pozo, "Applying Watershed Algorithms to the Segmentation of Clustered Nuclei," *Cytometry*, vol. 297, pp. 289–297, 1997.

[151] S. M. Golodetz, C. Nicholls, I. D. Voiculescu, and S. A. Cameron, "Two tree-based methods for the waterfall," *Pattern Recognition*, vol. 47, no. 10, pp. 3276–3292, 2014.

[152] J. V. Vogt, P. Soille, A. De Jager, E. Rimaviciute, W. Mehl, P. Haastrup, M. L. Paracchini, J. Dusart, K. Bódis, S. Foisneau, and C. Bamps, "Developing a pan-European Data Base of Drainage Networks and Catchment Boundaries from a 100 Metre DEM," *Proceedings of the AGILE Int. Conference on Geographic Information Science*, 2007.

[153] J. Lindblad, C. Wahlby, E. Bengtsson, and A. Zaltsman, "Image Analysis for Automatic Segmentation of Cytoplasms and Classification of Rac1 Activation," *Cytometry Part A*, vol. 33, pp. 22–33, 2004.

[154] F. Buggenthin, C. Marr, M. Schwarzfischer, P. S. Hoppe, O. Hilsenbeck, T. Schroeder, and F. J. Theis, "An automatic method for robust and fast cell detection in bright field images from high-throughput microscopy," *BMC Bioinformatics*, vol. 14, no. 297, 2013.

[155] C. Arteta, V. Lempitsky, J. A. Noble, and A. Zisserman, "Learning to Detect Cells Using Non-overlapping Extremal Regions," *Medical image computing and computer-assisted intervention*, no. Figure 1, pp. 1–8, 2012.

[156] C. Wählby, J. Lindblad, M. Vondrus, E. Bengtsson, and L. Björkesten, "Algorithms for cytoplasm segmentation of fluorescence labelled cells," *Analystical Cellular Pathology*, vol. 24, pp. 101–111, 2002.

[157] S. Chabrier, B. Emile, C. Rosenberger, and H. Laurent, "Unsupervised Performance Evaluation of Image Segmentation," *Journal on Applied Signal Processing*, vol. 2006, pp. 1–12, 2006.

[158] E. F. Rogers, F. R. Koniuszy, J. ShavelJr., and K. Folkers, "Plant Insecticides. I. Ryanodine, A New Alkaloid from Ryania Speciosa Vahl.," *Journal of the American Chemical Society*, vol. 70, no. 9, pp. 3086–3088, 1948.

[159] Y. Wang, J. Wu, M. J. Rowan, and R. Anwyl, "Ryanodine produces a low frequency stimulation-induced NMDA receptor-independent long-term potentiation in the rat dentate gyrus in vitro.," *The Journal of physiology*, vol. 495 (Pt 3, pp. 755–67, sep 1996.

[160] C. Legrand, E. Giacomello, C. Berthier, B. Allard, V. Sorrentino, and V. Jacquemond, "Spontaneous and voltage-activated Ca2+ release in adult mouse skeletal muscle fibres expressing the type 3 ryanodine receptor.," *The Journal of physiology*, vol. 586, pp. 441–57, jan 2008.

[161] M. A. Walker, T. Kohl, S. E. Lehnart, J. L. Greenstein, W. J. Lederer, and R. L. Winslow, "On the Adjacency Matrix of RyR2 Cluster Structures," *Plos Computational Biology*, pp. 1–21, 2015.

[162] T. M. Hoang-Trong, A. Ullah, and M. S. Jafri, "Calcium Sparks in the Heart: Dynamics and Regulation.," *Research and reports in biology*, vol. 6, pp. 203–214, 2015.

[163] P. Wu, W. Peng, H. Shen, J. Wu, W. Guo, X. Pan, R. Wang, S. R. W. Chen, and N. Yan, "Structural basis for the gating mechanism of the type 2 ryanodine receptor RyR2," *Science*, vol. 354, no. September, pp. 1–17, 2016.

[164] Z. Yuchi and F. Van Petegem, "Ryanodine receptors under the magnifying lens: Insights and limitations of cryo-electron microscopy and X-ray crystallography studies," *Cell Calcium*, vol. 59, no. 5, pp. 209–227, 2016.

[165] S. G. Priori and C. Napolitano, "Cardiac and skeletal muscle disorders caused by mutations in the intracellular Ca2+ release channels.," *The Journal of clinical investigation*, vol. 115, pp. 2033–8, aug 2005.

[166] M. Yano, T. Yamamoto, Y. Ikeda, and M. Matsuzaki, "Mechanisms of disease: ryanodine receptor defects in heart failure and fatal arrhythmia," *Nature Clinical Practice Cardiovascular Medicine*, vol. 3, no. 1, pp. 43–52, 2006.

[167] N. Macquaide, H.-t. M. Tuan, J.-i. Hotta, W. Sempels, I. Lenaerts, P. Holemans, J. Hofkens, M. S. Jafri, R. Willems, and K. R. Sipido, "Ryanodine receptor cluster fragmentation and redistribution in persistent atrial fibrillation enhance calcium release," *Cardiovascular Research*, no. 108, pp. 387–398, 2015.

[168] B. Zhang, J. Zerubia, and J. Olivo-Marin, "Gaussian approximations of fluorescence microscope point-spread function models," *Applied Optics*, vol. 46, pp. 1819–1829, 2007.

[169] C. R. Maurer, R. Qi, V. Raghavan, and S. Member, "A Linear Time Algorithm for Computing Exact Euclidean Distance Transforms of Binary Images in Arbitrary Dimensions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 2, pp. 265–270, 2003.

[170] F. Hiess, A. Vallmitjana, R. Wang, H. Cheng, H. E. D. J. ter Keurs, J. Chen, L. Hove-Madsen, R. Benítez, and S. R. W. Chen, "Distribution and function of cardiac ryanodine receptor clusters in live ventricular myocytes," *JOURNAL OF BIOLOGICAL CHEMISTRY*, vol. 290, pp. 20477–20487, AUG 2015.

[171] A. Vallmitjana, F. Hiess, S. R. W. Chen, L. Hove-Madsen, and R. Benítez, "Simultaneous detection and colocalization of calcium sparks and ryanodine receptor clusters in cardiac myocytes," *Biophysical journal*, vol. 108, pp. 262A–262A, JAN 2015. 59th Annual Meeting of the Biophysical-Society; FEB 7-11, 2014;Baltimore, MD, US.

[172] C. I. Danila and S. L. Hamilton, "Phosphorylation of ryanodine receptors," *Biological research*, vol. 37, no. 4, pp. 521–525, 2004.

[173] A. Vallmitjana, P. E. Molina, A. Herraiz, J. Montiel, J. Cinca, R. Benítez, and L. Hove-Madsen, "Detection, quantification and visualization of ryanodine receptor phosphorylation in human atrial myocytes using a novel ratiometric immunofluorescent analysis," *Cardiovascular Research*, JUL 4 2014. 3rd Congress of the ESC-Council-on-Basic-Cardiovascular-Science on Frontiers in Cardiovascular Biology; JUL 04-06, 2014;Barcelona, SPAIN.

[174] A. Vallmitjana, C. Nolla, L. Hove-Madsen, and R. Benítez, "Spatial localization of ryanodine receptors in human cardiac cells," pp. 6297–300, AUG 2015. 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Milan, ITALY.

[175] F. Hiess, R. Wang, A. Vallmitjana, D. R. L. Scriven, L. Hove-Madsen, R. Benítez, E. D. W. Moore, and S. R. W. Chen, "Superresolution imaging of ryr2 clusters in gfp ryr2 knock in mouse cardiomyocytes," *Biophysical journal*, vol. 106, pp. 398A–398A, JAN 28 2014. 58th Annual Meeting of the Biophysical-Society; FEB 15-19, 2014;San Francisco, CA, US.

[176] E. Niggli and N. Shirokova, "A guide to sparkology: The taxonomy of elementary cellular Ca2+ signaling events," *Cell Calcium*, vol. 42, no. 4-5, pp. 379–387, 2007.

[177] H. Cheng and W. J. Lederer, "Calcium sparks.," *Physiological reviews*, vol. 88, pp. 1491–545, oct 2008.

[178] E. Chudin, J. Goldhaber, A. Garfinkel, J. Weiss, and B. Kogan, "Intracellular Ca2+ dynamics and the stability of ventricular tachycardia," *Biophysical journal*, vol. 77, no. 6, pp. 2930–2941, 1999.

[179] H. Ishida, C. Genka, Y. Hirota, H. Nakazawa, and W. H. Barry, "Formation of Planar and Spiral Ca2+ Waves in Isolated Cardiac Myocytes," *Biophysical journal*, vol. 77, no. 4, pp. 2114–2122, 1999.

[180] Y. Bai, P. P. Jones, J. Guo, X. Zhong, R. B. Clark, Q. Zhou, R. Wang, A. Vallmitjana, R. Benítez, L. Hove-Madsen, L. Semeniuk, A. Guo, L. S. Song, H. J. Duff, and S. R. W. Chen, "Phospholamban knockout breaks arrhythmogenic Ca2+ waves and suppresses catecholaminergic polymorphic ventricular tachycardia in mice.," *CIRCULATION RESEARCH*, vol. 113, pp. 517–26, AUG 2013.

[181] C. Tarifa, A. Herraiz, A. Llach, N. Cabello, C. E. Molina, A. Vallmitjana, J. Montiel, R. Benítez, J. Cinca, and L. Hove-Madsen, "A higher frequency and proximity of calcium sparks to the cell membrane may potentiate arrhythmogenic afterdepolarizations in atrial myocytes from patients with atrial fibrillation," *European heart journal*, vol. 35, pp. 183–183, SEP 2014. ESC Congress; AUG 30-SEP 03, 2014; Barcelona, SPAIN.

[182] M. Barriga, A. Llach, E. Lacalle, N. Cabello, B. Ballester, A. Vallmitjana, C. Munoz, R. Benítez, J. Cinca, and L. Hove-Madsen, "Effects of atrial fibrillation on the distribution, dimensions, and frequency of calcium sparks in human atrial myocytes," *European heart journal*, vol. 32, pp. 468–468, AUG 2011. ESC Congress; AUG 28-31, 2011;Paris, FRANCE.

[183] M. Barriga, E. Alvarez-Lacalle, A. Llach, A. Vallmitjana, C. E. Munoz, J. Cinca, R. Benítez, and L. Hove-Madsen, "Effect of the resting membrane potential on the distribution, dimensions, and frequency of calcium sparks in human atrial myocytes," *European heart journal*, vol. 31, pp. 895–895, SEP 2010. ESC Congress; AUG 28-SEP 01, 2010;Stockholm, SWEDEN.

[184] P. P. Jones, Y. Bai, A. Vallmitjana, R. Wang, R. Clark, Q. Zhou, R. Benítez, L. Hove-Madsen, and S. Chen, "Disrupting Ca2+ waves, but not ca2+ leak, protects against cardiac arrhythmia," *Heart, Lung and Circulation*, vol. 21, 2012. 60th Annual scientific meeting of the Cardiac Society of Australia and New Zealand (CSANZ); AUG 16-19, 2012; Brisbane, QLD, AUSTRALIA.

[185] A. Herraiz-Martínez, J. Alvarez-Garcia, A. Llach, C. E. Molina, J. Fernandes, A. Ferrero-Gregori, C. Rodríguez, A. Vallmitjana, R. Benítez, R. Padró, J. Martínez-González, J. Cinca, and L. Hove-Madsen, "Ageing is associated with deterioration of calcium homeostasis in isolated human right atrial myocytes," *CARDIOVASCULAR RESEARCH*, vol. 106, pp. 76–86, APR 2015.

[186] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognition Letters*, vol. 27, pp. 861–874, 2006.

[187] A. Vallmitjana, M. Barriga, Z. Nenadic, A. Llach, E. Alvarez-Lacalle, L. Hove-Madsen, and R. Benítez, "Identification of intracellular calcium dynamics in stimulated cardiomyocytes," *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 68–71, SEP 2010. 32nd Annual International Conference of the IEEE Engineering-in-Medicine-and-Biology-Society; AUG 30-SEP 04, 2010; Buenos Aires, ARGENTINA.

[188] L. Breiman, "Random forests," *Machine Learning*, vol. 45, pp. 5–32, oct 2001.

[189] I. T. Jolliffe, *Principal Component Analysis*. Springer Series in Statistics, 2nd ed., 1986.

[190] C. E. Molina, A. Llach, A. Herraiz-Martinez, C. Tarifa, M. Barriga, R. F. Wiegerinck, J. Fernandes, N. Cabello, A. Vallmitjana, R. Benítez, J. Montiel, J. Cinca, and L. Hove-Madsen, "Prevention of adenosine A2A receptor activation diminishes beat-to-beat alternation in human atrial myocytes," *BASIC RESEARCH IN CARDIOLOGY*, vol. 111, p. 5, JAN 2016.

[191] X. Zhong, B. Sun, A. Vallmitjana, W. Gou, R. Wang, A. Guo, H. Duff, A. Gillis, L. S. Song, L. Hove-Madsen, R. Benitez, and S. R. W. Chen, "Suppression of Ryanodine Receptor Function Prolongs Ca2+ Release Refractoriness and Promotes Cardiac Alternans in Intact Hearts," *BIOCHEMICAL JOURNAL*, vol. 473, pp. 3951–3964, OCT 2016.

[192] L. M. Delbridge and K. P. Roos, "Optical methods to evaluate the contractile function of unloaded isolated cardiac myocytes.," *Journal of Molecular and Cellular Cardiology*, vol. 29, no. 1, pp. 11–25, 1997.

[193] C. Bazan, D. Torres Barba, P. Blomgren, and P. Paolini, "Image processing techniques for assessing contractility in isolated adult cardiac myocytes.," *International journal of biomedical imaging*, vol. 2009, p. 11, jan 2009.

[194] B. D. Lucas and T. Kanade, "An Iterative Image Registration Technique with an Application to Stereo Vision," *Proceedings of International Conference on Artificial Intelligence*, pp. 674–679, 1981.

[195] B. Horn and B. Schunck, ""Determining optical flow": A retrospective," *Artificial Intelligence*, vol. 59, no. 1-2, pp. 81–87, 1993.

[196] A. Vallmitjana, M. Barriga, L. Hove-Madsen, and R. Benítez, "Multilevel analysis of calcium dynamics in stimulated cultures of cardiomyocytes," pp. 6514–6517, July 2013. 2013 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Osaka, JAPAN.

[197] A. Vallmitjana, L. Hove-Madsen, and R. Benítez, "Método implementado por ordenador para caracterización dinámica de celulas en cultivos celulares y programas informáticos para llevar a cabo el método." Spanish patent number P201330974, sent 28 JUN 2013, published 14 JAN 2016.

[198] M. Barriga, C. Molina, A. Llach, N. Cabello, A. Vallmitjana, R. Benítez, J. Padro, J. Cinca, and L. Hove-Madsen, "Adenosine A2A receptor activation decreases beat-to-beat stability of intracellular calcium transients and their propagation in atrial myocytes," FEB 25-29 2012. 56th Annual Meeting of the Biophysical Society, San Diego, CA, US.

[199] R. Benítez, A. Vallmitjana, and L. Hove-Madsen, "A method to characterize calcium activity in stimulated cultures of cardiac myocytes," *Biophysical journal*, vol. 104, pp. 607A–607A, JAN 29 2013. 57th Annual Meeting of the Biophysical-Society; FEB 02-06, 2013; Philadelphia, PA, US.

[200] M. Barriga, R. Cal, N. Cabello, A. Llach, A. Vallmitjana, R. Benítez, L. Badimon, J. Cinca, V. Llorente-Cortes, and L. Hove-Madsen, "Low density lipoproteins promote unstable calcium handling accompanied by reduced SERCA2 and connexin-40 expression in cardiomyocytes," *PLOS ONE*, vol. 8, p. e58128, MAR 2013.

[201] Q. McNemar, "Note on the sampling error of the difference between correlated proportions or percentages ," *Psychometrika*, vol. 12, pp. 153–157, 1947.

[202] R. A. Fisher, "On the Interpretation of chi2 from Contingency Tables, and the Calculation of P," *Journal of the Royal Statistical Society*, pp. 87–94, 1922.

[203] A. Vallmitjana, C. Nolla, A. Herraiz, L. Hove-Madsen, and R. Benítez, "Multiscale detection of intracellular protein structures from microscope fluorescence images," APR 2015. ISBI - IEEE International Symposium on Biomedical Imaging, New York, US.

[204] X. Marimon, A. Vallmitjana, P. Jones, and R. Benítez, "A new user-friendly tool for localizing spots in fluorescence microscopy images," AUG 2015. 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Milan, ITALY.

[205] C. Tarifa, A. Herraiz-Martínez, A. Vallmitjana, S. A. Serra, D. Franco, R. Benítez, and L. Hove-Madsen, "Atrial spceific pitx2 insufficiencyincreases the frequency of calcium sparks, waves, and after-depolarizations in mouse atrial myocytes," *Biophysical journal*, vol. 112, p. 400a, Feb 14 2017. 61st Annual Meeting of the Biophysical-Society; FEB 11-15, 2017;New Orleans, LA, US.

[206] A. Herraiz-Martínez, A. Llach, C. Tarifa, E. Lozano-Velasco, S. A. Serra, J. Gandía, A. Vallmitjana, E. V. R. de Castroviejo, R. Benítez, A. Aranega, C. Muñoz-Guijosa, D. Franco, J. Cinca, , and L. Hove-Madsen, "The 4q25 variant rs13143308T links risk of atrial fibrillation to defective pitx2c and calcium homeostasis," 2017.

[207] R. Benítez, A. Vallmitjana, L. Hove-Madsen, and P. P. Jones, "Detection and classification of spontaneous calcium release events in cardiac myocytes," FEB 25-29 2012. 56th Annual Meeting of the Biophysical Society, San Diego, CA, US.

[208] A. Vallmitjana, E. Alvarez-LAcalle, L. Hove-Madsen, and R. Benítez, "Automatic method for the detection and characterization of calcium activity in cardiac cells," MAY 5 2012. ISBI - IEEE International Symposium on Biomedical Imaging, Barcelona, SPAIN.

[209] J. Rivera-Torres, A. Llach, G. Guzmán-Martínez, C. Calvo, R. Caballero, C. González-Gómez, L. J. Jiménez-Borreguero, J. A. Guadix, F. G. Osorio, C. Lopez-Otin, A. Herraiz-Martínez, N. Cabello, A. Vallmitjana, R. Benitez, L. B. Gordon, J. Jalife, J. M. Pérez-Pomares, J. Tamargo, E. Delpon, L. Hove-Madsen, D. Filgueiras-Rama, and V. Andres, "Cardiac electrical defects in progeroid mice and Hutchinson-Gilford progeria syndrome patients with nuclear lamina alterations," *PROCEEDINGS OF THE NATIONAL ACADEMY OF SCIENCES*, vol. 113, pp. E7250–E7259, OCT 2016.