CellPhe: a toolkit for cell phenotyping using time-lapse imaging and pattern recognition

Julie C Wilson (julie.wilson@york.ac.uk)
University of York
Laura Wiggins
University of York
Alice Lord
University of York
Peter O'Toole
University of York
William Brackenbury
University of York

Article

Keywords: cell phenotyping, pattern recognition, ensemble classification

Posted Date: October 28th, 2021

DOI: https://doi.org/10.21203/rs.3.rs-971415/v1

CellPhe: a toolkit for cell phenotyping using time-lapse imaging and pattern recognition

Laura Wiggins[1,2], Alice Lord[2], Peter J. O'Toole[1,2], William J. Brackenbury[1,2] and Julie Wilson[3*]

[1] York Biomedical Research Institute, University of York, York, UK
[2] Department of Biology, University of York, York, UK
[3] Department of Mathematics, University of York, York, UK

* Corresponding author
email: julie.wilson@york.ac.uk

Abstract. With phenotypic heterogeneity in whole cell populations widely recognised, the demand for quantitative and temporal analysis approaches to characterise single cell morphology and dynamics has increased. We present CellPhe, a pattern recognition toolkit for the characterisation of cellular phenotypes within time-lapse videos. To maximise data quality for downstream analysis, our toolkit includes automated recognition and removal of erroneous cell boundaries induced by inaccurate tracking and segmentation. We provide an extensive list of features extracted from individual cell time series, with custom feature selection to identify variables that provide greatest discrimination for the analysis in question. We demonstrate the use of ensemble classification for accurate prediction of cellular phenotype and clustering algorithms for the characterisation of heterogeneous subsets. We validate and prove adaptability using different cell types and experimental conditions. Our methods could be extended to other imaging modalities, such as fluorescence, and would be suitable for all time-lapse studies including clinical applications.

1 Introduction

Heterogeneity in whole cell populations is a long-standing area of interest[1,2,3] and previous studies have identified cell-to-cell phenotypic and genotypic diversity even within clonally derived populations.[4] The emergence of methods such as single-cell RNA sequencing has enabled characterisation of subsets within a population from gene expression profiles,[5] yet these methods involve collection of data at discrete time points, missing the subtle temporal changes in gene expression on a continuous scale. Such methods exclude information on single-cell morphology and dynamics, yet cellular phenotype plays a crucial role in determining cell function,[6,7] disease progression,[8] and response to treatment.[9] There remains a demand for quantitative and temporal analysis approaches to describe the subtleties of single-cell heterogeneity and the complexities of cell behaviour.

Modern microscopy advancements facilitate the ability to produce information-rich images of cells and tissue, at high-throughput and of high quality. Temporal changes in cell behaviour can be observed through time-lapse imaging, but the task of identifying individual cells and following them over time is an ongoing computer vision challenge.[10,11] Processing of cell time-lapse images can be broken down into two stages: segmentation, the ability to detect cells as regions of interest (ROIs) and distinguish them from background, and tracking, providing each identified cell with a unique identifier which is retained in consecutive frames. Imaging artefacts vary between experiments and finding solutions to issues such as background noise, inhomogeneity of cell size and texture and overlapping cells is still a challenge for biomedical research.[12] Reliable cell segmentation protocols are non-deterministic and experiment-specific.[13]

Recent years have seen the emergence of user-friendly software that uses machine learning algorithms to assist in objective, high-throughput cell segmentation and tracking, allowing users to focus their attention on data analysis.[14,15] However, to ensure precise quantification of cell morphology and motility, and to monitor major cellular events such as mitosis and apoptosis, it is vital that instances of erroneous segmentation and tracking are removed from data sets prior to downstream analysis methods.[16] Manual removal of such errors is heavily labour-intensive, particularly when time-lapses take place over several days.

Here we present CellPhe, a pattern recognition toolkit that uses the output of segmentation and tracking software such as CellProfiler[17] and Ilastik.[18] To maximise data quality for downstream analysis, CellPhe includes the recognition and removal of erroneous cell boundaries induced by inaccurate segmentation and tracking. Customised feature selection is used to identify the most discriminatory variables for a particular objective from an extensive list of features extracted from the time-course images. These extracted variables quantify cell morphology, texture and dynamics and describe temporal changes and can be used to reliably characterise and classify individual cells as well as cell populations. We demonstrate the use of ensemble classification for accurate prediction of cellular phenotype and clustering algorithms for identification of heterogeneous subsets. We exemplify CellPhe by characterising the behaviour of untreated and chemotherapy treated breast cancer cells from ptychographic time-lapse videos. Quantitative phase images (QPI)[19,20,21] do not make use of fluorescent labels, though their usage enhances the difference in intensity of cells and background which often improves cell segmentation accuracy. We show that our methods successfully recognise and remove a population of erroneously segmented cells, improving data set quality without fluorescence-induced perturbation.[22] Morphological and dynamical changes induced by chemotherapeutics, particularly at low drug concentration, are often more subtle than those that discriminate distinct cell types and we demonstrate the ability of CellPhe to automatically identify time series differences induced by chemotherapy treatment, with the chosen variables proving statistically significant even when not observable by eye.

The complexities of heterogeneous drug response and the problem of drug resistance further motivate our chosen application. The ability to identify discriminatory features between treated and untreated cells can allow automated detection of "non-conforming" cells such as those that possess cellular drug resistance. Further investigation of such features could elucidate the underlying biological mechanisms responsible for chemotherapy resistance and cancer recurrence. We validate the adaptability of CellPhe with both a different cell type and a different drug treatment and show that variables are selected according to experimental conditions, tailored to properties of the cell type and drug mechanism of action.

CellPhe would extend to other imaging modalities such as fluorescence images, where further extracted variables such as fluorescence intensity would complement our existing list of metrics. A comprehensive manual with a working example is provided to guide the user through the complete workflow.

2 Results

Overview of CellPhe

CellPhe is a toolkit for the characterisation and classification of cellular phenotypes from time-lapse videos, a diagrammatic summary of CellPhe is provided in **Figure 1**. Experimental design is determined by the user prior to image acquisition where seeded cell types and pharmacology are specific to the user's own analysis. Example uses are discrimination of cell types (e.g, neurons vs. astrocytes), characterisation of disease (e.g. healthy vs. cancer) , or assessment of drug response (e.g. untreated vs. treated). The user can then time-lapse image cells for the desired amount of time, using an imaging modality of their choice. Once images are acquired and segmentation and tracking of cells are complete, cell boundary coordinates are exported and used for calculation of an extensive list of morphology and texture features. These together with dynamical features and extracted time series variables are used to aid removal of erroneous segmentation by recognition of error-induced interruption to cell time series. Once all predicted segmentation errors have been removed from data sets, feature selection is performed and only features providing separation above an optimised threshold are retained. This identifies a list of most discriminatory features and allows the user to explore biological interpretation of these findings. The extracted data matrices are then used as input for ensemble classification, where the phenotype of new cells can be accurately predicted. Furthermore, clustering algorithms can be used to identify heterogeneous subsets of cells within the user's data, both inter- and intra-class.

The remaining results exemplify the use of CellPhe with a biological application, characterisation

and classification of chemotherapeutic drug response. We look at each of the CellPhe stages in detail (segmentation error removal, feature selection, ensemble classification and cluster analysis) and demonstrate that each step provides interpretable, biologically relevant results to answer experiment specific questions and aid further research.

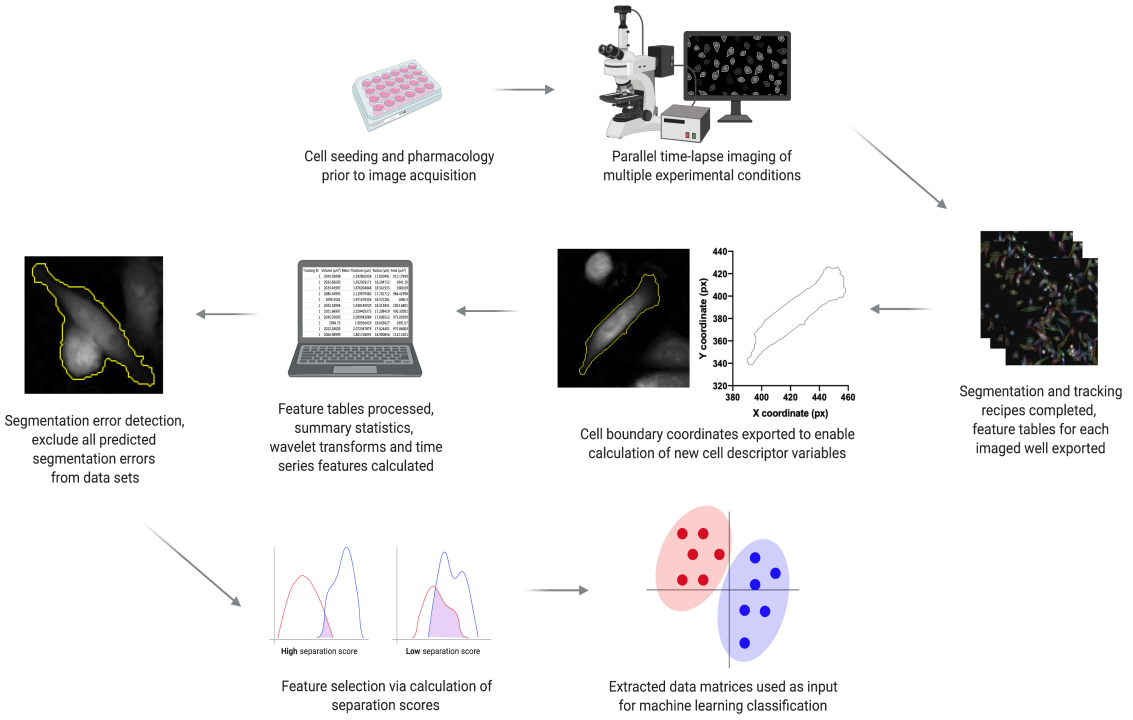

Figure 1: Summary of the CellPhe toolkit. Following time-lapse imaging, acquired images are processed and segmentation and tracking recipes implemented. Cell boundary coordinates are exported, features extracted for each tracked cell and the time series summarised by characteristic variables. Predicted segmentation errors are excluded and optimised feature selection performed using a threshold on the class separation achieved. Finally, multiple machine learning algorithms are combined for classification of cell phenotype and clustering algorithms utilised for identification of heterogeneous cell subsets.

CellPhe application: characterising chemotherapeutic drug response

The 231Docetaxel data set, obtained from multiple experiments involving MDA-MB-231 cells, both untreated and treated with 30µM docetaxel, is the main data set used to demonstrate our method. We show that the same analysis pipeline can be applied to other data sets by considering both a different cell line, MCF-7, in the MCF7Docetaxel data set, and a different drug, doxorubicin, with the 231Doxorubicin data set. In each case, we remove segmentation errors, as described in Section 2.5, before using feature selection (Section 2.6) to identify discriminatory variables tailored to the particular data set. We show that different variables are chosen depending on the inherent nature of the cell line and the effect of the drug in question. By using these features in classification algorithms, we aimed to characterise the behaviour over time of untreated and treated cells.

Segmentation Error Removal

The purpose of this analysis was to improve the quality of our data sets prior to untreated vs. treated cell classification by automating detection of segmentation errors and optimising the exclusion criteria of predicted errors.

Comparison of time series for cells with and without segmentation errors showed many of our features to be sensitive to such errors, motivating the need to remove these cells prior to treatment classification. Size metrics, such as volume, were particularly affected by segmentation errors as under- or over-segmentation could result in halving or doubling of cell volume respectively (**Figure**

2a). This noticeable disruption to the time series of several features suggested that reliable detection of segmentation errors would be possible.

After excluding 62 instances identified as tracked cell debris, a training data set for MDA-MB-231 cells (from the 231Docetaxel data set), was obtained, consisting of 1185 correctly segmented cells and 278 cells with segmentation errors. The number of cells in the segmentation error class was doubled using SMOTE and the resulting data set with 1741 observations used for the classification of segmentation errors as described in Section 2.5. The MDA-MB-231 cells (from 231Docetaxel and 231Doxorubicin, both untreated and treated) that were not used for training formed independent test sets (Table 1).

A total of 225 of the 1478 cells in the 231Docetaxel test set were predicted to be segmentation errors. Of these, 219 were confirmed by eye to be true segmentation errors, most of which were due to under- or over-segmentation throughout their time series. Other segmentation issues observed included background pickup, cells swapping cell ID, and cells repeatedly entering and exiting the field of view, all of which result in problem time series (**Figure 2b**). Although the 6 cells that were misclassified as segmentation errors were all large cells, further investigation showed that removal of these cells did not exclude an important subset from the data. Comparison of the area and volume of these cells with correctly segmented cells, confirmed that a population of very large cells were correctly classified as having no segmentation errors so that such cells were still represented in the test set (**Figure S1**).

This classifier was also used to identify a further 83 segmentation errors from the 1005 cells in the 231Doxorubicin data set, all 83 were confirmed by eye to be true segmentation errors (Table 1). It was necessary to train a new classifier for MCF-7 segmentation error detection due to differences between the cell lines. In this case 308 correctly segmented cells and 192 segmentation errors were identified by eye. After applying SMOTE to double the number of segmentation error observations, a classifier was trained with the resulting 692 observations as described in section 2.5. Just 4 of the 287 cells in the MCF7Docetaxel test set were classified as segmentation errors and were confirmed by eye to be true segmentation errors.

Data set	**TP**	**FP**
231Docetaxel (1478)	219	6
231Doxorubicin (1005)	83	0
MCF7Docetaxel (287)	4	0

Table 1: Segmentation error prediction on the test data. The number of correctly classified segmentation errors (True Positives, TP) and the number of correctly segmented time series incorrectly classified as segmentation errors (False Positives, FP) are shown. The number of cells in each test data before segmentation error removal is shown in parentheses.

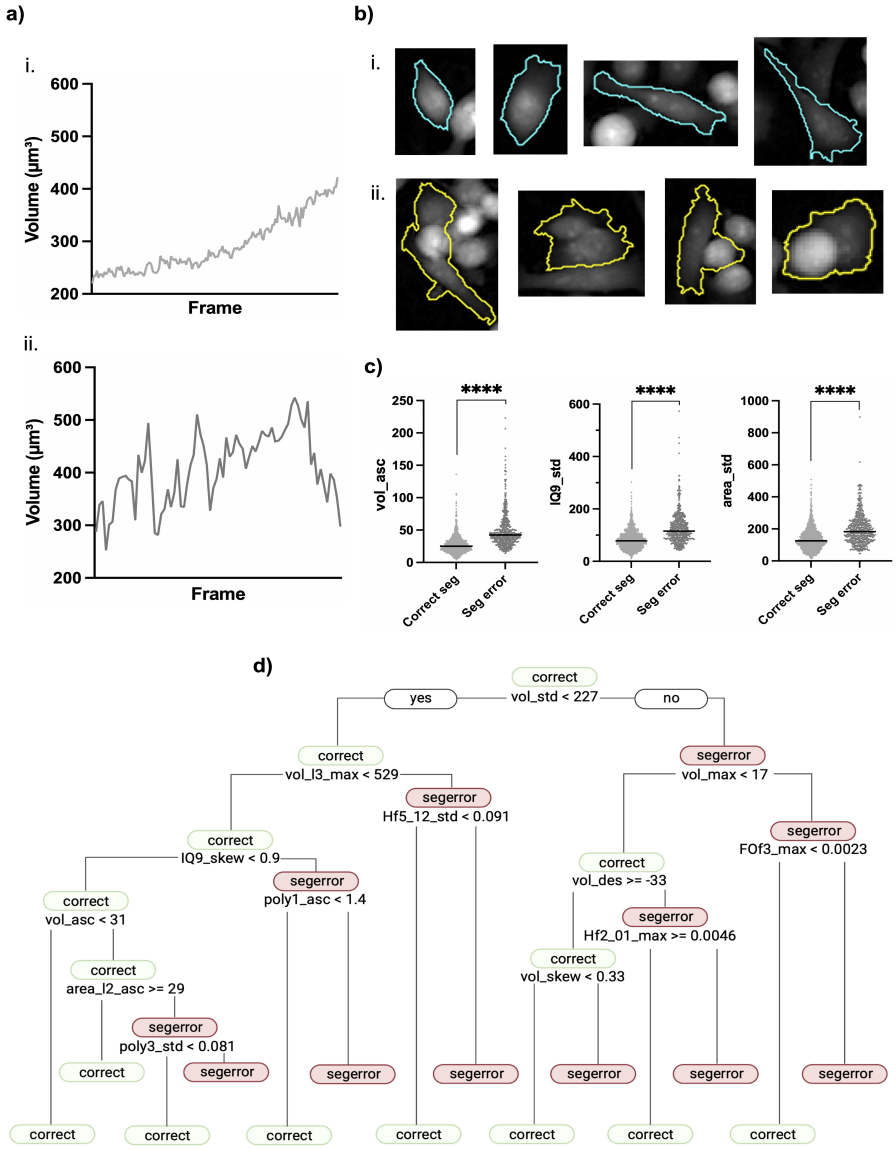

Figure 2: **(a)** Volume time series for **i.** a correctly segmented cell and **ii.** a cell experiencing segmentation errors, demonstrating greater fluctuation in volume when a cell experiences segmentation errors. **(b)** Examples of test set cells classified as **i.** correct segmentation and **ii.** segmentation error. **(c)** Beeswarm plots of features that are significant for identifying segmentation errors in the 231Docetaxel training set (****: $p < 0.0001$). Interruptions to time series induced by segmentation errors are identified by increased ascent and standard deviation for affected features. **(d)** A representative 231Docetaxel trained decision tree, demonstrating how volume, texture and polygon (shape) variables are used in combination to make classifications.

Although feature selection was not used in the identification of segmentation errors, we did calculate separation scores for the MDA-MB-231 training data to investigate the effect of such errors. As might be expected, volume was most affected, with segmentation errors resulting in larger standard deviation, ascent and maximum value. Other features with high separation scores included area as well as spatial distribution descriptors with the highest thresholds, features that detect the clustering of high intensity pixels, characteristic of cell overlap and over-segmentation (**Figure 2c**). Analysis of the trained decision trees showed that a combination of size, texture and shape variables

frequently formed the most important features for detecting segmentation errors with MDA-MB-231 cells, see **Figure 2d** for an example.

For the MCF7Docetaxel data set, velocity was found to be important in determining whether or not a cell experienced segmentation errors in addition to texture and shape variables. The cell centroid, used to determine position and hence velocity, is affected by boundary errors and so high velocity, uncharacteristic of MCF-7 cells, is a good indication of segmentation error for these cells.

Feature Selection

For the 231Docetaxel data set, the calculation of separation scores identified variables that provided good discrimination between untreated MDA-MB-231 cells and those treated with 30μM docetaxel. As separation scores do not provide information on how these variables work in combination, we performed Principal Component Analysis (PCA) to explore relationships between discriminatory variables.

Differences in the appearance of MDA-MB-231 cells induced by docetaxel treatment were observed by eye from cell timelapses. Untreated cells displayed a spindle-shaped morphology (a circular cross-section with tapering at both ends), with contractions and protrusions facilitating migration. Cells that received treatment were generally dense and spherical, and increased in size following a failed attempt at cytokinesis (**Figure 3a**). Discriminatory features identified by calculation of separation scores were consistent with differences observed by eye, the 20 variables that achieved greatest separation are shown in **Figure 3b**. The most discriminatory feature was the ratio of pixels within the cell boundary to the number of pixels within the minimal bounding box (A2B). Low values indicate a small cell area within a large bounding box, as is the case for spindle-shaped cells. Other morphological features that successfully characterised untreated MDA-MB-231 cell morphology were length, radius and polygonal representation variables, all of which discriminate between irregular untreated cell shape and the more rounded morphology of treated cells. Furthermore, separation scores highlighted differences in the texture of cells following treatment, with both first order and Haralick features providing good discrimination between untreated and treated cells.

Principal Component Analysis (PCA) demonstrated that the main variance within the data arises due to class differences, with separation of classes observed across PC1 which explains 54.2% of the total variance (**Figure 3c**). The dispersion of points within the scores plot illustrates heterogeneity of cells both inter- and intra-class. The non-conformity of some cells, for example treated cells behaving as untreated cells, is demonstrated by points clustering within the opposite class. The PC1 loadings with greatest absolute values are given in **Figure 3c**. Notably, only average area to minimal box ratio (A2B_mean) and variance in edge length of the polygonal approximation (poly2_mean) had positive PC1 loadings, indicating that treated cells had greater values for these variables in comparison to untreated cells. The remaining variables had negative PC1 loadings, including standard deviation, ascent and descent of morphological features such as length and minimal box, and several texture features. As the majority of untreated cells had negative PC1 scores we deduced that greater standard deviation, ascent and descent of features for untreated cells indicates that these cells experience increased fluctuation throughout their time series. As treated cells mainly had positive PC1 scores, they experience less fluctuation throughout their time series and instead display greater stability. Identified differences in feature time series are visualised in **Figure 3d**.

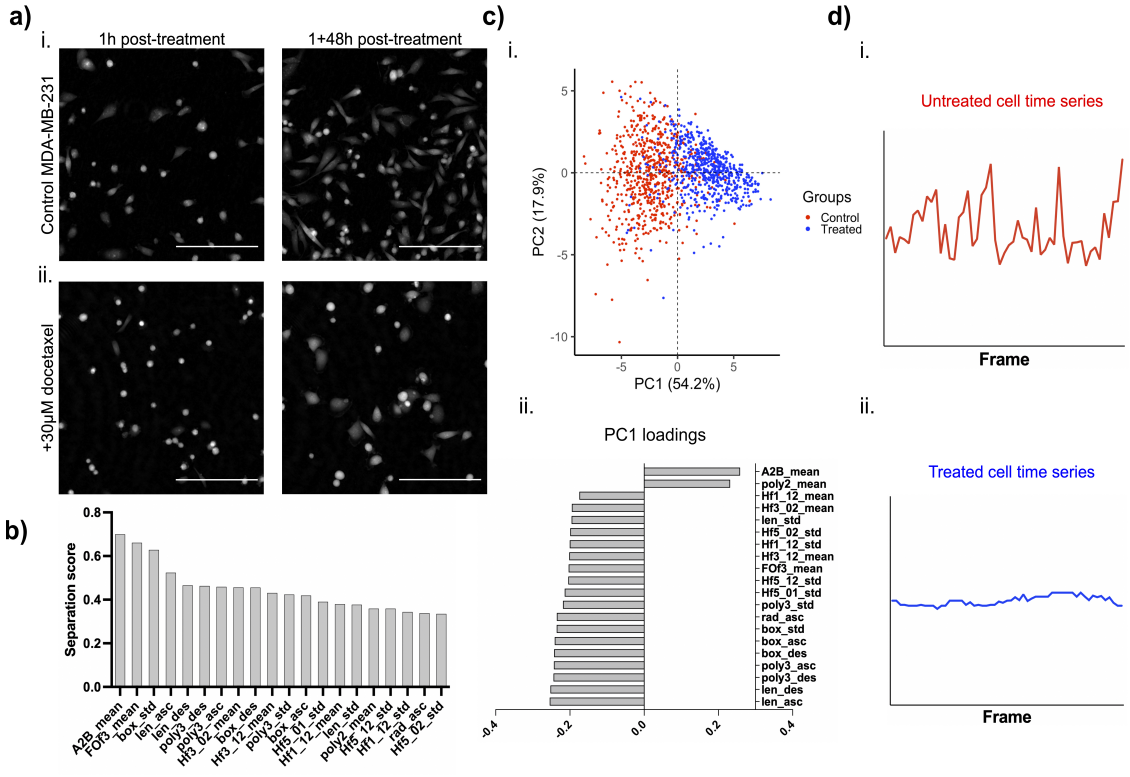

Figure 3: a) Images taken from cell timelapses of **i.** untreated MDA-MB-231 cells and **ii.** 30μM docetaxel treated MDA-MB-231 cells. Scale bar = 200μm. Increased cell count at 1+48h post treatment demonstrates healthy proliferation of untreated cells. Static cell count at 1+48h for treated cells is a result of cell cycle arrest and failed cytokinesis, leading to enlarged cell phenotype. **b)** The 20 most discriminatory features identified through calculation of 231Docetaxel separation scores. **c) i.** PCA scores plot with points colour-coded according to true class label. Observable separation of classes along PC1 demonstrates that the greatest source of variance within the data arises due to class differences. **ii.** PC1 loadings suggest increased activity within untreated cell time series, with standard deviation, ascent and descent variables obtaining negative PC1 loadings and scores. **d)** Representative feature time series plots for **i.** untreated MDA-MB-231 cells and **ii.** 30μM docetaxel treated MDA-MB-231 cells. Untreated cells experience greater fluctuation within their time series in comparison to treated cells where activity is more stabilised.

We assessed the adaptability of our feature selection method by calculating separation scores for both a different cell line and a different treatment, using PCA to evaluate the main sources of variance. We compared MCF-7 cells treated with 1μM docetaxel with untreated MCF-7 cells, and MDA-MB-231 cells that were treated with 1μM doxorubicin with untreated MDA-MB-231 cells and found that changes in the morphology and motility of cells upon treatment were both drug and cell-line specific with different variables selected. (**Figure 4**).

As was observed within the 231Docetaxel timelapses, cells increased in size due to failed cytokinesis. However, MCF-7 cells maintained a polygonal, epithelial-like morphology following treatment similar to that of the untreated population. Furthermore, no differences in movement were observed within the MCF7Docetaxel data set due to the poorly aggressive, non-invasive nature of MCF-7 cells described previously.[23] Conversely, remarkable differences in cellular dynamics were observed within the 231Doxorubicin data set, with motility of cells being severely hindered following treatment, particularly after the 24-hour time point. Only subtle differences in size and morphology of cells were observed by eye, with doxorubicin treated cells appearing slightly enlarged as a result of cell cycle arrest. Both untreated and treated sets contained examples of cells in G1 and G2, hence varied cell morphology can be observed within both (elongated and adherent cells in G1, round and dense morphology of cells in G2.)

The 20 variables that achieved greatest separation for each of the MCF7Docetaxel and 231Doxorubicin data sets are shown in **Figure 4b**. Primarily size and texture variables were identified as most discriminatory for MCF7Docetaxel with variables such as length, width and area characterising the enlarged cell shape of treated cells. Spatial distribution variables were chosen for several intensity thresholds, demonstrating differences in the clustering of pixels, following docetaxel treatment. Furthermore, mean cell density was also identified as a discriminatory variable with the untreated cell population having greater mean cell density than the treated population, likely a result of decreased cell proliferation and cell-cell adhesion for treated cells. As was observed by eye, movement features formed the majority of discriminatory variables for the 231Doxorubicin data set, with untreated cells having greater velocity, tracklength and displacement than treated cells. Differences in movement were also described through density ascent and descent, as cell density fluctuated more for untreated cells due to the increased likelihood of passing neighbouring cells when migrating. Subtle differences in cell shape and size observed by eye upon doxorubicin treatment were described by changes in the area to minimal box ratio and width and radius variables. Notably both data sets received lower separation scores than the 231Docetaxel data set, with MCF7Docetaxel having the lowest. This effectively provides a measure of class similarity, with high separation scores for 231Docetaxel indicative of significant changes to cells upon treatment and low separation scores for MCF7Docetaxel suggesting these changes are more subtle.

PCA scores plots obtained with the selected features are shown in **Figure 4c**. Differences between classes can be observed for the MCF7Docetaxel data set although the separation involves both PC1 and PC2 whilst the greatest source of variance, along PC1 (55% of the total variance), due to heterogeneity within treated cells. We found that the subset of treated cells with lowest PC1 scores were primarily cells tracked during a failed attempt at mitosis. This subset of treated cells is characterised by fluctuations in the times series of spatial distribution features indication changes in texture as cells enter and exit failed mitosis. On the other hand, the PCA scores plot for 231Doxorubicin shows the greatest source of variance to be due to class differences, with separation of classes along PC1 (69% of the total variance). All PCA scores plots demonstrated the potential to characterise untreated and treated cell behaviour, with feature selected variables providing good distinction of classes which was improved by using variables in combination.

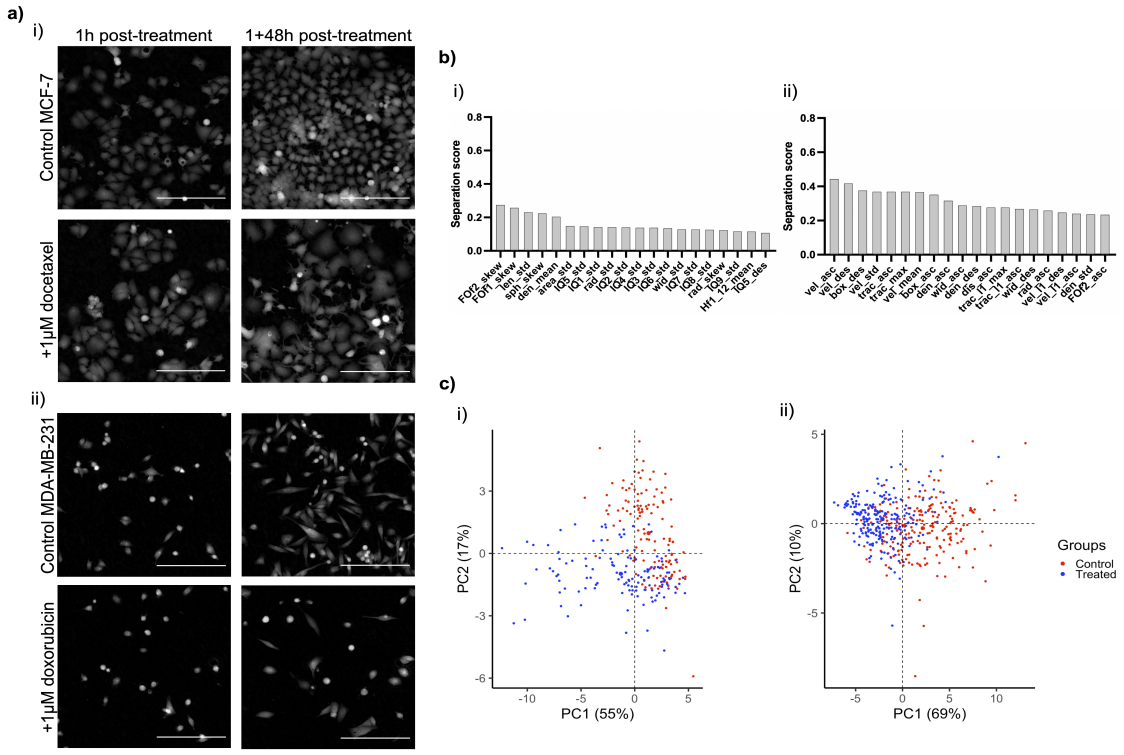

Figure 4: a) Images taken from cell timelapses of **i.** untreated and 1μM docetaxel treated MCF-7 cells and **ii.** untreated and 1μM doxorubicin treated MDA-MB-231 cells. Scale bar = 200μm. Differences in cell count following treatment can be observed for both due to cell cycle arrest induced by docetaxel or doxorubicin respectively. Docetaxel treated MCF-7 cells display enlarged cell phenotype at the 1+48h time point due to failed cytokinesis. In comparison, differences in morphology are more subtle for doxorubicin treated MDA-MB-231 cells at the 1+48h time point. **b)** The 20 most discriminatory features identified through calculation of separation scores for **i.** MCF7Docetaxel and **ii.** 231Doxorubicin. Low separation scores for MCF7Docetaxel indicate less discrimination between untreated and treated cell populations within this data set. Several movement features provide good separation for 231Doxorubicin demonstrating differences in cell motility following doxorubicin treatment. **c)** PCA scores plots for **i.** MCF7Docetaxel and **ii.** 231Doxorubicin, colour-coded according to true class label. Notably the greatest source of variance within MCF7Docetaxel is a result of treated cell variability, with separation of classes increased when PC1 and PC2 are used in combination. The greatest source of variance within the 231Doxorubicin data set arises due to class differences with separation of classes observable along PC1.

Classification of Treated and Untreated Cells

We found that the distribution of separation scores differed for each data set, with the 231Docetaxel set having the greatest number of variables achieving high separation, followed by 231Doxorubicin and with MCF7Docetaxel generally having much lower separation scores (**Figure 5a**). Optimal separation thresholds of 0.2, 0.05 and 0.025 were obtained for 231Docetaxel, 231Doxorubicin and MCF7Docetaxel respectively, resulting in 82, 162 and 195 variables (of a possible 702) being selected for classifier training.

Having chosen an optimal separation threshold, we trained an ensemble classifier for each data set as described in Section 2.6. Classification accuracy scores for training and test sets obtained using our ensemble classifier are provided in **Table 2**, whilst classification accuracy scores for the individual classification algorithms before their combination can be found in **Table S3.** Through visual inspection, we found that misclassifications formed subsets of cells whose behaviour deviated from the behaviour of the main population, we call this subset "non-conforming". (**Figure 5b**). For untreated cells, we found that healthy, proliferating cells were correctly classified whereas less motile cells, cell debris or large, non-motile mutant cells were instead classified as treated. For treated cells, we found that cells experiencing the drug-induced phenotypic differences identified through feature selection were classified as treated. However, treated cells displaying behaviour similar to that of an untreated cell, such as increased migration or fluctuation and elongation in cell shape, and were classified as untreated (**Figure 5c**).

We found that the proportion of non-conforming treated cells, those classified as untreated, decreased as drug concentration increased for all three data sets (**Figure 5d**). To explore the connection between the proportion of non-conforming treated cells and the population drug response of each treated set, we considered the total volume growth rate at each drug concentration in relation to the percentage of cells predicted as untreated (**Figure 5d**). We found that the overall growth rate decreased with increased drug concentration due to more cells responding at higher concentrations. This correlated positively with the percentage of cells predicted as untreated, with a greater percentage of cells predicted as untreated for high volume growth rate with proliferation still occurring.

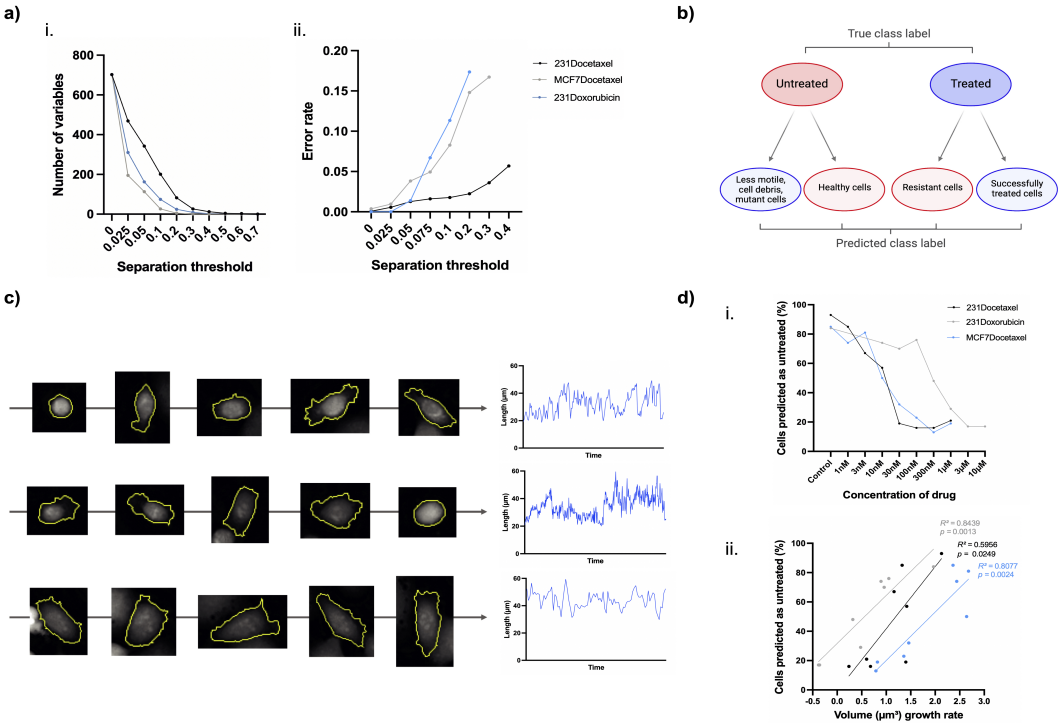

Figure 5: a) i. The number of variables with separation scores above different thresholds. A greater number of variables achieve high separation for 231Docetaxel in comparison to 231Doxorubicin and MCF7Docetaxel. **ii.** Optimisation of separation threshold for each data set. Thresholds of 0.2, 0.05 and 0.025 were selected for 231Docetaxel, 231Doxorubicin and MCF7Docetaxel respectively resulting in 82, 162 and 195 variables being used for classifier training. **b)** Sub-populations within each class, colour-coded according to the ideal final classification of each sub-population. Non-conforming cells for each class form a subset of misclassified cells. **c)** Examples of docetaxel treated MDA-MB-231 cells misclassified as untreated. Timelapse images demonstrate how these cells exhibit an elongated morphology characteristic of migratory untreated cells. Time series plots for cell length demonstrate the fluctuation in shape of these cells, typical of untreated cells. **d) i.** The percentage of cells predicted as untreated for a range of drug concentrations. For all three data sets, this percentage decreases as drug concentration increases due to a greater number of cells responding at higher concentrations. **ii.** Positive correlation between the total volume rate of growth and the percentage of cells predicted as untreated, with higher volume growth rates associated with a higher number of cells being predicted as untreated. Linear regression slopes were found to be significant (p values shown). R^2 correlation coefficients are also provided, demonstrating positive correlation for each data set.

	231Docetaxel	**MCF7Docetaxel**	**231Doxorubicin**
Train	Untreated: 96% Treated: 99% Overall: 98%	Untreated: 100% Treated: 99% Overall: 100%	Untreated: 97% Treated: 100% Overall: 99%
Test	Untreated: 97% Treated: 80% Overall: 93%	Untreated: 85% Treated: 81% Overall: 85%	Untreated: 84% Treated: 71% Overall: 81%

Table 2: Ensemble classification accuracy scores for each data set. All percentages have been rounded to the nearest whole number.

Subset Identification

Classification accuracy scores for the untreated cell population were notably greater than those of the treated population across all three of the data sets (**Table 2**), suggesting a greater proportion of non-conforming treated cells in comparison to non-conforming untreated cells. Imbalance of

classification accuracy scores in binary classification is often a result of hidden stratification,[24] where poor performance of one class is a result of misclassifications of important, unlabeled subsets. To investigate this phenomenon we performed hierarchical clustering on 231Docetaxel treated cells and the obtained dendogram is provided in **Figure 6a**, with examples of cells from each cluster in **Figure 6b**.

Figure 6c shows the distribution of mean volumes for each cluster in comparison to the untreated MDA-MB-231 population. Clusters 1, 2 and 4 span a similar range of volumes to the untreated set, whereas clusters 3 and 5 have greater mean volumes. Cluster 6 is formed primarily of cell debris as a result of cell death with mean volumes much lower than those of the untreated set.

Cells in the same cluster share similar properties (**Figure 6b**), and morphological differences between clusters of different cell cycle states can be observed. For example cells in clusters 1, 2 and 4 are much smaller and brighter than cells in clusters 3 and 5 as the cells are heading towards attempted mitosis, confirmed by visual inspection of cell timelapses, and hence resemble untreated mitotic cells. However, discrimination between clusters at similar stages of the cell cycle were not as readily identifiable by eye and we therefore calculated separation scores between these clusters to identify discriminatory variables (**Figure 6d**). Differences between cluster 2 and clusters 1 and 4 were primarily based on cell shape with cells in cluster 2 having greater fluctuation in variables such as length and width as they experienced repeated elongation and contraction during their tracking. A range of textural variables provided highest separation between cluster 1 and cluster 4, including many Haralick features, indicating differences in distribution of pixel intensities for cells within these clusters. Highest separation between clusters 3 and 5 was achieved by both shape and texture features. Cells in cluster 5 showed greater fluctuation in shape whereas the shape of cells in cluster 3 remained relatively spherical for the duration of tracking. Textural variables also achieved high separation between clusters 3 and 5, with cells in cluster 5 having greater changes in the distribution of interior pixel intensities.

Clusters also spanned a range of mean cell volumes beyond those of the untreated set when hierarchical clustering was repeated for MCF7Docetaxel treated cells. However, this was not the case for 231Doxorubicin treated cells and therefore k-means clustering was used to explore the connection between misclassifications and hidden subsets in the 231Doxorubicin treated cell test set. Two distinct clusters were obtained (**Figure 6e**i), where several movement variables achieved greatest separation between the two clusters. Cells in cluster 2 experienced increased migration in comparison to cells in cluster 1, with greater velocity, tracklength and displacement. We calculated classification accuracy scores for the two clusters individually and found that 94% of cells in cluster 1 were correctly classified as treated, but only 30% in cluster 2 (**Figure 6e**ii). The increased migration of cells in cluster 2 mean these cells have greater similarity to the untreated population. These non-conforming treated cells form the majority of treated cell misclassifications in the 231Doxorubicin test set.

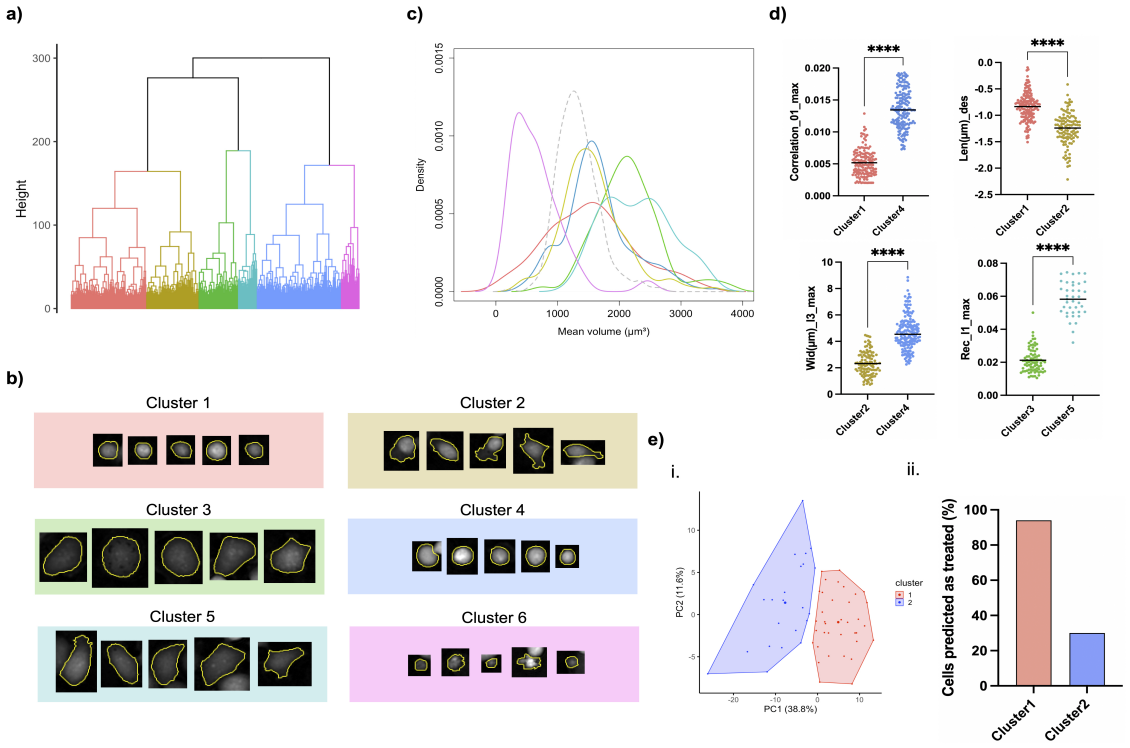

Figure 6: a) Dendogram obtained from hierarchical clustering of 231Docetaxel treated cells, with 6 clusters coloured. **b)** Examples of cells from each cluster with background colours identifying the cluster. Cells within a cluster share similar properties but differ to cells in other clusters. **c)** Density plots of mean cell volume, colour-coded according to cluster. The gray, dashed density plot represents 231Docetaxel untreated cells. Cluster 6 (cell debris cluster) has the greatest leftward shift due to cells losing volume upon cell death. Clusters 1, 2 and 4 primarily span the same range of volumes as the untreated set as cells in these clusters have not yet attempted cytokinesis. Clusters 3 and 5 have mean volumes greater than the untreated set as cells in these clusters have continued to grow following failed cytokinesis. **d)** Beeswarm plots of example discriminatory variables, identified from separation scores, for clusters in similar cell cycle stages (****: $p < 0.0001$). **e) i.** k-means clustering of 231Doxorubicin test set treated cells. Cells are colour-coded according to which cluster they were assigned. **ii.** The number of cells predicted as treated for each of the clusters. Cluster 1 was formed of successfully treated cells with 94% (34/36) of cells correctly classified as treated, whereas cluster 2 formed a subset of non-conforming treated cells, with only 30% (7/23) correctly classified as treated.

3 Discussion

The CellPhe toolkit complements existing software for automated cell segmentation and tracking, using their output as a starting point for time series feature extraction, cell classification and cluster analysis. Erroneous cell segmentation and tracking can significantly reduce data quality but such errors often go undetected and can negatively influence the results of automated pattern recognition. CellPhe's extensive feature extraction followed by bespoke feature selection not only allows the characterisation and classification of cellular phenotypes from time-lapse videos but provides a method for the identification and removal of erroneous cell tracks prior to these analyses. Attribute analysis showed that different features were chosen to identify segmentation errors for different cell lines. For example, sudden increases in movement resulting from large boundary changes can indicate segmentation errors for MCF-7 cells, contrasting with their innate low motility. On the other hand, size and texture variables provide better characterisation of the unexpected fluctuations in cell size and clusters of high intensity pixels induced by segmentation errors for MDA-MB-231 cells. Current approaches for removal of segmentation errors are subjective and labour-intensive, requiring

manual input of parameters such as expected cell size that need to be fine-tuned for different data sets. CellPhe provides an objective, automated approach to segmentation error removal with the ability to adapt to new data sets.

For cell characterisation, we have shown that CellPhe's feature selection method is able to adapt to different experimental conditions, providing discrimination between untreated and treated groups of two different breast cancer cell lines (MDA-MB-231 and MCF-7) and two different chemotherapy treatments (docetaxel and doxorubicin). The discriminatory variables identified here coincide with previously reported effects of docetaxel or doxorubicin treatment and can be interpreted in terms of the mechanism of action of each drug. Previous studies have identified a subset of polyploid, multinucleated cells following docetaxel treatment due to cell cycle arrest and occasionally cell cycle slippage.[25] Our findings support this with shape and size variables providing the greatest separation for docetaxel treatment in both MDA-MB-231 and MCF-7 cells. Many texture variables were also identified as discriminatory following docetaxel treatment, providing label-free identification of the multiple clusters of high intensity pixels in treated cells, likely a result of docetaxel-induced multinucleation. We found that at a higher, sub-lethal concentration of 1µM, migration of MDA-MB-231 cells was reduced with variables associated with movement providing greatest discrimination between untreated and doxorubicin treated cells. This is supported by studies that have identified changes in migration of doxorubicin treated cells, noting that low drug concentrations in fact facilitate increased invasion.[26,27]

We found an imbalance in untreated and treated classification accuracy scores, with a greater proportion of treated cells misclassified for all three data sets. This consistent imbalance suggests the misclassifications are in fact representative of a subset of non-conforming, and potentially chemoresistant, cells. The concept of hidden stratification, where an unlabelled subset performs poorly during classification, has been described previously[28] and poses a challenge in medical research as important subsets (such as rare forms of disease) could be overlooked. Here, the misclassified cells could be of most interest and the ability to identify non-conforming behaviour is precisely what is required from a classifier as treated cells that display behaviour similar to untreated cells could indicate a reduced response to drug treatment. The classification of cells treated with a range of concentrations supported this hypothesis as a greater proportion of cells were classified as untreated at lower drug concentrations, demonstrating that our trained ensemble classifier can be used to quantify drug response, at both single-cell and populational level.

Cluster analysis revealed cell subsets that appear to represent different responses to drug treatment. Heterogeneity of cellular drug response is a commonly reported phenomenon in cancer treatment, yet mechanisms underlying this are not well understood.[29] Analysis of cell volumes showed the mean volume of treated and untreated cells to be comparable for doxorubicin reflecting the fact that this treatment can induce G1, S or G2 cell cycle arrest.[30] However, for docetaxel treated cells, we found that clusters spanned a range of mean cell volumes beyond those of the untreated set for both cell lines. Clustering allowed identification of three general responses to docetaxel treatment: pre-"cytokinesis attempt", with cells having similar volumes to the untreated MDA-MB-231 population; post-"cytokinesis attempt", where cells were tracked following failed cytokinesis and therefore continued to grow to volumes beyond those of the late stages of the untreated cell cycle; and cell death, with a final cluster, composed primarily of cell debris. Furthermore, giant cell morphology has been linked with docetaxel resistance, a potential cause of relapse in breast cancer patients[9] and through cluster analysis we were able to identify a potentially resistant subset of very large, treated cells that could be isolated for further investigation.

Our chosen application demonstrated the breadth of quantification and biological insight that can be made by following our workflow, with characterisation of drug response and detection of potentially resistant cells just two of many potential applications for CellPhe. CellPhe offers several benefits for the quantification of cell behaviour from time-lapse images. First, errors in cell segmentation and tracking can be identified and removed, improving the quality of input for downstream data analysis. This is particularly important with machine learning where automation means that such errors can easily be missed, and algorithms consequently trained with poor data.

Second, cell behaviour is characterised over time by extracting variables from the time series of various features whereas many studies explore temporal changes by collecting data at discrete

time points (for example, 0 and 24 hours post-treatment) and using metrics from each static image, missing behavioural changes experienced by cells on a continuous level. With CellPhe, changes over time in features that provide information on morphology, movement and texture are quantified not just by summary statistics but by variables extracted from wavelet transformation of the time series allowing changes on different scales to be identified.

Third, whilst most studies use a limited number of metrics, assessed individually for discrimination between groups,[31,22] CellPhe provides an extensive list of novel metrics and automatically determines the combination that offers greatest discrimination. The bespoke feature selection frequently found the most discriminatory variables to be those with the ability to detect changes in cell behaviour over time. Previous research in this field has focused on identification of cell types from co-cultures[32] for use in automated diagnosis of disease such as cancer. Analysis methods for these studies are often cell line specific whereas CellPhe's feature selection method is successful in identifying discriminatory variables tailored to different experimental conditions.

Finally, CellPhe uses an ensemble of classifiers to predict cell status with high accuracy and we show that separation scores can be used to identify the variables associated with different cell subsets identified in cluster analysis to explore cell heterogeneity within a population, even when subtle differences are not readily visible by eye.

The interactive, interpretable, high-throughput nature of CellPhe deems it suitable for all cell time-lapse applications, including drug screening or prediction of disease prognosis. We provide a comprehensive manual with a working example and real data to guide users through the workflow step-by-step, where users can interact with each stage of the workflow and customise to suit their own experiments. Here we demonstrated the abundance of information and insight that can be made by following the CellPhe workflow to quantify cell behaviour from QPI images. CellPhe can also be extended to other imaging modalities and work is underway to determine further variables, such as fluorescence intensity, that would complement our existing list of metrics.

4 Acknowledgements

We would like to thank Dr. Jon Pitchford for his ongoing valuable advice and Dr. Fiona Frame for her useful contributions to the project. We would also like to thank the University of York Bioscience Technology Facility - Imaging and Cytometry Team for the helpful technical assistance they provided throughout the project. We express gratitude to Phasefocus UK for the Livecyte and CATbox systems that were used to acquire and export all time-lapse data presented here, and for their technical support throughout. Furthermore, we would like to sincerely thank BBSRC for their generosity in funding the project, grant number: BB/S507416/1.

5 Contributions

Conceptualisation: W.B., P.O'T., J.W., and L.W.; cell culture, pharmacology and imaging: L.W. and A.L.; data analysis and validation: L.W. and J.W.; software development J.W. and L.W; supervision: J.W., W.B. and P.O'T.; writing-original draft preparation, L.W. and J.W.; writing-review and editing, W.B. and P.O'T.

6 Competing interests

The authors declare no competing interests.

References

[1] S. Turajlic, A. Sottoriva, T. Graham, et al. Resolving genetic heterogeneity in cancer. *Nature Reviews Genetics*, 20:404–416, 2019.

[2] S. Goldman, M. MacKay, E. Afshinnekoo, et al. The impact of heterogeneity on single-cell sequencing. *Frontiers in Genetics*, 10:8, 2019.

[3] S.J. Altschuler and L.F. Wu. Cellular heterogeneity: do differences make a difference? *National Institute of Health, Cell*, 141:559–563, 2010.

[4] B. Carter and K. Zhao. The epigenetic basis of cellular heterogeneity. *Nature Reviews Genetics*, 22:235–250, 2021.

[5] F. Buettner, K. Natarajan, F. Casale, et al. Computational analysis of cell-to-cell heterogeneity in single-cell rna-sequencing data reveals hidden subpopulations of cells. *Nature Biotechnology*, 33:155–160, 2015.

[6] B.M. Davis, M. Salinas-Navarro, M.F. Cordeiro, et al. Characterizing microglia activation: a spatial statistics approach to maximize information extraction. *Scientific Reports*, 7, 2017.

[7] T. Henser-Brownhill, R.J. Ju, N.K. Haass, et al. Estimation of cell cycle states of human melanoma cells with quantitative phase imaging and deep learning. *IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 1617–1621, 2020.

[8] L.A. Tashireva, M.V. Zavyalova, O.E. Savelieva, et al. Single tumor cells with epithelial-like morphology are associated with breast cancer metastasis. *Frontiers in Oncology*, 10:50, 2020.

[9] R. Mirzayans, B. Andrais, and D. Murray. Roles of polyploid/multinucleated giant cancer cells in metastasis and disease relapse following anticancer treatment. *Cancers (Basel)*, 10(4):118, 2018.

[10] A. Voulodimos, N. Doulamis, A. Doulamis, et al. Deep learning for computer vision: A brief review. *Hindawi*, 2018(Computational Intelligence and Neuroscience), 2018.

[11] W. Chen, W. Li, X. Dong, and J. Pei. A review of biological image analysis. *Current Bioinformatics*, 12, 2017.

[12] H.E. Munim and A.A. Farag. A shape-based segmentation approach: an improved technique using level sets. *Tenth IEEE International Conference on Computer Vision, Volume 1*, 2:930–935, 2005.

[13] Z. Wang and H. Li. Generalizing cell segmentation and quantification. *BMC Bioinformatics*, 18, 2017.

[14] E. Gómez-de Mariscal, C. García-López-de Haro, W. Ouyang, et al. Deepimagej: A user-friendly environment to run deep learning models in imagej. *Nature Methods*, 18:1192–1195, 2021.

[15] D. Ershov, Minh-Son Phan, J. Pylvänäinen, et al. Bringing trackmate in the era of machine-learning and deep-learning. *bioRxiv*, 2021.

[16] R.F. Laine, I. Arganda-Carreras, R. Henriques, et al. Avoiding a replication crisis in deep-learning-based bioimage analysis. *Nature Methods*, 18:1136–1144, 2021.

[17] A. E. Carpenter, T. R. Jones, M. R. Lamprecht, et al. Cellprofiler: image analysis software for identifying and quantifying cell phenotypes. *Genome Biology*, 7(10), 2006.

[18] S. Berg, D. Kutra, T. Kroeger, et al. ilastik: interactive machine learning for (bio)image analysis. *Nature Methods*, 2019.

[19] J. Marrison, L. Räty, P. Marriott, et al. Ptychography - a label-free, high contrast imaging technique for live cells using quantitative phase information. *Scientific Reports*, 3(2369), 2013.

[20] Y. Rivenson, Y. Zhang, H. Günaydın, et al. Phase recovery and holographic image reconstruction using deep learning in neural networks. *Nature Light Sci Appl.*, 7(17141), 2018.

[21] Y. Park, C Depeursinge, and G. Popescu. Quantitative phase imaging in biomedicine. *Nature Photon*, 12:578–589, 2018.

[22] R. Suman, G. Smith, and K. E.A. Hazel et al. Label free imaging to study phenotypic behavioural traits of cells in complex co-cultures. *Scientific Reports*, 6(1):22–32, 2016.

[23] S. Comsa, A.M. Cimpean, and M. Raica. The story of mcf-7 breast cancer cell line: 40 years of experience in research. *BioMed Central Cancer, Anticancer Research*, 35:3147–3154, 2015.

[24] L. Oakden-Rayner, J. Dunnmon, G. Carneiro, et al. Hidden stratification causes clinically meaningful failures in machine learning for medical imaging. *arXiv*, 2019.

[25] H. Hernandez-Vargas, J. Palacios, and G. Moreno-Bueno. Molecular profiling of docetaxel cytotoxicity in breast cancer cells: uncoupling of aberrant mitosis and apoptosis. *Oncogene*, 26:2902–2913, 2007.

[26] J. Liu, L. Qu, L. Meng, et al. Topoisomerase inhibitors promote cancer cell motility via ros-mediated activation of jak2-stat1-cxcl1 pathway. *Journal of Experimental and Clinical Cancer Research*, 38:370, 2019.

[27] C.L. Liu, M.J. Chen, J.C. Lin, et al. Migration and invasion of breast cancer cells through the upregulation of the rhoa/mlc pathway. j breast cancer. *Journal of Breast Cancer*, 22:185–195, 2019.

[28] N.S. Sohoni, J.A. Dunnmon, G. Angus, et al. No subclass left behind: Fine-grained robustness in coarse-grained classification problems. *CoRR*, 2020.

[29] R. Wang, C. Jin, and X. Hu. Evidence of drug-response heterogeneity rapidly generated from a single cancer cell. *Oncotarget*, 8:25, 2017.

[30] X. Wang, Z. Chen, A.K. Mishra, et al. Chemotherapy-induced differential cell cycle arrest in b-cell lymphomas affects their sensitivity to wee1 inhibition. *Haematologica.*, 103(3):466–476, 2018.

[31] F. M. Frame, A. R. Noble, S. Klein, et al. Tumor heterogeneity and therapy resistance - implications for future treatments of prostate cancer. *Journal of Cancer Metastasis and Treatment*, 3:302–314, 2017.

[32] Y. Ozaki, H. Yamada, H. Kikuchi, et al. Label-free classification of cells based on supervised machine learning of subcellular structures. *PLoS One*, 14(1), 2019.

[33] M. Yang, D.J. Kozminski, L. Wold, et al. Therapeutic potential for phenytoin: targeting nav1.5 sodium channels to reduce migration and invasion in metastatic breast cancer. *Breast Cancer Research and Treatment*, 134(2):603–615, 2012.

[34] C. Uphoff, S. Gignac, and H. Drexler. Mycoplasma contamination in human leukemia cell lines. i. comparison of various detection methods. *Journal of Immunological Methods*, 149:43–53, 1992.

[35] L.K. Soh and C. Tsatsoulis. Texture analysis of sar sea ice imagery using gray level co-occurrence matrices. *IEEE Transactions on geoscience and remote sensing*, 37(2):780–795, 1999.

[36] S. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 7:674–693, 1989.

[37] R.M. Haralick, K. Shanmugam, and I. Dinstein. Textural features for image classification. *IEEE Transactions on systems, man, and cybernetics*, 6:610–621, 1973.

[38] A. Baddeley and R. Turner. spatstat: An R package for analyzing spatial point patterns. *Journal of Statistical Software*, 12(6):1–42, 2005.

[39] A. Haar. Zur theorie der orthogonalen funktionensysteme. *Mathematische Annalen*, 69(3):331–371, 1910.

[40] N.V. Chawla, K.W. Bowyer, L.O. Hall, et al. Smote: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002.

[41] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2019.

[42] W. N. Venables and B. D. Ripley. *Modern Applied Statistics with S*. Springer, New York, fourth edition, 2002.

[43] D. Meyer, E. Dimitriadou, et al. *e1071: Misc Functions of the Department of Statistics, Probability Theory Group (Formerly: E1071), TU Wien*, 2019. R package version 1.7-3.

[44] A. Liaw and M. Wiener. Classification and regression by randomforest. *R News*, 2(3):18–22, 2002.

[45] A. Kassambara and F. Mundt. *factoextra: Extract and Visualize the Results of Multivariate Data Analyses*, 2020. R package version 1.0.7.

Methods

Cell Culture. MDA-MB-231 cells and MCF-7 cells were cultured separately in Dulbecco's modified eagle medium supplemented with 5% fetal bovine serum and 4mM L-glutamine.[33] Fetal bovine serum was filtered using a 0.22μm syringe filter prior to use to reduce artefacts when imaging. Cells were incubated at 37°C in plastic filter-cap T-25 flasks and were split at a 1:6 ratio when passaged. No antibiotics were added to cell culture medium. Cells were confirmed to be mycoplasma-free by 4′,6-diamidino-2-phenylindole (DAPI) method.[34] To image the following day, cells were counted and then seeded in a Corning Costar plastic, flat bottom 24-well plate. Cells were seeded at a density of 8000 cells per well with a final volume of 500μL in each of the 24 wells.

Pharmacology. Docetaxel (Cayman Chemical Company) was prepared as 5mg/mL of DMSO and doxorubicin (AdooQ Bioscience) as 25mg/mL of DMSO, both were then frozen into aliquots. Once thawed, docetaxel and doxorubicin stock solutions were diluted in culture medium to give final working concentrations. Docetaxel dose response analysis for both MDA-MB-231 and MCF-7 cells involved imaging eight wells treated with the following concentrations of docetaxel: 0nM, 1nM, 3nM, 10nM, 30nM, 100nM, 300nM, 1μM, with additional concentrations 3μM, 10μM and 30μ imaged for MDA-MB-231 cells. Doxorubicin dose response analysis for MDA-MB-231 cells involved imaging eight wells treated with the following concentrations of doxorubicin: 0nM, 10nM, 30nM, 100nM, 300nM, 1μM, 3μM, 10μM.

Medium was removed from wells selected to receive treatment 30 minutes prior to image acquisition, and 500μL of desired drug concentration was added to each well. Control wells received a medium change and were treated with DMSO vehicle on the day of imaging to maintain consistent DMSO concentration throughout.

Image Acquisition and Exportation. Cells were placed onto the Phasefocus Livecyte 2 (Phasefocus Limited, Sheffield, UK) to incubate for 30 minutes prior to image acquisition to allow for temperature equilibration. One 500μm x 500μm field of view per well was imaged to capture as many cells, and therefore data observations, as possible. Selected wells were imaged in parallel for 48 hours at 20x magnification with 6 minute intervals between frames, resulting in full time-lapses of 481 frames per imaged well.

Phasefocus' Cell Analysis Toolbox® software was utilised for cell segmentation, cell tracking and data exportation. Segmentation thresholds were optimised for a range of image processing techniques such as rolling ball algorithm to remove background noise, image smoothing for cell edge detection and local pixel maxima detection to identify seed points for final consolidation.

The Phasefocus software outputs a feature table for each imaged well. This table consists of variables that describe the size and orientation of each cell on each frame as well as variables that describe the cell's movement up to the current frame. As some of these variables are highly correlated, for example dry mass and volume, we only include a subset of Phasefocus' features in our analyses to reduce the number of redundant variables. The first nine features in **Supplementary table S1** are those that were retained.

Implementation of CellPhe. The CellPhe toolkit for the calculation of variables from time series data and bespoke feature extraction is implemented in the C programming language and is freely available from the corresponding author. Further analysis, including classification and cluster analysis, were carried out using freely available R packages as detailed and R scripts can be supplied upon request.

New Features. Using the Regions of Interest (ROIs) produced by the Phasefocus software to identify each cell's boundary pixels, a range of additional morphological and texture features were extracted. In addition to shape descriptors calculated from the cell boundaries, a filling algorithm was used to determine the interior pixels from which texture and spatial features were extracted. The local density was also calculated as the proportion of the area, A, around the cell containing pixels from other cells. Here A is the annulus around the cell from its radius r to $r + a$ and a is the average radius of cells in the data set. A complete list of features together with their definitions is provided in **Supplementary table S1**.

Texture descriptors. Gray-level co-occurrence matrices (GLCMs) are widely used in texture analysis to investigate spatial relationships between the pixels with similar intensities.[35] Here, rather than consider the positions of pixels within a cell, we calculated GLCMs between the image of the cell at different resolutions to differentiate textures that are sharp and would be lost at lower resolution from those that are smooth and would remain. This was achieved by performing a two-level 2-D wavelet transform[36] on the pixels within the axis-aligned minimum rectangle containing a cell. GLCMs were then calculated between the original interior pixels and the corresponding values from the first and second levels of the transform as well as between the two sets of transformed pixels. Statistics first described by Haralick[37] were then calculated from each GLCM as follows:

$$\begin{aligned}
\text{Energy} &= \sum_{i,j=1}^{N} M_{ij}^2 \\
\text{Contrast} &= \sum_{n=0}^{N-1} \sum_{i,j=1|i-j|=n}^{N} M_{ij} \\
\text{Homogeneity} &= \sum_{i,j=1}^{N} \frac{M_{ij}}{1+(i+j)^2} \\
\text{Correlation} &= \sum_{i,j=1}^{N} M_{ij} \left(\frac{(i-\bar{i})(j-\bar{j})}{\sigma_i \sigma_j} \right) \\
\text{Entropy} &= -\sum_{i,j=1}^{N} M_{ij} \log\left(M_{ij}\right)
\end{aligned}$$

where M_{ij} is the (i, j)th entry in the normalised GLCM and $\bar{i}, \bar{j}, \sigma_i$ and σ_j are the means and standard deviations of the marginal densities $M_i = \sum_j^N M_{ij}$ and $M_j \sum_i^N M_{ij}$ respectively. With three co-occurrence matrices, this added 15 texture features.

Spatial distribution descriptors. We calculated spatial distribution descriptors to quantify the uniformity or clustering of cell interior pixels with intensities above 9 different thresholds. For IQn,

interior pixels with intensities greater than or equal to the $(n\times10)$th quantile are modelled as a spatial point process using the it ppp function in the R *spatstat* package.[38] The cell boundary coordinates form the boundary window of each spatial point process and we make use of Ripley's reduced second moment K function to compare the empirical distribution of pixels with a theoretical Poisson distribution. Each spatial distribution descriptor is then defined as follows:

$$\mathrm{IQ}n = K_{emp}(maxr) - K_{theo}(maxr)$$

where $maxr$ is one quarter of the smallest side of the enclosing rectangle within the cell boundary. The K functions, K_{emp} and K_{theo}, are evaluated for our empirical data and for a theoretical Poisson distribution respectively, such that:

$$K_{emp}(r) = \frac{a}{(n \times (n-1))} \times \sum_{i,j} I(d(i,j) \leq r),$$

and

$$K_{theo}(r) = \pi r^2,$$

where a is the area of the cell boundary, n is the number of pixels being considered, and the sum is taken over all ordered pairs of pixels i and j. $I(d(i,j) <= r)$ is an indicator variable that equals 1 if the distance $d(i,j)$ between pixels i and j is $\leq r$.

A negative IQn value indicates that the pixels are uniformly distributed, whereas positive indicates clustering of points, with large absolute values of IQn indicative of greater uniformity or clustering.

Summarising Cell Time Series. Cell tracking provides a time series for each of the 47 features extracted for a cell. The length of the time series depends on how many frames the cell has been tracked for and so differs between cells. In order to apply pattern recognition methods, we extracted a fixed number of characteristic variables for each cell from the time series for each feature. Statistical measures (mean, standard deviation and skewness) summarise time series of varying length, but may not be representative of changes throughout the time series. Therefore, in addition to summary statistics, we calculated variables inspired by elevation profiles in walking guides, that is, the sum of any increases between consecutive frames (total ascent), the sum of any decreases (total descent) and the maximum value of the time series (maximum altitude gain). These "elevation" variables were calculated for different levels of the wavelet transform of the time series to allow changes at different scales to be considered. The wavelet transform decomposes a time series to give a lower resolution approximation together with different levels of detail that need to be added to the approximation to restore the original time series. Using the Haar wavelet basis[39] with the multiresolution analysis of Mallat[36] allows increases and decreases in the values of the variables to be determined over different time scales.

Occasionally the automated cell tracking misses a frame or even several frames, for example when a cell temporarily leaves the field of view. To prevent jumps in the time series, we interpolated values for the missing frames, although these values were not used to calculate statistics. After interpolation, the three elevation variables were calculated from the original time series and three wavelet levels which, together with the summary statistics, provided 15 variables for each feature **Supplementary table S2**. This would have given 47 x 15 = 705 variables in total, but, as one feature, the tracklength or total distance travelled up to the current frame, is monotonically increasing, the total descent is always zero and therefore the 4 descent variables were not used.

One further variable was introduced to summarise cell movement as the area of the minimal bounding box around a cell's full trajectory. This area will be large for migratory cells and small for cells whose movement remains local for the duration of the time series. If, within a cell's trajectory, $minX$ and $minY$ are the minimal X and Y positions respectively with $maxX$ and $maxY$ the corresponding maximal positions, then the trajectory area is defined as

$$\text{trajectory area} = (maxX - minX) \times (maxY - minY).$$

Thus, a total of 702 characteristic variables were available for analysis and classification.

Segmentation Error Removal. To improve characterisation of cellular phenotype, we only included cells that were tracked for at least 50 frames in our analyses. Whilst the majority of these cells were correctly tracked, others had segmentation errors, with confusion between neighbouring cells, missing parts of a cell or multiple cells included.

In order to increase the reliability of our results, we developed a classification process to identify and remove such cells prior to further analysis. Cells (both treated and untreated) were classified by eye to provide a training data set. Due to class imbalance, with the number of segmentation errors far less than the number of correct segmentations, the Synthetic Minority Oversampling Technique (SMOTE)[40] was performed using the *smotefamily* package in R, with the number of neighbours K set to 3, to double the number of instances representing segmentation errors.

The resulting data set with all 702 variables was used to train a set of 50 decision trees using the *tree* package in R with default parameters. For each tree, the observations from cells with segmentation errors were used together with the same number of observations randomly selected from the correctly segmented cells to further address class imbalance. For each cell, a voting procedure was used to provide a classification from the predictions of the 50 decision trees. To minimise the number of correctly tracked cells being falsely classified as segmentation errors, this class was only assigned when it received at least 70% of the votes (i.e. 35). To add further stringency, the training of 50 decision trees was repeated ten times and a cell only given a final classification of segmentation error if predicted this label in at least five of the ten runs. MDA-MB-231 cells that were not used for training formed an independent test set. All cells either manually labelled as segmentation error or predicted as such were excluded from further analyses.

Classification of Untreated and Treated Cells. After removing segmentation errors, the remaining data were used to form training and test sets for the classification of untreated and treated cells. Training sets were balanced prior to classifier training to mitigate bias and data from cells in the independent test sets were never used during training.

A separate classifier was trained for each cell line - treatment combination, as shown in **Table 3** and feature selection performed to determine the most appropriate variables in each case. Each variable was assessed using the group separation, $S = V_B/V_W$, where V_B is the between-group variance:

$$V_B = \frac{n_1(\bar{x}_1 - \bar{\bar{x}})^2 + n_2(\bar{x}_2 - \bar{\bar{x}})^2}{(n_1 + n_2 - 2)}$$

and V_W is the within-group variance:

$$V_W = \frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{(n_1 + n_2 - 2)}.$$

Here n_1 and n_2 denote the sample size of group 1 and group 2 respectively, $\bar{x}_1$ and $\bar{x}_2$ are the sample means, $\bar{\bar{x}}$ the overall mean, and ${s_1}^2$ and ${s_2}^2$ are the sample variances. The most discriminatory variables were chosen for a particular data set by assessing the classification error on the training data to optimise the threshold on separation. Starting with a threshold of zero, the n^{th} separation threshold was maximised such that the classification error rate did not increase by more than 2% from that obtained for the $(n-1)^{\text{th}}$ threshold. The aim here was to reduce the risk of overfitting by only retaining variables achieving greater than or equal to this threshold for the next stage of classifier training.

Data were scaled to prevent large variables dominating the analysis and ensemble classification used to take advantage of different classifier properties. The predictions from three classification algorithms, Linear Discriminant Analysis (LDA), Random Forest (RF) and Support Vector Machine (SVM) with radial basis kernel were combined using the majority vote. Model performance was evaluated by classification accuracy, taking into account the number of false positives and false negatives. All classification was performed in RStudio[41] using open-source packages. LDA was performed using the *lda* function from the *MASS* library,[42] SVM classification used the *svm* function

from the *e1071* package[43] with a radial basis kernel and the *randomForest* package[44] was used to train random forest classifiers with 200 trees and 5 features randomly sampled as candidates at each split.

Cluster analysis. Both hierarchical clustering and k-means clustering were used to investigate subgroups within single-class data sets (i.e. treated and untreated cells separately). Data were scaled prior to clustering and analyses performed in R. Hierarchical clustering was implemented with the *factoextra* package[45] using the *hcut* function to cut the dendrogram into k clusters. Agglomerative nesting (AGNES) was used with Ward's minimum variance as the agglomeration method and the Euclidean distance metric to quantify similarity between cells. k-means clustering was performed using the R *stats* package, with the number of random initial configurations set to 50. The number of clusters k was chosen to obtain clusters with meaningful interpretation. Similarities and differences between clusters were identified through evaluation of separation scores to determine discriminatory features, as well as through observation of cells within each cluster by eye.

Data. Three data sets are used to demonstrate our pipeline for the classification of untreated and treated cells. For brevity we use abbreviations throughout to refer to each data set, for example "231Docetaxel" is a data set consisting of MDA-MB-231 cells, both untreated and treated with 30μM docetaxel. This is the main data set used to develop the methods, with a training data set compiled from 6 experiments performed on different days and an independent test data set compiled from a further 3 experiments, also performed on separate days and by a different individual.

We validate our methods using two further datasets, the 231Doxorubicin and MCF7Docetaxel data sets, details of which are given in **Table 3**. This table also includes details of the number of cells within each training and test set. We show that the classification pipeline can be successfully reproduced using fewer experimental repeats for the 231Doxorubicin and MCF7Docetaxel data sets. The 231Doxorubicin training set consists of data from one experiment with a further, independent experiment performed on a separate day used as a test set. Training and test sets for MCF7Docetaxel are from the same experiment, with random sampling used to produce independent training and test sets. Each training data set contains a balanced number of untreated and treated cells, treated with a single drug concentration. We selected 30μM docetaxel and 1μM doxorubicin for the experiments with MDA-MB-231 cells as the optimal doses with which to induce changes in cell morphology and migration without inducing cell death. However, a lower concentration (1μM) of docetaxel was used for MCF-7 cells as we found that this induced similar morphological and dynamical changes to those induced by higher concentrations but with reduced cell death (**Table 3**).

Data set	**Cell line**	**Treatment**	**Training set**	**Test set**
231Docetaxel	MDA-MB-231	30μM Docetaxel	Untreated: 646 Treated: 600	Untreated: 939 Treated: 314
231Doxorubicin	MDA-MB-231	1μM Doxorubicin	Untreated: 216 Treated: 216	Untreated: 185 Treated: 59
MCF7Docetaxel	MCF-7	1μM Docetaxel	Untreated: 130 Treated: 130	Untreated: 304 Treated: 27

Table 3: The three data sets used in this study with the number of cells in training and test sets used for untreated vs treated classification.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- CellPhemanual.pdf
- CellPhesupplementary.pdf
- WilsonCodeFlat.pdf
- WilsonEPCflat.pdf
- WilsonSuppInfo.pdf