CASUS
CENTER FOR ADVANCED SYSTEMS UNDERSTANDING

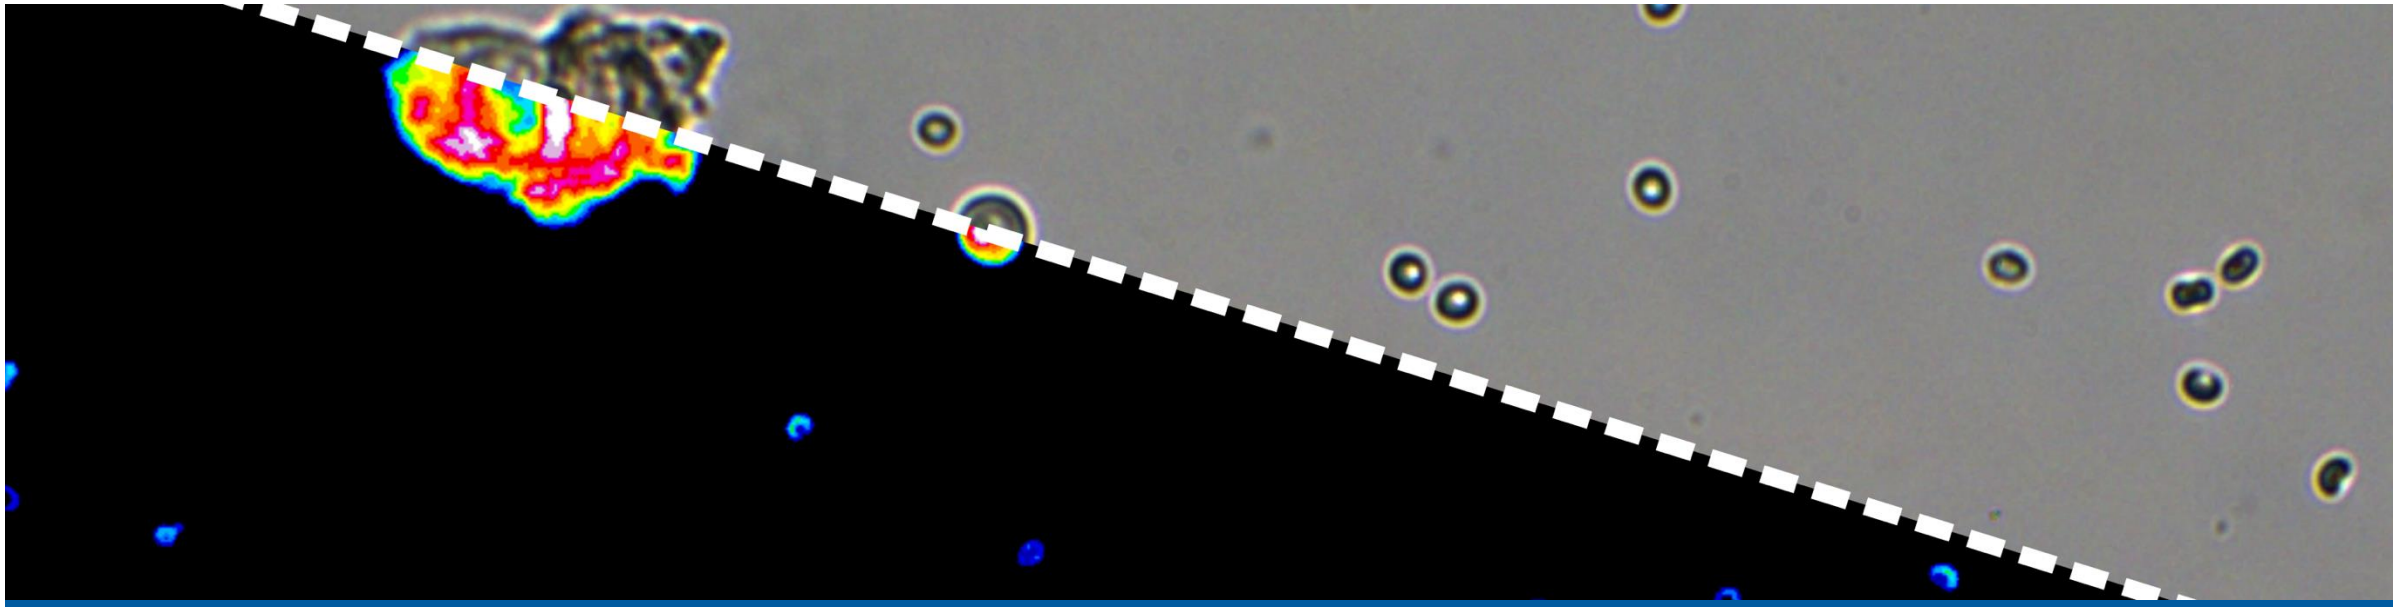

Generative AI for Inverse Problems in Biomedical Computational Microscopy

HZDR · CASUS · Prof. Dr. Artur Yakimovich · a.yakimovich@hzdr.de · www.casus.science

INSTITUTE OF

HZDR HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

PARTICIPATING INSTITUTIONS

UFZ HELMHOLTZ Centre for Environmental Research

TECHNISCHE UNIVERSITÄT DRESDEN

Uniwersytet Wrocławski

MAX PLANCK INSTITUTE OF MOLECULAR CELL BIOLOGY AND GENETICS

FUNDED BY

Federal Ministry of Education and Research

SACHSEN Diese Maßnahme wird mitfinanziert mit Steuermitteln auf Grundlage des vom Sächsischen Landtag beschlossenen Haushaltes.

The Mixed Blessing of a Microscope

Reality

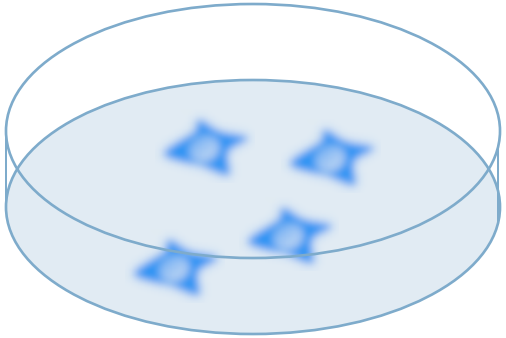

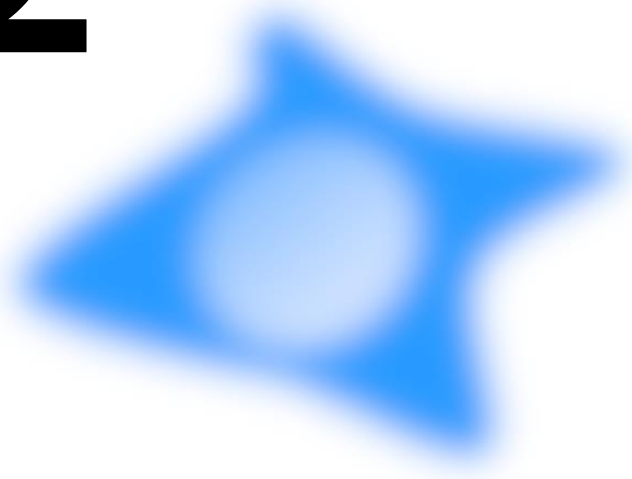

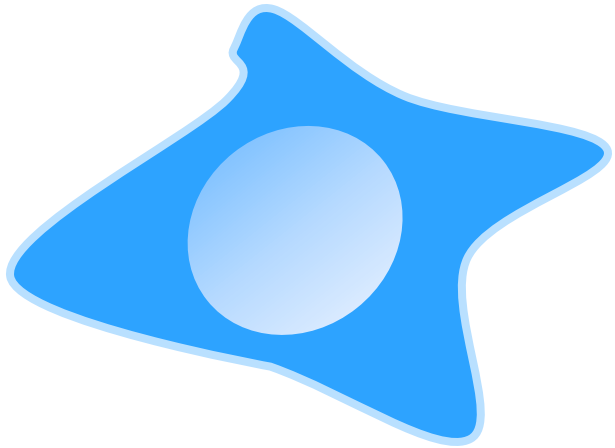

Small Objects
(e.g. cultured human cells)

Small Objects
Magnified by a Microscope

True Magnified
Small Objects

How Far can We See in a Microscope?

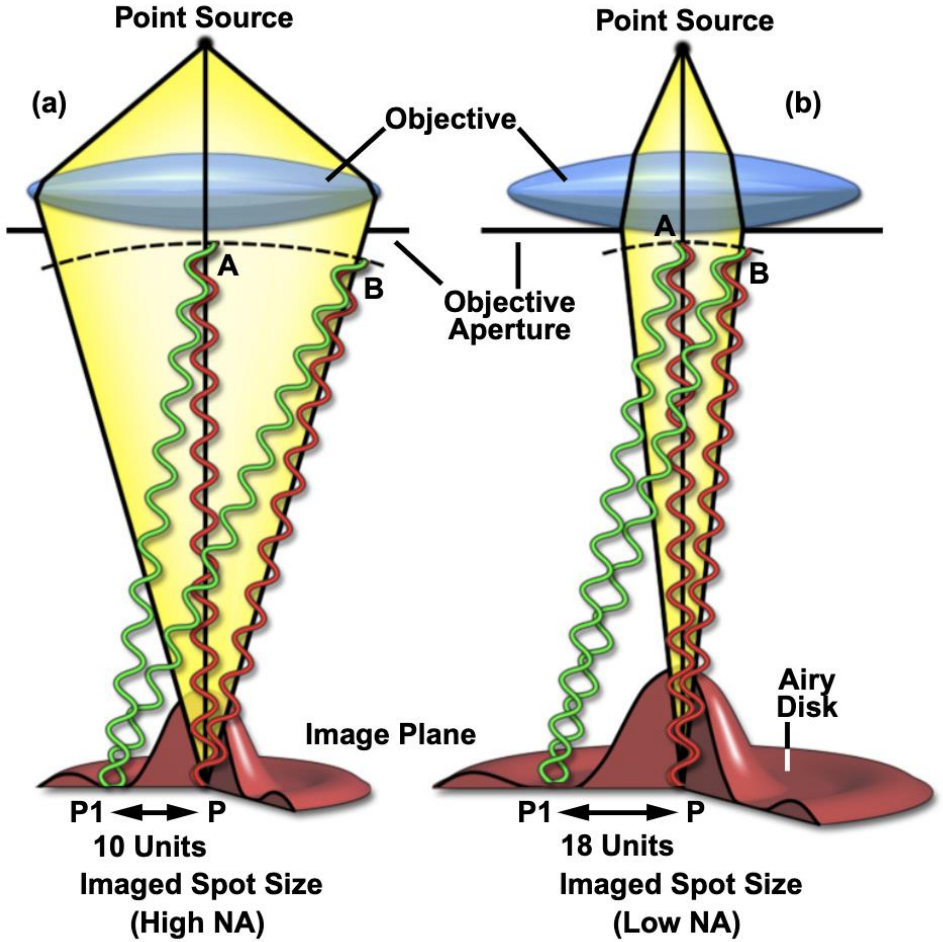

Source: microscopyu.com

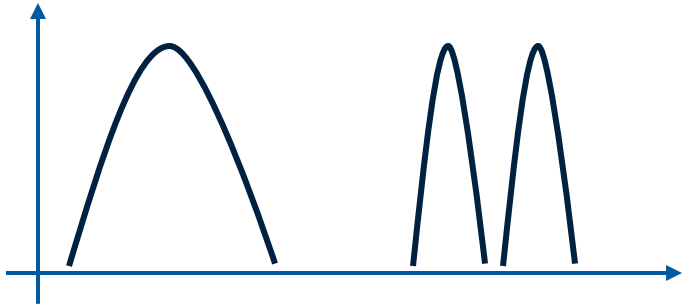

$$d = \frac{\lambda}{2n\sin\theta} = \frac{\lambda}{2\mathrm{NA}}$$

Refractive index

Wavelength

Numerical aperture

Resolvable feature size

half-angle subtended by the optical objective lens

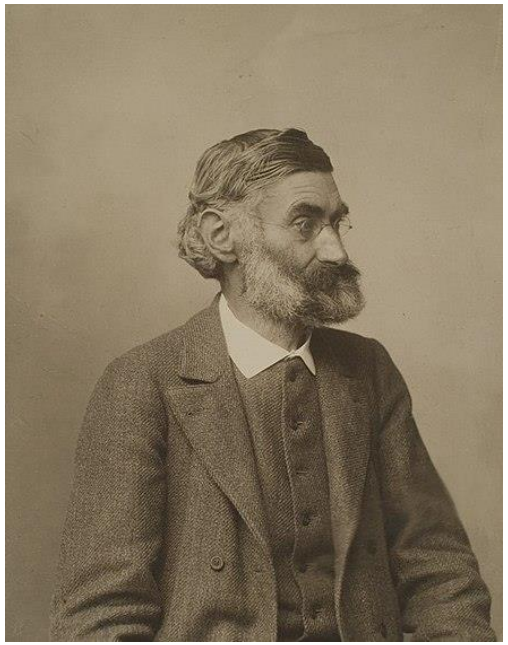

Ernst Karl Abbe

Super-resolution Microscopy

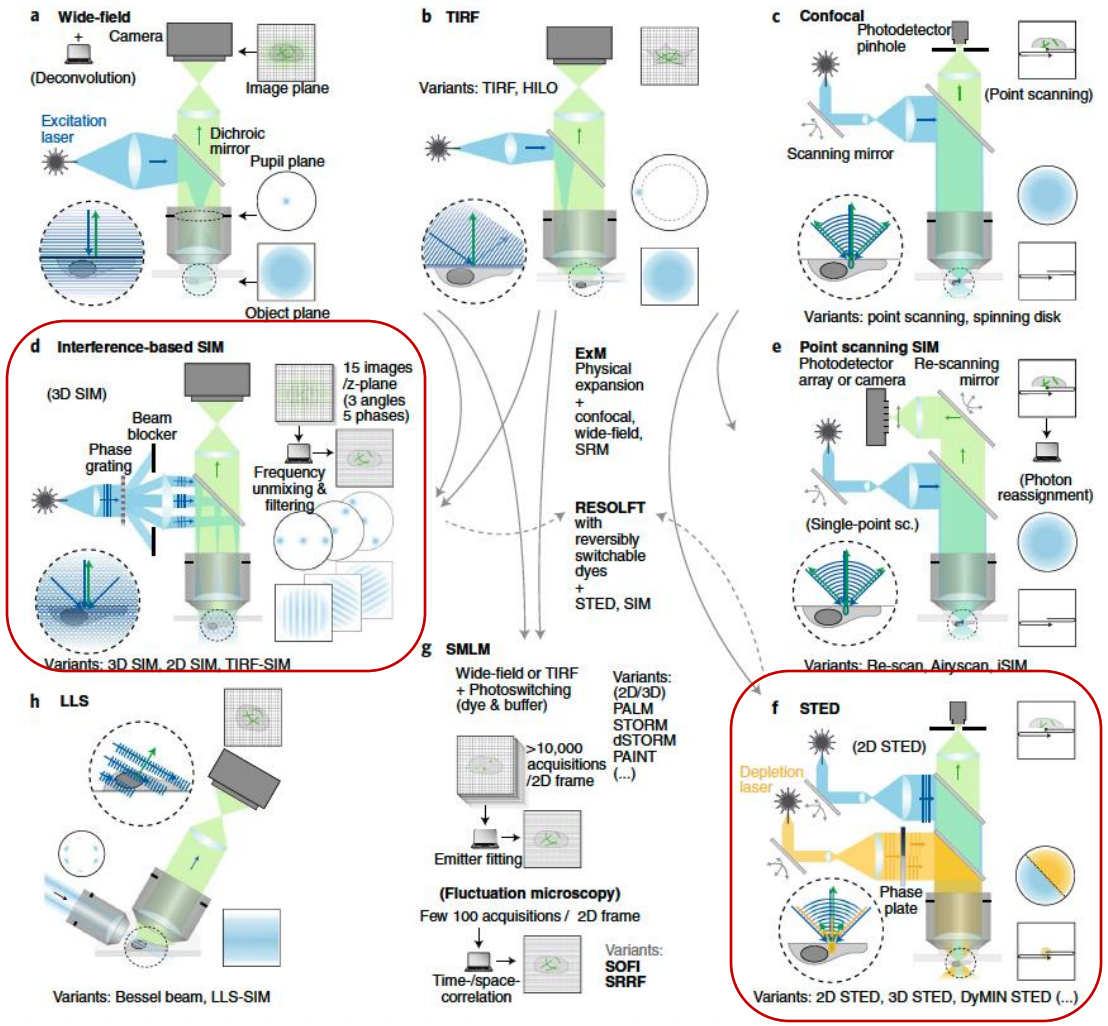

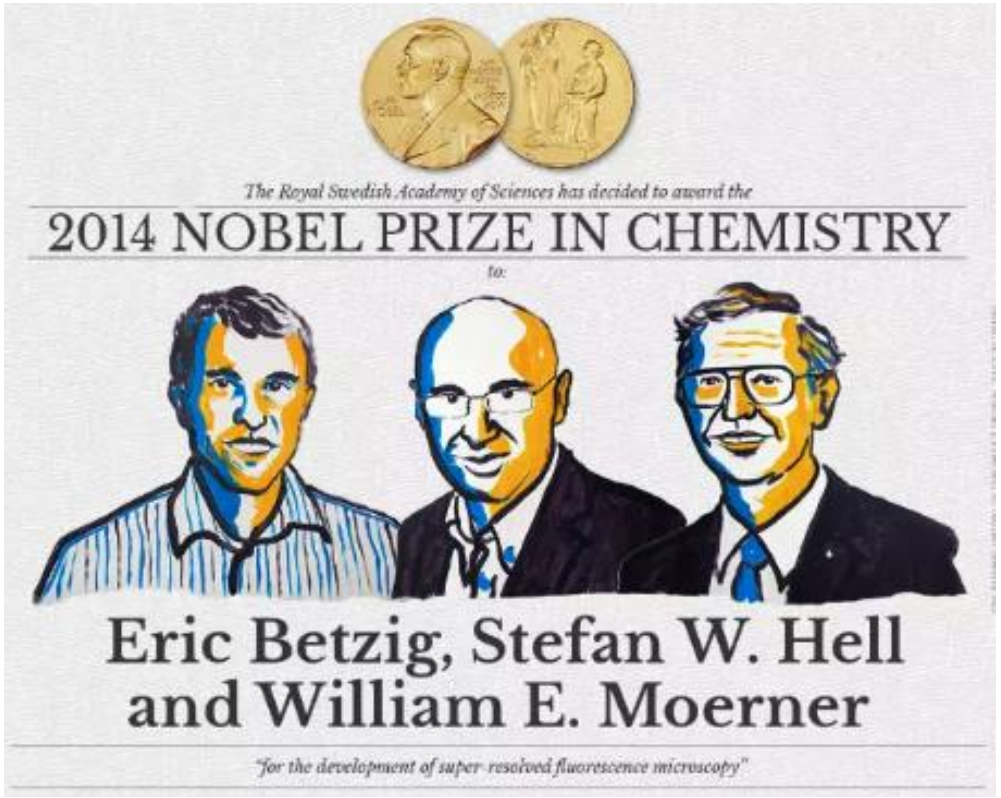

Schermelleh, Lothar, et al. "Super-resolution microscopy demystified." *Nature cell biology* 21.1 (2019): 72-84.

Super-resolution Microscopy as an Inverse Problem

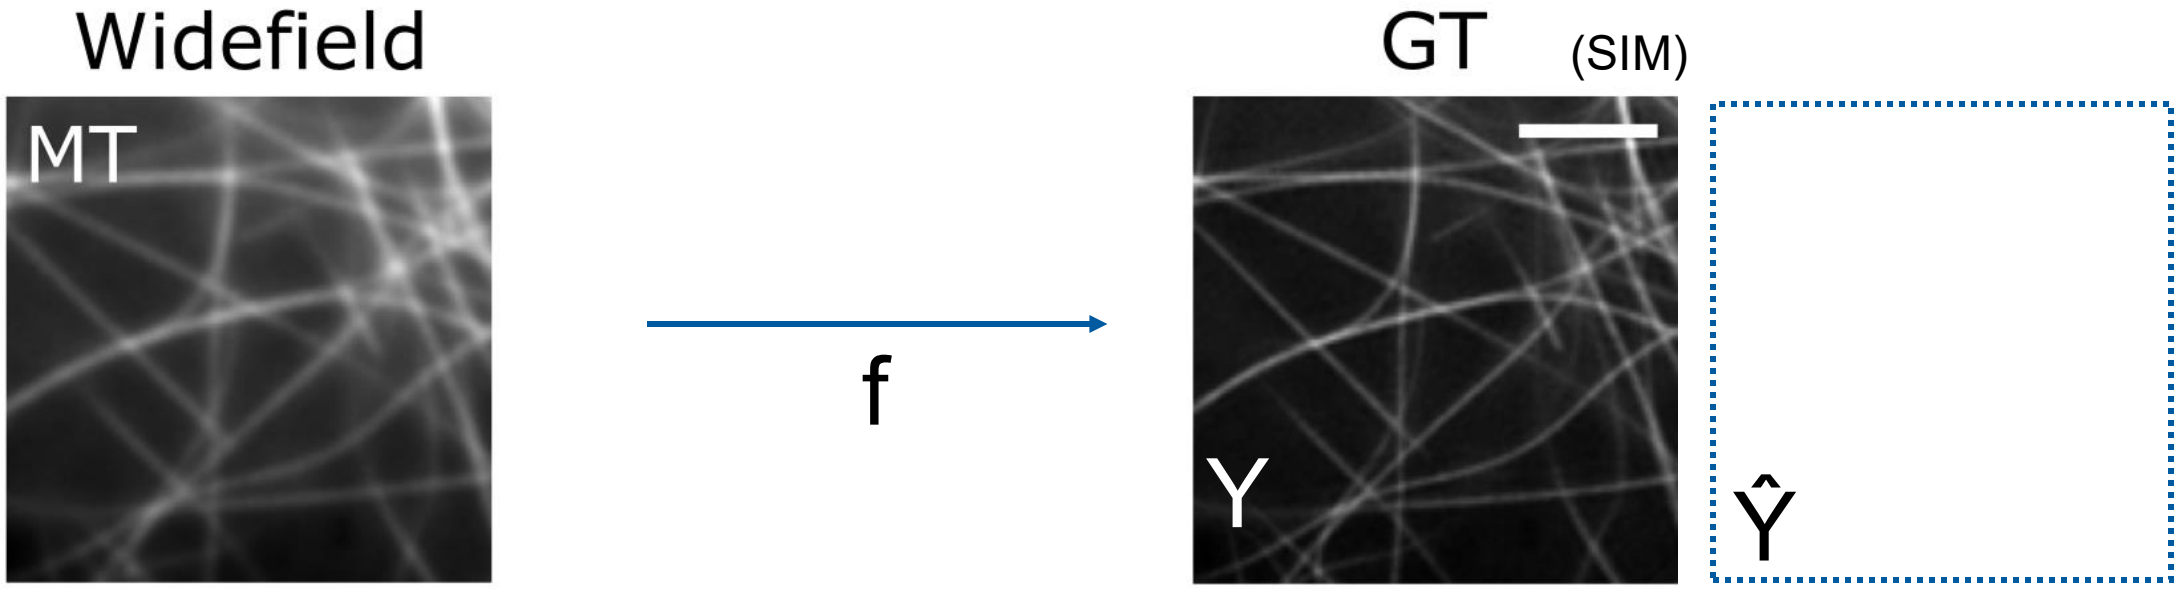

Machine Learning as a Solution to an Inverse Problem

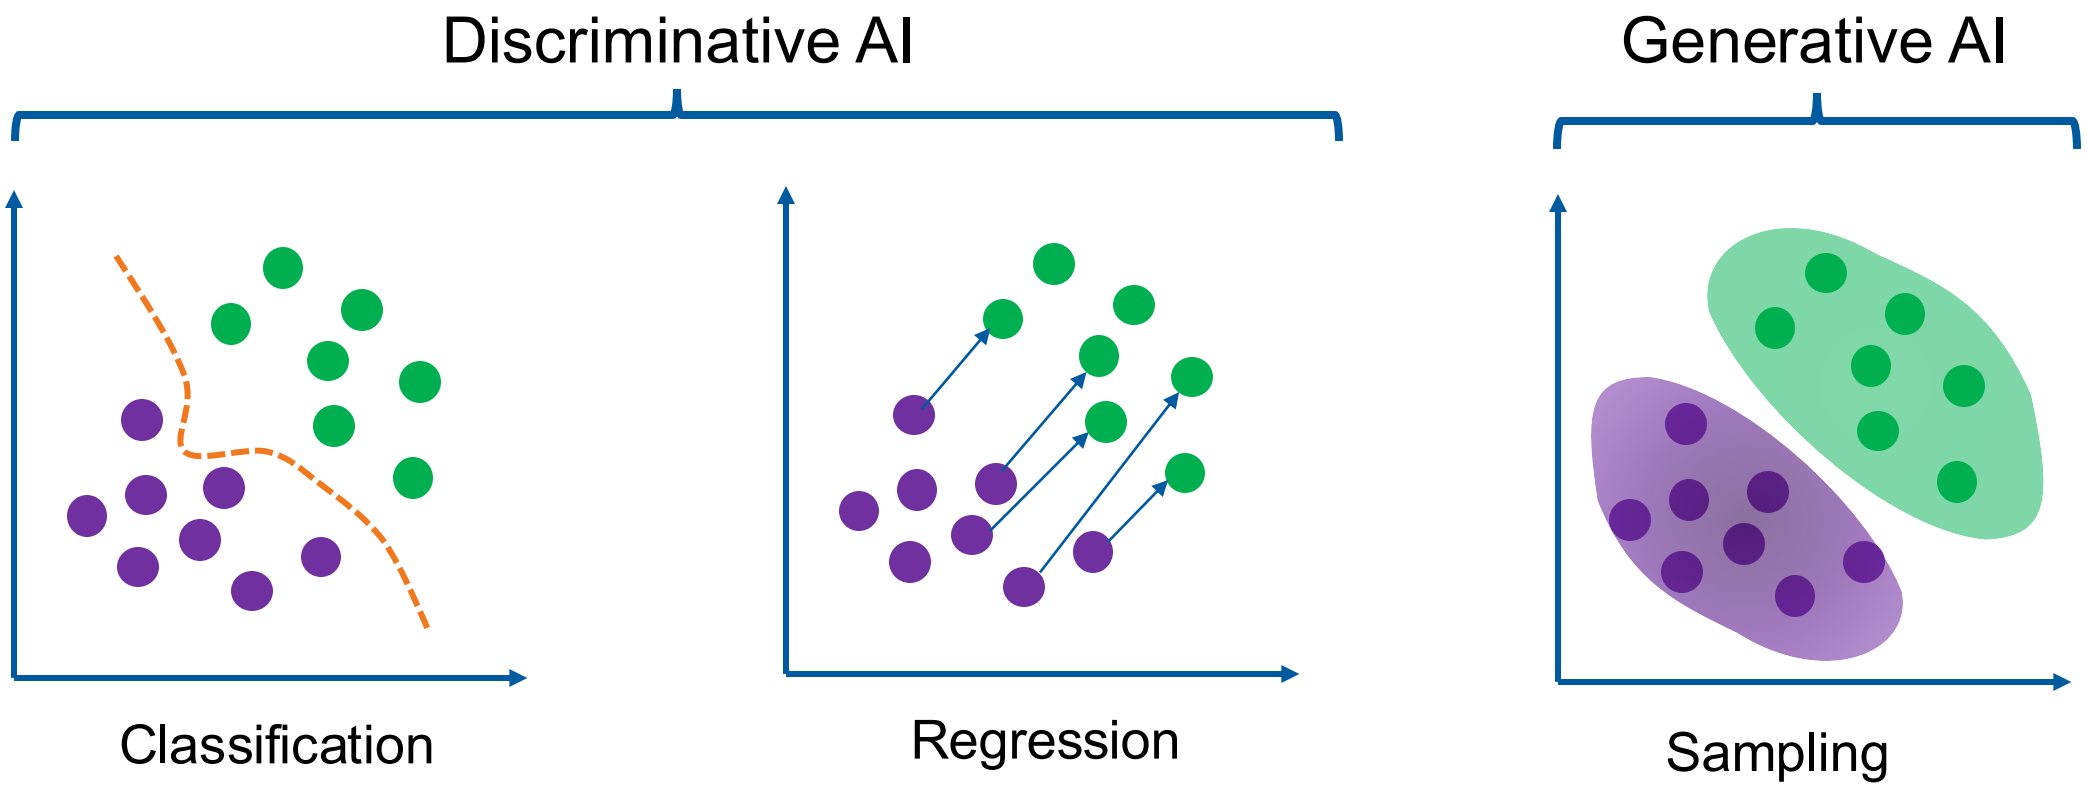

Machine Learning as a Solution to an Inverse Problem

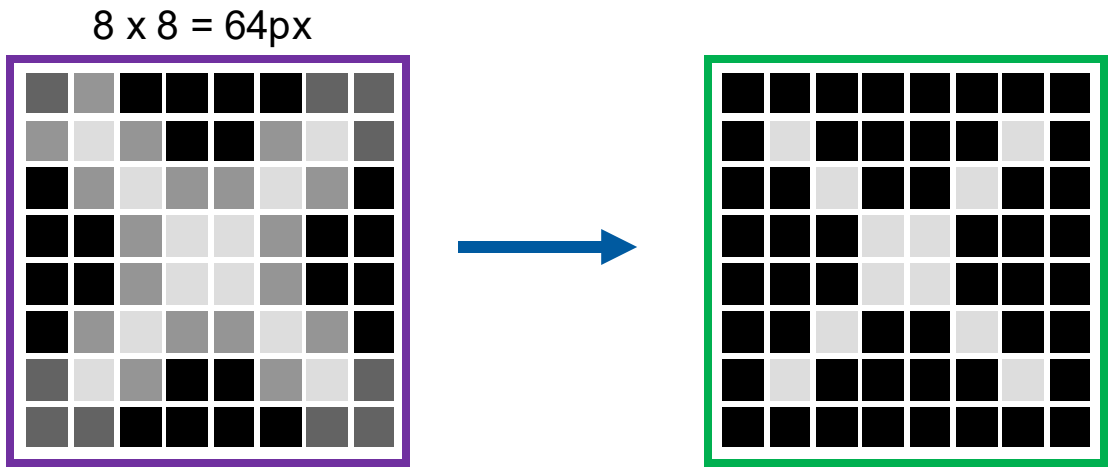

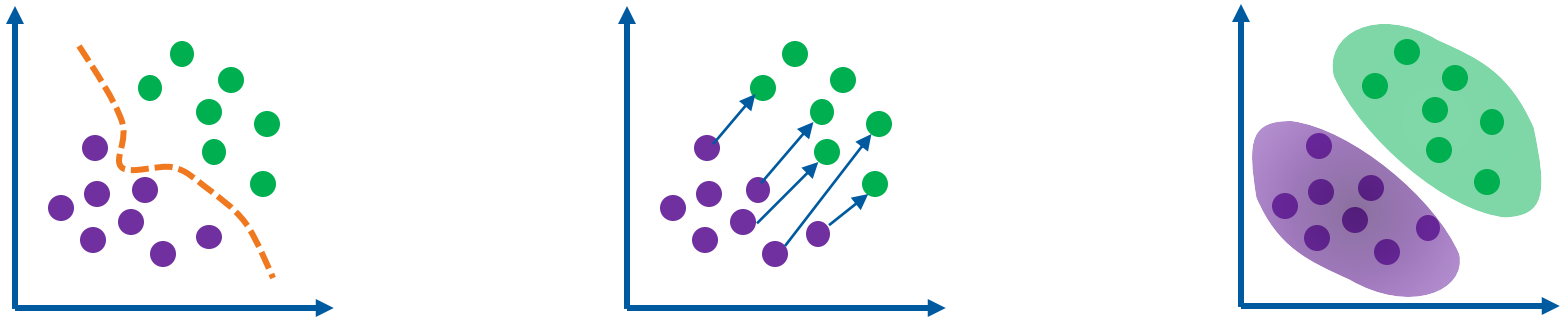

Classification | Regression | Sampling

Convolutional Neural Networks

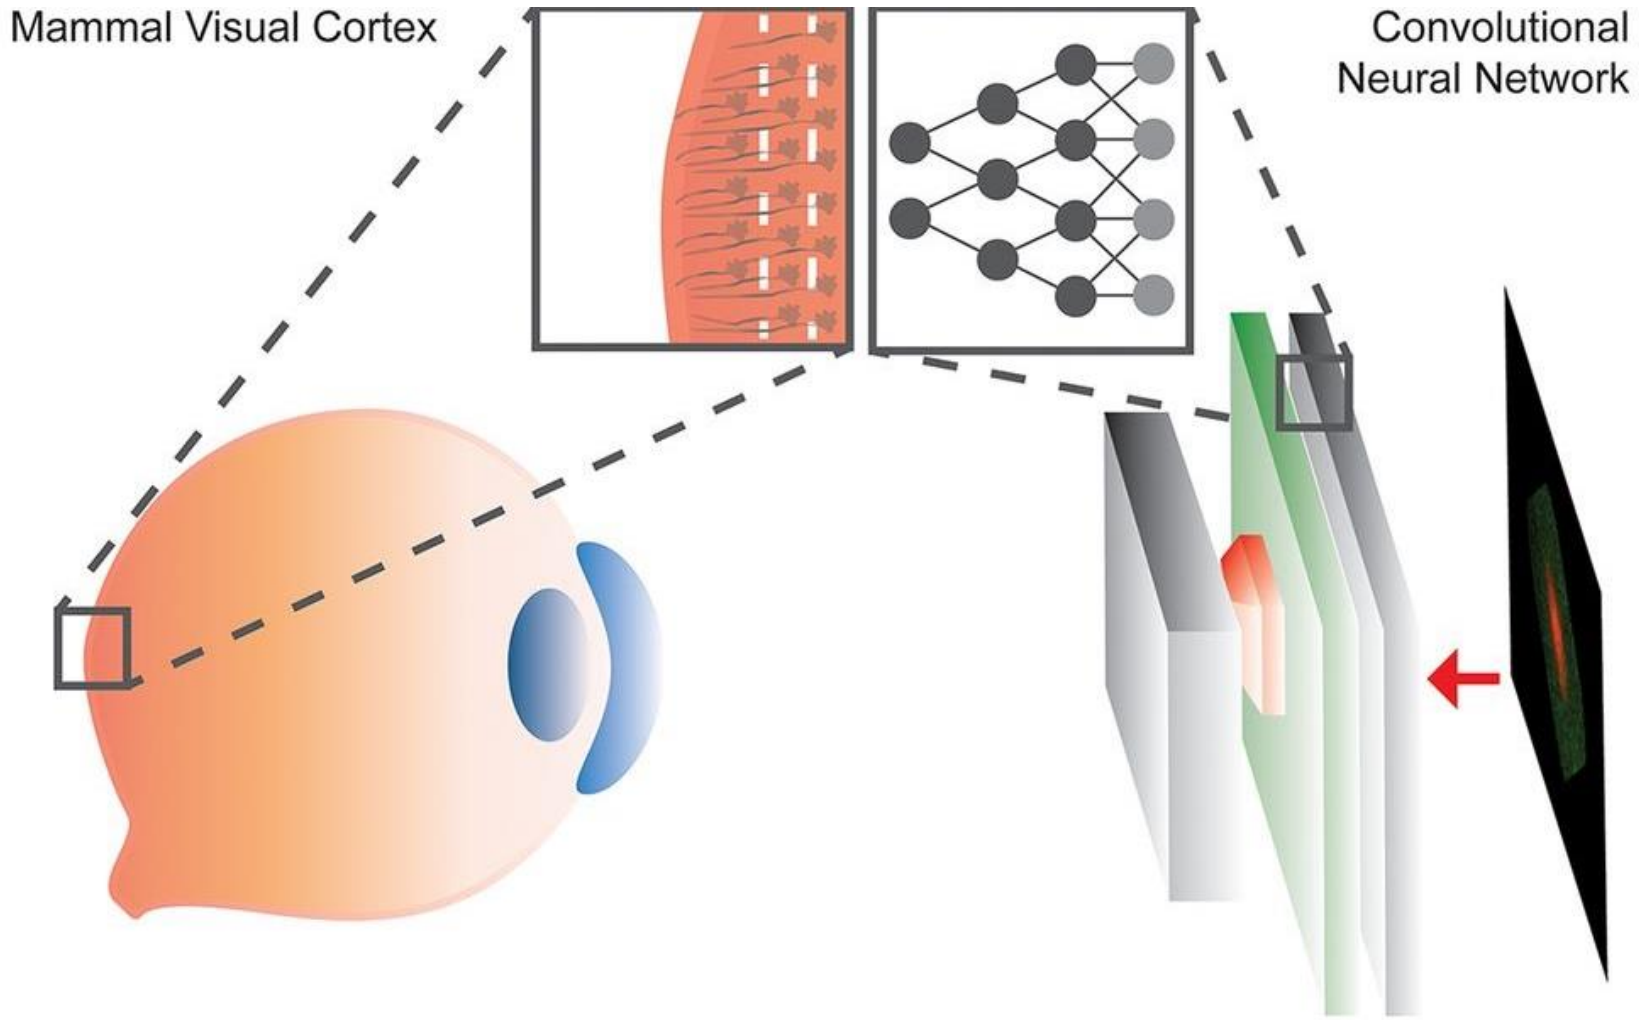

Super-resolution Microscopy as an Inverse Problem

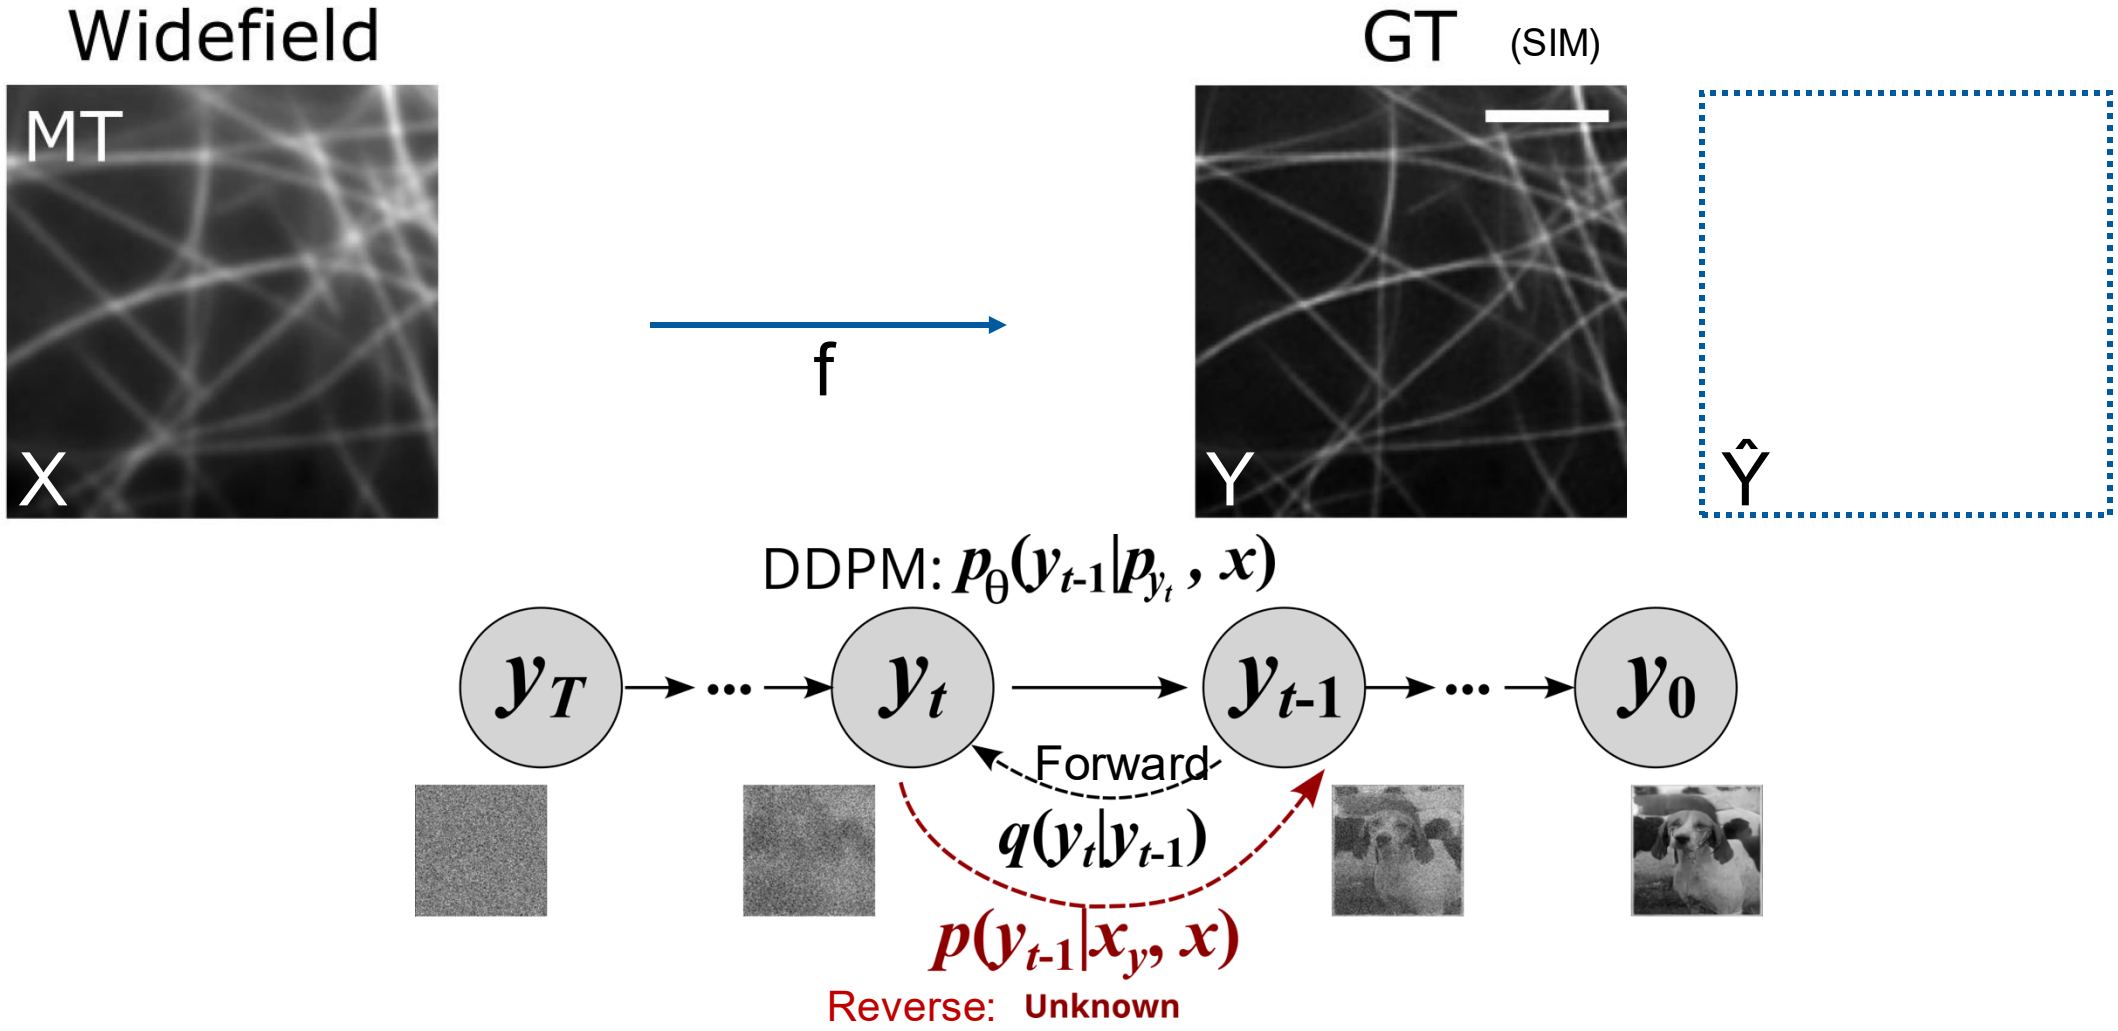

Ho, Jonathan, Ajay Jain, and Pieter Abbeel. "Denoising diffusion probabilistic models." *Advances in neural information processing systems* 33 (2020): 6840-6851.

Computational Super-resolution Microscopy: Finding a Suitable Dataset

BioSR ~20k

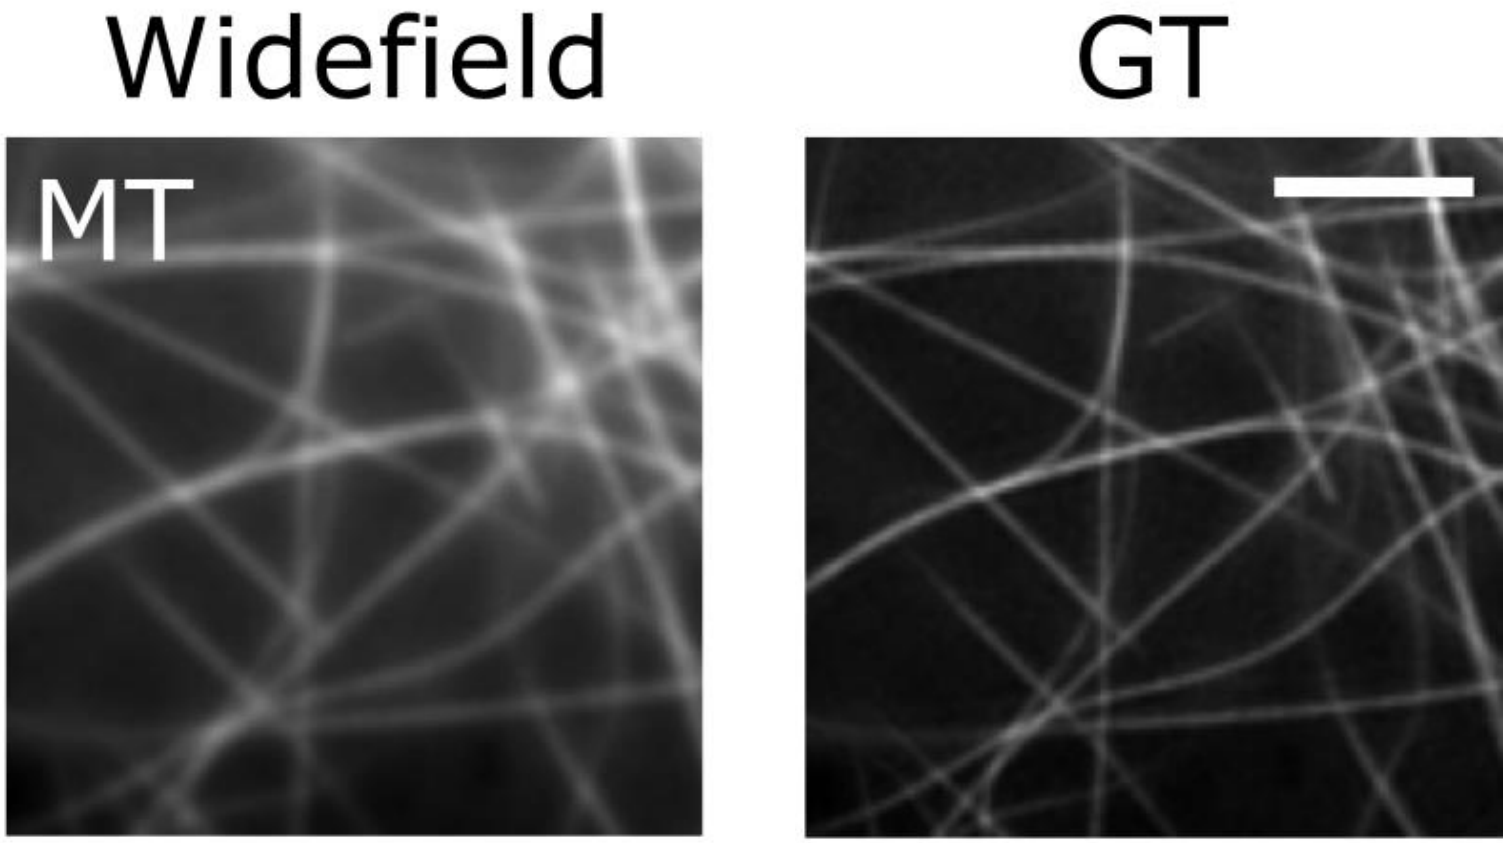

Not nearly enough!

Li, della Maggiora & Yakimovich et al., 2024, Coms. Eng.

Computational Super-resolution Microscopy: Finding a Suitable Dataset

~20k

BioSR

>1.2mln

ImageNet

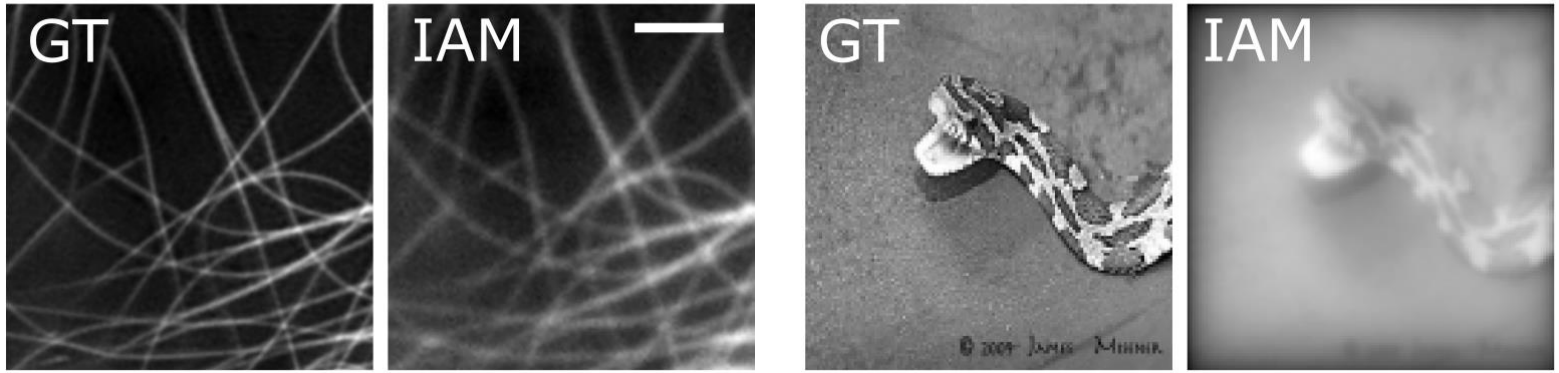

Rui Li

Gabriel della Maggiora

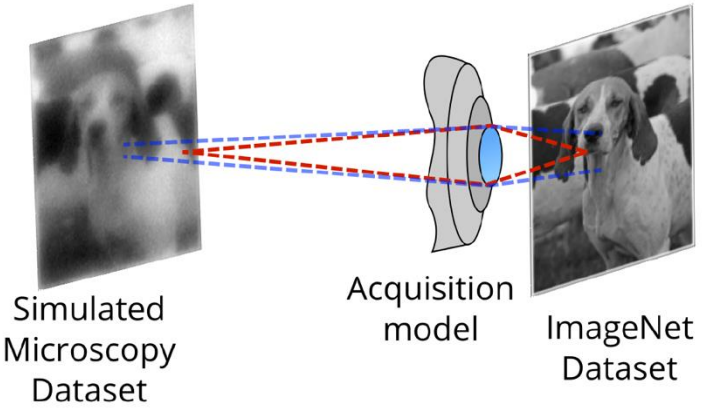

$$I = \phi(h * x + b)$$

, where the image *I*, is the result of the convolution between object *x* and system PSF *h* with the background signal noise *b* added

Arnison and Sheppard. A 3d vectorial optical transfer function suitable for arbitrary pupil functions. Optics communications, 211(1-6):53–63, 2002.

Li, della Maggiora & Yakimovich et al., 2024, Coms. Eng.

Super-resolution Inverse Problems are Ill-posed

Widefield

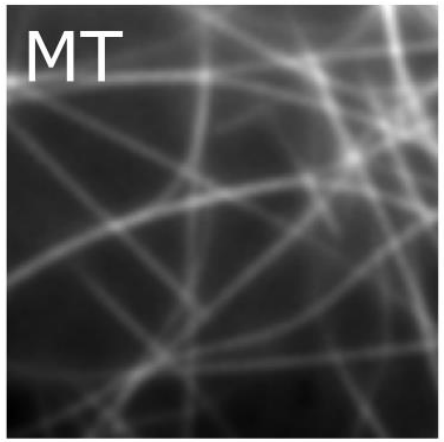

f

GT (SIM)

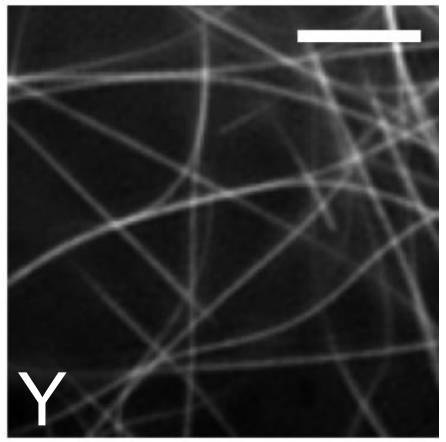

$\hat{Y}$

What colour are the dresses?

Valid option 1

Valid option 2

From Chang et al. 2023, CVPR

Super-resolution Microscopy using Physics-informed Diffusion Probabilistic Models

Reverse process:

Physics-informed regularisation term

$$p_\theta(y_{t-1}|y_t, x) - \nabla_{y_t}\left[||Ay_t - \tilde{x}||_2^2 + \lambda_t R(y_t)\right]$$

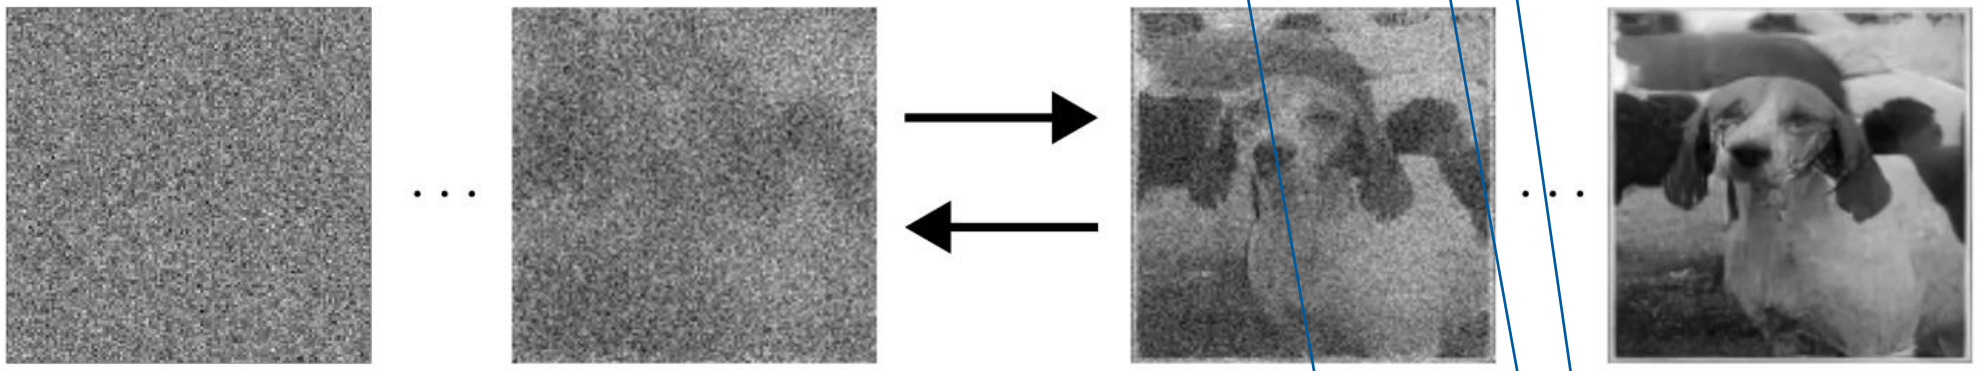

$y_T \sim \mathcal{N}(0, \mathbf{I})$ y_t y_{t-1} $y_0 \sim p(y|x)$

$q(y_t|y_{t-1})$

Forward process

Gaussian

Noisy input

Regulariser

Weight

Li, della Maggiora & Yakimovich et al., 2024, Coms. Eng.

PI-DDPM Performance on Benchmarks

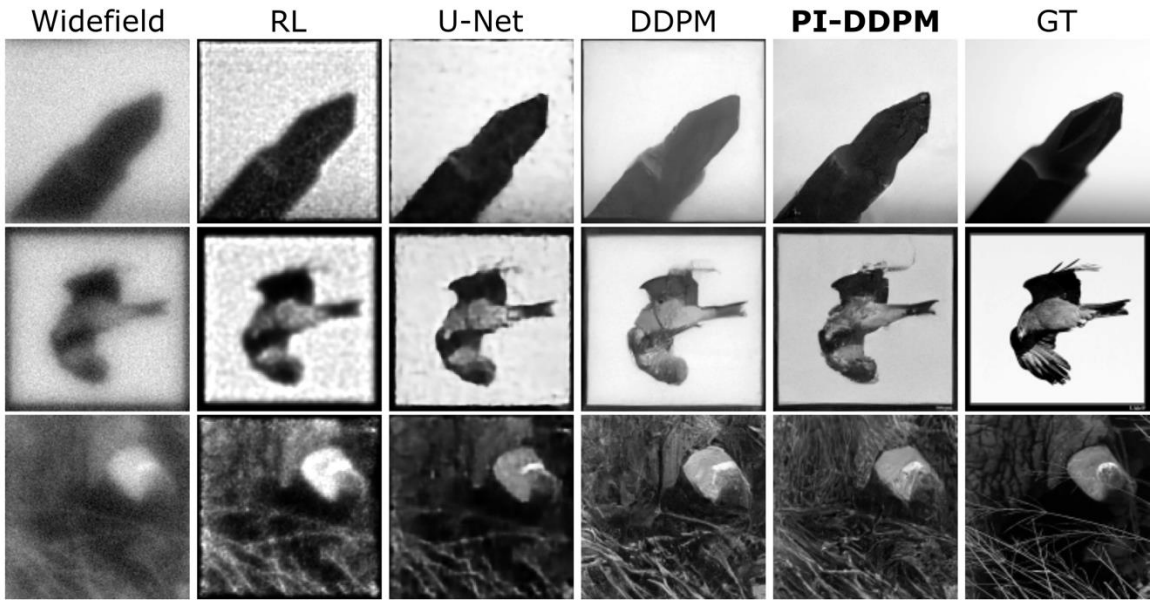

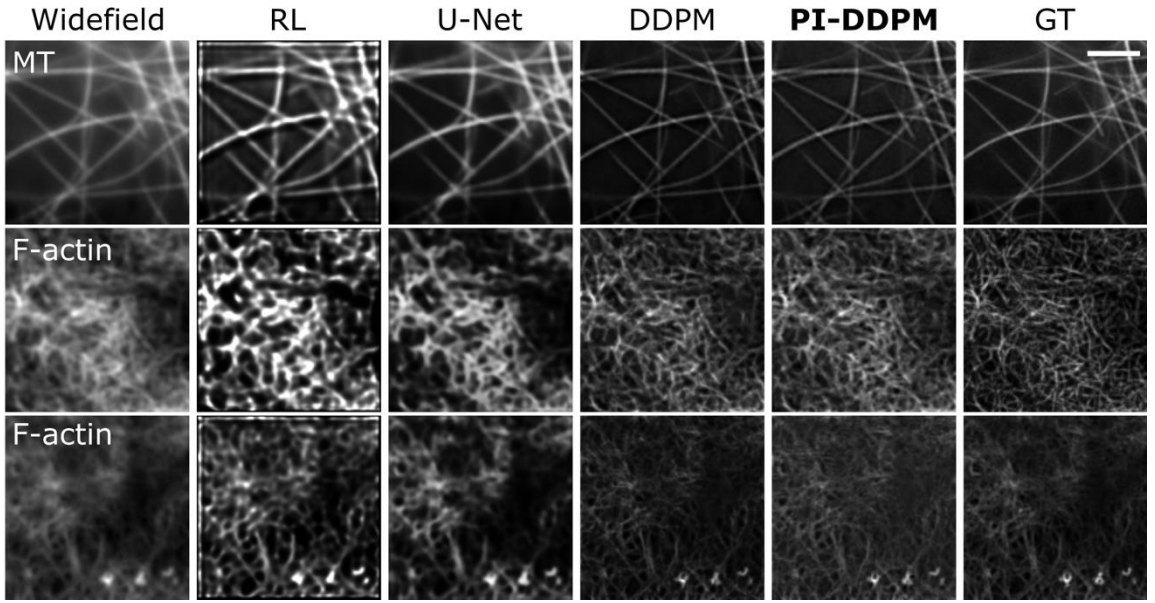

Table S1. Performance on Simulated Microscopy using BioSR Test Set

Metric / Model	Original	Richardson-Lucy	U-Net	DDPM	PI-DDPM
PSNR	16.127	13.394	18.301	**20.446**	**20.217**
MS-SSIM	0.745	0.684	0.812	**0.859**	**0.859**
NRMSE	0.156	0.214	0.122	**0.095**	**0.098**

Table 1: Performance on Widefield Microscopy using BioSR Test Set

Metric / Model	Original	U-Net	DDPM	PI-DDPM (ours)
PSNR	18.649	19.706	23.703	**23.974**
MS-SSIM	0.628	0.652	0.784	**0.795**
NRMSE	0.147	0.126	0.070	**0.069**

Li, della Maggiora & Yakimovich et al., 2024, Coms. Eng.

PI-DDPM Performance on Wild Data (dSTORM)

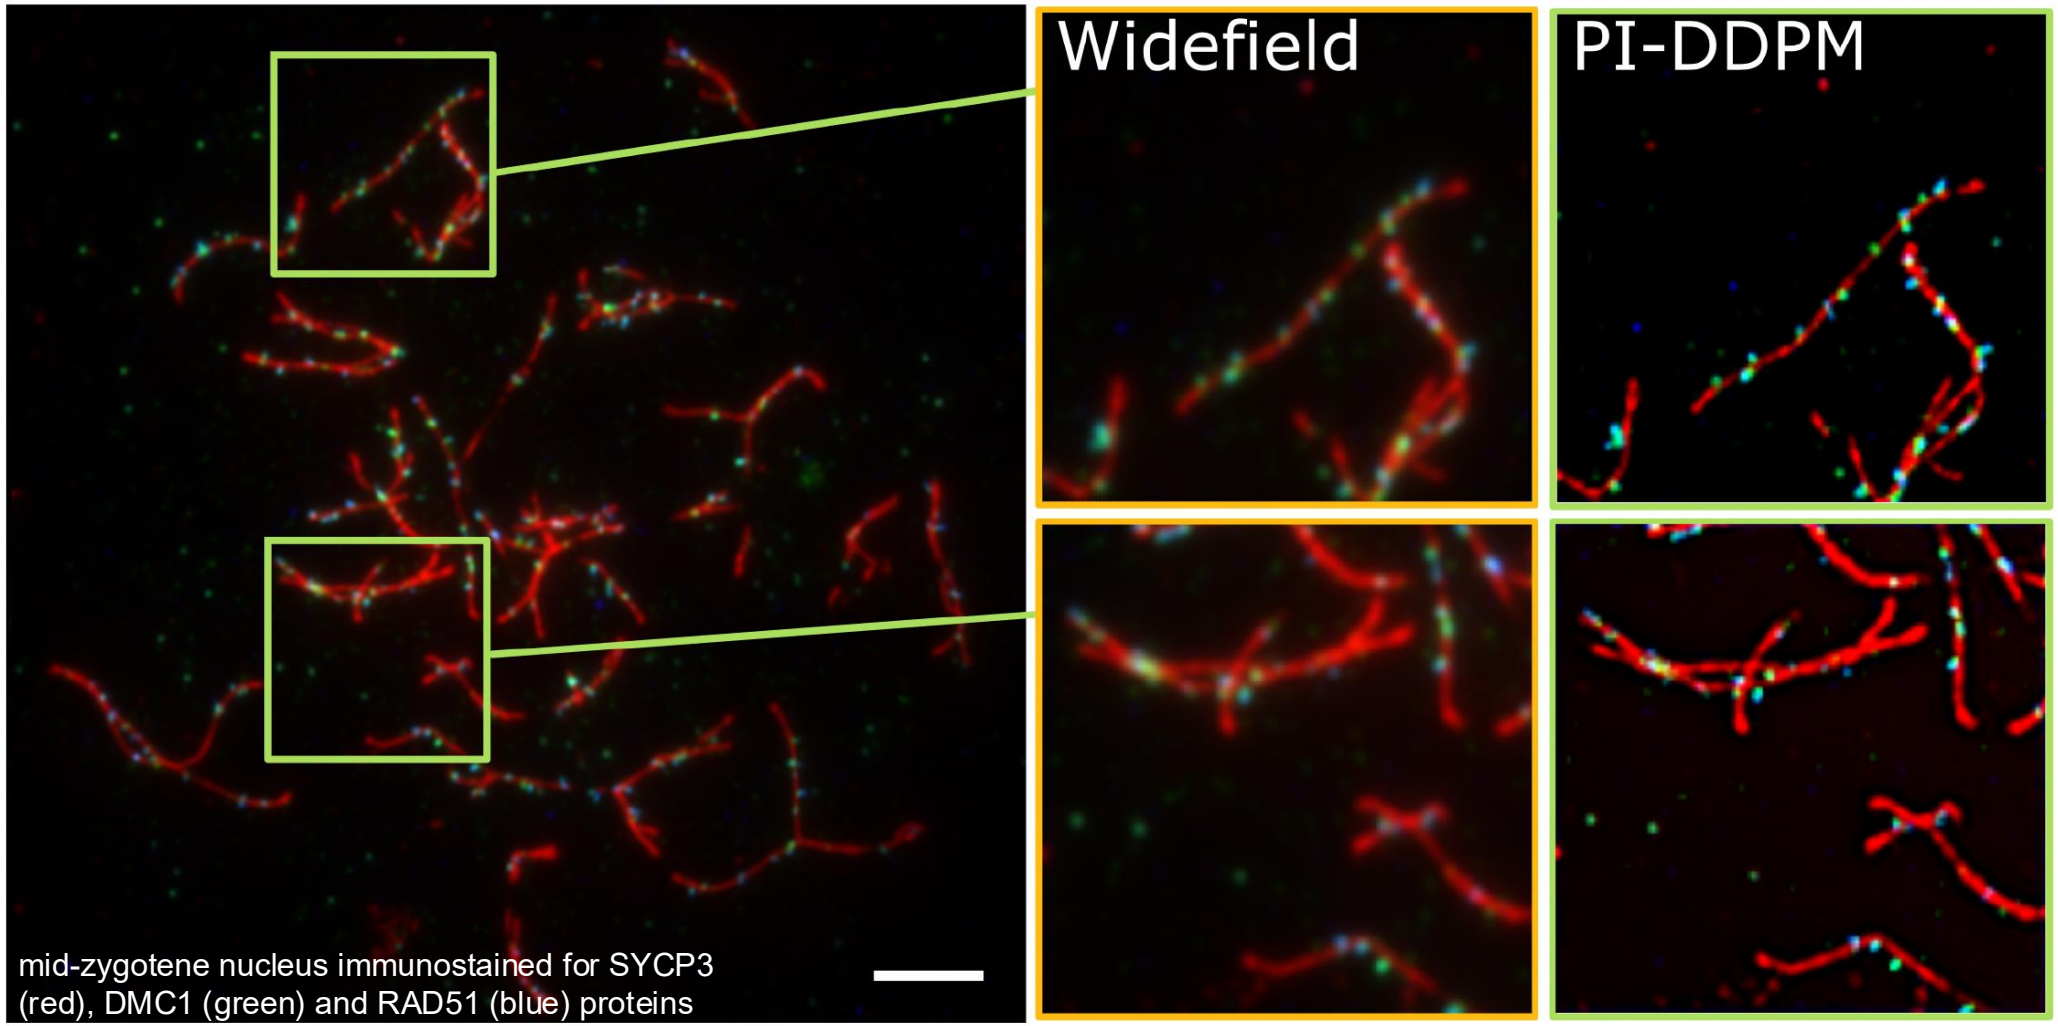

Li, della Maggiora & Yakimovich et al., 2024, Coms. Eng.

PI-DDPM Performance on Wild Data (dSTORM)

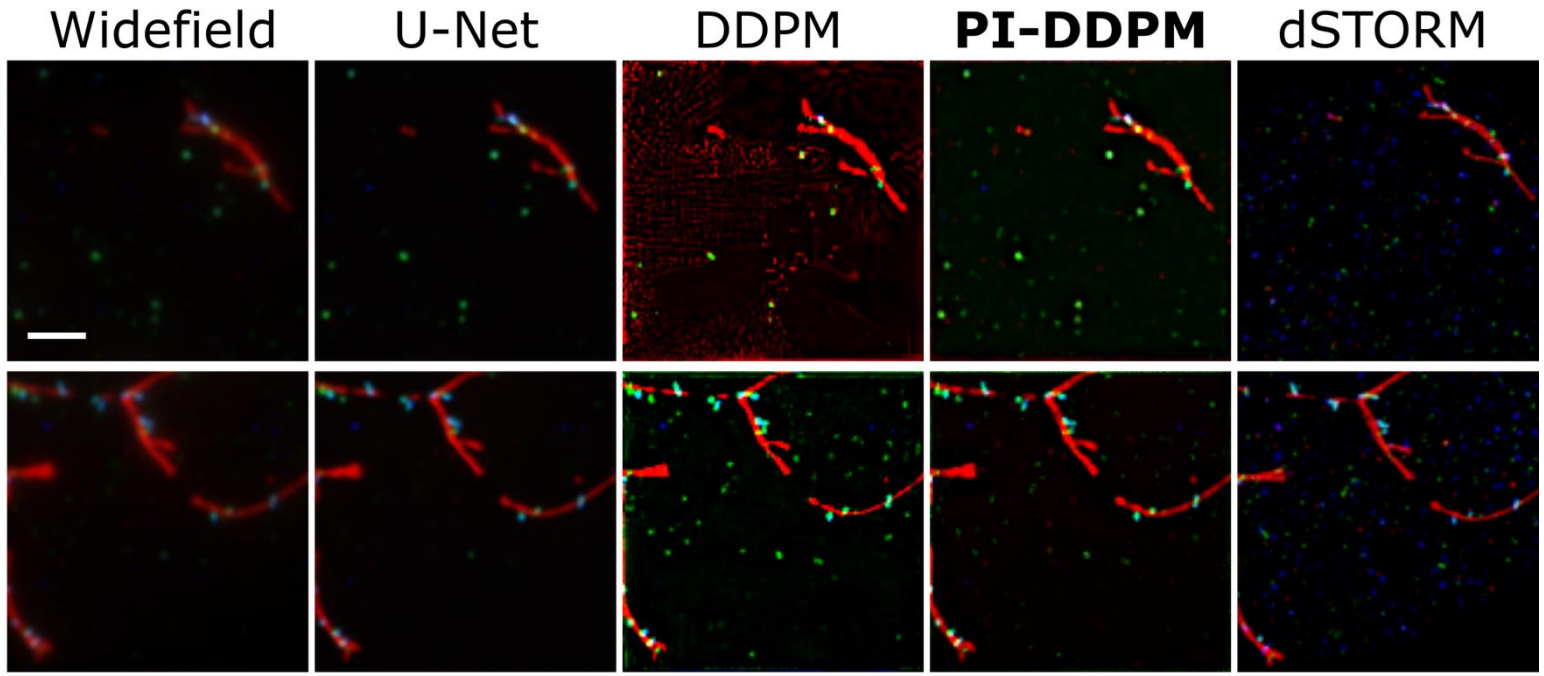

Table 2: Performance on dSTORM Test Set. Error stands for standard deviation.

Metric / Model	Original	U-Net	DDPM	PI-DDPM (ours)
PSNR	16.487	16.712	15.541±0.232	**16.778±0.807**
MS-SSIM	0.293	0.479	**0.638±0.007**	**0.612±0.039**
NRMSE	0.150	0.146	0.167±0.005	**0.145±0.015**

Li, della Maggiora & Yakimovich et al., 2024, Coms. Eng.

CVDM: Diffusion Schedule can be learned directly from the data

- noise added at each timestep is controlled by a variance schedule, instead of guessing it we propose to learn it from the data

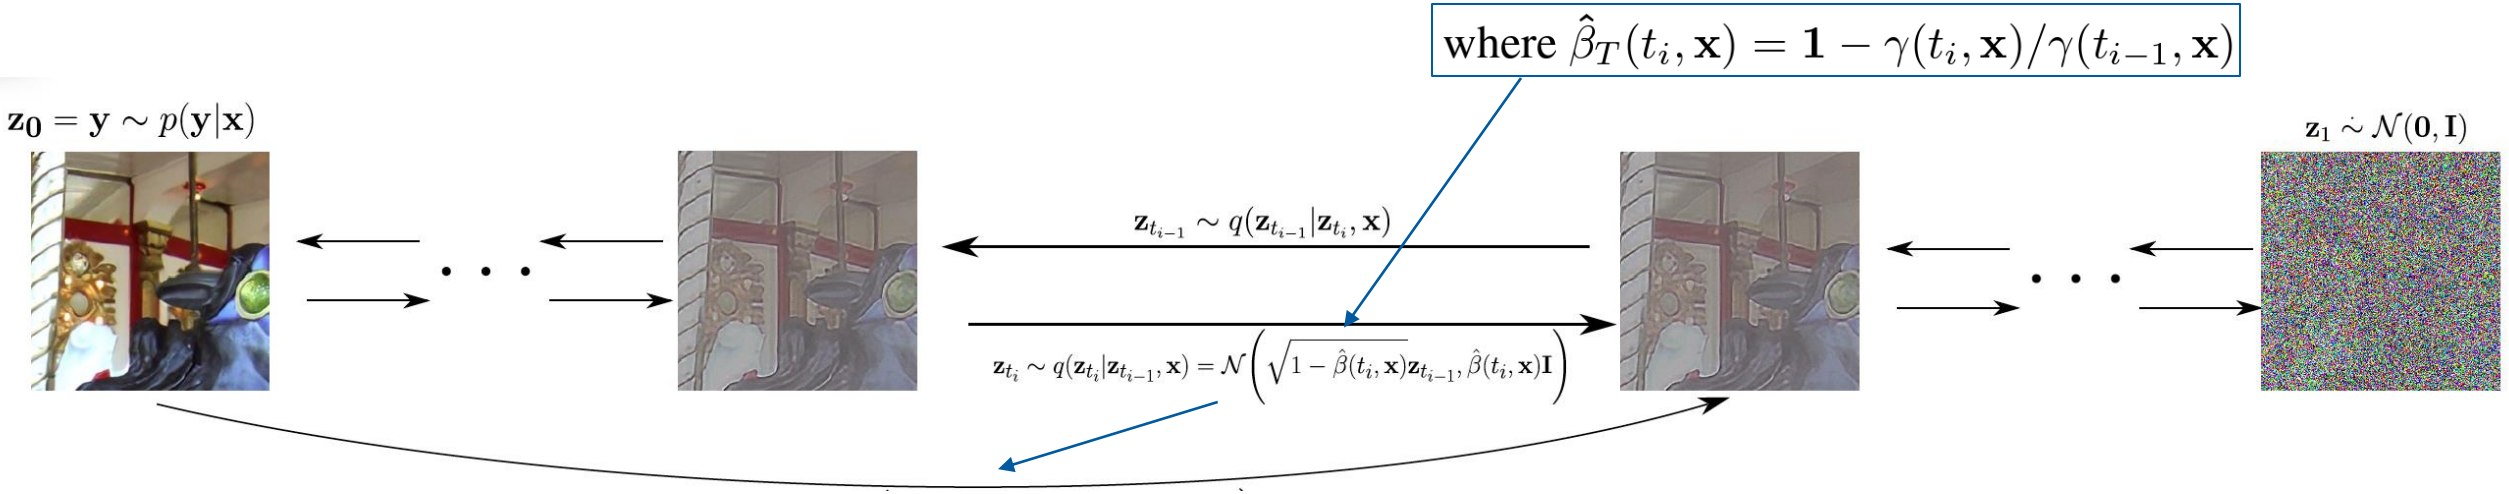

$$\mathbf{z}_{t_i} \sim q(\mathbf{z}_{t_i}|\mathbf{z}_0, \mathbf{x}) = \mathcal{N}\left(\sqrt{\gamma(t_i, \mathbf{x})}\mathbf{z}_0\,, (1-\gamma(t_i, \mathbf{x}))\mathbf{I}\right)$$

Gabriel della Maggiora

della Maggiora & Yakimovich et al., 2024, ICLR

CVDM: Diffusion Schedule can be learned directly from the data

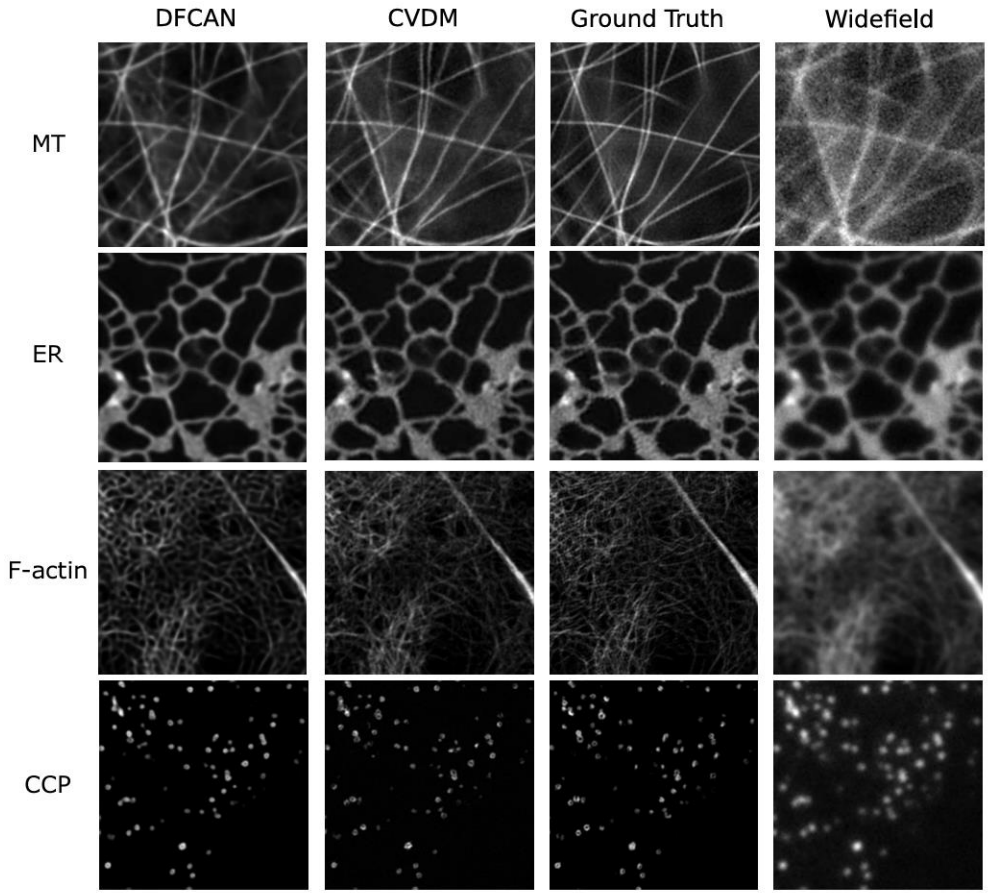

(a) CCP structures.

Metric / Model	DFCAN	CDDPM	CVDM (ours)
MS-SSIM (↑)	0.957	0.952	0.955
MAE (↓)	0.006	0.007	0.007
Resolution (nm) (↓)	107	100	**96**

(c) MT structures.

Metric / Model	DFCAN	CDDPM	CVDM (ours)
MS-SSIM (↑)	**0.901**	0.857	0.887
MAE (↓)	**0.033**	0.042	0.04
Resolution (nm) (↓)	127	101	**97**

(b) ER structures.

Metric / Model	DFCAN	CDDPM	CVDM (ours)
MS-SSIM (↑)	0.928	0.920	**0.934**
MAE (↓)	0.033	0.033	0.032
Resolution (nm) (↓)	165	157	**152**

(d) F-actin structures.

Metric / Model	DFCAN	CDDPM	CVDM (ours)
MS-SSIM (↑)	0.853	0.831	**0.863**
MAE (↓)	0.049	0.049	0.043
Resolution (nm) (↓)	151	104	**98**

ICLR

della Maggiora & Yakimovich et al., 2024, ICLR

HZDR CASUS CENTER FOR ADVANCED SYSTEMS UNDERSTANDING

CVDM: Diffusion Schedule can be learned directly from the data

- Schedule can be different for different parts of the image, which allows us to obtain uncertainty quantification

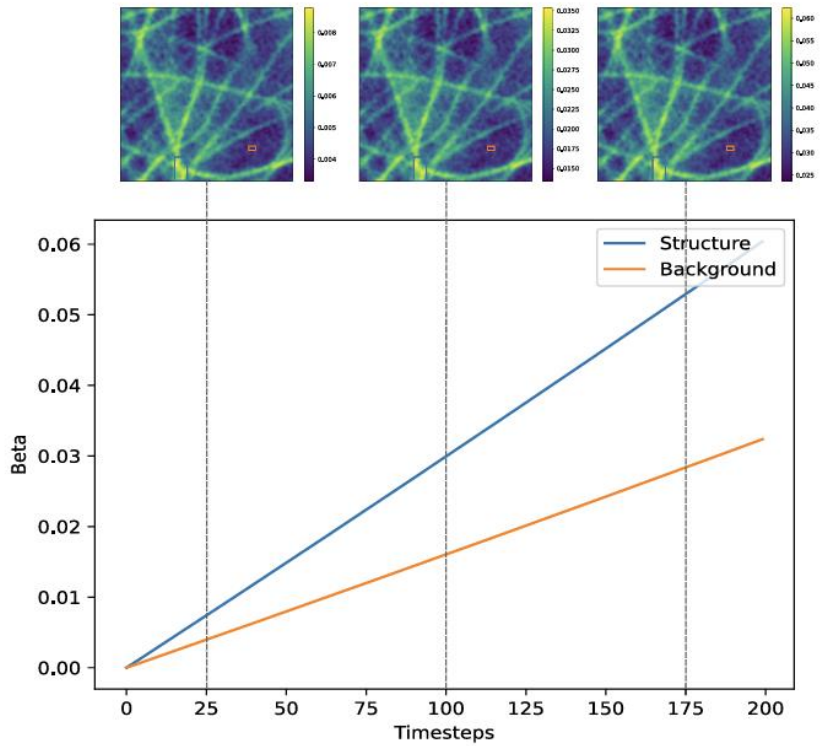

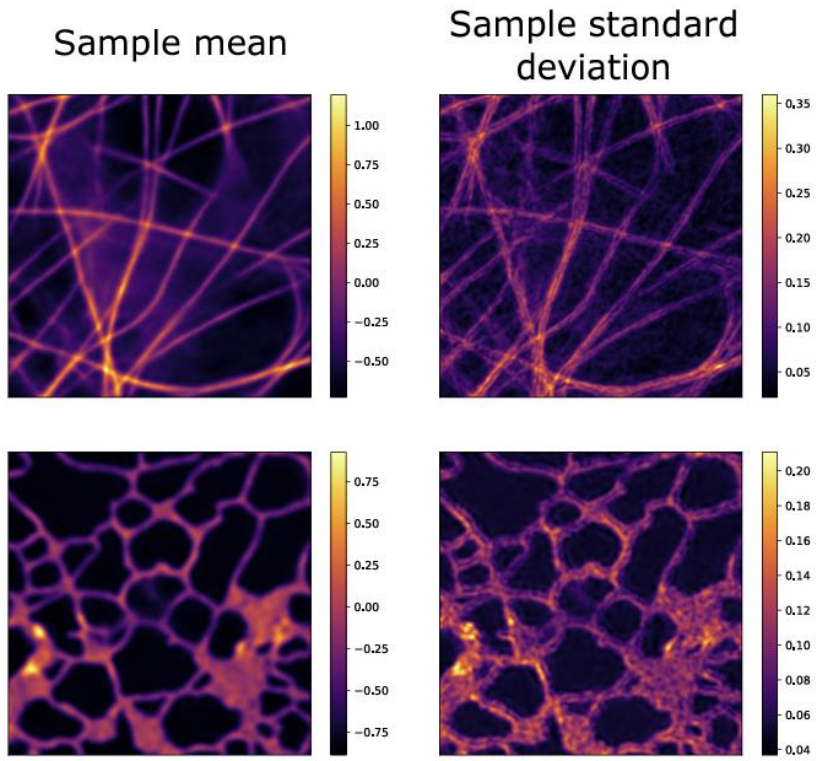

della Maggiora & Yakimovich et al., 2024, ICLR

Bringing Algorithms to Clinics with DL in UTI

HZDR HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

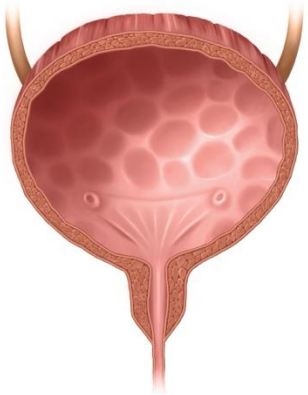

Nat. Rev. Urol.2010

- UTIs are among the most common bacterial infections worldwide
- 1 in 3 women will have at least 1 UTI by 24 years of age
- 40-50% of women will experience one UTI during their lifetime
- 44% experiencing recurrences
- Rate of emergency admissions due to UTI has almost doubled to 60/100,000 in the last five years.
- In 2012/13 unplanned admissions for UTIs cost £432 million per year.
- Leading cause of sepsis

Clinical Microscopy of Urine Samples

Adrian Urbański

Trina De

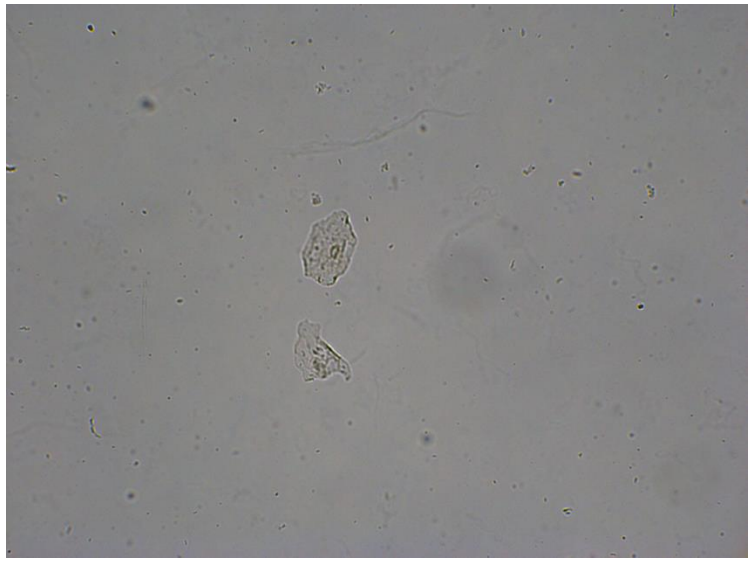

Barrier: Objects in urine are translucent!

UCL

Dr. Harry Horsley

HZDR CASUS CENTER FOR ADVANCED SYSTEMS UNDERSTANDING

Light passing through the Specimen Changes its Phase

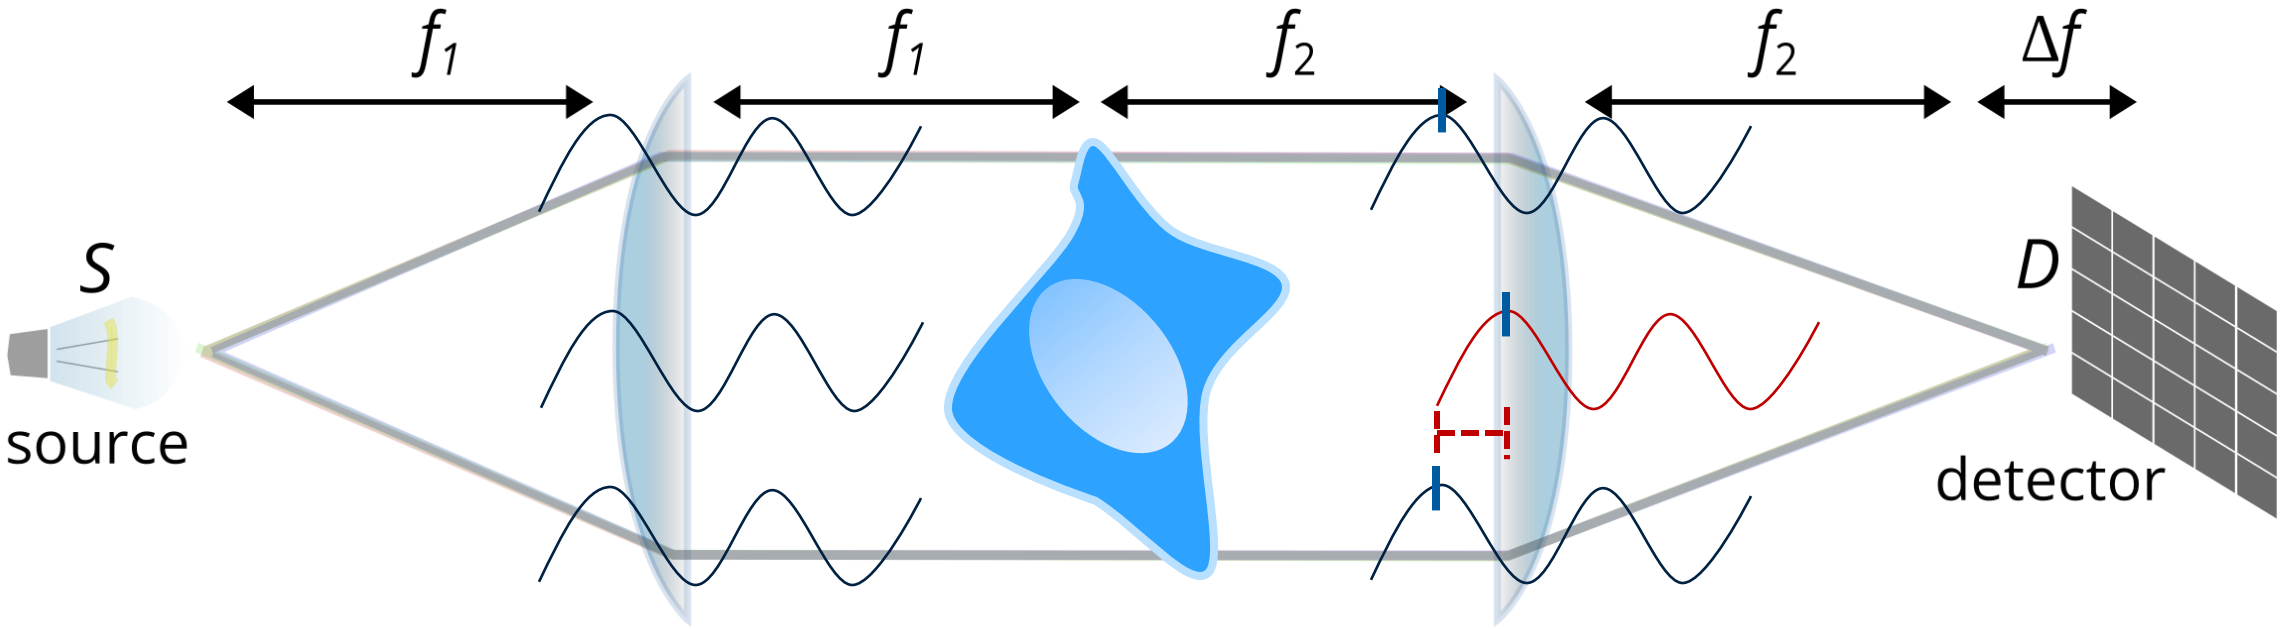

Gabriel della Maggiora

della Maggiora & Yakimovich et al., 2025, AAAI

Light passing through the Specimen Changes its Phase

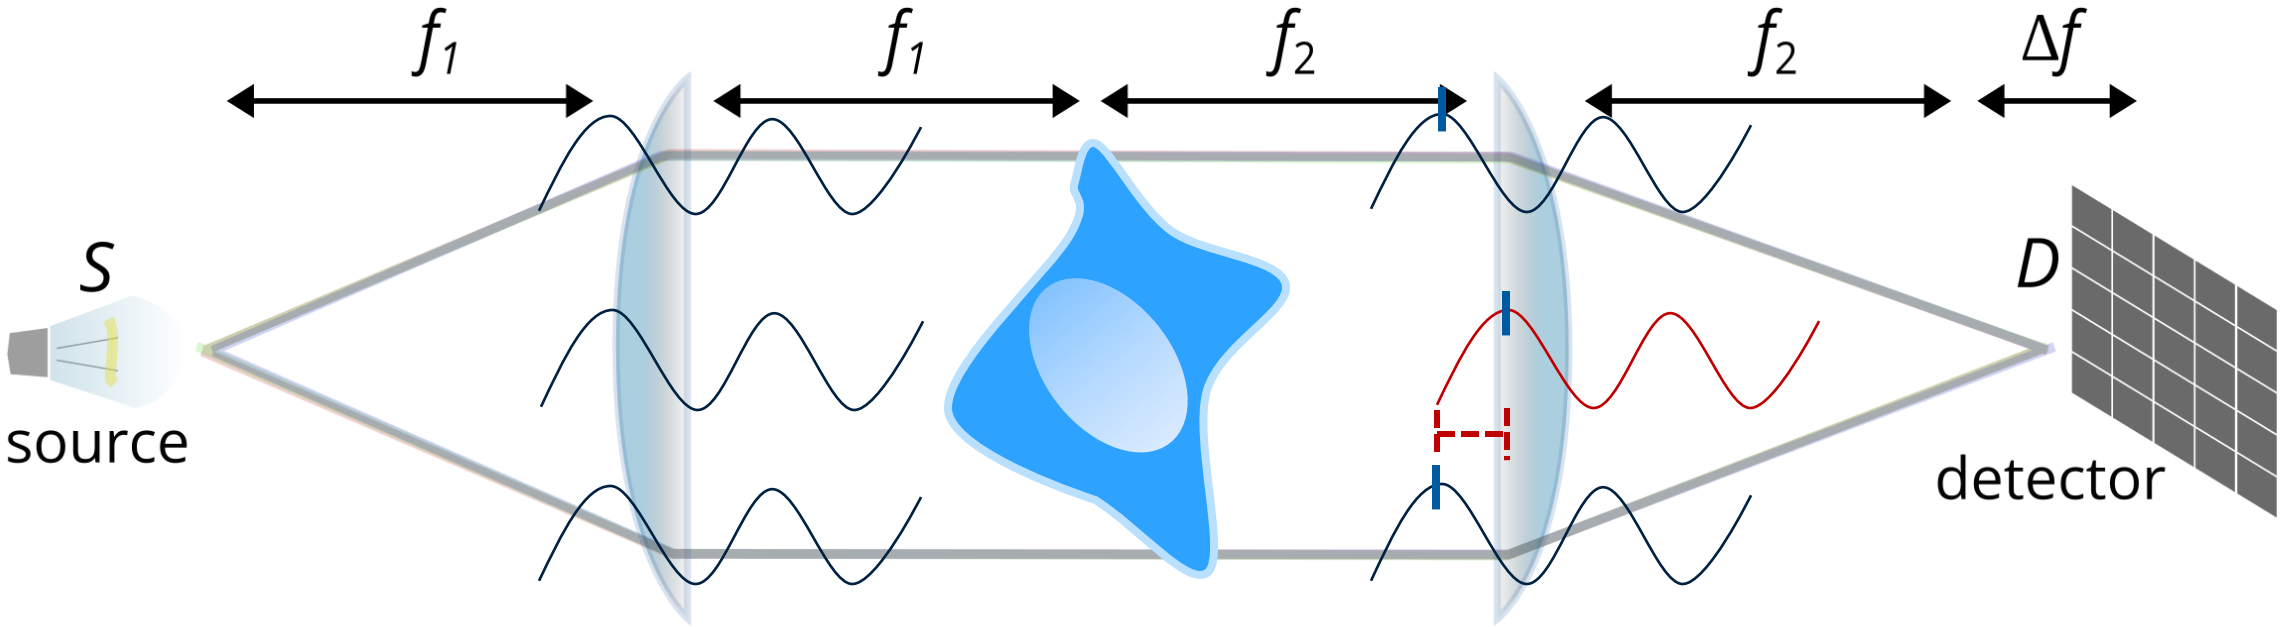

- The phase shift can be captured by

<u>Hardware</u>:

Phase-Contrast Microscopy

Quantitative Phase Imaging (QPI)

- <u>Software</u> (data from multiple Z-planes):

Transport-of-Intensity Equation (TIE):

$$-k\frac{\partial I(x,y;z)}{\partial z} = \nabla_{(x,y)}\cdot\left[I(x,y;z)\nabla_{(x,y)}\varphi(x,y;z)\right]$$

where (x,y) are spatial coordinates, z is the defocus distance, k is a wave number ($2\boldsymbol{\pi}/\boldsymbol{\lambda}$), $\boldsymbol{\varphi}$ – phase;

Quantitative Phase can be Retrieved from Chromatic Aberrations

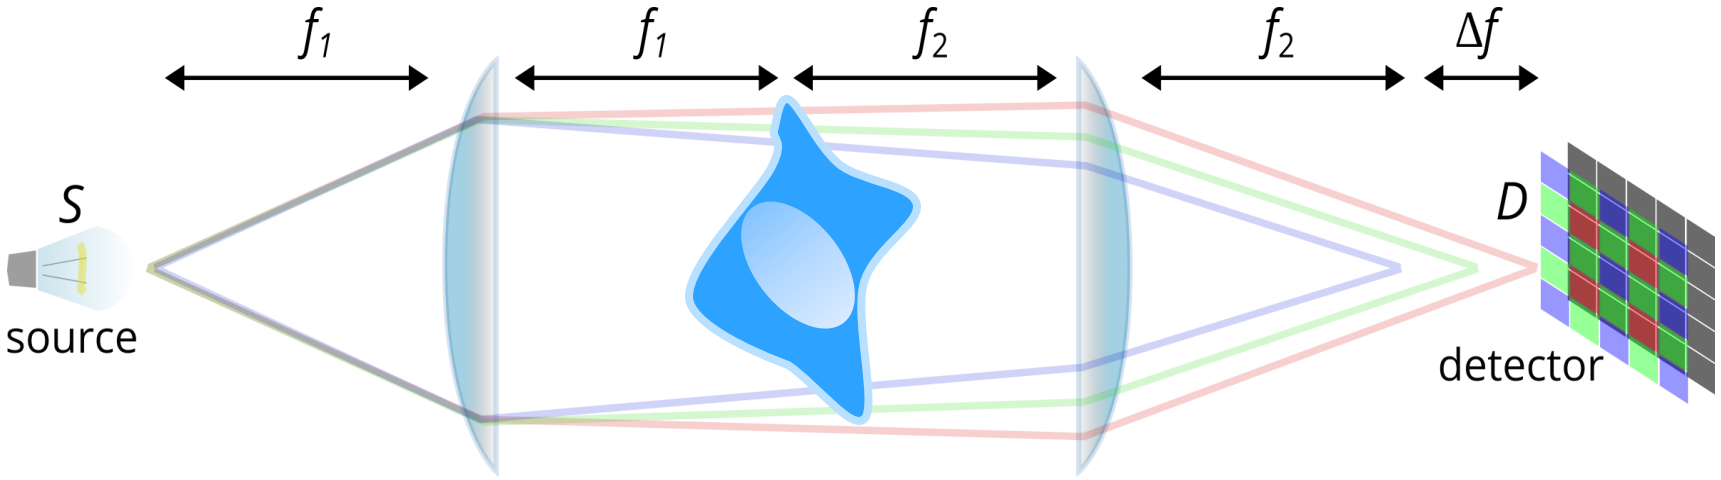

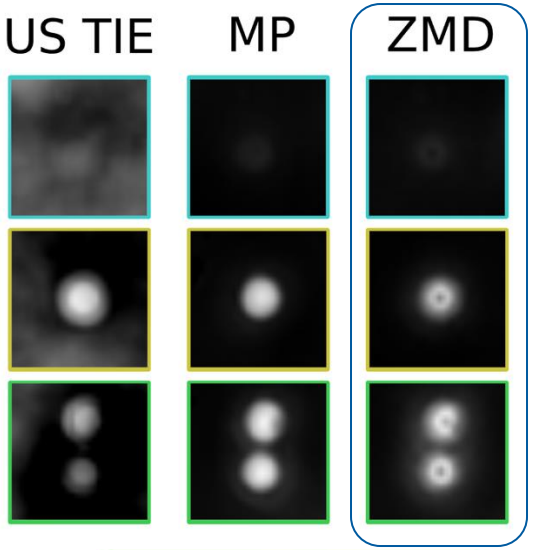

*Given that **z** and **λ** can be connected as **ξ** and x,y as **x** TIE can be reformulated as:*

$$-2\pi \frac{\partial I(\mathbf{x};\xi)}{\partial \xi} = \nabla_{\mathbf{x}} \cdot \left[I(\mathbf{x};\xi) \nabla_{\mathbf{x}} \varphi(\mathbf{x};\xi) \right]$$

Solution can be obtained using CVDM

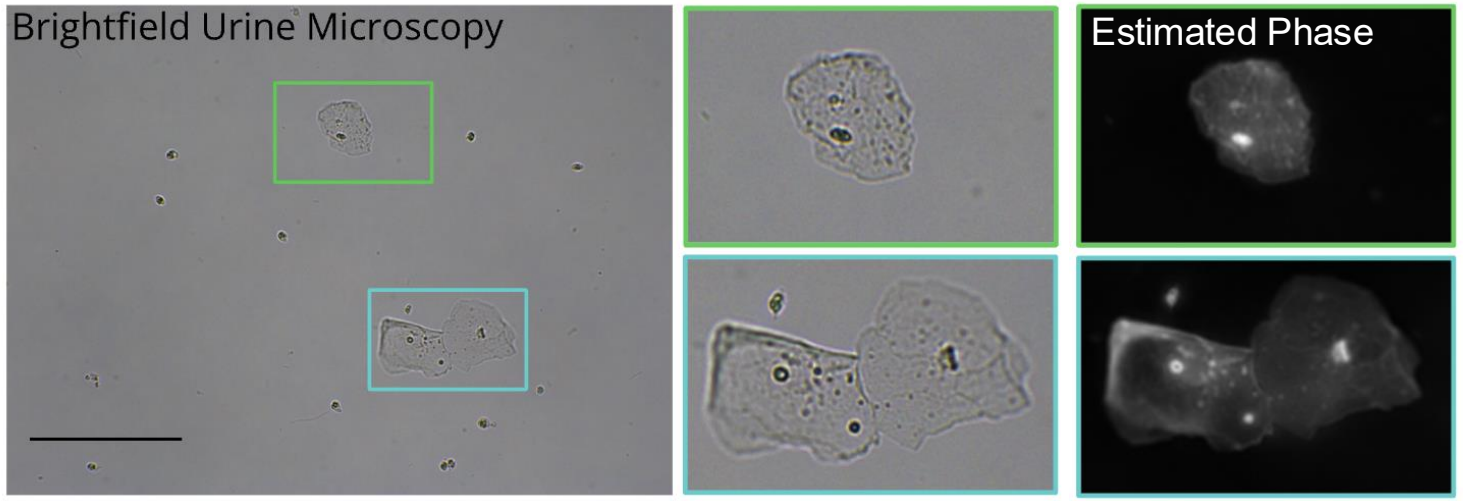

della Maggiora & Yakimovich et al., 2025, AAAI

Quantitative Phase can be Retrieved from Chromatic Aberrations

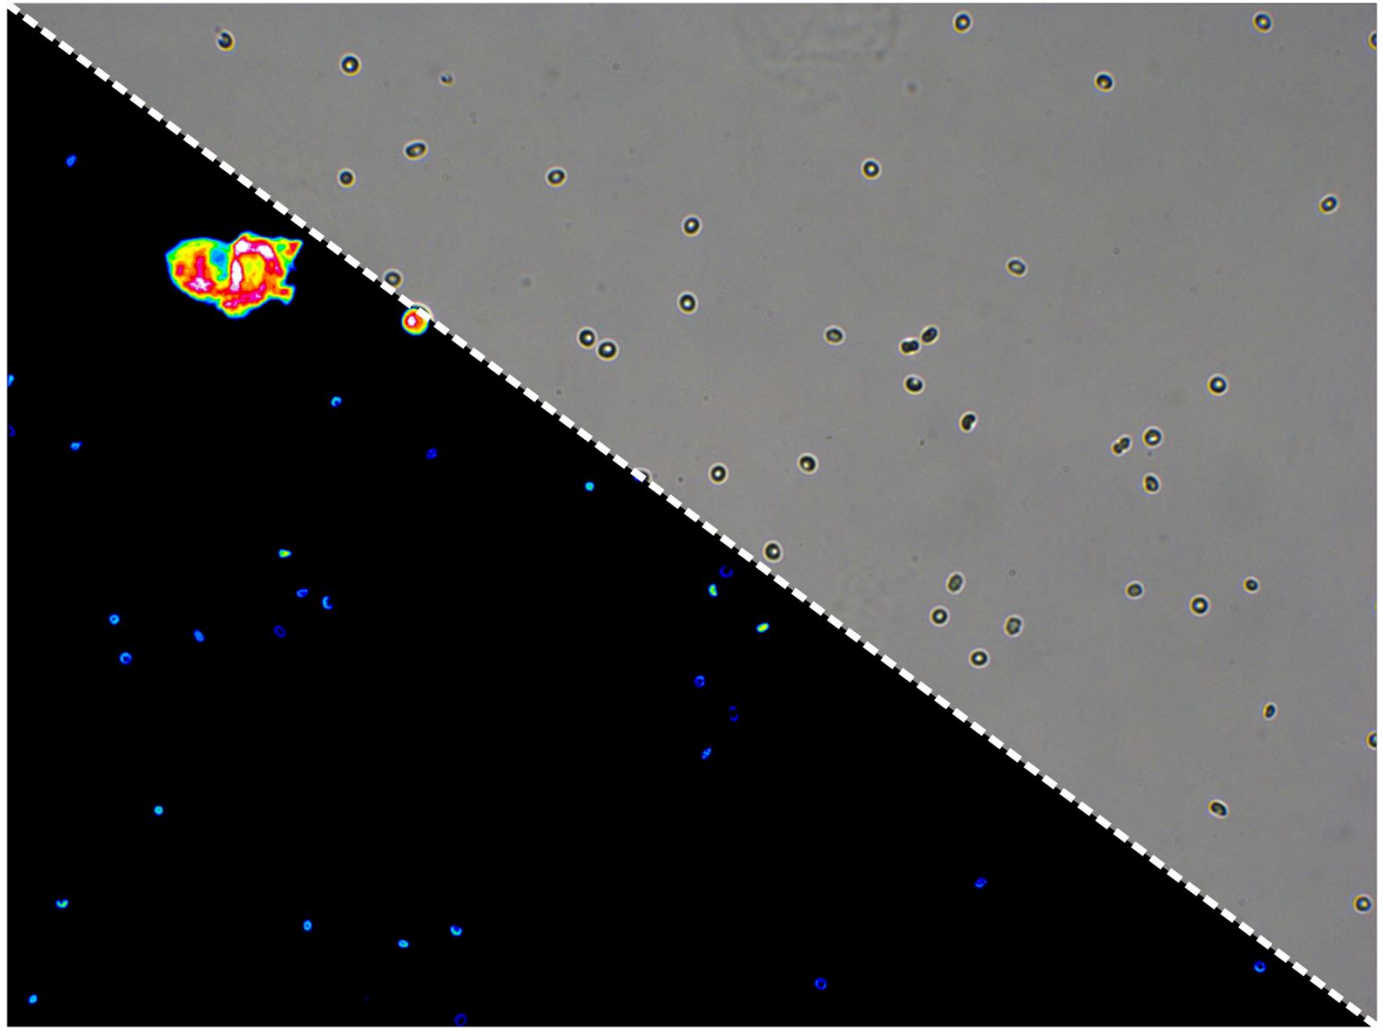

della Maggiora & Yakimovich et al., 2025, AAAI

Take-home Messages I:

- The transformation from widefield microscopy to SRM can be formulated as an inverse problem and solved using generative deep learning with SRM as GT.
- Diffusion models achieve state-of-the-art performance on this task.
- Their performance can be further improved by incorporating a physical model of image acquisition into the loss function.
- One of the trickiest hyperparameters in diffusion models – the noise schedule - can be learned from the data directly.
- As a side benefit, this parameter can provide uncertainty of the model on a per-pixel basis.
- Phase can be captured by combining generative AI and a simple RGB camera, allowing rapid application in clinics.

Recent AI Model Research has been Dominated by Scale

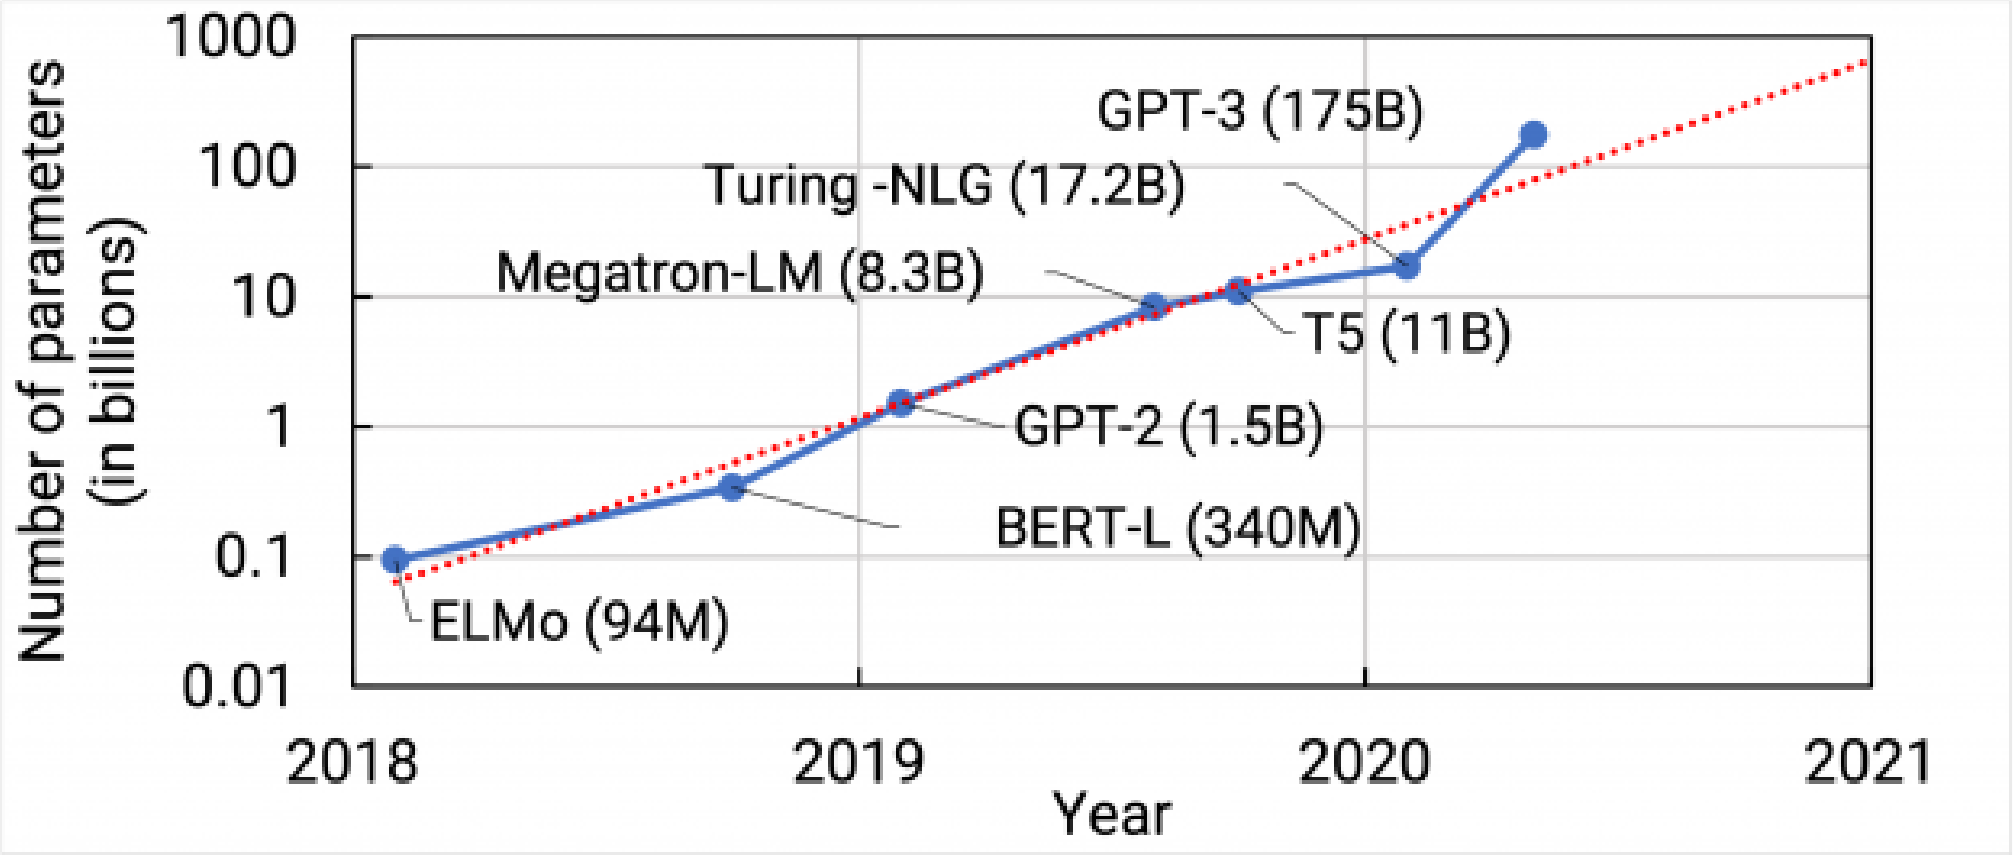

source: nvidia.com

Inverse Problems can also be Addressed by Regression

convolution noise

$$b = A(u) + \varepsilon$$

Widefield GT

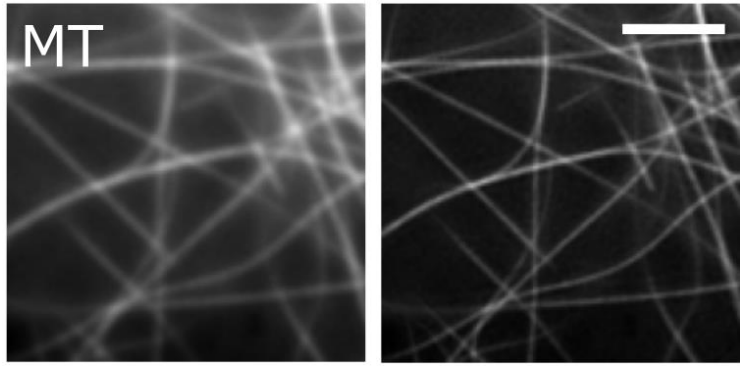

We can approximate

$$u \approx \overline{u} = \underbrace{\left(K^T K + \varepsilon I\right)^{-1}} K^T \overline{b}$$

Pseudo-differentiable operator

Inverse convolution becomes:

$$A^{-1} \approx \left(K^T K + \varepsilon I\right)^{-1} K^T$$

Next, we could use Beylkin, Coifman, and Rokhlin (BCR) Representation to perform wavelet decomposition:

$$A^{(l+1)} = W^{(l)} \begin{pmatrix} D_1^{(l)} & D_2^{(l)} \\ D_3^{(l)} & A^{(l)} \end{pmatrix} \left(W^{(l)}\right)^T$$

Multi-level decomposition:

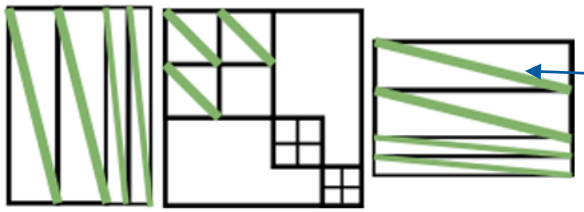

band matrices with values in green lines

Rui Li

Li & Yakimovich et al., ECCV 2024

BCR Decomposition can be Approximated by NN

Fan et al. 2019 have shown that

$$conv[2, 2p, 2](\xi)$$

can replace the multiplication with $\left(W^{(l)}\right)^T$

The multiplication with the matrix of $A(l)$

can approximated with local convolution operators

$$LC[2, nb, 1](\xi)$$

The truncation $u(L_0)$ can be represented as dense connection

$$Dense^2\,[1, 1]\left(V^{L_0}\right)$$

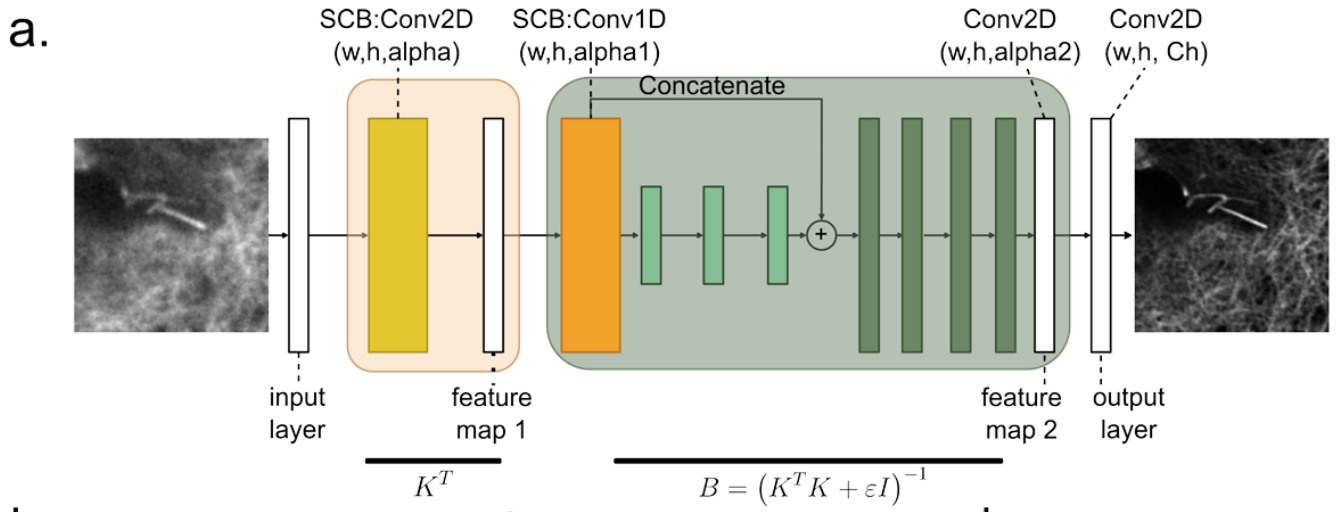

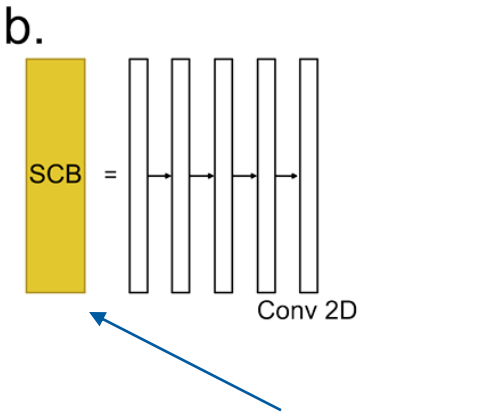

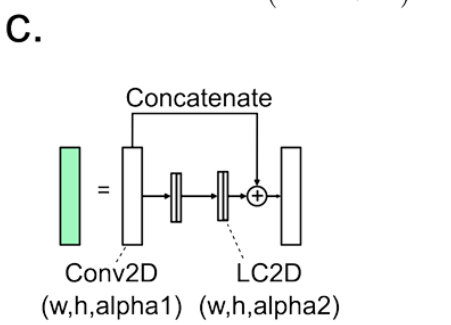

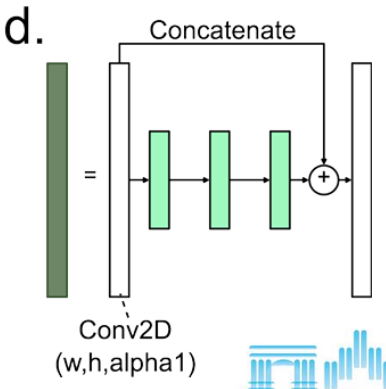

Rui Li

Li & Yakimovich et al., ECCV 2024

Multi-Stage Residual BCR Net (m-rBCR)

The m-rBCR model reintegrates the input image from other resolutions as posterior for the pseudo-differential operator

$$B = \left(K^T K + \varepsilon I\right)^{-1}$$

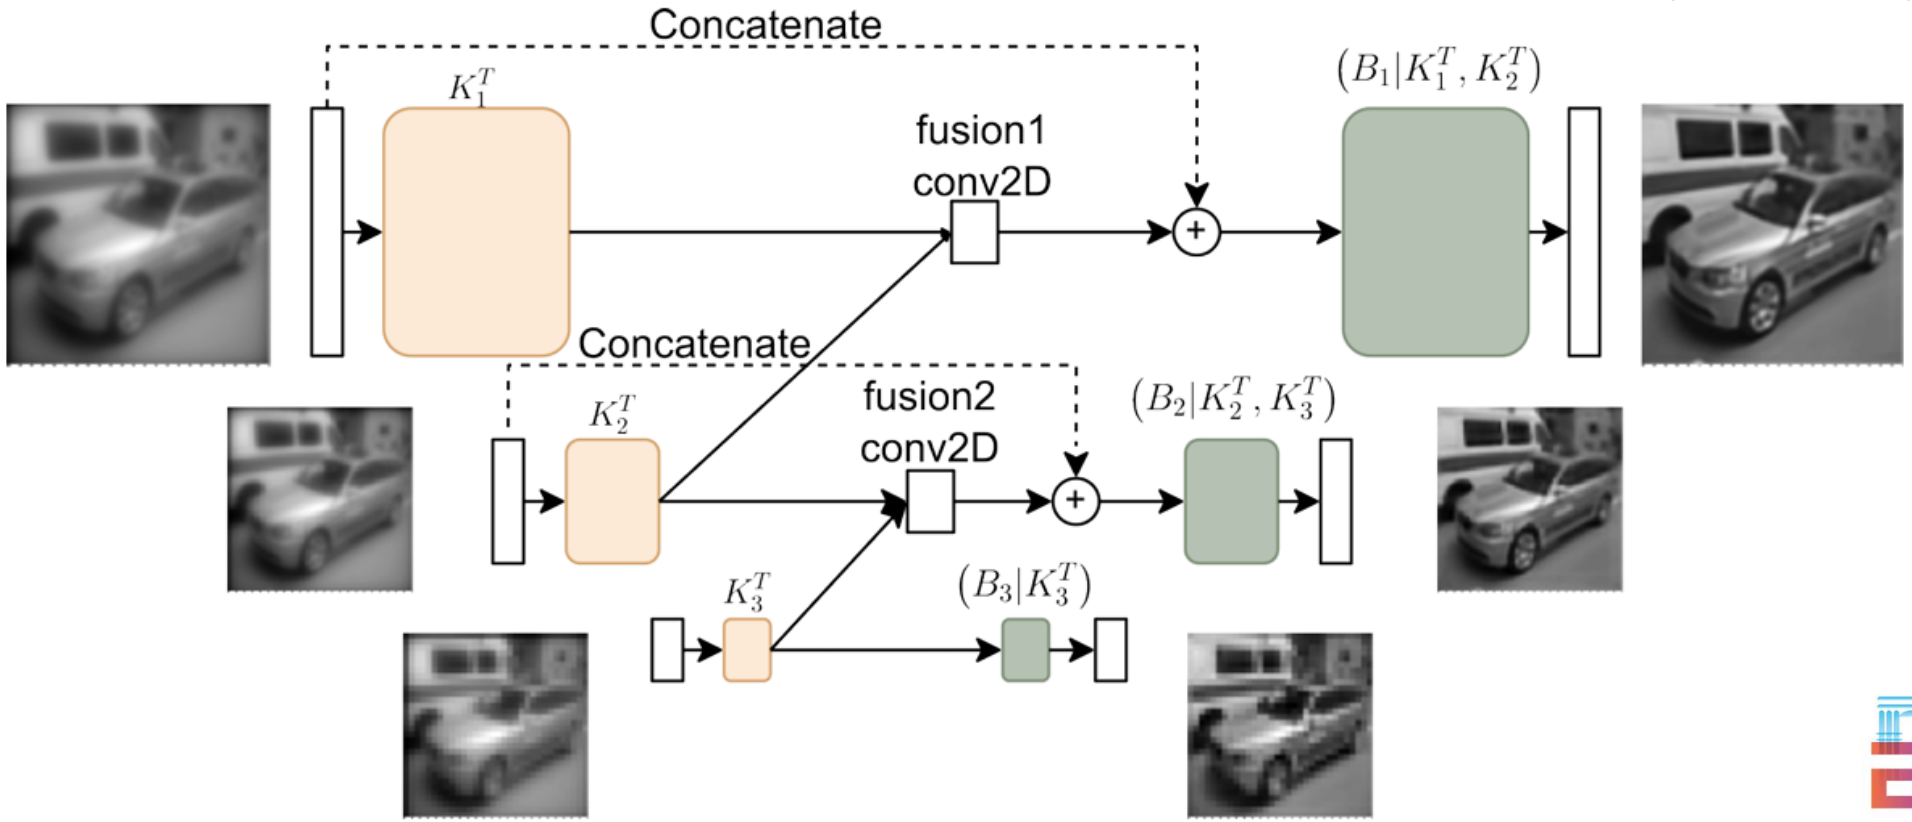

Rui Li

Li & Yakimovich et al., ECCV 2024

The Lean m-rBCR Performs 2nd Best with Fraction of Parameters

Model	Params. (↑)	Runtime	BioSR(sim.)		ImageNet(sim.)		W-C(real)		dSTORM(real)
			PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR
RL	N/A	0.0835	13.39	0.48	12.74	0.78	N/A	N/A	N/A
s-rBCR	0.086 (ours)	**0.0019**	20.28	0.59	20.84	0.83	17.67	0.57	19.55
m-rBCR	0.237 (ours)	0.0036	**24.89**	**0.78**	21.41	0.86	**23.10**	0.70	**20.13**
CARE	0.333	0.0039	22.15	0.65	18.47	0.74	21.53	0.66	19.62
DnCNN	0.556	0.0056	21.41	0.70	19.70	0.84	19.34	0.63	17.46
Noise2Noise	1.227	0.0068	16.07	0.57	16.24	0.60	15.07	0.57	18.06
MIMO-U-Net	6.807	0.0087	23.95	**0.78**	**22.35**	**0.88**	19.17	0.67	18.91
U-Net	7.780	0.0241	21.89	0.73	19.23	0.75	18.17	0.63	19.62
RCAN	15.334	0.0281	21.71	0.64	19.78	0.91	20.51	0.58	19.26
MPRNet	20.127	0.0173	21.44	0.63	20.12	0.83	21.53	0.55	18.54
DDPM L_2	23.988	1.0968	21.85	0.65	20.22	0.78	22.27	**0.73**	17.46
ESRGAN	49.841	0.0147	19.82	0.59	18.95	0.76	21.17	0.59	19.93

Rui Li

Li & Yakimovich et al., ECCV 2024

The Lean m-rBCR Performs 2nd Best with Fraction of Parameters

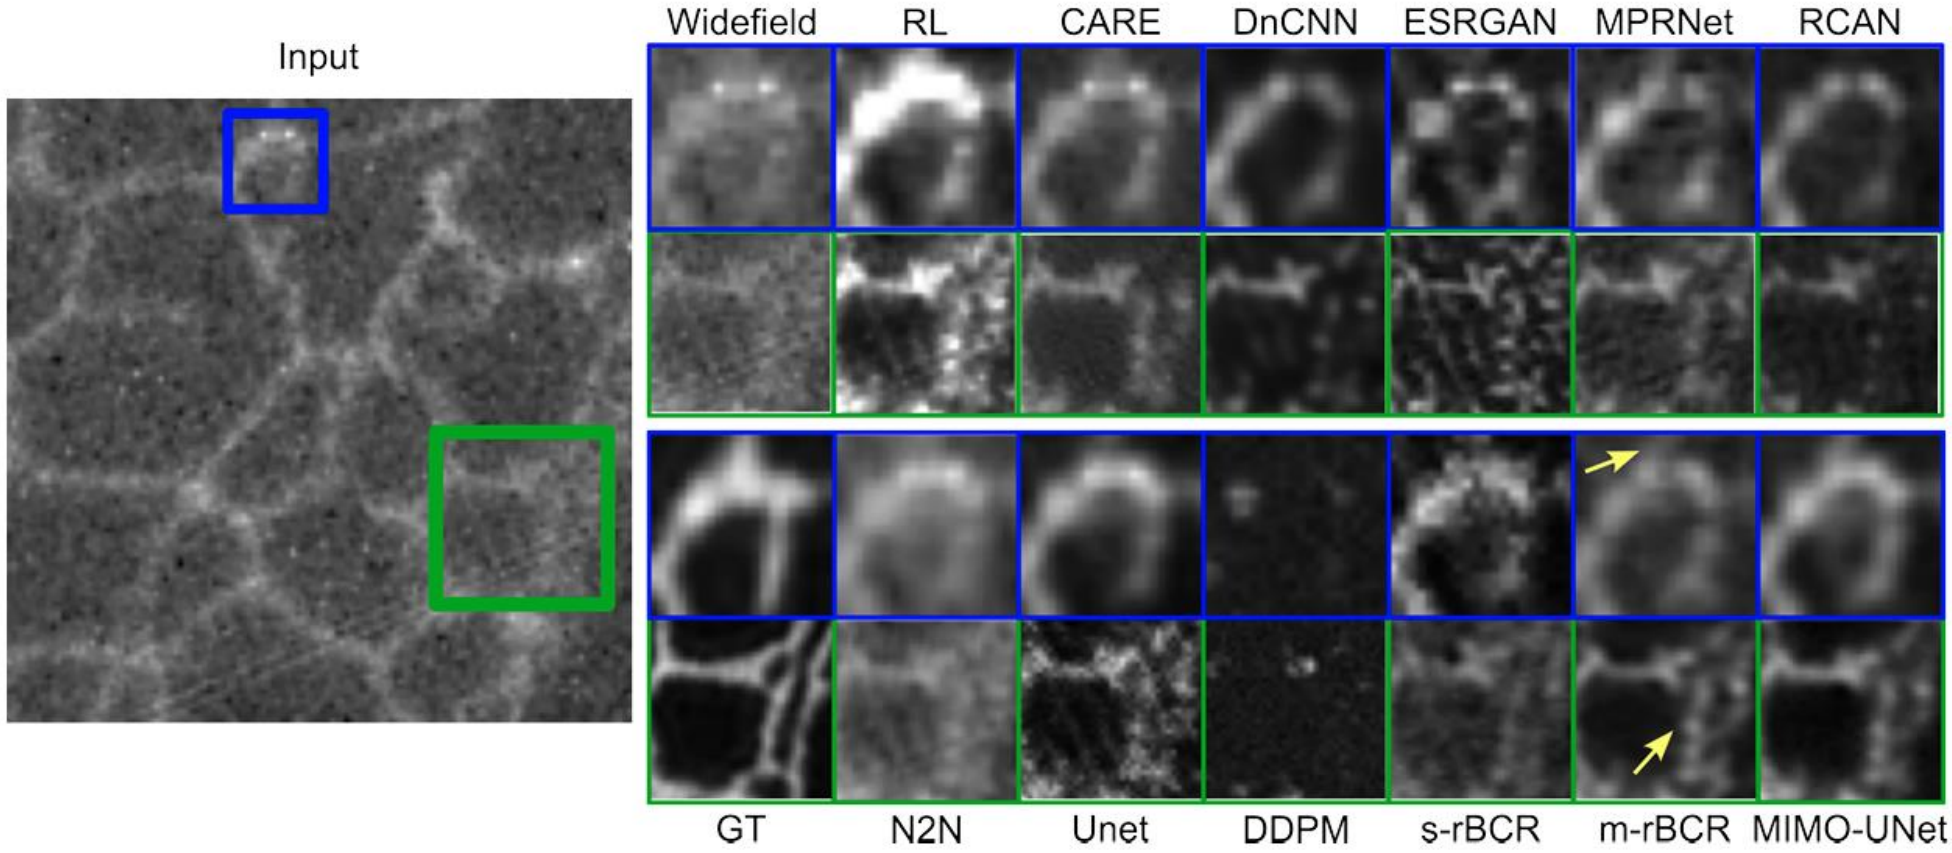

Rui Li

Li & Yakimovich et al., ECCV 2024

Take-home Messages II:

- Smaller Deep Learning models can be competitive while requiring a fraction of compute and training data.
- Microscopy images may require representations different from natural images, so alternative representations should be explored.

BioMed Image Analysis: Who has time for data annotation?

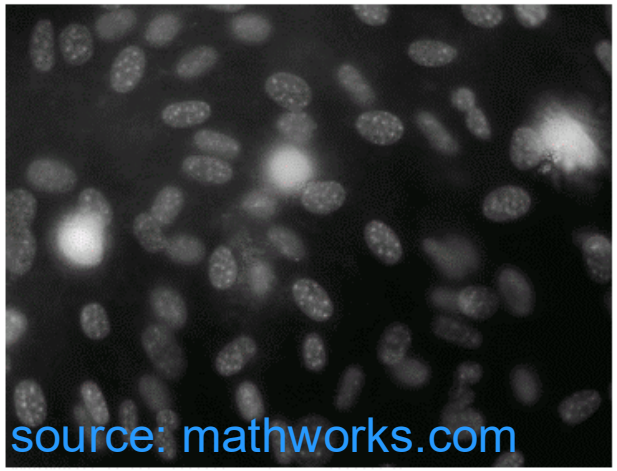

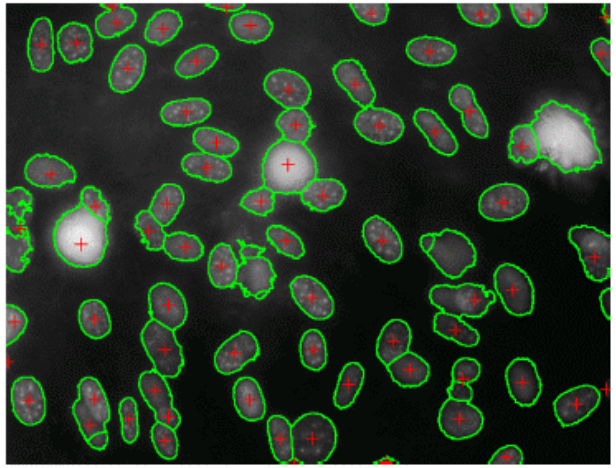

?

Binocular microscopy can turn into a 3D imaging with DL

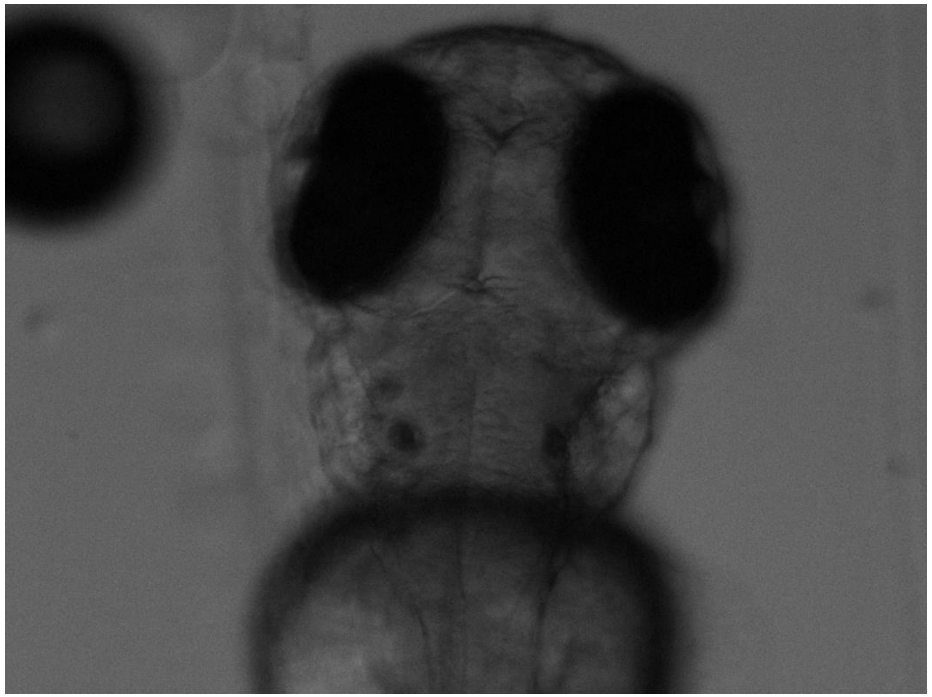

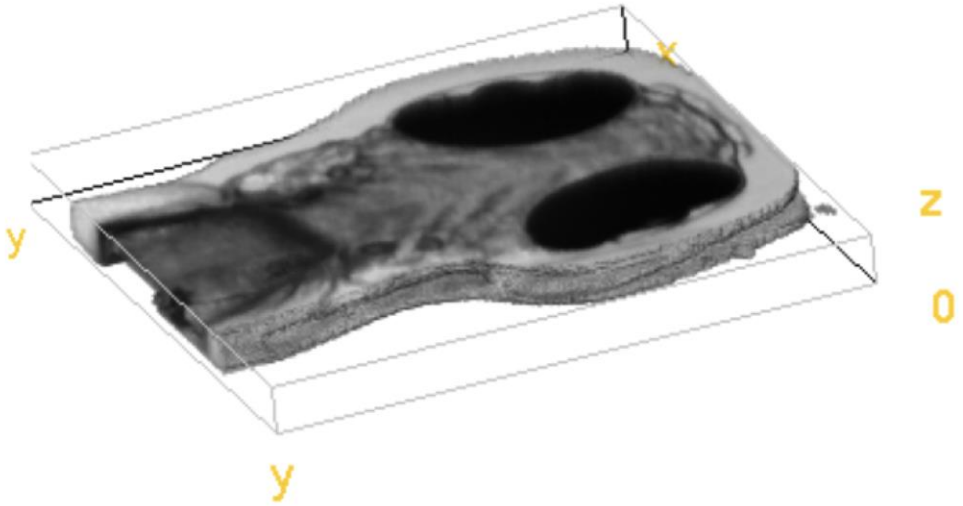

3D representation

Rui Li

Li & Yakimovich et al., 2024, Sci. Rep.

HZDR

CASUS
CENTER FOR ADVANCED SYSTEMS UNDERSTANDING

In-focus Pixels can be Selected using Advanced Microscopy

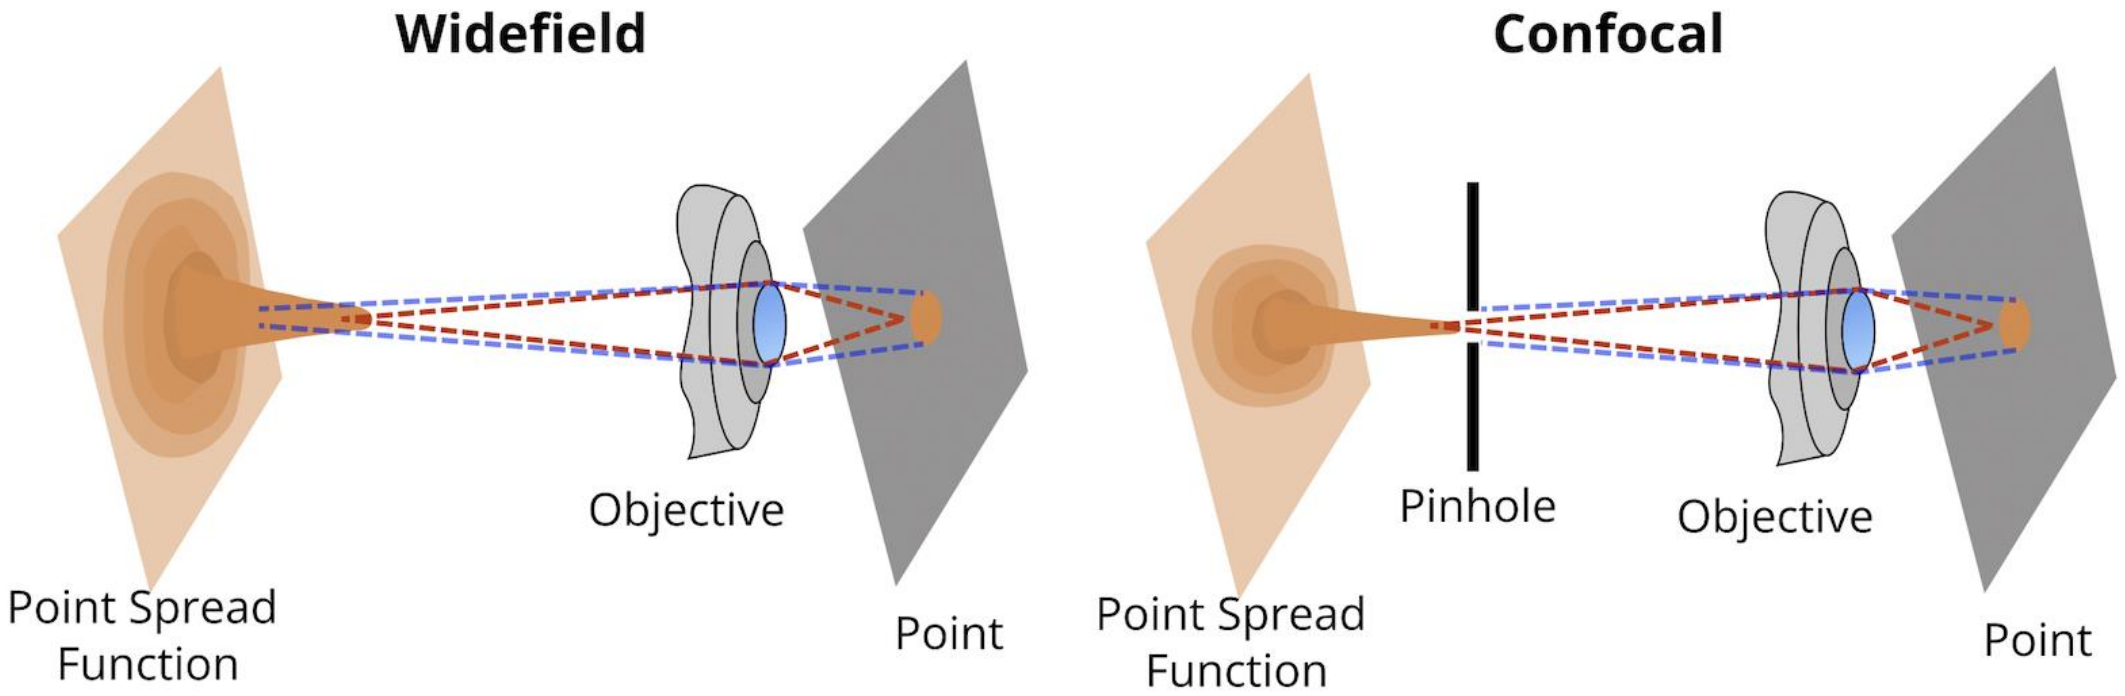

Rui Li

Li & Yakimovich et al., 2024, Sci. Rep.

Lateral change in contrast can be detected with Autofocus algorithms

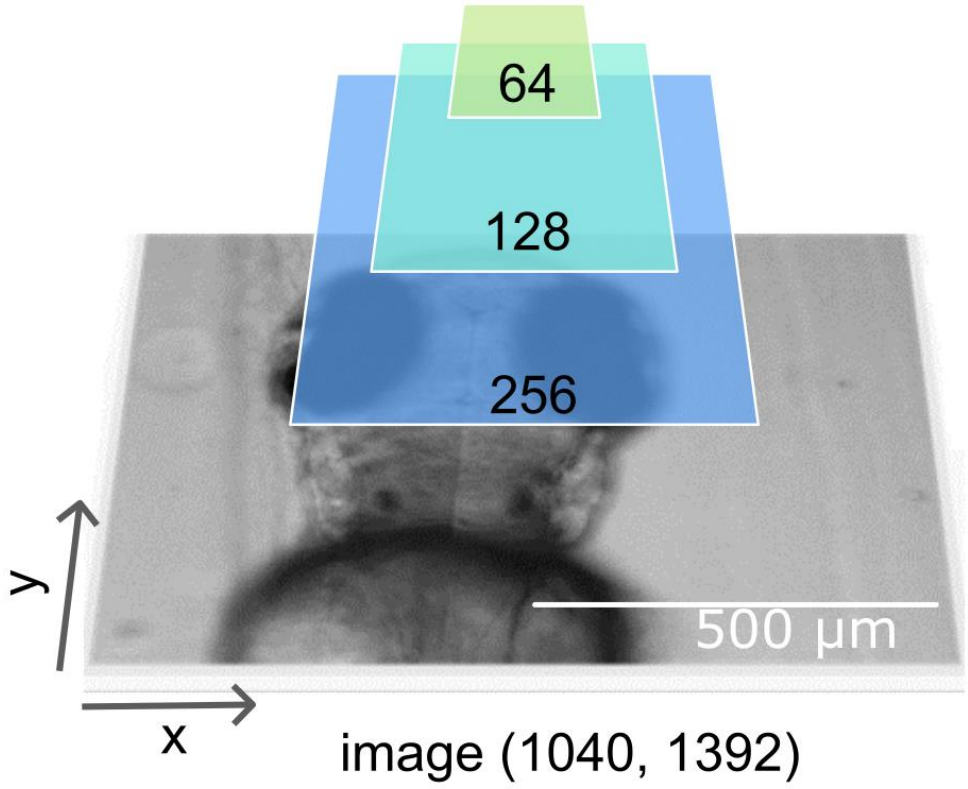

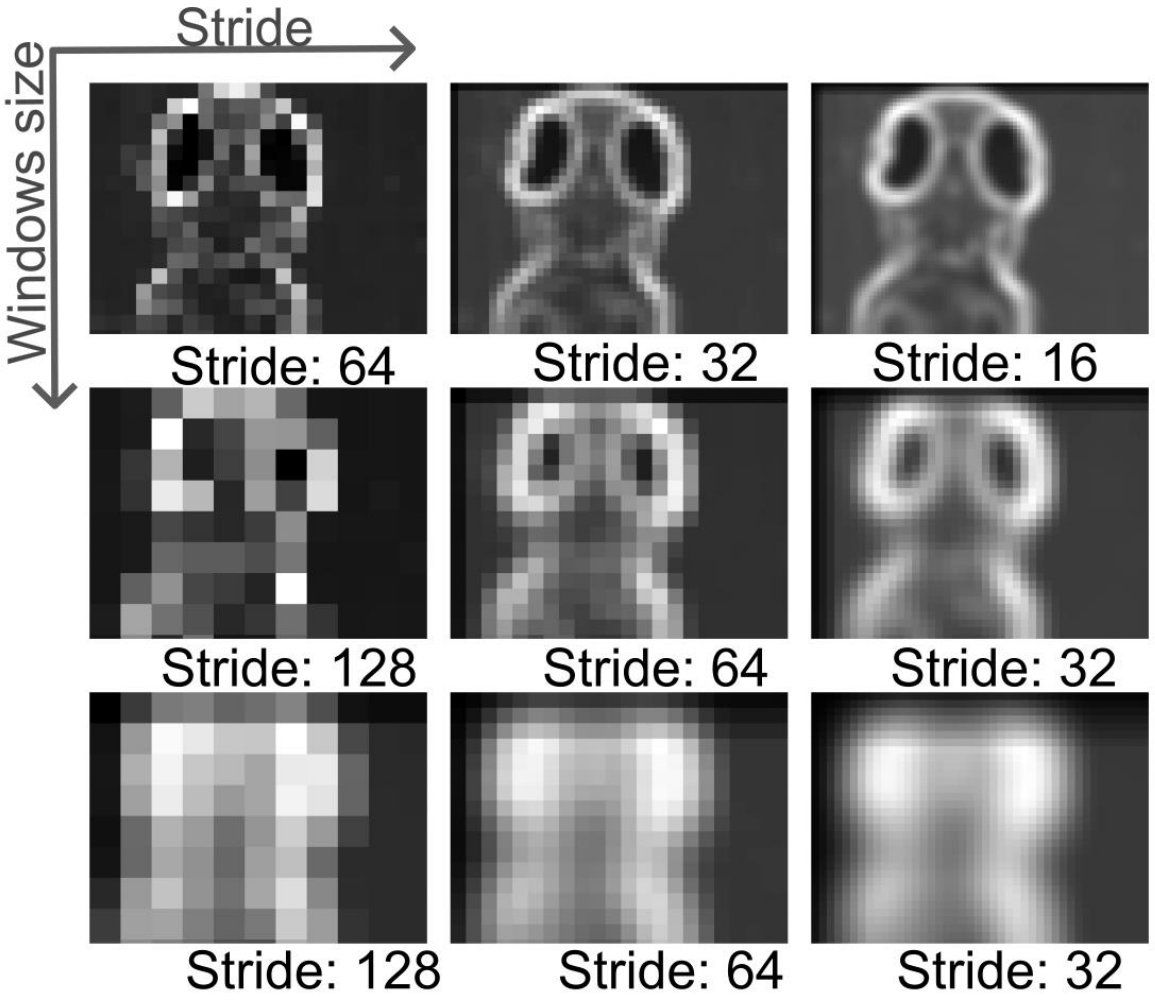

Li & Yakimovich et al., 2024, Sci. Rep.

Rui Li

Weak-labeling pipeline as a CNN

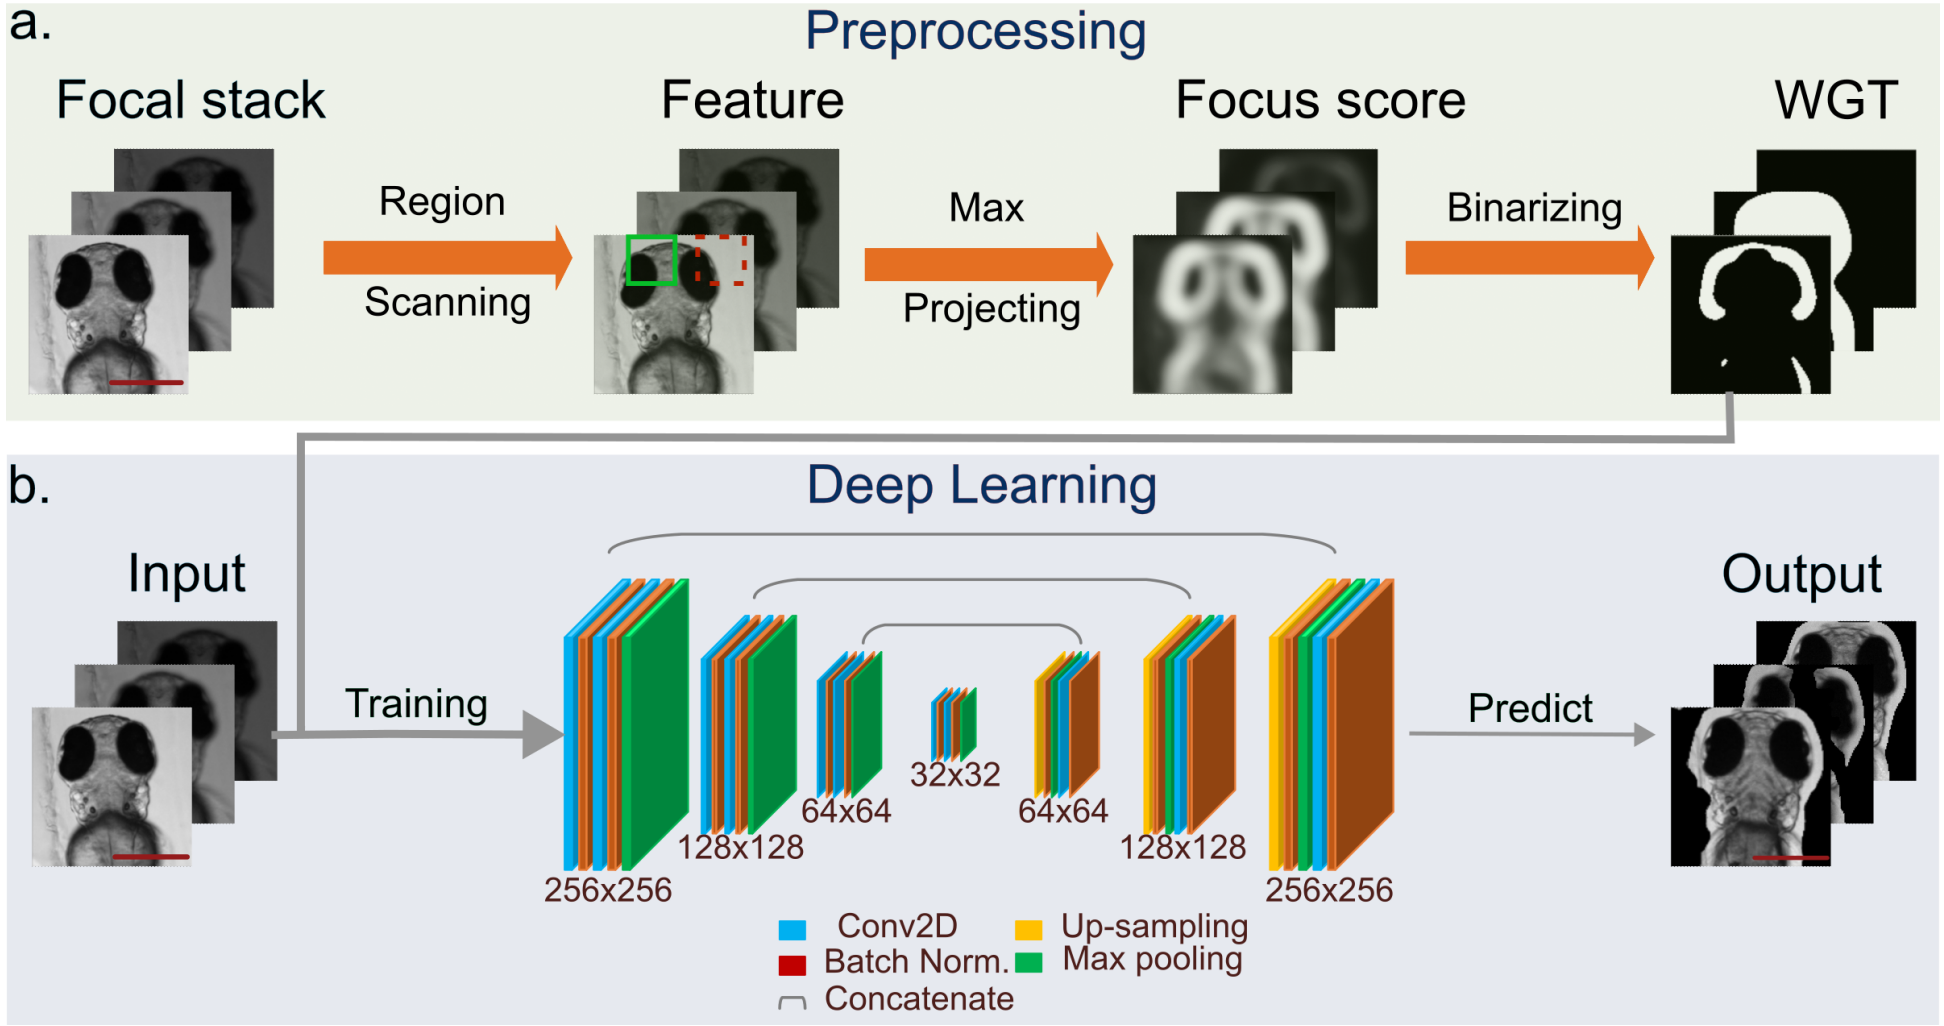

Li & Yakimovich et al., 2024, Sci. Rep.

Rui Li

Weak-labeling pipeline as a CNN

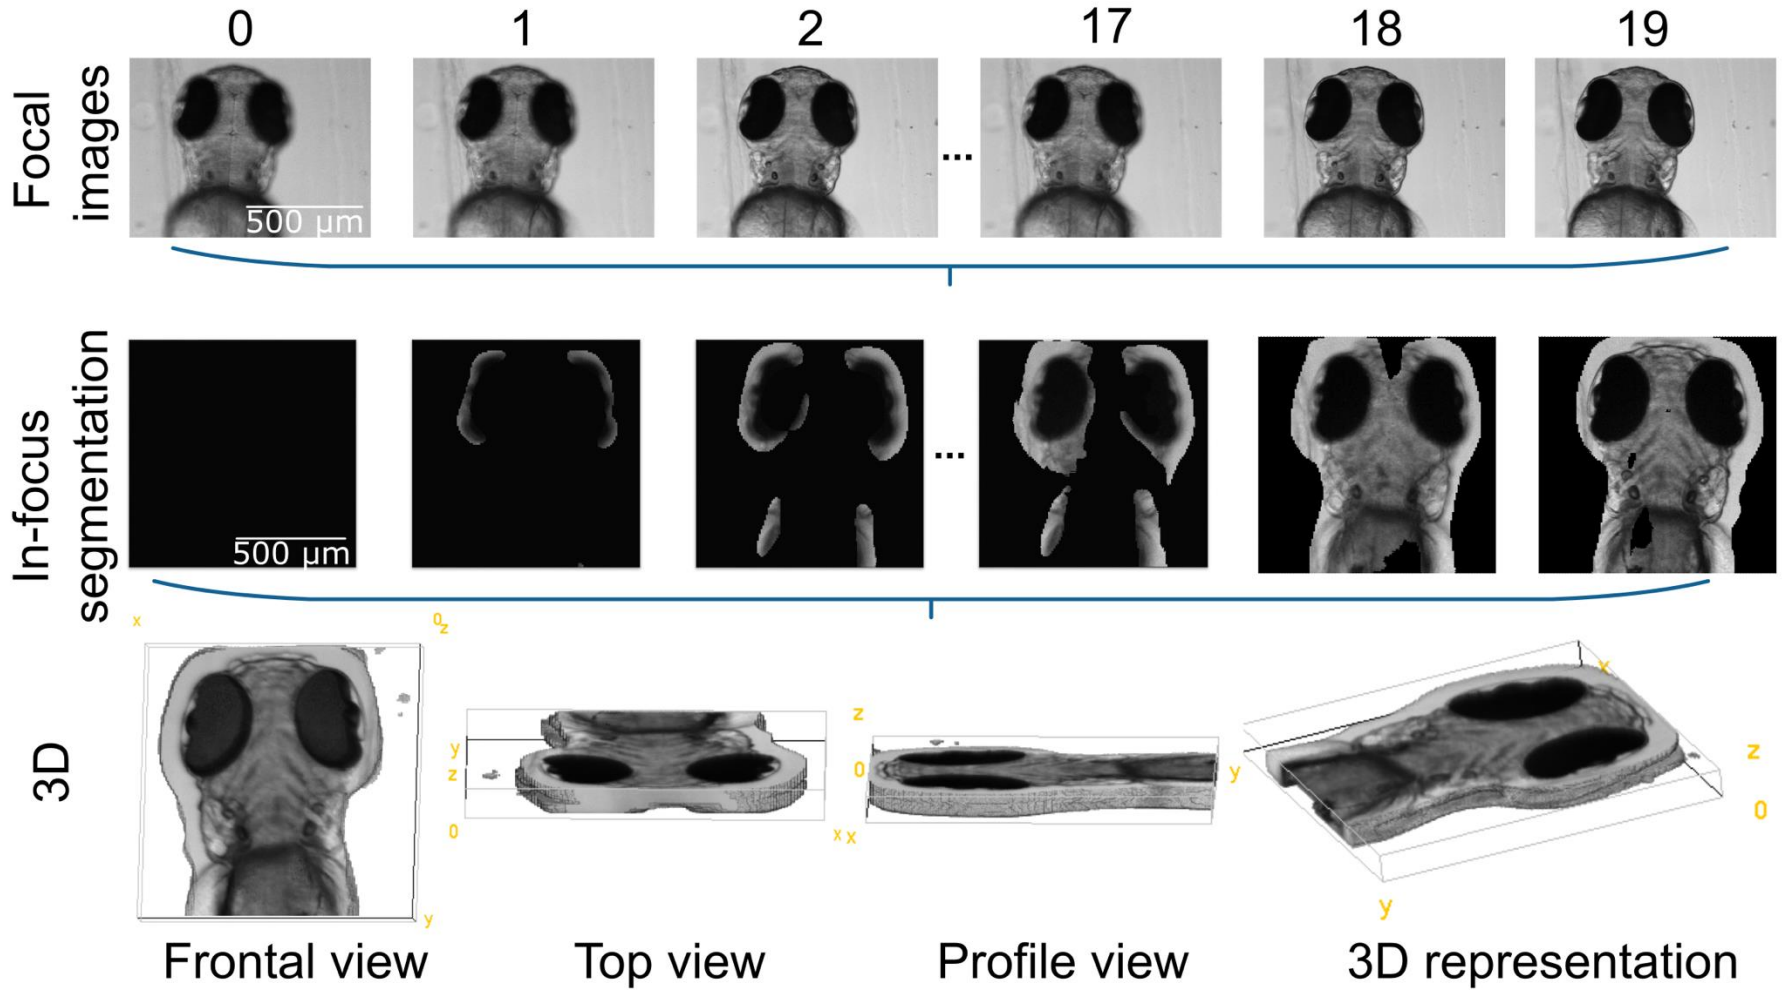

Li & Yakimovich et al., 2024, Sci. Rep.

Rui Li

Take-home Messages III:

- Deep learning is rapidly advancing Life Science research but requires large annotated datasets
- These labels can be used more efficiently using approaches like weak labelling, semi-supervised learning and self-supervised learning
- Conventional algorithms can be used to produce weakly labelled datasets in microscopy

Molecular labelling as Ground Truth

Fluorescence
upon Infection

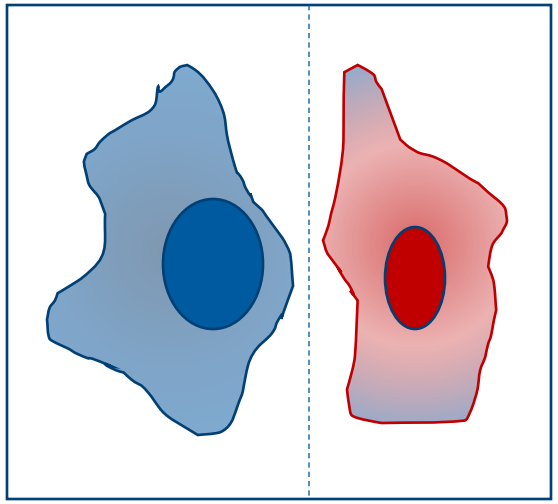

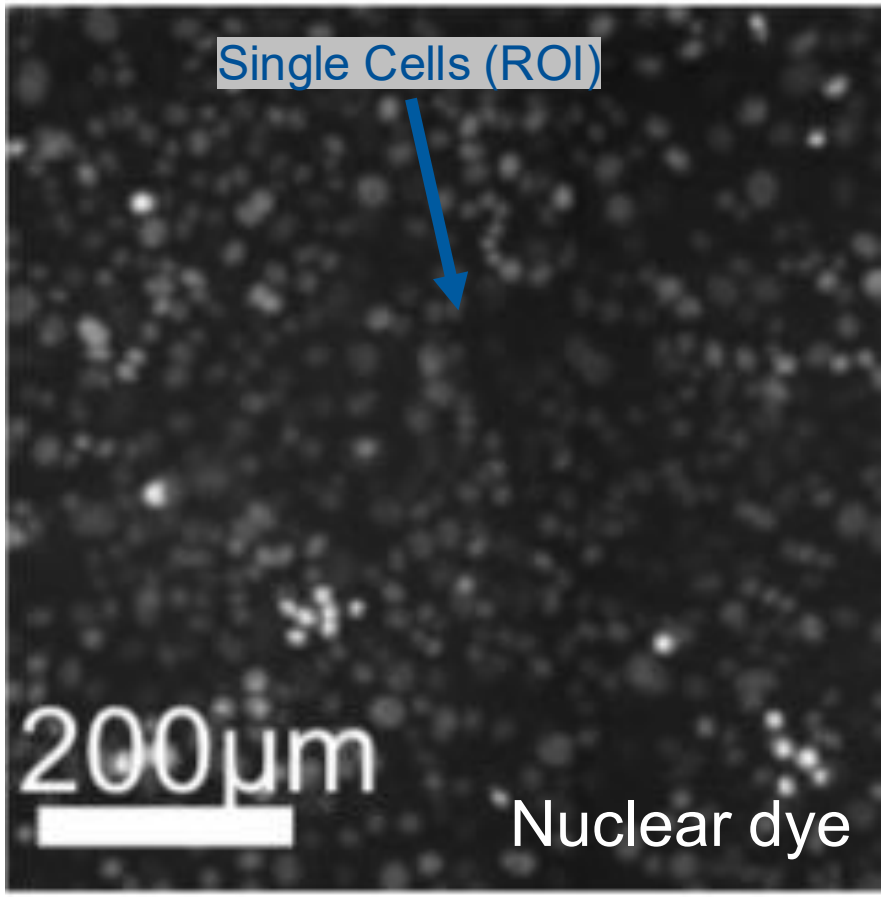

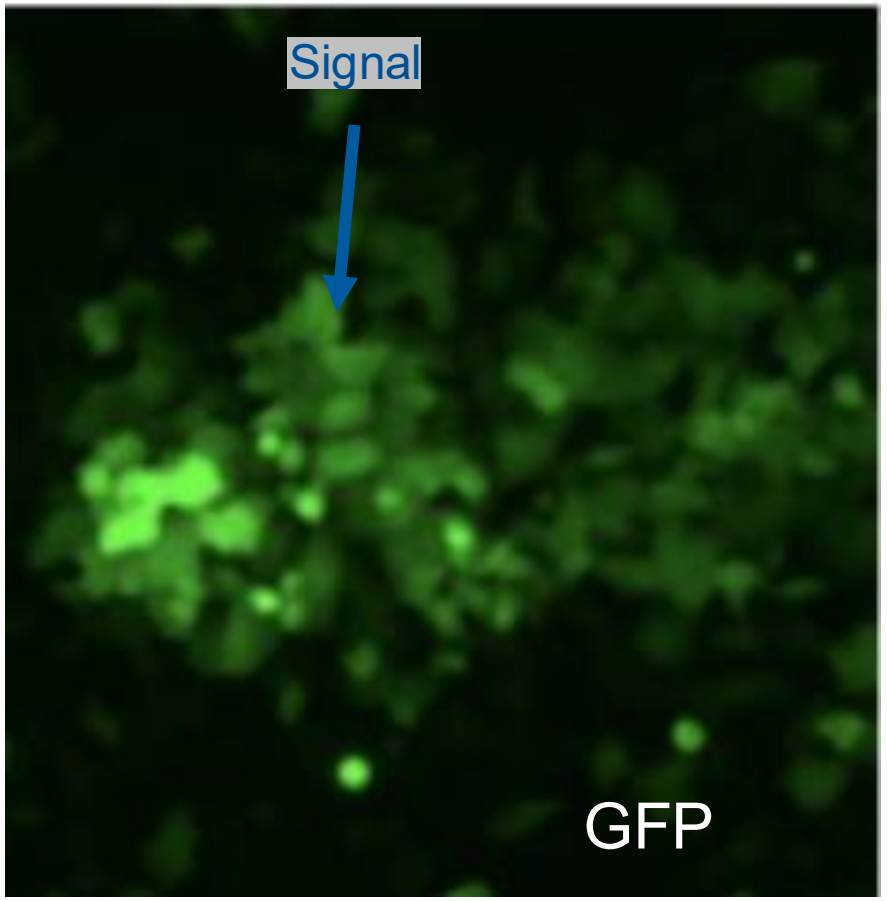

Signal Ground Truth

Deep Learning can be Used to Detect Infected Cells From Nuclear Signal

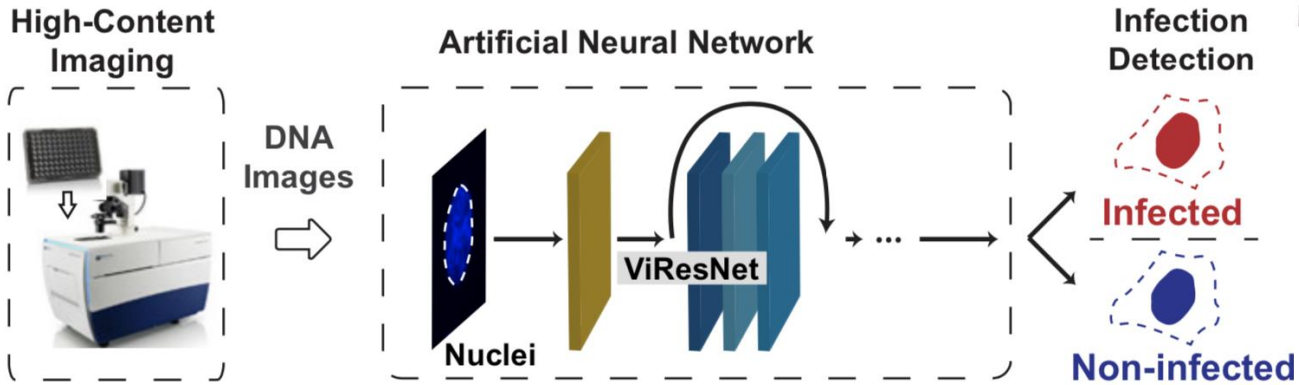

	Accuracy $=\frac{TP+TN}{TP+TN+FP+FN}$	Precision $=\frac{TP}{TP+FP}$	Recall $=\frac{TP}{TP+FN}$
ViresNet: HAdV	0.96	0.94	0.93
MLP: HAdV	0.81	0.75	0.78
ViresNet: HSV	0.95	0.94	0.93
MLP: HSV	0.78	0.74	0.73

- Deep Residual Networks outperform MLP

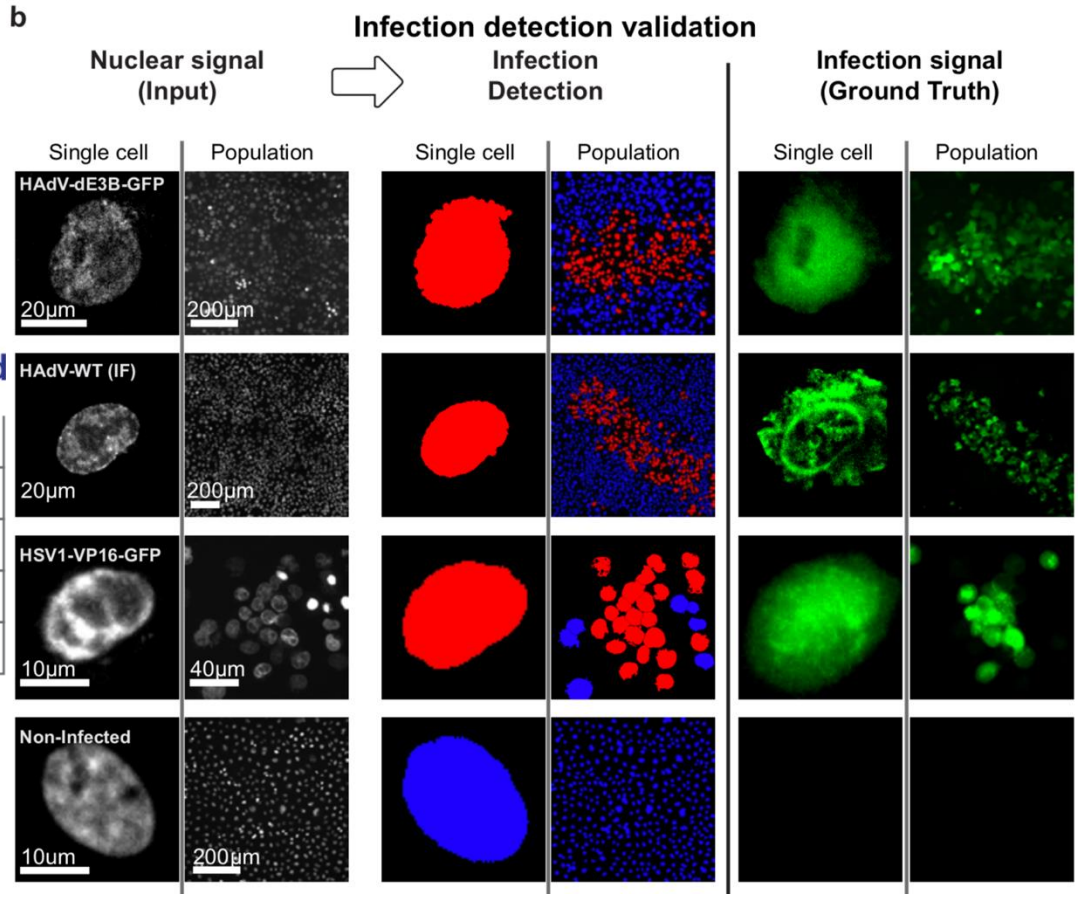

- Generalisation across stainings

Andriasyan, Yakimovich et al. iScience 2021

GFP Signal can be Predicted Directly using Generative AI or Regression DL (Virtual Staining)

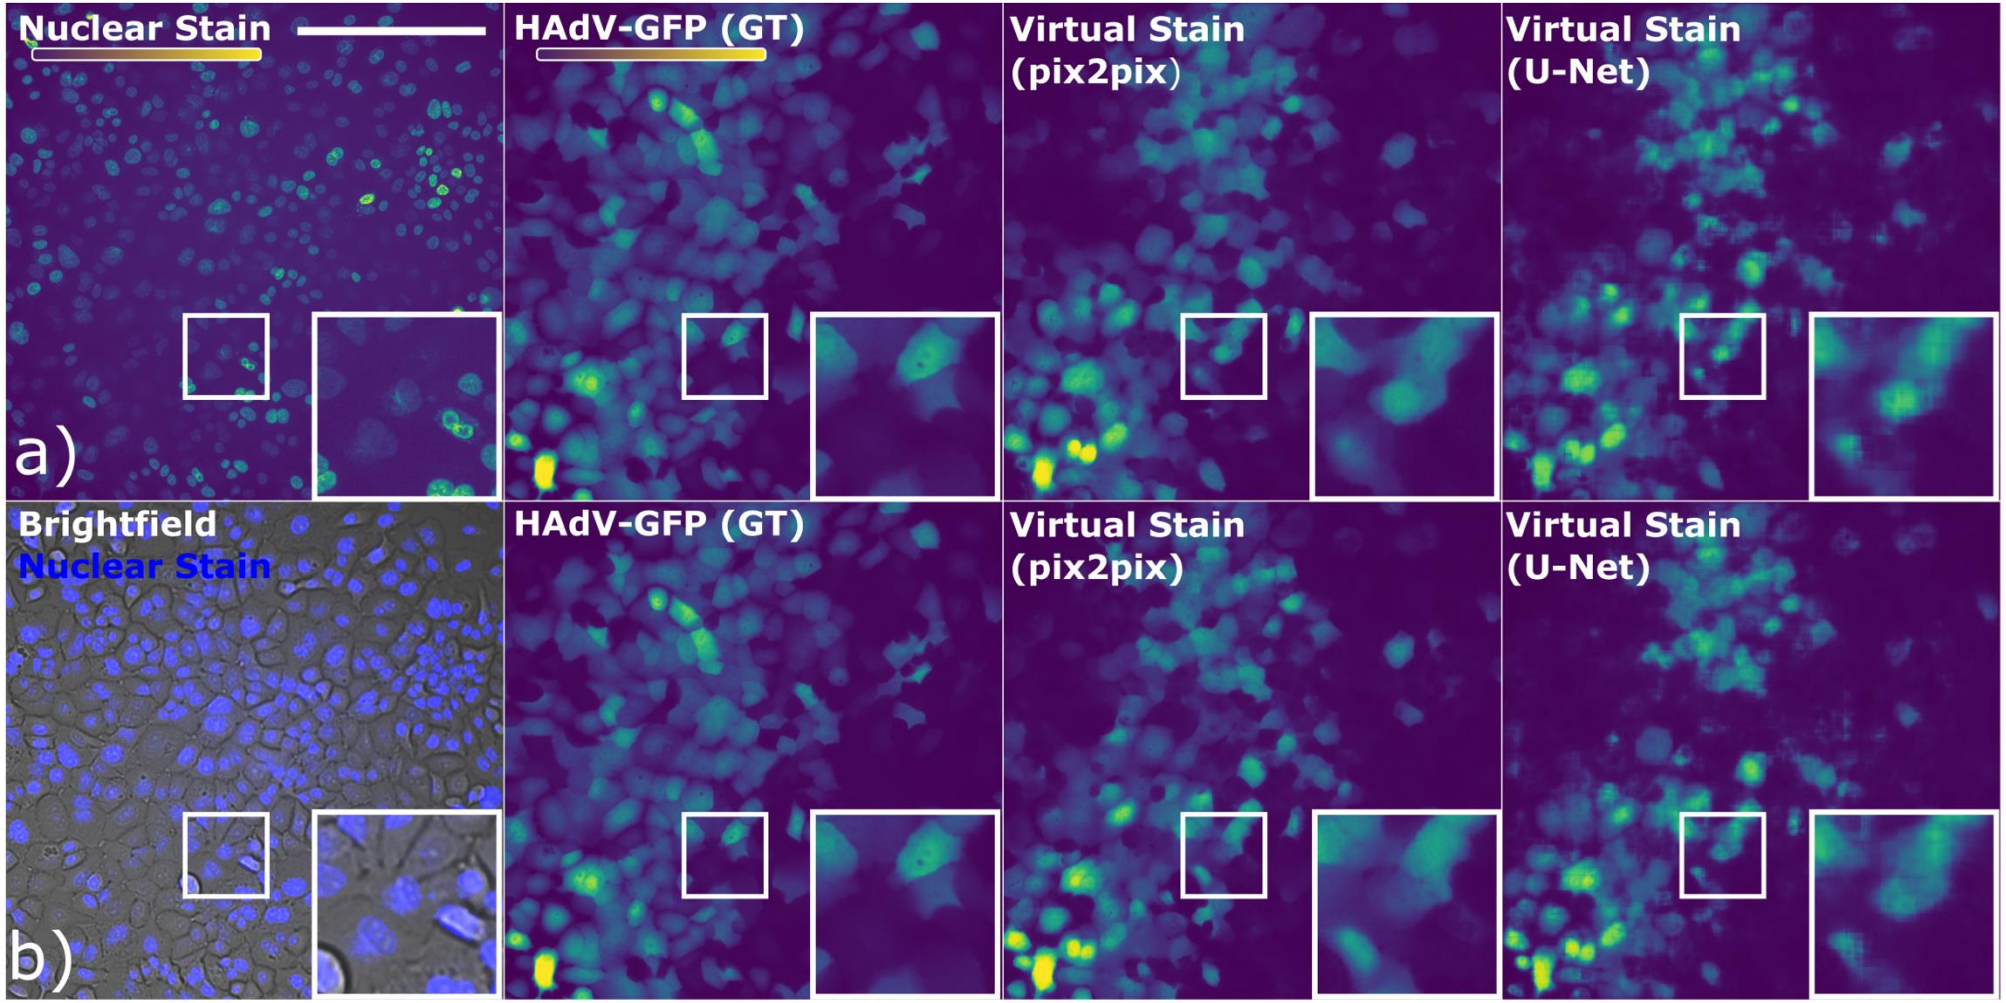

Wyrzykowska & Yakimovich. BioRxiv 2024

HZDR

CASUS
CENTER FOR ADVANCED SYSTEMS UNDERSTANDING

To expand this idea to other viruses and reporters, we proposed the Virus Infection Reporter Virtual Staining Benchmark

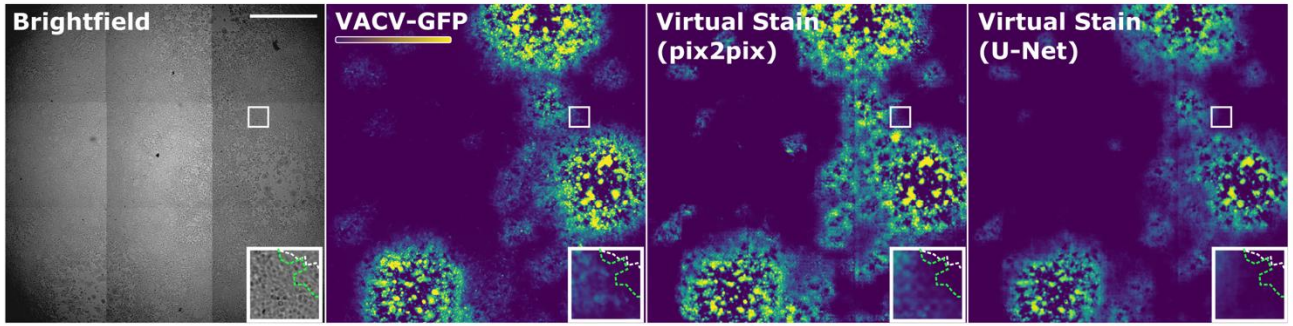

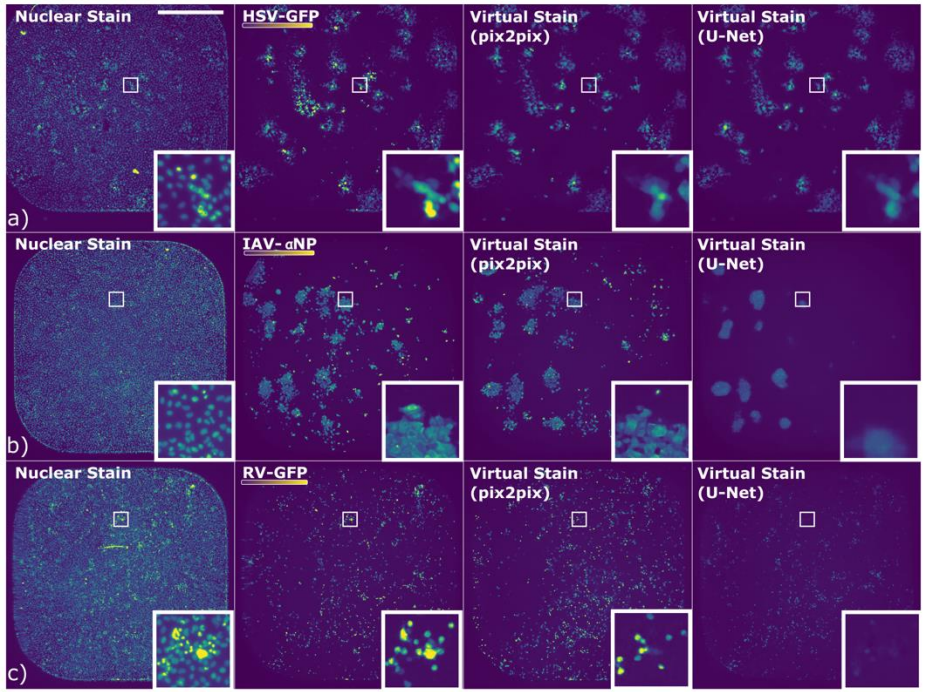

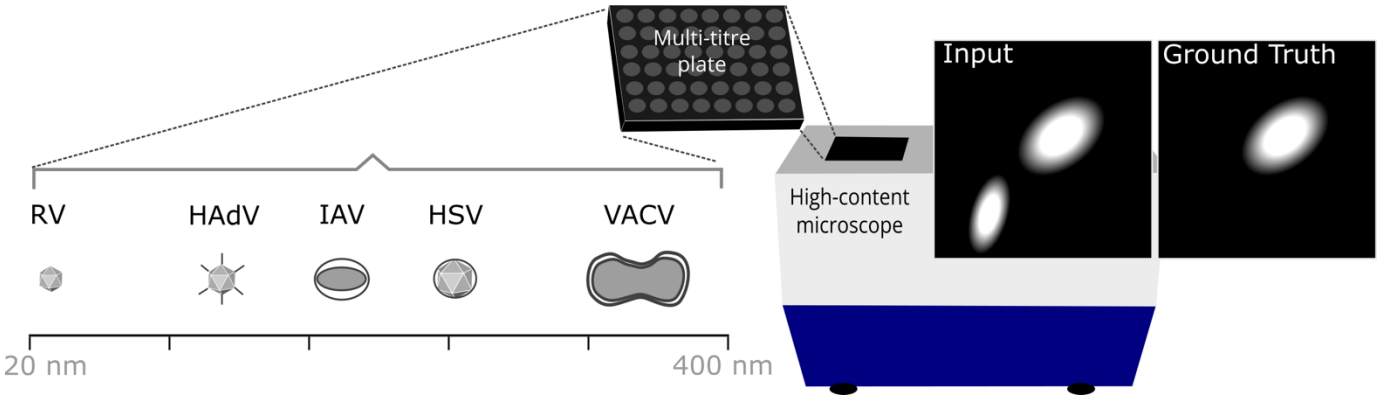

Wyrzykowska & Yakimovich. BioRxiv 2024

Take-home Messages IV:

- Molecular markers can be efficiently used as Ground Truth
- Virus infection can be detected quasi-label-free

Addressing Data-hungry Algorithms with Viruses

Computer Science

Modified **N**ational **I**nstitute of **S**tandards and **T**echnology database

- **60,000** training images
- **10,000** testing images

Cell Biology + Viruses

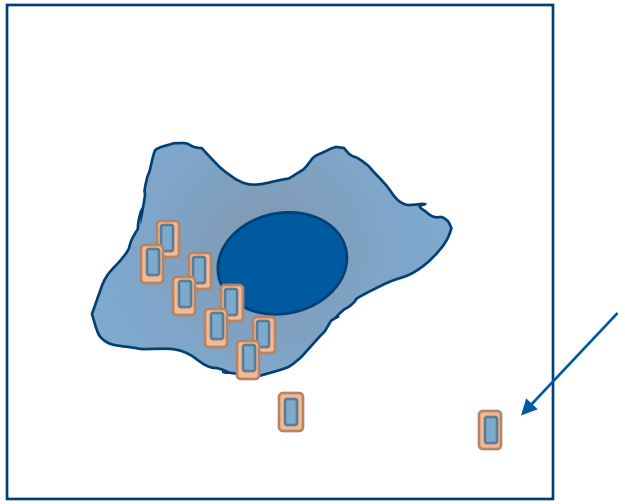

- ⩽ 50 Cells
- ~100-1000 particles / cell

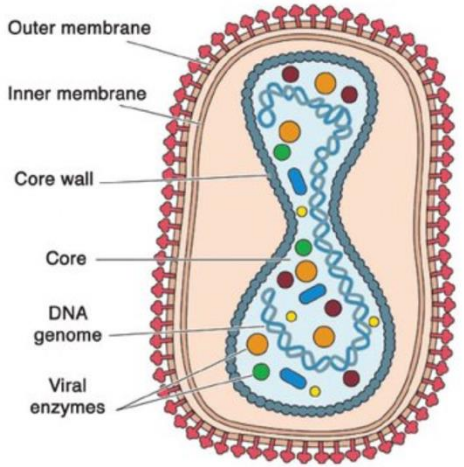

Harrison et al. 2004

Vaccinia virus :

- 220-450x140-260 nm
- Barrel-shaped core

- Focus on Subcellular Structures
- Facilitate Transfer Learning

HZDR CASUS CENTER FOR ADVANCED SYSTEMS UNDERSTANDING

Superior lateral resolution provides an excellent detection entry point

a

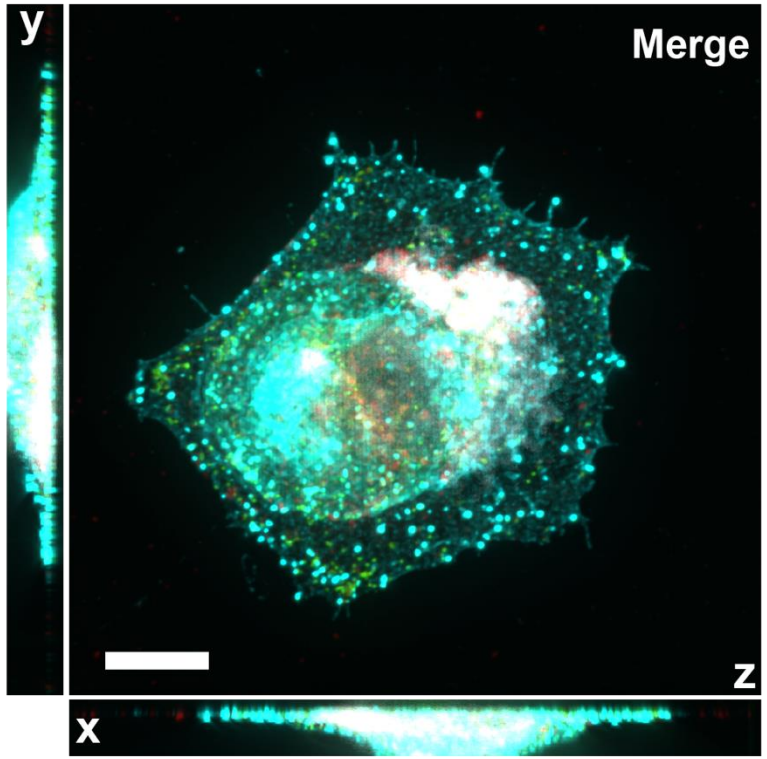

b

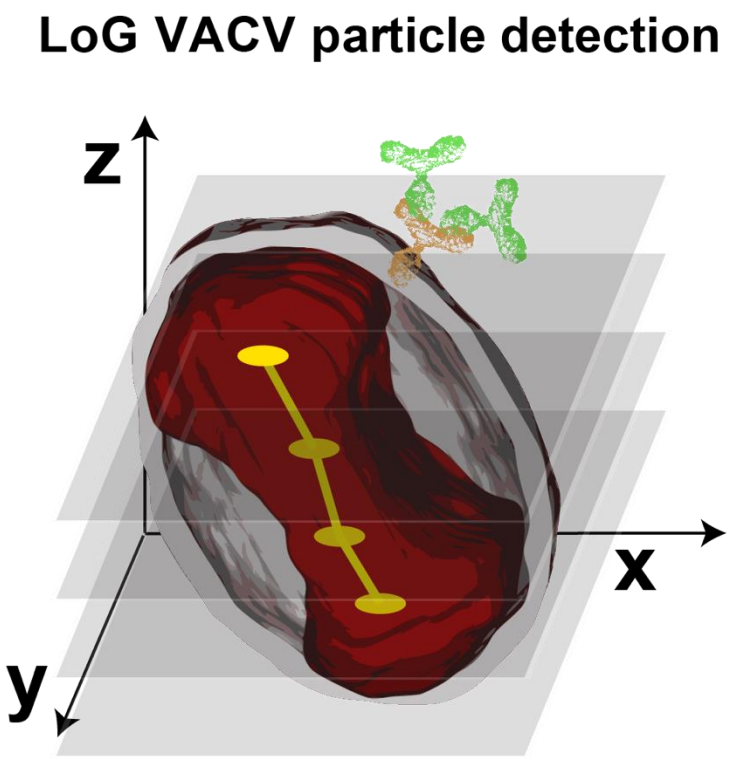

ZedMate plugin

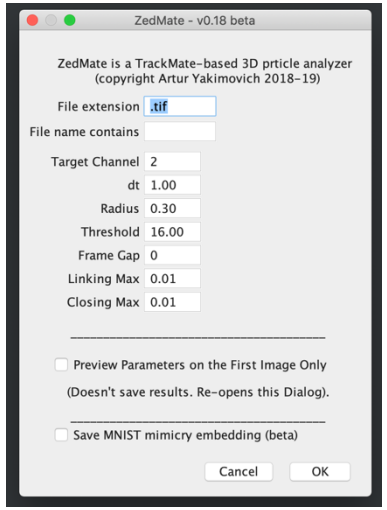

ImageJ - Fiji

Yakimovich et al. mSphere 2020

HZDR CASUS CENTER FOR ADVANCED SYSTEMS UNDERSTANDING

Once measured, Intensities can be analyzed in a linear fashion (binning)

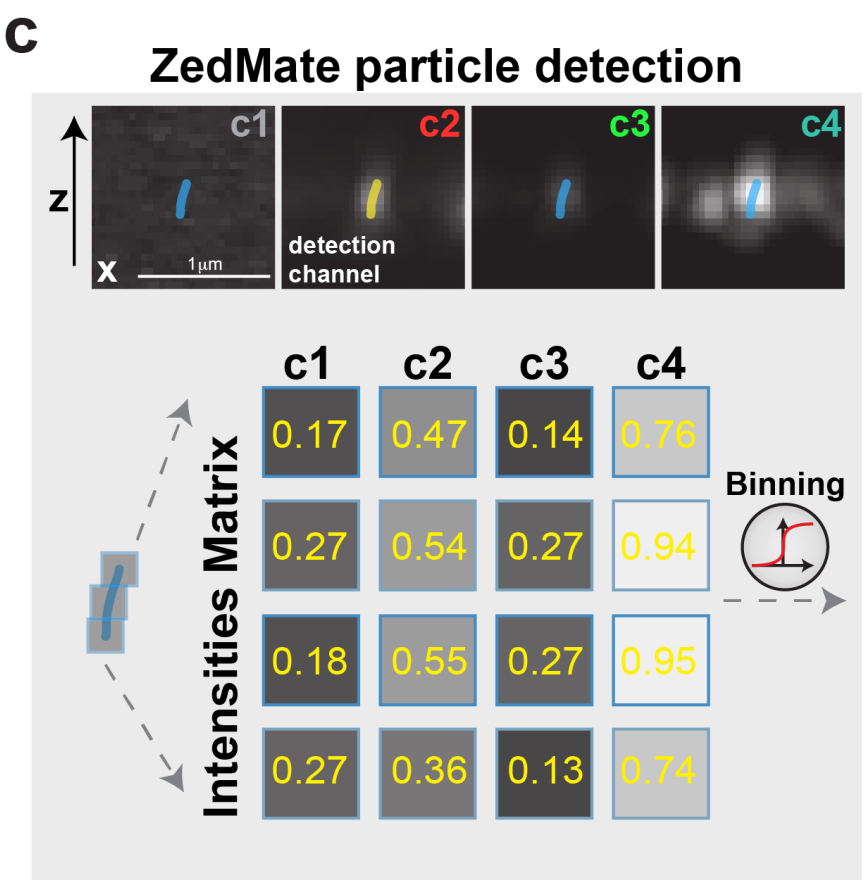

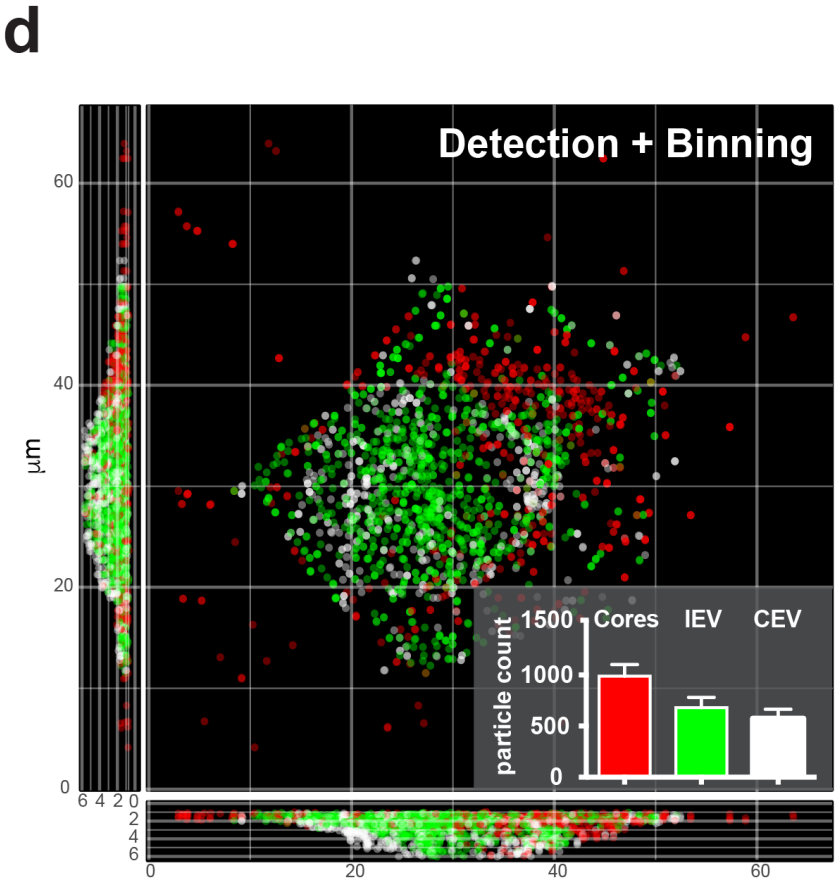

Yakimovich et al. mSphere 2020

MNIST can be repurposed for Transfer Learning

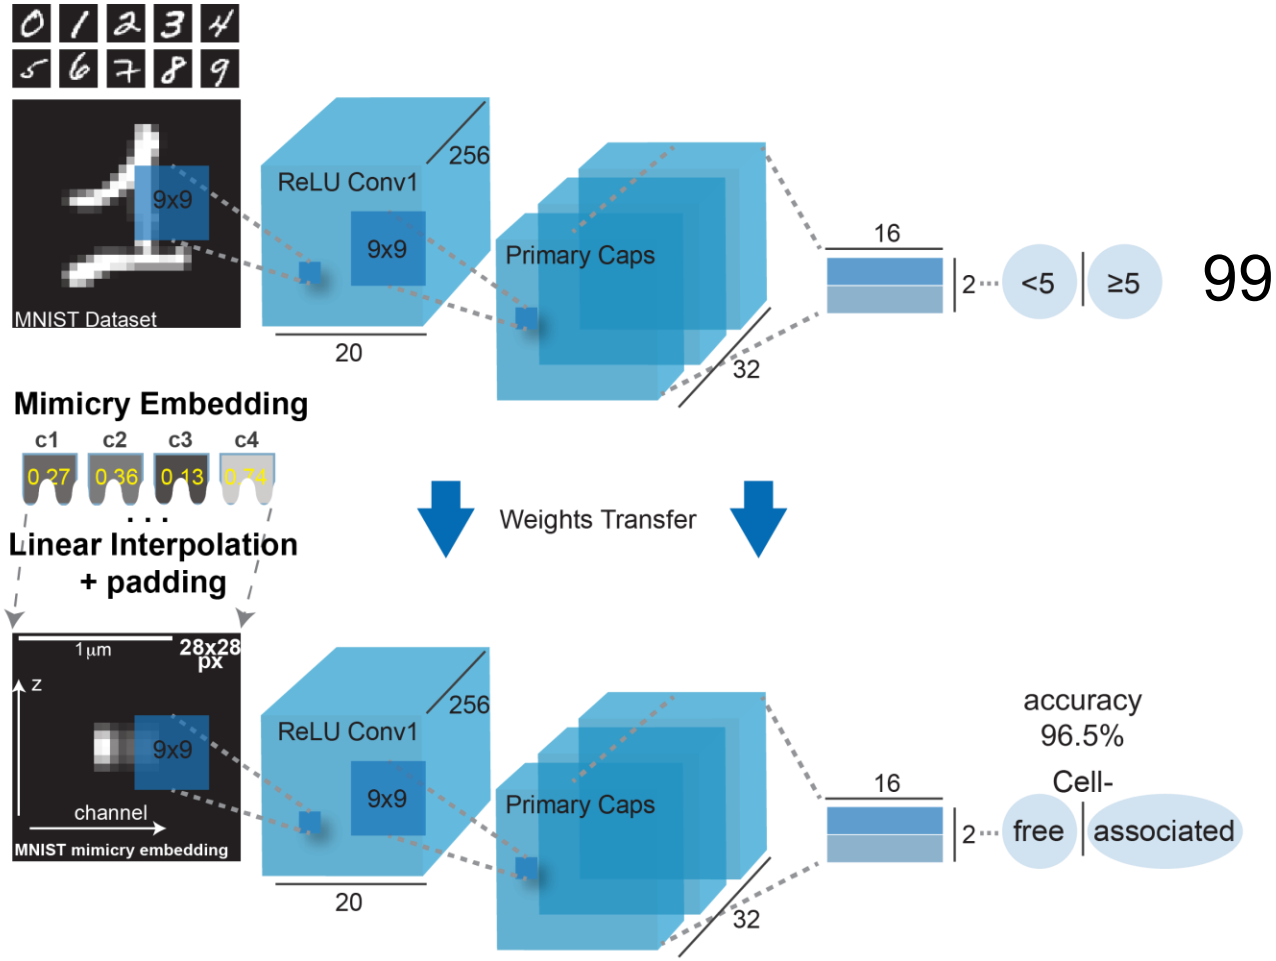

Dataset Annotation

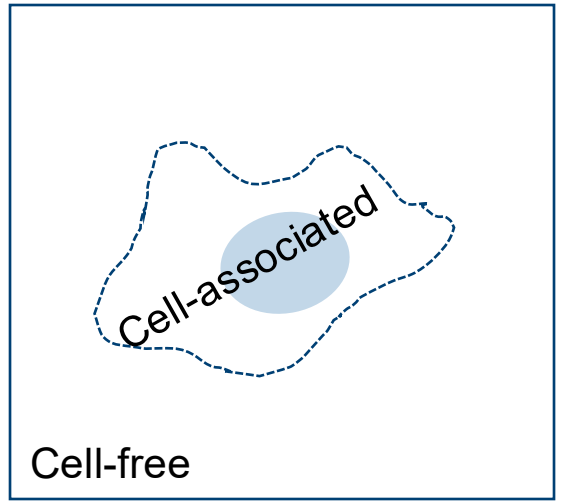

Yakimovich et al. mSphere 2020

HZDR CASUS CENTER FOR ADVANCED SYSTEMS UNDERSTANDING

CapsNet architecture allows visualization of features learned from the dataset

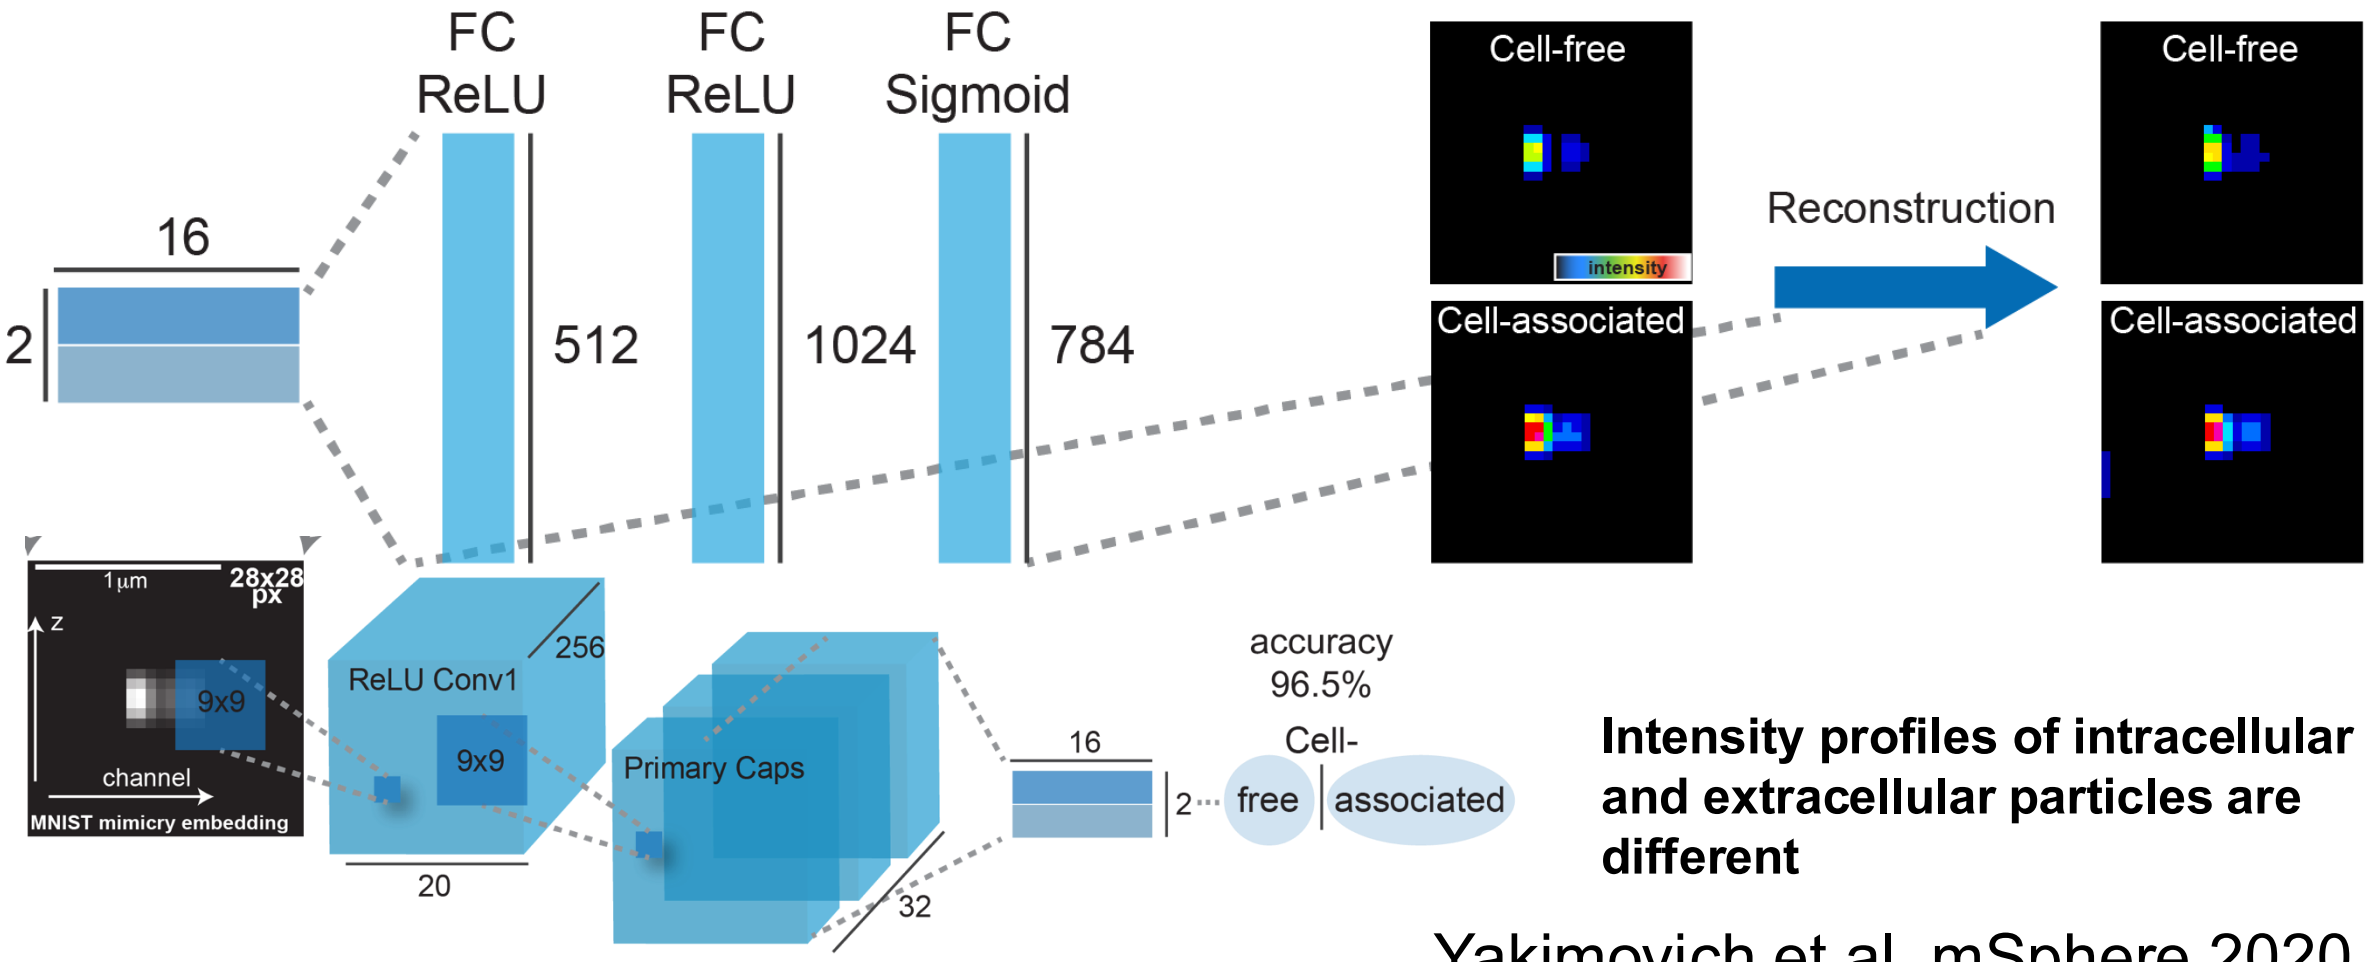

Intensity profiles of intracellular and extracellular particles are different

Yakimovich et al. mSphere 2020

HZDR

CASUS
CENTER FOR ADVANCED SYSTEMS UNDERSTANDING

Take-home Messages V:

- Deep learning is rapidly advancing Life Science research
- Changes in morphology picked up and summarised by DNN
- Discriminators combined with generators can assist discovery within the wild datasets

Acknowledgements: MLID group

MLID Group:

- Jacob Anter
- Trina De
- *Gabriel della Maggiora*
- Ashkan Mokarian
- Adrian Urbański
- *Rui Li*
- Martina Radoynova

PhD Alumni:

- Rui Li

MSc Alumni

- Maria Wyrzykowska
- Adrian Urbanski
- Nikita Deshpande
- Vaibhav Sharma
- Subasini Thangamani

Competitive Funding:

HELMHOLTZAI

CASUS **Open Project**

HELMHOLTZ IMAGING

Helmholtz Foundation Model Initiative

Competitive HPC Funding:

ScaDS.AI DRESDEN LEIPZIG

HZDR CASUS CENTER FOR ADVANCED SYSTEMS UNDERSTANDING

nature portfolio

Special Collection Call for Papers

Annotated Optical Microscopy Datasets

Scientific Data

Balazs Harangi, Artur Yakimovich

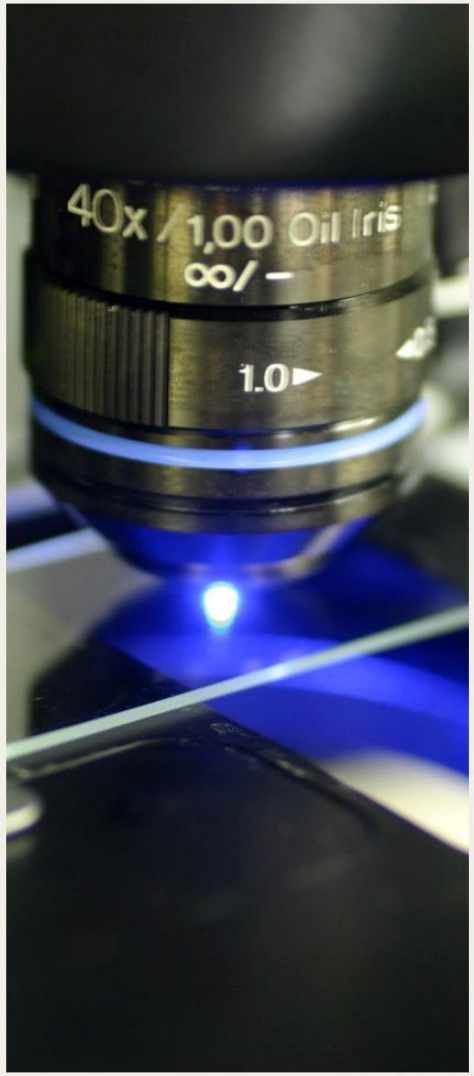

https://go.nature.com/4hyDnNK

HZDR

CASUS
CENTER FOR ADVANCED SYSTEMS UNDERSTANDING