Towards improving single-cell segmentation in heterogeneous configurations of cardiomyocyte networks *

Fabio Caraffini[1][0000−0001−9199−7368], Hassan Eshkiki[1][0000−0001−7795−453X], Mostafa Mohammadpour[2][0000−0002−9064−6504], Nikol Sullo[3][0000−0002−9885−5474], and Christopher H. George[3][0000−0001−9852−1135]

[1] Department of Computer Science, Swansea University, Swansea SA1 8EN, UK
{fabio.caraffini,h.g.eshkiki}@swansea.ac.uk
[2] Department of Computational Perception, Johannes Kepler University, Linz, Austria
mohammadpour@gtec.at
[3] Swansea University Medical School, Faculty of Medicine, Health and Life Sciences, Swansea SA2 8PP, UK
{nikol.sullo,christopher.george}@swansea.ac.uk

Abstract. To explore the formation and deterioration of cellular networks, we develop systems powered by Artificial Intelligence (AI) that accurately distinguish and quantify the differential configuration of cells in those networks (i.e. single cells, multicellular aggregates) as an initial proof-of-concept approach. We use image data acquired from self-organised cardiac cell networks formed in vitro which are difficult to segment using conventional methods. We used two data pre-processing approaches prior to the application of four segmentation algorithms (including two newly generated configurations of the Cellpose algorithm) for a total of eight segmentation pipelines. We demonstrate the effectiveness of a transfer learning capability in improving the accuracy of Cellpose in identifying discrete cells within complex (heterogeneous) cardiac cell network configurations. Our $\texttt{Cellpose}_3^{p_1}$ segmentation pipeline displays an F1-Score of 82.34%, a precision of 88.52% and an accuracy of 87.84%. Furthermore, in addition to our new method performing best in its ability to detect discrete cells in each network, it also avoided the problem of erroneously identifying cell boundaries in large multicellular aggregates. This preliminary work shows the feasibility of describing the physical and functional properties of cellular networks using accurate indices of cellular arrangement and heterogeneity.

Keywords: Cell Network· Cardiomyocytes, Segmentation· Cellpose· AI

* Supported by the Morgan Advanced Studies Institute, Wales, UK; the National Cardiovascular Research Network (funded by Health and Care Research Wales); the British Heart Foundation.

1 Introduction

AI has revolutionised our daily lives by serving as a central component in numerous contexts and fields, including the healthcare industry. There is increasing interest in AI to advance diagnosis, treatment, and patient care [7, 28]. Despite the new challenges and issues related to its use in these contexts [28], AI is making significant progress in healthcare applications by detecting patterns in large datasets in a way that humans alone cannot process or process efficiently. Machine Learning algorithms (ML), natural language processing techniques, and computer vision have greatly advanced the extraction of information from large amounts of medical data and help healthcare practitioners in decision-making processes. Typical applications include medical imaging (MRI, X-rays, etc.) [28], drug discovery [17], and electronic health record management [21] to name a few.

In this article, we take a first step toward understanding the formation and deterioration of cellular networks [27] with an AI-driven approach. We focus on cardiac cell networks as an interesting case study since, under experimental conditions in vitro, these networks comprise dynamic configurations of multi-cellular aggregations that arise from physical and functional coupling of individual cells (Fig. 1). At any given point in time, the temporal synchronisation of events across the particular network (e.g., Ca^{2+} signaling) is determined by the organisation and configuration of single cells with distinct boundaries (Fig. 1b, green regions) and those larger structures where it is not possible to define cellular boundaries (Fig. 1b, blue regions). It is therefore important that methods are developed to enable the accurate discrimination and quantification of network regions populated by discrete (boundarised) cells and larger multicellular structures. This initial step is key to understanding the influence and behaviour of cellular configurations in spatially organised cellular networks, which could lead to new knowledge on the link between cardiac cell network deterioration and the progression of cardiovascular disease [12, 27]. This is an unexplored domain that is ideally suited for the application of new AI-driven approaches.

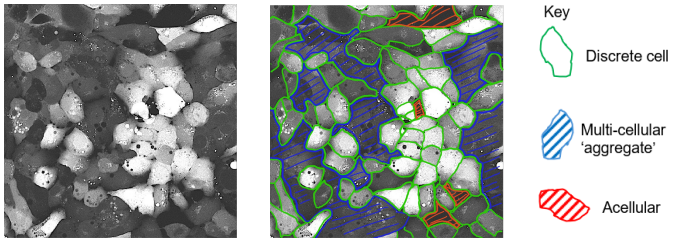

(a) (b)

Fig. 1: An original greyscaled image (a) and its manually segregated image (b).

In this proof-of-concept study, we focus on improving the detection of discrete single-cells. The remainder of this paper is structured as follows. Section 2

provides background information on existing cell segmentation systems; Section 3 describes the dataset and introduces the methods used to develop and assess the proposed segmentation pipelines to visualise and quantify the physical properties of discrete cells identified; Section 4 describes the experimental setup and reports on the validation results; Section 5 provides a biological perspective on the information acquired informed by our extensive experience in cardiac cell network research; Section 6 concludes this work.

2 Background

The primary goal in 2D cell segmentation is to identify the boundaries that demark cell edges. The expected result is a binary-valued mask showing a region of cells. In traditional computer vision, it is common to apply thresholding to convert a greyscale image to a binary image before segmenting it. This requires setting an intensity value, i.e. the threshold, that discriminates cells from the background by assigning only two possible values to pixels, whether their values are greater or smaller than the threshold. Several variants have been proposed to perform this task optimally, as the wrong choice of threshold value can lead to erroneous identification that introduces analytical artefacts. The Otsu method [18] is an iterative thresholding method that seeks an optimal threshold value by minimising intraclass variance, that is, the weighted variance between foreground and background, and appears to be suitable for white blood cell segmentation [24]. On the contrary, grid search thresholding compares the results of prefixed threshold values with the true label and selects the one that shows the best performance, e.g., in terms of F1-Score [29]. Active contours, see [1], is another 'energy-based' model for cell detection and segmentation that iteratively minimise an energy function while deforming a curve to fit the cell boundary.

In the last decade, Deep Learning (DL) has emerged as a new area of ML that takes advantage of multiple layers of non-linear information processing units for (supervised or unsupervised) feature extraction, pattern analysis, and classification. A Convolutional Neural Network (CNN) is a DL approach suitable for classification problems and image manipulation. Recently, CNNs have been adopted for semantic-based image segmentation tasks by exploiting their classification capabilities to binarily label each pixel of an image and thus segment it [4, 5]. Initially, CNNs were used in patch-based methods in biomedical image processing, classifying pixels based on their neighbourhood attributes to address segmentation challenges in, e.g., neural cell imagery. Methods such as the 'multi-task learning with a similarity interface' [22] were proposed to make them more suitable for segmentation tasks in microscopy images.

The U-net model proposed in [23] is another DL algorithm that features up- and down-sampling layers and the corresponding skip connections and is well known for its efficiency in segmentation tasks. U-net processes the entire image and directly generates a segmentation map that takes advantage of all contextual information from the image. This method offers advantages over patch-based CNNs. Since its publication, the literature has been populated with variants of

U-net, such as the V-net proposed in [16]. The latter improves upon U-net by including 3D convolution units that make it suitable for 3D segmentation. More complex structures are, e.g., those in [15], where a fully convolutional network is used to enable end-to-end semantic segmentation at the pixel level, and the M-Net model presented in [8], which features a multiscale input layer, a U-shaped convolutional network, a side output layer, and a multi-label loss function. An attentive cell instance segmentation method that uses the joint action of U-net and a single-shot multi-box detector is available in [31], while a versatile algorithm suitable for multiple tasks, including target classification, detection, semantic segmentation, instance segmentation, and recognition of human pose, among others, is known as the mask recurrent-CNN model [14].

Cellpose [26], is a widely used versatile DL algorithm with a U-Net-style architecture and residual blocks for cell and nucleus segmentation that was designed to process various types of images without requiring extensive training data or parameter tuning. Cellpose stands out for its adaptability (even to 3D images without needing specific 3D-labelled data) and ease of use. Cellpose developers trained this model using a variety of images sourced from the Internet from different datasets, including images of red and white blood cells, plasma cells, hanseniaspora, and animal tissue cells, by searching for keywords such as 'cytoplasm', 'cellular microscopy,' and 'fluorescent cells'. This vast training dataset gives Cellpose broad applicability and it is expected to work well even with noise, excessively bright fluorescence, insufficient brightness, or suboptimal stain quality. For the same reasons, it represents an obvious choice for many healthcare professionals in the analysis of clinical imaging data. However, despite its numerous advantages and rich featureset, Cellpose was developed to detect discrete cellular units clearly demarked by boundaries. As described above and shown explicitly in Fig. 1, these discrete cells are not the only component of physically and functionally coupled cardiac cell networks *in vitro*. The utility of cellular identification tools in this context needs to be rigorously evaluated, and in this preliminary work, we report on our findings to date.

3 Methodology

Conventional ML and complex DL approaches have strengths and weaknesses. The former methods generally exhibit inferior accuracy in segmenting microscopy images with complex backgrounds, especially in the presence of very small, atypically 'bright' elements; they are difficult to design; they face difficulties in handling extensive datasets. However, they have less demand for large training data; they are more explainable and computationally cheaper than the latter. Most CNNs typically require a significantly larger amount of manually-labelled training data and involve more computations than traditional ML approaches [2], all in search of superior performance.

We address the problem of segmenting heterogeneous cell configurations comprising 2D cardiac cell networks *in vitro* which comprise our foundational dataset. We focus on the application of the Cellpose algorithm, recognising it as

a state-of-the-art platform for cellular identification/quantification and due to its promising results in semantic segmentation [26, 19, 25].

We present the results obtained with the three best algorithmic configurations we found. We refer to them as $\texttt{Cellpose}_i$, with $i \in \{1, 2, 3\}$, and compare among themselves with another widely used segmentation model. We present two data preparation methods, which we apply to all models under study, thus obtaining eight possible segmentation strategies following the pipeline in Fig. 2. More details are provided in the following sections.

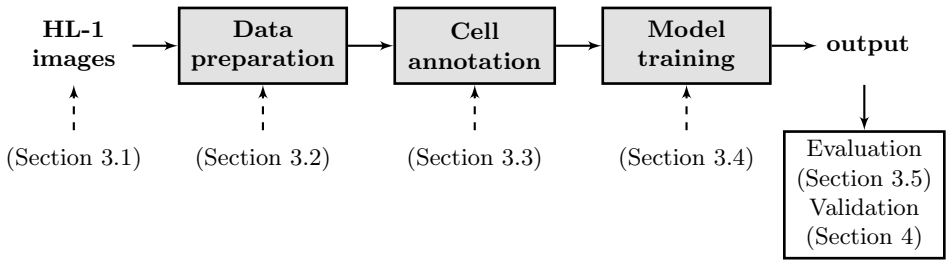

Fig. 2: General segmentation pipeline.

3.1 The HL-1 cardiac cell network imaging dataset

HL-1 cardiomyocytes are an immortalised mouse atrial-derived cell line that has been extensively characterised by us [6, 9, 3, 10]. Since HL-1 cells retain a remarkable proliferative capacity, the cellular density and interactions formed by these cells in vitro reflect 'random' self-organised network configurations. We analyse an image dataset comprising 92 HL-1 cell networks in which cells had been loaded with the fluorescent calcium reporter dye, fluo-3. The intense fluorescence of fluo-3 retained inside cells allows for a robust assessment of 'single' cellular boundaries although its distribution across all cells in the network is heterogeneous (Fig. 1a) which can make discrete cellular identification difficult. Each network was visualised using 180 images collected in 60 seconds (i.e., a frame capture rate of $0.3Hz$). Each image had dimensions of 512×512 pixels and defined a physical area of $61,504\mu m^2$.

3.2 Data preparation

To perform segmentation, it is required to produce a composite image from the image stack. We do this by 1) isolating frames for greyscale conversion and, if necessary, preprocessing them to improve the quality of the segmented result (as indicated in the next paragraph); 2) stacking them to compute their average; and 3) scaling the intensities of the average composite image within $0 - 255$[4].

[4] This value is imposed by the employed acquisition system.

After trying multiple approaches for Step 1), we empirically found two simple suitable methods that produced satisfactory results. The first approach is to use the greyscaled images and proceed directly to Step 2) (no further processing is needed). In the second method, histogram equalisation is performed on the greyscaled images. This is a technique used in image processing to enhance the contrast of an image by redistributing pixel intensities to cover the entire dynamic range evenly. This process helps to improve the visual appearance of images by making details more visible in both dark and bright areas [30]. Steps 2) and 3) are self-explanatory.

Note that averaging the images is a simple but beneficial approach, which helps filter out noise from (unprocessed) images. As can be seen in Fig. 3a, the unprocessed source experimental images are noisy. However, see Fig.s 3b and 3c, the effect of noise is mitigated after averaging the frames.

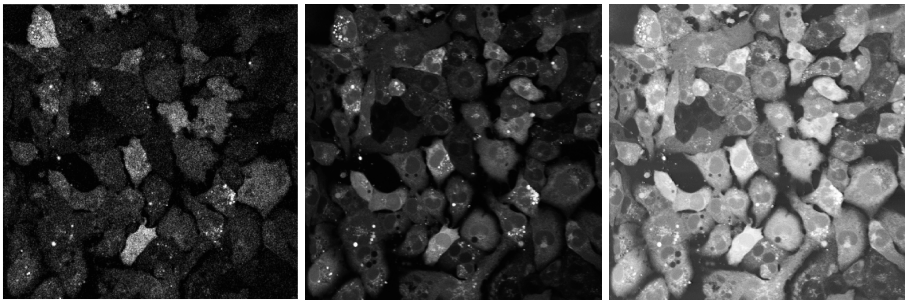

(a) Frame. (b) Greyscale (c) Histogram equalisation.

Fig. 3: Data (image) preparation - an example of a raw single-image frame (a) and its greyscale version without (b) and with (c) histogram equalisation.

For convenience, we refer to the first data preparation pipeline (without equalisation) as p_1 and to the second (with equalisation) as p_2.

3.3 Cell Annotation

The Cellpose Graphical User Interface (GUI) allows drawing Ground-Truth (GT) masks. We used it to label individual cells, as shown in Fig. 4.

In the present work, we focus exclusively on the segmentation/annotation of those cells with a clearly visible boundary (Fig. 1b, green outlines). All 92 composite images previously obtained following the processing described in 3.2 are annotated, and the corresponding masks are saved for the training process.

3.4 Segmentation algorithms and training

We employ Cellpose in three different configurations.

The first, $\texttt{Cellpose}_1$, simply refers to the built-in pre-trained Cyto2 model

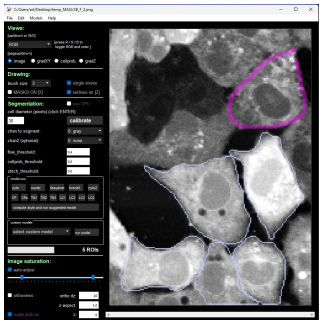

Fig. 4: Using Cellpose GUI to annotate the composite cellular network images.

available from the CellPose repository [26].

`Cellpose`$_2$ is instead obtained by retraining the Cellpose architecture from scratch with our dataset. This allows us to see the performance when the model is highly specialised in HL-1 cells.

`Cellpose`$_3$ is obtained by fine-tuning a pre-trained Cellpose model (i.e., `Cellpose`$_1$). In summary, we used existing Cyto2 weights to train a new model with our HL-1 data, thus adopting a particular kind of transfer learning [13].

We randomly divide the dataset into training and testing sets using a division ratio of 60 / 40%. With a total of 92 cell images available, 55 of these images are allocated to the training set and 37 to the test set, ensuring a robust evaluation of the model's performance on unseen data. The algorithms require image masks for training, and we use our annotations for this purpose. Each model is trained for 100 epochs. For optimal training performance and stability, all three versions of Cellpose are empirically optimised with the best configuration featuring a ReLu activation function, a learning rate of 10^{-3}, and a batch size of 8.

We also produce segmentation results with `StarDist` [25] for comparison. This is another widely used model based on the CNN architecture that detects individual cells by simultaneously generating a distance map and an object probability map. The success of this algorithm is due to the joint use of the maps that enables a precise delineation of cell boundaries for cells with irregular shapes and complex arrangements. `StarDist` represents cell boundaries using star-convex polygons, offering flexibility in capturing various cellular morphologies. This model also requires annotated data for the training data, typically comprising images paired with manually segmented cell outlines, which facilitates the network's ability to predict cell segmentations accurately. We used a `StarDist` pre-trained model on a subset of 497 fluorescence microscopy images.

3.5 Evaluation metrics

We use the five established evaluation metrics described in Table 1, which are calculated in terms of 'pixel-wise' True Positive (TP), False Positive (FP), False Negative (FN), and True Negative (TN), in the context of segmentation [11, 20].

Table 1: Performance Metrics in Image Segmentation

Metric	Description	Formula
Accuracy	Proportion of correctly classified samples	$\frac{TP+TN}{TP+TN+FP+FN}$
Precision	Proportion of TP predictions among all positive predictions	$\frac{TP}{TP+FP}$
Recall	Proportion of TP predictions among all actual positives	$\frac{TP}{TP+FN}$
F1-Score	Measure of overlap between predicted and GT regions	$\frac{2\times TP}{2\times TP+FP+FN}$
Intersection over Union (IoU)	Measure of overlap between predicted and GT regions	$\frac{TP}{TP+FP+FN}$

Note that the accuracy metric is often also commonly misused and overemphasised. F1-Score (also known as Dice similarity) can be of greater significance in segmentation problems, as it allows us to understand alignments with the GT, similarly, for IoU (also known as the Jaccard similarity index).

4 Validation results

Using the p_1 and p_2 data preparation methods outlined in Section 3.2, and applying the four models described in Section 3.4, we obtain eight algorithmic setups for the segmentation task. Some graphical results (validation) are shown in Fig. 5 and 6 for the preparation methods p_1 and p_2, respectively.

It can be immediately noticed that the original Cellpose algorithm, i.e., $\texttt{Cellpose}_1$, presents challenges in accurately distinguishing certain cell borders, particularly irregular ones, and when cellular density is visibly higher, leading to difficulties in correctly detecting cells. Furthermore, it occasionally identifies two separate cells instead of a single cell, which adds to its limitations in precise cell detection (Fig. 5b). $\texttt{Cellpose}_2$, see Fig. 5c, shows improved detection capabilities in cases with non-discernible boundaries compared to $\texttt{Cellpose}_1$, but suffers from undersegmentation in certain regions. On visual inspection, $\texttt{Cellpose}_3$ appears to be more accurate than the other Cellpose setups (Fig. 5d) and than `StarDist` (Fig. 5e), which seems to perform worse than all other models.

Numerically, the segmentation results are validated by the metrics in Table 1 computed throughout the test set, which comprises 37 images. Each evaluation metric is calculated for each cell, and then the average value (avg) and standard deviation (std) are calculated to provide an overall assessment of the performance of the model. Validation results are reported in Table 2, where the best value (i.e., highest avg and smallest std) per model and for each evaluation metric is reported in bold (except for cases with nonsignificant differences). The asterisk shows the algorithmic setup that performs best per metric.

Interestingly, the simpler p_1 data preparation method leads to better performance than p_2 in most cases. $\texttt{Cellpose}_3$ always outperforms the other methods

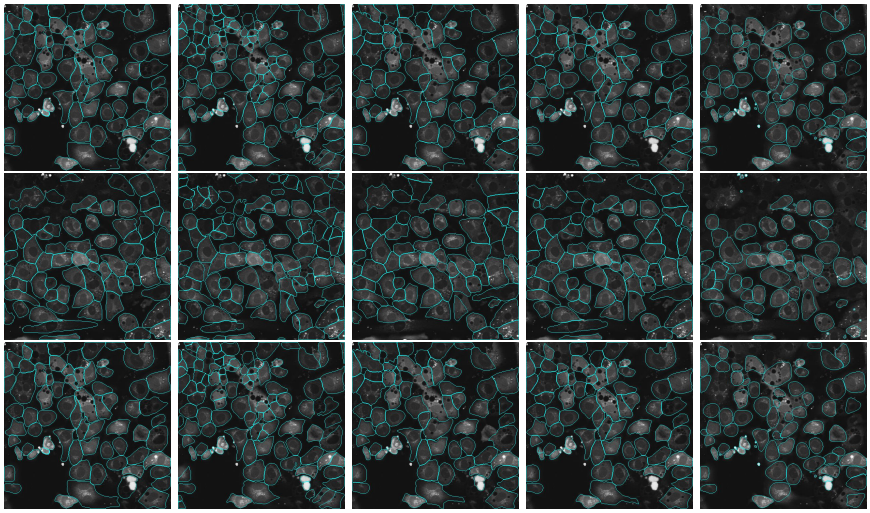

(a) GT (b) $\texttt{Cellpose}_1$ (c) $\texttt{Cellpose}_2$ (d) $\texttt{Cellpose}_3$ (e) `StarDist`

Fig. 5: Segmentation results with the p_1 image preparation strategy.

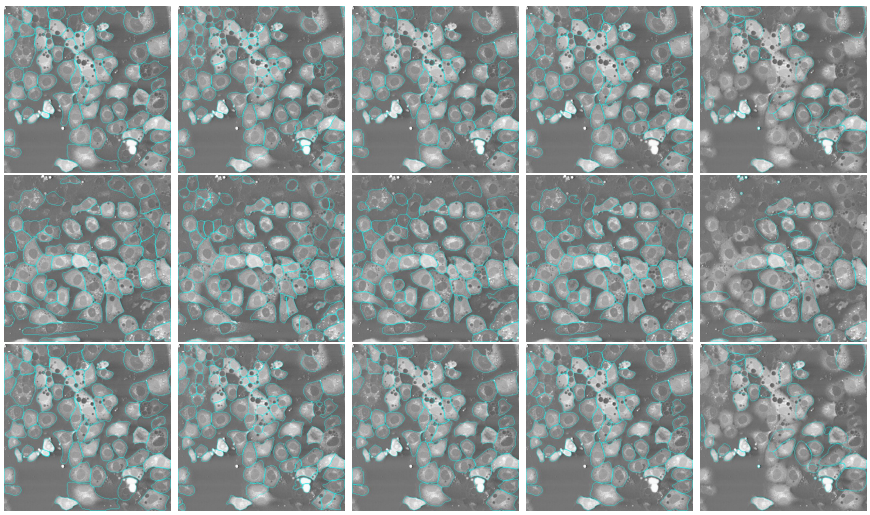

(a) GT (b) $\texttt{Cellpose}_1$ (c) $\texttt{Cellpose}_2$ (d) $\texttt{Cellpose}_3$ (e) `StarDist`

Fig. 6: Segmentation results with the p_2 image preparation strategy.

except in terms of recall, where $\texttt{Cellpose}_2$ displays marginally superior performance (of 1.19%). The results for $\texttt{Cellpose}_3$ are satisfactory, with an average precision of 87.84% with p_1 and an average precision of 87.84% with p_2. The IoU and F1-Score are also good, with an average of 71.98% and 82.34%, respectively.

Table 2: Models evaluation results.

Algorithmic setup		Evaluation metrics				
		IoU (%)	**F1-Score**(%)	**Precision**(%)	**Recall**(%)	**Accuracy**(%)
$\texttt{Cellpose}_1$	p_1	$\mathbf{55.44 \pm 22.58}$	$\mathbf{67.94 \pm 23.67}$	71.21 ± 24.60	$\mathbf{72.96 \pm 24.12}$	$\mathbf{78.34 \pm 9.77}$
	p_2	49.15 ± 23.10	62.11 ± 24.84	71.24 ± 24.08	63.66 ± 27.80	75.78 ± 10.58
$\texttt{Cellpose}_2$	p_1	$\mathbf{68.43 \pm 17.91}$	$\mathbf{79.42 \pm 17.63}$	79.29 ± 19.16	$^*\mathbf{82.43 \pm 17.90}$	$\mathbf{85.53 \pm 6.93}$
	p_2	65.82 ± 18.77	77.40 ± 17.94	$\mathbf{86.05 \pm 16.99}$	72.38 ± 21.16	84.56 ± 8.16
$\texttt{Cellpose}_3$	p_1	$^*\mathbf{71.98 \pm 16.25}$	$^*\mathbf{82.34 \pm 14.76}$	$^*\mathbf{88.52 \pm 9.45}$	80.24 ± 17.09	$^*\mathbf{87.84 \pm 5.69}$
	p_2	71.15 ± 16.62	81.79 ± 14.18	88.41 ± 11.88	79.16 ± 16.71	87.21 ± 6.38
$\texttt{StarDist}$	p_1	$\mathbf{43.75 \pm 16.42}$	$\mathbf{58.86 \pm 17.84}$	$\mathbf{81.74 \pm 23.46}$	$\mathbf{50.01 \pm 19.80}$	$\mathbf{73.19 \pm 13.48}$
	p_2	40.21 ± 20.84	53.94 ± 23.27	67.79 ± 22.57	49.70 ± 26.60	70.62 ± 12.22

The `StarDist` algorithm is always outperformed by the three Cellpose configurations, regardless of the preparation of input data. This algorithm displays the best results on round-shaped cells, which is not always the case. In addition, roundness is often a consequence of cell death, which introduces the risk of defining networks by the extent of death. However, the latter appears to be more 'precise' than $\texttt{Cellpose}_1$, as it is $\texttt{StarDist}^{p_1}$ than $\texttt{Cellpose}_2^{p_2}$ (the superscript denotes the data preparation method to simplify notation). Due to the high values of std, it can be argued that $\texttt{StarDist}^{p_1}$ is similar to $\texttt{Cellpose}_1^{p_2}$ in terms of IoU and F1-Score.

Hence, the segmentation pipeline that stands out is $\texttt{Cellpose}_3^{p_1}$, despite histogram equalisation visually resulting in clearer images, and graphical results may also favour p_2 for this reason. This suggests that $\texttt{Cellpose}_3^{p_1}$ detects features that we cannot easily see and which equalisation cannot reveal. It is important to note that the latter enhances contrast when there is a limited range of intensity values. However, in cases where spatial correlation is more significant than intensity, it may result in unrealistic effects that alter the shape of cells.

From the produced segmented images, we extrapolate[5] the area and eccentricity of each identified discrete cell. The first parameter is calculated as the total pixel count within a defined cellular boundary, while the second, indicating cell elongation or roundness, is calculated as the ratio of the best-fit ellipse's focal distance to the cell's shape. These are discussed in the next section.

5 Cellular Network Analysis

Extending some of the qualitative information provided by model validation (Section 3.5), we systematically evaluate the quantitative outputs of each model relating to three key metrics (cell number, cell area, and cell eccentricity) in 30 randomly selected networks. We find that `StarDist` is the algorithm most influenced by the preprocessing mode and performs poorly on all metrics when p_1 and p_2 were directly compared ($R^2 < 0.35$ in all instances) (Fig. 7).

[5] This is done in Python with `measure.regionprops` from `skimage`.

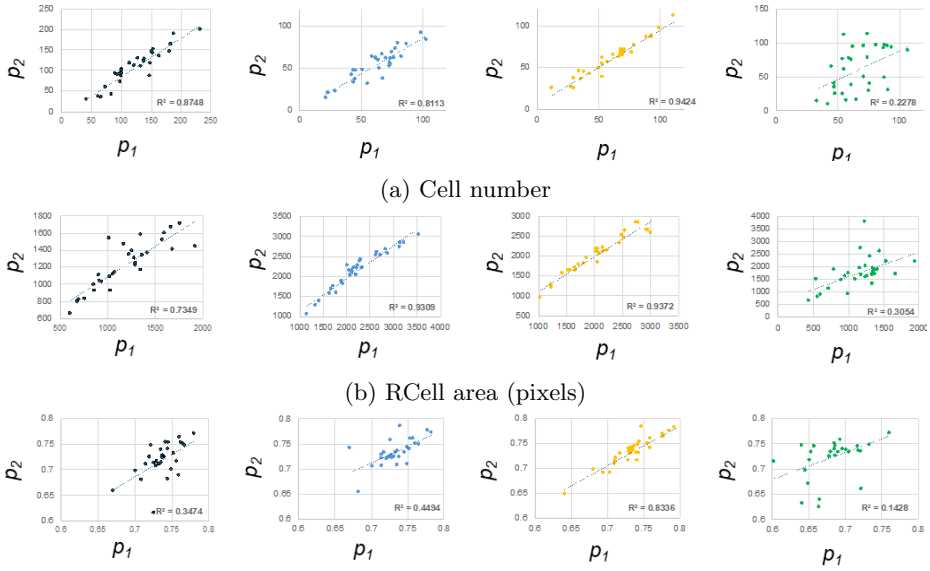

(a) Cell number

(b) RCell area (pixels)

(c) Cell eccentricity (normalised)

Fig. 7: Regression analysis of 30 chosen networks on the effect of P_1 and P_2 for (left to right) `Cellpose`$_1$, `Cellpose`$_2$, `Cellpose`$_3$ and StarDist.

`Cellpose`$_2$ and `Cellpose`$_3$ are much less influenced by the preprocessing mode, and `Cellpose`$_3$ overall exhibits a better performance ($R^2 > 0.8$ for all metrics). Benchmarking the performance of all methods against `Cellpose`$_1$, we found that `Cellpose`$_2$ and `Cellpose`$_3$ identified many fewer discrete cells (Fig. 8) that are physically larger cells (Fig. 7).

However, upon closer scrutiny of the segmented composite images of each of the 30 networks, we found that while `Cellpose`$_1$ identifies more 'cells', this algorithm tended to erroneously identify many 'discrete' cells (i.e., nucleated structures) within the multi-cellular aggregates (Figure 1, blue regions) resulting in the gross over-estimation of the number of discrete cells in the network. `Cellpose`$_2$ and `Cellpose`$_3$ does not reproduce this error, and visual inspection of the segmentation done by these methods reveals that their identification of single (discretely boundarised) cells in the network was remarkably accurate. Preliminary data suggest `Cellpose`$_3$ is the best algorithm for our purpose. The data also show that `StarDist` identifies numerous elements that are too small to be plausibly considered as cells (i.e., have pixel dimensions < 100; Table 3) and this undesirable feature, together with other limitations (Fig. 7 and 8) highlights that `StarDist` is not useful for our purposes.

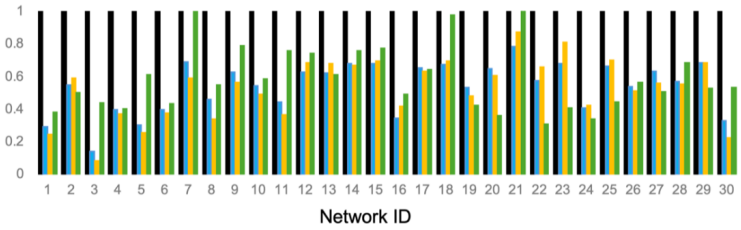

(a) Results with data preparation p_1

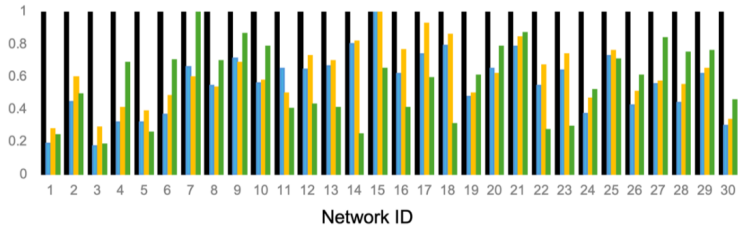

(b) Results with data preparation p_2

Fig. 8: Number of discrete cells segmented with $\texttt{Cellpose}_1$, $\texttt{Cellpose}_2$, $\texttt{Cellpose}_3$ and StarDist, normalised to $\texttt{Cellpose}_1$.

Table 3: Proportion of segregated non-cellular elements.

$\texttt{Cellpose}_1$		$\texttt{Cellpose}_2$		$\texttt{Cellpose}_3$		StarDist	
p_1	p_2	p_1	p_2	p_1	p_2	p_1	p_2
0.104%	0.213%	0.051%	0.117%	0.052%	0.053%	24.130%	10.597%

6 Conclusion

In this proof-of-concept work, we have shown the utility of a transfer learning algorithm to improve the accuracy of Cellpose to identify discrete cells in complex (heterogeneous) cardiac cell networks. This work should now be extended to the development of parallel methods that can quantify key features that describe the larger multicellular aggregates that are a hallmark feature of in vitro formed cardiac cell networks (e.g., HL-1 and induced pluripotent stem cell-derived cardiomyocyte networks). Pending this outcome, we will then be able to describe cardiac cell networks in terms of 'single cellularity' and 'multicellularity' and begin reconciling these indices with new information on the functional competency of the networks under test (e.g., normal 'healthy' networks or deteriorating 'disease-like' networks).

References

1. Bamford, P., Lovell, B.: Unsupervised cell nucleus segmentation with active contours. Signal Processing **71**(2), 203–213 (1998). https://doi.org/10.1016/S0165-

1684(98)00145-5
2. Caicedo, J.C., Roth, J., Goodman, A., Becker, T., Karhohs, K.W., Broisin, M., Molnar, C., McQuin, C., Singh, S., Theis, F.J., et al.: Evaluation of deep learning strategies for nucleus segmentation in fluorescence images. Cytometry Part A **95**(9), 952–965 (2019)
3. CH, G.: Arrhythmogenic mutation-linked defects in ryanodine receptor autoregulation reveal a novel mechanism of caˆ¡ 2+¿ release channel dysfunction. Circ Res **98**, 88–97 (2006)
4. Ciresan, D., Giusti, A., Gambardella, L., Schmidhuber, J.: Deep neural networks segment neuronal membranes in electron microscopy images. Advances in neural information processing systems **25** (2012)
5. Cireşan, D.C., Giusti, A., Gambardella, L.M., Schmidhuber, J.: Mitosis detection in breast cancer histology images with deep neural networks. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013: 16th International Conference, Nagoya, Japan, September 22-26, 2013, Proceedings, Part II 16. pp. 411–418. Springer (2013)
6. Claycomb, W.C., Lanson Jr, N.A., Stallworth, B.S., Egeland, D.B., Delcarpio, J.B., Bahinski, A., Izzo Jr, N.J.: Hl-1 cells: a cardiac muscle cell line that contracts and retains phenotypic characteristics of the adult cardiomyocyte. Proceedings of the National Academy of Sciences **95**(6), 2979–2984 (1998)
7. Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S.: Dermatologist-level classification of skin cancer with deep neural networks. nature **542**(7639), 115–118 (2017)
8. Fu, H., Cheng, J., Xu, Y., Wong, D.W.K., Liu, J., Cao, X.: Joint optic disc and cup segmentation based on multi-label deep network and polar transformation. IEEE transactions on medical imaging **37**(7), 1597–1605 (2018)
9. George, C.H., Higgs, G.V., Lai, F.A.: Ryanodine receptor mutations associated with stress-induced ventricular tachycardia mediate increased calcium release in stimulated cardiomyocytes. Circulation research **93**(6), 531–540 (2003)
10. George, C.H., Rogers, S.A., Bertrand, B.M., Tunwell, R.E., Thomas, N.L., Steele, D.S., Cox, E.V., Pepper, C., Hazeel, C.J., Claycomb, W.C., et al.: Alternative splicing of ryanodine receptors modulates cardiomyocyte ca2+ signaling and susceptibility to apoptosis. Circulation research **100**(6), 874–883 (2007)
11. Ghaznavi, A., Rychtáriková, R., Saberioon, M., Štys, D.: Cell segmentation from telecentric bright-field transmitted light microscopy images using a residual attention u-net: A case study on hela line. Computers in Biology and Medicine **147**, 105805 (2022)
12. Gintant, G.A., George, C.H.: Introduction to biological complexity as a missing link in drug discovery. Expert Opinion on Drug Discovery **13**(8), 753–763 (2018)
13. Hallou, A., Yevick, H.G., Dumitrascu, B., Uhlmann, V.: Deep learning for bioimage analysis in developmental biology. Development **148**(18), dev199616 (2021)
14. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)
15. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3431–3440 (2015)
16. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 fourth international conference on 3D vision (3DV). pp. 565–571. Ieee (2016)
17. Mock, M., Edavettal, S., Langmead, C., Russell, A.: Ai can help to speed up drug discovery—but only if we give it the right data. Nature **621**(7979), 467–470 (2023)

18. Otsu, N.: A threshold selection method from gray-level histograms. IEEE Transactions on Systems, Man, and Cybernetics **9**(1), 62–66 (1979). https://doi.org/10.1109/TSMC.1979.4310076
19. Pachitariu, M., Stringer, C.: Cellpose 2.0: how to train your own model. Nature methods **19**(12), 1634–1641 (2022)
20. Punn, N.S., Agarwal, S.: Modality specific u-net variants for biomedical image segmentation: a survey. Artificial Intelligence Review **55**(7), 5845–5889 (2022)
21. Rajkomar, A., Oren, E., Chen, K., Dai, A.M., Hajaj, N., Hardt, M., Liu, P.J., Liu, X., Marcus, J., Sun, M., et al.: Scalable and accurate deep learning with electronic health records. NPJ digital medicine **1**(1), 1–10 (2018)
22. Ramesh, N., Tasdizen, T.: Cell segmentation using a similarity interface with a multi-task convolutional neural network. IEEE journal of biomedical and health informatics **23**(4), 1457–1468 (2018)
23. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18. pp. 234–241. Springer (2015)
24. Salem, N., Sobhy, N.M., El Dosoky, M.: A comparative study of white blood cells segmentation using otsu threshold and watershed transformation. Journal of Biomedical Engineering and Medical Imaging **3**(3), 15 (2016)
25. Schmidt, U., Weigert, M., Broaddus, C., Myers, G.: Cell detection with star-convex polygons. In: Medical Image Computing and Computer Assisted Intervention - MICCAI 2018 - 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part II. pp. 265–273 (2018). https://doi.org/10.1007/978-3-030-00934-2_30
26. Stringer, C., Wang, T., Michaelos, M., Pachitariu, M.: Cellpose: a generalist algorithm for cellular segmentation. Nature methods **18**(1), 100–106 (2021)
27. Strogatz, S.H.: Exploring complex networks. nature **410**(6825), 268–276 (2001)
28. Topol, E.J.: High-performance medicine: the convergence of human and artificial intelligence. Nature medicine **25**(1), 44–56 (2019)
29. Vicar, T., Balvan, J., Jaros, J., Jug, F., Kolar, R., Masarik, M., Gumulec, J.: Cell segmentation methods for label-free contrast microscopy: review and comprehensive comparison. BMC bioinformatics **20**, 1–25 (2019)
30. Wang, X., Cheng, E., Burnett, I.S.: Improved (stem) cell segmentation with histogram matching image contrast enhancement. In: 2015 IEEE China Summit and International Conference on Signal and Information Processing (ChinaSIP). pp. 816–820. IEEE (2015)
31. Yi, J., Wu, P., Jiang, M., Huang, Q., Hoeppner, D.J., Metaxas, D.N.: Attentive neural cell instance segmentation. Medical image analysis **55**, 228–240 (2019)