

Methods for High-Content, High-Throughput Image-Based Cell Screening

Thouis R. Jones^{1,2}, Anne E. Carpenter², David M. Sabatini², Polina Golland¹

¹Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT, Cambridge, MA 02139, USA

²Whitehead Institute for Biomedical Research, Cambridge, MA 02142, USA

Abstract—Visual inspection of cells is a fundamental tool for discovery in biological science. Modern robotic microscopes are able to capture thousands of images from massively parallel experiments such as RNA interference (RNAi) or small-molecule screens. Such screens also benefit from lab automation, making large screens, e.g., genome-scale knockdown experiments, more feasible and common. As such, the bottleneck in large, image-based screens has shifted to visual inspection and scoring by experts.

In this paper, we describe the methods we have developed for automatic image cytometry. The paper demonstrates illumination normalization, foreground/background separation, cell segmentation, and shows the benefits of using a large number of individual cell measurements when exploring data from high-throughput screens.

I. INTRODUCTION

One of the most basic tools of modern biology is visual inspection of cells using a microscope. Modern techniques, such as immunofluorescent staining and robotic microscopes, have only magnified its importance for the elucidation of biological mechanisms. However, visual analysis has also become a major bottleneck in large, image-based screens, where tens to hundreds of thousands of individual cell populations are perturbed (genetically or chemically) and examined to find those populations yielding an interesting phenotype. Several genome-scale screens have relied on visual scoring by experts [1], [2]. There are benefits to manual scoring, such as the ability of a trained biologist to quickly intuit meaning from appearance, the robustness of the human visual system to irrelevant variations in illumination and contrast, as well as humans' ability to deal with the wide variety of phenotypes that cells can present.

However, automatic image cytometry has several advantages over manual scoring: simultaneous capture of a wide variety of measurements for each cell in each image (versus scoring a few features per image), quantitative rather than qualitative scoring, ease of reproducibility, detection of more subtle changes than is possible by eye, and the main benefits, elimination of tedious manual labor and much faster analysis of images.

Several groups have made use of automated cell image analysis [3], [4], [5], [6], [7], demonstrating the efficacy of such an approach. These groups have either made use of expensive and inflexible commercial systems, often bundled with a particular imaging platform, or they have developed their own software, seldom used outside of the original lab because of its specificity to a particular screen. In order to

reduce the duplication of effort in this area, and to make tools for automated cell-image analysis more widely available, we have created CellProfiler, an open-source, modular system for cell-image analysis [8], [9].

This paper describes the key algorithms in CellProfiler, and our overall strategies for accomplishing high-throughput image analysis. These include illumination correction to normalize for biases in the illumination and optical path of the microscope, identification of cells versus background, segmentation of individual cells, and capture of a wide variety of per-cell measurements (the “high-content” aspect of our work). We discuss methods and techniques for exploration and analysis of the resulting data and illustrate their application to real-world biological experiments.

II. CHALLENGES IN IMAGE-BASED HIGH-CONTENT SCREENING

We have analyzed several large screens with our system. This paper presents some of the challenges inherent to image-based screens, and the methods we use to address those difficulties. We will use two screens in particular as examples. The first is a set of cell microarrays, single glass slides with cells growing on an array of “spots” printed with gene-knockdown reagents [10]. The second is from an experiment screening ~ 5000 RNA-interference lentiviral vectors targeted to silence ~ 1000 human genes, run in a set of 384-well plates [11]. Both experiments produced thousands of high-resolution (512×512 pixels or larger) images, each containing hundreds of cells. Each image contains cells with a single gene's expression knocked down (decreased).

These experiments suffered from a variety of biases and sources of noise. Both showed illumination variation of around a factor of 1.5 within the field of view, swamping many measurements with noise if not corrected. The cell microarray experiment was performed with *Drosophila melanogaster* Kc167 cells, which are notoriously difficult to segment accurately [12]. Also in this experiment, significant post-measurement biases were detected based on spot position on the slide, due to variations in cell seeding density, concentrations of nutrients or stain, or other factors.

For both screens, discovery of unknown “interesting” phenotypes was and is an open-ended goal. We take a wide variety of per-cell measurements, because we do not know which measurements will be most useful or interesting *a priori*, both in the particular screen and for future explorations. Capturing a wide variety of measurements provides the most freedom

in post-cytometry analysis, but also leads to difficulties in finding which subset of hundreds of measurements can most effectively discriminate a particular phenotype.

Moreover, even in the more goal-directed screens, we are often focused on identifying cells that are different from the “usual” cell in ways that may not be completely specified. Algorithms and methods that work well on normal cells can fail completely when faced with cells that vary significantly in appearance. Robustness to wide variation in cell appearance is therefore an overarching concern in all of our work.

In the following section, we discuss how each of the issues discussed above arose during screens, and the methods we used to overcome these challenges.

III. METHODS

A. Illumination Normalization

Any image- and cell-based screen involves several devices whose physical limitations lead to biased measurements. One of the most pervasive of these is non-uniformity in the optical path of the microscope and the imager. It is typical for the overall illumination to vary by almost a factor of two across the field of view, making segmentation of individual cells more difficult, and seriously compromising intensity-based measurements. Since many such measurements vary less than two-fold in a group of cells, they will be useless unless the illumination is normalized. Fortunately, such variation is consistent from image to image within a single screen, provided as many elements as possible do not change within the screen, i.e., the microscope and optical components are kept the same, the same type of slide or plate is used consistently throughout the experiment, and the images are taken in as short a span as is feasible. We include uneven incoming illumination, sensor biases, and illumination variation due to lens and slide imperfections under the single term “illumination variation.”

We need to estimate the illumination variation in order to correct for it within each image. We model the image-forming process at pixel (x, y) in a particular image I as,

$$I_{x,y} = L_{x,y}(C_{x,y} + b), \quad (1)$$

where I is the image, L is the illumination function, C is an indicator function which is 1 if a cell overlaps pixel (x, y) , and 0 otherwise, and b is a term to account for background staining. Note that this model conflates the magnitude of $C + b$ and L , but since we lack any data that give L physical units, we only need to estimate it up to a scale factor. Lindblad and Bengtsson use a similar model for single image normalization (after log-transformation of the pixels), but without the background term [13]. In our experience, non-specific (background) staining is not always low enough to disregard during normalization.

In the cell microarray experiments, we found that the cell distribution was uniform in the field of view (figure 1) and, in this case, background staining b negligible (as judged from a histogram of pixel intensities). In this case, we estimate L as a smoothed per-channel average of I across all the images in the screen. The average intensities for the three channels (i.e., stains) and the (uniform) cell distribution are shown in

figure 1. Note that we smooth the intensity images to reduce sampling noise prior to using them for illumination correction.

In the well-based experiments, each well was imaged in four different locations. Each location had a significantly different cell distribution, but the background staining level b and the illumination function L were the same across locations (as judged by eye). We use a smoothed regression via equation (1) to estimate L for a range of values of b , taking the pair that best fit the position-wise averages. Cell distribution was estimated by smoothing DNA-stained images and adaptively thresholding to approximately locate nuclei (more accurate identification of nuclei is described in the next section).

Illumination correction is necessary for accurate segmentation and measurement of cells. We note that the optimal solution would combine the estimation of illumination variation and the segmentation steps into a single procedure, similarly to well-known EM-segmentation methods that simultaneously fit a smooth bias field and discrete segmentation labels to image data [14], [15]. In our initial implementation, we have focused on each step separately, with the goal of understanding the nature of the signal. In addition, the high-throughput nature of the experiments places substantial run-time limitations on the algorithms used for the analysis of individual images. We are currently working on a fast implementation for simultaneous illumination correction and segmentation of cell images.

B. Segmentation

The primary benefits of image-based assays are the capture of per-cell data, with a large number of per-cell measurements. This prevents the conflation of multimodal populations, as in expression profiling with gene-chips, and provides a much richer data source than other methods, such as flow cytometry. To exploit the full potential of this data, however, it is necessary to accurately segment individual cells within each image.

Unfortunately, the appearance of cells is highly variable from assay to assay. Experiments use different types of cells, different staining protocols, different growth substrates, and of course, different conditions within each assay. All of these prevent a single approach from being optimal for all cases. We have implemented several methods in our system in a modular fashion so we can easily adapt to new screens.

We have developed a successful, general approach for cell segmentation. Nuclei are more uniform in shape and more easily separated from one another than cells, so we first segment nuclei, then use segmented nuclei to seed the segmentation of individual cells. We threshold the nuclear image using a regularized version of Otsu’s method [16] or our own implementation that fits a Gaussian mixture to pixel intensities. After thresholding the nuclear channel, we separate nuclei that appear to abut or overlap by locating well-separated peaks in the intensity image, and use either a watershed transformation [17] or Voronoi regions of the peaks to place nuclear boundaries, as in related work [18], [19], [20]. Our thresholding and segmentation system are modular, so the user can experiment with different approaches on a small set

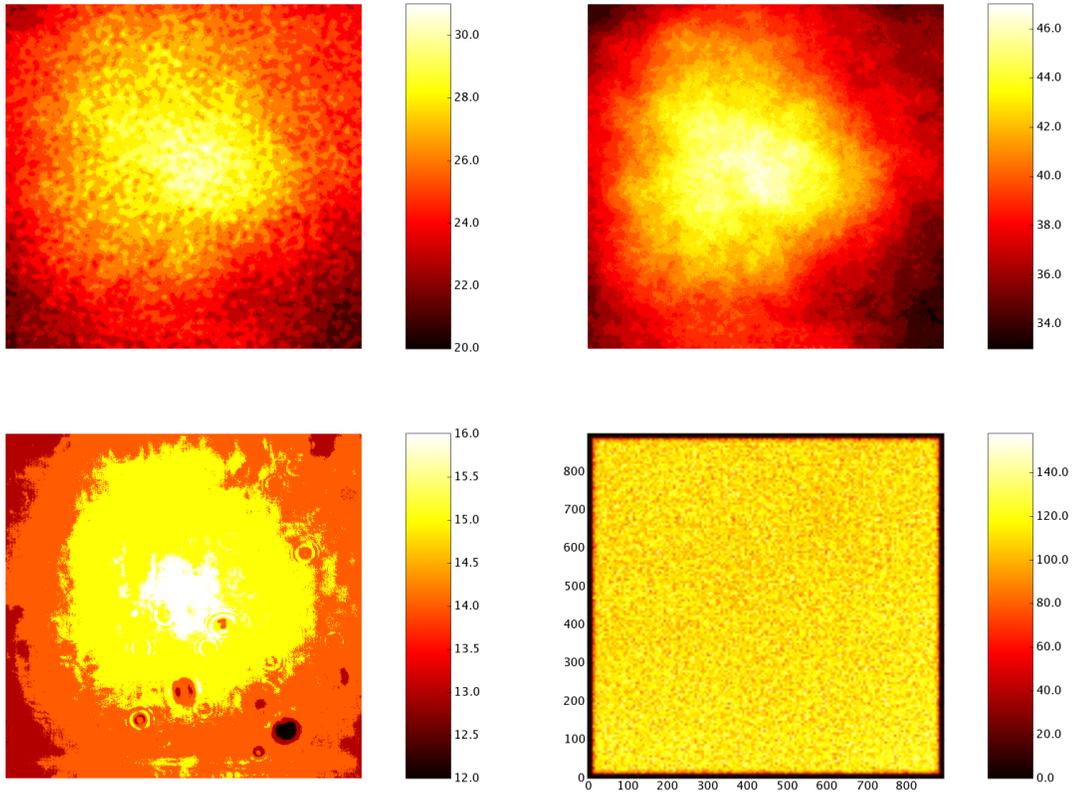


Fig. 1. Top: Mean intensity for DNA (nuclear) and Actin (cytoskeletal) stained channels in the cell microarray experiment. Bottom left: Mean intensity for phospho-Akt stained channel (a protein of interest in this screen). Bottom right: The flat distribution of nuclear centers. Nuclei that overlap the image boundary are eliminated before measurement. All images are false color.

of images to determine the best option, or modify existing modules for a particular experiment.

Given segmented nuclei, segmentation of individual cells is a matter of locating the borders between adjacent cells. The wide variety of cellular phenotypes discussed above prevents us from knowing the particular appearance of cell borders, and in fact, in many screens the borders may change significantly in response to a particular condition, such as a gene’s knockdown. For this reason, we use a very general method for placing cell borders.

A priori, we assume that a pixel we have classified as being “within some cell” is more likely to be associated with the closest nucleus in the image. This naturally leads to using the Voronoi regions of the nuclei to place borders between cells. Another approach is to assume the borders of the cells are brighter or darker, and use a watershed transformation to place boundaries. Both of these approaches are commonly used in image cytometry [21], [22]. However, the first approach makes no reference to the cytoskeletal stain (i.e., information on where the border of the cell is actually located), and the second relies on the borders of the cells being brighter or darker and is overly sensitive to noise in pixels at cell boundaries. In our experience, both of these methods provide poor results in

practice. We combine and extend these approaches by defining a distance between pixels that makes dissimilar pixels farther apart, and use this metric to compute nearest-neighbor regions.

We define similarity in terms of pixel neighborhoods. The distance between adjacent pixels at positions i and j is computed as

$$((i - j)^T \nabla g(I))^2 + \lambda \|i - j\|^2 \quad (2)$$

where $g(I)$ is a smoothed version of the image, $\|i - j\|$ is the Euclidean distance between pixels i and j , and λ is a regularization term that balances between image-based and Euclidean distances. Distances between non-adjacent pixels are computed as the shortest path stepping between adjacent pixels, and cells are segmented via Voronoi regions of nuclei under this metric. More details of this approach are given in our earlier work [23].

C. Measurements

After segmentation, it is possible to make per-cell measurements for each image. Even if the screen is very targeted and the staining protocol has been tuned to give a simple binary answer, we capture a wide variety of measurements in order to maximize our ability to make inferences from the data.

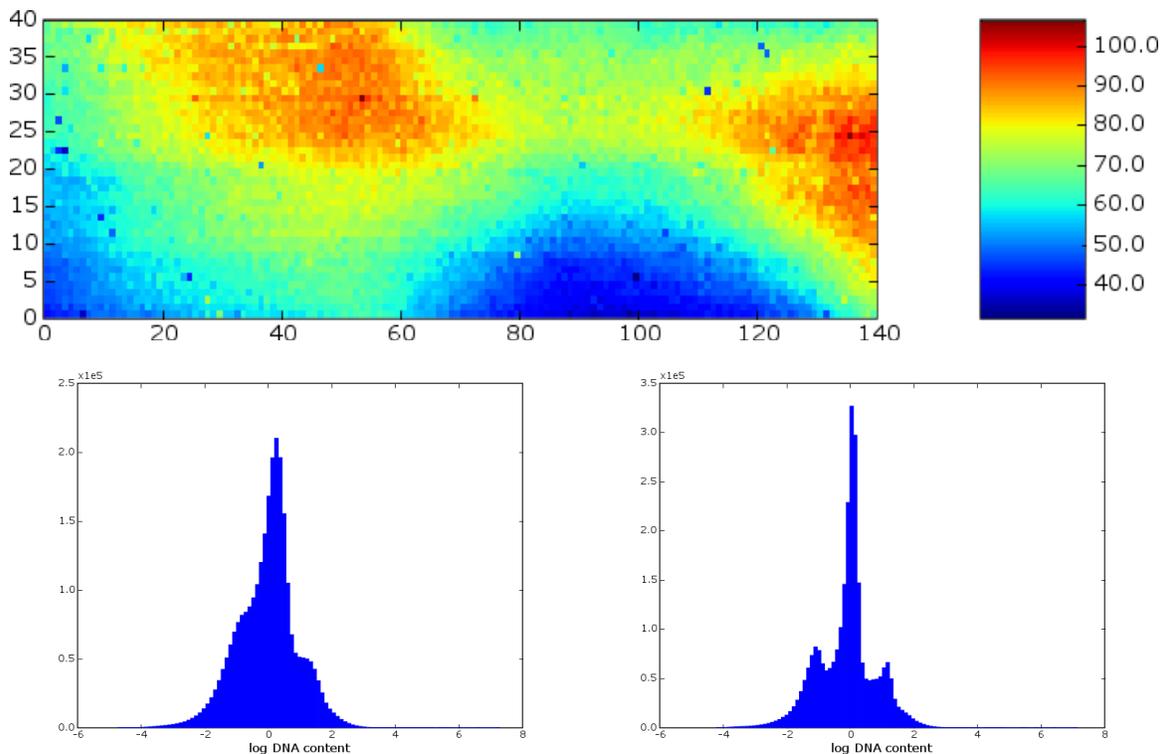


Fig. 2. Top: Median cell DNA content plotted on the physical layout of the slide in the cell microarray experiment. Bottom left: DNA content histogram for all cells on the slide, prior to spatial bias correction. Bottom right: DNA content histogram after bias correction. DNA content is measured by total intensity of the DNA stain within the nucleus, with unknown units. For this cell line, the dominant peak is made up of 4N cells, in which the DNA has duplicated, but the cells have not yet divided. The horizontal axis labels show relative values only.

For each cell, we make measurements of its morphology (e.g., area, perimeter, extent, convexity, and several Zernike moments), and intensity and texture of the various stains (e.g., mean and standard deviation of intensity, correlation of stains, and Gabor filter response at various scales). Measurements are also broken down by cellular compartment (nucleus, cytoplasm, and entire cell). A full discussion of which measurements to use in a given screen is not germane to this paper, but our guiding principle has been that, although it can make inference more difficult, taking too many measurements is better than taking too few. Adding new measurements to our system is simple because of its modular design.

Many of the measurements we capture have a clear biological meaning, such as cell size, or total DNA staining intensity in the nucleus (proportional to the amount of DNA present). Others have a less obvious connection to the biology of the cell, such as the eccentricity of the nucleus, or amount of variation in the cytoskeletal stain. Although we may not be able to assign meaning to every measurement, we can still make use of them when performing analyses or when classifying cells, as discussed in section III-E.

D. Spatial Bias Correction

Before measurements can be used to make biologically useful statements, we must control for systematic biases as much as possible. Biases in the data often come from variation

across the physical layout of the slide or multi-well plate in which the experiments were performed (a.k.a, “plate effects” and “edge effects” [24]).

Some measurements can be corrected by fitting a smooth function to the data on the physical layout, and dividing the corresponding per-cell measurements at each position by the smooth function. For example, if we plot median per-cell DNA intensity on the slide layout for a 5600-spot (140×40) slide (figure 2, top), we observe a spatially varying bias, most likely due to inhomogeneity in the stain for DNA. We correct for this bias by applying a 2D median filter to the 140×40 values and dividing each cell’s measurement by the smoothed value. The improvement in the per-cell DNA content histogram is obvious (figure 2, bottom left vs. right).

In some cases, it is difficult to determine how to correct a particular measurement or combination of measurements. Nonlinear interactions of cells with their environment makes it nearly impossible to remove all biases before making inferences from the data. Therefore, we make maximum use of nearby control spots or wells and check each measurement we use against the physical layout (as in figure 2). Bias correction is an active area of research [24], [25].

E. Exploration and Inference

We take several approaches to exploring data from high-throughput, high-content screens. For example, (1) per-cell

measurements can be combined to give per-gene values by taking means, medians, etc., or by using other data reduction techniques. (2) Pairs of populations produced by different gene knockdowns can be compared directly using distribution-based metrics. Or, (3) individual cells can be classified by their measurements, and gene knockdowns compared by how they change the balance of different classes of cells. We discuss each of these approaches below.

1) *Per-Gene Measurements*: Each per-cell measurement can be converted to a per-gene measurement by taking the mean, median, or otherwise reducing each measurement to a small set of parameters. This approach works particularly well when the screen focuses on a simple single parameter readout (e.g., presence of a given protein), or if the goal is to find gene knockdowns that have an easily measured effect (e.g., cause cells to grow larger). For example, figure 3 shows a scatterplot of per-gene mean cell size vs. mean nuclear size. Three replicates knocking down the gene *ial* are highlighted, in which cells and nuclei have grown larger than controls.

This approach is also effective for early, open-ended exploration, where identification of outliers is the primary task, especially since it can be applied to any measurement without prior knowledge about that measurement's biological implications.

Reducing the data in this way makes it weakly analogous to the data from gene-chips, in which mean expression level is measured for a large number of proteins. Like gene-chips, this approach can suffer from an over-reduction of measurements. For example, knocking down a gene may cause some cells to double in size, and an equivalent fraction to halve in size, but this would not affect the mean cell size [26]. In contrast, if we work with measurements' distributions directly, such differences are easily detected.

2) *Population comparisons*: To compare two populations' measurements directly without first reducing to a single per-gene value, we can apply distribution comparisons such as the Kolmogorov-Smirnov [3] or Kuiper [27] tests, or compute sample-based information-theoretic estimates, such as the KL-divergence between the two distributions [28]. These can be used to compare each sample against a set of positive or negative controls, or against the full slide-wide cell population, yielding a more experiment-specific per-gene measurement as discussed above.

Comparing gene knockdowns' populations via a single or small set of per-cell measurements, as in figure 3, top right, is similar to exploring data from flow cytometry, in which a few measurements are taken for a large number of cells. Flow cytometry is generally lower-throughput than image-cytometry. The number of measurements is also much more limited compared to automatic image-based cytometry.

3) *Per-Cell Classification*: To take full advantage of the large number of per-cell measurements, our primary method of exploration is via per-cell classifiers. We build or train classifiers that identify a phenotype of interest, and apply them to the full screen in order to determine which conditions or gene knockdowns cause enrichment or depletion in those

phenotypes. Our goal is to understand the function of genes, with the underlying assumption that gene knockdowns that cause similar changes in phenotype have similar functions in the cell.

In particular, we advocate the per-cell classifier approach because it detects very small changes in the percentage of cells falling into a particular class. Some phenotypes, such as mitotic (replicating) cells, are $<1\%$ of cells at the background level and increase only three-fold above this level in outliers and positive controls [8]. These changes are so small relative to the full population that they are swamped if measurements are blindly combined into per-gene values, or when comparing two otherwise similar distributions.

Given a classifier for cells showing a known phenotype, the list of gene knockdowns that enrich or deplete that phenotype can be used to impute function for those genes. For example, if we build a classifier for cells in metaphase, knockdowns that cause enrichment of that phenotype probably have a regulatory function in the metaphase to anaphase transition. Simplified examples of per-cell classifiers are shown in figure 4, in which classifiers were constructed to identify different phases of the cell cycle based on a pair of measurements, total nuclear DNA content (as measured by the DNA stain), and mean nuclear phospho-H3 content (a marker for mitosis). If a gene knockdown significantly changes the fraction of cells landing in one (or more) of these classifiers, it is likely to be a regulator for those phases of the cell cycle. Most classifiers are more complicated than this, involving a larger number of per-cell measurements [8].

To compute enrichments and p-values, we treat the output of classifiers as Bernoulli random variables. If negative controls are available in the screen, then enrichments are computed relative to those controls. Otherwise, we use the full screen-wide cell population as the control, the operative assumption being that for each phenotype, knockdown of most genes will not affect that phenotype. There are two justifications for this assumption: many genes are not expressed under experimental conditions, so they cannot be depleted by knockdown, and most genes' knockdown will have no effect on a particular phenotype.

The phenotypes targeted by the classifier can be biologically well-characterized, such as cells in a particular phase of the cell cycle (as above), or simply cells that have a novel appearance, without a well-defined biological interpretation attached. For an uncharacterized phenotype, the group of gene knockdowns causing enrichment or depletion in that phenotype can be informative depending on the group of gene knockdowns causing similar effects. For example, the genes in the group might share a physical or biochemical property, suggesting a mechanism for the phenotypic change. Or, if the group contains genes with a similar, known function, the uncharacterized genes in the group can be hypothesized to also share that function. This also allows for the identification of new, hypothetical cellular processes, rather than simply identifying genes involved in known processes.

The per-cell classifier approach can also be applied to a

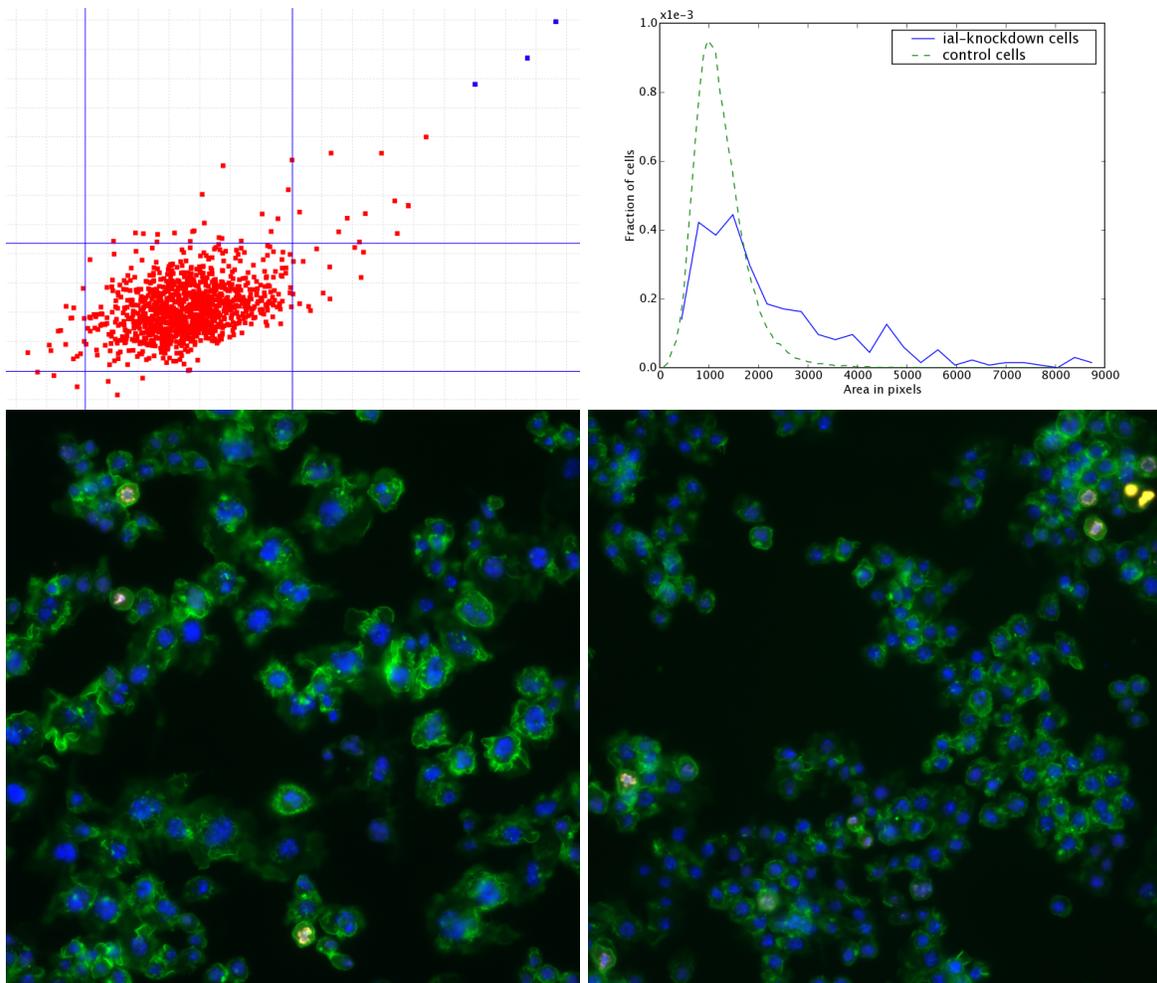


Fig. 3. Top left: Scatterplot of per-image mean cell area vs mean nuclear area in a cell microarray experiment. Three replicates knocking down the gene *ial* are highlighted in the upper right corner. The lines show two standard deviations around the mean. Top right: Per-cell histograms of cell area for the *ial* replicates compared to controls. Lower left: Cells with *ial* knocked down. Lower right: Control cells (Blue: DNA, Green: Actin, Red: phospho-Histone H3).

particular gene's knockdown which might not show a human-discernable phenotype, but for which we can still build a classifier. If the classifier is effective at separating the cells with the target gene knocked down from the cell population at large, the implication is that there is a measurable phenotype caused by the target gene's knockdown, and other knockdowns that cause the same phenotype have a similar function.

One of the benefits of the classifier-based approach is that it is less susceptible to spatial biases when the classifier is trained by a human, compared to data-reduction or full-population comparison methods, because of the robustness of the human visual system to these biases. Note, however, that the nonlinear effect of environment on cells can cause biases in the fraction of cells of a particular phenotype, so the results of applying the classifier should be checked for spatial bias, similar to figure 2.

The classifier approach is reminiscent of example-based image retrieval [29], [30]. However, rather than searching for images as the primary goal, we are using similar techniques

to quickly categorize subimages of cells, with the intent of determining the number and distribution of cells matching our query specification.

After finding hits in a particular screen, followup work may be necessary to validate the results. If there are a sufficient number of replicates, or data from other screens are available, it may be possible to make a categorical statement about a gene knockdown's effect without further experimental work. However, in almost all cases, biologists will investigate the mechanism of the effect in traditional followup experiments.

IV. DISCUSSION

This paper has presented several methods for high-content, high-throughput image-based screens of cells. Such screens are particularly valuable in biological and pharmaceutical research. We have developed CellProfiler [9], a modular, open-source system incorporating these methods.

High-throughput, high-content screening is a powerful technique for making discoveries about cellular processes, genetic

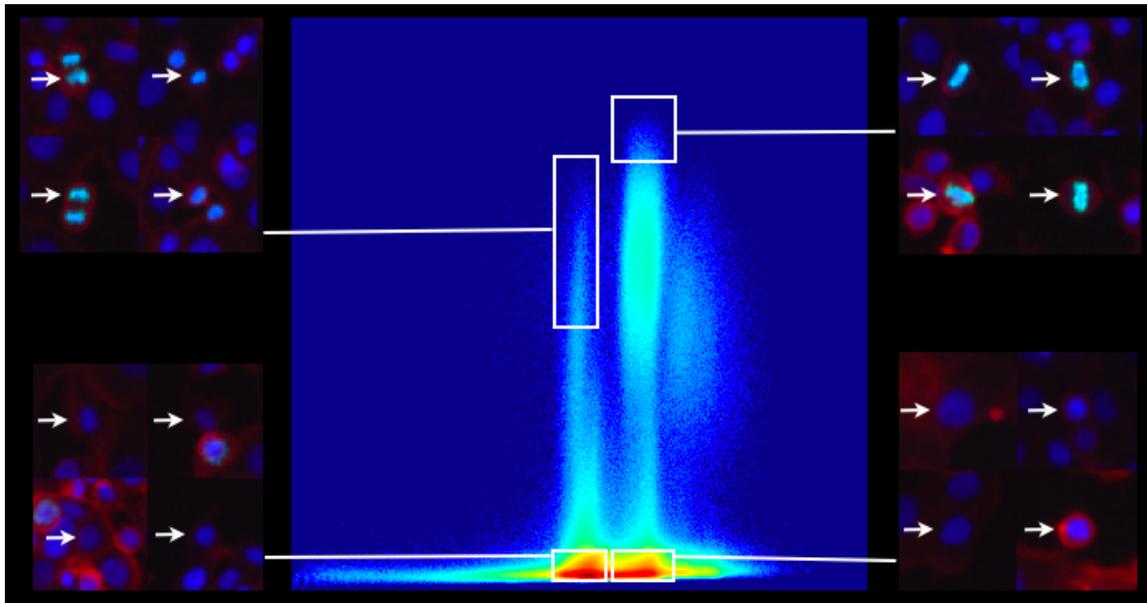


Fig. 4. Simplified examples of per-cell classifiers (from [8]). The central scatterplot shows total DNA content (horizontal axis) vs mean phospho-histone H3 staining intensity for *all* cells in the screen of human genes. Phospho-histone H3 is present in cells undergoing cell division (mitosis). Selecting different regions in the scatterplot selects different subpopulations of cells, as shown in the insets. Each inset shows 4 subimages; each subimage shows a random cell from the corresponding subpopulation (marked) and its surrounding image neighborhood. Counter-clockwise from lower left: 2N cells (normal complement of chromosomes), 4N cells (DNA duplicated), Metaphase (condensed DNA, preparing to separate), Anaphase/Telophase (daughter cells separating). This is also the progression of the cell cycle. A gene knockdown causing enrichment in any of these subpopulations relative to controls is likely a regulator of that phase of the cell cycle. Most classifiers involve many more measurements.

pathways, and drug candidates. It also poses new challenges and requires novel techniques to realize its full potential as a discovery tool. Algorithms for identifying, segmenting, and measuring individual cells must deal with noise and biases, be robust to a wide variety of cell appearances, and must be accurate enough to allow very small ($<1\%$) subpopulations to be identified accurately. However, the payoff for the increased effort is a dramatically more powerful method for detecting changes in cells under different experimental conditions, through the use of per-cell data and classifiers, compared to more traditional techniques. These methods have been proven in some of the first large-scale automatic screens to appear in the biological literature [8], [11].

For reasons of scope, this paper does not include a discussion of the architecture and design of our system implementing the techniques presented here [9]. We aimed to make the system as modular and extensible as possible, while maintaining a user-friendly interface. We believe it has been successful in these respects, particularly given its use in a variety of non-cell screening tasks (e.g., counting and classifying yeast colonies on Petri dishes, counting nuclear subcompartments/speckles, tumor measurement, etc. [31]).

In the future, we will incorporate methods for simultaneous illumination normalization and segmentation, and automatic methods for spatial bias correction. We have also started to explore general clustering based on per-cell measurements.

ACKNOWLEDGEMENTS

This work was supported by an academic grant from the Society for Biomolecular Sciences (AEC), a fellowship from the Merck/MIT computational and systems biology initiative (AEC), the MIT EECS/Whitehead/Broad Training Program in Computational Biology (NIH grant DK070069-01; TRJ), NIH grant R01 GM072555-01 (DMS), and a grant awarded by the U.S. Army Medical Research Acquisition Activity: W81XWH-05-1-0318-DS (DMS), although no official endorsement by the government should be inferred. AEC is a Novartis fellow of the Life Sciences Research Foundation.

REFERENCES

- [1] A. A. Kiger, B. Baum, S. Jones, M. R. Jones, A. Coulson, C. Echeverri, and N. Perrimon, "A functional genomic analysis of cell morphology using RNA interference," *J Biol*, vol. 2, no. 1475-4924 (Electronic), p. 27, 2003.
- [2] J. K. Kim, H. W. Gabel, R. S. Kamath, M. Tewari, A. Pasquinelli, J.-F. Rual, S. Kennedy, M. Dybbs, N. Bertin, J. M. Kaplan, M. Vidal, and G. Ruvkun, "Functional genomic analysis of RNA interference in *C. elegans*," *Science*, vol. 308, no. 1095-9203 (Electronic), pp. 1164-7, 2005.
- [3] Z. E. Perlman, M. D. Slack, Y. Feng, T. J. Mitchison, L. F. Wu, and S. J. Altschuler, "Multidimensional drug profiling by automated microscopy," *Science*, vol. 306, no. 1095-9203 (Electronic), pp. 1194-8, 2004.
- [4] J. A. Phillips, E. J. Rubin, and N. Perrimon, "Drosophila mrai screen reveals cd36 family member required for mycobacterial infection," *Science*, vol. 309, no. 1095-9203 (Electronic), pp. 1251-3, 2005.
- [5] J. N. Harada, K. E. Bower, A. P. Orth, S. Callaway, C. G. Nelson, C. Laris, J. B. Hogenesch, P. K. Vogt, and S. K. Chanda, "Identification of novel mammalian growth regulatory factors by genome-scale quantitative image analysis," *Genome Res*, vol. 15, no. 1088-9051 (Print), pp. 1136-44, 2005.

- [6] L. Pelkmans, E. Fava, H. Grabner, M. Hannus, B. Habermann, E. Krausz, and M. Zerial, "Genome-wide analysis of human kinases in clathrin- and caveolae/raft-mediated endocytosis," *Nature*, vol. 436, no. 1476-4687 (Electronic), pp. 78–86, 2005.
- [7] N. H. Pipalia, A. Huang, H. Ralph, M. Rujoi, and F. R. Maxfield, "Automated microscopy screening for compounds that partially revert cholesterol accumulation in niemann-pick c cells," *J Lipid Res*, vol. 47, no. 0022-2275 (Print), pp. 284–301, 2006.
- [8] A. E. Carpenter, T. R. Jones, M. Lamprecht, D. B. Wheeler, C. Clarke, I. H. Kang, O. Friman, D. A. Guertin, J. H. Chang, R. Lindquist, J. Moffat, P. Golland, and D. M. Sabatini, "CellProfiler: image analysis for high throughput microscopy," 2006, in preparation.
- [9] [Online]. Available: <http://cellprofiler.org>
- [10] D. B. Wheeler, S. N. Bailey, D. A. Guertin, A. E. Carpenter, C. O. Higgins, and D. M. Sabatini, "RNAi living-cell microarrays for loss-of-function screens in *Drosophila melanogaster* cells." pp. 127–32, 2004.
- [11] J. Moffat, D. A. Grueneberg, X. Yang, S. Y. Kim, A. M. Kloepper, G. Hinkle, B. Piqani, T. M. Eisenhaure, B. Luo, J. K. Grenier, A. E. Carpenter, S. Y. Foo, S. A. Stewart, B. R. Stockwell, N. Hacohen, W. C. Hahn, E. S. Lander, D. M. Sabatini, and D. E. Root, "A lentiviral RNAi library for human and mouse genes applied to an arrayed viral high-content screen," *Cell*, vol. 124, no. 0092-8674 (Print), pp. 1283–98, 2006.
- [12] S. Armknecht, M. Boutros, A. Kiger, K. Nybakken, B. Mathey-Prevot, and N. Perrimon, "High-throughput RNA interference screens in *Drosophila* tissue culture cells," *Methods Enzymol*, vol. 392, no. 0076-6879 (Print), pp. 55–73, 2005.
- [13] J. Lindblad and E. Bengtsson, "A comparison of methods for estimation of intensity nonuniformities in 2D and 3D microscope images of fluorescence stained cells." in *Proceedings of the 12th Scandinavian Conference on Image Analysis (SCIA)*, 2001, pp. 264–271, a.
- [14] W. Wells, W. Grimson, R. Kikinis, and F. Jolesz, "Adaptive segmentation of MRI data," in *IEEE TMI*, vol. 15, 1996, pp. 429–442.
- [15] K. V. Leemput, F. Maes, D. Vanermeulen, and P. Suetens, "Automated model-based bias field correction of MR images of the brain," in *IEEE TMI*, vol. 18, no. 10, 1999, pp. 885–895.
- [16] N. Otsu, "A threshold selection method from gray level histograms," *IEEE Trans. Systems, Man and Cybernetics*, vol. 9, pp. 62–66, Mar. 1979.
- [17] S. Beucher, "The watershed transformation applied to image segmentation," in *Scanning Microscopy International*, vol. 6, 1992, pp. 299–314.
- [18] N. Malpica, C. O. de Solorzano, J. J. Vaquero, A. Santos, I. Vallcorba, J. M. Garcia-Sagredo, and F. del Pozo, "Applying watershed algorithms to the segmentation of clustered nuclei," *Cytometry*, vol. 28, no. 0196-4763 (Print), pp. 289–97, 1997.
- [19] F. Meyer and S. Beucher, "Morphological segmentation," in *Journal of Visual Communication on Image Representation*, vol. 1, no. 1, 1990, pp. 21–46.
- [20] C. Wahlby, I.-M. Sintorn, F. Erlandsson, G. Borgfors, and E. Bengtsson, "Combining intensity, edge and shape information for 2d and 3d segmentation of cell nuclei in tissue sections," *J Microsc*, vol. 215, no. 0022-2720 (Print), pp. 67–76, 2004.
- [21] C. Wahlby, "Algorithms for applied digital image cytometry," Ph.D. dissertation, 2003.
- [22] S. Beucher, "The watershed transformation applied to image segmentation," in *Scanning Microscopy International*, vol. 6, 1992, pp. 299–314.
- [23] T. R. Jones, A. E. Carpenter, and P. Golland, "Voronoi-based segmentation of cells on image manifolds," in *CVBIA*, ser. Lecture Notes in Computer Science, Y. Liu, T. Jiang, and C. Zhang, Eds., vol. 3765. Springer, 2005, pp. 535–543.
- [24] D. E. Root, B. P. Kelley, and B. R. Stockwell, "Detecting spatial patterns in biological array experiments," *J Biomol Screen*, vol. 8, no. 1087-0571 (Print), pp. 393–8, 2003.
- [25] D. Kevorkov and V. Makarenkov, "Statistical analysis of systematic errors in high-throughput screening," *J Biomol Screen*, vol. 10, no. 1087-0571 (Print), pp. 557–67, 2005.
- [26] J. M. Levisky and R. H. Singer, "Gene expression and the myth of the average cell," *Trends Cell Biol*, vol. 13, no. 0962-8924 (Print), pp. 4–6, 2003.
- [27] N. H. Kuiper, "Tests concerning random points on a circle," in *Proc. K. Ned. Akad. Wet., Ser. A*, vol. 63, 1962, pp. 38–47.
- [28] S. Kullback and R. A. Leibler, "On information and sufficiency," in *Annals of Mathematical Statistics*, vol. 22, no. 1, 1951, pp. 79–86.
- [29] H. Muller, N. Michoux, D. Bandon, and A. Geissbuhler, "A review of content-based image retrieval systems in medical applications-clinical benefits and future directions," *Int J Med Inform*, vol. 73, no. 1386-5056 (Print), pp. 1–23, 2004.
- [30] K. Tieu and P. Viola, "Boosting image retrieval," *Int. J. Comput. Vision*, vol. 56, no. 1-2, pp. 17–36, 2004.
- [31] A. Carpenter, M. Lamprecht, and D. Sabatini, in preparation.