

Robust colony recognition for high-throughput growth analysis from suboptimal low-magnification brightfield micrographs

Yevgeniy Plavskin^{*1,2}, Shuang Li^{*1,2}, Naomi Ziv³, Sasha F. Levy^{4,5,6}, Mark L. Siegal^{†1,2}

* These authors contributed equally to this work

† Corresponding author: mark.siegal@nyu.edu

¹ Center for Genomics and Systems Biology, New York University, New York, 10003, USA

² Department of Biology, New York University, New York, 10003, USA

³ Department of Microbiology and Immunology, University of California, San Francisco, California, 94158, USA

⁴ Joint Initiative for Metrology in Biology, Stanford, California, 94305, USA

⁵ National Institute of Standards and Technology, Gaithersburg, Maryland, 20899, USA

⁶ Stanford Genome Technology Center, Stanford University, Palo Alto, CA, 94304, USA

Abstract

New technological advances have enabled high-throughput phenotyping at the single-cell level. However, analyzing the large amount of data automatically and accurately is a great challenge. Currently available software achieves cell and colony tracking through the use of either manual curation of images, which is time consuming, or high-resolution images requiring specialized microscopy setups or fluorescence, which limits applicability and results in greatly decreased experimental throughput. Here we introduce a new algorithm, Processing Images Easily (PIE), that automatically tracks colonies of the yeast *Saccharomyces cerevisiae* in low-magnification brightfield images by combining adaptive object-center detection with gradient-based object-outline detection. We tested the performance of PIE on low-magnification brightfield time-lapse images. PIE recognizes colony outlines very robustly and accurately across a wide range of image brightnesses and focal depths. We show that PIE allows for unbiased and precise measurement of growth rates in a large number (>90,000) of microcolonies in a single time-lapse experiment.

Introduction

The development of new technologies for high-throughput phenotyping has allowed increasingly precise measurements of many traits through vastly increased sample sizes, as well as the ability to explore previously intractable traits, such as the extent to which genetically identical individuals raised in the same environment vary (Geiler-Samerotte et al., 2013). Automated microscope stages can be combined with immobilized organisms to collect large amounts of data on cell shape, internal structure, and growth patterns (Levy et al., 2012; Ohya et al., 2005; Rusconi et al., 2014). However, subsequent data analysis often presents a significant challenge, as machine-vision approaches must be applied to automatically extract useful data from the large number of images created.

One particularly important phenotype is the rate of cell growth. In multicellular organisms, cell growth is key to normal development and to disease processes such as tumorigenesis. In microbes, growth rate is closely related to fitness. High-throughput microscopy allows population growth rate to be measured as a function of the growth of individual microcolonies (Bauer et al., 2015; Levy et al., 2012; Ziv et al., 2013, 2017) or the division times of individual cells (Cerulus et al., 2016; Di Talia et al., 2007) rather than a single population-wide average. Individual measurements, in turn, allow exploration of important properties of the growth-rate distribution within a population, such as dispersion and skewness (van Dijk et al., 2015; Levy et al., 2012; Ziv et al., 2017).

To automatically measure microcolony growth rates from microscope-based images, two computational problems need to be solved: recognition of colonies and tracking of colonies through subsequent time points. A commonly used software package for these purposes is CellProfiler, which identifies, segments and tracks cells in images (Carpenter et al., 2006). Other approaches to colony recognition have also been described, including a recent machine-learning based approach (fastER) that learns to recognize a wide range of cell types (Hilsenbeck et al., 2017). However, these approaches require user input (optimization of object-detection algorithm parameters for CellProfiler, and manually scored training data for fastER), which can make the approaches time consuming to customize and test for each application.

The budding yeast, *Saccharomyces cerevisiae*, is a powerful model organism in which to investigate growth variation at the level of individual cells or microcolonies (Levy et al., 2012; Ziv et al., 2013, 2017). Several methods automate the acquisition of growth data from microscopic images (e.g. (Bean et al., 2006; Cerulus et al., 2016; Di Talia et al., 2007)). However, these approaches use high-resolution phase contrast, differential interference contrast (DIC), or fluorescent data as input, which results in greatly decreased experimental throughput than can be achieved with low-resolution images. The utility of these approaches is further limited because they either require specialized microscopy setups (for phase contrast and DIC) or subject fluorescently tagged cells to possible phototoxicity during imaging, which could affect the measurement of colony growth.

We previously introduced a method (Levy et al., 2012; Ziv et al., 2013) that uses low-magnification brightfield images to achieve extremely high-throughput measurement of budding-yeast growth rates (on the order of 10^5 microcolonies per experiment). This method uses data collected periodically (typically once per hour) from cells grown in liquid medium in 96- or 384-well microscope plates coated with the lectin concanavalin A, to which yeast cells adhere, resulting in the formation of two-dimensional microcolonies as the experiment progresses. The image-processing algorithm (here called 'Lblob') identifies yeast colonies by performing blob detection of both the dark outlines and bright centers present in slightly defocused brightfield yeast images, based on user-supplied thresholds (Levy et al., 2012; Ziv et al., 2013). This method is useful not only for studying growth-rate variance as a trait in its own right (Levy et al., 2012; Ziv et al., 2013, 2017) but also, as a result of the large sample sizes the method affords, for estimating mean population growth rates with high precision (Bauer et al., 2015; Ziv et al., 2013, 2017).

Although our previous microcolony identification and tracking algorithm performs well, as we applied Lblob to an increasing number of experimental conditions and encountered variability in the quality of experimental equipment (e.g. microscope plates), we identified several specific sources of technical noise in growth-rate measurements. We therefore aimed to develop an alternative to Lblob that is more robust to variation both across and within experiments, that yields more precise growth-rate estimates, and that can potentially be more easily adapted to different growth conditions and even different species. Here, we introduce Processing Images Easily (PIE), an algorithm for tracking growing colonies in low-resolution brightfield images. Our image-processing procedure combines adaptive object-center detection based on relative brightness with gradient-based object-outline detection, resulting in a high level of robustness to image brightness and focal depth while minimizing the number of user-supplied parameters. The colony-tracking procedure, integrated into PIE from Lblob, then allows tracked colonies to be joined across subsequent time points, allowing simple calculation of growth rates, as well as simultaneous tracking of other colony properties (e.g. fluorescence) over time. We show that the increased robustness in colony outline identification as compared to Lblob results in decreased noise in growth-rate measurements in a yeast growth experiment.

Methods

Procedure Overview

The growth-rate measurement algorithm described here consists of two stages: images are processed to detect colonies and then colony objects are connected through images taken at subsequent time points. Key colony parameters for each time point are recorded for growth-rate estimation and downstream analyses.

The colony detection itself occurs as a three-stage process. First, an automated threshold on relative image brightness identifies 'cell center' objects within an image. Next, the image is split into 'gradient change' objects whose edges coincide with gradient direction changes in the image. Finally, colonies are identified as gradient-change objects overlapping a cell center. An

optional recursive 'cleanup' step removes spurious gradient-change objects that result from brightness fluctuations in the backgrounds of some images.

During the colony tracking stage, colonies are tracked forward through time to identify which colony or colonies in a preceding time point's image spatially overlap with each colony in the current time point's image (Levy et al., 2012). Colony fusion and splitting events are identified. In the case of fusion, tracking ends for both fusing colonies. In the case of splitting, colony areas are calculated by including both the major colony and the components that have split off. Colony areas and other properties called for by the experiment (e.g. colony fluorescence level in corresponding fluorescent-channel images) are then recorded for each time point.

Cell center recognition

Cell centers are recognized by identifying an appropriate intensity threshold on a locally background-corrected image. An original 16-bit image (**Fig 1A**) is first subjected to a top-hat filter (Solomon and Breckon, 2011) to correct for uneven illumination, thereby reducing background and making it more uniform while preserving the brightest parts of the image, which correspond to cell bodies. If the number of unique pixel intensities in the resulting top-hat image (**Fig 1B**) is fewer than 3, typically corresponding to dark images containing no cells, the image is flagged as lacking cells. Otherwise, the appropriate threshold for cell-center recognition is identified using a smoothed histogram of the log frequencies of pixel intensities of the top-hat-filtered image. The histogram bin centers are set to the unique pixel intensity values in the image, thereby avoiding the poor binning that often results from using evenly spaced bins on images in which the distribution of pixel intensities is far from continuous. The log-frequency histogram is then smoothed by applying the Savitzky-Golay filter (Orfanidis, 1996), which fits successive windows of adjacent data points to low-degree polynomials. This step eliminates small spiky peaks of the histogram, which would otherwise introduce error in threshold selection.

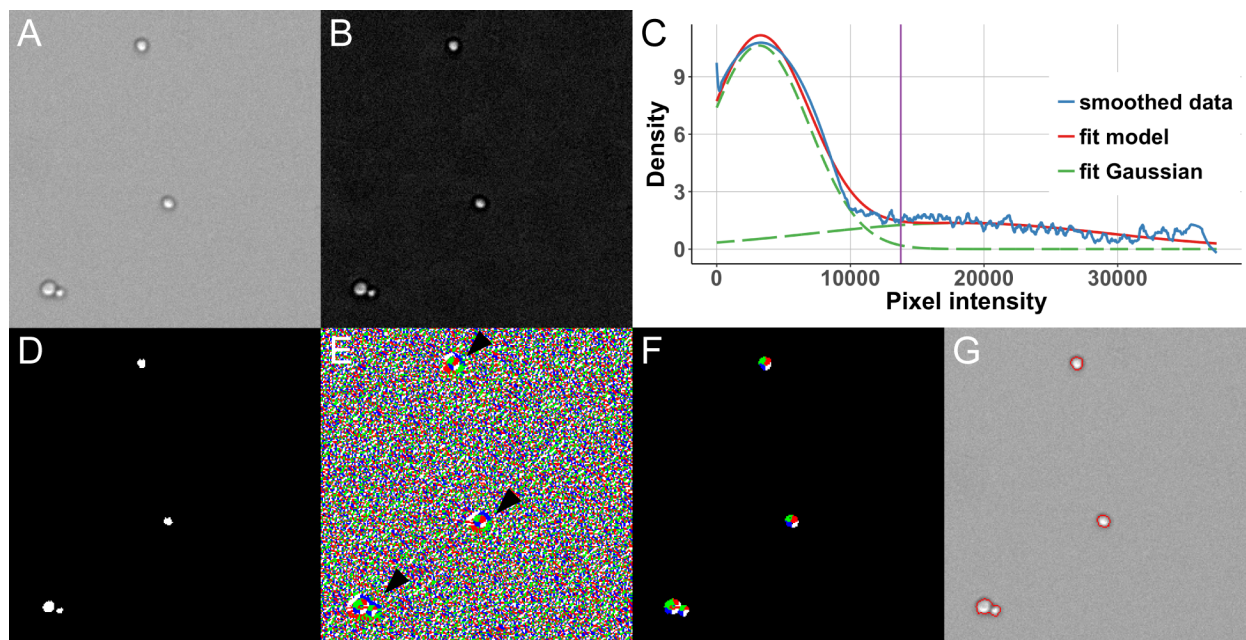


Fig 1. Steps of colony detection by Processing Images Easily (PIE).

(A) Raw brightfield image. To display each colony object clearly, only a representative part of the image is shown. (B) Top-hat filtered image. (C) Smoothed log-frequency histogram of the top-hat image with its two-Gaussian fit. Vertical axis is the density of the log frequencies of pixel intensities. Blue solid curve shows the smoothed log-frequency histogram of the top-hat image (smoothed data); green dashed curves are the two Gaussian fits; red solid curve is the fit combining the two Gaussians; purple vertical line is the automatically calculated threshold for cell centers in the top-hat image. (D) Cell centers after thresholding the top-hat image. (E) Combination of four gradient-changes of the raw brightfield image. Four gradient-changes are: increasing in x and increasing in y (green); increasing in x and decreasing in y (blue); decreasing in x and increasing in y (red); and decreasing in x and decreasing in y (white). (F) Gradient-change objects in (E) that overlap with cell centers (arrowheads) in (D). Color code is the same as in (E). (G) Outlines (red lines) of final colony objects recognized by PIE overlaid on raw brightfield image.

In a typical brightfield image of cells, the smoothed log-frequency histogram of the top-hat image resembles a truncated mixture of two Gaussian curves: one major Gaussian curve corresponds to background pixels, which have lower pixel intensities, and one minor Gaussian curve corresponds to bright cell bodies, which have higher pixel intensities. A good threshold to capture cell centers will sit in between the peaks of the two Gaussian curves. Because cell bodies appear as an intensity gradient, with the center being the brightest, a range of thresholds will produce satisfactory results.

Automated threshold identification is implemented by one of two methods. Our algorithm first attempts to fit a mixture of two Gaussians to the smoothed log-frequency histogram of the top-hat image (**Fig 1C**). If the fit is too poor — which might occur because of very high or low brightness, too many or few cells, capturing part of the well wall, or various combinations of these factors — a more time-consuming sliding-circle approach is used to identify a threshold separating the peak representing background pixels from the one representing bright cell center pixels.

The smoothed log-frequency histogram of the top-hat image is first fit, by non-linear least squares, to a two-Gaussian mixture model with one Gaussian constrained to have mean greater than 0. The second Gaussian is not constrained to have positive mean because the high-intensity (cell-body) values sometimes have such a broad distribution that their mean is estimated as negative. In most cases (93.11% of the images in the experiment described in this paper), the fit to this two-Gaussian mixture model is very good (adjusted $R^2 > 0.85$), and the Gaussian corresponding to the background pixels can be distinguished from the one corresponding to the cell-body pixels by simple comparisons of the two distributions' parameter values (**Sfig 1**). In such cases, the threshold is calculated as the mean plus twice the standard deviation of the Gaussian distribution of background pixels (**Fig 1C**).

Ideally, both distributions have positive means that are clearly distinguished (**Sfig 2A**), but this situation does not always hold. Although the core threshold-identification algorithm is robust against variation in brightness, variation in the proportion of the image corresponding to cell area (cell-to-background ratio) and potential interfering objects such as part of the well wall (**Sfig 2B,2C**), certain combinations of the aforementioned factors, such as low brightness and low cell-to-background ratio (**Sfig 2D**), can change the histogram shape such that the fit to the two-Gaussian mixture model is inadequate (6.89% of the images in the experiment described in this paper).

Improving the threshold-identification algorithm to be robust to the cases in which the two-Gaussian fit is inadequate is essential for two reasons. First, for experiments with a large number of samples and small replicate number for each sample, every image counts. Failure to process only a few images for a particular sample could substantially reduce statistical power to detect differences between samples. Furthermore, if failure to fit adequately does not happen randomly but instead correlates with subgroup identity (e.g. cells of a particular strain, or fast-growing cells), a bias in data collection would result.

To improve the robustness of automated threshold identification, we developed a supplemental algorithm to which images are passed if they fail to produce an adequate fit to the two-Gaussian mixture model with one mean constrained to be positive (adjusted $R^2 \leq 0.85$). To do this, we considered all possible shapes of the smoothed log-frequency histogram of the top-hat image, and subjected a wide range of problematic images to testing.

The supplemental algorithm first tests the source of the poor fit. If the major Gaussian distribution is a good fit for the histogram (i.e., the major Gaussian distribution has a similar amplitude and mean to the highest point in the histogram (**Sfig 2D**)), then the low R^2 is only caused by a poor fit to pixels corresponding to cell bodies. In this case (6.87% of the images in the experiment described in this paper, or nearly all of the images passed to the supplemental algorithm), the major Gaussian distribution's parameter values are used to determine a threshold, as this Gaussian is a good fit for the background pixels (**Sfig 1**). Otherwise, the two-Gaussian model is refit with no constraints on the means. When this fit is very good (adjusted $R^2 > 0.85$), the algorithm attempts to identify which Gaussian corresponds to background pixels by

simple comparisons of the two Gaussian distributions' parameter values (**Sfig 1**). In some cases, neither Gaussian can confidently be called the background-pixel distribution. In those cases, a sliding-circle method described below is applied to the fit model for threshold determination (**Sfig 1**). If the fit to the two-Gaussian model with unconstrained means is inadequate (adjusted $R^2 \leq 0.85$), then the sliding-circle method is applied not to the fit model but to the Savitzky-Golay filter-smoothed data for threshold determination (**Sfig 1, Sfig 2E**).

The sliding-circle method aims to find the nadir of the valley between the two peaks. To do this, it applies a circle mask that glides with its center moving along the smoothed data or fit model. At each position, the ratio of the area within the circle that falls below the curve to that above the curve is calculated. The threshold is determined as the position where this ratio, summed across a given position and two of its neighbours on each side, is maximized. Including neighbour points here reduces the impact of random noise in the pixel-intensity values on threshold determination. Because this sliding-circle method is more time-consuming than the curve-fitting method described above, it is only applied when the two-Gaussian fitting procedures fail. Full details of the fitting and threshold-calculation procedure can be found in the commented code included as Supplemental Information.

During the fit procedure, there are a number of circumstances where the threshold may be suspect, such as when the number of unique pixel-intensity values of the top-hat image is fewer than 200, and when the core threshold-identification algorithm fails (**Sfig 1**). These images are flagged as suspect in a special output .txt file and the corresponding plot containing the smoothed histogram with the fitted curve and the picked threshold is saved for post-analysis inspection of potentially problematic data. Additionally, a user-specified number of images for which the two-Gaussian fit is not suspect are randomly chosen for further inspection, which also aids in diagnosing the quality of the experiment.

To identify cell centers, the determined threshold is applied to the top-hat-filtered image (**Fig 1D**). The resulting binary image has 1s for all pixels that exceed the threshold and 0s otherwise. Therefore, it creates a mask that mostly corresponds to closed central regions of cells.

Colony outline detection

To detect colony outlines, PIE uses a gradient-based approach that provides robustness to differences in image brightness and, to some degree, image focal position. Rather than using edge detection directly to recognize colonies, we first divide images into objects whose outlines correspond to shifts in gradient direction. Objects that do not overlap the cell centers identified above are then removed, leaving colony objects.

To split the image based on gradient direction, a Sobel gradient filter (Solomon and Breckon, 2011) is applied to the image separately in the x and y directions; a threshold of 0 is then applied to the negative and positive versions of each of the gradient images, and pairs of resulting images are multiplied elementwise. The result is a set of four binary images, with objects representing each possible combination of gradient directions in each image: increasing

in x and increasing in y, increasing in x and decreasing in y, decreasing in x and increasing in y, and decreasing in x and decreasing in y. For an object with a bright center and a dark outline, such as a slightly defocused image of a cell, this procedure results in objects in four quadrants resembling pie pieces (**Fig 1E**). To identify pie-piece objects corresponding to cells, any objects not overlapping a cell center identified in the previous section are removed (**Fig 1F**). This leaves only those objects that correspond to cells. Contiguous pie-pieces are then merged into colony objects; as a result, any cells that are touching are considered to be part of a single colony (such as the mother cell and bud cell on the bottom left of **Fig 1F**). Any internal holes are filled to create the final colony objects (**Fig 1G**).

We find that for images of some strains, microcolonies are surrounded by a bright halo. This halo may occur due to refractory properties of cells, which vary with cell size and likely with internal structure. The cell center-finding algorithm described above identifies the halo region as a cell center, and as a result, pie pieces in the background are incorrectly identified as colonies. Our procedure allows for an optional recursive 'cleanup' procedure for sets of images in which this problem occurs frequently (**Sfig 3A**). This cleanup is optional because it substantially extends runtimes and could result in inconsistent tracking of small buds at early time points (**Sfig 3B**).

The cleanup procedure has two steps. The first step is based on the observation that, unlike most true pie pieces corresponding to cells, the spurious background pie pieces are often irregularly shaped and have long, exposed edges that are not in contact with neighboring pie pieces. The cleanup procedure takes advantage of this observation by removing any pie pieces in which the proportion of the perimeter pixels not in contact with another center-overlapping pie piece is less than a user-supplied value. The second step in the cleanup procedure takes advantage of the fact that, for an object with a bright center and dark edges (such as a cell), the gradient should radiate from the center out, meaning that the pie pieces that compose the cell should be in a predictable arrangement. Specifically, a negative x gradient and positive y gradient in the upper right quadrant should neighbor a negative x gradient and negative y gradient in the bottom right quadrant, and a positive x gradient and positive y gradient in the upper left quadrant (**Fig 1F**). Any center-overlapping pie pieces whose neighbors do not have the expected gradient direction are removed. The two cleanup steps are performed recursively, until no more pieces are removed.

Colony tracking through time

Colony tracking is performed in the same way as in Lblob (Levy et al., 2012), and described here for completeness, and summarized in **Sfig 4**. Colony identities and properties, such as area, centroid, pixel-coordinate list and bounding box (the smallest rectangle containing the identified colony) are processed for all fields and time points, then identified colonies in the same field are connected through time. If the x and y distances between the centroids of two colonies in the current and the subsequent time point are less than half the current or the subsequent bounding box width and height, then the two colonies are considered connected between these two time points. It is possible that one colony is not connected to any colony in the previous or subsequent time point, meaning it is a new or lost colony. A colony can also be connected to more than one colony in the previous or subsequent time point, meaning multiple colonies merge or one colony splits into several pieces, creating satellites. All colonies that are connected to at least one other colony are classified as connected objects. Connection serves as the guide for further colony tracking.

After the matrix of connections is created, for a particular field, properties for connected colonies are compiled through all time points. For experiments with only brightfield images, colony properties such as areas, centroid coordinates and identities of satellites are returned. For experiments with fluorescent channel(s), fluorescence-related properties such as the mean, median and upper-quartile fluorescent intensities of the colony and its local background in each channel are also returned, as described in more detail in the next paragraph. If a colony has only one-to-one connections for at least a user-defined minimum number of time points, its properties are recorded and added to the final data matrices. A user-supplied parameter, which we term the 'settle frame,' determines the latest time point at which a connected colony can first appear in order to still be tracked. There are two reasons why it might be important to include colonies appearing after the first time point. First, in experiments that are not performed at room temperature (e.g. yeast growth rate assays), the thermal expansion of the microscope plate in the imaging chamber can cause significant displacements of cells in the x-y plane during the initial phase of imaging. Tracking of such displaced cells can be poor between successive early time points, resulting in colonies that are 'lost' after the first time point and in 'new' colonies appearing at the next time point. Second, at early time points, colonies might not be recognized well because of smaller size or their particular focal plane. However, growth of these colonies often results in improvement of their recognition and stable tracking in future time points. If the new connected colony appears at or before the settle frame, its properties are added to data matrices with 0s for the beginning several time points in which the object is not recognized. For colonies that are connected initially but lost at later time points, their properties are added to data matrices with 0s for all time points after loss.

As described above, two possible scenarios cause colonies to lack one-to-one connections: merging and splitting. Upon merging, tracking stops: when multiple objects in one time point are connected to a single object in the subsequent time point, only properties before merging are recorded. Splitting occurs when a single colony in one time point is connected to multiple colonies in the subsequent time point; the largest colony in the latter time point is then selected

as the tracked colony object. This selection helps sustain tracking if only a few cells break off from a larger colony. Untracked colonies, including the smaller fragments of a split colony, whose distances to the closest tracked colony are below a user-defined proportion of the tracked colony's long axis length, are declared 'satellite' colonies. If a satellite colony is identified, its area is added to its corresponding tracked colony at that particular time point. All other properties, such as centroids and fluorescent intensities are still determined only based on the connected colony. For fluorescent intensities, satellites are not included because the risk of counting brightly autofluorescent debris as a satellite outweighs the benefit of having a small number of additional cells from which to measure a colony's overall per-pixel fluorescence (e.g., by the mean, median or upper-quartile fluorescence). If the smaller fragments of a split colony fail to be identified as satellites, a significant drop in apparent colony area over time might result. However, we apply a filter during the growth rate analysis step to remove colonies whose areas dip significantly in successive time points, thus eliminating any split colonies in which satellite assignment was not successful.

For fluorescence data, PIE improves on Lblob in two respects: the determination of the boundary within which to retrieve fluorescence data for a colony and the capacity to perform local background correction. In PIE, the boundary is placed conservatively within colonies to make fluorescence measurements more robust to slight errors in colony-outline recognition and to inconsistencies in fluorescence at cell edges. Specifically, the colony area is reconstructed as a binary mask and then the mask is eroded with a disk structure element with radius 2 to make the mask slightly smaller than the recognized colony area. The eroded mask is then used to retrieve fluorescent intensity statistics, such as mean, median and upper quartile, from each channel. Although this approach would be problematic for fluorescent markers of the cell wall or cell membrane, the underlying scripts can be easily modified to include the eroded border pixels in the fluorescence calculation. The local background is defined as pixels that are within an area contained by the colony bounding box extended by 5 pixels in all four directions, but are not within the area of the colony itself or of any other colony within the extended bounding box. The resulting background mask is eroded with a disk structure element with radius 2 for the same reason the colony mask is. Finally, the eroded background mask is applied to each fluorescent channel to retrieve the fluorescence statistics. With the statistics extracted from colony objects and their local backgrounds, various methods of local background correction can be applied.

We previously used Lblob to measure colony fluorescence, and therefore used neither mask erosion nor local background correction (Levy et al., 2012; Ziv et al., 2013). Although the lack of these enhancements likely rendered fluorescence measurements more noisy, it should not have biased them, particularly because we assayed fluorescent proteins that are uniformly expressed in the cytoplasm (Levy et al., 2012; Ziv et al., 2013).

Time-lapse microscopy

To test PIE's ability to identify and measure colonies at various focal planes and illumination levels, brightfield images of growing yeast cells were acquired using a Nikon Ti Eclipse inverted microscope with 10X objective, 1.5X magnifier and Perfect Focus System, which automatically corrects drifts and fluctuations in the Z axis during long-term imaging. Images were acquired with NIS-Elements software with gain set to 4 and saved as 11-bit .tif files, which were then converted to 16-bit images for subsequent processing. Each well of a 96-well microscope plate (clear bottom MatriPlates, 0.17 mm bottom glass thickness, catalog number MGB096-1-2-LG-L, Brooks Life Science Systems) was filled with 400 μ L SC medium and seeded (as previously described in (Levy et al., 2012; Ziv et al., 2013)) with cells of the prototrophic haploid strain FY4 (Winston et al., 1995) taken directly from saturation in SC media. A typical experiment will place different strains or media in different wells (Bauer et al., 2015; Levy et al., 2012; Ziv et al., 2013, 2017), but avoiding strain and medium variation in this experiment improves estimation of error associated with the imaging and analysis algorithm.

Colony outline comparison between Lblob and PIE

Two major potential sources of error were considered: illumination and focus. Two sets of images were taken for a field at one time point. Optimal focal plane (corresponding to slightly defocused cells) and brightness are manually determined to correspond to slightly defocused images where cells have a bright center and a dark outline. One set had the optimal focal plane (Z position of the stage determined to be 3702 for this experiment) but varied in brightness (exposure time 2–16 ms), which spans more than the typical illumination variation across one plate. The other set had the optimal brightness (exposure time of 12 ms for this experiment) but varied in focal planes (ranging between 3661 and 3737), which spans more than the typical range of focal-plane differences within one experiment. PIE code without cleanup was applied to process these images. The number of colonies identified for each image and the area for each colony within that image were saved for further analysis. The same images were also analyzed using our previous algorithm, Lblob (Levy *et al.*, 2012). For analyses of focal-plane and brightness ranges, parameters for Lblob were optimized for the images with optimal focal plane (Z position of the stage set at 3702) and optimal brightness (exposure set at 12 ms). Only objects with area larger than 30 pixels were counted, to avoid analyzing debris or poorly recognized cells.

Linear mixed effect modeling of growth rate

We used linear modeling to estimate effects on growth rate. In addition to an intercept representing the mean growth rate, the model included random effects attributed to a colony's well and image field within the well, as well as a residual error term capturing growth-rate heterogeneity among colonies:

$$y_{ijk} \sim \mu + b_i + b_{ij} + \epsilon_{ijk}$$

where y_{ijk} is the growth rate of colony k in field j of well i , μ is the mean growth rate, b_i is the effect of well i , b_{ij} is the effect of field j of well i , and ϵ_{ijk} is the residual error.

Parameter values corresponding to the average growth rate and the well-, field-, and colony-specific random variances, as well as standard errors on growth rate estimates, were computed with the lme4 package (Bates et al., 2015).

Growth-rate comparison between PIE and Lblob

To compare microcolony growth rates measured by PIE and Lblob, we performed a growth experiment similar to those done in previous studies (Bauer et al., 2015; Levy et al., 2012; Ziv et al., 2013, 2017). A microscope plate seeded with cells was subjected to time-lapse imaging for 10 hours with 1-hour intervals, using a manually determined optimal focal plane and brightness.

To compare growth rates for tracked colonies between PIE and Lblob, colonies were matched based on centroid position detected by each method. As previously, to eliminate potential imaging artifacts, colonies were required to be tracked for at least 8 time points, to grow at least 3 fold and not to have any drastic decreases (>2-fold or >500 pixels) in colony area, and regressions of $\log(\text{colony area})$ on time for calculating growth rates were required to have an $R^2 > 0.9$ (Ziv et al., 2013). For each image from time point 4, a distance matrix was constructed between colony centroid positions determined by Lblob and PIE. Colonies whose centroids were within 15 pixels of each other between PIE and Lblob were matched. Colonies from either PIE or Lblob that had more than a single matching colony detected by the other algorithm were removed from the matched list.

Recalculation of growth rates for colonies assigned growth rates by PIE but not Lblob

To investigate why some colonies were assigned slow growth rates by PIE but not assigned growth rates by Lblob, we recalculated growth rates for a random subset of PIE-tracked slow-growing colonies. We manually inspected colony tracking for 198 randomly selected colonies for which the PIE-estimated growth rate was below 0.3, but for which Lblob did not return a growth-rate estimate. We identified the earliest and latest time point for each colony at which PIE tracking accurately identified the outline of the colony, and measured the growth rate between these two time points using PIE-calculated colony areas.

Strains

FY4 *Saccharomyces cerevisiae* yeast were used for all experiments described here, with the exception of the demonstration of the successful ‘cleanup’ procedure shown in **Sfig 3A**, which used haploid strains derived from DBY4974/DBY4975 (Geiler-Samerotte et al., 2016; Joseph and Hall, 2004; Richardson et al., 2013).

Results

Robust colony recognition across focal planes and illuminations

For images acquired by high-throughput microscopy, focal and illumination differences exist within the same experiment. The glass bottom of the microscope plate cannot be perfectly even, so there will be focal-plane variations across the plate and even within one well or field. Illumination of the fields in a well's center is always more uniform than in its corners, where the light is partially shielded by the well wall. Illumination unevenness also exists within the same field, both spatially and temporally, which is why local background correction is a key component of cell-detection algorithms. An effective brightfield image-analysis algorithm should be able to: 1) identify objects even under non-optimal focal plane or brightness, and 2) perform consistent recognition and measurement of the same object across a range of focal planes and brightnesses.

To determine the degree to which PIE is able to decrease technical noise in colony measurements associated with variability in image illumination and focal plane, we first compared the number of identified cells and the consistency of identification across a range of illumination levels and focal planes for PIE and benchmarked these values against Lblob's performance on the same images. PIE not only performs highly consistently in colony identification across both the focal plane range and the brightness range, but also accurately identifies virtually all colonies within the images (**Fig 2-3**). In contrast, Lblob failed to identify any colony at low brightness and does very poorly with low focal planes (blurred images) (**Fig 2-3**). To see whether identified objects are actual colonies and not random debris, and how the identified objects overlap with colonies, we overlaid brightfield images with identified object outlines from both ends of the focal plane range or the brightness range (**Fig 4-5**). As we had hoped, across the focal plane range, PIE consistently recognizes colonies (**Fig 4**). The outlines corresponding to the higher Z position (blue) overlay well with those of the optimal Z position (green). The lower Z position outlines (red) tend to be larger, as images are quite blurred at this position, but outlines tend to center well with those of the optimal Z position. Lblob failed to recognize most colonies at the lower Z position (most cells missing red outline), and the outlines of the higher Z position (blue) do not overlap with those of the optimal position as well as in PIE (**Fig 4**). One limitation, which PIE shares with Lblob in this experiment, is its lower ability to detect cells at high Z position, which is suggested by the drop of identified object number for that focal plane (**Fig 2**). When the Z position is high enough, the inner part of the cell turns dark, which causes PIE's failure in recognition. Currently, PIE is only designed to recognize cells with a bright interior, as PIE uses the top-hat filtered image for cell-center recognition (**Methods**). However, in principle one could add cell-center recognition based on a bottom-hat filter, which would enable PIE to also detect cells that are dark inside. For brightness, PIE performs very consistently across the whole range in that the outlines determined at low brightness (red) and at high brightness (blue) largely overlap with those determined at optimal brightness (green), which is indicated by the mostly white cell outlines (**Fig 5**). Although Lblob failed to recognize any cells at low brightness (absence of red outlines), the high-brightness outlines do mostly overlap with those at the optimal brightness (**Fig 5**). Additionally, comparing

outlines of PIE and Lblob indicates that Lblob tends to include the shadows immediately adjacent to cell bodies and hence draws larger outlines than the actual colonies. The PIE outlines are more accurate, falling right on the edges of colonies (**Fig 4-5**).

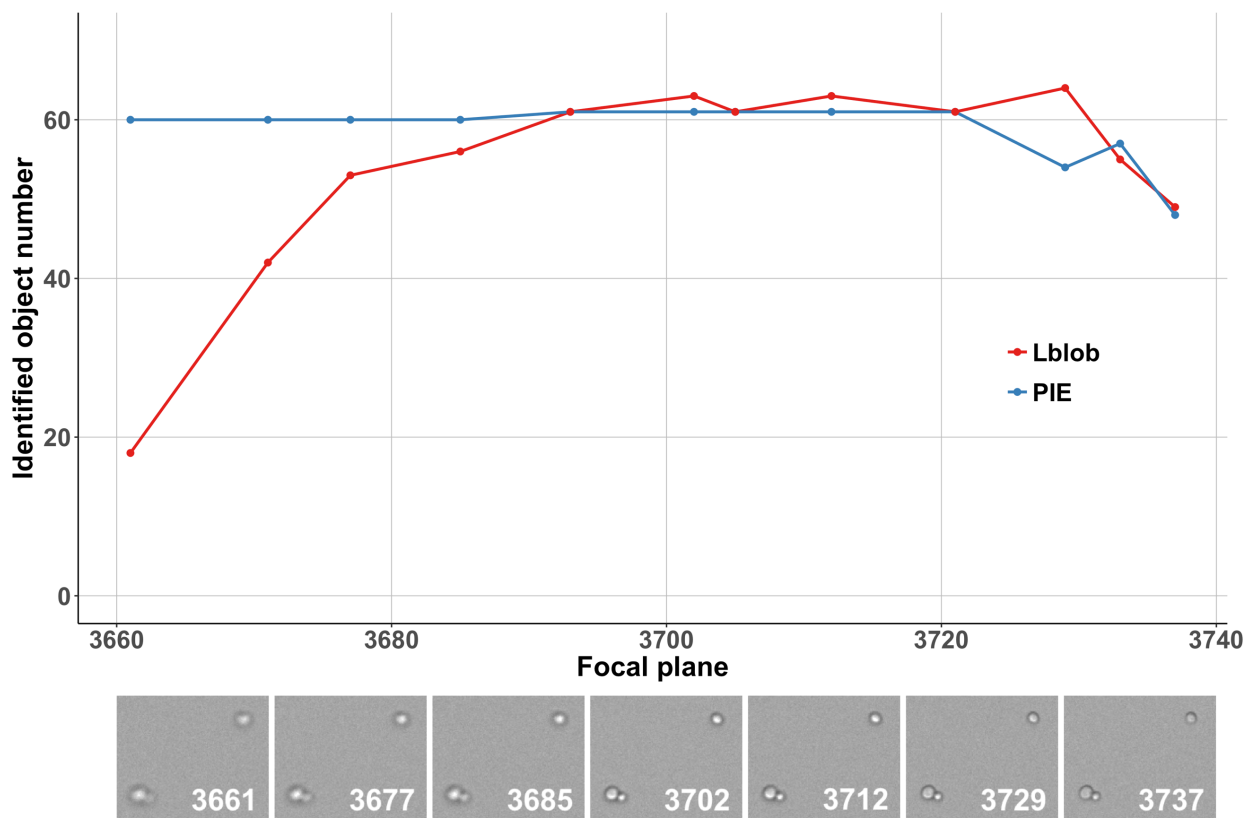


Fig 2. PIE performs consistently and identifies virtually all colony objects across various focal planes.

Upper panel shows identified object number across various focal planes by PIE (blue) and Lblob (red). Lower panel shows images of the same representative cropped region taken at different focal planes, with the corresponding Z position marked at the lower right corner.

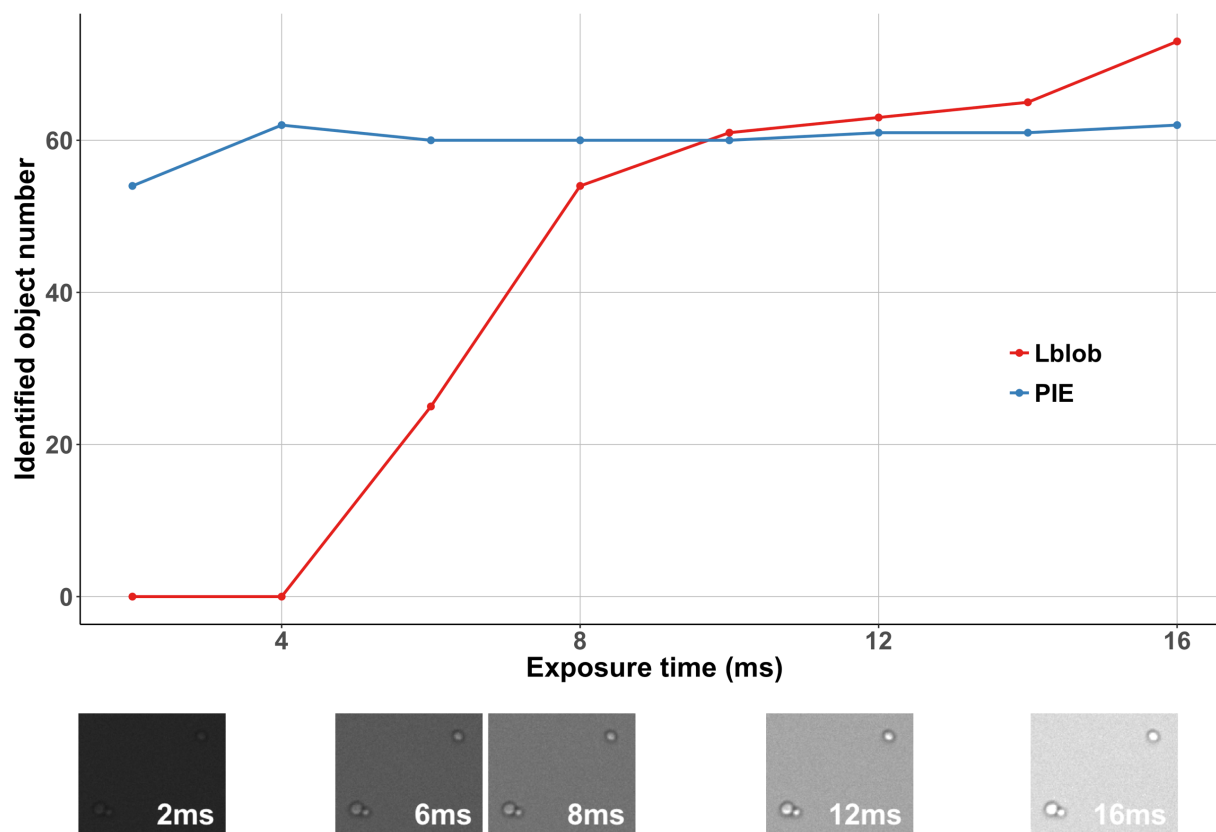


Fig 3. PIE performs consistently and identifies virtually all colony objects across various illumination levels.

Upper panel shows the number of objects identified by PIE (blue) and Lblob (red) across various illumination levels. Lower panel shows images of the same representative cropped region taken at different illumination levels, with the corresponding exposure time marked at the lower right corner.




Fig 4. PIE-recognized colony objects overlap well with real colonies on the raw brightfield image across various focal planes.

Outlines of colony objects recognized by PIE (left panel) and Lblob (right panel) at Z position 3661 (red), 3702 (green, the optimal Z position) and 3737 (blue) overlaid on a raw brightfield image (shown here is the image taken with the optimal setting: Z position at 3702 and exposure time at 12 ms).




Fig 5. PIE-recognized colony objects overlap well with real colonies on the raw brightfield image across various illumination levels.

Outlines of colony objects recognized by PIE (left panel) and Lblob (right panel) at exposure time 2 ms (red), 12 ms (green, the optimal exposure time) and 16 ms (blue) overlaid on a raw brightfield image (shown here is the image taken with the optimal setting: Z position at 3702 and exposure time at 12 ms).

Conclusions based on inspection of outlines on images are supported by quantitative comparisons of the colony area recognized for the same colony (at the initial time point when colonies are typically unbudded or budded single cells) across the focal plane range and the brightness range by PIE and Lblob. Both the focal plane and the brightness ranges were restricted to those where Lblob consistently recognizes most colony objects (Z position of stage between 3693 and 3721 for focal plane, and exposure between 10 and 16 ms for brightness), which are relevant ranges in real experiments. Colonies were matched by centroid location across the images with different focal planes or brightnesses. Then the mean, standard deviation and coefficient of variation across images were calculated for each identified colony. Mean colony areas calculated across the focal-plane range are generally larger in Lblob than PIE (mean across colonies: 202.5 pixels for Lblob, 106.6 pixels for PIE; Wilcoxon-Mann-Whitney test, $P < 2.2 \times 10^{-16}$), which is consistent with what had been observed in the outline overlay image (**Fig 4**). Standard deviations of colony area across the focal plane range are lower in PIE as well (mean standard deviation across colonies: 11.9 pixels for Lblob, 7.1 pixels for PIE; Wilcoxon-Mann-Whitney test, $P = 3.6 \times 10^{-10}$), which is also consistent with what had been observed in the outline overlay image (**Fig 5**). The association between mean and standard deviation in colony area is similar between Lblob and PIE, leading to comparable distributions of coefficient of variation (CV) (mean CV across colonies: 0.06 for Lblob, 0.07 for PIE; Wilcoxon-Mann-Whitney test, $P = 0.125$). As they were for the focal-plane range, mean colony areas calculated across the brightness range are larger in Lblob than PIE (mean across colonies: 208.3 pixels for Lblob, 105.6 pixels for PIE; Wilcoxon-Mann-Whitney test, $P < 2.2 \times 10^{-16}$). Moreover, standard deviations of colony area across the brightness range are much lower in PIE (mean standard deviation across colonies: 24.2 pixels for Lblob, 4.2 pixels for PIE; Wilcoxon-Mann-Whitney test, $P < 2.2 \times 10^{-16}$), suggesting that PIE performs very consistently across different brightnesses and this good performance is also consistent across different colonies. PIE has lower standard deviations even when taking into account differences in identified colony area, leading to lower CV per colony than Lblob (mean CV across colonies: 0.12 for Lblob, 0.04 for PIE; Wilcoxon-Mann-Whitney test, $P = 4.6 \times 10^{-16}$). In summary, PIE not only consistently identifies colonies across experimentally relevant brightness and focal ranges, but also performs much more consistently even within the range where both algorithms are able to detect most colonies. PIE is particularly robust against brightness differences relative to Lblob.

Improved robustness and sensitivity of growth-rate estimation from brightfield images

We next determined the effect that differences in image analysis had on growth-rate calculation using Lblob and PIE. To estimate growth rates from colony areas, we applied a sliding-window approach as previously described (Ziv et al., 2013). Not every identified object is assigned a growth rate (**Methods**): to be assigned a growth rate, colonies must be tracked for a minimum of eight time points, increase in area more than three-fold, have an R^2 of at least 0.9 for the regression of $\log(\text{area})$ on time, and not undergo drastic decreases (>2 -fold or >500 pixels) in area. These filters ensure that only well-tracked growing colonies are assigned growth rates, whereas debris and incorrectly tracked colonies are not. For the same set of raw images,

growth rates were returned for ~93,000 colonies using PIE and ~71,000 colonies using Lblob (Fig 6). An examination of the growth-rate distributions determined by the two approaches reveals two major differences: a higher growth rate measured by PIE relative to Lblob, and a heavier tail of slow-growing colonies identified by PIE (Fig 7, Sfig 5). The first of these discrepancies is consistent with the observation that Lblob tends to identify pixels surrounding cells as belonging to the colony (Fig 4-5, 8). As the colony grows, the pixels surrounding the colony that Lblob adds to the colony area increase proportionally to the colony perimeter, not its area. The increase in colony area resulting from the addition of pixels outside the colony is thus disproportionately higher at early time points, when the actual colony area is small. This inflation of early colony areas causes a depressed apparent growth rate (Fig 8-9).

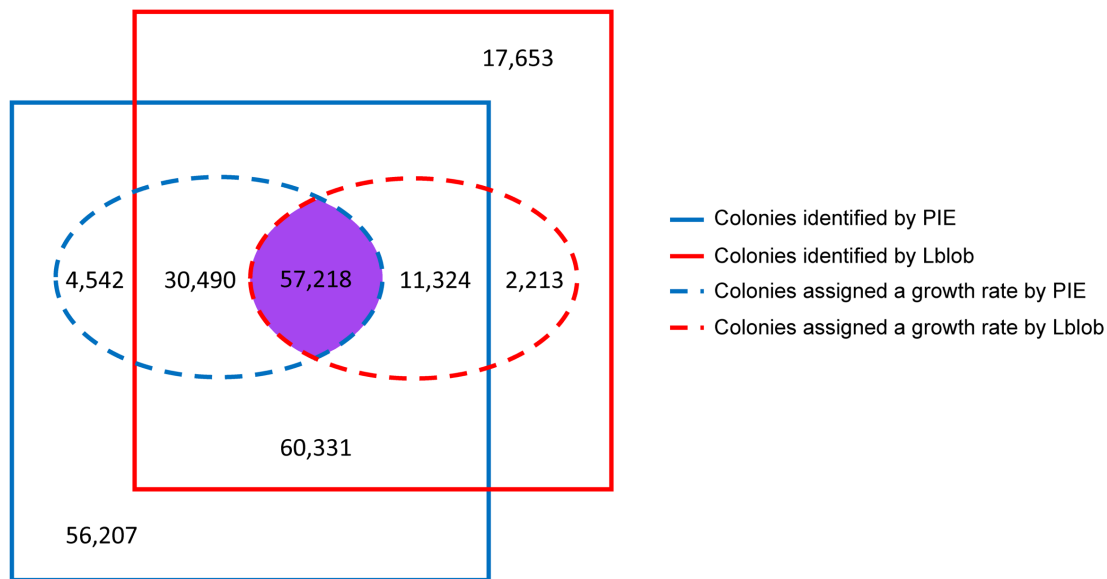


Figure 6. Number of colonies identified and assigned growth rates by PIE and Lblob.

Only objects that are tracked for at least 4 time points are counted as colonies. Colonies identified by Lblob (red) and PIE (blue) are represented as squares; colonies identified by both algorithms and matched between them at time point 4 are in the intersection between the squares. Colonies assigned a growth rate by PIE (blue) and Lblob (red) are represented by ellipses. Colonies that were assigned a growth rate by both algorithms and matched between them at time point 4 and are in the purple intersection between the ellipses.




Figure 7. A higher number of slow-growing colonies is assigned growth rates by PIE.

Growth rates for all colonies measured by PIE (blue) and Lblob (red), including colonies not matched between the two algorithms. Growth-rate distribution for colonies measured by PIE have a larger tail of slow-growing colonies; however, the mode of the PIE growth-rate distribution is higher than that of Lblob.





Figure 8. PIE accurately tracks colony area across time.

Upper panel shows $\log(\text{area})$ changes over time for a representative colony that was assigned a growth rate in both PIE (blue) and Lblob (red). Lower panel shows the corresponding images of the colony at each time point with PIE-derived (blue) and Lblob-derived (red) outlines. Because Lblob parameters were optimized for the image brightness level at time point 4, colony tracking at time point 1 (when the image is brighter) is poor. Lblob tracked colony area is disproportionately inflated at early time points, resulting in a lower estimated growth rate. The trendlines in the upper panel show $\log(\text{area})$ as a function of time for the eight time points that result in the fastest growth rate for each algorithm (Ziv et al., 2013); PIE growth rate = 0.40, Lblob growth rate = 0.37.

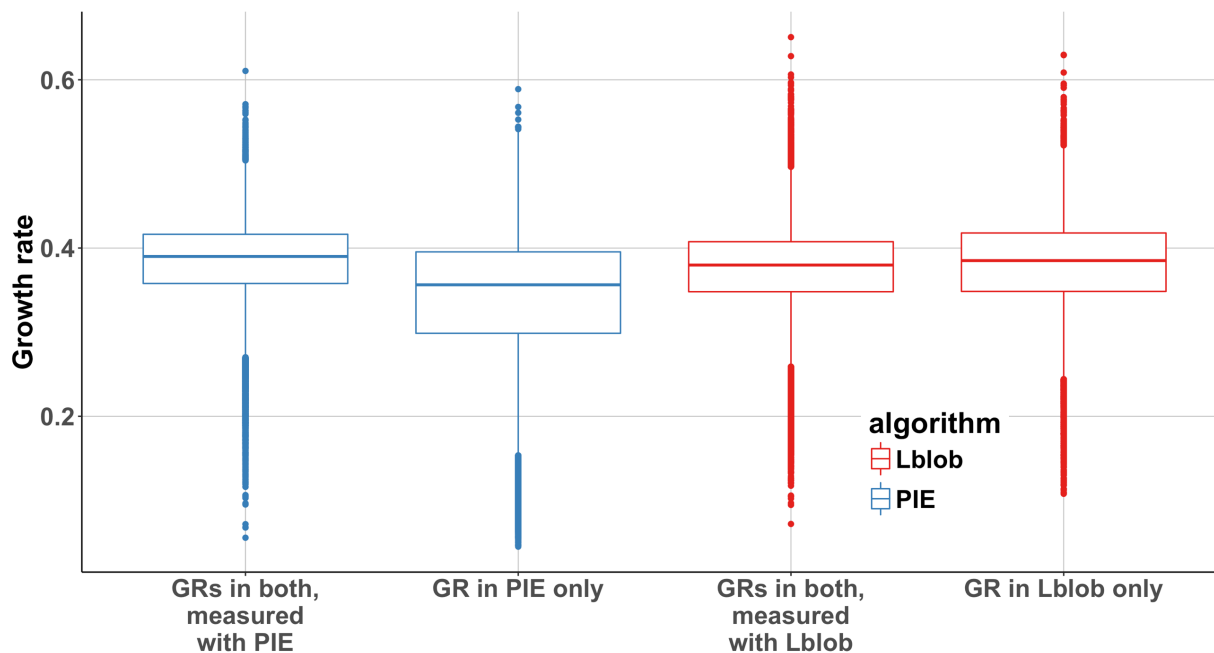


Figure 9. Colonies to which PIE fails to assign growth rates are largely not biased in growth rate. Comparison of growth rates for colonies assigned growth rates by both Lblob (red) and PIE (blue), or uniquely by one or the other algorithm.

To further investigate differences in colony tracking between the two algorithms, we defined colonies as the objects that are tracked for at least 4 time points and matched colonies tracked using PIE to those tracked using Lblob. Among the ~159,000 colonies matched between PIE and Lblob, ~57,000 colonies have growth rates assigned by both algorithms, ~60,000 colonies do not have a growth rate assigned by either algorithm, ~30,000 colonies have growth rates assigned by PIE but not Lblob, and ~11,000 colonies have growth rates assigned by Lblob but not PIE (**Fig 6**). The high number of colonies not assigned a growth rate in this experiment is in part due to the stringent requirement for colonies to be tracked for at least 8 time points. Aside from the shift in measured growth rates described above, colonies that have a growth rate assigned by both algorithms have generally similar growth rate distributions (**Fig 9, Sfig 5**). The colonies that have growth rates assigned exclusively by Lblob have a significant yet slight shift towards fast growth (**Fig 9**). The similarity between the growth-rate distributions of colonies that have growth rates assigned by both algorithms and those that have growth rates assigned exclusively by Lblobs suggests that the reason those colonies were filtered out by PIE is not correlated with growth rate. On the other hand, the colonies that have growth rates assigned exclusively by PIE include a large number of slow-growing ones (**Fig 9**).

To verify that the increased number of slow-growing colonies identified by PIE was not due to errors in PIE's measurement of growth rates, we manually inspected the tracking of a random subset of 198 colonies with PIE-specific, low growth rates (which either had a matched colony in Lblob that was not assigned a growth rate, or did not have a corresponding Lblob colony at all). We found that these colonies were well tracked by PIE, and the low growth rates assigned to

them were not deflated, in that the rates did not differ substantially from those calculated simply for each colony from two time points manually confirmed to have accurate outlines (**Methods**). When we examined why these colonies were not assigned a growth rate by Lblob, we found that in many cases, colonies were discarded due to inconsistent tracking by Lblob over time and Lblob's tendency to overestimate initial colony area. Slow-growing colonies are more sensitive to inaccuracies in tracking. Thus, we find that improved colony tracking by PIE generally results in an increased number of real slow-growing colonies that are assigned a growth rate as compared to Lblob.

To determine whether the increased robustness of PIE to image brightness and focal-plane variation resulted in elimination of more technical variation and a more precise growth-rate estimate, we fit a linear model. The model partitions growth-rate variance into a component that captures the actual biological heterogeneity of colony growth rates, as well as components that capture technical variation among fields and among wells, to which differences in brightness and focal plane contribute (**Methods**). We analyzed growth rates for only those colonies assigned a growth rate by both algorithms (**Table 1**). We find that although within-field (residual) variance is similar between the two algorithms, PIE reduces between-field variance approximately four-fold and between-well variance approximately two-fold. Correspondingly, the standard error of the mean growth rate calculated using PIE is 0.25% of the mean growth rate, compared to 0.35% when using Lblob. Thus, although growth-rate estimation is extremely precise for both algorithms we find that PIE's increased robustness to imaging variability results in increased precision in growth-rate measurement.

	PIE	Lblob
mean growth rate	0.381	0.377
SE of growth rate mean	0.95×10^{-3}	1.30×10^{-3}
between-well variance	0.78×10^{-4}	1.42×10^{-4}
between-field variance	0.96×10^{-4}	4.26×10^{-4}
residual variance	2.84×10^{-3}	2.56×10^{-3}

Table 1. Linear model results for growth rates of ~57,000 colonies analyzed by PIE or Lblob

Discussion

PIE combines adaptive cell-center recognition and gradient-based edge detection to recognize outlines of yeast colonies in low-resolution brightfield images and tracks colonies through time. As we have shown, PIE's colony detection is robust to variation in image brightness and, to some degree, focal plane. Compared to our previously published brightfield colony-tracking algorithm (Levy et al., 2012), PIE both increases robustness of growth-rate measurement in detected colonies and has increased sensitivity that allows it to successfully track a larger

proportion of slow-growing colonies. We have no reason to suspect that any conclusions drawn using Lblob in our prior work need revision, but the use of PIE will be beneficial going forward, especially in experiments that aim to detect very subtle differences in growth rate or that assay numerous genotypes or conditions, each allocated to a small number of wells per plate.

Unlike a number of previously used methods for tracking yeast colony growth (Bean et al., 2006; Cerulus et al., 2016; Di Talia et al., 2007), PIE follows Lblob in allowing yeast colonies to be robustly tracked from low-resolution brightfield images. This advance is important for a number of reasons. First, the ability to use brightfield images has significant advantages over approaches based on fluorescence, phase contrast, or DIC images. Although the basic requirements for any high-throughput growth rate experiment, such as the one described in this paper, include a microscope with a mechanized stage, an automatic focusing system, and computer software that can automate image acquisition, brightfield imaging minimizes the amount of additional specialized microscopy equipment required. The use of brightfield images also eliminates concerns about phototoxicity, which can affect growth measurements when cells are repeatedly exposed to fluorescence excitation over the course of an experiment. However, PIE allows for simultaneous analysis of fluorescence and brightfield images if desired; in this way, no single fluorescence channel is required to be set aside for cell recognition. Second, the ability to accurately track colonies in low-resolution images vastly increases the sample size of any experiment. For example, for the growth-rate data described here, acquired with a 10X objective lens and 1.5X multiplier, we tracked and measured growth rates for >90,000 colonies in a single experiment. Increasing the objective to 40X would have resulted in a significant improvement in image resolution, but a >16-fold decrease in sample size.

PIE also contains relatively few parameters that need tuning and, for budding yeast, requires absolutely no user input into the image-analysis portion of the algorithm, with the exception of a decision about whether to run the algorithm in 'cleanup' mode. We hope that its corresponding ease of use, along with its robustness to experimental variation and the very high sample sizes it affords, encourages others, including those without image-analysis expertise, to use PIE to integrate high-throughput image-based growth assays into their work.

The PIE pipeline presented here and available from the authors upon request can perform brightfield growth-rate measurements and, if needed, simultaneous measurements of colony fluorescence at each time point. A pipeline for running additional experiment types that may be of interest to users is currently under development. Future directions include the ability to integrate fluorescent images taken at only a single time point (either before or after brightfield imaging), as well as experiments that include a treatment (such as a heat shock) in the middle of the experiment (as described, for example, in (Levy et al., 2012)). A pipeline to run PIE analysis in parallel on high-performance clusters is also currently under development, which will greatly expedite PIE-based analysis. Finally, although PIE was originally developed for tracking budding yeast colony growth, we are currently exploring the extent to which PIE may be applicable to images of other colony-forming cells.

References

Bates, D., Maechler, M., Bolker, B., and Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *J. Stat. Softw.* 67, 1–48.

Bauer, C.R., Li, S., and Siegal, M.L. (2015). Essential gene disruptions reveal complex relationships between phenotypic robustness, pleiotropy, and fitness. *Mol. Syst. Biol.* 11, 773.

Bean, J.M., Siggia, E.D., and Cross, F.R. (2006). Coherence and timing of cell cycle start examined at single-cell resolution. *Mol. Cell* 21, 3–14.

Carpenter, A.E., Jones, T.R., Lamprecht, M.R., Clarke, C., Kang, I.H., Friman, O., Guertin, D.A., Chang, J.H., Lindquist, R.A., Moffat, J., et al. (2006). CellProfiler: image analysis software for identifying and quantifying cell phenotypes. *Genome Biol.* 7, R100.

Cerulus, B., New, A.M., Pougach, K., and Verstrepen, K.J. (2016). Noise and Epigenetic Inheritance of Single-Cell Division Times Influence Population Fitness. *Curr. Biol.* 26, 1138–1147.

van Dijk, D., Dhar, R., Missarova, A.M., Espinar, L., Blevins, W.R., Lehner, B., and Carey, L.B. (2015). Slow-growing cells within isogenic populations have increased RNA polymerase error rates and DNA damage. *Nat. Commun.* 6, 7972.

Di Talia, S., Skotheim, J.M., Bean, J.M., Siggia, E.D., and Cross, F.R. (2007). The effects of molecular noise and size control on variability in the budding yeast cell cycle. *Nature* 448, 947–951.

Geiler-Samerotte, K.A., Bauer, C.R., Li, S., Ziv, N., Gresham, D., and Siegal, M.L. (2013). The details in the distributions: why and how to study phenotypic variability. *Curr. Opin. Biotechnol.* 24, 752–759.

Geiler-Samerotte, K.A., Zhu, Y.O., Goulet, B.E., Hall, D.W., and Siegal, M.L. (2016). Selection Transforms the Landscape of Genetic Variation Interacting with Hsp90. *PLoS Biol.* 14, e2000465.

Hilsenbeck, O., Schwarzfischer, M., Loeffler, D., Dimopoulos, S., Hastreiter, S., Marr, C., Theis, F.J., and Schroeder, T. (2017). fastER: a user-friendly tool for ultrafast and robust cell segmentation in large-scale microscopy. *Bioinformatics*.

Joseph, S.B., and Hall, D.W. (2004). Spontaneous Mutations in Diploid *Saccharomyces cerevisiae*. *Genetics* 168, 1817–1825.

Levy, S.F., Ziv, N., and Siegal, M.L. (2012). Bet hedging in yeast by heterogeneous, age-correlated expression of a stress protectant. *PLoS Biol.* 10, e1001325.

Ohya, Y., Sese, J., Yukawa, M., Sano, F., Nakatani, Y., Saito, T.L., Saka, A., Fukuda, T., Ishihara, S., Oka, S., et al. (2005). High-dimensional and large-scale phenotyping of yeast mutants. *Proc. Natl. Acad. Sci. U. S. A.* 102, 19015–19020.

Orfanidis, S.J. (1996). *Introduction to Signal Processing*.

Richardson, J.B., Uppendahl, L.D., Traficante, M.K., Levy, S.F., and Siegal, M.L. (2013). Histone variant HTZ1 shows extensive epistasis with, but does not increase robustness to, new mutations. *PLoS Genet.* 9, e1003733.

Rusconi, R., Garren, M., and Stocker, R. (2014). Microfluidics expanding the frontiers of microbial ecology. *Annu. Rev. Biophys.* 43, 65–91.

Solomon, C., and Breckon, T. (2011). *Fundamentals of Digital Image Processing: A Practical Approach with Examples in Matlab* (John Wiley & Sons).

Winston, F., Dollard, C., and Ricupero-Hovasse, S.L. (1995). Construction of a set of convenient *Saccharomyces cerevisiae* strains that are isogenic to S288C. *Yeast* 11, 53–55.

Ziv, N., Siegal, M.L., and Gresham, D. (2013). Genetic and nongenetic determinants of cell growth variation assessed by high-throughput microscopy. *Mol. Biol. Evol.* 30, 2568–2578.

Ziv, N., Shuster, B.M., Siegal, M.L., and Gresham, D. (2017). Resolving the Complex Genetic Basis of Phenotypic Variation and Variability of Cellular Growth. *Genetics* 206, 1645–1657.