Supplementary Material

Generative Adversarial Networks for Augmenting Training Data of Microscopic Cell Images

Piotr Baniukiewicz, E. Josiah Lutton, Sharon Collier, Till Bretschneider*

Department of Computer Science & Zeeman Institute, University of Warwick, Coventry CV4 7AL, UK

Software and Training Data

We have been using Python 3.6.8 and the pix2pix implementation from https://github.com/affinelayer/pix2pix-tensorflow using the default parameters.

All training data and the network models are available from the following repository: https://pilip.lnx.warwick.ac.uk/Frontiers_2019/.

Examples of the training and testing of the three 2D networks:

```
# Training

python pix2pix.py --mode train --output_dir Labels/ADB-GFP/train_output --max_epochs 200 --input_dir Labels/ABD-GFP/originals_processed --which_direction BtoA --display_freq 500 --save_freq 1200

# Testing: Generate synthetic cells from binary masks

python pix2pix.py --mode test --output_dir Labels/cross-tests/ADB-GFP/test_output --input_dir Labels/cross-tests/ABD-GFP/originals_processed --checkpoint Labels/ABD-GFP/train_output
```

The modified pix2pix implementation for generating pseudo 3D cell data can be found here:

https://pilip.lnx.warwick.ac.uk/Frontiers_2019/Net_3D/Figure%205/pix2pix_3D_multichannel.py

Supplementary Figures

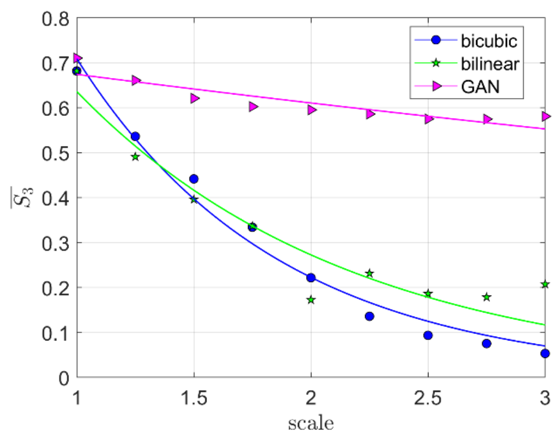

Supplementary Figure S1. S_3 metric versus image scale for different interpolation algorithms, including proposed GANs. Compared to bicubic and bilinear scaling of the original image which result in a marked decrease of the S_3 values, scaling of masks that are used as input to GANs and subsequent texturing preserve the local perceived sharpness better.

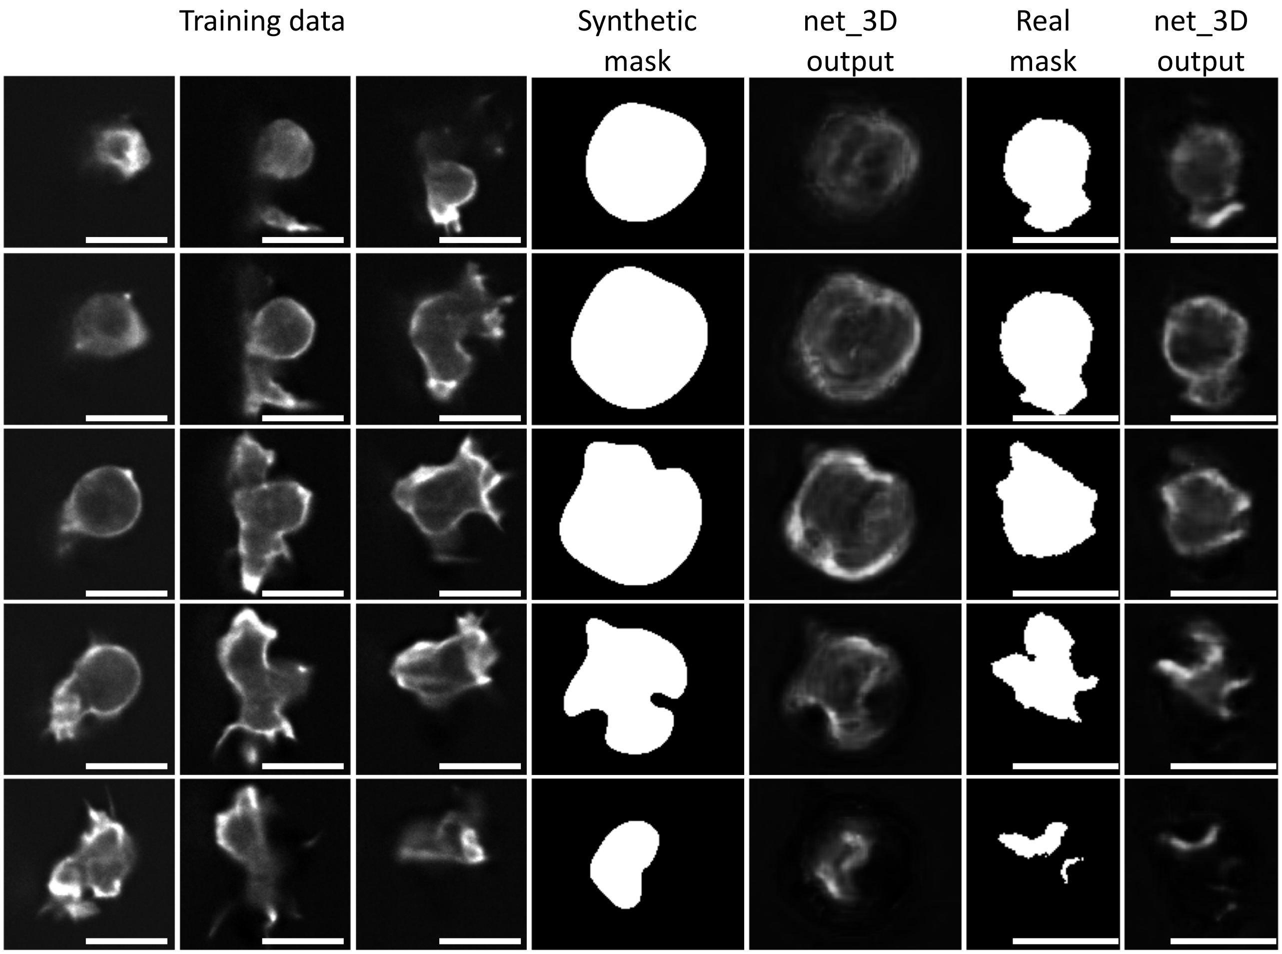

Supplementary Figure S2. Slices taken from: net_3D training data (columns 1-3), synthetic mask (column 4), output of net_3D applied to the synthetic mask (column 5), manually segmented mask (column 6), output of net_3D applied to the manually segmented mask (column 7). Each column shows a sequence of evenly-spaced slices through a volume, with slice positions in column 5 corresponding to those in column 4, and slice positions in column 7 corresponding to those in column 6. Scale bars represent 10 µm.

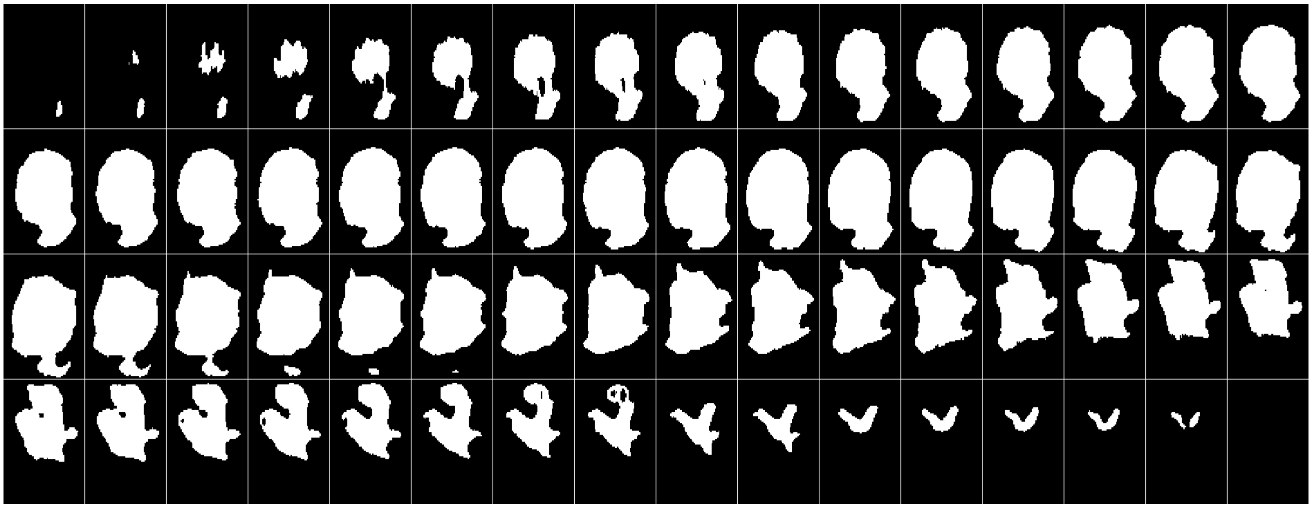

Supplementary Figure S3. All slices of an example input binary mask for net_ABD. Binary masks were obtained from real 3D cell images through manual segmentation, yielding precise cell shapes with varied structural features.

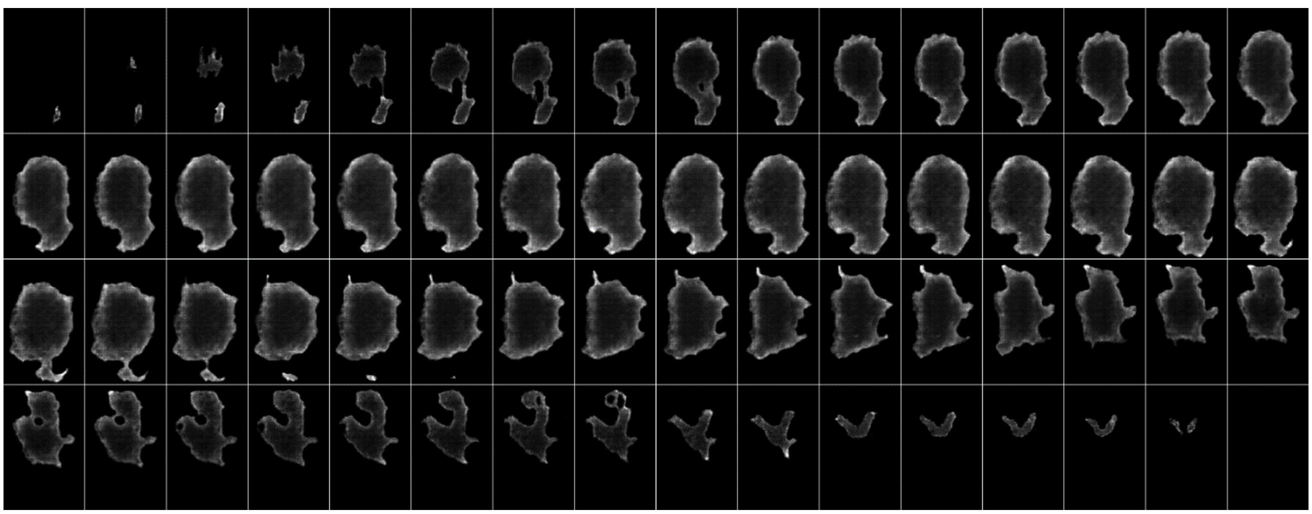

Supplementary Figure S4. All slices of the output of net_ABD when applied to the binary mask shown in Supplementary Figure S3. Similar to the 2D case (Figures 2-3), each slice shows increased intensity at the boundary of the mask, with the highest intensities at the tips of protruding structures. Adjacent slices also show a high level of correlation in intensity, meaning that the 3D nature of the data is still preserved despite processing each slice separately.

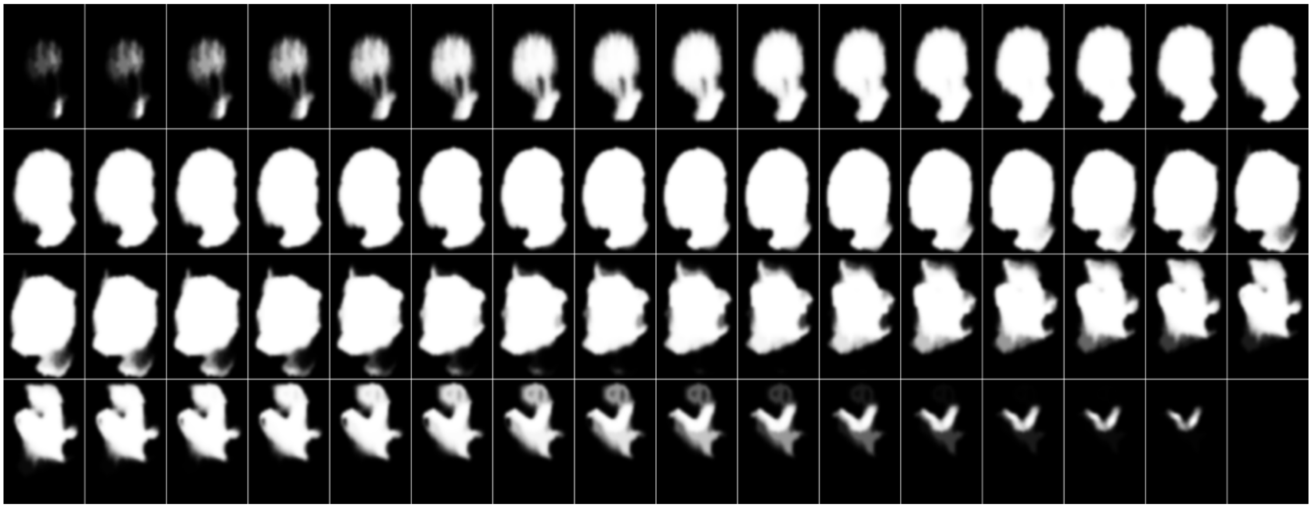

Supplementary Figure S5. Blurred input mask for net_ABD, obtained by applying a Gaussian filter to the mask shown in Supplementary Figure S3.

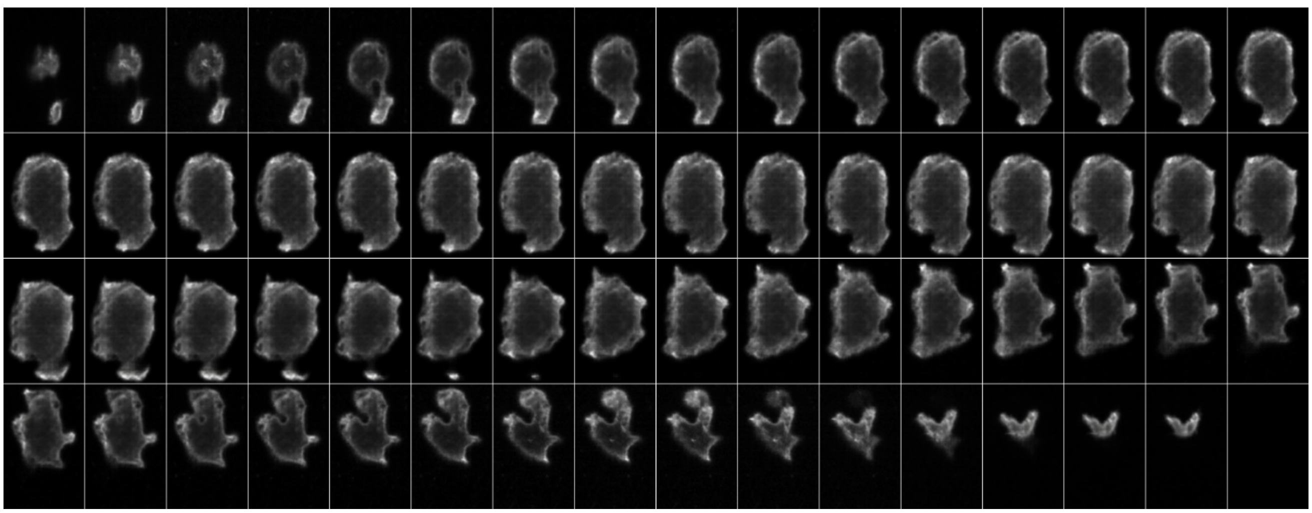

Supplementary Figure S6. All slices of the output of net_ABD when applied to the blurred mask shown in Supplementary Figure S5. This shows a similar pattern of localization to Supplementary Figure S4, but the sharpness of the boundary is reduced, which more closely represents 3D-captured cell images (see Figure 8 for further details).

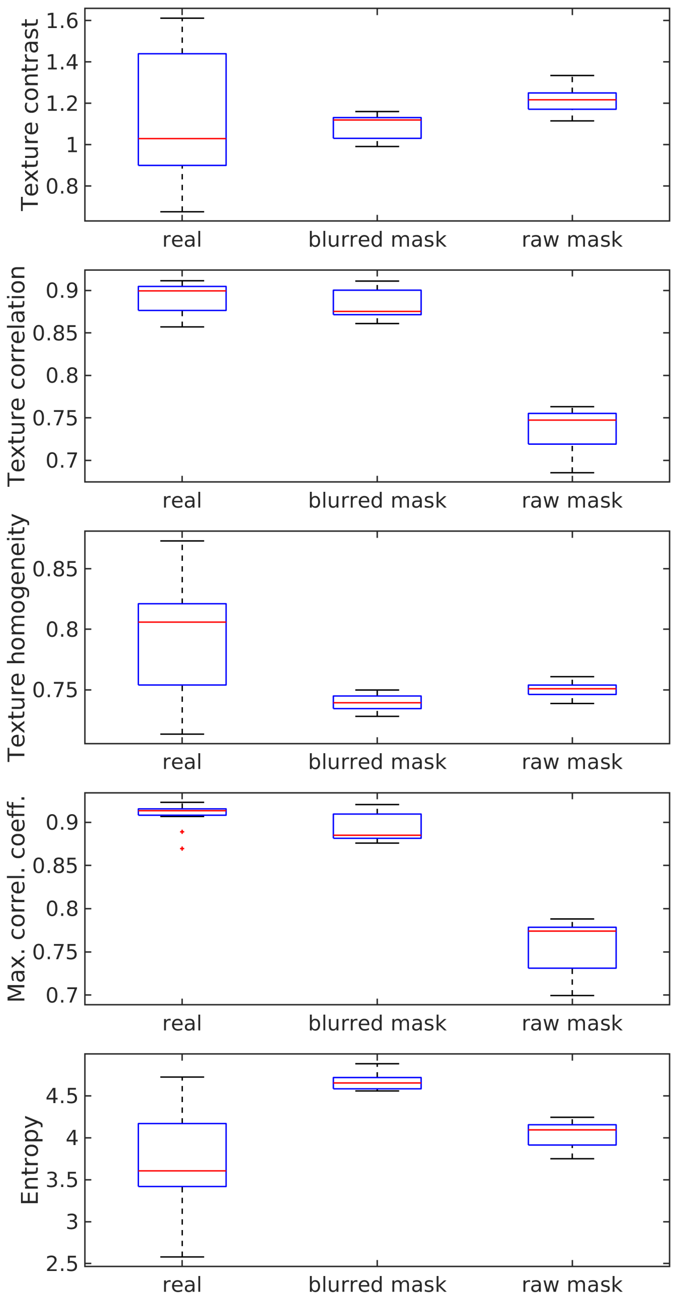

Supplementary Figure S7. Box plots comparing the distribution of 5 Haralick texture features in real data and synthetic data generated by applying net_ABD to blurred binary masks and raw binary masks (pseudo-3D). In all cases, the median value for the features of the blurred mask dataset lie within the range of values for the real data, but are higher in contrast and entropy, and lower in correlation, homogeneity and maximum correlation coefficient. This is likely due to low-level fluctuations in intensity in the synthetic data that are absent in real data. Blurring the mask yields correlation, contrast and maximum correlation coefficient descriptors that are closer to those in the real data, which shows that blurring the mask improves these descriptors, despite only acting on the cell boundary.