CRISP: Correlation-Refined Image Segmentation Process

Jennifer K. Briggs[1], Erli Jin[2], Matthew J. Merrins[2], Richard K.P. Benninger[1]
[1]Department of Bioengineering, University of Colorado Anschutz Medical Campus, United States; Barbara Davis Center for Childhood Diabetes, University of Colorado Anschutz Medical Campus, United States

[2]Department of Medicine, Division of Endocrinology, Diabetes & Metabolism, University of Wisconsin-Madison, Madison, WI, United States
Correspondence: Jennifer.briggs@cuanshutz.edu, richard.benninger@cuanshutz.edu

Abstract:
Calcium imaging offers the ability to observe cellular activities in real-time across various contexts. However, the manual outlining of cells in calcium imaging data introduces potential errors. This paper introduces the Correlation-Refined Image Segmentation Process (CRISP), an automated algorithm designed to enhance the accuracy of cell mask refinement and to assist in the identification of cell boundaries. CRISP leverages intracell correlations to refine manually drawn masks and automate the detection of the largest cell area that contains only pixels from within the cell. The algorithm not only enhances the precision of calcium trace data but also improves the reliability of functional network analyses.

Introduction:
Calcium imaging is a widely used for quantifying cellular activity across many different fields of biology[1]. A benefit of calcium imaging is the ability to capture the activity of multiple cells at the same time. For example, the pancreatic islet, which is central to diabetes, is composed of thousands of beta-cells that communicate electrically to secrete insulin in a coordinated, pulsatile fashion. Therefore, the activity and coordination of many beta-cells at one time has become a rich topic to study[2–4].

Typically, to study activity and coordination of cells using calcium imaging, a researcher takes a series of time-stacked calcium images and manually outlines each cell of interest to create a mask. The image intensity of all pixels within those masks are then averaged to create the calcium time course for each cell. This methodology presents two possible problems. First, when assessing the coordination between multiple cells, it is essential that the pixels within each cell outline are within the cell of interest. If the researcher accidently included pixels within the cell mask that were actually a part of the neighboring cell, the coordination between the two cells would be over reported because the calcium trace for the cell of interest would actually represent pixels from the cell of interest and its neighboring cell. Alternatively, in practice, researchers typically outline cell masks from either the first calcium image or the average over all images in time. If the cells move slightly over the time course, a cell mask could contain pixels from neighboring cells at various time points when the cells moved. Both of these cases are common for expert and careful researchers. The second methodical problem is that manually circling cells in a careful enough manner to avoid including neighboring pixels is fatiguing for the researcher. Especially if there are many cells to manually identify, researcher fatigue cause less accurate outlines, therefore making the error in the cell mask heteroskedastic (more error as time and fatigue increases), and possible confounding results.

In this manuscript, we develop two automated algorithms based on a single principle to ameliorate these methodological problems arising from using manually identified cell masks for timeseries analysis.

Methods:

Ca^{2+} Imaging: Islets were isolated from *Ins1-Cre:ROSA26*$^{GCaMP6s/H2B\text{-}mCherry}$ mice as in[8]. The following day, islets were loaded in a QE1 quick exchange chamber (Warner Instruments) mounted on a Nikon A1R confocal microscope equipped with a 60×/1.42 NA oil objective. The imaging solution contained, in mM: 135 NaCl, 4.8 KCl, 2.5 $CaCl_2$, 1.2 $MgCl_2$, 20 HEPES, 10 glucose. Excitation of GCaMP6s (488 nm, 20% power) and H2B-mCherry (561 nm, 10% power) was used with a 405/488/561 multi-band dichroic and emission filters (GCaMP6s: 525/50; H2-B-mCherry, 595/50) to report simultaneously report cytosolic Ca^{2+} levels and nuclear location. Timelapse imaging was performed at 0.125 Hz for 20 minutes.

Testing set, first presented in[9]*:* Islet isolation from a Connexin 36 knockout mouse was performed following the methods of Scharp et al. and Stefan et al.,[10] and the islets were maintained in RPMI medium with 10% fetal bovine serum and 11 mM glucose at 37°C in 5% CO2 for 24-48 hours before imaging.
Isolated islets were stained with 4 mM Fluo-4 AM in imaging medium (125 mM NaCl, 5.7 mM KCl, 2.5 mM CaCl2, 1.2 mM MgCl2, 10 mM HEPES, 2 mM glucose, 0.1% BSA, pH 7.4) at room temperature for 1-3 hours. Imaging was conducted using a PDMS microfluidic device that stabilizes the islet and allows for rapid reagent changes. Fluorescence imaging was performed 15 minutes after increasing glucose from 2 mM to 11 mM using an LSM5Live microscope with a 20x 0.8 NA Fluar Objective and a 488 nm diode laser. The device was maintained at 37°C in a humidified, temperature-controlled chamber, with images acquired at 4-6 frames per second, and average power minimized to 200 mW/cm².

Cell Masking: To test the utility of CRISP for improving accuracy of manually drawn cell masks, manual masks were drawn around cells in three ways. Good masks were drawn by zooming in closely on the cell of interest and carefully outlining the exact cell. Each good mask took approximately 15 seconds to draw. These masks were used to define the "true cell". Medium masks were drawn after zooming in to the cell of interest but drawing around the cell without being careful. Each medium mask took approximately 5 seconds to draw. Bad masks were drawn without zooming into the cell. Each bad mask took approximately 1 second to draw.

CRISP: Automated Cell Mask Refinement: The CRISP algorithm compares the correlation coefficient between pixels with a baseline timeseries. This baseline timeseries is defined as the 10% of the pixels nearest to the center of the cell mask. The CRISP score for each pixel was defined as the correlation coefficient between the pixel and the baseline timeseries. To ensure generalizability across systems with different within cell correlation (r), pixels were removed if they had a CRISP score less than the average CRISP score minus the standard deviation of all CRISP scores times a threshold (R_{th}).

To identify the optimal threshold, bad and medium masks were run through CRISP using correlation thresholds between -2 and 0.3. Positive labels were defined as pixels that were labeled to be in the cell mask. For example, true positives were defined as pixels in the "good" cell mask and correctly labeled by CRISP to stay in the cell mask, whereas false positives were defined as pixels that were considered to be in the cell mask (e.g. not removed) by CRISP, but were not present in the "good" cell mask.

CRISP: Semi-minor Axis Identification: The baseline timeseries was defined as all cells within a circular annulus within a 2-pixel radius of the nucleus. The annulus was given a score function (eq 1) based on the number of pixels in the annulus with correlation coefficient less than the threshold (p_+) minus the annulus radius (ϕ) divided by 100. This penalty function ensures that the radius does not get to large in case the cell of interest is directly located next to another highly synchronized cell.

$$Score = p_+ - \phi/100 \quad (1)$$

To identify the optimal score, CRISP algorithm was run for thresholds between 0.7 and 0.95.
All calculations were conducted in MATLAB.
All code will be made publicly available at: https://github.com/jenniferkbriggs/CRISP.

Results

Pixels within on cell are highly correlated compared to outside of cell. To test our hypothesis that pixels within a cell were more highly correlated with each other than with pixels outside the cell, we conducted calcium imaging of a beta cell islet using a Nikon A1R confocal microscope (**Fig 1a**). We then drew a circular mask with radius of 30 pixels centered at a beta cell nucleus and extracted the calcium time course **(Fig 1b).** Strikingly, a colormap showing the correlation coefficient between the center pixels and all other pixels can be used to distinguish a single cell, indicating that pixels within a cell are more correlated with each other than those outside of the cell (**Fig 1c**). Using this finding, we developed CRISP – Correlation Refined Imaged Segmentation Process. Crisp can be used for two tasks. Automated cell mask refinement (**Fig 1c, right top)** uses CRISP to refine manually drawn cell masks. Automated semi-minor axis identification (**Fig 1c, right bottom)** uses CRISP to identify the largest circle that contains only pixels within a cell.

Automatic cell mask refinement with CRISP. To assess CRISP accuracy and identify the optimal correlation threshold for removing pixels, "bad", "medium", and "good" masks were manually around 45 cells (15 cells for three islets) **(Fig 2a).** CRISP was used to refine medium and bad masks and assessed performance by comparing CRISP-masks to good mask. Visual inspection for a bad mask shows that CRISP correctly identifies mostly correct pixels to remove from the manually outlined mask (**Fig 2b).** CRISP mask refinement resulted in an area under the ROC of 0.84 ± 0.03 for bad masks and 0.81 ± 0.05 **(Fig 2c),** indicating that CRISP is effective at mask refinement. An optimal correlation threshold of 0.25 was chosen to maximize true positive rate (correctly keeping pixels that belong to the cell), while minimizing false positive rate (incorrectly keeping pixels that belong to the cell). Before CRISP, there was high variance in the time courses between pixels within cell masks **(Fig 2d)**. CRISP successfully removed these pixels with high variance (or low correlation), thereby decreasing the noise in individual cell masks **(Fig 2e)**.

To ensure generalizability, we tested CRISP using a separate data set, originally reported in[9]. This data set utilize a different microscope (LSM5Live) with lower resolution and microscope settings and a separate mouse genotype (Connexin 36 knockout mice). We tested CRISP on three islets, with 10 masked each (n = 30). For both medium (**Fig 2f)** and bad masks **(Fig 2g)**, CRISP correctly estimated (true positive + true negative) 77% of pixels. Considering the total number of pixels characterized by CRISP was ~1000 and that the testing set came from a completely different mouse model and microscope, this indicates that CRISP generalizes well.

CRISP cell refinement improves correlation coefficient accuracy. Functional network analysis is a common analysis used in fields of islet biology and neuroscience to understand how synchronized different cells are to one another. CRISP automatic cell refinement was run on medium and bad masks for 60 cells from four islets **(Fig 3a).** To test the if CRISP could improve the accuracy of network analysis results, we calculated the correlation coefficient each cell within the islet **(Fig 3b).** Medium masks resulted in an average correlation coefficient error (compared to good masks) of 0.0052. This error resulted in cells have an average difference of 0.93 connections, with a max of 6 connections difference. CRISP significantly decreased this error to 0.005 **(Fig 3c)**. Bad masks resulted in an average correlation coefficient error of 0.0067. This error resulted in an average difference of 1.3 connections, with a max of 6 connections difference. CRISP significantly decreased this correlation error to 0.0057 **(Fig 3d)**. Therefore,

CRISP may be able to improve accuracy of functional network theory calculations, though only slightly.

Automatic semi-minor axis identification with CRISP

To determine the largest radius of a circular cell mask centered at the nucleus that contained only pixels from within the cell, we developed a two-step algorithm. We manually drew masks around 10 cells from 5 islets to serve as training data. In the first algorithmic step, we identified the optimal correlation coefficient that would maximize the correctly labeled cells within a large radius (40 pixels). For each cell in the training set, we created a circular annulus (or field of view) with radius of 40 pixels around the cell. This radius was to ensure that the full cell would be well within the annulus. We set all pixels within a 2-pixel radius from the center of the cell as the baseline to compare to. This baseline is defined differently than for automated cell mask refinement because the approximate cell mask size is not known. We then calculated the correlation coefficient between every cell in the annulus and the baseline pixels over many fixed thresholds. We use a fixed correlation threshold (as opposed to standard deviation-based threshold used in automatic cell refinement) because the large annulus likely contains pixels from other cells, and thus a standard-deviation based threshold maybe biased towards other cells in the annulus. The optimal correlation threshold was 0.818 with an area under the receiver operating curve of 0.8 **(Fig 4a).** Upon visual inspection in some annulus', the cell was accurately identified with only the correlation threshold **(Fig 4b).** However, in other annulus', large numbers of pixels were incorrectly labeled to be inside of the cell **(Fig 4c).** To increase algorithm accuracy, we created an iterative anulus expansion step. In this step, the anulus is initialized at 2 pixels **(Fig 4d, left)**. The score function for that annulus is computed and the anulus radius is increased. To identify the optimal score threshold, we calculated the score for all training cells up to an anulus radius of 40 pixels **(Fig 4d, right)**. The area under the ROC curve of 0.989, indicating that CRISP is highly accurate at semi-minor axis identification **(Fig 4e).** CRISP did over- and underestimated multiple cell semi-minor axes, but only by one or two pixels (**Fig 4f).**

We used CRISP to assess the average semi-minor axis length (radius) of 150 beta cells from five islets. Each of the five islets resulted in a roughly similar radius distribution, which was positively skewed with an average radius of 4.0866 microns (**Fig 4g)**. Only 30% of cells were larger than 4.0866 microns. For very high throughput analyses, such as 3D light-sheet imaging, running CRISP for each cell may be unnecessarily computationally expensive. Instead, it would be optimal to identify a set radius that consistently captures the maximum possible number of pixels. After estimating the average radius, we identified an optimal radius of 4.65 microns, which underestimates approximately 10% of cells **(Fig 4h)**.

Discussion

Time course imaging is used throughout fields of biology. The increase in microscope resolution has provided the opportunity to study activity from individual cells. In nearly all cases, when studying this activity, a user manually masks each cell. This act of manual masking induces user error and can be very time consuming. We sought to remove some of this potential error by developing an algorithm: Correlation-Refined Image Segmentation Process (CRISP) that utilizes intracell correlations to refine cell masks and automatically circle cells.

CRISP cell mask refinement algorithm resulted in what is typically considered an "excellent" area under the curve[11]. When testing CRISP on an independent dataset, the average correct labeling was approximately 77%. Considering that the labels were on individual pixels, and the resultant calcium trace is an average of all pixels, 23% of incorrectly labeled pixels will have a negligible impact on the resultant calcium trace. Further, our testing set was calcium time courses taken between 2007-2008 with very different technology than our training set (which was collected in 2023) and a completely different mouse model. This was intentional to fully test the generalizability of CRISP. With these considerations in mind, 77% accuracy was particularly impressive.

Methods exist for fully automating cell masking using clustering algorithms[5]. These algorithms rely on differences in intensity for clustering and work well on systems were cells are sparsely packed (e.g. neurons). However, they tend to not perform well in systems where cells are densely clustered, such as islets or cardiomyocytes. Correlation-based methods also exist[6,7]. However, these methods utilize an iterative approach where individual pixels are added if under some correlation threshold, similar to Step 1 of Automated Semi-minor axis identification. We showed that these algorithms also do not work well in dense images because multiple highly correlated cells will be lumped together **(Fig 4c).** By including an iterative anulus growth approach, we greatly increase the algorithm accuracy (AUC = 0.989).

This CRISP algorithm can be used to accomplish two tasks: automated cell mask refinement and automated semi-minor axis cell masking. The first task could be beneficial to remove errors induced by fatigue or multiple scientists masking a data set. We do not recommend that CRISP be used as a replacement for careful and thoughtful cell masking. The second task would be beneficial for times when the data are extremely plentiful and perfect masks are not necessary. In this case, CRISP can be used for rapid identification of the majority of pixels within each cell, removing the need for user drawn masks.

Author Contributions

E.J. and R.K.P.B constructed the microscope and performed islet imaging experiments. J.K.B. developed algorithm and conducted analysis. J.K.B drafted the paper. All authors edited the paper. R.K.P.B and M.J.M provided resources

Acknowledgements

We thank Barak Blum at the University of Wisconsin-Madison for providing Rosa26$^{H2B\text{-}mCherry}$ mice and the University of Wisconsin Optical Imaging Core for use of the spinning disk confocal.

The Merrins laboratory gratefully acknowledges support from the NIH/NIDDK (R01DK113103 and R01DK127637to M.J.M., and R01DK106412 to R.K.P.B.) and the United States Department of Veterans AffairsBiomedical Laboratory Research and Development Service (I01BX005113 to M.J.M.). The Benninger laboratory gratefully acknowledges support from the NIH/NIDDK (R01DK106412, R01DK102950, R01DK140904 to R.K.P.B.) and the University of Colorado Diabetes Research center (P30 DK116073). Jennifer K Briggs greatfully acknowledges support from NSF GRFP (DGE-1938058_Briggs).

References

Figures

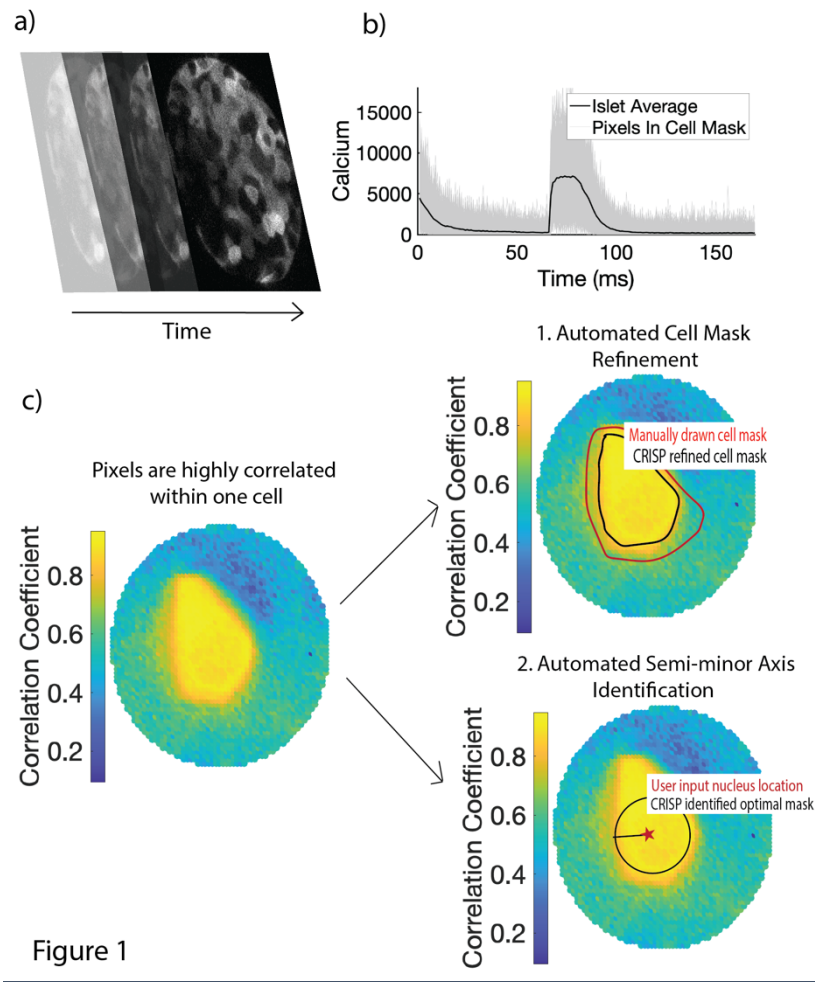

Figure 1: (a) Example of calcium image stacks of pancreatic islet. (b) Resultant calcium time course from (a). (c) Left: Color plot showing the correlation between the nucleus and the rest of the pixels in the region of interest. Cell outline can be visualized only based on the correlation color plot. Using this observation, two algorithms are created. Right: 1. Automated mask refinement takes user defined mask (drawn in red) and removes uncorrelated pixels that are not a part of the cell to create the CRISP refined cell mask (black). 2. Automated semi-minor axis identification begins with the user input nucleus location (red star) and identifies the largest circle radius that captures pixels only within the cell of interest.

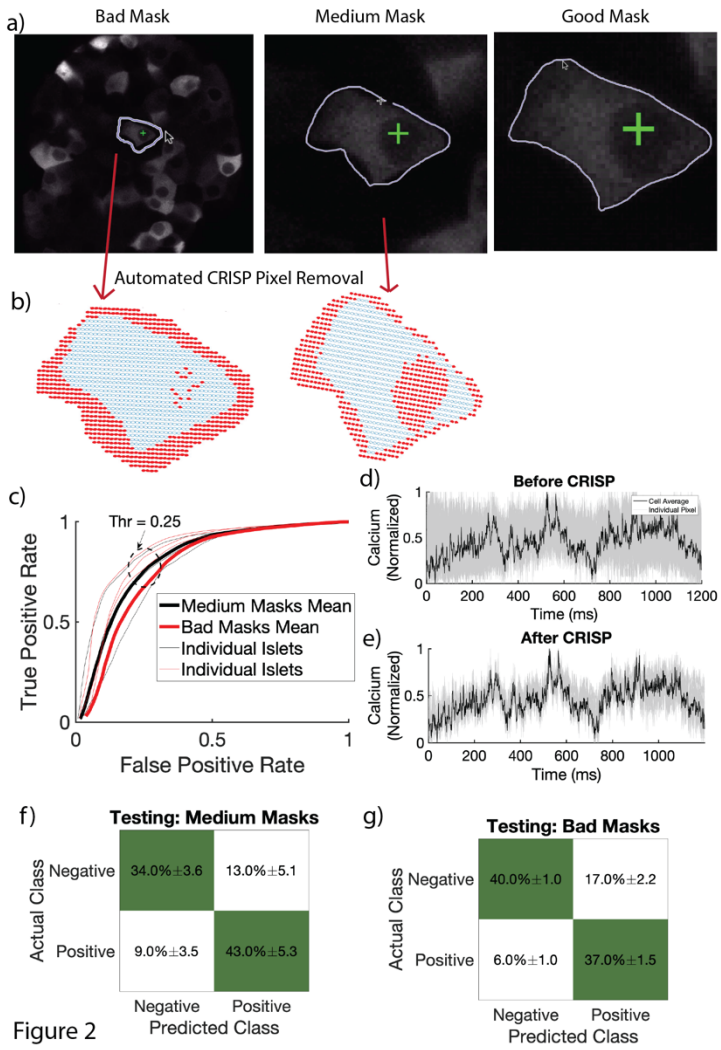

Figure 2: Automated Cell Mask Refinement. a) Representative bad (circled without zooming), medium (circled with zoom), and good (circled carefully with zoom) masks. b) CRISP removed pixels (red) from representative bad and medium masks in (a). c) Receiver operating characteristic curve (ROC) of four islets with fifteen beta cells circled per islet. Red lines indicate ROC curves for bad user defined masks. Black lines indicate ROC curves for medium user defined masks. Dark lines indicate average over four islets, while light lines indicate individual islets. Light lines indicate individual islets. Average area under the curve ***0.83.*** *d) Calcium time course of all pixels in a bad mask before CRISP. Grey lines indicate individual pixels. Black line indicates average calcium over the full cell. e) As in c but after CRISP. f) Confusion matrix showing performance of testing data set for medium masks. g) As in e for bad masks.*

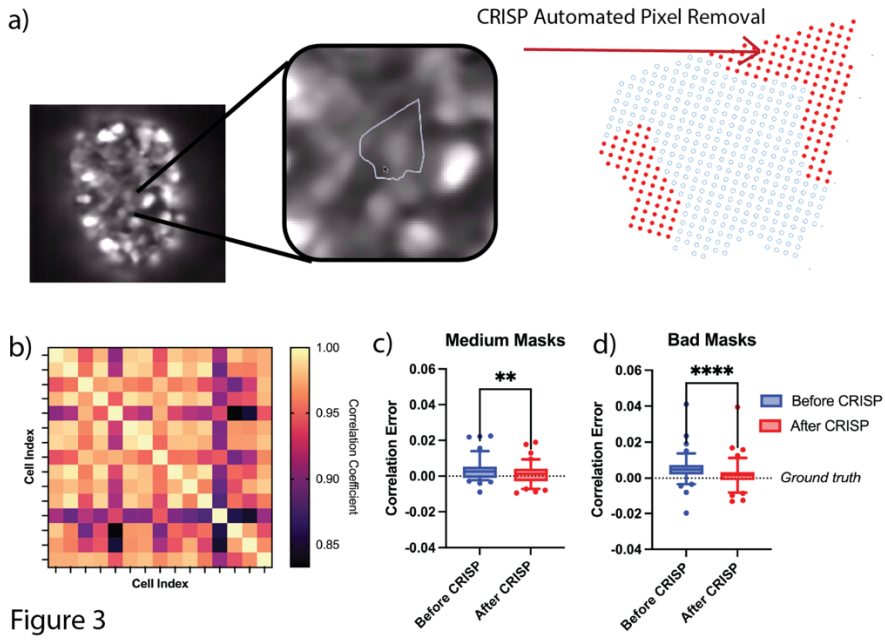

Figure 3: a) Example figure of islet with bad user defined mask. Red pixels (right) are pixels correctly removed from mask by CRISP. b) Heat map showing correlation coefficient between 15 cells in an islet. c) Difference between correlation coefficient from medium and good masks before CRISP (blue) and after CRISP (red) $p = 0.001$, d) As in h but for bad masks $p < 0.001$.

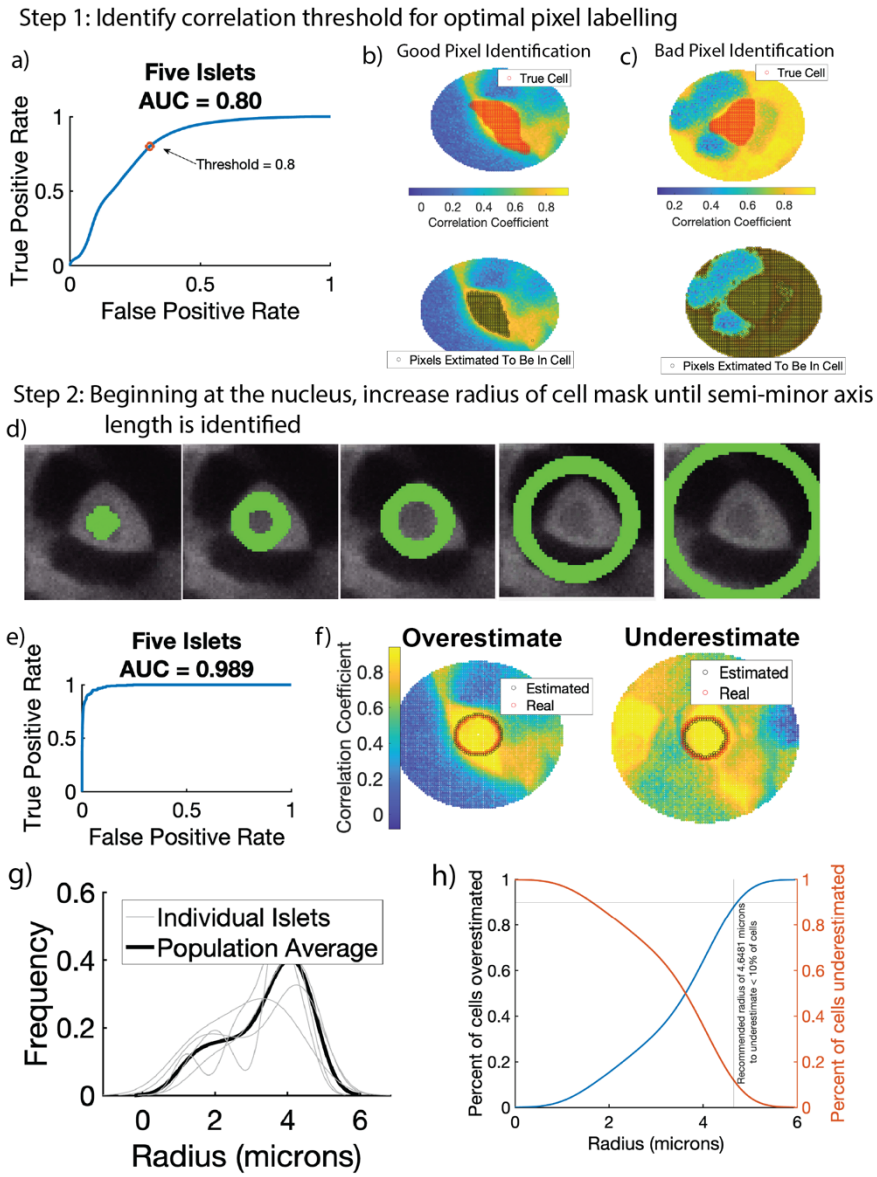

Figure 4

Figure 4: Automated Semi-minor axis identification in two steps. In step 1, the correlation between every pixel in a large region and the user-defined nucleus is calculated. Receiver operating characteristic curve (ROC) had an area under the curve of 0.8. Threshold of 0.8 was identified for step 2. b) Representative example of bad pixel identification. Top shows correct cell mask (red) overlaid on a correlation color map. Bottom shows pixels estimated to be inside the cell (black). c) Representative example of good pixel identification. Correct cell mask (red, top) is closely estimated by CRISP (black, bottom). In step 2, after threshold of 0.8 is identified in step 1, CRISP begins with a small annulus surrounding the nucleus and calculates a score (see methods). CRISP then iteratively increases annulus radius until score threshold is reached. d) Left shows initial anulus radius, middle shows annulus with where CRISP identified optimal radius, right shows annulus that are too large and minimize score. e) ROC curve for CRISP semi-minor axis identification algorithm. Area under the curve is 0.989. f) Example of annulus that is underestimated by CRISP. g) Example of annulus that is overestimated by CRISP. Using CRSIP, the semi-minor axis length was estimated for five islets with approximately 150 cells. a) Distribution of radius of beta cells. b) Recommendation for high throughput beta cell calcium analysis where a single radius is defined for all beta cells.