

# BARcode DEmixing through Non-negative Spatial Regression (BarDensr)

Shuonan Chen<sup>3,4,\*</sup>, Jackson Loper<sup>2,3</sup>, Xiaoyin Chen<sup>5</sup>, Tony Zador<sup>5</sup>, and Liam Paninski<sup>1,2,3</sup>

<sup>1</sup>Mortimer B. Zuckerman Mind Brain Behavior Institute, Columbia University

<sup>2</sup>Department of Statistics, Columbia University

<sup>3</sup>Center for Theoretical Neuroscience, Columbia University

<sup>4</sup>Department of Systems Biology, Columbia University

<sup>5</sup>Cold Spring Harbor Laboratory

\*Corresponding author.

August 11, 2020

## Abstract

Modern spatial transcriptomics methods can target thousands of different types of RNA transcripts in a single slice of tissue. Many biological applications demand a high spatial density of transcripts relative to the imaging resolution, leading to partial mixing of transcript colonies in many pixels; unfortunately, current analysis methods do not perform robustly in this highly-mixed setting. Here we develop a new analysis approach, *BARcode DEmixing through Non-negative Spatial Regression* (BarDensr): we start with a generative model of the physical process that leads to the observed image data and then apply sparse convex optimization methods to estimate the underlying (demixed) colony densities. We apply BarDensr to simulated and real data and find that it achieves state of the art signal recovery, particularly in densely-labeled regions or data with low spatial resolution. Finally, BarDensr is fast and parallelizable. We provide open-source code as well as an implementation for the ‘NeuroCAAS’ cloud platform.

## Author Summary

Spatial transcriptomics technologies allow us to simultaneously detect multiple molecular targets in the context of intact tissues. These experiments yield images that answer two questions: which kinds of molecules are present, and where are they located in the tissue? In many experiments (e.g., mapping RNA expression in fine neuronal processes), it is desirable to increase the signal density relative to the imaging resolution. This may lead to mixing of signals from multiple RNA molecules into single imaging pixels; thus we need to *demix* the signals from these images. Here we introduce BarDensr, a new computational method to perform this demixing. The method is based on a forward model of the imaging process, followed by a convex optimization approach to approximately ‘invert’ mixing induced during imaging. This new approach leads to significantly improved performance in demixing imaging data with dense expression and/or low spatial resolution.

# 1 Introduction

Understanding the spatial context of gene expression in intact tissue can facilitate our understanding of cell identities and cellular interactions. How do neighboring cells' gene expressions relate to each other? How are different cell types with different activity patterns positioned in relation to each other? Is the subcellular distribution of gene expression informative about cell type or state? Multiplexed spatial transcriptomics methods offer a promising path forward to investigate these questions, allowing us to spatially resolve gene expression patterns. These assays can measure thousands of different genes simultaneously by looking at the same slice of tissue multiple times through multiple rounds of imaging. Using small barcoded sequences ('probes') which bind to target transcripts and amplify (generating easily detectable 'rolonies'), we can get exponentially more information about the nature of the tissue in each imaging round.

However, fully exploiting this new data type can be challenging, for many reasons. Insufficient optical resolution can cause parts of multiple rolonies to appear in the same imaging voxel, resulting in a 'mixed' signal (Chen et al., 2015; Alon et al., 2020). Tissue can deform or drift over multiple rounds of imaging (Qian et al., 2020), and the signal from individual rolonies can vary slightly between imaging rounds (Moffitt et al., 2016). The chemical washes may fail to complete their work in a given round, such that the imaging in the next round contains residual signal from the previous round (leading to a 'ghosting' effect). Some rolonies may entirely fail to bind to any probes in a given round (Lubeck et al., 2014; Chen et al., 2015). Most of these problems are rare, but they combine to yield a complex relationship between the signal of interest and the observed data.

Traditional techniques for extracting meaning from these images rely on good image preprocessing and clever heuristics; there are two main approaches that we are aware of. Both work well in ideal conditions. One school of thought ('blobs-first') begins by trying to identify regions in the tissue where a rolony may be present, and then tries to use the imaging data to guess the barcode identity of each rolony (Shah et al., 2016; Wang et al., 2018; Qian et al., 2020; Gyllborg et al., 2020; Alon et al., 2020). Another school of thought ('barcodes-first') begins by looking at each voxel and trying to determine whether the fluorescence signal emitted in that voxel over all the rounds is consistent with one of the barcodes (Lee et al., 2014; Moffitt et al., 2016, 2018). These two approaches are implemented in e.g. the 'starfish' (<https://github.com/spacetx/starfish>) package (under the names of 'spot-based' and 'pixel-based' approaches, respectively).

Both of these general approaches face difficulties whenever different rolonies make contributions to the same voxel. This can happen in regions of high expression density, and/or insufficient optical resolution. In many cases it is desirable to maximize the signal density, to increase the number of transcripts detected per cell and therefore the power of any downstream statistical analyses — while conversely, for practical reasons, we would like to minimize imaging time and file size, encouraging lower imaging resolution. To correctly identify rolony positions and identities in images with overlap, it is then necessary to perform some kind of 'demixing.' Because of this challenge, many current methods simply discard any blobs in regions where strong mixing occurs (Chen et al., 2015; Wang et al., 2018; Gyllborg et al., 2020).

To overcome this challenge, we sought to address the multiplexing problem directly. *BARcode DEMixing through Non-negative Spatial Regression* (BarDensr) is a new approach for detecting and demixing rolonies. This approach directly models the physical process which gives rise to the observations (Figure 1), including background-noise components, color-mixing, the point-spread function of the optics, and several other features. By directly modeling these physical processes, we are able to accurately estimate overall transcript expression levels — even when the transcript density is so high that it is very difficult to isolate and decode the identity of individual rolonies.

We provide a Python package for implementing these methods on either CPU or GPU architectures (<https://github.com/jacksonloper/bardensr>). The method requires about two minutes of compute time on a p2.xlarge Amazon GPU instance to process a seven-round, four-channels 1000 × 1000-pixel field of view from an experiment targeting 79 different transcripts. We also provide an implementation for the NeuroCAAS web-service (Abe et al., 2020), which can be used in a drag-and-drop fashion, with no installation required. We compared this method with three alter-

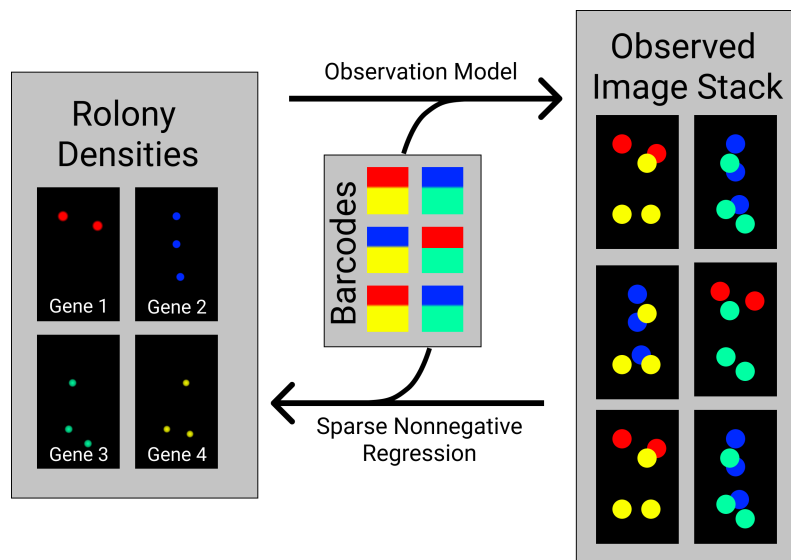


Figure 1: **BarDensr uses non-negative regression to demix and deconvolve the observed image stack, yielding a sparse intensity image for each barcode.** To find rolonies, we write down an observation model. We then use sparse non-negative regression to invert the observation model, yielding demixed and deconvolved intensities.

54 natives: the spot-based method of *starfish*; another ‘blobs-first’ approach (Single Round Matching,  
 55 or SRM, based on methods from (Wang et al., 2018; Qian et al., 2020)); and a ‘barcodes-first’  
 56 approach (Correlation approach, or ‘corr,’ based on (Lee et al., 2014; Moffitt et al., 2016, 2018)).  
 57 Both in simulation and real data, BarDensr improves on the state of the art in demixing accuracy.

## 58 2 Methods

### 59 Data

60 The experimental images were obtained using an improved version of BARseq (Chen et al., 2019)  
 61 to detect 79 endogenous mRNAs in the mouse primary visual cortex. The Cold Spring Harbor  
 62 Laboratory Animal Care and Use Committee approved all animal procedures and experiments.  
 63 Gene identities were read out using a seven-nucleotide gene identification index (GII), which were  
 64 designed with a minimal hamming distance of three nucleotides between each pair of GIIs.

65 Rolonies were prepared as described by (Sun et al., 2020). Imaging was performed on an  
 66 Olympus IX81 inverted scope with a Crest Xlight2 spinning disk confocal, a Photometrics BSI  
 67 Prime camera, and an 89 North LDI 7-line laser source. All images were acquired using an  
 68 Olympus UPLFLN 40x 0.75 NA objective. The microscope was controlled by micro-manager  
 69 (Edelstein et al., 2014).

70 See Appendix A for the preprocessing steps for this data, and Appendix E for the process of  
 71 generating the simulation data.

### 72 Notation and Observation Model

73 Formally speaking, what is the result of a spatial transcriptomics imaging experiment? For each  
 74 voxel ( $m$ ) in the tissue, at each imaging round ( $r$ ), in each color-channel ( $c$ ), we record a fluores-  
 75 cence intensity. We will use  $\mathbf{X}_{m,r,c}$  to denote this fluorescence intensity. Our task is to use  $\mathbf{X}$  to

76 uncover the presence and identity of rolonies in the tissue.<sup>1</sup> Below we describe the parameters  
77 used to model the physical process that yields these intensities:

78 **The rolonies,  $\mathbf{F}$ .** The transcripts in the tissue are amplified in place into a ‘rolony’ structure  
79 which is easy for fluorophores to bind to (Shah et al., 2016). Each voxel  $m$  may contain a  
80 different amount of roloniy material, and hence a varying level of fluorescence signal. We refer  
81 to the amount of material in voxel  $m$  for rolonies associated with barcode  $j$  as the **rolony**  
82 **density**. We denote this density by  $\mathbf{F}_{m,j}$ . The variable  $\mathbf{F}$  indicates where rolonies are  
83 and how bright we should expect them to be. This density should always be non-negative.  
84 Note that  $\mathbf{F}$  cannot be observed directly – instead, we observe fluorescence signal in different  
85 rounds and channels, and must use these signal observations to estimate the roloniy densities.

86 **The codebook,  $\mathbf{B}$ .** In each imaging round  $r$ , the rolonies associated with gene  $j$  will bind to  
87 specific fluorescently labeled detection probes. We use the binary variable  $\mathbf{B}$  to indicate  
88 which imaging rounds and fluorescent probes each gene is associated with. Specifically, we  
89 let  $\mathbf{B}_{r,c,j} = 1$  whenever a roloniy with barcode  $j$  should bind to a fluorescent probe associated  
90 with specific color-channel  $c$  in imaging round  $r$  (and 0 otherwise). Here we assume  $\mathbf{B}$  is  
91 known. The vector of values of  $\mathbf{B}$  for a particular gene  $j$  is known as the ‘barcode’ for that  
92 gene, and the collection  $\mathbf{B}$  of all the barcodes is known as the ‘codebook.’

93 **The probe response functions,  $\mathbf{K}, \varphi$ .** If a probe centered at a particular voxel is illuminated  
94 with a particular wavelength, the probe will emit a certain amount of signal which we can  
95 record at the corresponding voxel. We may also observe dimmer responses at neighboring  
96 voxels, due to the possible spreading of the single point object in the optical system. We use  
97 a non-negative matrix  $\mathbf{K}$  to denote the *point-spread function*, i.e., the typical fluorescence  
98 signal-levels produced at each voxels in the neighborhood of a probe. We use the matrix  $\varphi$   
99 to represent the responsiveness of each type of fluorescent probe to each wavelength; each  
100 element of this matrix lies in the range of  $[0, 1]$ . Here we assume that the number of types of  
101 fluorescent probes is the same as the number of color-channels measured (though this could  
102 be relaxed). We further assume that the voxel-resolution of the roloniy density is the same  
103 as the voxel-resolution of the original images.

104 **Phasing,  $\rho$ .** A washing process is applied after each round of imaging. However, in practice this  
105 washing step may not completely remove all of the reagents from every voxel. This can  
106 result in a ‘ghost’ of one round appearing in the next rounds. For each color-channel  $c$ , we  
107 let  $\rho_c \in [0, 1]$  indicate the fraction of activity which appears as a ‘ghost’ signal in the next  
108 round.

109 **Background,  $a$ .** The images we obtain may also include background fluorescence from the tissue.  
110 We assume that the background is constant across rounds. We model this effect using a non-  
111 negative per-voxel value  $a_m$  for each voxel  $m$ .

112 **Per-round per-wavelength gain,  $\alpha$ , and baseline,  $b$ .** The brightness observed from all rolonies  
113 at a particular color-channel in a particular round may have an associated gain factor. We  
114 model this gain factor with a non-negative per-round ( $r$ ) per-channel ( $c$ ) multiplier  $\alpha_{r,c}$  and  
115 non-negative intercept  $b_{r,c}$ .

Putting all these pieces together, we obtain an *observation model*. This model states that the  
observed brightnesses  $\mathbf{X}_{m,r,c}$  should be given by the formulae

$$\mathbf{X}_{m,r,c} \approx a_m + b_{r,c} + \alpha_{r,c} \sum_{j,m',c'}^{J,M,C} \mathbf{K}_{m,m'} \mathbf{F}_{m',j} \varphi_{c,c'} \mathbf{Z}_{r,c',j},$$
$$\mathbf{Z}_{r,c,j} = \rho_c \mathbf{Z}_{r-1,c,j} + \mathbf{B}_{r,c,j}.$$

<sup>1</sup>Throughout we assume that  $\mathbf{X}$  is preprocessed, including background removal and image registration (see Appendix A for more detail), hence that there are no systematic shifts of the image between imaging rounds.

116 Here the variable  $\mathbf{Z}$  is used to incorporate the round-phasing effects; i.e.,  $\mathbf{Z}_{r,c,j}$  measures the  
 117 concentration of probes of type  $c$  which we would expect at round  $r$ , arising from a rolonny with  
 118 barcode  $j$ . We will also find it convenient to define

$$\mathbf{G}_{r,c,j} = \alpha_{r,c} \sum_{c'} \varphi_{c,c'} \mathbf{Z}_{r,c',j}.$$

119 This represents the total contribution of fluorescence signal expected to arise in round  $r$  and  
 120 channel  $c$  from a rolonny of type  $j$ . A summary of notation can be found in Table 1.

121 Overall, the model introduced above could certainly be expanded to model the physical imaging  
 122 process more accurately, but we found that it was sufficient for our purposes: detecting and  
 123 demixing rolonies.

## 124 Inference

125 Our task is to uncover the positions and barcodes of rolonies in the tissue. According to the  
 126 model in the previous section, this information can be obtained from the rolonny density variable,  
 127  $\mathbf{F}$ . However,  $\mathbf{F}$  cannot be directly measured; thus our primary task is to estimate  $\mathbf{F}$  from the  
 128 original image data. To do this we must in a sense invert the observation model specified above:  
 129 the observation model tells us how rolonny densities gives rise to the fluorescence signal, but we  
 130 would like to use observations of the fluorescence signal to estimate the rolonny densities.

### 131 Using the observation model to estimate the rolonny densities $\mathbf{F}$

132 We use a sparse non-negative regression framework to estimate the unknown parameters. In this  
 133 estimation we are guided by three ideas:

- 134 • We believe our observation model is *approximately* correct. We formalize this by saying that  
 135 we believe our squared ‘reconstruction loss’ can be made small. We define this loss by

$$L_{\text{reconstruction}} = \sum_{m,r,c} \left( \mathbf{X}_{m,r,c} - \left( a_m + b_{r,c} + \alpha_{r,c} \sum_{j,m',c'}^{J,M,C} \mathbf{K}_{m,m'} \mathbf{F}_{m',j} \varphi_{c,c'} \mathbf{Z}_{r,c',j} \right) \right)^2.$$

|              | Description                           | Dimensions            | Support        |
|--------------|---------------------------------------|-----------------------|----------------|
| $M$          | number of voxels                      | scalar                | $\mathbb{N}$   |
| $C$          | number of types of probes/wavelengths | scalar                | $\mathbb{N}$   |
| $R$          | number of rounds                      | scalar                | $\mathbb{N}$   |
| $J$          | number of barcodes                    | scalar                | $\mathbb{N}$   |
| $\mathbf{X}$ | observed imaging intensities          | $M \times R \times C$ | $\mathbb{R}^+$ |
| $\mathbf{F}$ | spatial rolonny density               | $M \times J$          | $\mathbb{R}^+$ |
| $\mathbf{B}$ | (known) binary codebook matrix        | $R \times C \times J$ | $\{0, 1\}$     |
| $\rho$       | per-channel phasing factor            | $C$                   | $[0, 1]$       |
| $\mathbf{Z}$ | phased barcodes                       | $R \times C \times J$ | $\mathbb{R}^+$ |
| $\alpha$     | per-round per-channel scale factor    | $R \times C$          | $\mathbb{R}^+$ |
| $\mathbf{G}$ | scaled color-mixed phased barcodes    | $R \times C \times J$ | $\mathbb{R}^+$ |
| $b$          | per-round per-channel offset          | $R \times C$          | $\mathbb{R}^+$ |
| $a$          | per-pixel baseline intercept term     | $M$                   | $\mathbb{R}^+$ |
| $\mathbf{K}$ | spatial point-spread function         | $M \times M$          | $\mathbb{R}^+$ |
| $\varphi$    | probe wavelength-response matrix      | $C \times C$          | $[0, 1]$       |
| $\omega$     | tolerated reconstruction error        | scalar                | $\mathbb{R}^+$ |

Table 1: Notation

- 136 • We believe that all of our parameters are non-negative. For example, we do not believe it is  
137 possible to have *negative* densities for rolonies at a particular voxel. Likewise, we expect the  
138 per-round per-channel scaling factors ( $\alpha$ ) and probe-response terms ( $\varphi$ ) to be non-negative.
- 139 • We believe that the rolony densities,  $\mathbf{F}$ , are sparse: many voxels will not contain any rolony  
140 at all. Ideally we would formalize this idea by putting a penalty on the number of voxels  
141 with nonzero rolony amplification. However, this penalty is difficult to optimize in practice.  
142 Instead, following a long history of work in sparse estimation theory ([Hastie et al., 2015](#)),  
143 we enforce this sparsity by placing a linear penalty on the total summed density. We define  
144 this penalty by

$$L_{\text{sparsity}} = \sum_{m,r,c,c',j} \alpha_{r,c} \mathbf{F}_{m,j} \varphi_{c,c'} \mathbf{Z}_{r,c',j}.$$

145 (Note that for a general sparse estimation problem, this penalty would be defined using a  
146 summed absolute value term; however, in our case all parameters are already constrained to  
147 be non-negative, so this is not necessary.)

148 Together, these three ideas suggest constrained optimization as a natural way to estimate our  
149 parameters. We will seek the non-negative parameters that give the smallest possible value of  
150  $L_{\text{sparsity}}$ , subject to the constraint that  $L_{\text{reconstruction}}$  falls below a noise threshold  $\omega$ . We provide  
151 an automatic way to select this noise threshold (see Appendix I), as well as an interactive process  
152 for the user to select this threshold so that the reconstruction loss appears satisfactory.

153 Assuming that  $\mathbf{B}, \mathbf{K}$  are known, this constrained optimization problem can be written as:

$$\begin{aligned} \min_{\mathbf{F}, \rho, \alpha, b, a, \varphi \geq 0} \quad & L_{\text{sparsity}}, \\ \text{subject to} \quad & L_{\text{reconstruction}} \leq \omega. \end{aligned} \tag{1}$$

154 To solve this optimization problem, we use a projected gradient descent approach. The linear  
155 structure of the problem makes it possible to pick all learning rates automatically; for example,  
156 the resulting algorithm reaches convergence for a single  $1000 \times 1000$  field of view (with a total of 28  
157 images, with seven rounds and four color-channels) and 81 different barcodes (79 from the original  
158 experiment, and two additional unused barcodes as described below) in about two minutes on a  
159 `p2.xlarge` Amazon GPU instance. Details can be found in Appendix I.

Before concluding this section, we will address an issue of what is known as ‘identifiability.’  
Let us say we have learned a model via our inference method, i.e. we have learned  $\mathbf{F}, \rho, \alpha, b, a, \varphi$ .  
Now let us consider a new model,  $\mathbf{F}', \rho', \alpha', b', a', \varphi'$ , such that

$$\begin{aligned} \mathbf{F}' &= 4\mathbf{F} & \rho' &= \rho \\ \alpha' &= \alpha/2 & b' &= b \\ a' &= a & \varphi' &= \varphi/2. \end{aligned}$$

160 Under this new model, the reconstruction loss is the same and the sparsity loss is the same.  
161 As far as our inference method is concerned, the two models are identical. It follows that our  
162 inference procedure simply cannot hope to learn overall scaling factors of this kind. Thus, any  
163 learned parameters should be understood as being known up to overall scale factors. To resolve  
164 this ambiguity we normalize  $\alpha$  by dividing by its sum (recall that  $\alpha$  is non-negative, so this sum  
165 will be positive) and multiply  $\mathbf{F}$  by the same factor. Similarly, we divide each row of  $\varphi$  by its  
166 diagonal value and multiply the corresponding column of  $\alpha$  by the same value.

## 167 Finding rolonies

168 Let us now assume we have used the non-negative regression framework to estimate  $\mathbf{F}$  (the collec-  
169 tion of rolony density images, one for each barcode). These per-barcode density images indicate  
170 the positions of rolonies that belong to a particular barcode; see the left side of Figure 1 for a






Figure 2: **Rolony densities make it easier to detect rolonies.** The left plot shows the max-projection of the original image across all rounds and channels; detecting blob-like structures in this image can be challenging, especially when two rolonies are in close proximity. By contrast, the rolony densities for particular genes are sparser, so it is easier to identify the positions of individual rolonies in the tissue. The middle and right plots show examples of these rolony densities (in fact, to make it possible to even visualize the results, we here show the rolonies densities *after* applying the point-spread function  $\mathbf{K}$ ; the true rolony densities  $\mathbf{F}$  are even sparser). The orange marks represent rolonies detected by a hand-curated approach. Note that the rolony densities appear to show several rolonies which were not detected by the hand-curated approach (in particular, we see several bright spots with no orange marks). In Figure 27 and 28, we show that these additional rolonies do indeed seem to be valid, and were simply missed by the original hand-curated approach.

171 schematic. We can then apply a blob-finding algorithm to these per-barcode images to find the  
172 rolonies for each barcode; in practice we simply find local maxima in the per-barcode images.

173 Finding rolonies, or ‘blobs,’ in the per-barcode images is easier than finding blobs in the original  
174 images. See Figure 2 as an example. The per-barcode images include fewer blobs and the blobs  
175 are smaller, so there are fewer problems with overlapping blobs. More specifically:

- 176 • There are fewer blobs in each rolony density than in the original image stack. In the observed  
177 images, the intensity measured for each voxel for each wavelength at each round is a sum of  
178 contributions from all nearby rolonies which emit signal at that wavelength in that round.  
179 By contrast, the intensity measured at a particular voxel in the per-barcode images is only  
180 the sum of contributions from rolonies with that one specific barcode.
- 181 • The blobs are smaller in the rolony density than in the original image stack. In the observed  
182 images, the intensity at a voxel is a contribution from all rolonies which are within the radius  
183 of the point-spread function  $\mathbf{K}$ . Recall that this function smears signal from a single voxel  
184 across all nearby voxels. By contrast, the intensity of a per-barcode image at a particular  
185 voxel represents the amplification level of rolonies in that one voxel. In this sense, the  
186 inference process attempts to invert the point-spread function (i.e., perform deconvolution).  
187 On its own, this inversion process would not be numerically stable; however, the sparsity  
188 penalty and non-negativity constraint ensures it is numerically well-behaved (Hastie et al.,  
189 2015).

190 The spatial rolony variable  $\mathbf{F}$  thus represents a *demixed and deconvolved* version of the raw  
191 data. The original data is mixed, insofar as each raw intensity represents contributions from many  
192 barcodes. It is also convolved, insofar as each raw intensity represents contributions from many  
193 positions in space via the point-spread function. The non-negative sparse regression allows us to  
194 simultaneously demix and deconvolve, yielding per-barcode images which are cleaner and easier  
195 to understand.

196 Although it is easier to find blobs in the rolony densities, there is still one obstacle to be  
197 overcome: the threshold. Any blob-finding algorithm must specify an intensity above which a blob

198 is considered real. How can this threshold be chosen? Here we make use of ‘unused barcodes.’  
199 There could be as many as  $C^R$  unique barcodes in a codebook for an experiment with  $R$  rounds and  
200  $C$  channels of measurement (assuming only one channel emits signal in each round, which is the  
201 case in the experiments we studied). However, most of these barcodes are not used in the actual  
202 experiment. These unused barcodes give us a way to pick a sensible threshold. Along with the real  
203 codebook, we additionally include several unused barcodes; we enumerated all possible barcodes  
204 such that each round contained exactly one active channel, then selected uniformly at random  
205 from the set of barcodes such that each barcode differed from every other barcode in at least  
206 three rounds. We then run BarDensr on this augmented codebook. Blobs in the rolonity densities  
207 associated with the unused barcodes must correspond to noise, since the true data-generating  
208 process did not include any signal from such barcodes. We therefore set the threshold to be the  
209 smallest value which guarantees that no blobs were detected in the unused barcodes. (In practice,  
210 using just two unused barcodes sufficed to estimate a stable and accurate threshold.)

### 211 **Accelerating computation**

212 The time required to apply BarDensr scales roughly linearly with the number of voxels in the data.  
213 There are several approaches the BarDensr package uses to relieve the computational burdens of  
214 working with large datasets:

- 215 1. **Exploiting barcode sparsity.** In any given patch of the data, many of the barcodes may  
216 not appear at all. If we can use a cheap method to detect genes which are completely missing  
217 from a given patch, we can then remove these genes from consideration in that patch, yielding  
218 faster operations. We call this ‘sparsifying’ the barcodes.
- 219 2. **Coarse-to-fine.** As we will see below, BarDensr is effective even when the data has low  
220 resolution. This suggests a simple way to accelerate computation: downsample the data, run  
221 BarDensr on the downsampled data (which will have fewer voxels), and then use the result  
222 to initialize the original fine-scale problem. If this initialization is good, fewer iterations of  
223 the optimization will be necessary to complete the algorithm.
- 224 3. **Parallelization.** BarDensr can use multiple CPU cores or GPUs (when available) to speed  
225 up parallel aspects of the optimization (e.g., processing data in spatial patches).

226 Details on these methods (which can be used in combination with each other) can be found in  
227 Appendix [H](#).

### 228 **Code availability**

229 The BarDensr Python package is available from <https://github.com/jacksonloper/bardensr>. The  
230 NeuroCAAS implementation of BarDensr can be found at <http://www.neurocaas.com/analysis/8>.  
231 This NeuroCAAS implementation requires no software or hardware installation by the user. The  
232 BarDensr NeuroCAAS app has a simple input-output model. As input, the user must upload a  
233 stack of images, a codebook, and a configuration file specifying parameters such as the radius of  
234 the smallest rolonities of interest (see the NeuroCAAS link above for further details regarding the  
235 data format.) We assume that the images have been registered and background-subtracted before  
236 input into NeuroCAAS. There are two outputs from BarDensr NeuroCAAS implementation. The  
237 first output takes the form of a comma-separated-value file listing all entries in the rolonity density  
238  $\mathbf{F}$  which have signal greater than zero. The second output is a structured HDF5 file, which stores  
239 the results of singular value composition (SVD) on the cleaned images for each spot detected; this  
240 helps the user assess the quality of the spots detected by the algorithm (see the next section as  
241 well as Figures [9](#) - [10](#) for detail). See the NeuroCAAS link provided above for full details. Also  
242 see Appendix [B](#) for further details on the AWS hardware selected here.



### 243 3 Results

#### 244 The rolony densities estimated by BarDensr provide sparse, single images 245 to detect spots for individual barcodes.

246 As emphasized in Section 2, the sparse non-negative regression approach aims to yield per-gene  
247 rolony density images which are easy to work with. The cartoon in Figure 1 may help illustrate  
248 this idea. Our belief is that the true per-gene rolony densities will be sparse images, so the learned  
249 rolony densities should also be sparse images.

250 To test this belief, we applied BarDensr to the experimental data described in Section 2. Figure  
251 2 compares the raw data with the learned rolony densities for *Nrgn* and *Slc17a7* in a small region  
252 of the tissue. As hoped, the spatial rolony densities are indeed quite sparse compared to the raw  
253 data. This ensures that blob-detection is relatively easy. This figure also shows that many of the  
254 bright spots in the rolony density images appear next to rolony locations found by a hand-curated  
255 method (see Appendix C for details). For visualization purposes, this figure shows the blurred  
256 version of the spatial rolony densities (i.e. **KF**); these make it easier to see the bright spots.

257 To get a sense for what all the different genes look like, we also examined the rolony densities  
258 for all the barcodes (81 in total in this dataset, including two unused barcodes); see Figures 11  
259 and 12. These sparse images enable us to identify the rolony location easily for each barcode.

#### 260 BarDensr provides improved demixing and detection accuracy compared 261 to existing approaches.

262 To benchmark BarDensr against other methods, we generated simulated data with rolony density,  
263 gene expression levels, and noise levels matched to the experimental data, as shown in Figure  
264 3, and then examined how well we could recover the ‘true’ rolonies from the simulated data.  
265 Qualitative results for several different genes are shown in Figures 20 - 23. Quantitatively, we  
266 present a Receiver Operating Characteristic curve (ROC curve) in Figure 4, which summarizes  
267 the percentage of true detected rolonies (also known as ‘1-FNR’, the complement of the False  
268 Negative Rate (FNR)). Depending on the False Positive Rate (FPR) we are willing to tolerate,  
269 different detection rates can be achieved; the ROC curve summarizes this relationship.

270 We compare BarDensr to several other approaches. *Starfish* is one package developed for  
271 analyzing spatial transcriptomics data. This method has many hyperparameters. To give this  
272 method its best chance, we first tried to find the best parameters manually, and additionally used  
273 the *BayesianOptimization* package (Nogueira, 2014) to find the hyperparameters which allowed  
274 it to perform as well as possible on the simulated data. Figure 4 shows that this performance  
275 falls short of the detection rates achieved by BarDensr. We also investigated SRM (see Appendix  
276 C) and a correlation-based method (‘corr’, see Appendix D) for comparison. These two methods  
277 represent ‘blobs-first’ and ‘barcodes-first’ approaches. BarDensr has better recovery prediction  
278 than either of these.

279 Our simulated data here do not capture the full biological content of the real observed data.  
280 For example, in real data, the tissue often has some regions with dense rolony concentrations  
281 (e.g. nuclei) and other regions which are more sparse. In order to quantify performance in more  
282 realistic biological contexts, we performed a ‘hybrid’ simulation, a la (Pachitariu et al., 2016).  
283 We started with the original experimental data and injected varying numbers of spots at random  
284 locations in the image with varying peak intensities (cf. Appendix E). To test if the model is able  
285 to recover these injected spots with the original image background, we computed the FNR (FPR  
286 could not be computed here since we do not know the ground truth in the original experimental  
287 data). We ran two variants on this simulation: one ordinary simulation and one simulation with  
288 ‘dropout,’ in which some rolonies emit a strong bright signal in most of the rounds but simply  
289 vanishes in one or more rounds (see Appendix E). The results of the dropout and non-dropout  
290 experiments are shown in Figure 5. As expected, the performance decreases when the intensity of  
291 the injected spots is smaller. However, as long as the intensity of injected spots was at least half  
292 the maximum intensity of the original image, BarDensr was able to find all the spots, even in the

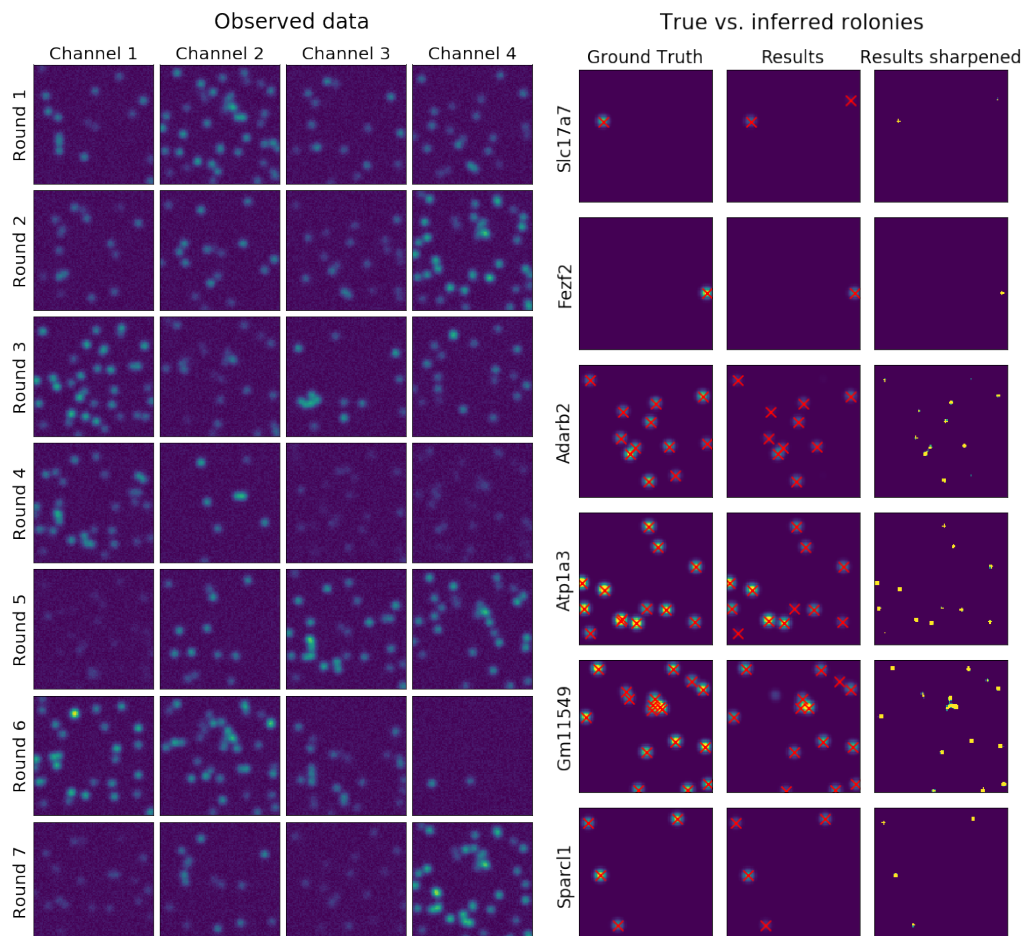


Figure 3: **BarDensr accurately recovers the ground truth in simulated data.** The left plot shows the simulated data in all rounds and channels. Our simulator uses the same barcodes as the true data, the same distribution of gene prevalence found by a hand-curated approach, a similar point-spread function and per-round-and-channel scaling  $\alpha$ , and a similar density of spots and noise level. In the right plot, we applied BarDensr to this simulated data, and found that we were able to largely recover the true rolonies in this simulation. The first column of plots shows the true positions of rolonies which were used to generate the simulated data. Each plot shows rolonies for a particular gene. The final column of plots shows the rolony densities learned by BarDensr. The middle column of plots shows a blurred version of the rolony densities (which are a bit easier to see) and the spots discovered from these rolony densities. The algorithm accurately recovers most of the simulated ground truth rolonies, with a few mistakes. In some cases, multiple rolonies of exactly the same barcode lie right next to each other, but the algorithm identified a single large rolony instead of several small rolonies. There are also rare false positives (where we detect a spot that did not exist in the ground truth) and false negatives (where we failed to detect a spot that did exist in the ground truth).

293 simulation with dropout; by contrast, the SRM approach was unable to find all the injected spots  
294 in the hybrid experiment, especially in the dropout variant.




Figure 4: **BarDensr discovers more correct rolonies in simulated data.** What percentage of rolonies are correctly detected? We use the Receiver Operating Characteristic curve (ROC curve) to look at this percentage (the complement of False Negative Rates, or 1-FNR) as a function of the tolerated False Positive Rate (FPR), for BarDensr (red), *starfish* (orange), Single Round Matching (SRM, green), as well as the correlation-based method (‘corr’, gray); cf. Appendix C and D for details on these other methods. Figure 20 and 21 (for the simulation with sparser spots, top plots of this figure), as well as 22 and 23 (for the simulation with denser spots, bottom plots of this figure) illustrate these simulation data. In drawing these curves, we consider two qualitatively different kinds of errors: errors because a roloniy isn’t detected at all, and errors because a roloniy is detected but it is assigned the wrong barcode. The dotted lines reflect ROC for the former, the solid lines reflect ROC for the latter. The left plots show these curves for simulated data. The right plots show these curves for simulated data with ‘dropout’ – a form of noise present in some spatial transcriptomic methods (cf. Appendix E for details). For all four kinds of simulations, we found BarDensr is able to find significantly more spots.

295 **Errors are mostly mis-identification on the barcodes, not missed detec-**  
 296 **tions.**

297 We used simulated data to investigate the failures represented by the FPR and FNR described  
 298 above: are they caused by failure in assigning the rolonies to the correct barcodes (‘barcode mis-

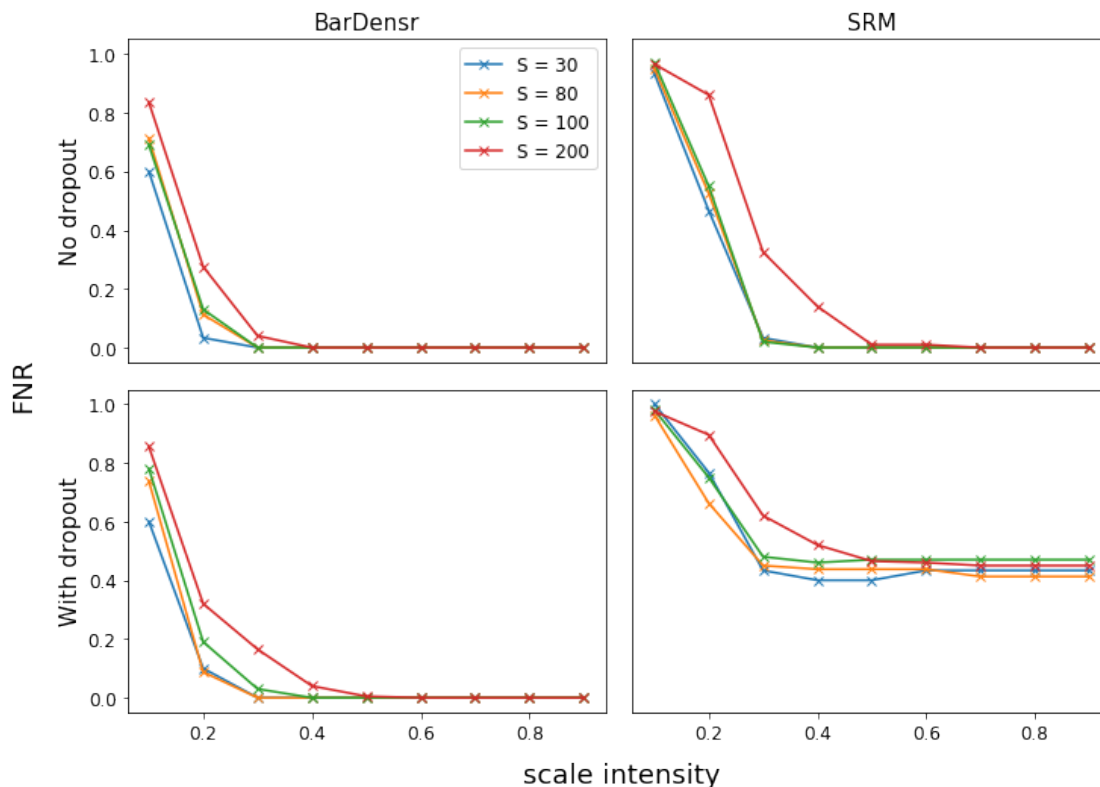


Figure 5: **Hybrid simulation.** Showing the False Negative Rates (FNR, y-axis) as the function of scale intensity (x-axis) and spot number ( $S$ , colored lines), without (left) and with (right) dropout, using BarDensr (top) and SRM (bottom). Scale intensity indicates the intensity of the injected spots in the simulation, relative to the maximum intensity in each frame in the original data. See Appendix E for detail.

299 identification'), or failure in detecting rolonies? To find out, we computed how the failure rates  
 300 would change if mis-identified barcodes were not considered 'errors.' We denote this the 'total hit  
 301 rate' analysis (cf. Figure 4, dotted lines); both BarDensr and SRM have very high total hit rates  
 302 for the simulated data examined here, indicating that both of these methods detect spots well,  
 303 but sometimes mis-classify the spot identity. See Appendix F for further details.

### 304 BarDensr remains effective on data with low spatial resolution.

305 High-resolution imaging can be expensive and time-consuming. BarDensr can also work on low-  
 306 resolution images. To show this, we spatially downsampled the experimental images for each  
 307 frame (each round and each channel). We then fit BarDensr to these lower-resolution images. An  
 308 example is shown in Figure 6 (additional examples with  $5\times$  and  $10\times$  lower resolutions can be seen  
 309 in Figure 24). These figures show that BarDensr correctly detects the overall expression levels of  
 310 each gene in low-resolution images – even when the downsampling is so extreme that picking out  
 311 individual rolonies is not feasible.

312 To test if BarDensr can recover the correct gene expression level when applied on the low-  
 313 resolution data, we also quantified the cell-level gene activity on a larger region where 43 cells are  
 314 detected using a seeded watershed algorithm (see Appendix G for detail). The bottom plots of  
 315 Figure 6 suggest that with  $5\times$  downsampled data, the cell-level gene expression, as well as the cell  
 316 clusters, are preserved with high consistency compared to the results of applying the method to  
 317 the original fine scale.

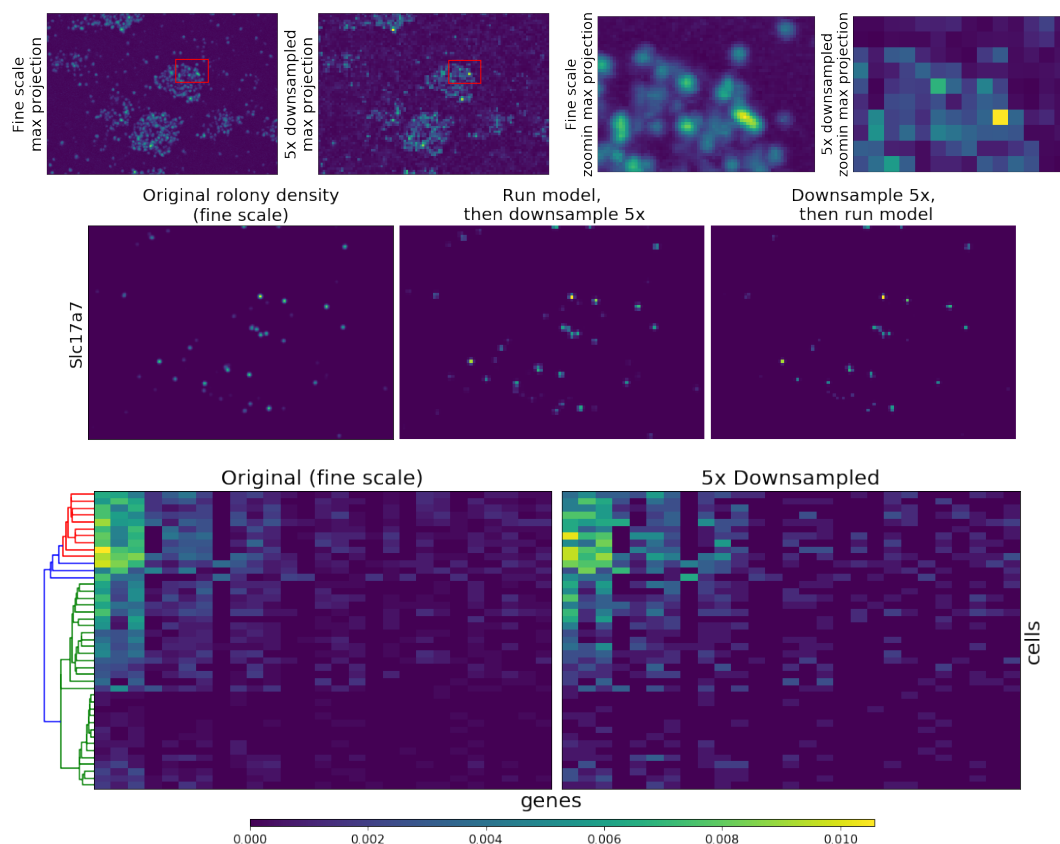


Figure 6: **BarDensr works on low-resolution data.** **Top row.** The  $5\times$  downsampled image is compared to the original ‘fine scale’ image. All these plots show the max-projection across all rounds and channels, with the right two showing the zoomed region indicated by the red rectangles in the left two. Note that it is difficult to visually isolate single spots from the downsampled image. To test the performance of BarDensr on this low-resolution data, we first run the model on the original data (i.e., top left), obtain rolyny densities, and then finally downsample the rolyny densities (‘run-then-downsample’). Next, we run BarDensr on downsampled data (i.e., the second plot on the top row) and examine the estimated rolyny densities (‘downsample-then-run’). **Middle row.** The rolyny densities for a selected gene (*Slc17a7*) estimated using the original fine scale (left), as well as these two approaches (middle for ‘run-then-downsample’ and right for ‘downsample-then-run’). For a more complete example, see Figure 24. **Bottom row.** The cell-level gene expression quantification, for those genes that have more than four spots in the fine scale in a  $1000 \times 1000$  region. The color of the heatmap indicates the proportion of gene counts (i.e., the total counts of each gene divided by the total counts of all genes detected in the region). The x-axis represents the 24 genes that were chosen, ordered based on the counts in the fine scale. The y-axis represents the cells, ordered based on the hierarchical clustering result from the fine scale, as shown in the dendrogram on the left. A total of 43 cells are segmented from the original image using a seeded watershed algorithm (cf. Appendix G). The two different results yield nearly identical clusterings, indicating that BarDensr recovers gene activity with accuracy sufficient to cluster cells even given low-resolution images.

318 **BarDensr computations can be scaled up to tens of thousands of barcodes**  
 319 **via sparsifying and coarsening accelerations.**

320 In Section 2, we described how the barcode sparsity could help us potentially apply the method to  
 321 a large dataset with more barcodes. To test if we can use a much larger dataset, we considered a




Figure 7: **Speed-up of BarDensr using coarse-to-fine method.** Area Under the ROC curve (AUROC) as a function of wall clock time. The red curve (‘coarse-to-fine’) is the result where we fit the model to the two-times-downsampled ‘coarse’ data, followed by running the model at the original fine scale using the parameters learned from the coarse scale as the initial conditions. The black curve (‘fine only’) indicates the result of running the model on the original fine scale for all the iterations. The total number of updates in the model is 20 and 10 for ‘coarse-to-fine’ (the gray line indicates the end of the 20 coarse updates), and is 15 for ‘fine only.’

322 simulated example with more unique barcodes (53,000 unique barcodes and 17 sequencing rounds).  
323 With so many barcodes, naively running BarDensr is prohibitively expensive (in both compute  
324 time and memory) on large datasets. However, we also expect such datasets are extremely sparse  
325 in terms of barcodes – any given small region of the image is quite unlikely to include rolonies  
326 from all 53,000 barcodes. This is particularly true when each barcode corresponds to a unique  
327 cell instead of a unique gene (Chen et al., 2019): a small region of tissue may contain many  
328 different transcripts, but it will only contain a small number of different cells. Thus we should  
329 be able to take advantage of this sparsity to speed up BarDensr. We simulated a  $50 \times 80$  small  
330 region where 40 rolonies were present in total. We then obtained a coarse, downsampled image,  
331 and then ran BarDensr and learned the parameters for this low-resolution data. If the learned  
332 parameters from the coarse scale indicated a particular barcode did not appear, then we assumed  
333 that this barcode should be absent even if we used the data at original resolution. The result in  
334 Figure 25 shows nearly perfect prediction performance. This problem was quite small, so we could  
335 also run the method without using any sparsity-based acceleration techniques; we found that the  
336 unaccelerated version did not outperform the accelerated version, suggesting that BarDensr can be  
337 used for datasets of this kind with larger number of molecular or cellular barcodes (cf. (Kebschull  
338 et al., 2016; Han et al., 2018; Chen et al., 2018, 2019)).

339 Finally, given a small number of barcodes, BarDensr can run without these acceleration tech-  
340 niques – but these accelerations are still worth applying, to help cut down on computation times  
341 and reduce memory usage. We found that these techniques reduced runtime by a factor of four  
342 (Appendix H). Figure 7 shows the speed-up of the BarDensr using ‘coarse-to-fine’ accelerations.  
343 Further, as shown in Figure 8, BarDensr performs well while taking advantage of the gene-sparsity  
344 for each small region after coarsening.

### 345 BarDensr recovers interpretable parameters.

346 BarDensr uses a data-driven approach to estimate all the relevant features of the physical model:  
347 the per-channel phasing factor, the per-round per-channel scale factor, the per-round per-channel



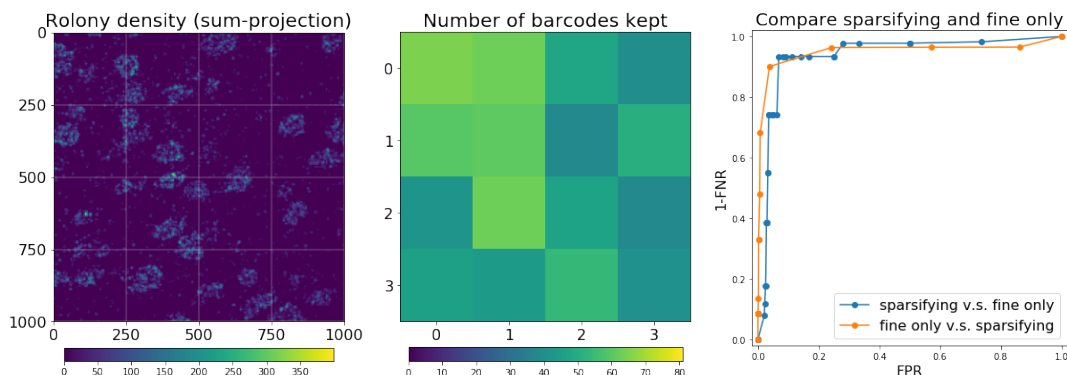


Figure 8: **BarDensr can take advantage of gene-sparsity.** Here we used two different approaches to analyze a  $1000 \times 1000$  region of the experimental data. The first approach uses BarDensr naively, applying it directly to the image. The second approach is illustrated on the left and middle plots. This approach accelerates the method using a ‘coarse-to-fine’ method by taking advantage of ‘gene-sparsity’, i.e., the fact that many barcodes do not appear in any given small region. Specifically, we split this region into  $4 \times 4$  patches (the borders of these patches are indicated as the white lines on the left plot). After the relatively fast ‘coarse’ step, the barcodes that have very low maximum rolony densities were removed before the following ‘fine’ step (cf. Appendix H for more detail). This keeps only a relatively small number of barcodes to consider for each patch (ranging from 38 to 65 out of 81 barcodes, as shown in the middle plot), therefore reducing the computation time and the memory usage for the ‘fine’ step later. Since we are here analyzing real experimental data, there is no ground truth we can use to compare the efficacy of the two methods. However, we here show that both methods yield nearly the same result, as shown in the ROC curves on the right plot. In particular, we can treat one method as the ‘truth’ and construct an ROC curve indicating the accuracy of the other method. We can then do the reverse, treating the other method as ‘truth.’

348 offset, the per-pixel background, the per-wavelength response matrix, and the spatial rolony den-  
 349 sities (the latter of which have already been described in detail above). In the data analyzed here,  
 350 we found that the per-channel phasing factor was relatively small, suggesting very little ‘ghosting’  
 351 in this data. The wavelength-response matrix was almost diagonal, although we found some slight  
 352 color-mixing from channel 2 to channel 1, consistent with visual inspection (see the fifth round in  
 353 Figure 16 as an example). This indicates that our model is able to correctly recover the color-  
 354 mixing effects. We also investigated whether all of the features of our model were necessary for the  
 355 purposes of finding rolonies. For each feature of the model, we tried removing that aspect of the  
 356 model and seeing whether the method still performed well. For the data analyzed here, we found  
 357 that the  $\varphi$  and  $\rho$  parameters were not essential (though they did seem to improve the performance,  
 358 at least qualitatively). By contrast, all of the other parameters were essential; removing any of  
 359 them yielded nonsensical results.

360 **BarDensr is able to capture the important signal based on the assessment**  
 361 **on the predicted signal intensities.**

362 Our algorithm is based upon a physical model of how this data is generated. Rolonies appear at  
 363 different positions in the tissue, they emit fluorescence signal in different conditions, the fluores-  
 364 cence signal is smeared by a point-spread function, and finally we observe this signal, together with  
 365 certain background signal and noise. As long as this model captures all the important features  
 366 of the physical process, observed intensities should match the predicted intensities at each voxel  
 367 in each round and in each channel. To think about this more clearly, let’s define these predicted

368 intensities as the ‘reconstruction’:

$$\text{reconstruction}_{m,r,c} \triangleq a_m + b_{r,c} + \alpha_{r,c} \sum_{j,m',c'}^{J,M,C} \mathbf{K}_{m,m'} \mathbf{F}_{m',j} \varphi_{c,c'} \mathbf{Z}_{r,c',j}.$$

369 To test our model, we can visually compare the reconstruction to the observed data. If the residual  
 370 between the two includes significant highly-structured noise, then it is likely that we are missing  
 371 important aspects of the data. Figures 13 - 18 show the results of these comparisons. They appear  
 372 fairly promising, but certain structured features do appear in the residual. Most strikingly, we  
 373 have found that a minority of rolonies ‘dropout’ for one or more rounds: a rolony may give a  
 374 strong bright signal in most of the rounds but simply vanish in one round. Our current physical  
 375 model does not accommodate this, and this limitation appears in the residual as bright and dark  
 376 spots. However, as mentioned above and shown in the hybrid simulation data in Figure 5, our  
 377 method is robust to these ‘dropout’ effects; it is still able to capture the correct rolony positions  
 378 when it occurs on a small number of rounds.

### 379 Diagnostics based on ‘cleaned’ images are useful to check the accuracy 380 of BarDensr.

381 The reconstruction is made up of many parts: it has the background component  $a$ , the per-round  
 382 per-channel offset and scale terms  $(\alpha, b)$ , and rolony contributions arising from  $\mathbf{F}, \varphi, \mathbf{Z}$ . As shown  
 383 above, it is straightforward to compare the total reconstruction to the observed data. However,  
 384 this does not isolate the contributions of individual estimated rolonies.

Therefore we adapted a partial subtraction approach from (Lee et al., 2020). We pick one  
 barcode,  $j^*$ , and focus only on the contributions to the reconstruction from this one barcode. In  
 particular, we assume that all other aspects of the model are exactly correct. We assume that  
 $a, \alpha, b, \varphi$  and  $\mathbf{Z}$  are all exactly right. We further assume that  $\mathbf{F}_j$  is exactly correct for every  $j \neq j^*$ .  
 Assuming all these aspects of the model were perfect, we can look at what the data *would have*  
*looked like* if it had only included one type of barcode, namely  $j^*$ . We call this counterfactual  
 simulation the ‘cleaned image’:

$$\mathbf{X}_{m,r,c}^{(j^*)} = \mathbf{X}_{m,r,c} - a_m - b_{r,c} - \sum_{j \neq j^*, m'} \mathbf{K}_{m,m'} \mathbf{F}_{m',j} \mathbf{G}_{r,c,j}. \quad (2)$$

385 This is the data with all aspects of the model subtracted away – *except* for the contributions from  
 386 barcode  $j^*$  (see Figure 9 as an example). The cleaned image for the barcode  $j^*$  has much in  
 387 common with the rolony density for  $j^*$ . However,  $\mathbf{X}^{(j^*)}$  differs from  $\mathbf{F}_{j^*}$  in one crucial way. For  
 388 each voxel  $m$ ,  $\mathbf{F}_{j^*}$  gives exactly one value. However, for each voxel  $m$ ,  $\mathbf{X}^{(j^*)}$  gives  $R \times C$  values  
 389 – one for each round and channel of the experiment. According to our model, however, it should  
 390 be possible to express all these values in terms of a mathematical ‘outer product’:

$$\mathbf{X}_{m,r,c}^{(j^*)} \approx \mathbf{F}_{j^*,m} \mathbf{G}_{j^*,r,c}.$$

391 In this outer product we see that  $\mathbf{X}_{m,r,c}^{(j^*)}$  (which varies across voxels, rounds, and channels) is the  
 392 product of two objects: the rolony density (which varies across voxels) and the transformed barcode  
 393  $\mathbf{G}$  (which varies across rounds and channels) for  $j^*$ . This is actually a very strong assumption;  
 394 most tensors would not exhibit this kind of structure. We can empirically check for this ‘rank-one’  
 395 structure by computing the singular value decomposition (SVD) of  $\mathbf{X}_{m,r,c}^{(j^*)}$ . If the SVD yields only  
 396 one strong singular value, then  $\mathbf{X}_{m,r,c}^{(j^*)}$  can be well-approximated by this rank-one outer product,  
 397 and furthermore the SVD yields the correct values for  $\mathbf{F}_{j^*,m}$  and  $\mathbf{G}_{j^*,r,c}$ . We can compare the  
 398 values for these quantities (as returned by the SVD analysis) to the estimated values (as returned  
 399 by BarDensr). We show some examples in Figure 10 comparing the estimated value of  $\mathbf{G}_{j^*}$  with  
 400 SVD results (a similar but more complete set of the spots can be seen in Figure 19). Note that  
 401 the match isn’t quite perfect (the temporal singular vector of the corresponding cleaned images

varies a bit from our estimate). In future work we hope to investigate whether these differences could be accounted for by a more accurate physical model. For now, we content ourselves that the method is accurate enough to provide a useful diagnostic for the detected rolonies.

We can also use these cleaned images to help us compare BarDensr with other methods by eye. Figure 26 investigates cleaned images for gene *Arpp19*, comparing the results of our method to the hand-curated results. In cases where the results of the two approaches disagree, these cleaned images suggest that our results are often reasonable.

## 4 Conclusion and future work

By directly modeling the physical process that gives rise to spatial transcriptomics imaging data, we found that BarDensr can correctly detect transcriptomic activity – even when rolonies are densely packed in tissue or optical resolution is limited.

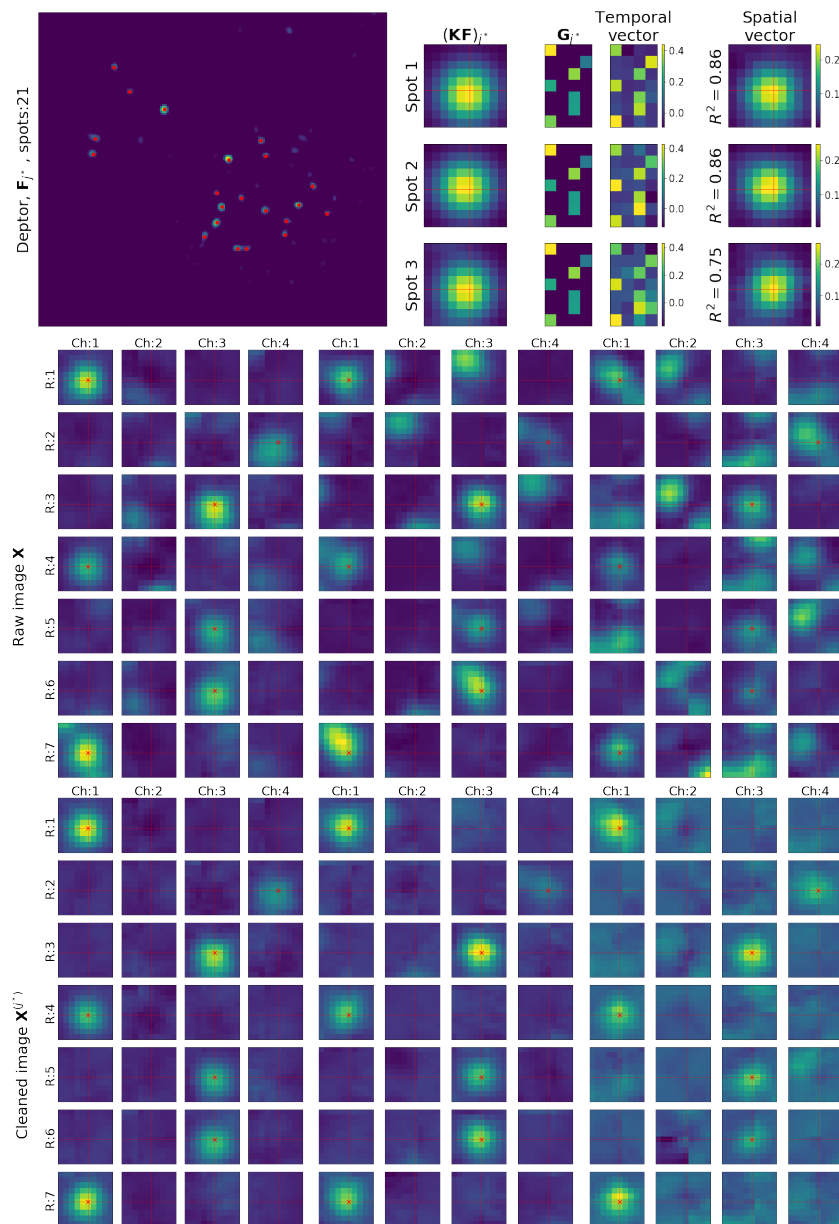
BarDensr is computationally scalable, but so far we have only investigated real-world transcriptomic experiments with less than a thousand barcodes. To scale to larger barcode libraries we need to address the possibility that the barcode library may be unknown or corrupted. In experiments with tens of thousands of barcodes, some barcodes present in the data may be unknown to the experimentalist. If these barcodes are ignored, the performance of our method may be negatively impacted. In the future we hope to adapt our method to learn these barcodes directly, using the model outlined in this paper. Together with the computational acceleration approaches used in this paper, this would extend BarDensr to larger-scale data with potentially corrupted barcode libraries.

## Acknowledgements

We thank Abbas Rizvi, Li Yuan, Daniel Soudry, Ruoxi Sun, Darcy Peterka, and Ian Kinsella for many helpful discussions. This work was supported by the National Institutes of Health [NIH 5RO1NS073129, 5RO1DA036913, RF1MH114132, U19MH114821, and U01MH109113 to A.M.Z., and 1U19NS107613 to L.P.], the Brain Research Foundation (BRF-SIA-2014-03 to A.M.Z.), IARPA MICrONS [D16PC0008 to A.M.Z. and D16PC0003 to L.P.], Paul Allen Distinguished Investigator Award [to A.M.Z.], Simons Foundation [350789 to X.C.], Chan Zuckerberg Initiative (2017-0530 ZADOR/ALLEN INST(SVCF) SUB awarded to A.M.Z and 2018-183188 to L.P.), and Robert Lourie (to A.M.Z.). This work was additionally supported by the Assistant Secretary of Defense for Health Affairs endorsed by the Department of Defense, 1120 Fort Detrick, Fort Detrick, MD 21702 through the FY18 PRMP Discovery Award Program W81XWH1910083 (to X.C). Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the U.S. Army. In conducting research using animals, the investigator adheres to the laws of the United States and regulations of the Department of Agriculture.

## References

- Taiga Abe, Ian Kinsella, Shreya Saxena, Liam Paninski, and John P Cunningham. Neuroscience cloud analysis as a service. *bioRxiv*, 2020.
- Shahar Alon, Daniel R Goodwin, Anubhav Sinha, Asmamaw T Wassie, Fei Chen, Evan R Daugherty, Yosuke Bando, Atsushi Kajita, Andrew G Xue, Karl Marrett, et al. Expansion sequencing: Spatially precise in situ transcriptomics in intact biological systems. *bioRxiv*, 2020.
- E Kelly Buchanan, Ian Kinsella, Ding Zhou, Rong Zhu, Pengcheng Zhou, Felipe Gerhard, John Ferrante, Ying Ma, Sharon Kim, Mohammed Shaik, et al. Penalized matrix decomposition for denoising, compression, and improved demixing of functional imaging data. *arXiv preprint arXiv:1807.06203*, 2018.



**Figure 9: Using cleaned images and SVD to examine model fit quality and variability.** Top left: spots are identified in  $\mathbf{F}_{j^*}$  for each barcode  $j^*$  using local-max-peak-finding. For the gene barcode (*Deptor*) shown here, 19 spots were detected (red dots) and three spots with highest accuracy are shown on the right panel. The middle and bottom panels show the zoomed-in  $R \times C$  plots of the raw image  $\mathbf{X}$  (middle) and ‘cleaned’ image  $\mathbf{X}^{(j^*)}$  (bottom) at these three spot locations for barcode  $j^*$ . Note that ‘cleaned’ images are significantly sparser than the raw images, as desired. Top right: we applied SVD to the cleaned image  $\mathbf{X}^{(j^*)}$  at these three spot locations. The first two columns show the zoomed-in image of the original spot  $(\mathbf{KF})_{j^*}$  and the learned weighted barcode matrix  $(\mathbf{G}_{j^*})$  corresponding to this gene barcode  $j^*$ . The top singular vectors are plotted in the last two columns (showing a good match with  $\mathbf{G}_{j^*}$  and the cropped  $(\mathbf{KF})_{j^*}$ ).  $R^2$  is the squared correlation coefficient between  $\mathbf{X}^{(j^*)}$  and the outer product of these two singular vectors; the high  $R^2$  values seen here indicate that the model accurately summarizes  $\mathbf{X}^{(j^*)}$ .

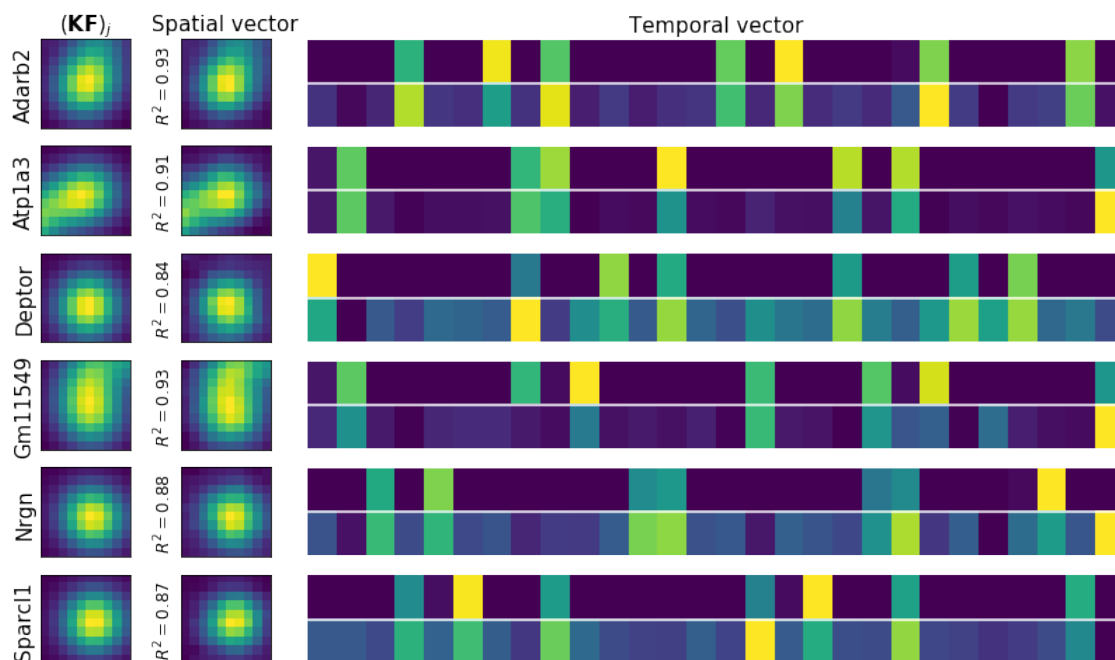


Figure 10: **Results of SVD analysis of cleaned images for the top high- $R^2$  spots.** This plot summarizes the results of the analysis illustrated in Figure 9. The first column shows  $(\mathbf{KF})_{j^*}$  cropped around the brightest spots; the second column shows the top spatial singular vectors for the same crops, and the last column shows the top temporal singular vectors for these spots. For the last column, the top row (above the white line) shows the scaled  $\mathbf{G}_{j^*}$  learned from the model, and the bottom row shows the corresponding top temporal singular vectors for these spots. Note that there is some variability visible in these temporal singular vectors.  $R^2$  is computed as in Figure 9. Only six barcodes that are most abundant in the selected region are shown here; Figure 19 provides a more complete illustration.

- 446 Kok Hao Chen, Alistair N Boettiger, Jeffrey R Moffitt, Siyuan Wang, and Xiaowei Zhuang. Spatially resolved, highly multiplexed rna profiling in single cells. *Science*, 348(6233):aaa6090, 2015.
- 447
- 448 Xiaoyin Chen, Yu-Chi Sun, George M Church, Je Hyuk Lee, and Anthony M Zador. Efficient in situ barcode sequencing using padlock probe-based baristaseq. *Nucleic acids research*, 46(4):e22–e22, 2018.
- 449
- 450
- 451 Xiaoyin Chen, Yu-Chi Sun, Huiqing Zhan, Justus M Kecsull, Stephan Fischer, Katherine Matho, Z Josh Huang, Jesse Gillis, and Anthony M Zador. High-throughput mapping of long-range neuronal projection using in situ sequencing. *Cell*, 179(3):772–786, 2019.
- 452
- 453
- 454 Arthur D Edelstein, Mark A Tsuchida, Nenad Amodaj, Henry Pinkard, Ronald D Vale, and Nico Stuurman. Advanced methods of microscope control using  $\mu$ manager software. *Journal of biological methods*, 1(2), 2014.
- 455
- 456
- 457 Georgios D Evangelidis and Emmanouil Z Psarakis. Parametric image alignment using enhanced correlation coefficient maximization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(10):1858–1865, 2008.
- 458
- 459
- 460 Daniel Gyllborg, Christoffer Mattsson Langseth, Xiaoyan Qian, Sergio Marco Salas, Markus M Hilscher, Ed S Lein, and Mats Nilsson. Hybridization-based in situ sequencing (hybiss): spatial transcriptomic detection in human and mouse brain tissue. *bioRxiv*, 2020.
- 461
- 462

- 463 Yunyun Han, Justus M Kechschull, Robert AA Campbell, Devon Cowan, Fabia Imhof, Anthony M  
464 Zador, and Thomas D Mrsic-Flogel. The logic of single-cell projections from visual cortex.  
465 *Nature*, 556(7699):51–56, 2018.
- 466 T. Hastie, R. Tibshirani, and M. Wainwright. *Statistical Learning with Sparsity: The Lasso  
467 and Generalizations*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability.  
468 Taylor & Francis, 2015. ISBN 9781498712163. URL [https://books.google.com/books?id=](https://books.google.com/books?id=LnUIrgEACAAJ)  
469 [LnUIrgEACAAJ](https://books.google.com/books?id=LnUIrgEACAAJ).
- 470 Justus M Kechschull, Pedro Garcia da Silva, Ashlan P Reid, Ian D Peikon, Dinu F Albeanu, and  
471 Anthony M Zador. High-throughput mapping of single-neuron projections by sequencing of  
472 barcoded rna. *Neuron*, 91(5):975–987, 2016.
- 473 Dongmin Kim, Suvrit Sra, and Inderjit S Dhillon. A non-monotonic method for large-scale non-  
474 negative least squares. *Optimization Methods and Software*, 28(5):1012–1039, 2013.
- 475 Je Hyuk Lee, Evan R Daugharthy, Jonathan Scheiman, Reza Kalhor, Joyce L Yang, Thomas C  
476 Ferrante, Richard Terry, Sauveur SF Jeanty, Chao Li, Ryoji Amamoto, et al. Highly multiplexed  
477 subcellular rna sequencing in situ. *Science*, 343(6177):1360–1363, 2014.
- 478 JinHyung Lee, Catalin Mitelut, Hooshmand Shokri, Ian Kinsella, Nishchal Dethe, Shenghao Wu,  
479 Kevin Li, Eduardo B Reyes, Denis Turcu, Eleanor Batty, et al. Yass: Yet another spike sorter  
480 applied to large-scale multi-electrode array recordings in primate retina. *bioRxiv*, 2020.
- 481 Eric Lubeck, Ahmet F Coskun, Timur Zhiyentayev, Mubhij Ahmad, and Long Cai. Single-cell in  
482 situ rna profiling by sequential hybridization. *Nature methods*, 11(4):360, 2014.
- 483 Jeffrey R Moffitt, Junjie Hao, Guiping Wang, Kok Hao Chen, Hazen P Babcock, and Xiaowei  
484 Zhuang. High-throughput single-cell gene-expression profiling with multiplexed error-robust  
485 fluorescence in situ hybridization. *Proceedings of the National Academy of Sciences*, 113(39):  
486 11046–11051, 2016.
- 487 Jeffrey R Moffitt, Dhananjay Bambah-Mukku, Stephen W Eichhorn, Eric Vaughn, Karthik  
488 Shekhar, Julio D Perez, Nimrod D Rubinstein, Junjie Hao, Aviv Regev, Catherine Dulac, et al.  
489 Molecular, spatial, and functional single-cell profiling of the hypothalamic preoptic region. *Sci-  
490 ence*, 362(6416):eaau5324, 2018.
- 491 Fernando Nogueira. Bayesian Optimization: Open source constrained global optimization tool for  
492 Python, 2014. URL <https://github.com/fmfn/BayesianOptimization>.
- 493 Marius Pachitariu, Nicholas A Steinmetz, Shabnam N Kadir, Matteo Carandini, and Kenneth D  
494 Harris. Fast and accurate spike sorting of high-channel count probes with kilosort. In *Advances  
495 in neural information processing systems*, pages 4448–4456, 2016.
- 496 Xiaoyan Qian, Kenneth D Harris, Thomas Hauling, Dimitris Nicoloutsopoulos, Ana B Muñoz-  
497 Machado, Nathan Skene, Jens Hjerling-Leffler, and Mats Nilsson. Probabilistic cell typing  
498 enables fine mapping of closely related cell types in situ. *Nature methods*, 17(1):101–106, 2020.
- 499 Sheel Shah, Eric Lubeck, Wen Zhou, and Long Cai. In situ transcription profiling of single cells  
500 reveals spatial organization of cells in the mouse hippocampus. *Neuron*, 92(2):342–357, 2016.
- 501 Stanley R Sternberg. Biomedical image processing. *Computer*, (1):22–34, 1983.
- 502 Yu-Chi Sun, Xiaoyin Chen, Stephan Fischer, Shaina Lu, Jesse Gillis, and Anthony M Zador.  
503 Integrating barcoded neuroanatomy with spatial transcriptomics reveals the molecular logic of  
504 cortical projections. Unpublished, 2020.
- 505 Xiao Wang, William E Allen, Matthew A Wright, Emily L Sylwestrak, Nikolay Samusik, Sam  
506 Vesuna, Kathryn Evans, Cindy Liu, Charu Ramakrishnan, Jia Liu, et al. Three-dimensional  
507 intact-tissue sequencing of single-cell transcriptional states. *Science*, 361(6400):eaat5691, 2018.






Figure 11: **Spatial rolonity density  $F_j$  of all 81 barcodes.** These images are the supplement to Figure 2 in the main text. The rolonity densities represent a *demixed* view of the data. Each plot corresponds to a single barcode, and indicates the rolonity density at different spatial locations. Above we show these rolonity densities for one region in the experimental data. The title for the plots above indicates the gene associated with the barcode as well as the maximum intensity of the plot. The orange dots represent rolonies detected by a hand-curated approach.

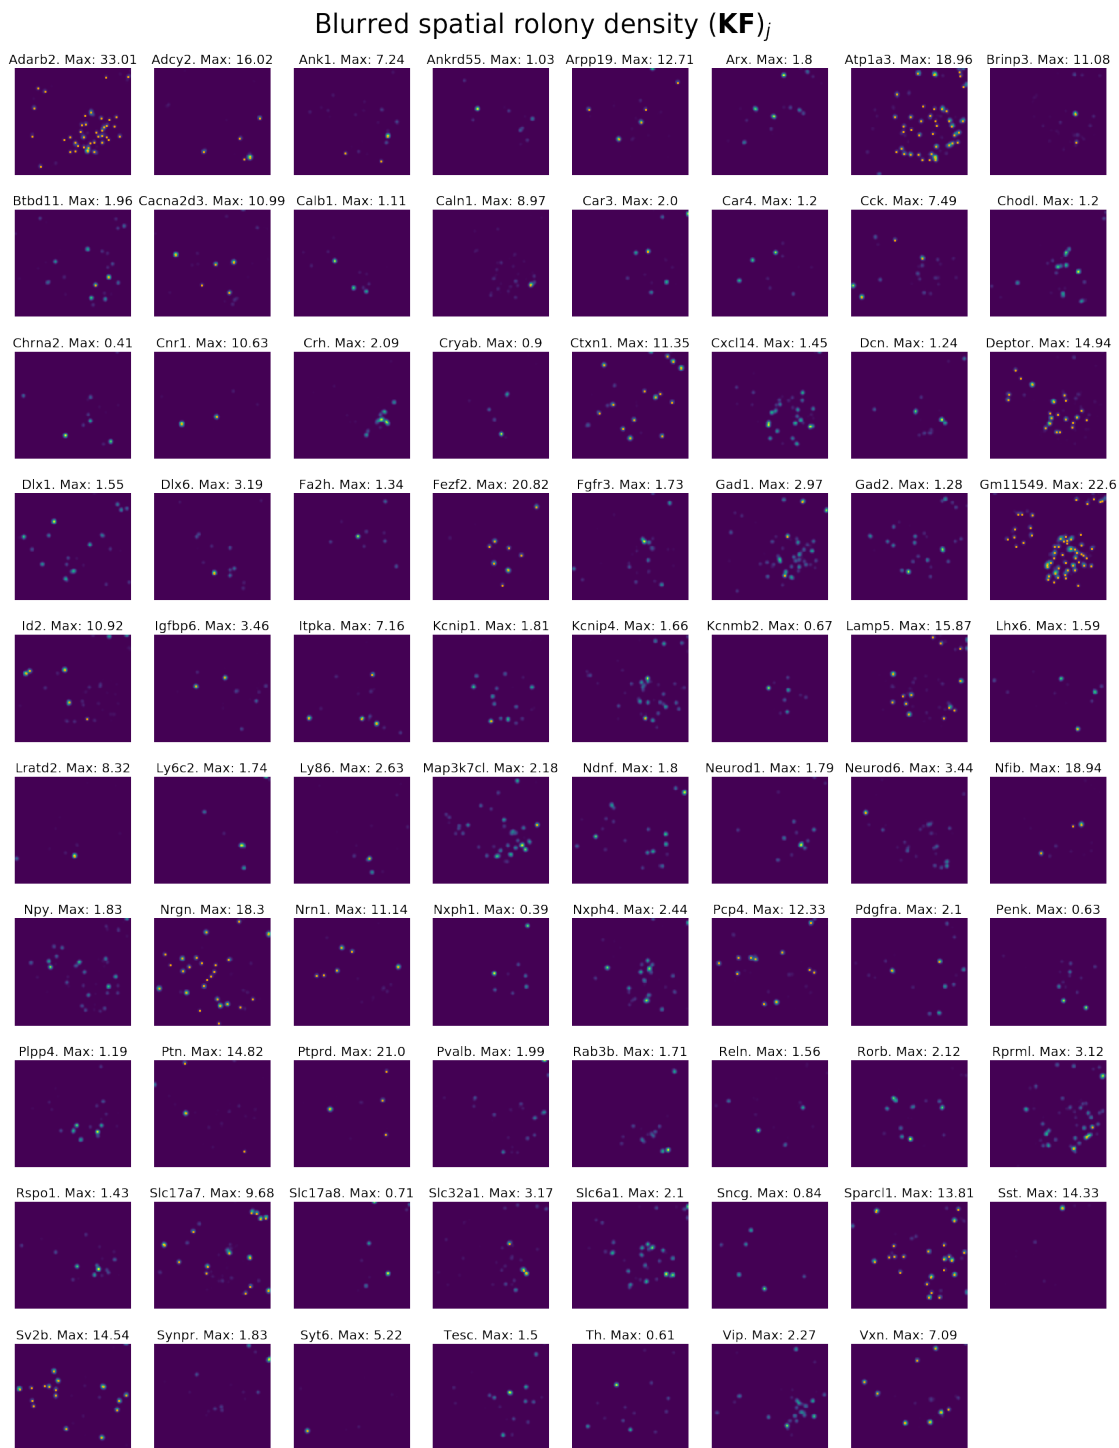


Figure 12: Spatial rolonity density ( $\mathbf{KF}$ )<sub>*j*</sub> of all 81 barcodes, after applying the point-spread function. As of Figure 11, these images are the supplement to Figure 2 in the main text, except we display ( $\mathbf{KF}$ )<sub>*j*</sub> instead of ( $\mathbf{F}$ )<sub>*j*</sub> for each barcode *j*. Recall that the point-spread function  $\mathbf{K}$  has the effect of smearing signal over a spatially localized area. It represents physical processes which blur the signal of interest. Under the BarDensr model, the signal intensities observed at each voxel *m* from a given barcode will arise directly from linear combinations of ( $\mathbf{KF}$ )<sub>*m,j*</sub> over different barcodes *j*.




Figure 13: **Original data (X) after normalization for each round and channel.** In order to create clearer visualizations, we noise-normalized the data as described in Appendix A, so that images from all rounds and channel are on the same scale.




Figure 14: **Data reconstructed from BarDensr.** Under the BarDensr model, the fluorescence signal observed at each voxel in Figure 13 should be approximately given by the equations from Section 2. We here plot the results of those equations, visualized using the same colormap-intensity scale as used Figure 13. At least by eye, we see excellent agreement between the data and the model's predictions.






Figure 15: **Residuals.** As mentioned in Figure 14, the BarDensr model makes predictions about what the observed data should look like. There is broad agreement, but there is some disagreement. Here we highlight the the residual between the predictions and the data. Note the difference in scale compared to the previous two figures.




Figure 16: **Zoomed original data (X) after normalization for each frame.** Zoomed in for one of the target spots (a  $20 \times 20$  region). See Figure 13 for more details.






Figure 17: **Zoomed reconstructed data.** Zoomed in for the target spots ( $20 \times 20$ ), plotted at the same intensity scale as Figure 16. See Figure 14 for more details.




Figure 18: **Zoomed residual plot.** Zoomed in for the target spots ( $20 \times 20$ ). See Figure 15 for more details.




Figure 19: **Further SVD analysis of cleaned images for the top high- $R^2$  spots.** The figure supplements Figure 10 in the main text, and is structured in the same way, except that this plot shows more examples (with more barcodes and spots). Each row shows two spots for a given barcode. The first two columns show  $(\mathbf{KF})_{j^*}$  cropped around the two spots; the third and fourth columns show the top spatial singular vectors for the same crops. The final wide column shows the top temporal singular vectors for these spots, with the first row (above the thin white line) showing the scaled  $\mathbf{G}_{j^*}$  learned from the model, and the following two rows showing the corresponding top temporal singular vectors for these spots. The two spots are ordered by  $R^2$ , which is computed as in Figure 9.




Figure 20: **Benchmarking results with no dropout.** Comparing *starfish*, SRM, ‘corr’, and BarDensr results, to the ground truth. Showing the top six barcodes with highest density (the gene density was generated randomly, see Appendix E). This figure corresponds to the top left plot in Figure 4 in the main text. Without dropout, BarDensr accurately detects the barcodes in the original data.




Figure 21: **Benchmarking results with 50% dropout.** Similar to Figure 20 but with dropout for 50% of the simulated spots. This figure corresponds to the top right plot in Figure 4 in the main text.




Figure 22: **Benchmarking results with no dropout, using five times denser simulation.** This is the similar process except that the spots density is five times denser than Figure 20 and 21. This figure corresponds to the bottom left plot in Figure 4 in the main text.






Figure 23: **Benchmarking results with 50% dropout, using five times denser simulation.** Similar to Figure 22 but with dropout for 50% of the simulated spots, for the denser simulation. Some missing spots (FN) can be observed from our model as well as other two methods (e.g., see the fifth row *Kif5a*). False discovery (FP) can also be seen in this plot for SRM (e.g., see the third row *Rbfox3*). This figure corresponds to the bottom right plot in Figure 4 in the main text.




Figure 24: **Further demonstration of BarDensr applied to low-resolution data (supplementing Figure 6.)** To test BarDensr's performance on low-resolution data, we first run BarDensr on the original data, obtain rolny densities, and then finally downsample the rolny densities ('run-then-downsample'). Next, we run BarDensr on downsampled data and look at the learned rolny densities ('downsample-then-run'). For highly-expressed genes, these two results are nearly indistinguishable.

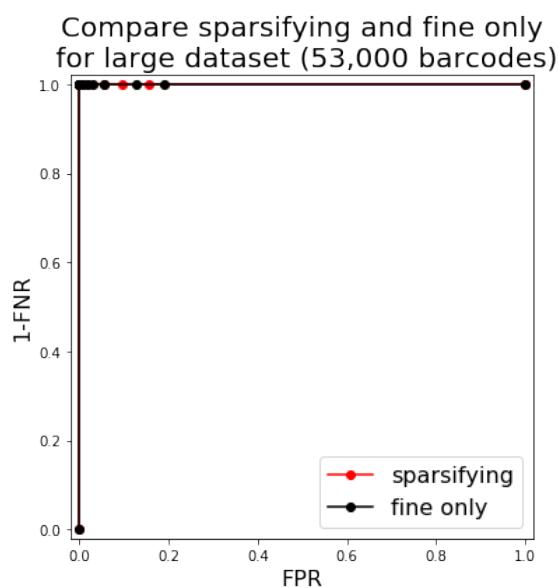


Figure 25: **BarDensr can be scaled up to a larger number of barcodes by using sparsified image.** To test if we can scale up BarDensr, we computed an ROC curve for the method using a simulated dataset with 53,000 barcodes and 17 sequencing rounds. After running the model on a  $5\times$  downsampled  $50 \times 80$  pixels simulated image, barcodes that are set to zero at the coarse scale were removed and the model was run at the original scale, with the parameters learned from the downsampled image as the initial conditions. See also Appendix H.

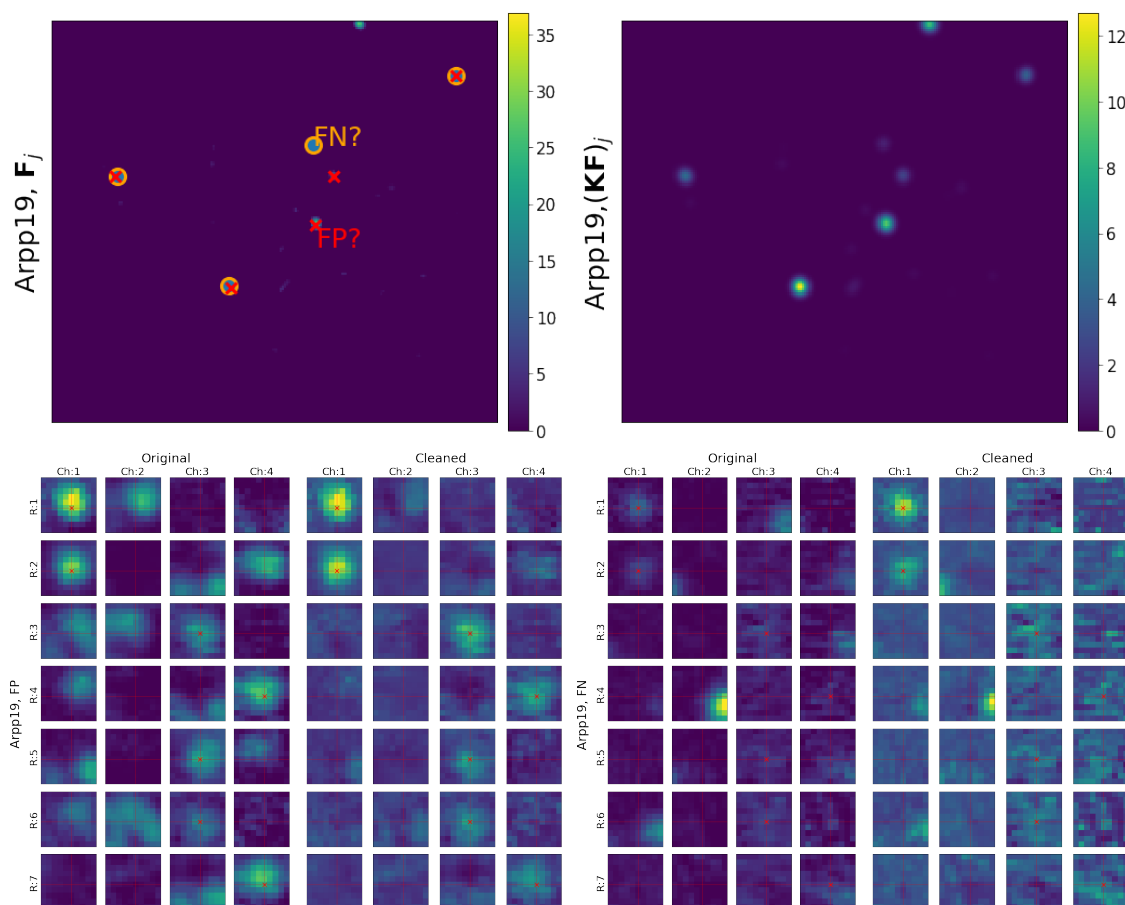


Figure 26: **Spatial rolonity density comparison against hand-curated results.** On the top two plots, we show the rolonity density  $F_j$  (left) and the blurred rolonity density  $(\mathbf{KF})_j$  (right) for gene *Arpp19*, derived from the experimental data. These rolonity densities indicate the presence of *Arpp19*-rolonies. However, they might be incorrect, indicating that these detected rolonies might not be present in the real data. In this figure we investigate this question qualitatively. First, we compare with the rolonity positions detected by a hand-curated method (as represented by orange circles on the left top plot) with the rolonies suggested by the rolonity densities (as indicated by red crosses on the left top plot). We see a broad agreement. Where there is a point of disagreement, we can visualize the signal intensities in all the voxels near that point. The two plots on the bottom-left show the original data from a spot that was detected by BarDensr, but not detected in the hand-curated results (as indicated as False Positive (FP) in the top left plot); the left columns show the original image and the right columns show the ‘cleaned’ image (similar to Figure 9, see Equation 2 for details). The red cross in each round indicates the channels that are activated by this barcode. These crosses line up well with the observed signal, suggesting BarDensr has correctly identified a new rolony. It appears that the hand-curated method failed to detect this rolony because of the presence of nearby rolonies, leading to a mixed signal; BarDensr is specifically designed to handle these kinds of confusing situations. The two plots on the bottom-right show a spot which is detected in the hand-curated result but not detected by BarDensr (as indicated as False Negative (FN) in the top left plot). We show both the original data and the cleaned data, as in the bottom left plots. In this case, the data do not appear to support the presence of a rolony, suggesting BarDensr correctly rejected this region as a rolony and the the hand-curated approach labelled it incorrectly. We conjecture that the hand-curated approach misidentified this as a spot because of the signal arising from a nearby rolony in round 7, channel 4; this again created a mixture of signals BarDensr was better equipped to recognize.




Figure 27: **Supplement to Figure 2 for *Nrgn***. The top plot shows the same spatial rolon density of *Nrgn* as in Figure 2. The orange crosses indicate the spots detected in the hand-curated results. The three spots highlighted with red are further zoomed in the bottom, with their indices on the titles. These spots were detected to have large signal intensities by BarDensr, but were not detected in the hand-curated results. The correct barcode frames for *Nrgn* are indicated with red crosses in the bottom plots, suggesting that each of these spots appear to be well-modeled as *Nrgn* spots.

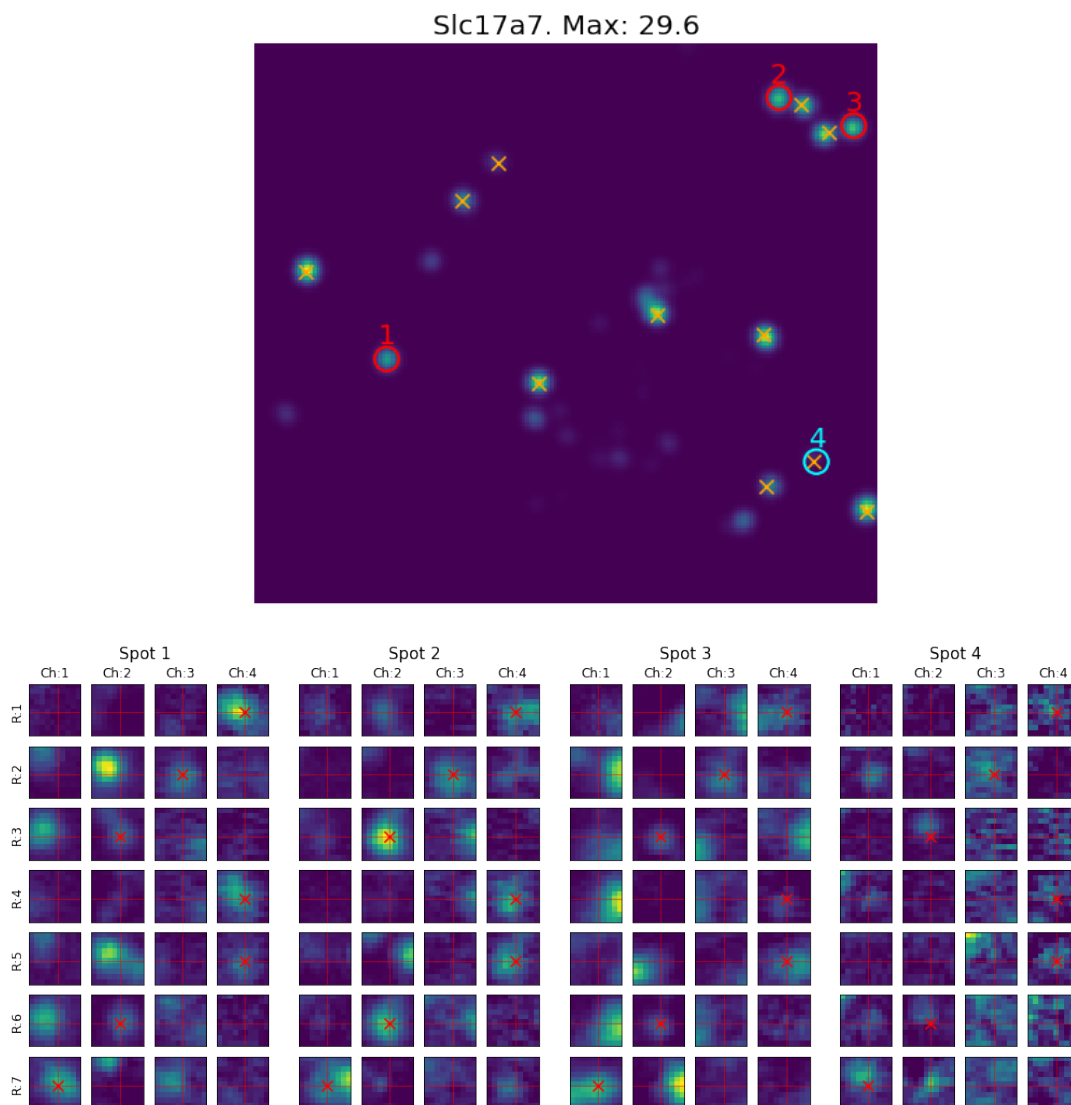


Figure 28: **Supplement to Figure 2 for *Slc17a7***. The top plot shows the same spatial rolon density of *Slac17a7* as in Figure 2. The orange crosses indicate spots that detected by hand-curated method. The four spots highlighted with red or cyan are further zoomed in the bottom, with their index on the titles. The first three spots (Spot 1 - 3, shown in red) were found by BarDensr but were not detected by hand-curated results, as in Figure 27. The fourth spot (Spot 4, shown in cyan) is the spot that is detected by hand-curated results, but no signal detected in BarDensr. The correct barcode frames for *Slc17a7* are indicated with red crosses in the bottom plots.



## 508 A Data preprocessing

509 The data is preprocessed before input into the model as follows: first the data was max-projected  
510 across all z-stacks. The channel color-mixing was corrected and the background was removed  
511 using rolling-ball background subtraction (Sternberg, 1983). Then the different image stacks  
512 were registered to the same voxels, using the Image Alignment Toolbox (ECC image alignment  
513 algorithm) (Evangelidis and Psarakis, 2008).

514 Finally, we performed a crude noise-normalization on each frame. First we estimated the  
515 noise level on each frame by spatially high-pass filtering (i.e., original image minus a Gaussian-  
516 filtered image, with a sigma of 2 pixels) to isolate spatially-uncorrelated noise, then computing  
517 the standard deviation. (See e.g. (Buchanan et al., 2018) for a related approach applied in the  
518 temporal domain.) Then we divided each original frame by its estimated noise scale to obtain the  
519 noise-normalized images.

## 520 B Hardware time and cost comparisons

521 To develop an efficient implementation of BarDensr on the NeuroCAAS cloud platform (Abe  
522 et al., 2020), we needed to find the most cost-effective hardware for the job. Using a  $1000 \times 1000$   
523 sized image from the experimental data described in the main text, we ran the model on several  
524 different AWS instance types. The most cost-effective machine was `m5.2xlarge`, which completed  
525 the analysis in three and a half minutes with a total cost of two cents. On the other extreme, the  
526 `p3.2xlarge` machine completed the analysis in one minute with a total cost of five cents. As a  
527 compromise between speed and cost, we settled on the `p2.xlarge` machine, which completes the  
528 analysis in two minutes with a total cost of three cents.

## 529 C Single Round Matching (SRM) and ‘hand-curated’ method

530 We compared BarDensr against several different alternative methods, including one we call ‘SRM.’  
531 This method is an implementation of the widely-used ‘blobs-first’ algorithms suggested in the  
532 literature (Wang et al., 2018; Qian et al., 2020). First, blobs were detected in every channel in  
533 the first round by finding local maxima on a per-channel basis. These blobs were then used as  
534 a reference in understanding subsequent rounds (the first round is used as the reference since it  
535 usually has the least corruption by noise and artifacts, such as phasing and photo-bleaching).

536 At each detected rolon position, SRM then read out the signal intensities from all chan-  
537 nels/rounds as a vector of length  $R \times C$ . This vector was compared against each barcode in the  
538 library. Each detected rolon was assigned to the barcode with the greatest similarity (as measured  
539 by a dot-product). For some rolonies the similarity was low to all barcodes; these rolonies were  
540 filtered out. Thresholds were determined using the Bayes optimization method from (Nogueira,  
541 2014).

542 The ‘hand-curated’ method in the main text corresponds to SRM described here. After the  
543 process of SRM with the chosen threshold, we manually checked the detected rolonies and the  
544 assigned barcodes to make sure that the results were reasonable.

## 545 D Correlation-based method

546 We also compared BarDensr against a ‘correlation-based method’ (Moffitt et al., 2016, 2018). This  
547 approach begins by computing a vector of length  $R \times C$  for every voxel, indicating the fluorescence  
548 signal in each round and channel at that voxel. At each voxel, for each barcode, the cosine  
549 distance between this vector and the barcode was computed. The barcode with the minimum  
550 cosine distance was assigned to be a *potential* gene identity for each voxel. Finally, a ‘minimum  
551 distance image’ was constructed: for each voxel, this image contains the cosine distance between  
552 that voxel’s  $R \times C$  vector and the barcode which it is most similar to. Coordinates of blobs were

553 found by a seeking local minima in this image. Thresholds were again determined via (Nogueira,  
554 2014).

## 555 E Simulation

### 556 Generating arbitrary distribution for genes

557 For the simulation benchmarking in Figures 4 and 20-23, we used barcodes from a STARmap  
558 experiment (developed in (Wang et al., 2018), unpublished data), with total of 57 genes. This  
559 data is similar to our original experimental data in that it has six rounds and four channels in  
560 total, and the scale of the number of barcodes is also similar. This data was chosen instead of  
561 our experimental data in order to directly apply *starfish* method (the *starfish* application on our  
562 original experimental data was not available at the time this analysis was conducted). In creating  
563 simulations, we wanted to accurately represent the uneven distribution of genes; in real data some  
564 genes are more abundant than others. Therefore, we began by randomly selecting 10 out of 57  
565 genes to be ‘abundant’ genes. In generating a dataset with simulated rolonies, we created rolonies  
566 with these abundant genes roughly ten time more often than the other rolonies.

### 567 Dropout

568 We used two setups to generate simulated testing data: without dropout and with dropout. In the  
569 experimental data, it is commonly observed that a small portion of rolonies disappear/diminish in  
570 some rounds. (Based on our visual inspection of the experimental data, qualitative dropout events  
571 were observed in  $< 5\%$  of the rolonies detected in this data, but we did not attempt to estimate  
572 this dropout rate precisely.) In the ‘Dropout’ simulations, we tried to mimic this phenomenon.

573 Specifically, for the ‘no dropout’ simulations, we generated the data with the following process.  
574 1. Generate the spot position with a uniform distribution across the voxels. 2. For each spot  
575 position, generate the spot identity (gene) using a prespecified gene distribution (as discussed  
576 above). 3. For each position  $m$  and gene  $j$  pair from steps 1 and 2, the *magnitude* of the rolony  
577 density at  $(m, j)$  was generated from a uniform distribution in the range  $(10, 40)$ . We use these  
578 values to fill in the rolony density,  $\mathbf{F}$ . 4. We then generate synthetic data according to the  
579 BarDensr model. Finally, we add some speckle noise. Note that in our simulation, parameters  
580 such as the per-frame intensity ( $\alpha$ ), phasing ( $\rho$ ), and color-mixing ( $\varphi$ ) were left out for simplicity.

581 For the ‘dropout’ simulations, 50% of the simulated spots were randomly selected to be the  
582 ‘dropout spots.’ For each ‘dropout spot’, one round is selected randomly and the signal intensity  
583 for that spot for that round is diminished to 10% of the original signal. The simulation process is  
584 otherwise the same.

### 585 Hybrid simulation

586 To test the efficacy on our model on even more realistic data, we used a ‘hybrid simulation,’  
587 as in e.g. (Pachitariu et al., 2016). In essence, the hybrid simulation creates a fake dataset by  
588 superimposing the real data with additional synthetic rolonies to the data. The question is whether  
589 the algorithm can at least find the synthetic rolonies which were added. The results are shown  
590 in Figure 5. In generating the synthetic rolonies, we used the codebook used in the original  
591 experiment, and the genes of the synthetic rolonies were given by the observed gene distribution  
592 from the hand-curated analysis of the same dataset.

593 We ran a few different versions of these simulations. There were several key parameters which  
594 we varied:

- 595 • We had both ‘dropout’ and ‘no dropout’ versions; in the dropout versions some of the  
596 synthetic rolonies had signal in one round diminished.
- 597 • We could vary the number of synthetic spots which were injected ( $S$ ). According to the  
598 hand-curated analysis of the real data, the real data contained approximately 400 spots in

599 this field of view. We investigated how the number of spots affected the results, looking at  
600  $S \in \{30, 80, 100, 200\}$ .

- 601 • We could vary the intensity of the synthetic spots, relative to the maximum intensity observed  
602 in the data. We varied this between 10% and 90%.

## 603 F Error analysis

604 In simulated data, we can exactly quantify the different kinds of errors that BarDensr makes, by  
605 comparing against the true rolon positions used to make the simulated data. We first examined  
606 the ‘total hit rate’ in our Receiver Operating Characteristic curve (ROC curve, shown as the  
607 dotted lines in Figure 4). For this ROC, we consider a spot to be successfully detected by the  
608 algorithm as long as the algorithm finds *any rolon* near the site of a true rolon – even if the  
609 algorithm incorrectly assigns the gene associated with that true rolon. The ROC for BarDensr  
610 clings closely to the upper left side of the plot, suggesting nearly perfect performance. We also  
611 looked at what we call the ‘hit rate’ – for this ROC we consider a spot to be successfully detected  
612 only if the algorithm detects a rolon in the right place and of the right gene. Figure 4 shows the  
613 results, suggesting most errors were caused by gene mis-identification.

## 614 G Cell segmentation

615 For the bottom plots on Figure 6, we first segmented the cells in the selected region with the  
616 following process. We first obtained the max projection across  $R \times C$  frames from the the image  
617 stacks. After applying a Gaussian filter with a sigma of 8 pixels, all the pixels with intensity  
618 lower than 10% of the maximum intensity were assigned to be zero. Finally, we used a watershed  
619 segmentation algorithm to identify contiguous cellular regions. This results in 47 segmented cells  
620 in the region. Four of these occupied less than 100 pixels in total and were removed from the  
621 analysis.

## 622 H Sparsifying and coarse-to-fine

### 623 H.1 Handling tens of thousands of barcodes with sparsifying and coarse- 624 to-fine

625 To scale up BarDensr, we tested if eliminating unnecessary barcodes could help accelerate compu-  
626 tation. For this purpose, we set up a simulation with 53,000 barcodes and 17 sequencing rounds,  
627 similar to the setup in a larger scale experiment such as (Chen et al., 2018) (Figure 25). We  
628 generated a dataset with  $50 \times 80$  voxels and a total of 40 spots. We then processed this data in  
629 two steps.

630 In the first step (the ‘coarse’ step), the image was five times downsampled, and BarDensr was  
631 applied to the downsampled data. For each gene, if the maximum intensity for a rolon density  
632 was lower than  $10^{-5}$ , that gene was considered to be absent.

633 In the second step (the ‘sparsified fine’ step), we then applied BarDensr to the original, full-  
634 resolution data – but only using those barcodes that weren’t ‘absent’ in the previous step. To  
635 make this approach even faster, we also used the learned parameters from the downsampled data  
636 as initial conditions for the algorithm’s run on the full-resolution data. Moreover, in this second  
637 step, the parameter  $b$  and  $\alpha$  were not updated at all, since we found that they were learned quite  
638 accurately in the first step.

639 After the first step, 71 out of 53,000 barcodes were kept to be used in the second step. This  
640 sparsified coarse-to-fine approach sped up our analysis by more than a factor of 10.

## 641 H.2 Coarse-to-fine

642 The method described above involves both sparsifying and using a kind of coarse-to-fine approach.  
643 We also investigated performance using only the coarse-to-fine aspect. These investigations are  
644 summarized in Figure 7. First, a  $1000 \times 1000$  image was simulated with 20,000 spots following  
645 the simulation process described above, with no dropout (Appendix E). The image was then  
646 downsampled to  $500 \times 500$  and BarDensr was applied to this downsampled data to estimate  
647  $\mathbf{F}$ ,  $\alpha$ ,  $a$  and  $b$ . with 20 iterations. These parameters were then used as the initial conditions to run  
648 the model with the full size image (note that in order to use the parameters from the downsampled  
649 image to initialize the full-resolution model, we needed to upsample  $\mathbf{F}$  and  $a$ ). Finally, we tried  
650 running the model directly on the full-resolution data (without using the downsampled data to get  
651 initial conditions). We then compared the results. Both approaches work better if they are allowed  
652 to run longer, because they use an iterative approach to optimize the loss function. Eventually,  
653 both approaches yield the same results. However, Figure 7 shows that the coarse-to-fine approach  
654 is able to achieve the best possible performance three times faster.

## 655 H.3 Sparsifying leads to speedups even in the case of a small barcode 656 library

657 We also tested the sparsifying approach on the experimental data (Figure 8). The original  
658  $1000 \times 1000$  image was first five times downsampled to obtain a  $200 \times 200$  ‘coarse’ image, and  
659 the parameters were learned from this downsampled image (‘coarse’ process). After five times  
660 upsampling of learned  $\mathbf{F}$  to obtain the image on the original scale, both the original image and the  
661 upsampled  $\mathbf{F}$  were split into  $4 \times 4$  patches. Each patch is of size  $250 \times 250$  (plus 20 pixel edges in  
662 the end of both dimension, whenever the coverage of the patch does not exceed the image region).  
663 These 16 patches cover the entire  $1000 \times 1000$  image with overlaps on the edge regions. For each  
664 patch, the barcodes that have the maximum intensity lower than the maximum intensity of the  
665 two unused barcodes in the upsampled  $\mathbf{F}$  were considered to be absent from the region. Each  
666 patch was then used to fit the model at the original scale to learn  $\mathbf{F}$  and  $a$ , but using a smaller  
667 amount of barcodes for the binary codebook matrix  $\mathbf{B}$ . For this ‘fine’ process, the parameter  $b$   
668 and  $\alpha$  were not updated but the ones learned from the coarse process were used. After the model  
669 was run on all the 16 patches, we needed to stitch the results back together into a single result  
670 for the entire field of view. This is slightly involved, because the patches concerned overlapping  
671 regions of voxels. Indeed, we insured that the border-regions between any two patches contained  
672 20 pixels of overlap. For each voxel in these overlap regions, we used the signal from the patch  
673 whose center was closest to that voxel. The results were compared to the ‘fine only’ approach. We  
674 filled  $\mathbf{F}$  with zero for the removed barcodes in each patch, in order to keep the original dimension  
675 for the following process.

676 In order to compare the agreement between ‘sparsifying’ approach and the ‘fine only’ approach,  
677 we computed two ROC curves. In each curve, one of the results was used as the gold standard for  
678 the other. Specifically, when using the ‘fine only’ as the gold standard, a threshold was determined  
679 based on the maximum intensity of the unused barcodes in the ‘fine only’ results, and a binary  
680 matrix of the size of the original image ( $1000 \times 1000$ ) was generated for each barcode (‘0’ indicates  
681 there is no signal, ‘1’ indicates there is signal, for each pixel), which is the final gold standard to  
682 compare with the sparsifying result. The same process applies when using the sparsifying result as  
683 the gold standard. For computing ROC, colonies within 3 pixel radius of  $\mathbf{F}$  with the same barcode  
684  $j$  were considered to belong to the same colony.

## 685 I Algorithm Details

The primary computational challenge of this method is to solve a constrained optimization problem.

$$\begin{aligned} & \min_{\theta \in \Theta} L_{\text{sparsity}}(\theta) \\ & \text{subject to } L_{\text{reconstruction}}(\theta) \leq \omega, \end{aligned}$$

where

$$\begin{aligned} L_{\text{reconstruction}}(\theta) &\triangleq \sum_{m,r,c} \left( X_{m,r,c} - \left( a_m + b_{r,c} + \alpha_{r,c} \sum_{j,m',c'}^{J,M,C} \mathbf{K}_{m,m'} \mathbf{F}_{m',j} \varphi_{c,c'} \mathbf{Z}_{r,c',j} \right) \right)^2 \\ L_{\text{sparsity}}(\theta) &\triangleq \sum_{m,r,c,c',j} \alpha_{r,c} \mathbf{F}_{m,j} \phi_{c,c'} \mathbf{Z}_{r,c',j} \\ \theta &= (\mathbf{F}, \rho, \alpha, b, a, \varphi). \end{aligned}$$

686 Here  $\Theta$  denotes the set of feasible parameters; in our case,  $\Theta$  simply requires that all variables  
687 are at least  $10^{-10}$ . The threshold  $10^{-10}$  was chosen somewhat arbitrarily and serves to ensure  
688 numerical stability of the optimization process. Technically we also believe that  $\rho_c < 1$  for every  
689  $c$ , but in practice we found it unnecessary to enforce this constraint.

We approach the reconstruction constraint using Lagrange multipliers. We define:

$$\begin{aligned} \mathcal{L}(\theta, \lambda) &= L_{\text{reconstruction}}(\theta) + \lambda L_{\text{sparsity}}(\theta) \\ \theta^*(\lambda) &= \arg \min_{\theta \in \Theta} \mathcal{L}(\theta, \lambda) \\ L_{\text{reconstruction}}^*(\lambda) &= L(\theta^*(\lambda)). \end{aligned}$$

690 Assuming we can evaluate  $L_{\text{reconstruction}}^*(\lambda)$ , we can solve the overall constrained optimization  
691 problem by taking

$$\lambda^* = \max \{ \lambda : L_{\text{reconstruction}}^*(\lambda) \leq \omega \}$$

692 and taking our final parameters to be  $\theta^*(\lambda^*)$ . It is unclear whether strong duality holds in this  
693 case, so the resulting parameters may not be optimal. However, in practice we find that they give  
694 useful results for uncovering rolonies.

In conclusion, to solve our overall problem it suffices to be able to solve  $\min_{\theta \in \Theta} \mathcal{L}(\theta, \lambda)$  for any fixed  $\lambda$ . We approach this problem via a blockwise coordinate descent approach. Specifically, we start with an initial guess and iterate through a variety of updates until convergence is achieved. Throughout, we will use the notations

$$\begin{aligned} [x]_+ &= \begin{cases} x & \text{if } x > 10^{-10} \\ 10^{-10} & \text{otherwise} \end{cases} \\ \text{xmab}_{m,r,c} &= X_{m,r,c} - a_m - b_{r,c} \\ X_{m,r,c}^{(j)} &= X_{m,r,c} - a_m - b_{r,c} - \alpha_{r,c} \sum_{j' \neq j} \sum_{m',c'} \mathbf{K}_{m,m'} \mathbf{F}_{m',j'} \varphi_{c,c'} \mathbf{Z}_{r,c',j'} \\ \tilde{\mathbf{F}}_{m,j} &= \sum_{m'} \mathbf{K}_{m,m'} \mathbf{F}_{m',j}. \end{aligned}$$

**$\alpha$  update.** For each  $r, c$ , the relevant portion of the loss for  $\alpha_{r,c}$  is given by

$$\begin{aligned} \mathcal{L}(\alpha_{r,c}) &= \sum_m \left( X_{m,r,c} - \left( a_m + b_{r,c} + \alpha_{r,c} \sum_{j,c'} \tilde{\mathbf{F}}_{m,j} \varphi_{c,c'} \mathbf{Z}_{r,c',j} \right) \right)^2 \\ &\quad + \lambda \sum_{m,c',j} \alpha_{r,c} \mathbf{F}_{m,j} \varphi_{c,c'} \mathbf{Z}_{r,c',j} + \dots \end{aligned}$$

695 Fixing all other variables, subject to the constraint that  $\alpha_{r,c} \geq 10^{-10}$ , the lowest possible  
696 value of this loss is given by

$$\alpha_{r,c} \leftarrow \left[ \frac{\sum_m \left( \mathbf{x} \mathbf{m} \mathbf{a} \mathbf{b}_{m,r,c} \sum_{j,c'} \tilde{\mathbf{F}}_{m,j} \varphi_{c,c'} \mathbf{Z}_{r,c',j} \right) - \sum_{m,c',j} \mathbf{F}_{m,j} \varphi_{c,c'} \mathbf{Z}_{r,c',j}}{\sum_m \left( \sum_{j,c'} \tilde{\mathbf{F}}_{m,j} \varphi_{c,c'} \mathbf{Z}_{r,c',j} \right)^2} \right]_+.$$

697 Note that this update can be done in parallel across all  $r, c$ .

698  **$\rho$  update.** We update the variable  $\rho$  via a line-search. One at a time, we look at  $\rho_c$  and consider  
699 possible values for this parameter in the interval  $[\rho_c/2, 3\rho_c/2]$ . We search for values of  $\rho_c$  in  
700 this interval which minimizes the loss.

**$a, b$  updates.** Fixing all other variables, the best possible values for  $a$  are easy to find. The same  
goes for  $b$ . These values are given by

$$a_m \leftarrow \left[ \frac{1}{RC} \sum_{r,c} \left( X_{m,r,c} - b_{r,c} - \alpha_{r,c} \sum_{j,c'} \tilde{\mathbf{F}}_{m,j} \varphi_{c,c'} \mathbf{Z}_{r,c',j} \right) \right]_+$$

$$b_{rc} \leftarrow \left[ \frac{1}{M} \sum_m \left( X_{m,r,c} - a_m - \alpha_{r,c} \sum_{j,c'} \tilde{\mathbf{F}}_{m,j} \varphi_{c,c'} \mathbf{Z}_{r,c',j} \right) \right]_+.$$

**$\mathbf{F}$  update.** We make updates to  $\mathbf{F}$  one column at a time. We select a random column,  $j^*$ , and then  
update the values of  $\{\mathbf{F}_{m,j^*}\}_{m \in \{1 \dots M\}}$ . We update these values via a projected coordinate  
descent algorithm. Let  $f$  denote the  $j^*$ th column of  $\mathbf{F}$  and define

$$g_{r,c} = \alpha_{r,c} \sum_{m',c'} \varphi_{c,c'} \mathbf{Z}_{r,c',j^*}$$

$$\|g\|_2^2 = \sum_{r,c} g_{r,c}^2$$

$$\phi_m = \sum_{r,c} \left( \sum_{m'} \left( K_{m,m'} X_{m',r,c}^{(j)} \right) - \frac{1}{2} \lambda \right) g_{r,c}.$$

701 In terms of these objects, it is straightforward to show that the relevant portion of the loss  
702 can be written as

$$\mathcal{L}(f) = \|g\|^2 \|Kf\|^2 - 2f^T \phi + \dots.$$

703 We would like to minimize this subject to the constraint that  $f_m \geq 10^{-10}$ . To approach  
704 this problem we use a projected gradient descent approach. We start by selecting a search  
705 direction, namely the gradient of the Lagrangian:

$$\Delta = \phi - \|g\|^2 K^T K f.$$

706 We then zero out the coordinates of this search directions which point negatively along the  
707 active constraints:

$$\tilde{\Delta}_m = \begin{cases} 0 & \text{if } f_m = 10^{-10} \text{ and } \Delta_m \leq 0 \\ \Delta_m & \text{otherwise.} \end{cases}$$

708 We then update  $f$  by moving it somewhat in this search direction and then forcing it to be  
709 positive. How far should we move in the search direction? Following (Kim et al., 2013), we  
710 use the following carefully-chosen step-size:

$$f \leftarrow \left[ f + \frac{\phi^T \tilde{\Delta} - \|g\|^2 K^T K \tilde{\Delta}}{\|K \tilde{\Delta}\|^2} \tilde{\Delta} \right]_+.$$



711 If we did not force  $f$  to be positive, such updates would yield the best possible distance to  
 712 travel along the search direction. However, due to the positivity-enforcement, one can find  
 713 pathological examples where applying this update actually makes the loss worse. To be safe,  
 714 we use a backtracking procedure; as long as the loss is made actively worse by this step, we  
 715 cut the learning rate in half and try again.

$\varphi$  **update** Fix  $c^*$ . Let us look at the loss with respect to  $\varphi_{c^*,1}, \varphi_{c^*,2} \cdots \varphi_{c^*,C}$ . We find that it is given by

$$\begin{aligned} \mathcal{L}(\varphi_{c^*}) = & \sum_{m,r} \left( \mathbf{x} \mathbf{m} \mathbf{a} \mathbf{b}_{m,r,c^*} - \sum_{c'} \varphi_{c^*,c'} \left( \sum_{j,m'} \alpha_{r,c^*} \mathbf{K}_{m,m'} \mathbf{F}_{m',j} \mathbf{Z}_{r,c',j} \right) \right)^2 \\ & + \lambda \sum_c \varphi_{c^*,c'} \left( \sum_{m,j} \alpha_{r,c} \mathbf{F}_{m,j} \mathbf{Z}_{r,c',j} \right) + \cdots . \end{aligned}$$

Define

$$\begin{aligned} \Gamma_{c_1,c_2} = & \sum_{m,r} \left( \sum_{j,m'} \alpha_{r,c^*} \mathbf{K}_{m,m'} \mathbf{F}_{m',j} \mathbf{Z}_{r,c_1,j} \right) \left( \sum_{j,m'} \alpha_{r,c^*} \mathbf{K}_{m,m'} \mathbf{F}_{m',j} \mathbf{Z}_{r,c_2,j} \right) \\ \phi_c = & \sum_{m,r} \mathbf{x} \mathbf{m} \mathbf{a} \mathbf{b}_{m,r,c^*} \left( \sum_{j,m'} \alpha_{r,c^*} \mathbf{K}_{m,m'} \mathbf{F}_{m',j} \mathbf{Z}_{r,c,j} \right) - \frac{1}{2} \lambda \left( \sum_{m,j} \alpha_{r,c^*} \mathbf{F}_{m,j} \mathbf{Z}_{r,c,j} \right) . \end{aligned}$$

Fixing all other variables, the problem of minimizing the loss with respect to  $\varphi_{c^*}$  can then be understood as a quadratic programming problem.

$$\begin{aligned} \min_{\varphi_{c^*}} \quad & \frac{1}{2} \varphi_{c^*}^T \Gamma \varphi_{c^*} - \varphi_{c^*}^T \phi \\ \text{subject to} \quad & \varphi_{c^*} \geq 10^{-10} \end{aligned}$$

716 This problem is low dimensional and easy to solve using an off-the-shelf package. We use  
 717 `scipy.optimize.nnls`.

## 718 I.1 Selecting $\omega$

719 So far, we have assumed that  $\omega$  used in Equation 1 is a user-provided parameter. This  $\omega$  repre-  
 720 sents the the maximum tolerated reconstruction error. There are three methods for choosing this  
 721 parameter which we can suggest:

722 **Interactively.** If the observation model is correct, the predicted values of  $\mathbf{X}$  should be ‘close’  
 723 to the observed values. To discern this, our package provides an interactive method for  
 724 selecting an  $\omega$  which is satisfactory. This function starts with very large error tolerance  
 725 (specifically we take  $\omega$  to be half the maximum observed intensity. The function then allows  
 726 the user to visually compare the true observations with the predicted values estimated with  
 727 this value of  $\omega$ . If the predicted values appear to miss important features of the observed  
 728 data, the user can then reduce  $\omega$ . The optimization will be re-run (warm-starting from the  
 729 old initial condition, so this does not require very much time), and new predicted values are  
 730 displayed. The function allows the user to continually reduce  $\omega$  until the user deems that  
 731 all the important features of the observed data are captured by the predicted values.

732 **Automatically.** An automatic method can be achieved by starting with the original data, slightly  
 733 blur it, and take average squared magnitude of the difference. This magnitude can be used  
 734 to estimate the amount of speckle noise in the image. We can then choose  $\omega$  so that the  
 735 average reconstruction loss at each voxel is less than twice the value of this speckle noise.

736 **Via manually-labeled data.** If the user is willing to annotate a portion of their data with their  
737 beliefs about which colonies are located at which positions, this annotated data can be used  
738 to select the  $\omega$ . Specifically, one can select the  $\omega$  which yields the most accurate colony  
739 detection.

740 In practice, we find that the interactive method is the most straightforward to use.