



Semi-supervised 3D Neural Networks to Track iPS Cell Division in Label-free Phase Contrast Time Series Images¹

Adele Peskin[†]
NIST
Boulder, CO USA
adele.peskin@nist.gov

Joe Chalfoun
NIST
Gaithersburg, MD USA
joe.chalfoun@nist.gov

Michael Halter
NIST
Gaithersburg, MD USA
michael.halter@nist.gov

Anne Plant
NIST
Gaithersburg, MD USA
anne.plant@nist.gov

ABSTRACT

In order to predict cell population behavior, it is important to understand the dynamic characteristics of individual cells. Individual induced pluripotent stem (iPS) cells in colonies have been difficult to track over long times, both because segmentation is challenging due to close proximity of cells and because cell morphology at the time of cell division does not change dramatically in phase contrast images; image features do not provide sufficient discrimination for 2D neural network models of label-free images. However, these cells do not move significantly during division, and they display a distinct temporal pattern of morphologies. As a result, we can detect cell division with images overlaid in time. Using a combination of a 3D neural network applied over time-lapse data to find regions of cell division activity, followed by a 2D neural network for images in these selected regions to find individual dividing cells, we developed a robust detector of iPS cell division. We created an initial 3D neural network to find 3D image regions in (x,y,t) in which identified cell divisions occurred, then used semi-supervised training with additional stacks of images to create a more refined 3D model. These regions were then inferred with our 2D neural network to find the location and time immediately before cells divide when they contain two sets of chromatin, information needed to track the cells after division. False positives from the 3D inferred results were identified and removed with the addition of the 2D model. We successfully identified 37 of the 38 cell division events in our manually annotated test image stack, and specified the time and (x,y) location of each cell just before division within an accuracy of 10 pixels.

CCS CONCEPTS

• Applied computing -> Life and medical sciences -> Computational biology -> Imaging; Artificial Intelligence -> Learning;

KEYWORDS

Cell division, neural networks, semi-supervised learning

ACM Reference format:

Adele Peskin, Joe Chalfoun, Michael Halter and Anne Plant. 2022. Semi-supervised 3D Neural Network to Track iPS Cell Division in Label-free Phase Contrast Time Series Images. In 13th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics (ACM-BCB '22), August 7–10, 2022, Chicago, IL, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3535508.3545532>.

1 Introduction

Characterizing and predicting cell population behavior is critical to the efficiency and flexibility in manufacturing cell-based products. Characterization is challenging because individual cells within populations demonstrate heterogeneous and dynamic characteristics. Individual cells express their properties slightly differently and at different rates, and data quantifying the dynamic characteristics of large numbers of individual cells over time are needed for meaningful characterization. Image segmentation of individual induced pluripotent stem (iPS) cells can be difficult because they reside very close to one another in colonies. In addition, each individual cell must be tracked over time as it divides into two daughter cells, which is difficult because the morphology of these cells does not change significantly while they are dividing. Phase contrast images of iPS cells during cell division exhibit very different morphology than other stem cell types previously studied, such as the C3H10 mesenchymal stem cell line [1,2], in which the parent cell rounds up and then splits into two daughter cells. In this iPS cell line, the rounding of cells is not obvious (see Figure 1), and while features of metaphase appear in phase images, they are difficult to capture due to the high level of pixel noise.

We have previously created AI models that locate the nuclei of individual stem cells in phase contrast images of iPS cells, by training the models with cells containing a fluorescent nuclear protein [3, 4] to segment nuclei. However, these models cannot capture the nuclei of cells as they divide, and so tracking of cells is interrupted at cell division. Recognizing when a cell is undergoing division so that its daughter cells can continue to be tracked and linked to the parent cell is a critical challenge that still needs to be

¹ This contribution of NIST, an agency of the U.S. government, is not subject to copyright.

resolved for the automated tracking of individual stem cells within a colony.

Efficient automation is essential since we aim to collect many images of each of thousands of cells over long periods of time, resulting in volumes of data that are too large to rely on manual processing. In this study, our intention is to develop a reliable mitosis detector that enables tracking of label-free iPS cells in phase contrast images. While fluorescence probes have been used to study mitosis, exposure to the wavelengths and intensities of light needed to excite fluorescence can be damaging to cells, especially when such exposure at short time intervals and for long times is required [4].

We begin with a small set of manually annotated images, and then greatly expand our training set using semi-supervised learning. A large number of different techniques are used in the literature to create semi-supervised training data [5-9], which includes both labeled and unlabeled training data. Each method is formulated from the specific image features of each data set, either by using a student-teacher model [5], segmentation of image regions by methods other than neural networks [6], including clustering [7] and interactive cell segmentation [8], and by using generative adversarial networks [9].

Here we use a combination of 2D and 3D neural networks to create semi-supervised training data: a 3D neural network that locates regions of cell divisions, and then a 2D network trained with images found by the 3D model to locate dividing cells.

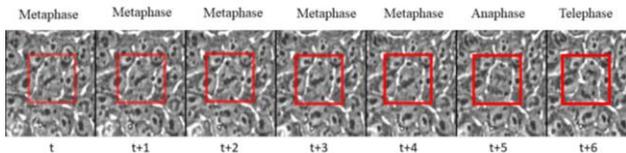


Figure 1: Seven time steps that show a single iPS cell (at the center of each image tile in the red boxes) in the course of cell division. A phase-dense line of chromatin seen in the first 5 frames is associated with the metaphase stage, which is followed by anaphase and telophase stages of cell division. The cell does not change shape during these time steps.

1 Methods

2.1 Culture and Image Acquisition Protocols

All images were acquired using a human iPS cell line in which Lamin B1 has been endogenously tagged with mEGFP (LamB1:mEGFP) using CRISPR/Cas9 technology generated at the Allen Institute for Cell Sciences (WTC-mEGFP-LMNB1-cl210), and was obtained from Coriell Institute for Medical Research (Catalog # AICS-0013, Camden, NJ). Cells were regularly maintained using complete mTeSR medium supplemented with Pen/Strep in six well tissue culture treated plates (TPP, Product # 92006, Switzerland) coated with Matrigel (hESC certified, from Corning). Generally, cells were passaged using Accutase when 70 to 80% confluent and re-plated at 100k to 200k cells per well.

Immediately prior to imaging, the cell culture media was aspirated, the cell culture plate was placed on the microscope stage (Ludl Electronic Products, Hawthorne, NY) and maintained at 37 °C in a custom built incubation chamber (Kairos Instruments, Pittsburgh, PA). Time lapse image capture was performed on a Zeiss 200M microscope (Carl Zeiss USA, Thornwood, NY) using a Zeiss 10 \times , 0.3NA objective (Zeiss part number 420341-9911- 000) and a CoolSNAP HQ2 CCD camera (Photometrics, Tucson, Arizona). Stage, filters and shutters were controlled with μ Manager1 open source software. The stage was programmed to move from field to field with an overlap of adjacent fields of 10%. A summary of the data sets used in this paper are given in Table 1. The time interval between images for data sets 1 and 2 was 125 s and the time interval between images in data set 3 was 160 s. The sample was exposed to light from a low-power LED (centered at 525 nm, Thorlabs, Newton, NJ) with Kohler aligned Zernike phase contrast optics. The illumination power was 26 μ W over a 10 ms exposure time. A spatial calibration target was used to determine that each pixel is equivalent to an area of 0.394 μ m².

Table 1: Data sets: Images were collected on two separate days, Day 1 and Day 2 below. Images from the first day were used to create 20 different image stacks in time which were used for model training. The smaller test set from Day 2 is used as our test data set.

Set		1	2	3
Day collected		1	1	2
Use		training	training	testing
Tile #		Tile 0	Tiles 1-19	
Used to train		Initial 3D; 2D models	Final 3D model	
Label		manual	automatic	manual
Tile Size		256x256 x 500	256x256 x 500	256x256 x 200
# cell divisions		46		38
# human identified		40		26
Interval between images		125 s	125 s	160 s

2.2 Manually annotated training data for our 3D and 2D models

A summary of the neural network models used in this study is given in Table 2. Initial training data for all models were created from data set 1 from Table 1, a stack of 256x256 pixel images over 500 time frames (256x256x500, in x,y,t). Manual annotations generated a list of 40 (x,y,t) positions at which a cell division was seen. The following steps (see Figure 2) were carried out to create training

data from these coordinates, which were used to train both our 3D and 2D models:

Table 2: Neural network models

#	Model	use	Training data	Training data set
1	Initial 3D model	Finds mitotic regions in 3D image data	manual	1
2	Final 3D model	Finds mitotic regions in 3D image data	automatic	2
3	2D sister chromatid model	Finds dividing cells with sister chromatids	manual	1
4	2D all chromatin model	Find all condensed chromosomes of division	manual	1

- At each (x,y,t) location in space and time corresponding to a cell division event, a stack of sub-images was processed. Each sub-image contained an individual dividing cell in (x,y) centered around (x,y) with 21 layers in time, 10 before and 10 after the labeled point of division. The size of the pixel tile was chosen to contain the entire cell as well as the two daughter cells. Figure 2a shows 11 of the 21 time steps for a dividing cell that was manually annotated.
- Manual annotations of dividing cells assigned a pixel value of 2 for the phase-dense chromatin, as shown in Figure 2, for 8-12 time steps per division, depending upon the number of steps for which dense chromatin was present. Nuclear pixels were given a pixel value of 1. Labeled nuclear pixels overlap with those found using our nuclear segmentation model [1,2].

2.3 Training the initial 3D model (Model 1 of Table 2)

The following steps were taken to train our initial model:

- Pixels corresponding to the nuclear mask in the $256 \times 256 \times 500$ image stack were labeled as class 1 using the previously published nuclear segmentation model [1,2].
- Class 1 and class 2 labels from the manually labeled image of each mitotic event (steps 1-2 above) were then added into the $256 \times 256 \times 500$ mask overlaying the class 1 nuclear pixels, shown in grey for a single slice in Figure 2c.
- Due to GPU memory constraints, the 500 time step image was sub-divided into tiles with a smaller number of time steps to feed into the network. Overlapping tiles of size $256 \times 256 \times 16$ were created at 4 time step intervals of 16 consecutive steps.
- The 3D tiles fed into the network went through an augmentation process. Images were randomly flipped in (x,y) and randomly rotated in (x,y) , and then a smaller tile size of $128 \times 128 \times 16$ was randomly selected to send into the network. Each batch of images (batch size=4, 128, 128, 16) is slightly

different and the dataset covers the entire $256 \times 256 \times 500$ stack.

- A 3D U-Net [10] network was then trained on the $128 \times 128 \times 16$ tiles to create Model 1 in Table 2. Full 3D convolutions were used in the model.

The 3D U-Net was run using Tensorflow python code. Each phase image in time was normalized individually. We use a z-score normalization that puts each image in the same pixel value scale, and also handles outliers. All pixel values in the normalized images are capped at ± 5 standard deviations. Weights are applied to the loss terms for the three classes (background, nuclei, chromatin), in the ratio of [1.0, 2.0, 40.0]; i.e., a very heavy weight is assigned to loss terms for the infrequent chromatin pixels. We used an Adam optimizer, and an original learning rate of $1.0e-4$. This rate is reduced by a factor of 10 for the first epoch and then replaced with the original value.

Mitotic events are obtained from the output of the 3D model as clusters of pixels labeled as chromatin (pixel value of 2). Interestingly, when we used the trained 3D model (x,y,t) to inference the original $256 \times 256 \times 500$ stack of phase images (data set 1), we found 6 additional mitotic events that were not included in the 40 that were manually identified, increasing our confidence that the model was performing well. We added the 6 new events, along with the original 40, to the training image stack and retrained the 3D model. We refer to this updated model as our initial 3D model (Model 1).

2.4 Training the 2D models (Models 3,4 of Table 2) for chromatin and chromatid classification

Given the limited amount of data, our 3D mitosis model was able to roughly identify regions in space and time where mitotic events take place (see Figure 3, part 2). This resolution is not sufficient for accurately tracking parent and daughter cells in time. To overcome this limitation, we include 2D models to refine the (x,y,t) estimate of cell division. (Starting out by training a 2D model to find the cell divisions was not successful because of the rarity of the events; events are not rare in the limited images found by the 3D model.) Training data for the 2D models were created in two different ways, as shown in Figure 2 to generate two separate models. For the first model, called the sister chromatid model (Model 3 in Table 2), pairs of adjacent areas of phase-dense chromatin (sister chromatids) were labeled in corresponding masks of 2D phase slices (as in Figure 2c, middle). For the second model, called the all-chromatin model (Model 4 in Table 2), we labeled both the sister chromatids and the corresponding dense line of metaphase chromatin that appeared for several preceding time steps (Figure 2c, bottom row). The output from both models were used because the sister chromatid model found more of the sister chromatids than the all-chromatin model. The 2D model training data is created with these steps:

- Training data for the 2D models have only 2 classes: dense chromatin as class 1, and the background as class 0. Section 2.2 and Figure 2c describe how the chromatin was labeled

for the 2D models. 128x128 2D phase images centered around each manually assigned dividing cell were used for training, covering 30 time steps centered at the dividing cell in time (the areas inferred by the 3D model cover a wider (x,y,t) range than the annotated data).

- Augmentation included image flipping, random rotation, and image blurring.
- Due to the relative sparsity of division events, the loss terms are weighted for the 2 classes in the ratio of 1 for the background and 20 for the class 1 chromatin.
- Both of the 2D models (3,4 in Table 2) are used in the next step to make semi-supervised training data.

2.5 Creation and testing of semi-supervised training data for the final 3D model (Model 2)

To create more 3D training data and improve generalizability without additional manual annotation, we developed a semi-supervised training data creation process, which is completely automated and uses both the 3D and 2D models, and some classical image processing (thresholding of phase images, discussed below). We created a new set of training data (see Figure 3, parts 2 and 3) using 19 additional 256x256x500 image tiles (Data set 2 of Table 1) with the following steps:

- Each new image stack is inferred with the 3D model 1 (Figure 3, part 2). This creates 3D stacks of class 2 (chromatin) pixels, representing the mitotic events. The 3D model finds the regions in space where mitotic events occur.
- To be used in the training data, class 2 pixel clusters of this 3D inference have to extend for at least 6 time steps and have a total area over all time steps greater than 1000 pixels (Figure 3, part 3A).
- From each 3D pixel cluster of step 2, we create a set of phase contrast images 128x128 in (x,y) and covering 30 time steps, centered at the center of the cluster in (x,y,t) (Figure 3, part 3B). We then use the 2D models (Models 3,4) to identify the dense chromatid material associated with a cell division, metaphase chromatin and individual sister chromatids in dividing cells.
- We use the test outlined in Figure 4 to determine if inferred chromatin pixel clusters belong to true mitotic events. The cells in all of our data sets do not move significantly while they are in metaphase, so we test an (x,y) stack over time to see if there is a stack of static condensed chromatin. To do this test, corresponding phase contrast tiles are thresholded to find their darkest phase-dense regions, below the 20th percentile intensity. We find the center value of the sister chromatid pixels in the inferred cell. Resulting dark pixel clusters (Figure 4b) found within 10 pixels of the sister chromatid (x,y) location are selected from previous time steps. (As more data is collected, we will verify the assumption that the cells do not move as they divide.)
- False positive events from the 3D model are identified when the 2D inferencing finds sister chromatids, but we cannot trace the sister chromatids in the thresholded phase image for at least 4 time steps.

- Chromatin pixels that are part of a timed event are then overlaid onto nuclear masks (Figure 3, part 3C) and labeled as class 2 (Figure 3, part 3D).
- Tiling and running the final 3D network are done using the same steps as for the initial model.

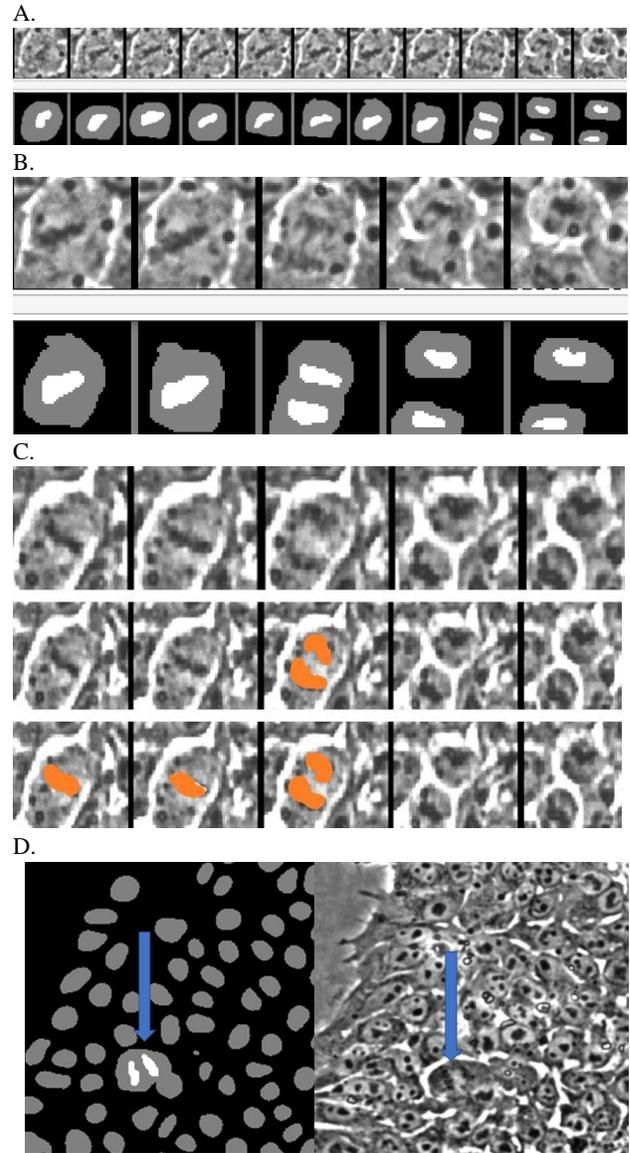


Figure 2: Manual annotations for the 3D (A-B) and 2D (C) models, using all of the events that were manually selected in our 256x256x500 tile. A: Below each time frame of a dividing cell is the additional annotation of that cell as nucleus (grey, class 1) and chromatin (white, class 2). B: an enlarged version of A. C: Labeling of chromatin for the 2D models: first row, the phase images in time, second row, labels of sister chromatids prior to division of the cell in orange, third row, label of all chromatin prior to separation of sister chromatid and division, in orange. D: An example cell division event labels put into the 3D image stack.

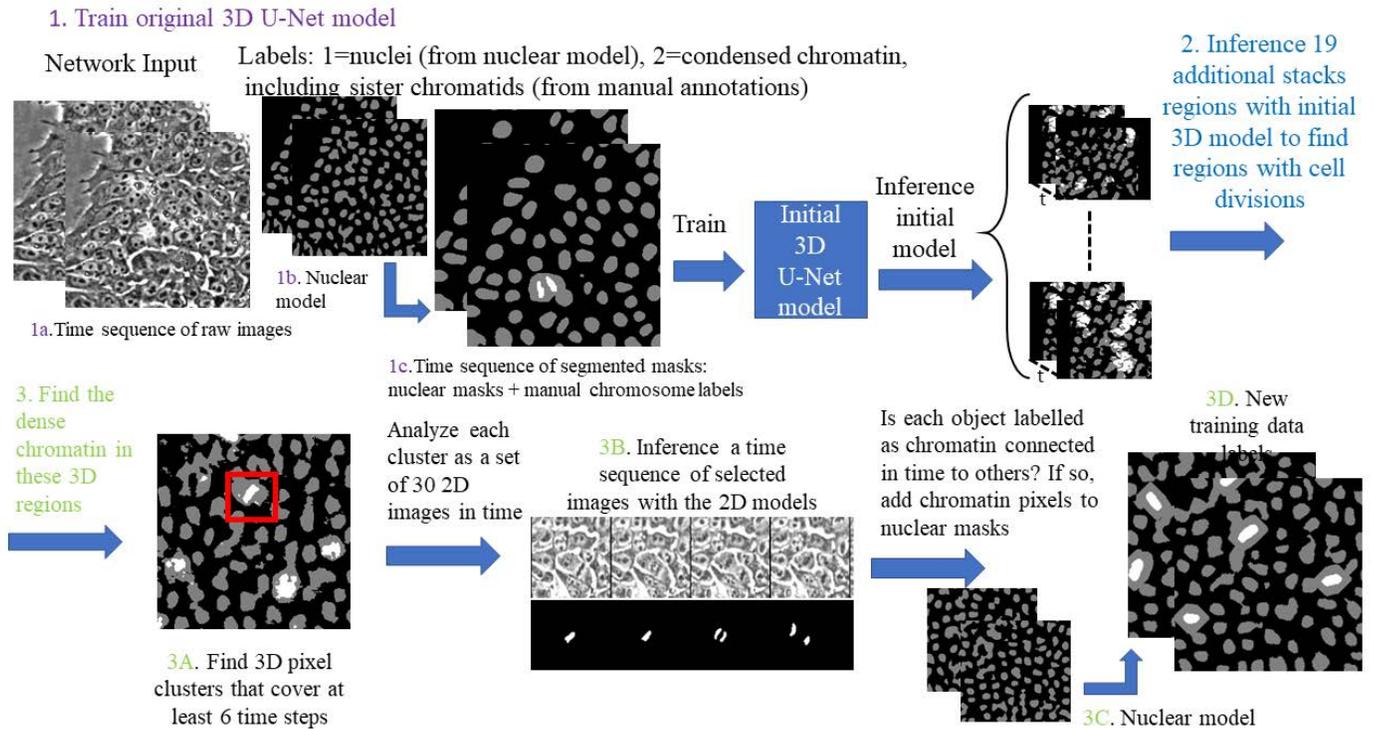


Figure 3: Creating semi-supervised training data: Data are manually annotated to find phase-dense chromatins for both the 3D and 2D models (see Figure 2); Top row: an initial 3D model is trained using nuclear masks and manual chromatin labels; 19 additional image stacks are inferred with this initial model to find 3D regions that contain a cell division; Bottom row: 3D clusters are collected that span at least 6 time steps; 2D slices from those regions are inferred with the 2D models to find chromatin objects; sister chromatid objects that persist over 4 time steps are added as new training data labels, overlaid over nuclear masks.

2.6 Evaluation of the pipeline that identifies (x,y,t) positions for each cell division

We evaluated our models for cell division using a separate small test image set collected on a different day than our training data set (Set 3 in Table 1). The test set is a $256 \times 256 \times 200$ image stack. We used this data set to test both the initial 3D model, trained from the manually annotated image stack of $(256 \times 256 \times 500)$ time steps, and the final 3D model, which was made from the 19 stacks of semi-supervised data, together with the 2D models, and compared results.

The test process follows similar steps to creating the semi-supervised training data. We first inferred the test tile with our 3D models to find regions of mitotic activity, then selected out phase images found by the 3D models and inferred with the 2D models to find dividing cells. To find the (x,y,t) of the specific cell that is dividing we find the end of the event, which is the last time

step before the divided cells start moving away from one another. The identification of cell coordinates at this stage is critical to provide input to the program for tracking cells, including dividing cells, over long times. As an additional check, for each dividing cell that we located with the 2D models, we also make sure the pixels that we believe indicate division were found by the 3D model. For this, we count the number of class 2 pixels from the 3D model inference in a small region around the dividing cell, 21×21 pixels in (x,y) and 5 pixels in time. It is possible that the 2D model can incorrectly inference cells that are not dividing but were included in the collection of $128 \times 128 \times 30$ images. Cells for which there are not corresponding pixels in the 3D inferred stack are not considered in the final count of true and false positive events.

3. Results

We inferred both of our 3D models, Model 1 and Model 2, on a separate test image stack, dataset 3. The pipeline for testing both models included inferring with the 3D model (Models 1 or 2) to find regions of mitotic activity and then using the 2D models (Models 3 and 4) to find the dividing cells.

Both models found 38 dividing cells in our test tile set 3. Our models found 12 events that were not identified by a human but were later validated to be true division events. Taking into account that the time steps between images was longer in the test data set (125 sec. vs 160 sec.) we were very encouraged by these results. Visually, Model 2 was able to locate the areas associated with phase-dense chromatin with smoother boundaries, as shown in Figure 5. Although the shapes of the inferred regions differed,

both models found the same mitotic regions after filtering out the noise. Future work in which more training data are used to train the model should be able to locate the chromatin of dividing cells more precisely, perhaps even in the absence of the 2D models. Identifying the chromatin precisely is essential for the ultimate ability to track dividing cells over time.

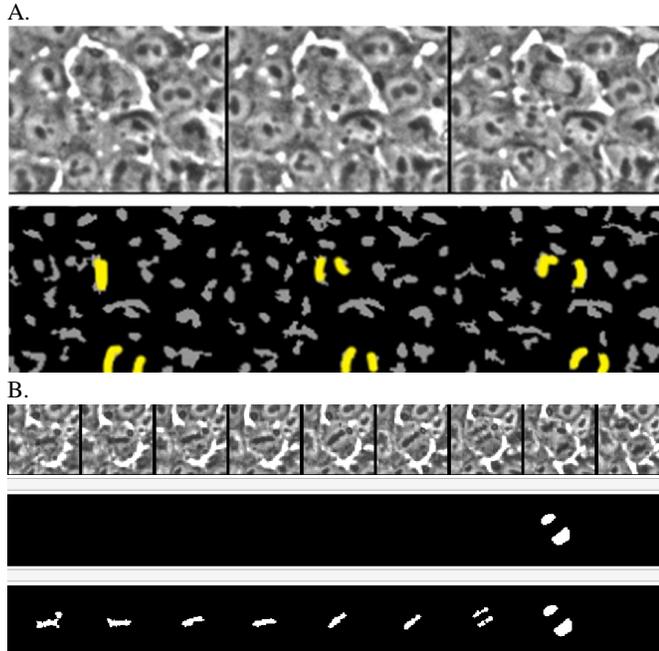


Figure 4 Testing to see if a presumed phase-dense region is part of a mitotic event. **A:** phase images are thresholded at 20 % intensity. A true event will have metaphase chromatin and/or sister chromatids in these dark regions of the phase images that largely maintain their(x,y) coordinates over the time of the mitotic event (seen in the bottom thresholded image colored in yellow). **B:** An example of an event inferred by each 2D model, in which the chromatin remained at nearly the same (x,y) position over 8 time steps.

For each model, we use the evaluation process in Section 2.6, locate the dividing cells, and then compare that list with the list of dividing cells from the manual annotation. If the location of the dividing cell we found in the model is within 10 pixels in the (x,y) plane and within 3 time steps of the annotated value, we consider that cell to be a true positive in our test. If a cell is found from the 2D model but is not present in the 3D inferencing (defined as fewer than 20 pixels in the same region surrounding the cell location), we consider that to be a false negative for the overall model. The overall results of testing both (the initial 3D model + 2D models) and (the final 3D model + 2D models) using test data set 3, are shown in Table 3. The results are similar for both models, with a few more false positive cells for the initial model. There was only a single set of sister chromatids that was found using the 2D models in the 3D rectangular area defined by an inferred 3D pixel cluster that was not part of that 3D cluster. A few time steps of this particular event that was missed by both 3D models is shown in

Figure 6. The sister chromatids are not distinct, and the metaphase chromatin is only present in two of the timed images. This type of event will be included in the training of future models.

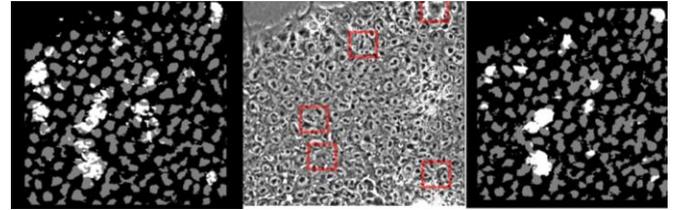


Figure 5: Differences in the inferred test stack between the initial and final models. The center image shows the initial time step for the test image stack. Several mitotic events (shown in red boxes) can be seen at this time step. On the left is the outcome of inferencing with the initial 3D model (Model 1), and on the right, with the final 3D model (Model 2). The final model defines the mitotic regions more clearly.

Table 3: Tabulated results: Numbers of correctly located dividing cells of test set 3, using each 3D model, Model 1 and Model 2, combined with our 2D models (#3,4): TP=true positive, FP=false positive, TN=true negative

	Cells located by 2D model	Cells located by 3D model	TP	FP	TN
Initial model	37	38	37	10	1
Final model	37	38	37	7	1

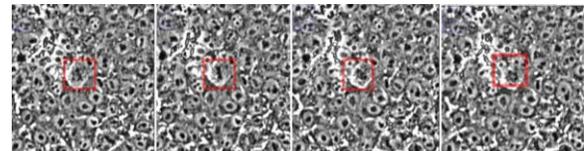


Figure 6: 4 time steps leading up to the one event not found by either 3D model. The splitting cell is within the red boxes.

4. Discussion

Due to the heterogeneous and dynamic aspects of living cells, adequate characterization of cell populations requires large image datasets that contain many cell images and many time frames. Human detection of cell behavior and tracking of cell lineage during division is not adequate because of the large amount of data required, and because of the subtle nature of the features of interest.

We approached the development of a robust automated method of analysis of iPS cell division by starting with a small set of manually annotated data that identified dividing cells. To improve and

generalize the model, sequential additions of unannotated datasets containing dividing cells were analyzed with the 3D U-Net model and were further optimized in a semi-supervised fashion that included the addition of a 2D network model. These inferred data were added to retrain the 3D model. Although the same mitotic regions were found from both the initial and final 3D models, the regions were more defined and there were fewer false positive region detections using the final 3D+2D combination of models. Future work will involve improving our model as more data becomes available.

Providing a human annotated dataset for training is time consuming and error prone. For example, an initial 3D U-Net model based on manual annotation identified 6 division events that were not previously found by human observation alone. Our models also found 12 events in our test tile set 3 that were not identified by a human. Part of the challenge is that the dataset that consists of dividing cells is relatively sparse. To give some context, the total number of cells at each time step is between 260 and 300 as determined by applying nuclear segmentation [3,4] and FogBank cell separation [11] to each frame. Over the 200 time steps, this is equivalent to $5.6e4$ cell image objects. A total of 38 division events were found and each event was spread out over approximately 10 time frames, so approximately only 380 image objects out of $5.6e4$ were associated with a dividing cell.

Another significant challenge in this study is validating the accuracy of these image analysis methods. In this study, we have relied on manual observations to identify the accuracy of the cell-division detection pipeline, but it is apparent that the human eye is far from adequate at finding these events. Once the network model has identified events, the human eye is guided and can confirm that a division event occurred. It is possible that even with careful examination we missed events in annotations.

While we consider human observation to be ground truth in this study, we also acknowledge that the development of training data, and evaluation of accuracy of a pipeline ideally should not be left to the human eye. Thus, future work will involve the development of more definitive methods for the identification of dividing and divided cells. Recent development of an iPS cell line that expresses multiple fluorescent proteins at different levels in individual cells [12] suggests an approach for a more robust and automatable way of creating training sets and evaluating pipelines for cell division.

5. Conclusions

We have successfully built a combination of 2D and 3D neural network models that identify cell division events in iPS cell images taken over time. We have shown that the model does a better job of detecting mitotic events than the humans who performed manual annotations of the events.

Our method of creating semi-supervised 3D+2D training data is completely automated and can continue to be used to improve our model as more data are collected. Because the reliability of identifying a division event can occur in an automated fashion, manual annotation of new training data will not be required. We have seen that the addition of 19 image sets to the initial training

data resulted in a 3D+2D model that more accurately located the mitotic events.

The addition of the 2D models to the pipeline made it possible to improve the accuracy of the analysis with a limited amount of training data. The 2D models were used both to make new semi-supervised training data, and to find the location of each dividing cell after using the 3D+2D combination of models. This combination of models and the creation of semi-supervised 3D training data were an effective way to create an analytical pipeline while minimizing the difficult and error-prone collection of ground truth data by humans.

REFERENCES

- [1] Y. Su, Y. Lu, J. Liu, M. Chen, A. Liu, Spatio-Temporal Mitosis Detection in Time-Lapse Phase-Contrast Microscopy Image Sequences: A Benchmark, *IEEE Trans. On Medical Imaging*, 40 (5), May 2021.
- [2] Y. Su, Y. Lu, J. Liu, M. Chen, A. Liu, Spatiotemporal Joint Mitosis Detection Using CNN-LSTM Network in Time-Lapse Phase Contrast Microscopy Images, *IEEE Trans. On Medical Imaging*, 5, 2017.
- [3] N.J. Schaub, NA Hotaling, P. Manescu, S. Padi, Q. Wan, R. Sharma, A. George, J. Chalfoun, M. Simon, M. Ouladi, et al, Deep learning predicts function of live retinal pigment epithelium from quantitative microscopy. *J Clin Invest* 2020, **130** (2):1010-1023.
- [4] C. Ling, M.Halter, A. Plant, M. Majurski, J. Stinson, J. Chalfoun, Analyzing U-Net Robustness for Single Cell Nucleus Segmentation from Phase Contrast Images. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW): 14-19 June 2020. 4157-4163.
- [5] W. Burton, C. Myers, P. Rullkoetter, Semi-supervised learning for automatic segmentation of the knee from MRI with convolutional neural networks, *Computer Methods and Programs in Biomedicine*, 189, 2020.
- [6] E. Takaya, Y. Takeichi, M. Ozaki, S. Kurihara, Sequential semi-supervised segmentation for serial electron microscopy image with small number of labels, *J. Neuroscience Methods*, 351, 2021.
- [7] M. Peikari, S. Salama, S. Nofech-Mozes, A. Martel, A cluster-then-label semi-supervised learning approach for pathology image classification, *Scientific Reports*, 8:7193, 2018.
- [8] H. Su, Z. Yin, S. Huh, T. Kanade, J. Zhu, Interactive cell segmentation based on active and semi-supervised learning, *IEEE Trans. Medical Imaging*, 35, 3, 2016.
- [9] K. Pasupa, S. Tungjitnob, S. Vathanavaro, Semi-supervised learning with deep convolutional generative adversarial networks for canine red blood cells morphology classification, *Multimedia Tools and Applications*, 79, 2020.
- [10] O. Ronneberger, P. Fischer, T. Brox. U-Net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 234-241, 2015.
- [11] J. Chalfoun, M. Majurski, A. Dima, C. Stuelten, A. Peskin, M. Brady, FogBank: a single cell segmentation across multiple cell lines and image modalities. *BMC Bioinformatics*, 15:431-442, 2014.
- [12] D. El-Nachef, K. Shi, K. Beussman, R. Martinez, M. Regier, G. Everett, C. Murry, K. Stevens, J. Young, N. Sniadecki, J. Davis, A rainbow reporter tracks single cells and reveals heterogeneous cellular dynamics among pluripotent stem cells and their differentiated derivatives, *Stem Cell Reports*, 15, 1, 226-241, July 2020.