



HAL
open science

DeepSpot: a deep neural network for RNA spot enhancement in smFISH microscopy images

Emmanuel Bouilhol, Anca Flavia Savulescu, Edgar Lefevre, Benjamin Dartigues, Robyn Brackin, Macha Nikolski

► **To cite this version:**

Emmanuel Bouilhol, Anca Flavia Savulescu, Edgar Lefevre, Benjamin Dartigues, Robyn Brackin, et al.. DeepSpot: a deep neural network for RNA spot enhancement in smFISH microscopy images. Biological Imaging, 2022, pp.1-18. 10.1017/S2633903X22000034 . hal-03659677

HAL Id: hal-03659677

<https://hal.science/hal-03659677v1>

Submitted on 5 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives licence (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is unaltered and is properly cited. The written permission of Cambridge University Press must be obtained for commercial re-use or in order to create a derivative work.

This peer-reviewed article has been accepted for publication but not yet copyedited or typeset, and so may be subject to change during the production process. The article is considered published and may be cited using its DOI.

10.1017/S2633903X22000034



Biological Imaging (2021), 1: 1–18

doi:10.1017/blg.2021.xx



RESEARCH ARTICLE

DeepSpot: a deep neural network for RNA spot enhancement in smFISH microscopy images

Emmanuel Bouilhol^{1,2*} , Anca Flavia Savulescu³, Edgar Lefevre², Benjamin Dartigues², Robyn Brackin⁴ and Macha Nikolski^{1,2*} 

¹Université de Bordeaux, CNRS, IBGC, UMR 5095, Bordeaux, 33000, France

²Université de Bordeaux, Bordeaux Bioinformatics Center, Bordeaux, 33000, France

³University of Cape Town, IDM, Faculty of Health Sciences, Cape Town, 7935, South Africa

⁴Advanced Medical Bioimaging CF, Charité – Universitätsmedizin, Berlin, 10117, Germany

*Corresponding author. Email: emmanuel.bouilhol@gmail.com; macha.nikolski@u-bordeaux.fr

Received: April 6, 2022

Keywords: smFISH image analysis, mRNA spot detection, Deep learning

Abstract

Detection of RNA spots in single-molecule FISH microscopy images remains a difficult task, especially when applied to large volumes of data. The variable intensity of RNA spots combined with the high noise level of the images often requires manual adjustment of the spot detection thresholds for each image. In this work, we introduce DeepSpot, a Deep Learning-based tool specifically designed for RNA spot enhancement that enables spot detection without the need to resort to image per image parameter tuning. We show how our method can enable downstream accurate spot detection. DeepSpot's architecture is inspired by small object detection approaches. It incorporates dilated convolutions into a module specifically designed for Context Aggregation for Small Object (CASO) and uses Residual Convolutions to propagate this information along the network. This enables DeepSpot to enhance all RNA spots to the same intensity, and thus circumvents the need for parameter tuning. We evaluated how easily spots can be detected in images enhanced with our method by testing DeepSpot on 20 simulated and 3 experimental datasets, and showed that accuracy of more than 97% is achieved. Moreover, comparison with alternative deep learning approaches for mRNA spot detection (deepBlink) indicated that DeepSpot provides more precise mRNA detection. In addition, we generated smFISH images of mouse fibroblasts in a wound healing assay to evaluate whether DeepSpot enhancement can enable seamless mRNA spot detection and thus streamline studies of localized mRNA expression in cells.

Impact Statement

Our paper introduces DeepSpot, a Deep Learning-based tool specifically designed to enhance RNA spots which enables downstream spot detection without the need to resort to image per image parameter tuning. DeepSpot's architecture is inspired by small object detection approaches by integrating dilated convolutions into a module specifically designed for Context Aggregation for Small Object (CASO) and using Residual Convolutions to propagate this information along the network.

© The Author(s) 2021. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. Introduction

Single cell microscopy together with RNA single molecule fluorescence in-situ hybridization (smFISH) technologies allow gene expression profiling at subcellular precision for determining molecular states of various cell types⁽¹⁾ and that in high-throughput fashion⁽²⁾. The repertoire of mRNA expression quantification methods is large and includes smFISH, clamp-FISH, amp-FISH and multiplexed versions such as e.g. MerFISH, all allowing the localization of RNA at sub-cellular level. There are technological differences between these methods in terms of number of detected RNAs and number of processed cells, however all produce imaging data with mRNA spots that can be further matched to the spots' x, y coordinates. With such increased image acquisition automation and the consequent growing number of high-throughput projects focused on spatially resolved transcriptomics, the need for automated and highly accurate detection of mRNA spots in fluorescent microscopy images has become increasingly important.

Single Molecule FISH (smFISH) is a method for visualizing individual RNA transcripts in fixed cells. smFISH is based on targeting RNA molecules with a set of 24-48 oligonucleotide probes, each individually labeled with one fluorophore. The combined fluorescence intensity level obtained from this high number of probes makes each RNA transcript visible as a spot that can be computationally identified and quantified⁽³⁾.

Despite the progress made in recent years, it is still difficult to detect the localization of spots corresponding to different mRNAs in a fully automated manner. First, background intensity is often irregular due to various factors, including autofluorescence that can be caused by intrinsic cell properties and fixative-induced fluorescence phenomenon⁽⁴⁾ or growth medium and buffers. An additional contributor to the background noise is off-target binding of probes, which depends on a various number of parameters, including the length of the transcript, its sequence, the cell type used and others. Second, spot detection is affected by the non-homogeneous intensity distribution and indistinct spot boundaries relative to the background. Moreover, FISH images may have a low signal-to-noise ratio (SNR). Additionally the boundary between background (noise) and signal (spots) is usually not evident⁽⁵⁾.

The main drawback of classical mRNA spot detection methods is the requirement of a strong human input to determine the best parameters to handle variable image to image properties such as SNR and presence of artifacts. Even small differences in these characteristics lead to the necessity for parameter fine-tuning⁽⁶⁾. Other than being time-consuming, the quality of detection largely depends on the capacity of the user to correctly choose the method's parameters according to each image properties (contrast, spots, artifacts, noise). Some recent deep-learning based approaches for mRNA spot detection try to circumvent this limitation, such as deepBlink⁽⁷⁾.

Here we introduce DeepSpot, a Convolutional Neural Network method dedicated to the enhancement of fluorescent spots in microscopy images and thus enabling downstream mRNA spot detection by conventional widely used tools without need for parameter fine-tuning. With DeepSpot we show that it is possible to avoid the manual parameter tuning steps by enhancing the signal of all spots so that they have the same intensity throughout all images regardless to the contrast, noise or spots shape. All the code as well as the pretrained model is available on GitHub and a plugin for the image analysis tool Napari, is also distributed <https://github.com/cbib/DeepSpot>.

DeepSpot gives a new twist to the ResNet network architecture and learns to automatically enhance the mRNA spots, bringing them all to the same intensity. In parallel, a multi-network architecture is integrated, trained by minimizing the binary cross-entropy (BCE) while providing context for mRNA spots thanks to the atrous convolutions. We evaluated the impact of the spot enhancement on the downstream mRNA spot detection, by performing spot detection using ICY with fixed parameters on both simulated images and experimental images manually annotated. Moreover, we compared the quality of mRNA spot detection from images enhanced by DeepSpot with deepBlink, and have shown that our method achieves greater generalization to effectively handle full variability of smFISH data. Finally,




Figure 1. RNA spots on a noisy background (image A). Spots' intensity is increased after the enhancement by e (image B).

to illustrate the end-to-end use of DeepSpot in projects where detecting sub-cellular localization with high precision is essential, we generated smFISH images from mouse fibroblasts in a wound healing assay, where enrichment of expression of β -Actin towards the location of the wound is expected in the migrating 3T3 mouse fibroblasts.

2. Related Work

Methodologically mRNA spot detection can be related to the detection of small objects topic in image analysis. The goal of spot detection is to find small regions in an image that differ from the surroundings with respect to certain properties, such as brightness or shape, more precisely, regions with at least one local extremum⁽⁸⁾. Spots can be considered as a particular case of more or less circular objects of small extent. Object detection has been one of key topics in computer vision which goal is to find the position of objects. However, small object detection, such as mRNA spots, remains difficult because of low-resolution and limited pixels⁽⁹⁾.

In this work we propose a deep learning network inspired by small object detection approaches for mRNA spot enhancement and we show how it can enable the downstream accurate detection of spots.

2.1. mRNA spot detection

In the case of FISH images, mRNA spots are small, compact and smaller than the resolution limit of the microscope⁽¹⁰⁾, therefore images of mRNA spots correspond to the maximum intensity pixel surrounded by the diffraction of the fluorescence signal defined by the Point Spread Function (PSF) which can be modeled by a Gaussian within small radius disk (see Figure 1, image B). This radius depends on several imaging parameters and optical properties of the microscope such as the diffraction limit or the excitation state of the fluorophore.

While there is no universal solution to the detection of small objects such as mRNA spots in fluorescent cellular images, a large plethora of work is available on the subject. A number of approaches have gained wide popularity thanks to the development of software tools embedding the algorithms and providing user with a graphical interface. In particular, ImageJ/Fiji^(11,12) is widely used, largely due to the plugin-based architecture, recordable macro language and programmable Java API. Another popular tool is CellProfiler⁽¹³⁾, based on similar paradigms and FISH-Quant⁽¹⁴⁾. A more recent platform, ICY⁽¹⁵⁾, also provides the possibility to develop new algorithms as well as a user interface for image

analysis, including the ICY spot detector for the detection of mRNA spots, method based on wavelet transform decomposition⁽¹⁰⁾.

Deep-learning networks have been introduced for mRNA spot detection^(16,17), and more recently deepBlink⁽⁷⁾. The latter focuses on a fully convolutional neural network based on the U-Net architecture. deepBlink not only provides the code, but also annotated smFISH data and implements a threshold-independent localization of spots.

2.2. Detection and enhancement of small objects

Consistently with the mRNA spot detection difficulty, detection of small objects remains a challenging part of the general object detection problem due to the limited information contained in small regions of interest. For instance, it has been shown that the object size has a major impact on the accuracy of Deep Learning object detection networks such as VGG, ResNet or Inception V2⁽⁹⁾. Indeed, small objects do not contain sufficient semantic information⁽¹⁸⁾ and thus the challenge is to capture semantic features while minimizing spatial information attenuation⁽¹⁹⁾.

Expectedly, adding more context improves the detection of small objects^(19–21). An elegant solution is to use the dilated convolution (a.k.a. atrous convolution), because the receptive field can be expanded without loss of resolution and thus capture additional context without loss of spatial information^(22,23).

Of particular interest to our work is the signal enhancement, an image processing technique aiming to reinforce the signal only in those regions of the image where the actual objects of interest are, and potentially to weaken the noise or the signal from other structures⁽²⁴⁾. In our case, the objects of interest are mRNA spots.

Image enhancement is the transformation of one image X into another image $e(X)$ (see Figure 1). Pixel values (intensities) of X at spot locations are modified according to the transformation function e , the resulting pixel values in image $e(X)$ measuring the certainty of mRNA presence at that position. Thus, $e(X)$ can be considered as a probability map that describes possible mRNA spot locations⁽²⁴⁾.

Small object enhancement has been developed in other fields than mRNA spots in fluorescent imaging, such as in astronomy to enhance stars or galaxies over the cosmic microwave background⁽²⁵⁾ or in the biomedical imaging to facilitate the human detection of larger spots such as nodules⁽²⁶⁾. In the microscopy field, Deep-STORM⁽²⁷⁾ propose to enhance the resolution of Single Molecule Localization Microscopy (SMLM) mRNA spots, with deep learning. However, these images don't have the same characteristics as smFISH data in terms of noise or signal. For example, human nodules are much larger objects than typical mRNA spots. As for the star enhancement method proposed by⁽²⁵⁾, it is not suited for low intensity spots⁽²⁸⁾, which is a major concern in spot detection for smFISH images.

3. Materials and Methods

3.1. Materials

3.1.1. smFISH datasets with alternatively established spots' localization

We constructed a dataset DS_{exp} of 1553 images from the experimental smFISH data acquired in⁽²⁹⁾, by applying 256×256 pixels patches to better fit in the GPU memory (see Table 1). The authors have performed spot detection using image analysis techniques such as local maximum detection, implemented in the BIG-FISH pipeline⁽³⁰⁾. Along with the scripts, the authors provided a list of 57 parameter combinations that were used to detect the spots in 57 different image acquisition series of 32 genes used in their study. We ran the pipeline with these parameters and performed an additional manual curation to keep patches with number of spots between 10 and 150 and remove those with a visually obvious over- or under-identification of spots. The resulting DS_{exp} dataset of 1553 images representing 27 different genes is thus experimentally generated and guarantees high confidence in the ground truth annotation of mRNA spot coordinates.

We have also downloaded the dataset DS_{dB} from the deepBlink publication⁽⁷⁾, composed of 129 smFISH images acquired from 4 different cell culture conditions of the HeLa 11ht cell line. The authors provided their annotation of mRNA spots' locations which was performed using TrackMate and curated by experts.

To estimate the variability of spot intensity, we also calculated the Coefficient of Variation (CV) of spot intensities for each gene of DS_{exp} as well as for the entire DS_{exp} and DS_{dB} datasets (Supplementary Figure S1). While for most images and genes, the CV ranges from 0.25 to 0.50, some have a very high CV. The global CV for all the spots of all the images in the DS_{exp} dataset is 0.72 and 0.22 for DS_{dB} .

3.1.2. Novel smFISH dataset from a wound healing assay

To evaluate whether DeepSpot enables precise mRNA spot detection in a biological context, we made use of the wound healing assay⁽³¹⁾ to generate a second experimental dataset DS_{wound} . In the wound healing assay, migrating cells, such as fibroblasts, are grown on a coverslip and serum starved for synchronization. A scratch in the middle of the coverslip is then generated, mimicking a wound, followed by induction of the cells to polarize and migrate towards the wound, to generate wound closure, done using replacement of serum starved medium with 10 % FBS-containing medium. We used 3T3 mouse fibroblasts in a wound healing assay, followed by cell fixation. Fixed samples were taken for single molecule FISH experiments to visualize and quantify β -Actin mRNA and imaged on a custom built Nikon Ti Eclipse widefield TIRF microscope. β -Actin has been previously shown to be enriched in neuronal growth cones of extending axons, as well as the leading edges of migrating cells and this enrichment has typically been associated with cell polarity and neuronal plasticity^(32–34). Based on this, we hypothesized that β -Actin would be enriched in the leading edge of migrating 3T3 fibroblasts. The dataset is composed of 96 images, 48 images of non migrating cells (control) and 48 images of migrating cells. Each image was divided into 4 patches, yielding a total of 384 patches of 256×256 pixels size.

3.1.3. Simulated and hybrid datasets

In addition to the experimental datasets, we have built 20 simulated datasets with images 256×256 pixels of width and height, same as the patches of DS_{exp} . Briefly, the background was generated by a combination of Poisson noise and Perlin noise with a random intensity between [80, 150]. Elastic transformations were added to this noise to approximate the variety of textured background noise in the experimental images. Spots were generated as circles, randomly placed in the image, and then convolved with a Gaussian function that approximates PSF. Their size randomly ranges from 4 to 9 pixels in diameter, including Gaussian smoothing.

Two different types of simulated datasets were generated DS_{fixed} and DS_{var} , each containing 10 datasets defined according to the number of spots per image $i \in \{[10..100] \bmod 10\}$, see Table 1. For example, in the DS_{fixed}^{20} dataset each image contains 20 spots and in the DS_{var}^{70} dataset each image contains 70 spots. Each image in the DS_{fixed} dataset has the same fixed spot intensity for all spots randomly chosen in the interval [160, 220], while in the DS_{var} dataset the spot intensity is randomly chosen from the same interval for each spot, resulting in images with variable spot intensity.

In addition to the experimental and simulated datasets, we have built a hybrid dataset DS_{hybrid} where the experimental data from DS_{exp} is augmented by appending an additional 25% of simulated images generated with both variable spots' intensity and variable spots' number per image, within the [160, 220] and [10, 100] intervals, respectively. The intensity of the spots is calculated with respect to the intensity of the background noise, so that the generated images have a SNR between 10 and 40. These values correspond to the minimum and median SNR values in our experimental images respectively (Supplementary Figure S2).

3.1.4. Ground truth

Since the goal of our network is to learn to transform an image X into $e(X)$, where intensity at spots' location is enhanced, the training set has to be provided with the enhanced counterpart of each image

Table 1. List of datasets used for training (22 training datasets) and evaluation (21 test datasets) of the DeepSpot network. Images in the DS_{var}^i datasets have spot intensities between 160 and 220 for each image. For the fixed intensity datasets DS_{fixed}^i the spot intensity is set to one value within $[160..220]$ for a given image, but varies from image to image. 10 variable intensity DS_{var}^i and 10 fixed intensity DS_{fixed}^i datasets are named according to the number of spots in the images, i . The dataset DS_{hybrid} combines DS_{exp} with 25% simulated images.

Datasets	#spots per image	intensity	#images train	#images test	#datasets train/test
DS_{exp}	$10 \leq i \leq 150$	variable	1453	100	1/1
DS_{var}^i	$i \in \{[10..100] \bmod 10\}$	variable	2000	400	10/10
DS_{fixed}^i	$i \in \{[10..100] \bmod 10\}$	fixed	2000	400	10/10
DS_{hybrid}	$10 \leq i \leq 150$	variable	2000	NA	1/NA

in the training set. That is, training sets include pairs of images $\langle X, e(X) \rangle$ where e is the procedure that is used to produce ground truth enhanced images: for each spot (Figure 1, image A), the ground truth enhancement procedure e is applied at the spots' locations $A(X)$, resulting in images where spots are enhanced as shown in Figure 1, image B.

In this work we implement e as a kernel of 3×3 pixels at all locations where the spots were annotated in the experimental dataset DS_{exp} or generated for DS_{fixed} , DS_{var} and DS_{hybrid} (see Table 1). The kernel has the same pixel values for all the spots, in order to drive the network to learn to enhance all spots up to the same level of intensity, regardless of the initial intensity in the acquired data. The enhancement kernel of DeepSpot is smaller than the smallest spot size in our datasets, therefore it is not expected to augment the size spot. This is particularly important for spatially close spots. Moreover, background is kept the same between X and $e(X)$ so the transformation does not affect the background.

3.2. Method

In this section, we present the DeepSpot enhancement network in detail. We first overview the network architecture and then we discuss the custom loss function.

3.2.1. Network architecture

DeepSpot network is composed of two main components as presented in Figure 2. The first component is a multi-path network shown in panel A. The second component is an adapted residual network as shown in panel B.

Context Aggregation Module

As pointed out in⁽²²⁾, finding small objects is fundamentally more challenging than large objects, because the signal is necessarily weaker. To solve this problem in the context of mRNA spot detection, we developed a new module, that we called the Context Aggregation for Small Objects module (CASO).⁽³⁵⁾ demonstrated that using image evidence beyond the object extent (context) always enhances small object detection results, we therefore developed our CASO module to aggregate context around the mRNA spots. The CASO module is a multi-path network as shown in panel A. It takes the input image and processes it along three different paths each with different types of convolution blocks to collect specific information from the input image. Each path contains 3 convolution blocks.

1. The first path is composed of traditional convolution blocks (2D convolution, batch normalization, activation and max pooling). These blocks, often used in CNNs, are particularly efficient to reinforce the semantic information at the expense of spatial information.
2. The second path uses only 2D convolutions, batch normalization and activation. As Max Pooling is known to keep mostly the maximum intensities in images, some of the faint spots




Figure 2. DeepSpot network architecture is composed of the Context Aggregation for Small Objects module (CASO) constituted of a multi-path network (panel A) and a customized ResNet component (panel B). A custom loss function is used for training the network.

- may be eliminated during the max pooling operation. In this path we used strided 2D convolutions instead of a max pooling layer, to keep the information of low intensity of spots. To make sure to end up with the same receptive field as the first path, we set the stride to 2.
3. The third path makes use of the atrous convolution pooling⁽²²⁾, implemented as a 2D convolution with the dilatation rate of 2. The following layers are batch normalization, activation and max pooling.

The CASO module is a multi-path neural network and can learn more comprehensive and complementary features than a single path. In particular, the goal of the atrous convolution is to bring more context around the small spots (see section 2.2), while the two other paths aggregate the semantic information of the bright spots and faint spots for the first and second path respectively. The results of the three encoding paths are then concatenated to construct a longer feature vector containing information extracted by each path. For all convolutional blocks, the activation function is the Rectified Linear Unit (ReLU). The number of filters for the 2D convolutions in the are 32, 64, 128 for the first, second and third block, respectively.

Custom ResNet

For the second component (panel B) we customized the ResNet architecture to create a residual neural network composed of ten consecutive convolutional residual blocks (ResBlock), using full pre-activation blocks described in⁽³⁶⁾, where the authors suggested that the better results obtained by the full pre-activation blocks are due to the pre-activation by the batch normalization that improves regularization of the model due to the fact that the inputs to all weight layers have been normalized. Each ResBlock is composed of three sub-blocks as presented in Figure 3. A sub-block is constituted of a batch normalization followed by an activation (ReLU) and a 2D convolution. After the 3 sub-blocks a spatial dropout layer with rate 0.2 is applied. Each ResBlock ends by the residual connection.

To obtain an output image with the same size as the input image, we used a particular type of up convolutional blocks. Recently,⁽³⁷⁾ demonstrated that the use of Up Residual Blocks (UpResBlocks)




Figure 3. Full pre-activation residual block, composed of Batch normalisation, Activation, Convolution, repeated three time before Dropout and residual connection.

instead of classic up-convolution blocks improves the performance of generative networks by preserving the effective features from the low dimensional feature space to the high dimensional feature space. Our decoding path is composed of three UpResBlocks and reconstitutes an output with the same size as the input while propagating low dimension feature information of the enhanced spots from the custom ResNet to the last layer. Each UpResBlock is constituted of three sub-blocks containing a 2D transposed convolution, batch normalization and activation (ReLU). The three sub-blocks are followed by a spatial dropout layer with rate 0.2. UpResBlocks then end with a residual connection. A sigmoid activation function is applied to the last convolution, so that all the pixels have values in the $[0, 1]$ interval. The final image is obtained by normalizing the pixel intensities between 0 and 255.

3.2.2. Loss

We defined our custom loss function as a combination of binary cross-entropy (BCE) and Mean Squared Error (MSE) functions. The main term of the loss function is the BCE loss, defined by $\mathcal{L}_{\text{BCE}}(x, \hat{x}) = -(x \log(\hat{x}) + (1 - x) \log(1 - \hat{x}))$ that measures the difference between the images predicted by the network \hat{x} and the ground truth images x . While mostly used for classification, it can also be used for segmentation and enhancement due to its performance for pixel-level classification⁽³⁸⁾.

To this main \mathcal{L}_{BCE} term we added a regularization term defined by Mean Squared Error, $MSE = \frac{1}{n} \sum_{i=1}^n (\max(x_i) - \max(\hat{x}_i))^2$ that is computed between the maximum value of the predicted image $\max(\hat{x}_i)$ and the maximum value of the ground truth image $\max(x_i)$. This regularization drives the network to produce spots whose intensity is close to 255 (see Table 2), and therefore standardizes the signal enhancement intensity in the output images, which in its turn facilitates the downstream automatic detection of the spots. The total loss function is $\mathcal{L}_{\text{BCE}} + \mathcal{L}_{\text{MSE}}$.

4. Results

We trained the DeepSpot network on our 20 simulated training datasets DS_{fixed} and DS_{var} as well as on the experimental and hybrid datasets DS_{exp} and $\text{DS}_{\text{hybrid}}$, resulting in 22 models M_{fixed} , M_{var} , M_{exp} and M_{hybrid} . Training parameters were optimized with HyperOpt algorithm⁽³⁹⁾ and the ASHA scheduler⁽⁴⁰⁾. The best configuration obtained and used for further trainings had the learning rate of 0.0001, the dropout rate of 0.2 and the batch size of 32 and 128 filters per convolution.

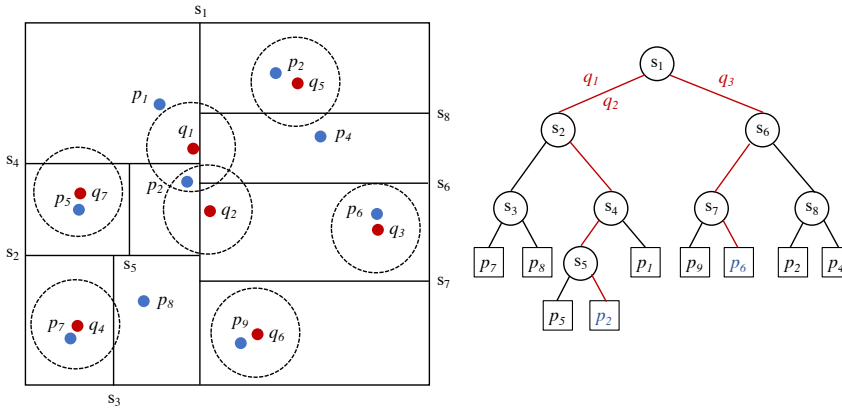


Figure 4. Spot matching by 1-neighbor k -d tree between the detected mRNA spots $\{p_1, \dots, p_9\}$ depicted in blue and annotated spots $\{q_1, \dots, q_7\}$ depicted in red. The k -d tree construction for $\{p_1, \dots, p_9\}$ is shown on the left. Using the matching radius depicted by circles, the k -d tree queries for q_1 and q_2 shown in red, lead to the same leaf p_2 and correspond to an ambiguous match, while query for q_3 leads to a unique match. mRNA spots p_1, p_4 and p_8 are the False Negatives.

Each of the resulting 22 models was evaluated on the 21 test data sets from Table 1, yielding 462 enhanced test datasets (8100 images total). To assess whether DeepSpot enhancement enables easy spot detection, we applied the ICY spot detector⁽¹⁰⁾ to the images enhanced by different models. Moreover, we defined an unique set of ICY parameters that matches the shape and intensity of the enhancing kernel of the DeepSpot network: scale 3, sensitivity 20 and scale 7, sensitivity 100. We then evaluated whether the detected spots from the enhanced images matched well with the annotated ground truth of spots' coordinates.

4.1. Evaluation procedure

We denote by $D(X) = \{p_1, \dots, p_n\}$ the point pattern detected by ICY from the enhancement of of an image X by DeepSpot and by $A(X) = \{q_1, \dots, q_m\}$ the ground truth annotation of spots coordinates. Notice that m is not necessarily equal to n , corresponding to under- or over-detection and even for a well-detected spot, the coordinates in A and D may slightly differ. To account for these remarks, we used the k -d tree algorithm⁽⁴¹⁾ to query the detection D for nearest neighbors in A as proposed in^(42,43). Number of neighbors was set to 1 and matching radius t to 3 (coherent with the enhancement kernel for ground truth images).

This allows to establish a matching for all annotated points under $t = 3$ and thus also defines the number of False Negatives or False Positives corresponding to the missing matches from $D(X)$ or $A(X)$, respectively. True Negatives are defined by all pixels p of the confusion matrix such that $p \in X \setminus \{A(X) \cup D(X)\}$. However, given that $|X| \gg |X \setminus \{A(X) \cup D(X)\}|$ implies inflated TN values, this makes measures such as accuracy, AUC and ROC curve irrelevant.

The drawback of matching m versus n points is the possibility of ambiguous matching. With the k -d tree approach it happens when two points $q_i, q_j \in A$ can match to one $p_k \in D$ (see Figure 4). This can happen if annotated spots q_i, q_j are close and the detected matching point for both of them, p_k , lies within the same distance t , which can correspond to an over enhancement and thus blurring between the two spots in the enhanced image. While alternative solutions such as the Linear Assignment Problem can be used, they do not avoid the problem of matching two different numbers of point. The k -d tree approach has the advantage to keep the ambiguous matches explicitly to measure this effect. We thus also report the number of ambiguous matches (AM).

4.2. DeepSpot enhances the mRNA spot signal to the same intensity

To avoid the manual selection of the detection threshold, it is imperative to have a homogeneous spot intensity for the whole dataset in order to use unique set of parameters for all images. Table 2 summarizes the intensities obtained after enhancement by DeepSpot for each category of datasets described in Table 1 (experimental, hybrid, simulated with variable and fixed intensities).

Table 2. Spot enhancement performance in terms of resulting spot intensity. The measures displayed correspond to the spot intensity between [0,255] after enhancement by the neural network and averaged by category of models and datasets. Between brackets are shown the 95% confidence intervals. Models categories are listed in rows, while columns correspond to the dataset categories on which the different models were applied.

Test \ Train	DS _{exp}	DS _{fixed}	DS _{var}
M _{exp}	249.55 [249.29, 249.82]	250.81 [250.68, 250.93]	247.07 [246.68, 247.45]
M _{hybrid}	241.15 [240.79, 241.52]	248.71 [248.57, 248.85]	243.3 [243.14, 243.45]
M _{var}	248.92 [248.47, 249.37]	251.99 [251.93, 252.05]	254.57 [254.53, 254.62]
M _{fixed}	242.93 [242.16, 243.71]	254.73 [254.71, 254.75]	249.33 [249.09, 249.57]

As expected, intensities were closer to 255 when training and test datasets belong to the same category. For example, models trained on data with fixed intensities M_{fixed} and applied to data with fixed intensities DS_{fixed} produced enhanced spot intensities very close to 255. Similarly enhancement close to 255 could be observed when models M_{var} were evaluated on DS_{var}. Of particular interest are the enhancement results from the M_{exp} training, that are close to 250 for every dataset, experimental or simulated. Hybrid model enhances intensities to 241 on the experimental dataset. In general, the enhanced spot intensities were between 241 and 255, representing a variation of only 5.4% from the maximum intensity which is sufficient to fully separate smFISH spots from the background in the enhanced images.

Table 3. Models' performance per model type. Metrics (F1-score, precision, recall, Ambiguous Matches (AM)) were calculated by averaging the values obtained for each image of the 21 test datasets. Top values in cells correspond to the mean value, bottom values between brackets show the 95% confidence interval. Best values are highlighted in bold.

Model	F1-Score(%)	Precision(%)	Recall(%)	AM(mean/img)
M _{exp}	87.8 [87.6, 88.1]	99.6 [99.5, 99.7]	80.1 [79.8, 80.5]	2.7 [1.4, 4.0]
M _{hybrid}	97.3 [97.2, 97.3]	99.3 [99.2, 99.4]	95.5 [95.4, 95.6]	0.94 [0.56, 1.32]
M _{var}	94.3 [94.3, 94.4]	99.5 [99.5, 99.5]	91.4 [91.3, 91.5]	1.19 [1.03, 1.35]
M _{fixed}	91.6 [91.5, 91.7]	99.9 [99.9, 99.9]	85.7 [85.6, 85.8]	2.05 [1.77, 2.32]

4.3. DeepSpot enables accurate mRNA spot detection

Summary statistics of performance of each model type are shown in Table 3 including the mean F1-score, precision, recall and ambiguous matches, with 95% confidence interval. For a given model M each metric was computed for all enhanced images (8100 total). The mean metric value \bar{x} and 95% confidence interval C were then calculated separately for each model type. Due to the high prevalence of True Negatives, instead of the Accuracy measure, we calculated the F1-Score, which gives an indication of the model accuracy with a better balance between classes than the actual accuracy measure, by not including True Negatives. We compared the F1-scores for each of the 14 genes present in the test dataset, no major difference in DeepSpot performance can be observed for these genes (Supplementary Figure S3). This consistency in performance across different SNR values (Supplementary Figure S2) shows that DeepSpot has the capacity to enhance spots in images with varying characteristics.

The results in Table 3 indicate that the number of FP is very low for all models, given that both precision and recall are high. Importantly, M_{hybrid} has shown best overall performance in terms of precision, recall and F1-score, thus indicating that the mRNA spot enhancement by M_{hybrid} leads to the least FP and FN counts in the downstream mRNA spot detection.




Figure 5. Heatmap of the F1-scores obtained by each of the 22 models when evaluated on the 21 test datasets described in Table 1.

Figure 5 shows the mean F1-scores for each of the 22 models evaluated on the 21 datasets. This heatmap indicates that (i) models trained on the DS_{fixed}^i training sets perform better on the corresponding DS_{fixed}^i test sets rather than on variable intensity test sets, (ii) models trained on DS_{var}^i training sets show better performance on the corresponding test sets rather than on fixed intensity test sets. It also shows that M_{var} models are globally better than M_{fixed} models. A plausible hypothesis is that training on

variable intensities makes models better at generalizing on other data. Finally, M_{hybrid} is the model that has the best overall performance, including the experimental dataset. Again, the diversity of training data drives this model to be more robust to newly encountered data.

More generally, given that the F1-score is above 90% for all the models, we can conclude that the architecture of the DeepSpot neural network is particularly suited for the task of mRNA spot enhancement.

4.4. DeepSpot enables more accurate spot detection compared with deepBlink

The main objective of our DeepSpot method is to circumvent the parameter fine-tuning and enable the downstream spot detection with an unique parameter set. As such, this objective fits well with the one that the authors of deepBlink⁽⁷⁾ have pursued, despite the fact that the latter proposes a new spot detection method, while our goal is to fit a spot enhancement step into commonly used workflows. Consequently, deepBlink constitutes a relevant comparison target.




Figure 6. Examples of the results obtained with DeepSpot and deepBlink on the experimental test datasets DSexp (first row) and deepBlink (second row), and the simulated spots with fixed (DSfixed) and variable (DSvar) intensity datasets for the third and fourth rows, respectively. Colored circles indicate where the spots were detected by DeepSpot and deepBlink (blue and green, respectively). In the ground truth column, pink circles indicate the spots that were previously annotated as ground truth by alternative methods. The last two columns show the magnification at the positions indicated by the colored rectangles for DeepSpot and deepBlink, respectively.

We compared the accuracy of spot detection by the model M_{dB} made available on deepBlink associated GitHub with that of DeepSpot when trained on hybrid data M_{hybrid} , both on our datasets as well as on the dataset provided by the authors of deepBlink, DS_{dB} . Table 4 shows the F1-scores for each dataset

category. It should be noted that deepBlink's precision is close to 99% while its recall is very low for certain datasets (Supplementary Figure S4), meaning that the main drawback of deepBlink is in terms of false negatives (missing true spots). We illustrate this on Figure 6 where we provide a comparison of deepBlink and DeepSpot on both experimental and simulated data example.

Table 4. Models' performance for deepBlink and DeepSpot for smFISH spot detection. Overall F1-scores are calculated by averaging the values obtained for each image of the test datasets corresponding to each dataset category. Top values in each cell correspond to the mean value, bottom values between brackets show the 95% confidence interval. Best values are highlighted in bold.

Model \ F1-score(%)	DS _{exp}	DS _{var}	DS _{fixed}	DS _{dB}
M _{dB}	78.7 [75.26, 82.14]	47.97 [46.86, 49.07]	67.4 [66.28, 68.51]	94.2 +- 0.141(std)
M _{hybrid}	94.43 [93.41, 95.45]	96.67 [96.58, 96.75]	98.04 [97.96, 98.12]	87.92 [85.36, 90.49]

DeepSpot clearly outperformed deepBlink on our datasets (Table 4). However, we found that M_{dB} performed better on experimental data than on simulated data, presumably because deepBlink model has only been trained on experimental data. M_{hybrid} results were consistent on all of our datasets. We have also applied our M_{hybrid} model to the smFISH test dataset DS_{dB} provided by the authors of deepBlink. Results reported in Table 4 for the detection of spots by M_{dB} from the images of the DS_{dB} dataset are those obtained by the authors and originally reported in⁽⁷⁾. Not surprisingly, deepBlink model performance is better on their own smFISH images, however DeepSpot managed to have an F1-score of nearly 88%, a noticeable achievement since the model has not been trained on the DS_{dB} data. Together, the results of Table 4 indicated that DeepSpot is a robust methodology that offers a generalist model for mRNA spot enhancement and ensures high quality downstream spot detection without parameter tuning.

4.5. DeepSpot's use in an end-to-end smFISH experiment

To evaluate whether DeepSpot can be effectively used in an end-to-end smFISH experiment, we have performed a wound healing assay in which cells migrate towards a wound to close it. To investigate whether β -Actin was enriched at the leading edges of 3T3 migrating fibroblast cells (see 3.1.1). The wound location was manually annotated as shown on a typical image example in panel A, Figure 7.

All the cell images were segmented by manually, yielding cell and nucleus masks. Wound location defined the cell migration direction as schematically shown in panel C, Figure 7. We partitioned the cell masks into three sections by computing 120° section centered at the nucleus centroid, and anchoring one of these sections as oriented towards the wound location S_{wound} at 60° angle to the left and to the right of the line between the nucleus centroid and the wound location. This cell segmentation allowed us to compute the normalised number of detected mRNA spots in the cytoplasmic part of each section S₁, S₂, S₃.

Using the DypFISH framework⁽⁴⁴⁾, we further compared the cytoplasmic mRNA relative density in S_{wound} (light blue) and in sections that are not oriented towards the wound (dark blue), as shown in Figure 7, Panel B. We used the Polarity Index⁽⁴⁴⁾, that measures the enrichment of mRNA in different sections. Briefly, the Polarity Index measures how frequently the relative concentration within the wound section is higher than in the non-wound section. The Polarity Index lies between [-1, 1], a positive value implying a wound-correlated enrichment of RNA transcripts while a negative values implies enrichment away from the wound and a value of zero implies no detectable enrichment.




Figure 7. Processing steps for the end-to-end wound healing assay. Panel A is a typical smFISH image of the wound healing assay is shown. Panel B shows the cell quantization procedure in three sections and the direction of the wound (red arrow). S_{wound} section is oriented towards the wound and is shown in light blue, other sections are in dark blue. Panel C represent how the cell and nucleus were manually segmented and mRNA spots were counted after enhancement within the cytoplasmic portion of each section. Panel D presents the Wound Polarity Index (WPI) of cytoplasmic mRNA transcripts in S_{wound} compared to other sections for the β -Actin RNA. WPI was calculated in migrating and non-migrating cells. The bars correspond to the median and the error bars to the standard deviation from the median for 100 bootstrapped WPI estimates.

β -Actin mRNA was highly enriched in the leading edge of migrating cells, whereas almost no detectable enrichment of β -Actin was found in the leading edge of control cells. This is in line with previously published data, showing enrichment of β -Actin in leading edges of migrating fibroblasts⁽⁴⁵⁾.

5. Discussion and conclusion

Recent FISH microscopy methods are capable of generating thousands of images, it has thus become imperative to introduce algorithms capable to streamline the detection of mRNA spots and in particular to avoid manual fine-tuning of numerous parameters.

5.1. Limitation

The impact of the non-specifically bound probes was not evaluated in this work. Consequently, even if DeepSpot is trained to distinguish between multiple signal intensity variations, in cases where the signal intensity of the non-specifically bound probes is too similar in shape and intensity to the signal of the specifically bound probes, then DeepSpot will not be able to enhance only spots corresponding to specifically bound probes. This issue, common to all detection methods deserves further investigation.

Moreover, it remains to be investigated whether DeepSpot would be suitable for more complex systems such as tissues and which adaptations in terms of its architecture and training data it would require. Moreover, even if DeepSpot architecture was designed to be easily extensible to 3D by changing the 2D convolution layers to 3D layers in Tensorflow, the performance of DeepSpot in 3D has not been evaluated.

5.2. Discussion

In this work we introduced DeepSpot, a novel CNN architecture specifically designed to enhance RNA spots in FISH images, thus enabling the downstream use of well known spot detection algorithms, such as the ICY spot detector, without parameter tuning. In particular, the architecture of our network introduces the Context Aggregation for Small Object module that relies on sparse convolution to provide more context for enhancement of small objects corresponding to mRNA spots. DeepSpot network has been trained and tested on 21 simulated datasets, all with different signal and noise characteristics, as well as on a previously published experimental dataset that was annotated for spot locations. We have shown (i) that our approach achieves better performance when the training is performed on data with highly variable intensity and (ii) that performing training on a combination of experimental and simulated data is a viable approach in real-life setting.

Furthermore, we compared the performance of combining DeepSpot and ICY to that of the state-of-the-art deep learning-based method deepBlink and have shown that on average, DeepSpot enables a substantially better detection of mRNA spots than deepBlink. We found that DeepSpot / ICY workflow provided excellent quality spot detection on the test datasets corresponding to the datasets on which it has been trained, with the average F1-score above above 97%, but also achieved high precision results on fully unknown datasets with the F1-score of 88% for the datasets provided with the deepBlink publication. Taken together, the good results on both known and unknown data indicate that DeepSpot is a more generalist model than deepBlink and that it achieves a good balance between overfitting and underfitting. We hypothesize that this generalization capacity is possibly due to both strong regularization within the network and the diversity of signal provided by the carefully constructed training data.

To evaluate how well our method is suited for end-to-end biological investigations, we have shown the efficiency of the DeepSpot model trained on the combination of experimental and simulated data in the context of an independent study of cell migration. We have performed single molecule FISH to detect β -Actin in mouse fibroblasts in a wound healing assay and enhanced the resulting images using our combination model, which allowed us to detect that the β -Actin mRNA enrichment is specific to leading edge of migrating cells as contrasted by its expression in non-migrating cells.

To conclude, we have shown that DeepSpot enhancement enables automated detection and accurate localization of mRNA spots for downstream analysis methods and can thus be a useful tool to streamline not only spot detection, but also studies of localized mRNA enrichment within cells.

Acknowledgements. We thank Dr. Arthur Imbert for sharing the experimental smFISH data and helping to run the BIG-FISH pipeline.

Competing Interests. None

Data Availability Statement. DeepSpot network along with the code for training and for mRNA spot enhancement is fully open source and available on GitHub at <https://github.com/cbib/DeepSpot>. Our pretrained models used in this study are also available on our GitHub page, as well as the Napari plugin.

Data for simulated images for DS_{var}, DS_{fixed} and DS_{hybrid} are available on Xenodo at <https://doi.org/10.5281/zenodo.5724466>

Hardware and framework. All the analysis presented in this paper are performed on 1 GPU, Tesla T4 with 16GB of memory, on a dedicated machine with two 2 CPU Intel Xeon Silver 4114 and 128Go RAM. Training curves and comparison with a vanilla U-Net architecture are available in the Supplementary Figure S5 and Supplementary Table 1. This work is implemented in Python 3.8 and we used TensorFlow 2.4 for creating and training the neural networks.

Ethical Standards. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

Funding Statement. This research received no specific grant from any funding agency, commercial or not-for-profit sectors.

Author Contributions. Emmanuel Bouilhol: Conceptualization, Methodology, Investigation, Software, Writing, Visualization Edgar Lefevre: Software, Validation Benjamin Dartigues: Software, Validation Robyn Brackin: Resources Anca Flavia Savulescu: Data Curation, Resources, Writing, Validation Macha Nikolski: Conceptualization, Methodology, Investigation, Writing, Supervision, Validation All authors approved the final submitted draft.

References

1. R. Ke, M. Mignardi, A. Pacureanu, J. Svedlund, J. Botling, C. Wählby, and M. Nilsson, "In situ sequencing for rna analysis in preserved tissue and cells," *Nature methods*, vol. 10, no. 9, pp. 857–860, 2013.
2. N. Battich, T. Stoeger, and L. Pelkmans, "Image-based transcriptomics in thousands of single human cells at single-molecule resolution," *Nature methods*, vol. 10, no. 11, pp. 1127–1133, 2013.
3. A. Raj, P. Van Den Bogaard, S. A. Rifkin, A. Van Oudenaarden, and S. Tyagi, "Imaging individual mrna molecules using multiple singly labeled probes," *Nature methods*, vol. 5, no. 10, pp. 877–879, 2008.
4. R. M. Rich, D. L. Stankowska, B. P. Maliwal, T. J. Sørensen, B. W. Laursen, R. R. Krishnamoorthy, Z. Gryczynski, J. Borejdo, I. Gryczynski, and R. Fudala, "Elimination of autofluorescence background from fluorescence tissue images by use of time-gated detection and the azadioxatriangulenium (adota) fluorophore," *Analytical and bioanalytical chemistry*, vol. 405, no. 6, pp. 2065–2075, 2013.
5. M. Zhang, *Small Blob Detection in Medical Images*. Arizona State University, 2015.
6. J. C. Caicedo, S. Cooper, F. Heigwer, S. Warchal, P. Qiu, C. Molnar, A. S. Vasilevich, J. D. Barry, H. S. Bansal, O. Kraus, et al., "Data-analysis strategies for image-based cell profiling," *Nature methods*, vol. 14, no. 9, p. 849, 2017.
7. B. T. Eichenberger, Y. Zhan, M. Rempfler, L. Giorgetti, and J. A. Chao, "deepblink: threshold-independent detection and localization of diffraction-limited spots," *Nucleic Acids Research*, vol. 49, no. 13, pp. 7292–7297, 2021.
8. T. Lindeberg, "Detecting salient blob-like image structures and their scales with a scale-space primal sketch: A method for focus-of-attention," *International Journal of Computer Vision*, vol. 11, no. 3, pp. 283–318, 1993.
9. J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, et al., "Speed/accuracy trade-offs for modern convolutional object detectors," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7310–7311, 2017.
10. J.-C. Olivo-Marin, "Extraction of spots in biological images using multiscale products," *Pattern recognition*, vol. 35, no. 9, pp. 1989–1996, 2002.
11. M. D. Abramoff, P. J. Magalhães, and S. J. Ram, "Image processing with imagej," *Biophotonics international*, vol. 11, no. 7, pp. 36–42, 2004.
12. J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid, et al., "Fiji: an open-source platform for biological-image analysis," *Nature methods*, vol. 9, no. 7, p. 676, 2012.
13. L. Kamentsky, T. R. Jones, A. Fraser, M.-A. Bray, D. J. Logan, K. L. Madden, V. Ljosa, C. Rueden, K. W. Eliceiri, and A. E. Carpenter, "Improved structure, function and compatibility for cellprofiler: modular high-throughput image analysis software," *Bioinformatics*, vol. 27, no. 8, pp. 1179–1180, 2011.
14. F. Mueller, A. Senecal, K. Tantale, H. Marie-Nelly, N. Ly, O. Collin, E. Basyuk, E. Bertrand, X. Darzacq, and C. Zimmer, "Fish-quant: automatic counting of transcripts in 3d fish images," *Nature methods*, vol. 10, no. 4, pp. 277–278, 2013.
15. F. De Chaumont, S. Dallongeville, N. Chenouard, N. Hervé, S. Pop, T. Provoost, V. Meas-Yedid, P. Pankajakshan, T. Lecomte, Y. Le Montagner, et al., "Icy: an open bioimage informatics platform for extended reproducible research," *Nature methods*, vol. 9, no. 7, p. 690, 2012.

16. P. R. Gudla, K. Nakayama, G. Pegoraro, and T. Misteli, "Spotlearn: Convolutional neural network for detection of fluorescence in situ hybridization (fish) signals in high-throughput imaging approaches," in *Cold Spring Harbor symposia on quantitative biology*, vol. 82, pp. 57–70, Cold Spring Harbor Laboratory Press, 2017.
17. M. A. Mabaso, D. J. Withey, and B. Twala, "Spot detection in microscopy images using convolutional neural network with sliding-window approach," 2018.
18. Z. Liu, G. Gao, L. Sun, and Z. Fang, "Hrdnet: high-resolution detection network for small objects," in *2021 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1–6, IEEE, 2021.
19. K. Fu, J. Li, L. Ma, K. Mu, and Y. Tian, "Intrinsic relationship reasoning for small object detection," *arXiv preprint arXiv:2009.00833*, 2020.
20. S. Bell, C. L. Zitnick, K. Bala, and R. Girshick, "Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2874–2883, 2016.
21. J. Noh, W. Bae, W. Lee, J. Seo, and G. Kim, "Better to follow, follow to be better: Towards precise supervision of feature super-resolution for small object detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9725–9734, 2019.
22. R. Hamaguchi, K. Fujita, A. and Nemoto, T. Imaizumi, and S. Hikosaka, "Effective use of dilated convolutions for segmenting small object instances in remote sensing imagery," *arXiv preprint, arXiv:1709.00179*, 2017.
23. F. Yu and V. Koltun, "Multi-scale context aggregation by dilated convolutions," *arXiv preprint arXiv:1511.07122*, 2015.
24. I. Smal, M. Loog, W. Niessen, and E. Meijering, "Quantitative comparison of spot detection methods in fluorescence microscopy," *IEEE transactions on medical imaging*, vol. 29, no. 2, pp. 282–301, 2009.
25. A. Sadr, E. E. Vos, B. A. Bassett, Z. Hosenie, N. Oozeer, and M. Lochner, "Deepsource: point source detection using deep learning," *Monthly Notices of the Royal Astronomical Society*, vol. 484, no. 2, pp. 2793–2806, 2019.
26. J. Liu, J. M. White, and R. M. Summers, "Automated detection of blob structures by hessian analysis and object scale," in *2010 IEEE International Conference on Image Processing*, pp. 841–844, IEEE, 2010.
27. E. Nehme, L. E. Weiss, T. Michaeli, and Y. Shechtman, "Deep-storm: super-resolution single-molecule microscopy by deep learning," *Optica*, vol. 5, no. 4, pp. 458–464, 2018.
28. C. Pino, R. Sortino, E. Sciacca, S. Riggi, and C. Spampinato, "Semantic segmentation of radio-astronomical images," in *Proc. of Progress in Artificial Intelligence and Pattern Recognition*, (Cham), pp. 393–403, Springer International Publishing, 2021.
29. R. Chouaib, A. Safieddine, X. Pichon, A. Imbert, O. S. Kwon, A. Samacoits, A.-M. Traboulsi, M.-C. Robert, N. Tzanov, E. Coleno, I. Poser, C. Zimmer, A. Hyman, H. Le Hir, K. Zibara, M. Peter, F. Mueller, T. Walter, and E. Bertrand, "A dual protein-mrna localization screen reveals compartmentalized translation and widespread co-translational rna targeting," *Developmental Cell*, vol. 54, no. 6, pp. 773–791.e5, 2020.
30. A. Imbert, W. Ouyang, A. Safieddine, E. Coleno, C. Zimmer, E. Bertrand, T. Walter, and F. Mueller, "Fish-quant v2: a scalable and modular analysis tool for smfish image analysis," *bioRxiv*, 2021.
31. K. A. Moutasim, M. L. Nystrom, and G. J. Thomas, "Cell migration and invasion assays," in *Cancer cell culture*, pp. 333–343, Springer, 2011.
32. H. Zhang, R. Singer, and G. J. Bassell, "Neurotrophin regulation of β -actin mrna and protein localization within growth cones," *The Journal of cell biology*, vol. 147, no. 1, pp. 59–70, 1999.
33. K. Lapidus, J. Wyckoff, G. Mouneimne, M. Lorenz, L. Soon, J. S. Condeelis, and R. H. Singer, "Zbp1 enhances cell polarity and reduces chemotaxis," *Journal of cell science*, vol. 120, no. 18, pp. 3173–3178, 2007.
34. J. Condeelis and R. H. Singer, "How and why does β -actin mrna target?," *Biology of the Cell*, vol. 97, no. 1, pp. 97–110, 2005.
35. P. Hu and D. Ramanan, "Finding tiny faces," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 951–959, 2017.
36. K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," in *European conference on computer vision*, pp. 630–645, Springer, 2016.
37. Y. Wang, X. Guo, P. Liu, and B. Wei, "Up and down residual blocks for convolutional generative adversarial networks," *IEEE Access*, vol. 9, pp. 26051–26058, 2021.
38. S. Jadon, "A survey of loss functions for semantic segmentation," in *2020 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB)*, pp. 1–7, IEEE, 2020.
39. J. Bergstra, D. Yamins, and D. Cox, "Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures," in *International conference on machine learning*, pp. 115–123, PMLR, 2013.
40. L. Li, K. Jamieson, A. Rostamizadeh, E. Gonina, M. Hardt, B. Recht, and A. Talwalkar, "A system for massively parallel hyperparameter tuning," 2020.
41. J. Bentley, "Multidimensional binary search trees used for associative searching," *Communications of the ACM*, vol. 18, no. 9, pp. 509–517, 1975.
42. H. Samet, *The design and analysis of spatial data structures*, vol. 85. Addison-Wesley Reading, MA, 1990.
43. M. De Berg, O. Cheong, M. Van Kreveld, and M. Overmars, *Computational geometry: algorithms and applications*. Springer, 2008.

44. A. F. Savulescu, R. Brackin, E. Bouilhol, B. Dartigues, J. H. Warrell, M. R. Pimentel, N. Beaume, I. C. Fortunato, S. Dal-longeville, M. Boule, *et al.*, “Interrogating rna and protein spatial subcellular distribution in smfish data with dypfish,” *Cell Reports Methods*, vol. 1, no. 5, p. 100068, 2021.
45. E. H. Kislauskis, Z. Li, R. H. Singer, and K. L. Taneja, “Isoform-specific 3'-untranslated sequences sort alpha-cardiac and beta-cytoplasmic actin messenger rnas to different cytoplasmic compartments.,” *The Journal of cell biology*, vol. 123, no. 1, pp. 165–172, 1993.