Methods 85 (2015) 44-53

Contents lists available at ScienceDirect

Methods

journal homepage: www.elsevier.com/locate/ymeth





CrossMark

Computer vision for image-based transcriptomics

Thomas Stoeger^{a,b,1}, Nico Battich^{a,b,1}, Markus D. Herrmann^{a,b}, Yauhen Yakimovich^a, Lucas Pelkmans^{a,*}

^a Faculty of Sciences, Institute of Molecular Life Sciences, University of Zurich, Winterthurerstrasse 190, CH-8057 Zurich, Switzerland ^b Life Science Zurich Graduate School, Ph.D. program in Systems Biology, Switzerland

A R T I C L E I N F O

ABSTRACT

Article history: Received 19 February 2015 Received in revised form 13 April 2015 Accepted 17 May 2015 Available online 23 May 2015

Keywords: Image-based transcriptomics Single-molecule Single-cell Segmentation Localization Subcellular High-throughput FISH In situ hybridization Single-cell transcriptomics has recently emerged as one of the most promising tools for understanding the diversity of the transcriptome among single cells. Image-based transcriptomics is unique compared to other methods as it does not require conversion of RNA to cDNA prior to signal amplification and transcript quantification. Thus, its efficiency in transcript detection is unmatched by other methods. In addition, image-based transcriptomics allows the study of the spatial organization of the transcriptome in single cells at single-molecule, and, when combined with superresolution microscopy, nanometer resolution. However, in order to unlock the full power of image-based transcriptomics, robust computer vision of single molecules and cells is required. Here, we shortly discuss the setup of the experimental pipeline for image-based transcriptomics, and then describe in detail the algorithms that we developed to extract, at high-throughput, robust multivariate feature sets of transcriptome. These computer vision algorithms and pipelines can be downloaded from: https://github.com/pelkmanslab/ImageBasedTranscriptomics.

© 2015 Published by Elsevier Inc.

1. Image-based transcriptomics is unique in several ways

In the past few years a wealth of techniques have been developed to study genome-wide transcriptional output at the single-cell level [1–7]. In contrast to methods relying on sequencing or PCR, image-based transcriptomics visualizes single transcripts in a population of single cells in situ. This allows not only the absolute quantification of transcript copy numbers, but also the spatial mapping of transcript molecules to the sub-cellular microenvironment [4]. Being an *in situ* technology, it does not require homogenization of cells and therefore minimizes the loss of material, thus achieving very high detection efficiency [4]. Another advantage of image-based transcriptomics is that it can be combined with the phenotypic characterisation of each single cell and its context within a population of cells or tissue, by microscopic assays and stainings commonly used in cell and developmental biology. This makes image-based transcriptomics of particular interest when studying the localization dynamics of the transcriptome in response to stimuli or perturbations and to identify sources of cell-to-cell variability in these processes [8,9]. While establishing image-based transcriptomics, we soon realized

* Corresponding author.

¹ These authors contributed equally to this work.

that a robust computer vision pipeline was as important as the experimental platform for accurately identifying and characterizing each single transcript molecule within a cell. Therefore, we here describe in detail our recent computer vision algorithms that result in accurate detection of objects in spinning disk confocal microscopy images. Besides providing a robust guide for identifying billions of individual transcript molecules with little hands-on user time, we describe how to unlock functionally important parameters of gene expression, which are impossible to grasp without the power of computer vision. For instance, multivariate descriptors of the position of each single transcript molecule enable an unsupervised characterization of the localization of transcripts of every cell.

1.1. General outline

Image-based transcriptomics employs multi-well plates to stain cells in parallel with specific probes against a transcript of interest (Fig. 1). Within single wells of a multi-well plate, the transcripts of different genes are stained by an automated experimental procedure. Each single transcript molecule is detected by highthroughput microscopy and computer vision. Experimental and computational steps can be performed with equipment that is commonly used for image-based high-throughput assays.

E-mail address: lucas.pelkmans@imls.uzh.ch (L. Pelkmans).



Fig. 1. Outline of image-based transcriptomics using bDNA sm-FISH.

Each single transcript molecule is stained by branched DNA single-molecule in situ hybridization (bDNA sm-FISH). This technology, which is commercially available from Affymetrix and Advanced Cell Diagnostics, applies a series of consecutive in situ hybridizations, which visualize each single transcript molecule as a bright fluorescent spot. In a first round of in situ hybridization, two epitope-specific primary probes bind next to each other on the same transcript molecule. While it is technically possible to implement bDNA FISH with only one epitope-specific probe [10], requiring the simultaneous binding of two probes in direct spatial adjacency should reduce unspecific signal [11]. Targeting 15 different epitopes of each transcript in a single hybridization reaction ensures that at least one epitope is accessible to the detection reagents without the need to denature the specimen. The subsequent rounds of *in situ* hybridization create a docking platform for \sim 500 fluorescently labelled probes per single epitope. This level of fluorescence is sufficiently high to enable the specific, rapid and robust detection of single transcript molecules by high-throughput imaging.

1.2. Alternative methods for RNA detection in imaging

Another method for directly visualizing single transcript molecules *in situ* is oligonucleotide-based single molecule FISH (o-nuc sm-FISH). This approach targets individual transcripts by up to 40 different oligonucleotides, which are directly conjugated to fluorophores. While a recent study achieved to monitor 61 different ncRNAs, it had to restrict itself to "a few dozen cells ... due to limited imaging throughput" [12]. Possibly, this reflects the lower signal-to-noise ratio of single fluorescent spots of o-nuc sm-FISH and their need for a 600 times longer illumination time [4].

Alternatively, transcripts can be visualized indirectly via reverse transcription to cDNA that can be sequenced *in situ* by padlock probes [13] or oligonucleotide ligation and detection [14,15]. While the former sequencing approach can presently detect 31 different genes simultaneously in thousands of single cells within a tissue slide [13], the latter approach can currently read around 200 mRNAs simultaneously for 40 different cells [15]. The efficiency for detecting single transcript molecules has been estimated to be 30% [13,16] and 3% [15] respectively, which is much lower than the 85% of hybridization efficiency in sm-FISH [4,17]. Such low efficiencies currently prevent these alternative methods from surveying the transcriptome with single-molecule sensitivity and resolution *in situ* [18,19].

2. Establishing image-based transcriptomics with single molecule resolution

The detailed experimental protocol for high-throughput bDNA sm-FISH has been published previously [4] and therefore, we here mainly provide additional assistance for setting up a robust

automated experimental platform. As a general introduction to high-throughput image-based assays and the infrastructure and software supporting such experiments we highly recommend the excellent essay by Buchser and colleagues [20].

Table 1 contains an overview of potential problems occurring during the detection of single transcripts. The most critical factor in getting reliable results is to use an automated incubator that contains rotating towers for the individual storing of multi-well plates during hybridization reactions. This prevents the occurrence of different hybridization efficiencies in different wells of a multi-well plate (data not shown). Table 2 highlights potential pitfalls, which could affect the biological interpretation of accurate single-molecule measurements. We recommend repeating the control experiments suggested in Table 1 and Table 2 in different weeks to ensure that your setup of image-based transcriptomics functions robustly.

3. Establishing the image analysis pipeline

A robust image analysis pipeline is required for accurate measurements of absolute transcript levels as well as measurements of transcript localization in the cytoplasm of single cells, and extraction of features that describe the cellular phenotype. First, homogeneous intensity values throughout the images in all channels must be ensured, and then object segmentation must be performed minimizing errors. To ensure this, we developed four algorithms to perform high-throughput illumination correction of raw images, robust nuclei and cell segmentation, and robust spot detection. They can be downloaded from https://github.com/pelkmanslab/ImageBasedTranscriptomics and applied on an example dataset available on https://image-based-transcriptomics.org. The algorithms presented in this manuscript do not intend to replace single-cell quality control. For the latter we recommend interactive user-guided supervised machine learning, which has been implemented before by our group [23] and others [24]. Supervised machine learning not only readily identifies rare cells that have not been correctly segmented, but also allows the selection of a group of cellular objects that is relevant for a specific biological question (e.g., interphase cells).

The algorithms presented in this manuscript intend to reduce human hands on time and increase the amount of high-quality primary data after computational image-analysis (Table 3). Computational running time has not emerged as a practical issue for image-based transcriptomics. The algorithms are robust in the sense that their input parameters rarely have to be adjusted for individual experimental plates.

While the principles of the algorithms presented in this manuscript have been sketched in one of our earlier publications, the description beneath provide a detailed guide for using those algorithms. Moreover we here include implementations of these

Table 1

Suggested controls for the detection of transcript molecules.

Possible artefact	Experiment	Hints
Inability to detect single molecules	Assay with probes against a single epitope of HPRT1 [4]	Exposure time during imaging; protease concentration
False positive detection	Probe against bacterial gene dapB. Less than 1 spot per cell should be detected	Protease concentration; cells without cytoplasmic DNA
Spill-overs	Stain adjacent wells for the negative control (bacterial dapB) and the highly abundant ACTB transcripts. Test full plates	Liquid handling
Efficiency of single molecule detection	Stain same transcript on two different sets of epitopes by two different sets of amplification reagents, which can be visualized by two different fluorophores. Efficiency of detection should be $\sim 85\%$ [4]	Protease concentration; amount of targeted epitopes per transcript; computational spot detection
Positional bias between wells (staining reaction)	Stain all wells of a multi-well plate with probes against the non-abundant housekeeping gene HPRT1	Always use incubator with rotating towers for hybridization; never skip protocolled in-solution mixing
Low reproducibility	Multiple independent assays across different weeks	Aberration of liquid handling < 1%; cell seeding
Tearing of signal of single molecules	Compare signal obtained by multiple units and types of objectives	Choose best objective and remove remaining effect computationally (see below)

algorithms for MATLAB and implementations as modules for CellProfiler to segment single nuclei and cells.

3.1. Illumination correction

Illumination correction of raw images is essential for subsequent steps in the image analysis pipeline. It ensures correct object detection and accurate measurements of intensity features, reducing biases due to uneven illumination of the sample as well as positional differences in the signal gain resulting from the detection system. During image-based transcriptomics, we exploit the statistical power of the large number of images acquired per channel to learn pixel-wise illumination and signal gain biases (Battich et al. [4], Fig 2). Briefly, we calculate the standard deviation and mean intensity values per pixel for every pixel position of a given channel. To correct the illumination bias, per-pixel *z*-scoring is performed as shown in Fig. 2(Eq. (1)). The *z*-scored values are then reversed to intensity values as shown in Fig. 2(Eq. (2)).

3.2. Nucleus segmentation

Pixels belonging to nuclei objects can be easily distinguished from background pixels by thresholding an image of a nuclei-specific stain such as DAPI. However, this often results in

Table 2

Suggeste	d contro	ls for	the	proper	interpretation	ot	single	e mol	lecul	le measurements.	
----------	----------	--------	-----	--------	----------------	----	--------	-------	-------	------------------	--

Possible artefact	Experiment	Hints	Illuminat
Positional bias	Compare the number of	Avoid "edge-effects" by	manus
between	cells and local cell density	following the cell seeding	Nucleus s
wells	[21] across individual wells	protocol of Lundholt et al.	(CellPi
(biological)		[22]	Nucleus s
Variable number	Monitor and potentially	Trypsinization time;	manus
of cells per	adjust cell dissociation	repeatedly shear cells	Cell segm
seed	protocol such that, on	through pipette pressed	(CellPr
	average, a cell aggregation	towards plastic dish	Cell segm
	score of less than 1.2 is		manus
	achieved for each seed		Cell segm
Loss of cells	Perform live-imaging of	Slowly pipet to side of well	segme
during assay	cells with Hoechst dye	(most steps) or center of	Inference
	prior to the assay and	well (only for in-solution	param
	compare with presence of	mixing)	manus
	cells after image-based		Optional
	transcriptomics		correc
Staining of cell-	Repeat and time	Time-dependent decay of	Spot dete
outline varies	succinimidyl ester staining	carboxylic acid,	manus
between	with multiple freshly	succinimidyl ester, in	Measurin
experiments	prepared staining solutions	aqueous solutions	transc

clumps of several nuclei because a single, image-wide threshold value is generally not sufficient to separate nuclei that lie very close to each other. Such clumped objects are relatively large and display multiple concave regions. Generally, at the intersection of individual objects, a line of low intensity pixels connects two concave regions, which can be found by the watershed algorithm [25]. Thus, we identify single nuclei with an algorithm consisting of two parts: first, intensity thresholding by the Otsu method [26] identifies primary objects; and, secondly, objects consisting of multiple nuclei are separated along the best identified watershed line (Fig. 3).

The algorithm (algorithm 1) uses illumination-corrected images and processes them as follows:

- 1) Initial objects are identified by simple thresholding.
- 2) Clumped objects are selected on the basis of size and shape features: area, solidity, and form factor.
- 3) The perimeter of selected objects is analysed and concave regions along the boundary of objects are identified.
- Putative watershed lines connecting two concave regions are determined using the Dijkstra shortest-path algorithm [27].

Table 3

Time for image-based transcriptomics on ten 384-well plates to obtain results, whose quality appeared acceptable to us. Time estimates are based on our experience and depend upon the specific computational infrastructure.

	Hands-on time	Computational time	Computers used by us
Illumination correction (this manuscript)	5 min	2–5 h	4
Nucleus segmentation (CellProfiler)	30 min	<<1 h	1500
Nucleus segmentation (this manuscript)	30 min	<<1 h	1500
Cell segmentation (CellProfiler)	5–10 h	<<1 h	1500
Cell segmentation (this manuscript)	5 min	2–10 h	1500
Cell segmentation (manual segmentation)	>> 1 month (expected)		
Inference of spot detection parameters (this manuscript)	1–2 h	2-8 h	10
Optional lens aberration correction	1 h	1–2 h	1500
Spot detection (this manuscript)	30 min	<<1 h	1500
Measuring localization of transcripts	5 min	1–2 h	1500

- 5) All possible cuts along the selected watershed lines are considered and features of each potential cutting line (intensity along the line, angle between concave regions) as well as features of the resulting objects (area/shape) are measured.
- 6) An "optimal" cut line is finally chosen by minimizing a cost function that evaluates the measured features. The resulting objects have a minimal size and are as round as possible, while the separating line is as straight and as short as possible and the intensity along the line as low as possible.

Algorithm 1 IdentifyPrimaryIterative

- Initialize() // initialize objects by thresholding the input intensity image;
- 2. *InitialSegmentation() //* recognize objects as *segmented objects* without cutting them first;
- 3. Repeat
- 4. **For each** object **in** segmented objects
- 5. **If** lower size threshold < size of object < upper size threshold
- 6. **and** solidity of object < solidity threshold
- 7. **and** transformed form factor of object > *form factor threshold*
- 8. then
- 9 Add *object* to the collection of *clumped objects* to be cut;
- 10. end
- 11. End
- PerimeterAnalysis() // analyze perimeter of selected clumped objects and calculate the curvature along their boundary [see PerimeterAnalysis.m];
- PerimeterWatershedSegmentation() // cut selected clumped objects along watershed lines between concave regions [see PerimeterWatershedSegmentation.m];
- 14. For each object in clumped objects
- 15. *IdentifyConcaveRegions()* of the *object*, **where** *region* is concave
- 16. **If** angle of circle segment of *region* > *equivalent* segment threshold
- 17. **and** radius of region < equivalent radius threshold
- IdentifyLinesAndNodes() of the object // find all watershed lines and nodes, where each node is a single pixel on the line that overlaps with the object boundary;
- 19. **Select** watershed nodes **If** node lies within concave regions;
- 20. **Select** all *watershed lines* **If** *line* connects two *watershed nodes* from different concave regions;
- 21. For each line in watershed lines
- 22. Measure line *length* and *straightness* and the *intensity profile* along the line;
- 23. Measure the *angle* between normal vectors at watershed nodes;
- 24. Measure *area*, *solidity* and *form factor* of the cut object, i.e. the smaller of the two objects that would result from a cut along the line.
- 25. If size of the object < threshold of object being too small
- 26. **then**

discard such cutting *line* from selected *watershed lines*;

- 27. end
- 28. **Select** "optimal" watershed *line* by **minimizing** the cost function

Algorithm (continued)	Algorithm	(continued)
-----------------------	-----------	-------------

Algo	orithm 1 IdentifyPrimaryIterative
29.	cost (a, b, c, d, e, f, g, h) $\leftarrow a - 2 * b - c - d - e + 2$
	*f - g - 2 * h,
30.	where
	a is a solidity of cut object,
	b is a form factor of cut object,
	c is a mean intensity along the <i>line</i> ,
	d is a max intensity along the <i>line</i> ,
	e is a 0.75 quantile intensity along the <i>line</i> ,
	f is an angle between two watershed nodes,
	g is a line straightness,
	h is a line length.
31.	End // of for-each-loop at line 21
32.	End // of for-each-loop at line 14
33.	Until no more <i>clumped objects</i> can be found.

Whenever attempting to identify individual nuclei of a novel cell line or whenever changing imaging conditions, we recommend to empirically test different schemes and parameters for segmentation of nuclei. Good settings can usually be found empirically by using the inbuilt testing mode of IdentifyPrimaryIterative. Contrasting CellProfiler's default options for separating objects, which are part of the IdentifyPrimaryAutomatic module, IdentifyPrimaryIterative can simultaneously consider the local intensity of the DAPI stain and the geometry of identified objects to separate them. In practice we never had to adjust the threshold value suggested by the Otsu method [26]. Depending on the biological question of interest, one might choose settings for the separation of objects, which under-segmentation. favour overor For instance over-segmentation increases the fraction of emerging cells during anaphase cells that are already considered as individual objects. Under-segmentation on the other hand facilitates the correct segmentation and thus quantification of multinucleate cells.

Frequently not every object, which can be identified in image-based assays, should be considered in subsequent analysis. For instance we preclude the analysis of DAPI positive cellular debris, apoptotic bodies and mitotic cells by a dual strategy, which is independent of IdentifyPrimaryIterative. First the DiscardObjectsBySize.m module removes small objects within CellProfiler. Second, supervised machine learning identifies debris, and apoptotic and mitotic cells [23].

3.3. Cell segmentation

The segmentation of cells uses the segmentation of nuclei as seeds [28]. It is imperative to ensure correct segmentation of the cellular cytoplasm as this will not only have a major impact in the number of spots (or transcript molecules) allocated to each cell, but will also drastically affect measurements of transcript localization. To achieve the high accuracy in cell segmentation required for image-based transcriptomics, we developed an algorithm that performs sequential rounds of watershedding, rather than the one round of watershedding typically applied [28]. This iterative algorithm allows accurate identification of the boundary between cells with relatively minimal user input.

In the algorithm, an arbitrary amount of different segmentations are combined in such a way that the allocation of single pixels to their correct seeds (nuclei) never becomes worse and thus becomes optimal by iteratively performing many different segmentations (Fig. 4). Besides largely eliminating human hands-on time, this strategy generally yields superior results compared to a single segmentation: different parts of a single cell can be segmented by opposing segmentation settings, which only yield



Fig. 2. (A and B) Method for illumination correction of images. For each channel the mean intensity μ_i the standard deviation σ_i are calculated for each pixel p_i in the field of view. Then an overall mean intensity $\bar{\mu}$ as well as the mean standard deviation $\bar{\sigma}$ of all pixels is derived from the "mean" and "std." matrices. Illumination correction is performed by per-pixel *z*-scoring (Eq. (1)), where z_i the *z*-scored value for pixel p_i and ln_i is the original intensity value for pixel p_i in a given image. The corrected intensity value C_i for pixel p_i in an image was then calculated as in Eq. (2). C) Illumination correction examples for the DAPI channel. D) As in C but for Alexa Fluor succinimidyl ester (a general protein stain).

optimal segmentation accuracy in a subpart of the cell, but perform sub-optimally in other subparts.

Briefly, the algorithm (Algorithm 2) treats the input images as follows:

- 1) Calculate the watershed cell segmentation at different thresholds.
- 2) One label image is constructed. If a pixel is part of different objects at a given threshold (which is likely in cell-rich regions), it will be allocated to the object of the higher threshold (e.g. if threshold specifications were 1 and 0.5, it would be attributed to the object identified with a threshold of 1).
- Define background pixels by a single user-provided microscope-specific threshold, which can be determined manually once.
- 4) Re-label pixels of prospective objects (cells), which are not connected to their original seed (nucleus), as pixels belonging to the background.

Algorithm 2 IdentifySecondaryIterative

- 1. 1. Initialize empty FinalSegmentation matrix;
- Load OrigInputImage matrix;DefineThresholds(OrigInputImage) // defines a

Algorithm (continued)

lgorithm 2 IdentifySecondaryIterative	
sorted list of <i>thresholds</i> $\{T_i\}$, where T_{min} is a minimalthreshold [e.g. chosen by CPthreshol $T_i < T_{i-1}$	d.m] and

- 4. For each threshold T_i in thresholds $\{T_i\}$
- 5. Obtain binary *ThresholdedImage* of pixels, where each *pixel* = 1
 If pixel intensity > T_i and *pixel* not in *foreground* else *pixel* = 0;
- 6. Find labeled segmentation
 S_i ← WatershedMethod(ThresholdedImage,
 SobelGradient(OrigInputImage), SeedMarkersImage); //
 [see IdentifySecondary.m for 'Watershed' choice of
 method];
- Select *indexes* of all non-background pixels in segmentation S_i;
- 8. FinalSegmentation(indexes) $\leftarrow S_i$ (indexes); // overwrite selected pixels within FinalSegmentation with the label values in S_i ;

А



Fig. 3. (A) Scheme for nuclei segmentation and iterative correction of primary segmented objects. (B) Strategy for selection of objects to be separated by combining the object solidity, form factor and area. All features measured as in the CellProfiler module "MeasureObjectAreaShape.m".

Algorithm (continued)

Algorithm 2 IdentifySecondaryIterative

- 9. End
- 10. *CleanSegmentation(FinalSegmentation) //* Make sure *FinalSegmentation* complies to CellProfiler expectations.

Function CleanSegmentation(FinalSegmentation):

1. Identify borders between different object labels as the non-zero values upon Sobel-filtering, i.e.

Algorithm (continued)

Alg	gorithm 2 IdentifySecondaryIterative
	SobelGradient(FinalSegmentation);
2.	Set the identified borders between different label values
	of objects to background;
3.	Set pixels with the object labels to background value, if
	other pixels with the same object labels do not connect
	them to the seed with the same label values of objects.



Fig. 4. Improvement of segmentation of cells by iterative correction. Several different and partially overlapping segmentations are combined to a single optimal segmentation (Panel A). Detection of single cells stained by carboxylic acid, succinimidyl ester (Panel B).

In our experience IdentifySecondaryIterative has never been performing worse than segmentation by a single round of wathershedding. The few remaining miss-segmented cells can be identified by supervised machine learning. As with any image-based assay the ability to resolve fine structures of the cellular periphery depends upon their size and the resolution of the microscopic images. Like other algorithms that segment 2D images to segment cells, IdentifySecondaryIterative works best on cells that do not grow on top of each other, such as HeLa Kyoto cells, RPEI cells and primary keratinocytes. If cells can grow on top of each other, it is not always possible to allocate a single pixel to a single cell (e.g.: A431, NIH 3T3, HEK293), though supervised machine learning could be used to discard those cells, which grow on top of each other.

3.4. Spot detection and correction of lens aberrations

The basic strategy for detecting single transcripts as spots has been developed by Jiri Matas [32] and Arjun Raj and their colleagues [17]. After emphasizing spot-like signal by a Laplacian of Gaussian filter (Fig. 5A), a threshold for the detection of objects is chosen such that, on each single image, the specific value of the threshold only mildly affects the number of detected transcripts (Fig. 5B and C). As the numerical value of the threshold will partially depend upon the absolute intensity of the acquired images, we rescale the intensities of individual images such that they are comparable between different images and a single numerical value for the threshold can be chosen (Fig. 5C). This seemingly minor, but essential, detail of our image-analysis pipeline contrasts the most common high-throughput implementation of spot detection algorithms, which rescales the intensities of any image according to the intensities of its dimmest and brightest pixel [17,28,33]. While the accompanying code supports additional refinement of the spot detection, these additional parameters (2D/3D, minimal intensity of pixels, size of spots) have a negligible effect on the detection of transcripts once robust imaging conditions have been established experimentally.

For identifying the settings for detecting spots, we include on each experimental plate 4 wells in which bacterial dapB transcripts are probed (a negative control for mammalian cells), and 4 wells for probing transcripts of the housekeeping gene HPRT1 (Hypoxanthine Phosphoribosyltransferase 1, which plays a central role in purine nucleotide synthesis). In a first computational step, we find the upper- and lower-image intensity boundaries for those two reference transcripts (see Exp_getIntensitiesOfReferences.m). The next step detects spots at varying threshold values while rescaling the intensities of single images according to the previously identified bounds (see Exp_getSpotCountsOfReferences.m). Upon completion of the computation, a threshold is chosen manually such that its specific numerical value only mildly affects the number of detected transcripts (see Exp_selectDetectionThreshol d.m). In practice, a fast manual choice and optimization of settings is as good as a fully computational procedure, but offers the advantage of being a first quality control of the data. The number of spots in the dapB negative control should be much lower and



Fig. 5. Detection of single transcripts as spots. Application of a Laplacian of Gaussian (LoG) filter emphasizes spot-like signals (Panel A). Workflow for detecting transcripts as spots (Panel B). The specific numerical value for the detection threshold only mildly affects the number of spots once the intensities of individual images are rescaled similarly. Lines represent five different, randomly chosen images; arrows and asterisk indicate suggested thresholds (Panel C). The signal of individual transcripts is slightly teared in the corner of an image (Panel D).



Fig. 6. Readouts of single-transcript localization (Panel A). Pipeline of converting single-transcript readouts to single-cell readouts (Panel B). Inspecting expected behavior of basic measurements of the localization of transcripts. The distance of transcripts to the nucleus is shorter for transcripts translated at the ER (Panel C). Median distance of all transcripts is normalized by *z*-scoring against 100 relocalizations of the transcripts to random pixels of the cytoplasm. Median of all cells over all single-cell medians is shown (Panel D). Differing distances to the nucleus become apparent to humans in large cells upon visualizing transcripts (green), the nucleus (blue) and the cell outline (white lines) (Panel D, numbers as in Panel C).

more sensitive towards changes in the numerical value of the threshold [4].

Optical aberrations, which tear the signal of individual transcript molecules in the corners of an image, make the signal less spot-shaped. This creates a spatial bias in the detection of transcript molecules of approximately \pm 3% at different positions of an image [4]. While it is best to reduce this effect experimentally (see above), it can be optionally attenuated further by computationally modifying the threshold for the spot detection at different positions of an image. Use the ScanSpotThresholds.m CellProfiler module to test multiple different thresholds surrounding the previously identified reference threshold. Inclusion of all images of a plate (recommended: approximately 10,000 images), allows computing the spatial bias of the detection of spots, which can be used to construct a correction matrix that will modify the spot detection threshold for each pixel (see Exp_computeCorrectionMatrix.m).

You can now identify spots with a CellProfiler pipeline containing the IdentifySpots2D.m module; optionally apply a correction matrix against the spatial bias, which can be loaded by the LoadSingleMatrix.m module; and, insert the parameters for the spot detection determined above. Additionally, we recommend enabling the deblending option, an algorithm from astrophysics [34], which can spatially resolve individual transcript molecules below the optical diffraction limit. If a correction matrix for the spatial bias has been applied, monitor its impact on the spatial bias of the spot detection (see Exp_checkBiasCorrection.m) and potentially restrict or expand the range of thresholds that have been considered for the construction of the correction matrix. In addition to the algorithm outlined above, which provides highly reproducible and specific measurements of the number of transcripts in a high-throughput experimental setup with bDNA sm-FISH [4], we would like to note several excellent algorithms, that have been used with o-nuc sm-FISH to identify those fluorescent spots that likely indicate single transcripts [29–31].

3.5. Quantification of spot localization

Being an *in situ* technology, image-based transcriptomics can quantify the localization of each single transcript molecule. Although the subcellular localization of transcripts and it variability across single cells can hold more biological information than single-cell transcript abundance [4], it is not yet used routinely in functional genomics studies due to technical limitations. This section describes how this powerful source of information can be unlocked from image-based transcriptomics data.

Each single transcript molecule can be characterized by a set of measurements (Algorithm 3), which describe its distance to the centroid or edge of an organelle or the cell [4]. In addition, the position of each transcript molecule can be placed in relation to other molecules, for instance by measuring the variance of its pairwise distances to all other molecules, or by counting the number of transcript molecules within a certain area. Such readouts of single molecules are created by the MeasureLocalizationOfSpots.m CellProfiler module [4]. By choosing an arbitrary amount of differently sized areas, different scales of subcellular crowding can be compared.

Algorithm 3 CPgetSpotLocalizations(LookupImage, VectorWithDistancesForFractions, VectorWithDistanceContainingFractionOfSpots)

- Initialize *Results* // a key-value array, containing all measurements;
 Define South Distance of University (artacia)
- 2. Define SpotDistances as all Euclidean distances between spot pairs (cartesian product);
- 3. // Determine fractions of spots within given distance and distances for given fractions of spots;
- 4. For each spot in spots
- 5. For each DistanceOfFraction in VectorWithDistancesForFractions
- 6. *Results*[*FractionOfSpotsAtDistance*] ← *normalize* over
- 7. **Select** all spots within given *DistanceOfFraction* **exclude** the spot itself;
- 8. End
- 9. For each DistanceContainingFractionOfSpots in VectorWithDistanceContainingFractionOfSpots
- 10. *Results*[*DistanceContainingFractionOfSpots*] **Select** min(
- 11. all SpotDistances for a given spot within DistanceContainingFractionOfSpots exclude the spot itself);
- 12. End
- 13. End
- 14. *Results*[*MeanDistance*] ← *mean*(columns of *SpotDistances*);
- 15. *Results*[*VarianceDistance*] ← *variance*(columns of *SpotDistances*);
- 16. *Results*[*StandardDeviationDistance*] ~ *sqrt*(*ResultsVarianceDistance*);
- 17. *Results*[*DistanceToCellCentroid*] ← measure distances of all *spots* to centroid of the cell;
- 18. // Treat spots at the cell membrane specially.
- 19. For each spot in spots
- 20. Determine coordinate of the *closest membrane pixel*;
- 21. Results[ShortestDistanceToMembrane] ← EuclidianDistance(centroid of spot, closest membrane pixel);
- $\label{eq:linear} 22. \quad Results [DistanceToNucleus] \leftarrow EuclidianDistance(centroid of spot, centroid of nucleus);$
- 23. **If** *EuclidianDistance(centroid of spot, closest membrane pixel) > sqrt(2)* **then**//spot is not at the membrane;Construct a projection line connecting the *centroid of the nucleus* and *centroid of the spot*;
- 24. *Results*[*DistanceAlongProjection*] *EuclidianDistance*(*centroid of spot, closest membrane pixel*) // membrane pixel is picked along the projection line;
- 25. else
- 26. $Results[DistanceAlongProjection] \leftarrow ResultsShortestDistanceToMembrane]$
- 27. **end**
- 28. End
- 29. *Results*[*MembraneBorderingCell*] ← look up pixel at position within *LookupImage* // *LookupImage* is an image indicating for each pixel, whether closest membrane is adjacent to a cell);
- 30. return Results.

Cellular readouts of transcript localization can be derived from the readouts of single transcript molecules. For instance, one may compute the first central moments of the distribution of every readout across all single transcript molecules within a single cell with the accompanying MeasureChildren.m CellProfiler module [4], and subsequently quantify properties of the single-cell distributions of these central moments. In practice, these information-rich multivariate readouts for each single cell, generated for thousands of cells in a single population, rarely lend themselves to ready interpretation or presentation. Therefore, we have previously developed and documented [4] an unsupervised clustering scheme that uses selected cellular statistics to identify a small number of main patterns in single cell subcellular transcript localization. This analysis has been well described by us [4] and can be computed independently of CellProfiler by our locpatterns package (https://github.com/pelkmanslab/locpatterns). Briefly, this package uses the per-cell mean and standard deviation of the single-transcript localization features to first identify a number of different patterns, by clustering random subsets of cells, such that the clusters are most reproducible. In a second step, it determines the similarity of each single cell to each of the identified patterns.

Supervised machine learning can be further applied to classify cells with a distinctive subcellular localization of transcripts [23].

One convenient way to evaluate the basic computational quantification of the localization of transcript molecules is the median distance of all transcript molecules to the nucleus. Plotting the median of this single-cell readout for multiple genes should yield a bimodal distribution (Fig. 6A): transcripts, which become translated at the endoplasmic reticulum (ER), should have a shorter distance to the nucleus compared to transcripts with a cytoplasmic translation. For instance, we noticed that transcripts of RAB13, which have previously been described to enrich in filopodia [35], tended to be furthest from the nucleus (Fig. 6B). One way of controlling finer details of the localization of transcripts is the unbiased clustering of genes by multiple readouts of the localization of transcripts. Mitochondrially-encoded transcripts should be identified as a group of colocalizing transcripts even when mitochondria are not stained [4]. Furthermore, at least in HeLa cells, one should observe a further sub-clustering of different groups of mitochondrially-encoded transcripts reflecting different positions within the mitochondria [4]. In addition, this may reveal further subclustering of transcripts translated at the ER [4], as well as transcripts translated in the cytoplasm. Such findings indicate extensive functional subcompartmentalization of the transcriptome, both on organelles and in the cytoplasm, which are properties of posttranscriptional control of gene expression that have remained hidden thus far.

4. Conclusion

Image-based transcriptomics combines precise counting of transcript molecules with a unique multivariate quantification of the subcellular position of each single transcript molecule for thousands of genes in tens of thousands of single cells. Being an image-based *in situ* technology it can be readily combined with image-based assays, which monitor additional specific biological markers of interest. To enable such lines of research, every experimental and computational step of image-based transcriptomics needs to be highly reproducible across different weeks and geared towards the quantification of single molecules. To enable image-based transcriptomics to reach its full potential, we developed computer vision algorithms that build on and improve those currently used to detect objects in confocal images. By using iterative watershedding we have improved the segmentations of nuclei and cells. In addition, we describe how to perform spot detection for transcript identification in an automated way for thousands of images. Accurate detection of nuclear outlines, cell outlines, and transcript molecules are essential for the correct quantification of a high-dimensional multivariate feature space of each transcript and to reveal bona fide novel properties of the spatial organization of the transcriptome [4]. The computer vision pipeline presented here complements our earlier work [4], and can be used independently of transcripts in other image-based approaches. It also forms a practical guide on how to extend image-based-assays to mapping small particles relative to spatial hallmarks of single cells. Indeed, the highly robust and automated protocol of the underlying computer vision pipeline has been instrumental for uncovering parameters of gene expression, which remain otherwise hidden.

Acknowledgments

We would like to acknowledge A. Schwab for help on the development of the IdentifyPrimaryIterative.m module, Q. Nguyen and S. Lai from Affymetrix for helpful comments on experimental procedures, and V. Green for useful comments on the manuscript. L.P. acknowledges financial support for this project from the Swiss National Science Foundation, the University of Zurich and the University of Zurich Research Priority Program in Systems Biology and Functional Genomics.

References

- [1] F. Tang et al., Nat. Methods 6 (2009) 377–382.
- [2] S. Islam et al., Genome Res. 21 (2011) 1160–1167.
- [3] T. Hashimshony, F. Wagner, N. Sher, I. Yanai, Cell Rep. 2 (2012) 666-673.
- [4] N. Battich, T. Stoeger, L. Pelkmans, Nat. Methods 10 (2013) 1127–1133.
- [5] S. Picelli et al., Nat. Protoc. 9 (2014) 171–181.
- [6] A.R. Wu et al., Nat. Methods 11 (2014) 41-46
- [7] H.C. Fan, G.K. Fu, S.P. Fodor, Science 347 (2015) 1258367.
- [8] P. Liberali, B. Snijder, L. Pelkmans, Nat. Rev. Genet. 16 (2015) 18–32.
- [9] N. Crosetto, M. Bienko, A. van Oudenaarden, Nat. Rev. Genet. 16 (2015) 57–66.
- [10] J.R. Sinnamon, K. Czaplinski, Methods Mol. Biol. 1206 (2015) 137–148.
- [11] F. Wang et al., J. Mol. Diagn. JMD 14 (2012) 22–29.
- [12] M.N. Cabili, Localization and abundance analysis of human lncRNAs at singlecell and single-molecule resolution, Genome Biol. (2015).
- [13] R. Ke et al., Nat. Methods 10 (2013) 857-860.
- [14] J.H. Lee et al., Science 343 (2014) 1360-1363.
- [15] J.H. Lee et al., Nat. Protoc. 10 (2015) 442-458.
- [16] C. Larsson, I. Grundberg, O. Soderberg, M. Nilsson, Nat. Methods 7 (2010) 395-397
- [17] A. Raj, P. van den Bogaard, S.A. Rifkin, A. van Oudenaarden, S. Tyagi, Nat. Methods 5 (2008) 877–879.
- [18] E. Shapiro, T. Biezuner, S. Linnarsson, Nat. Rev. Genet. 14 (2013) 618-630.
- [19] A.K. Shalek et al., Nature 510 (2014) 363-369.
- [20] Buchser, W. et al. in Assay Guidance Manual (eds G. S. Sittampalam et al.) (2004).
- [21] B. Snijder et al., Nature 461 (2009) 520-523.
- [22] B.K. Lundholt, K.M. Scudder, L. Pagliaro, J. Biomol. Screen. 8 (2003) 566–570.
- [23] P. Ramo, R. Sacher, B. Snijder, B. Begemann, L. Pelkmans, Bioinformatics 25 (2009) 3028–3030.
- [24] T.R. Jones et al., BMC Bioinformatics 9 (2008) 482.
- [25] L. Vincent, P. Soille, IEEE Trans. Pattern Anal. Mach. Intell. 13 (1991) 383–598.
- [26] N. Otsu, IEEE Trans. Syst. Man. Cybern. 9 (1979) 62–66.
- [27] E.W. Dijkstra, Numer. Math. 1 (1959) 269-271
- [28] A.E. Carpenter et al., Genome Biol. 7 (2006) R100.
- [29] S.A. Rifkin, Methods Mol. Biol. 772 (2011) 329–348.
- [30] T. Trcek et al., Nat. Protoc. 7 (2012) 408-419.
- [31] F. Mueller et al., Nat. Methods 10 (2013) 277-278.
- [32] J. Matas, O. Chum, M. Urban, T. Pajdla, Image Vis. Comput. 22 (2002) 761-767.
- [33] P. Ruusuvuori et al., BMC Bioinformatics 11 (2010) 248.
- [34] E. Bertin, S. Arnouts, Astron. Astrophys. Sup. 117 (1996) 393-404.
- [35] S. Mili, K. Moissoglu, I.G. Macara, Nature 453 (2008) 115–119.