Megan L. Gelsinger¹ / Laura L. Tupper² / David S. Matteson¹

Cell Line Classification Using Electric Cell-Substrate Impedance Sensing (ECIS)

¹ Department of Statistics and Data Science, Cornell University, NY Ithaca, USA, E-mail: mlg276@cornell.edu, matteson@cornell.edu. <https://orcid.org/0000-0002-7573-5592>.

² Department of Mathematics and Statistics, Williams College, MA Williamstown, USA, E-mail: llt1@williams.edu

Abstract:

We present new methods for cell line classification using multivariate time series bioimpedance data obtained from electric cell-substrate impedance sensing (ECIS) technology. The ECIS technology, which monitors the attachment and spreading of mammalian cells in real time through the collection of electrical impedance data, has historically been used to study one cell line at a time. However, we show that if applied to data from multiple cell lines, ECIS can be used to classify unknown or potentially mislabeled cells, factors which have previously been associated with the reproducibility crisis in the biological literature. We assess a range of approaches to this new problem, testing different classification methods and deriving a dictionary of 29 features to characterize ECIS data. Most notably, our analysis enriches the current field by making use of simultaneous multi-frequency ECIS data, where previous studies have focused on only one frequency; using classification methods to distinguish multiple cell lines, rather than simple statistical tests that compare only two cell lines; and assessing a range of features derived from ECIS data based on their classification performance. In classification tests on fifteen mammalian cell lines, we obtain very high out-of-sample predictive accuracy. These preliminary findings provide a baseline for future large-scale studies in this field.

Keywords: biophysics, classification analysis, supervised learning

DOI: 10.1515/ijb-2018-0083

Received: August 20, 2018; **Revised:** September 10, 2019; **Accepted:** November 13, 2019

1 Introduction

Bioelectrical impedance cell culture platforms, such as ECIS, offer scientists the valuable opportunity to study a variety of cell characteristics, such as attachment, growth, and death, in a non-invasive, real-time environment [1]. Although ECIS technology has been in use since the 1990's, a major application of ECIS data has not yet been formally explored: its potential in identifying unknown or mislabeled cells. Better cell line identification can help address the current scientific reproducibility crisis that results, in part, from the use of misidentified cells in biological research. We introduce the foundations for such a tool, leveraging the multivariate data provided by ECIS technology and insights gained through statistical classification analysis. While the majority of this manuscript is devoted to our analysis, we begin with a brief introduction to ECIS and to previous work using ECIS technology and data.

In the 1990's, ECIS was introduced as a label-free, real-time, electrical impedance-based method for studying cell behavior in tissue culture [2]. In the introductory paper, researchers provided both the models and a basic blueprint for the device needed to implement ECIS. The basic structure they proposed remains in use today. It consists of a typical cell culture plate, in which each well is fitted with a gold electrode, allowing for the passage of electrical currents at various AC frequencies. After application of such currents, the device measures the electrical impedance associated with each current as a function of time and divides it into its components, resistance and capacitance, allowing for live monitoring of cell activity.

In seeking to expand the scope of ECIS technology to include cell identification and classification, we continue a long history of finding innovative ways to use the originally proposed design [2]. Since its introduction, researchers in biotechnology have studied ways of improving the overall efficiency and effectiveness of ECIS devices through experimenting with the design scheme [3]. Others in the field have constructed specially tuned versions of an ECIS device for specific applications in cancer cell and clinical diagnostic research [4, 5].

Other studies have used the technology to monitor the activity of cells grown in tissue culture. Early research focused on showing that ECIS measurements could be associated with known biological characteristics of a particular cell line from the attachment and spreading phases of growth [6, 7]. More recently, others have

Megan L. Gelsinger is the corresponding author.

© 2020 Walter de Gruyter GmbH, Berlin/Boston.

used ECIS to study a cell line before and after an event, such as exposure to cytotoxins, and evaluate induced differences [8, 9]. A more comprehensive overview of the use of ECIS in the biological literature has been detailed elsewhere [10].

Previous studies have focused on only one or two cell lines at a time, and their analysis has neglected the full potential of multiple frequency-time measurements available through a standard ECIS device. One group considered two different cell lines, one cancerous and one healthy, but their analysis relied on only one frequency of data and did not extend to additional cell lines [11]. Others considered multiple frequencies in their analysis of cytotoxic effects on cell morphology, but they limited their use of the multivariate data to a single time index and only one cell line [8]. While these approaches have been sufficient for the scope of studies to date, we consider the full breadth of information available from parallel measurements at multiple frequencies. This additional information allows new types of analysis, including new approaches for cell identification.

The main application of ECIS data we consider is cell line classification. Biological studies often rely upon cultured mammalian cells, which are known to mutate or become misidentified during the life of an experiment. When such anomalies go unnoticed, studies report erroneous results, contributing to the multi-billion dollar irreproducibility problem [12]. We focus on ECIS, rather than alternative technologies such as short tandem repeat (STR) analysis, due to ECIS' advantages in terms of time and money. ECIS data requires only 24–48 hours to collect and analyze, as opposed to the weeks required by alternative methods that involve shipping out samples for analysis at outside facilities. In the long run, ECIS devices also prove cost effective: after an up-front purchase of the device (\$3000–\$5000), the user must only invest \$10 per test for a disposable tray. In contrast, STR costs around \$100 for each test, in addition to the cost of lost time waiting for the test results [13].

While previous studies, such as those cited above, have suggested specific features to help characterize single cell lines using ECIS data, none have employed statistical techniques to quantify the ability of potential characteristic features to classify multiple cell lines at once; the typical statistical analysis in these studies only considered *t*-tests and *F*-tests to measure differences in the mean value of a specific feature across populations. This is fundamentally different than performing classification based on a feature, as the latter technique simultaneously measures the characteristic across all populations, and quantifies how well it separates and identifies the groups. Our work determines which features of multi-frequency ECIS data are most effective in a given classification problem; this also suggests features that may be most useful in future analyses.

While we ultimately aim to develop an integrated system to perform cell line classification, that will require much more data collection, and we limit the scope of this manuscript to a preliminary analysis. We extract many previously considered cell-specific ECIS characteristics from a standard ECIS® device, measure them across both time and AC frequency, and quantify their classification performance through established statistical techniques, including classification trees; regularized discriminant analysis (RDA); and two special cases of RDA, linear discriminant analysis (LDA) and quadratic discriminant analysis (QDA).

The remainder of this paper provides the details of our study. In Sections 2 and 3, we describe our dataset and feature specification, respectively. In Section 4, we discuss the classification analysis results, and in Section 5 we offer conclusions, and discuss avenues for future work.

2 ECIS data and experimental design

The dataset considered throughout this paper was provided by Applied BioPhysics, the proprietor of ECIS®, as recorded by one of their 96-well ECIS® Zθ devices. It consists of fourteen replicates of fifteen different mammalian cell lines grown on two different serums (gel and bovine serum albumin (BSA)). The data will be grouped (and compared) by serum type in our forthcoming analysis. Each plate was exposed to nine different electrical current frequencies over the course of 20 hours, resulting in 64 approximately evenly spaced measurement times. The fifteen different cell lines are: HUTU80, CHOK1, MDCKII, LCELL, WI-38, VA-13, NRK, PPC-1, DuPro, Glomotel, BSC1, DU145, H1299, NIH3T3, and PC-3. The nine AC frequencies measured are: 250, 500, 1000, 2000, 4000, 8000, 16000, 32000, and 64000 Hz. Each 96-well plate held three different cell lines; for each line, we considered data from one plate. Impedance measurements were recorded approximately once every 20 minutes across all nine AC frequencies (see Figure 1 (B)). The impedance measurements were decomposed into resistance and capacitance according to the originally proposed model [2].

Figure 1: Visualizing ECIS Data: Resistance measurements as recorded by an ECIS® device. (A) A single response for each of three different cell lines, each exposed to two different AC frequencies. (B) Response of the H1299 cell line to all nine AC frequencies.

Notationally, we can represent the data as follows: $Y_i^{k,s}(t_\ell, \nu_r)$. Here, Y is the resistance value at the given parameter settings. We enumerate all cell cultures in the dataset, regardless of type, with an overall index $i = 1, \dots, 420$. Next, $k = 1, \dots, 15$ represents the index of the cell line in the vector (HUTU80, CHOK1, ..., PC-3), as listed above, and $s = 1, 2$ represents the serum (gel, BSA). We note that the cell line and serum type are implicitly encoded in the overall index, i ; likewise, we may suppress k and s without loss of information. The data depends functionally on t_ℓ the discretized time index, with $\ell = 1, \dots, 64$ as described above. It also depends functionally on the discretized input frequency ν_r , with $r = 1, \dots, 9$, corresponding to the index of the frequency vector (250, 500, ..., 64000), as listed above. This notation helps us understand subsets of the data, which we will reference when describing our classification features. For example, $\{Y_i^{11,2}(t_9, \nu_5)\} \equiv Y_{1:14}(t_9, \nu_5)$ corresponds to the resistance of the BSA-treated BSC1 cells at time index nine (hour three) and 4000 Hz.

3 Feature specification

3.1 Technical review

Many of the characteristics identified in the literature as varying by cell line or cell condition were extracted from data obtained during confluence, or the steady-state portion of growth following the initial growth and spreading stages. During this phase, the well is completely covered with cells, allowing for only minute movements. Data collection commonly began after 20–24 hours, once cells had reached confluence, and continued for about 20 hours at a fine sampling frequency, ranging from every two minutes to every second depending on the study. With this resolution of data, researchers were able to extract features such as:

- the slope parameter β characterizing a least-squares straight-line fit to a log-log plot of power spectrum versus frequency ν [7–9, 11, 14]; e.g. Brownian noise displays an ν^{-2} power law, with $\beta = 2$;
- the sample kurtosis of the first differenced resistance time series, a measure of distributional shape [11];
- the first e^{-1} crossing of the autocorrelation function of the noise time series from confluence onward (extracted from the resistance time series over the same time period), to estimate its exponential decay rate [9, 11].

Two other post-confluence features suggested in previous works were R_b and α . R_b reflects the barrier resistance between cells, whereas α reflects the constraint on current flow beneath the cells (see Figure 2 for visualization). R_b and α are defined by the following relationships:

Figure 2: Visualizing the ECIS Model: Image obtained from Applied BioPhysics' website (biophysics.com). In the ECIS model, cells are viewed as disks, characterized by parameters R_b , α , r_c , C_m , and h .

$$\frac{1}{Z_c(\nu)} = \frac{1}{Z_n(\nu)} \left(\frac{Z_n(\nu)}{Z_n(\nu) + Z_m(\nu)} + \frac{\frac{Z_n(\nu)}{Z_n(\nu) + Z_m(\nu)}}{\frac{\gamma r_c}{2} \frac{I_0(\gamma r_c)}{I_1(\gamma r_c)} + R_b \left(\frac{1}{Z_n(\nu)} + \frac{1}{Z_m(\nu)} \right)} \right); \quad (1)$$

$$\gamma r_c = r_c \sqrt{\frac{\rho}{h} \left(\frac{1}{Z_n(\nu)} + \frac{1}{Z_m(\nu)} \right)} = \alpha \sqrt{\frac{1}{Z_n(\nu)} + \frac{1}{Z_m(\nu)}}. \quad (2)$$

For frequency ν , $Z_c(\nu) > 0$ is the impedance (per unit area) of the cell-covered electrode, $Z_n(\nu)$ is the impedance of the cell-free (empty, reference) electrode, $Z_m(\nu)$ is the membrane impedance of the cells, r_c is the radius of the cell, ρ is the resistivity of the solution, h is the height of the space between the ventral surface of the cell and the substrate, and I_0 and I_1 are the modified Bessel functions of the first kind of order 0 and 1 [2].

An ECIS[®] device is incapable of recording all of the parameters in eqs. (1)–(2). Instead, it records only $Z_c(\nu)$ and $Z_n(\nu)$ over a set of frequencies $\{\nu_1, \dots, \nu_w\}$, with $w \geq 3$, and estimates the values of $Z_m(\nu)$, R_b and α which fit the data best. Since these optimal values cannot be computed analytically, the ECIS[®] software uses the Nelder-Mead downhill simplex method to select a series of parameter values. From there, the program calculates the mean squared error between the measured and modeled frequency scan values, and returns the set of parameter values which minimize this error. The absolute error is normalized to the cell-free reference, so the frequencies are appropriately weighted. It is important to note that underlying these equations and calculations is the assumption that the ECIS model can be understood as a series resistor-capacitor circuit, with cells represented by circular discs hovering above (not directly touching) a gold electrode. We consider R_b and α as characteristics which may differentiate cell lines, as they were considered in other previous studies [2, 8].

Several previous studies also evaluated features that were specific to the attachment and spreading phases of growth. In particular, some researchers found that certain cell lines “peaked higher” and “increased more rapidly” during these initial stages of growth [15, 16]. Additional remarks are very limited as the majority of prior studies concentrated less on attachment and spreading and more on confluence behavior.

The statistical analysis involved in these studies focused on t -tests and F -tests to look for significant differences in a feature's mean value across populations of interest. For example, when one group considered whether a particular cell population inoculated with varying degrees of a cytotoxin had differing mean R_b values for each dosage level, they performed Student's t -test for two samples with unequal variance, for pairs of dosages, such as 5 μ M versus 10 μ M of cytotoxin [8]. Another example came from the comparison of cancerous and noncancerous cells, in which an F -test was conducted to assess the likelihood of the two cell lines coming from the same distribution, given the means and variances of their power slopes β [11].

3.2 Feature space specification

In our study, we included R_b and α as our confluence phase features. Computationally only one time index worth of data during the confluence phase is needed to approximate these values, making them feasible for our analysis; we obtained twelve observations per frequency during this phase of growth, following the sampling scheme described previously. Typically, when using a single time index, the data from all measured frequencies are combined to estimate R_b and α . Since we had twelve time indices worth of data, though, we averaged the R_b and α at each of these twelve points to get more stable estimates of R_b and α .

We note that since the design for our study did not emphasize the confluence stage of growth, we did not extract any of the other confluence region features mentioned in Section 3.1. Studies which did look at these additional features had thousands of observations available for a single frequency during confluence for their analysis, compared to our twelve, leading to our decision to concentrate on R_b and α .

To extend our feature space, we also included proxies for those features from the attachment and spreading portions of cell growth mentioned in Section 3.1 using the $Z_c(\nu)$ data collected from the ECIS[®] device. These included, at each frequency ν_r :

- i. End of run resistance value: $\frac{1}{5} \sum_{\ell=60}^{64} Y_i(t_\ell, \nu_r)$ (EOR);
- ii. Maximum resistance of moving average smoothed time series value: $\max_k \left[\frac{1}{5} \sum_{\ell=k-2}^{k+2} Y_i(t_\ell, \nu_r), \text{ for } k = 3, \dots, 62 \right]$ (MR);

iii. Resistance at two hours (t_6): $\frac{1}{5} \sum_{\ell=4}^8 Y_i(t_\ell, \nu_r)$ (R2h);

where smoothing was leveraged for stability purposes.

The utility of using multi-frequency data in our analysis stems from the physics underpinning the system. We note that for the cell layer, the membrane capacitance and the solution resistance are fixed values and not frequency dependent. The path the current will flow, though, is frequency dependent. Hence, changing the frequency changes the path through which it will flow, and thus the overall impedance of the system ($Z_c(\nu)$). It is known that at lower frequencies, more current flows under and between cells, making it easier to quantify the resistance between the cells. At higher frequencies, the reactance of the cell membrane is low, with the majority of the current flowing capacitively through the membrane, allowing for an easier exploration of the cell-substrate interaction. Likewise, $Z_c(\nu)$ at different frequencies captures different characteristics of a cell's morphology, making multi-frequency data potentially more useful for classification than data from just a single frequency. It is also important to keep in mind the presence and effect of the gold electrode that the cells are grown on. In particular, at low frequencies the gold interface resistance becomes very large, potentially contributing a large portion of the overall resistance of the system. For these reasons, we use a range of AC frequencies to create a more accurate picture of the various sources of impedance and resistance in the system, especially those that relate to the cell membrane [2].

Previous studies involving these features evaluated them at only one frequency, typically 4000 Hz. One of the main advantages of ECIS[®] Z θ , though, is its ability to record multi-frequency data. Given the physical justification for using this multi-frequency data, we include relevant features (EOR, MR, R2h) at all nine measurement frequencies given in Section 2. This gives us a total of $(3 \times 9) + 2 = 29$ characteristics in our feature space which we utilize in our classification analysis below. In particular, we will investigate whether single, pairs of, and/or trios of these features applied to various classification techniques sufficiently differentiate the cell lines. We focus on this collection of features as we can assess their performance both quantitatively, through numeric output, and qualitatively through visualizations.

4 ECIS-based cell line classification analysis

Classification algorithms provide predictive accuracy rates which reflect quantitatively how well certain characteristics differentiate the data. In previous ECIS-related works, none of these algorithms were used to assess features for general cell line classification; instead, F -tests and t -tests were relied upon to determine whether the mean value of the feature in question differed by group. We consider a variety of classifiers in our study to provide this missing quantitative evaluation of potential features for cell line classification.

Classification analysis was performed using several different supervised learning methods on several different combinations of features. The first "grouping" of features considered was similar to that of previous works, where each feature was analyzed on an individual basis as a tool for cell line differentiation. We also considered all pairs and trios of our 29 features to determine if any provided better separation of the data, and if so, by how much. All combinations of features were evaluated using both classification trees and RDA (with LDA and QDA as special cases).

As mentioned in Section 2, we examined those cells inoculated in gel and BSA serums separately looking to gain insight on the conditions which promote optimal cell line classification. The analysis that follows provides an in-depth description and visualization of the classification procedures performed on the gel-based cells. The same procedures were performed on BSA-based cells; we provide a highlighted summary of these results and comparison to the gel-based cell results at the conclusion of this section.

4.1 Classification trees

We began our analysis of the gel-based cells with classification trees, as they use the most elementary division of the feature space to classify observations into groups. Classification trees perform an iterated series of binary splits on the data space to create several regions, then ultimately assign each observation in a given region to the most commonly occurring class of training observations in that region. To create these splits, we select the predictor, or feature, X_j and cutpoint b such that splitting the feature space into regions $\{X|X_j < b\}$ and $\{X|X_j \geq b\}$ leads to the greatest possible improvement in a particular metric. We note that based on this definition, the splits in classification trees divide the feature space into a series of rectangles in two-dimensional space.

One of the most common metrics for assessing the quality of a particular split is cross-entropy or deviance, which is defined for a particular region m by

$$D_m = - \sum_{k=1}^K \hat{p}_{mk} \log \hat{p}_{mk}.$$

Here we assume the data can be divided into K classes, with $\hat{p}_{mk}$ representing the portion of training observations in the m th region that are from the k th class. Cross-entropy, therefore, represents a measure of node purity. A region m is “pure” when D_m assumes a small value. This occurs when $\hat{p}_{mk}$ is near zero or one for each k , meaning nearly none or nearly all of the training observations in the m th region are from the k th class, respectively. So, to construct a binary split, we consider all predictors $X_1, \dots, X_c$ at all possible values b , and select the feature and cutpoint such that the sum of the D_m ’s associated with the two new regions generated by the split is minimized. A classification tree is built by iteratively performing these splits on training data until reaching a specified node purity threshold or minimum number of observations per node. A test set can then be classified using the tree, and a predictive accuracy rate computed [17].

Traditionally, all potential features are passed to the classification tree algorithm, which selects those features and associated splits which lead to optimal node purity. It is often the case in this traditional approach, though, that more than one, two, or three features appear in the splitting rules for the resultant tree, given there are at least that many features provided during the construction phase. Given our desire to provide a visual accompaniment to our quantitative results, we opt for a non-traditional approach to constructing our classification trees, only providing them with one, two, or three potential features. This way we can visualize the resultant splits of the features in 1D, 2D, or 3D space.

In our study, $K = 15$, corresponding to the 15 different cell lines in the dataset. We construct all possible trees that use one, two, or three of the features described in Section 3.2, such as EOR@2000Hz ($(1/5) \sum_{\ell=60}^{64} Y_i(t_\ell, \nu_4)$) and R2H@16000Hz ($(1/5) \sum_{\ell=4}^8 Y_i(t_\ell, \nu_7)$). For example, one single-feature tree is based on $X_1 = \text{EOR@2000Hz}$; each split in the tree is based on an optimal cutpoint corresponding to a particular value of this feature. As there are 29 features, we obtain 29 single-feature trees. Next we build trees using each pair of features: in one such tree, for example, we allow node splits on either $X_1 = \text{EOR@2000Hz}$ or $X_2 = \text{R2H@16000Hz}$. We construct $\binom{29}{2} = 406$ such trees, and similarly $\binom{29}{3} = 3654$ trees involving three features.

The training data used to grow the trees is obtained by randomly dividing the entire dataset for each cell line in half, using seven of the fourteen available observations as the training set and the remaining seven as the test set. We construct the classification trees using the `tree()` function in R on the training set and assess the out-of-sample classification predictive accuracy rate using `predict.tree()` on the test set. We repeat this analysis on twenty different random splits of the data into training and testing sets to lessen the effects of overfitting to any particular sample. Reported accuracy rates for each feature, feature pair, or trio of features are the average over all twenty of these trials.

Table 1 reflects the best feature(s) for classification based on the average accuracy rate. We note that the features highlighted in Table 1 are the best by a small margin; at least 10 % of all the single, pairs and trios of features were within 5 % of the best respective classification rate. This demonstrates that we have not only one ultimate best feature, but a set of high-performing features that we could use to differentiate the cell lines; it is possible that one such feature in the set could be more easily obtainable than the others, making it more ideal to use than the designated “best.” This group of high-performing features is important to consider when interpreting a further implication of Table 1 — the results apply to more than just the features highlighted in Table 1; they apply to the entire class of high performing features which cannot be reproduced for space’s sake. The table reveals a large increase in the out-of-sample classification accuracy when using an informative pair of features as opposed to an informative single feature. There is a less marked improvement when using a strong trio of features versus a strong pair of features. Lastly, the table suggests that two or three informative features are sufficient to capture the group dynamics of fifteen different cell lines, as the accuracy rates exceed 90 %.

Table 1: Classification Tree Analytical Results: Best feature(s) for cell line classification based on twenty trials of classification tree construction using cells grown on gel serum. Abbreviations are EOR: End of run resistance; MR: Maximum resistance; R2h: Resistance at two hours. The out-of-sample classification rate is the average over all twenty values obtained for each feature. The standard error of this average value is reported in column four. These results show that when classifying cell lines, a 2D feature space is significantly more informative than a 1D feature space. Using three features as opposed to two results in a less dramatic improvement in the out-of-sample classification rate, but an improvement nonetheless.

Feature Space Dimension	Selected Classification Feature(s)	Out-of-Sample Classification Accuracy	Approximate Standard Error
1	EOR @ 2000 Hz	0.717	0.009
2	R2h @ 32000 Hz	0.952	0.005
3	MR @ 32000 Hz R2h @ 32000 Hz MR @ 16000 Hz Rb	0.968	0.004

Figure 3 offers a visualization of the best pair of features proposed in Table 1, and their ability to separate the data with respect to classification trees. In panel (A), we see that most of the cells from the same line (represented by points of the same color) reside in the same rectangular region, indicating correct classification. This behavior reflects the high out-of-sample classification rate (95 %) observed in Table 1.

Figure 3: Classification Tree Visual Results: Best pair of features for cell line classification as determined through classification trees using cells grown on gel serum. (A) Example of training/testing set used to conduct classification analysis. Training points are represented by hollow circles, test points represented by x's. Most observations of the same color lie in the same rectangular region, indicating good separation. This behavior corresponds to the high classification accuracy (95 %) presented in Table 1. (B) Boundaries created when using all 14 observations in training set (full information).

4.2 Linear and quadratic discriminant analysis

While classification trees, and their orthogonal divisions of the feature space, yielded high out-of-sample classification accuracy, we wanted to see if a more flexible separation of the feature space could produce even better results. Linear and quadratic discriminant analyses are appropriate extensions, as they separate the feature space with linear (not necessarily orthogonal to an axis) and quadratic divisions, respectively. As both LDA and QDA are special cases of RDA, we reserve the details of their results to Section 4.3 where we discuss RDA at length; this section is meant to provide the mathematical formulation of LDA and QDA for those unfamiliar with the methods. As a brief note, a normality assumption is associated with each of the following discriminant methods — that the feature data within each cell line class is normally distributed. In practice, it has been found that these methods are robust to mild violations of this assumption, i.e. so long as the data is not heavily skewed or multimodal [18]. We graphically assessed our data for approximate normality within each class and found no major violations.

Formally, LDA assumes that observations X of the k th class are drawn from the multivariate Gaussian distribution with mean vector μ_k and covariance matrix Σ : ($X \sim \text{Normal}(\mu_k, \Sigma)$). Note the lack of subscript on the covariance matrix; LDA assumes that all K classes have the same covariance structure Σ , a characteristic which distinguishes this method from other classifiers (see QDA below). The LDA algorithm itself is based on Bayes's theorem. Assuming a Gaussian density for the k th class, the Bayes classifier assigns an observation $X = x$ to the class κ for which the posterior probability of x belonging to the k th class is maximized. This corresponds to assigning an observation $X = x$ to the class κ for which

$$\kappa_{\text{Bayes}} = \arg \max_k \delta_k(x) = \arg \max_k [x^T \Sigma^{-1} \mu_k - \frac{1}{2} \mu_k^T \Sigma^{-1} \mu_k + \log \pi_k],$$

where π_k is the prior probability of an observation belonging to the k th class. The LDA algorithm estimates μ_k , Σ and π_k from the training set, which yields estimates for $\hat{\delta}_k(x)$. Typically, $\hat{\mu}_k$ and $\hat{\Sigma}$ are traditional sample estimates of the mean and covariance, the latter of which might have a scaling constant to mitigate bias. The prior π_k is usually estimated as the proportion of the training data that belongs to the k th class. Likewise, when the sample proportions of each class are the same, as is the case in our study, the $\log \hat{\pi}_k$ terms cancel, leading to

$$\kappa_{\text{LDA}} = \arg \max_k \hat{\delta}_k(x) = \arg \max_k [x^T \hat{\Sigma}^{-1} \hat{\mu}_k - \frac{1}{2} \hat{\mu}_k^T \hat{\Sigma}^{-1} \hat{\mu}_k].$$

An observation $X = x^*$ from the testing set is then assigned to the κ th class if $\kappa = \arg \max_k \hat{\delta}_k(x^*)$. The predicted classes using this scheme can then be compared to the true classes to assess classification accuracy [17].

QDA is similar to LDA; QDA assumes that observations X of the k th class are drawn from the multivariate Gaussian distribution with mean vector μ_k and covariance matrix Σ_k : ($X \sim \text{Normal}(\mu_k, \Sigma_k)$). Here too the Bayes classifier assigns an observation $X = x$ to the class κ for which the posterior probability of x belonging to the k th class is maximized. This corresponds to assigning an observation $X = x$ to the class κ for which

$$\kappa_{\text{Bayes}} = \arg \max_k \delta_k(x) = \arg \max_k \left[-\frac{1}{2}(x - \mu_k)^T \Sigma_k^{-1}(x - \mu_k) + \log \pi_k \right].$$

Notice here that each of the $k = 1, \dots, K$ classes has its own distinct covariance function, and that the observation x appears in a quadratic form in the Bayes classifier. These are the defining characteristics of QDA. The QDA algorithm estimates μ_k , Σ_k and π_k from the training set similarly to LDA estimates, which yields estimates for $\hat{\delta}_k(x)$,

$$\kappa_{\text{QDA}} = \arg \max_k \hat{\delta}_k(x) = \arg \max_k \left[-\frac{1}{2}(x - \hat{\mu}_k)^T \hat{\Sigma}_k^{-1}(x - \hat{\mu}_k) \right].$$

An observation $X = x^*$ from the testing set is then assigned to the κ th class if $\kappa = \arg \max_k \hat{\delta}_k(x^*)$. As above, the predicted classes using the QDA scheme can then be compared to the true classes to assess classification accuracy [17].

4.3 Regularized discriminant analysis

To see if a more flexible separation of the feature space would yield higher classification rates, we performed classification analysis through the RDA algorithm on all combinations of features for the gel-based cells [19]. This form of discriminant analysis was appealing for two reasons: first, it allows for a “path” between LDA and QDA which is computationally feasible in high dimensions, and secondly, it has LDA and QDA as special cases. Not all data conforms exactly to the complete homoscedasticity that LDA demands nor the complete heteroscedasticity that QDA demands. Instead, their covariance structure decomposes into a mixture of the two extremes, with some shared structure across all K classes and some unique structure within the k th class. RDA analysis allows for such composite covariance structures, including both LDA and QDA as special cases. The RDA algorithm begins with computing a convex combination of the group-specific QDA and pooled (LDA) sample covariance matrices over the k groups in the sample via

$$\hat{\Sigma}_k^{\rho_1} = (1 - \rho_1)\hat{\Sigma}_k + \rho_1\hat{\Sigma}$$

where $0 \leq \rho_1 \leq 1$ is a regularization parameter. The second step shrinks this estimator towards a multiple of the identity matrix in the following manner:

$$\hat{\Sigma}_{k,\text{RDA}} = (1 - \rho_2)\hat{\Sigma}_k^{\rho_1} + \frac{\rho_2}{p} \text{tr}(\hat{\Sigma}_k^{\rho_1})I_p, \quad (3)$$

where $0 \leq \rho_2 \leq 1$ is a second regularization parameter. Once $\hat{\Sigma}_{k,\text{RDA}}$ is calculated, performing RDA is analogous to QDA, but with $\hat{\Sigma}_{k,\text{RDA}}$ substituted for $\hat{\Sigma}_k$. In eq. (3), we note that LDA is achieved when $\rho_1 = 1$ and $\rho_2 = 0$, which corresponds to creating a series of linear divisions of the feature space. We also note that QDA is achieved when $\rho_1 = \rho_2 = 0$, which corresponds to creating a series of quadratic divisions of the feature space. Technically, $0 \leq \rho_1, \rho_2 \leq 2$; however, constructing a model between LDA and QDA, whose divisions of the feature space are somewhere in between linear and quadratic, is only achieved when $0 < \rho_1 < 1$ and $\rho_2 = 0$ [19]. Likewise, we restrict our analysis to this setting, exploring a variety of values for ρ_1 , including the special cases of $\rho_1 \in \{0, 1\}$, fixing $\rho_2 = 0$.

To gain an understanding of the effect of different values of ρ_1 on the classification accuracy, we selected values from 0.05 to 0.95 to use in the RDA algorithm on all pairs and trios of features, leaving $\rho_2 = 0$ fixed. To run this algorithm in R, we first visually confirmed that our feature data for each cell line was approximately normal. We next modified the function `RDA.R` by Aerts and Wilms to accommodate specific values of ρ_1 and ρ_2 [20]. The same random sampling scheme used during the classification tree analysis was adopted here, with half of the data constituting the training set, and the other half forming the test set. Here too we repeated our analysis on twenty different random samplings of the data to mitigate any sensitivity to initialization. Reported accuracy rates for each feature are the average over all twenty of these trials.

As seen in Supplemental Figure 1, the pair of features with highest out-of-sample classification accuracy is associated with $\rho_1 = 0.1$. Even so, RDA fit with $0.05 \leq \rho_1 \leq 0.60$ appears to attain a comparable degree of out-of-sample classification accuracy, using the best pair of features associated with each individual ρ_1 value. Thus, it does not appear as though RDA applied to this data set and these pairs of features is extremely sensitive to the ρ_1 value, as illustrated in Supplemental Figure 2. For trios of features, we see similar behavior. In Supplemental Figure 1, the trio of features with the highest out-of-sample classification accuracy is associated with $\rho_1 = 0.05$; however, RDA seems to perform nearly as well with a wide range of ρ_1 values, indicating a lack of sensitivity to this value, apart from it not being zero or one. It appears as though any form of regularization between LDA and QDA is advantageous to increasing the out-of-sample classification accuracy, with the particular value of ρ_1 of lesser importance.

Table 2 contains the numerical results associated with the best pair and trio of features identified in Supplemental Figure 1, while Figure 4 offers a visualization of the feature space characteristic of RDA, somewhere between a linear and quadratic division. Note that these divisions are quite different from those in Figure 3. In Table 2 we see these differences correspond to about a 2 percentage point increase in the out-of-sample classification rate for the best pair of features using RDA versus classification trees. We also see an increase of about 2 percentage points for the best trio of features using RDA versus classification trees. We emphasize again, though, that the features highlighted in Table 2 are merely members of a larger set of high-performing features. For example, the pairs (R2h @ 32000 Hz, MR @ 16000 Hz) and (R2h @ 32000 Hz, MR @ 8000 Hz) have classification accuracies only one percentage point worse than the best pair. So, while not all pairs and trios of features could yield classification rates as high as 97 % and 99 %, respectively, there are a number which could perform within a few percentage points of these values.

Table 2: RDA Analytical Results: Best features for cell line classification based on twenty trials of RDA with ρ_1 as indicated below and $\rho_2 = 0$. The cells were grown on gel serum. The out-of-sample classification rate is the average over all twenty values obtained for each feature. The standard error of this average value is reported in column five. The ρ_1 values used in the RDA analysis and the associated best classification features correspond to the settings with the highest out-of-sample predictive accuracy in Supplemental Figure 1. We see an improvement of several percentage points in the classification rates for both pairs and trios of features, as compared to those from the classification trees.

Feature Space Dimension	ρ_1	Selected Classification Feature(s)	Out-of-Sample Classification Accuracy	Approximate Standard Error
2	0.10	R2h @ 64000 Hz MR @ 16000 Hz	0.971	0.004
3	0.05	R2h @ 32000 Hz MR @ 4000 Hz MR @ 64000 Hz	0.993	0.002

Figure 4: RDA Visual Results: Best pair of features for cell line classification with cells grown on gel serum as determined through RDA with $\rho_1 = 0.10$ and $\rho_2 = 0$ (as indicated in Table 2). (A) Example of training/testing set used to conduct classification analysis. Training points are represented by hollow circles, test points represented by x's. This separation corresponds to about 97 % predictive accuracy, as seen in Table 2. (B) Boundaries created when using all 14 observations in training set (full information).

Based on our preliminary analysis on the effect of the ρ_1 value on the classification rate, and the low parameter value sensitivity observed, we suspected that the rates associated with LDA and QDA would be similar to those when we specified $\rho_1 \in \{0.05, 0.1\}$, the optimal parameter values for RDA on trios and pairs of features, respectively. As LDA and QDA are common classification techniques, with desirable visualization properties, we extended our RDA analysis to account for these two special cases. The same random sampling and averaging techniques were used here as before. Instead of using the `RDA.R` code, we used the `lda()`, `predict.lda()`,

`qda()` and `predict.qda()` functions to allow for evaluation of single features, as well as pairs and trios. We verified that the results using the modified `RDA.R` code had the same LDA and QDA results for pairs and trios of features as the `lda()` and `qda()` functions. The numerical and visual results can be found in Supplemental Tables 1–2 and Supplemental Figure 3–Figure 4, with highlights summarized below.

Compared to classification trees, we see that LDA and QDA offer nearly uniformly higher out-of-sample classification rates for each combination of features. RDA, however, outperformed both LDA and QDA. High-performing features showed some consistency: the best single feature and feature pair under LDA also appeared in the top 10 selected by RDA. The best single feature and feature trio under QDA also appeared in the top 10 selected by RDA.

The supplemental figures show how LDA and QDA separate the feature space into linear and quadratic regions, respectively. A visual comparison of the LDA, QDA, and RDA regions (in Supplemental Figure 3–Figure 4 and manuscript Figure 4) shows some similarities; but the classification rates in the corresponding tables in the manuscript and supplemental materials indicate that RDA is still the superior classification method. Hence, a more flexible classification algorithm, which allows for a division of the feature space somewhere between linear and quadratic, seems to yield the best results for this dataset.

4.4 Serum effect on classification rates: BSA vs. Gel

The same analysis that was applied to the cells inoculated in gel serum, as described in Sections 4.1–4.3, was repeated on those cells inoculated in BSA serum. Supplemental Table 3 provides the numeric results of the analysis, while Supplemental Figure 5 – Figure 6 provide visualizations. The ordering of best to worst classification methods was slightly different for the BSA-based cells, with LDA performing better than QDA, but ultimately, RDA with the optimal ρ_1 performing the best overall.

The BSA-based cells had lower classification rates than the gel-based cells for all classification methods. Supplemental Figure 6 indicates why BSA-based cells were harder to classify; the plot shows that the observations overlapped more between groups, and, within a particular group, the observations were generally more spread out. This behavior may also explain why the value of ρ_1 in the RDA algorithm had a greater effect on the out-of-sample classification rate for the BSA-based cells (see Supplemental Figure 5); changing the boundaries in Supplemental Figure 6, as dictated by ρ_1 , would yield more differences than in Supplemental Figure 2. Supplemental Table 3 also suggests that lower frequencies of data, such as 2000 Hz and 4000 Hz, are more informative for BSA-based cells, as opposed to the higher frequencies, such as 32000 Hz and 64000 Hz, which were selected as the best features for the gel-based cells. Although the sample size in this study was too small to draw definitive conclusions, these results suggest a slight advantage to using cells grown on gel serum, as opposed to BSA serum. Biologically, we know that cells bind to surfaces using transmembrane proteins called integrins. Whether a cell binds differently to gelatin or BSA depends upon a variety of factors including the cell's available integrins and whether serum is a component of the medium (which it was in this study). Given this biological behavior and the initial results in this manuscript, it may be worthwhile to continue to evaluate the differences between these substrates in future experiments.

4.5 Alternative approaches

We considered several additional analyses on the gel-based data to assess the sufficiency of the classification methods and features used in Sections 4.1–4.3. To test the performance of the EOR feature versus its component features, the resistance values at time indices (TI) 60 - 64, we created a new feature set replacing EOR with five separate features TI60 through TI64; the remaining original features R2h, MR, Rb, and Alpha were retained. Conducting classification tree, LDA, and QDA analyses on this new feature set, we found the substitution of TI60 – TI64 for EOR inferior in almost every case (see Supplemental Table 4). We also performed an analysis using just TI60 – TI64 as features, but this performed uniformly worse (see Supplemental Table 5). Given this reduced performance, and the fact that the resistance values at the individual time indices are less stable and reliable than the aggregate feature EOR, we further advocate for the use of the latter over the former.

While trained on a subset of data, classification trees, LDA, QDA, and RDA algorithms are all attractive to use in our application since they all return a set of classification rules that can act on new data without the retention of the training data. Likewise, once generated, it is feasible to share the features with others interested in differentiating their own cell lines into types. This is less straightforward using another popular classification technique, the k-Nearest Neighbor (KNN) method. Nevertheless, we ran a cross-validated KNN algorithm on our data, following a procedure similar to the one detailed above for the other classification methods. Results and more details of our procedure are detailed in Supplemental Table 6. Overall, we found KNN underperformed those methods detailed in Sections 4.1–4.3.

5 Conclusion

These results have established that previous studies have overlooked valuable information by only relying on ECIS data obtained from one frequency, typically 4000 Hz. While the tables in this manuscript and the supplemental materials provide only examples of the best feature(s), it was clear that neither 4000 Hz, nor any single frequency for that matter, was the sole frequency appearing in the high performing feature sets. Instead, we found that the most effective trees used feature sets involving multiple frequencies. We also found that the “one-at-a-time” feature analysis used in previous work is not sufficient to solve the multiple classification problem. While only incremental improvement in accuracy was gained from considering trios of features, as opposed to pairs of features, there was a marked improvement when considering pairs versus individual features in this study. As for the best classification method, our results suggest that RDA with a value of ρ_1 away from the boundary yields the most accurate results. We also note that LDA and QDA offer only small improvements over basic classification trees, a margin which decreases as the number of features increases.

Overall, results are encouraging for future work. We saw clear separation of the various cell lines using our classification features, which previous work had not achieved. It is possible that the differences in resistance signatures are smaller between a healthy and contaminated version of a cell line, which was the focus of previous work, versus two distinct cell lines, making ours a simpler data set to work with. It is also possible, though, that our features were superior to those used in previous studies since ours were not limited to the confluence region, nor were they restricted to one AC frequency. In general, we found the resistance curves more distinct during the growth phase than the confluence phase, which our features were able to capture and use to better inform the classification procedure.

Since we have demonstrated the value of multi-frequency ECIS data for cell line classification, we plan to extend the experimental design to allow for a broader feature space. In particular, we will continue to look at cells grown on different serums, and also examine cells wounded by a high-frequency current after reaching confluence [10]. We will also obtain more data at a finer temporal resolution during confluence to study the post-confluence features proposed in previous studies that we were unable to address in this work. As in this study, we will evaluate each feature at multiple frequencies. We note, though, that in order to apply the discriminant methods used here, any future data must abide by the (approximate) normality condition and have sufficient sample size relative to the feature dimension. We will also consider the use of functional data features as extensions of this current work, recognizing that examining the time series in their entirety (as curves) may be informative. Adding these features may enrich the classification algorithm and provide even better cell line separation. Finally, since our results indicate that two or three features are sufficient for good classification accuracy, a two- or three-dimensional visualization tool could be developed to allow users to view the feature space and classification regions for their specific analysis.

Supplementary materials

Supplemental Tables and Figures referenced in Sections 4.3–4.5 are available with this paper. Six Supplemental Tables and three Supplemental Figures provide similar visualizations to those tables and figures present in this manuscript, describing the LDA and QDA techniques applied to the gel-treated cell data, all the classification techniques applied to the BSA-treated cell data, and the alternative approaches applied to the gel-treated cells. Three Supplemental Figures also provide visualizations of the effect of different values of ρ_1 on RDA, in terms of the classification rate and the boundaries of the feature space division.

All code used to generate the analysis, figures and tables in this manuscript and the Supplementary Materials is publicly available on GitHub: https://github.com/megangelsinger/IJBS_ECIS_2019.git

Funding

The authors thank the editors and the referees for very helpful comments. The authors also wish thank Dr. Maryclare Griffin and Dr. Charlie Keese for stimulating discussions and helpful comments, as well as other current and former members of Applied BioPhysics, Inc. for introducing us to ECIS and providing the data for our analysis. Financial support from a Xerox PARC Faculty Research Award, NSF grants 1455172 and 1934985, Cornell University Atkinson Center for a Sustainable Future AVF-2017, USAID, and the Cornell University Institute of Biotechnology and NYSTAR is gratefully acknowledged.

References

- [1] Bennet D, Kim S, Ramasamy S. Drug and bioactive molecule screening based on a bioelectrical impedance cell culture platform. *Int J Nanomed*. 2014;9:5789.
- [2] Giaever I, Keese CR. Micromotion of mammalian cells measured electrically. *Proc Nat Acad Sci*. 1991;88:7896–900.
- [3] Das S, Mitra A, Pradhan R. Characterization of electrode/electrolyte interface of ecis devices. *Electroanalysis*. 2012;24:2405–14.
- [4] Isaacson M, Lawrence DA, Mishra NN, Mousseau DE, Retterer S, Szarowski D, et al. On-chip micro-biosensor for the detection of human cd4+ cells based on ac impedance and optical analysis. *Biosens Bioelectron*. 2005;21:696–704.
- [5] Das S, Mandal M, Mitra A, Pradhan R, Rajput S. Electric cell–substrate impedance sensing technique to monitor cellular behaviours of cancer cells. *RSC Adv*. 2014;4:9432–8.
- [6] Giaever I, Keese CR, Mitra P. Electric measurements can be used to monitor the attachment and spreading of cells in tissue culture. *BioTech*. 1991;11:504–10.
- [7] Giaever I, Keese CR, Lo CM. Monitoring motion of confluent cells in tissue culture. *Exp Cell Res*. 1993;204:102–9.
- [8] Huang E, Lim J, Lo CM, Lo JC, Opp D, Wafula B. Use of electric cell–substrate impedance sensing to assess in vitro cytotoxicity. *Biosens Bioelectron*. 2009;24:2625–9.
- [9] Friedman J, Lo CM, Lovelady DC, Patel S, Rabson DA. Detecting effects of low levels of cytochalasin b in 3t3 fibroblast cultures by analysis of electrical noise obtained from cellular micromotion. *Biosens Bioelectron*. 2009;24:2250–4.
- [10] Lukic S, Wegener J. Impedimetric monitoring of cell-based assays. *eLS*. 2015;1–8. DOI: <https://doi.org/10.1002/9780470015902.a0025710>.
- [11] Lo CM, Lovelady DC, Maggi AN, Rabson DA, Richmond TC. Distinguishing cancerous from noncancerous cells through analysis of electrical noise. *Phys Rev E*. 2007;76:041908.
- [12] Cockburn IM, Freedman LP, Simcoe TS. The economics of reproducibility in preclinical research. *PLoS Biol*. 2015;13:e1002165.
- [13] Keese CR, Gelsinger ML. Interview: an overview of ecis technology, Oct 2017.
- [14] Janshoff A, Sapper A, Wegener J. Cell motility probed by noise analysis of thickness shear mode resonators. *Anal Chem*. 2006;78:5184–91.
- [15] Brandenburg SM, Heijink IH, Noordhoek JA, Postma DS, Slebos DJ, van Oosterhout AJM. Characterisation of cell adhesion in airway epithelial cell types using electric cell–substrate impedance sensing. *Eur Respir J*. 2010;35:894–903.
- [16] Choi CK, English AE, Park G, Sparer TE. Electrical impedance measurements predict cellular transformation. *Cell Biol Int*. 2009;33:429–33.
- [17] James G, Witten D, Hastie T, Tibshirani R. An introduction to statistical learning. Springer Texts Stat. 2013. DOI: <https://doi.org/10.1007/978-1-4614-7138-7>.
- [18] Odom-Maryon T, Langholz B, Niland J, Azen S. Generalization of normal discriminant analysis using fourier series density estimators. *Stat Med*. 1991;10:473–85.
- [19] Friedman JH. Regularized discriminant analysis. *J Am Stat Assoc*. 1989;84:165–75.
- [20] Aerts S, Wilms I. Code: Cellwise robust regularized discriminant analysis. <https://feb.kuleuven.be/ines.wilms/code%20>, 2017. [Online; accessed: 01 Nov 2017].

Supplementary Material: The online version of this article offers supplementary material (DOI:<https://doi.org/10.1515/ijb-2018-0083>).