Photonics-Workshop:
«Artificial Intelligence in Photonics»

Fachanlass Swissmem-Fachgruppe Photonics in Zusammenarbeit mit Swissphotonics
Dienstag, 3. September 2019, 13.30 bis ca. 17.30 Uhr

Fachhochschule Nordwestschweiz FHNW, Campus Brugg-Windisch, Raum 6.-1D09 im Gebäude 6, 5210 Windisch

Photonics for AI or AI for Photonics

Andrea Dunbar
FHNW, Campus Brugg-Windisch
3 September 2019

Technologies **in focus** at CSEM

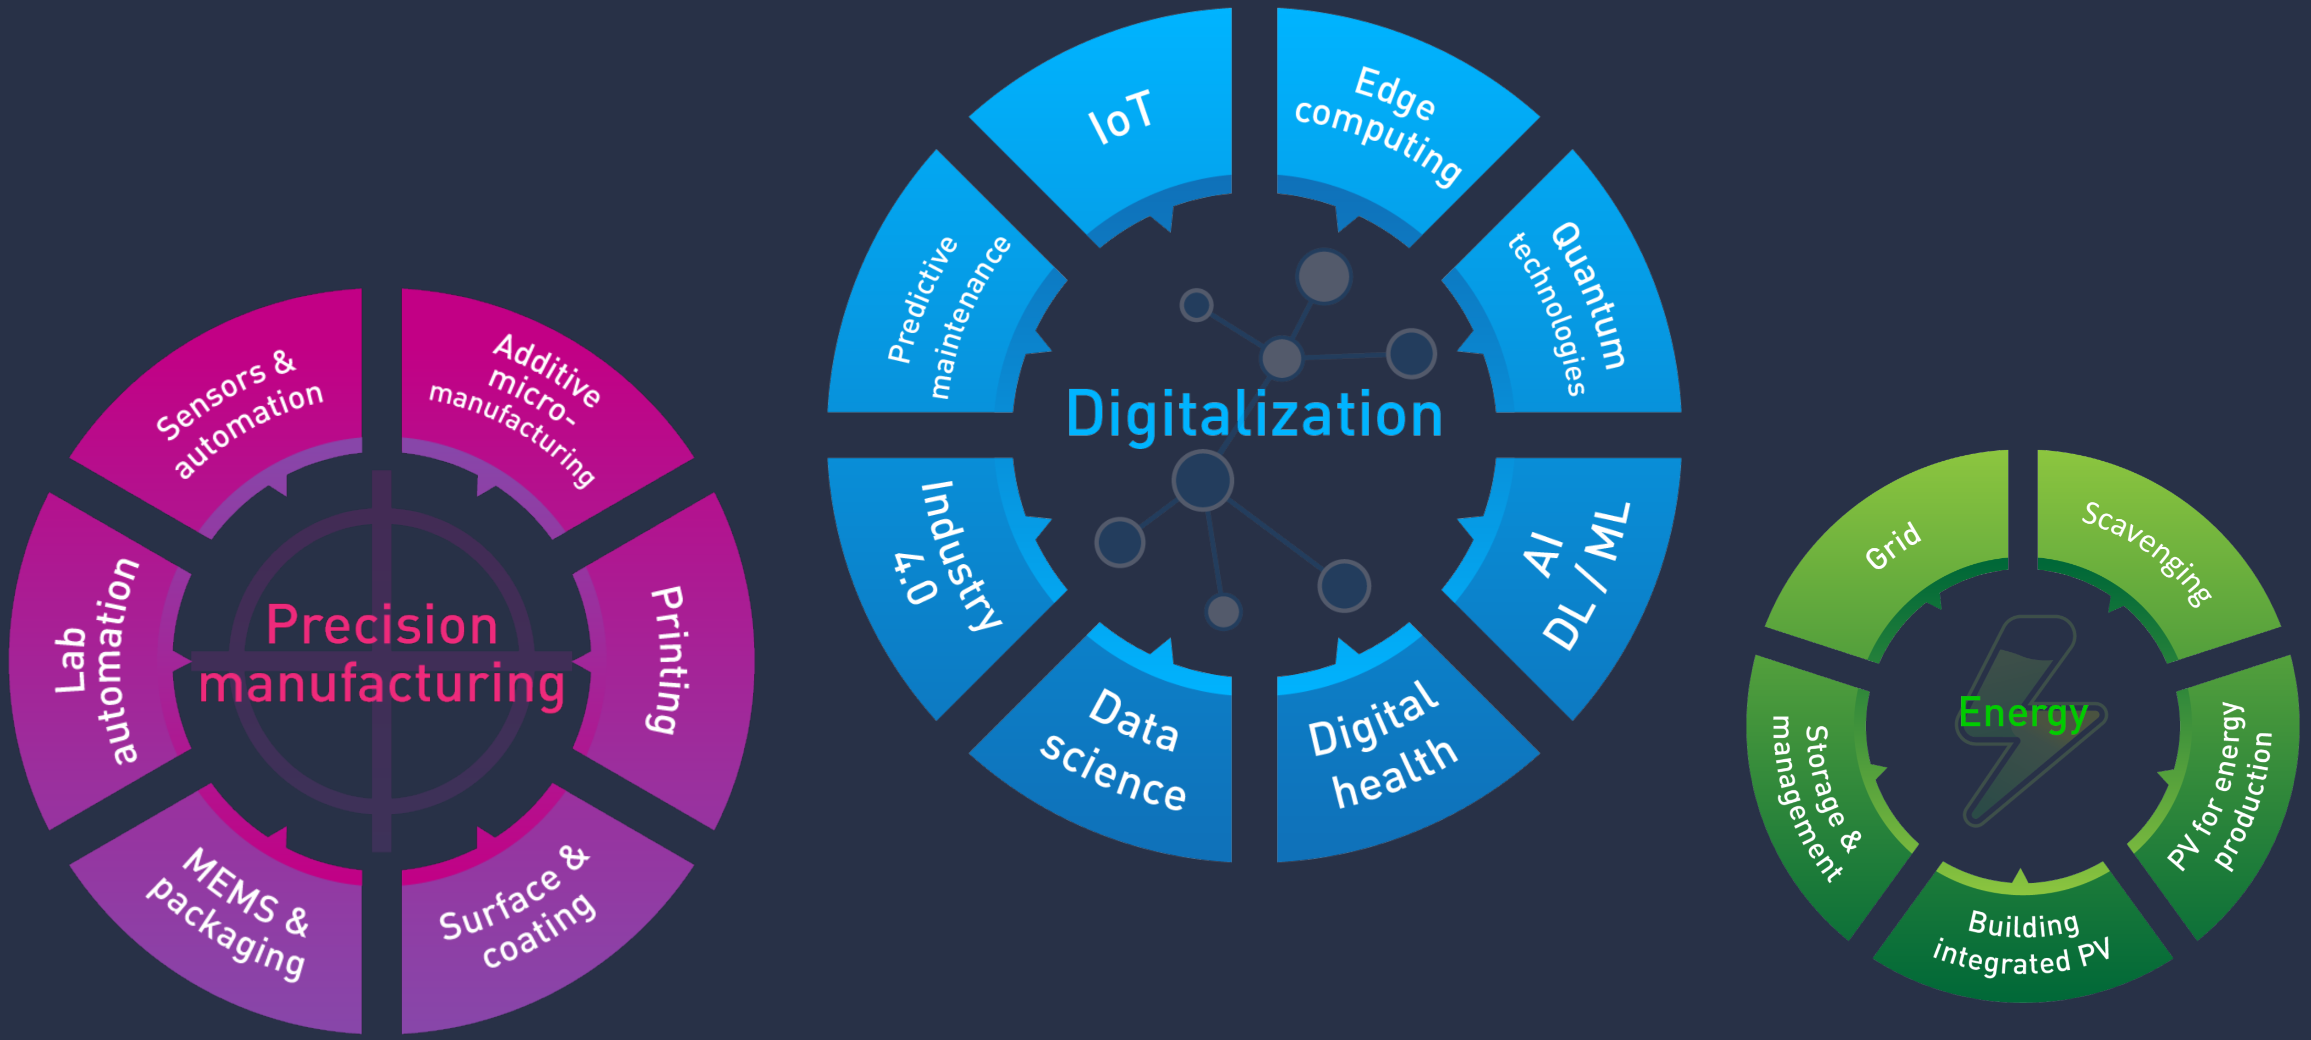

csem

CSEM at a glance

83.0
Turnover
(mio CHF)

43
Nationalities

437
Persons

43
New
ventures

175
Industrial
clients

203
Patent
families

64
European
projects

A public-private **partnership**

- Flexible and professional
- Industry friendly IP approach

14 %
Swiss confederation
(EPFL)

73 %
Private Organizations

13 %
Neuchâtel – City & Canton

From photon in to information out

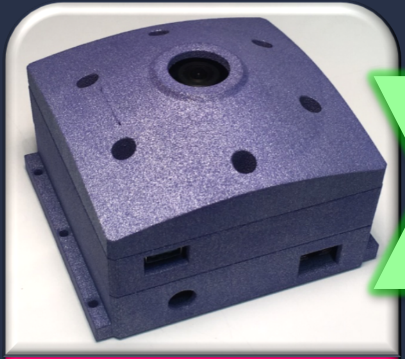

Modal

0.5 dm^3

Tegra **TX2**

Max **7 W**
Wired / Battery

Cost: ~$1000 / 10 U

Pupil

8 cm^3

ARM **H7**

Max **0.3W**
Wired / Battery

Cost: ~$200 / 100 U

Witness

2 cm^3

Risc-V + Cust ML

Max **2mW**
Battery / Solar

Cost: ~$10 / 10 kU

csem showcase

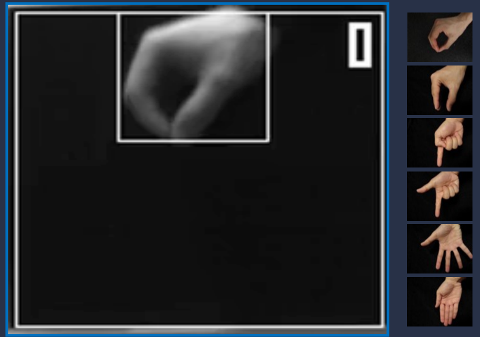

REDO slide with Hierachicial computing

50 billion IoT : Now is the time for unlimited autonomy

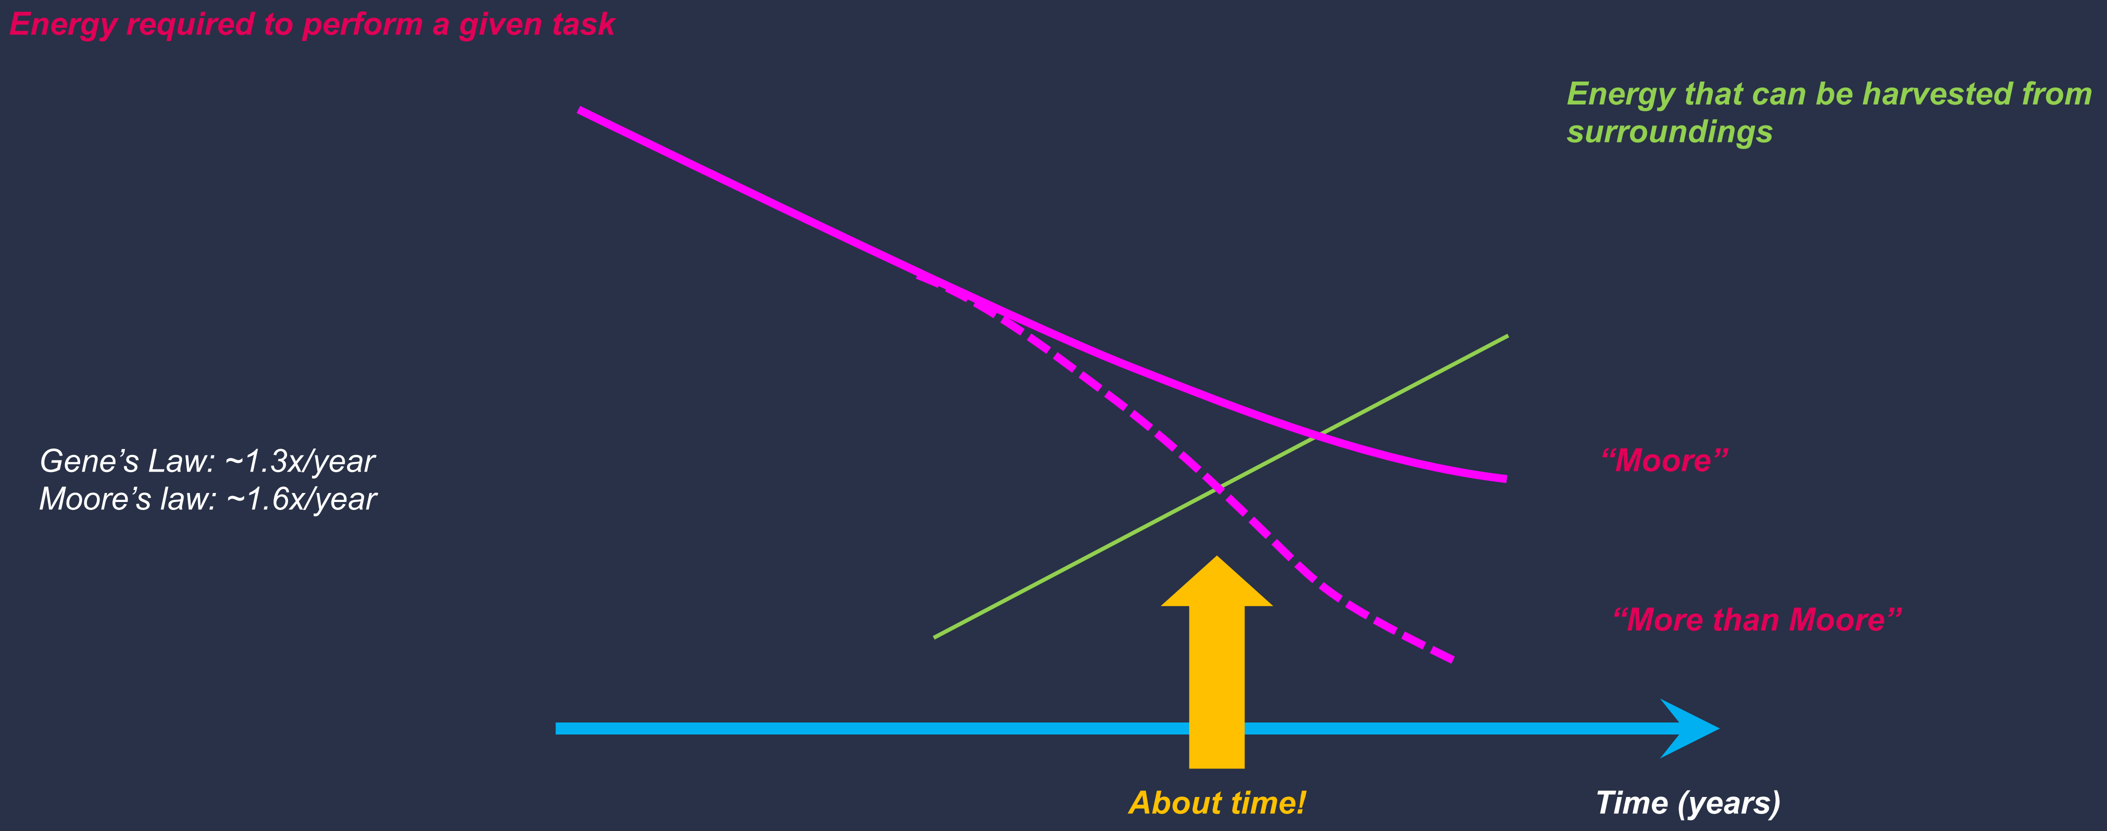

Gene's Law: ~1.3x/year
Moore's law: ~1.6x/year

csem

VISION ML NEEDS PHOTONICS

Advanced miniaturized IoT systems

Witness: Image recorder powered by ambient li

Scene capture:

- QVGA (320x320), 107° FOV
- When light available @ max 1
- Based on motion detec

Mechanical

- 8cm
- 0.6n ple sticker
- 4mm camera coin

1mW Power Budget

Expected

Achieved

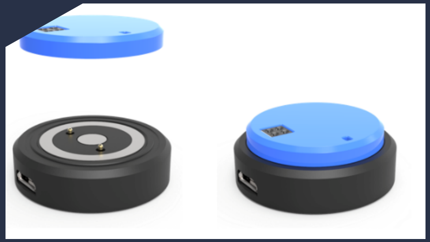

USB Dock
to download
captured images

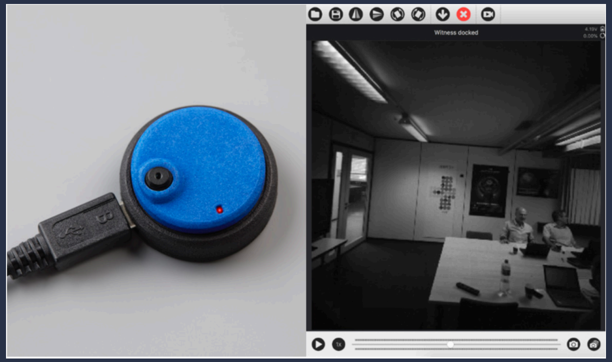

Optics photonics to help Machine Learning

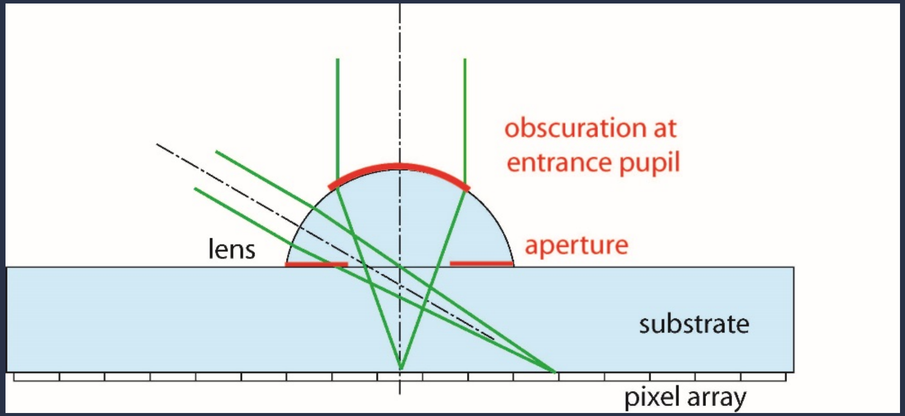

Equalize illumination for all AOI, avoid double images and choose lowest aberration rays by applying an obscuration and an aperture.

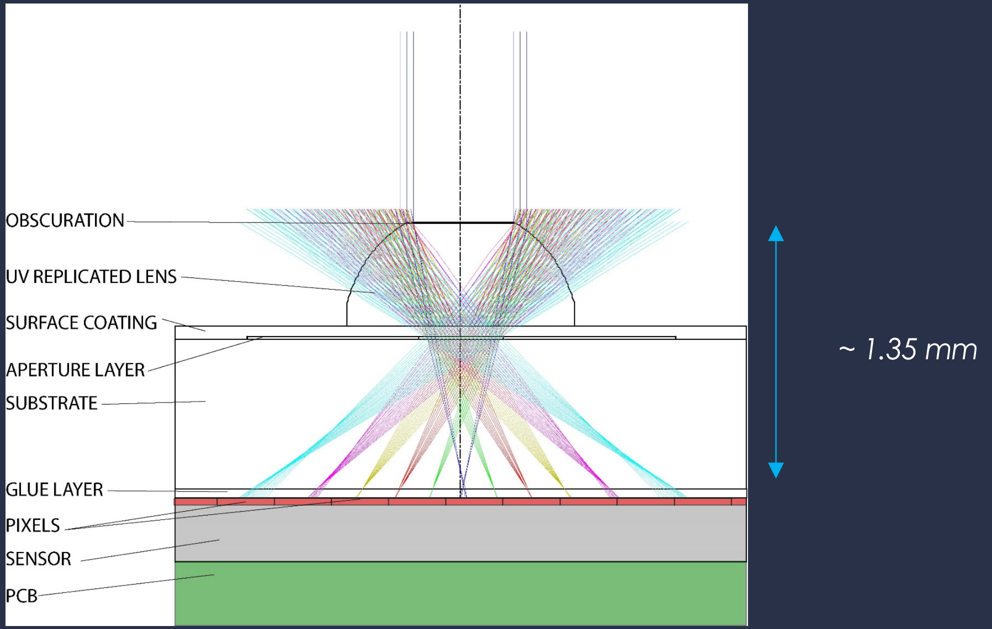

For a first prototype fabricated by diamond turning a total track length of TOTR = 1.35 mm is demonstrated

csem

CTI Feasability ASICs → Ergo: ULP HDR IoT Imager (2020Q1)

- Resolution: 320 x 320 (640 x 640) pixels
- Dynamic range: at least 120 dB intra-scene
- Power budget: < 700 µW (1 mW) at 10 fps
- SNR: 20 dB at 5 lux and 25 °C
- Interface: 1 to 8 bits SPI (HyperBus)
- Persistence: Image stored on-chip
- Cost in production: < USD 2

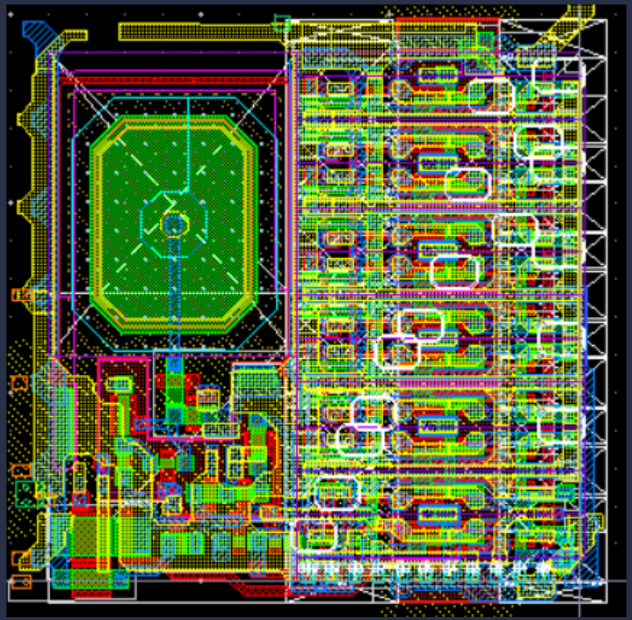

Witness: Software

The **ERGO** imager stores the picture internally: there is no need for external memory

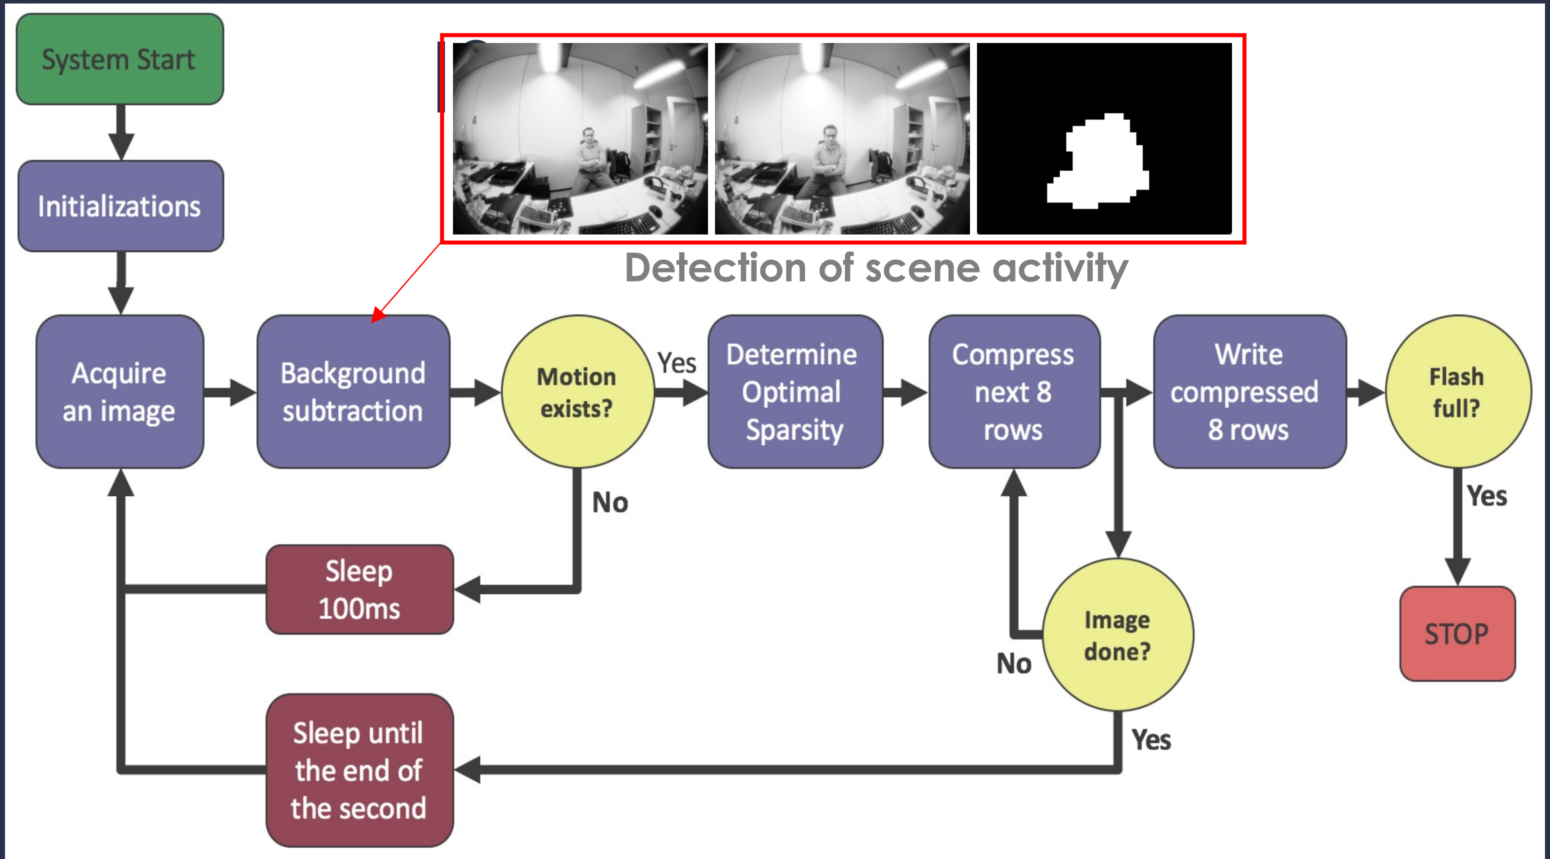

Motion detection is performed
on subsampled images

Full resolution image is processed by chunks of 8 lines (MCU internal memory is 64kB only)

Only changed parts are compressed and stored in the flash memory

ML for OPTICS

Image Recognition

Image restoration

Observations are degraded compared to sharp (clean) images

Natural image

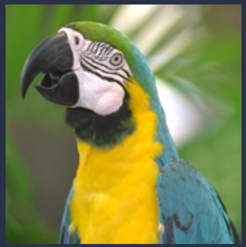

Low-res

Bayer pattern (mosaic)

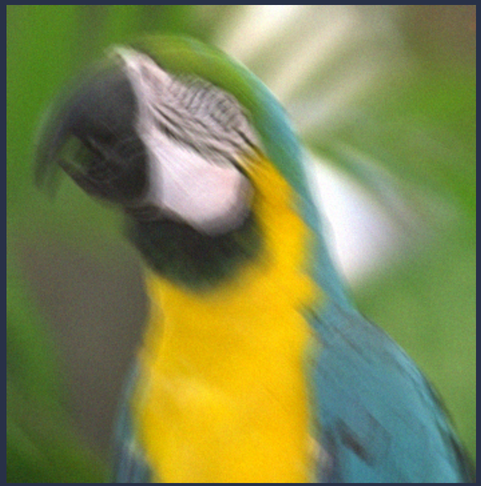

Blurry

Noisy

Image restoration: Supervised, End-to-End

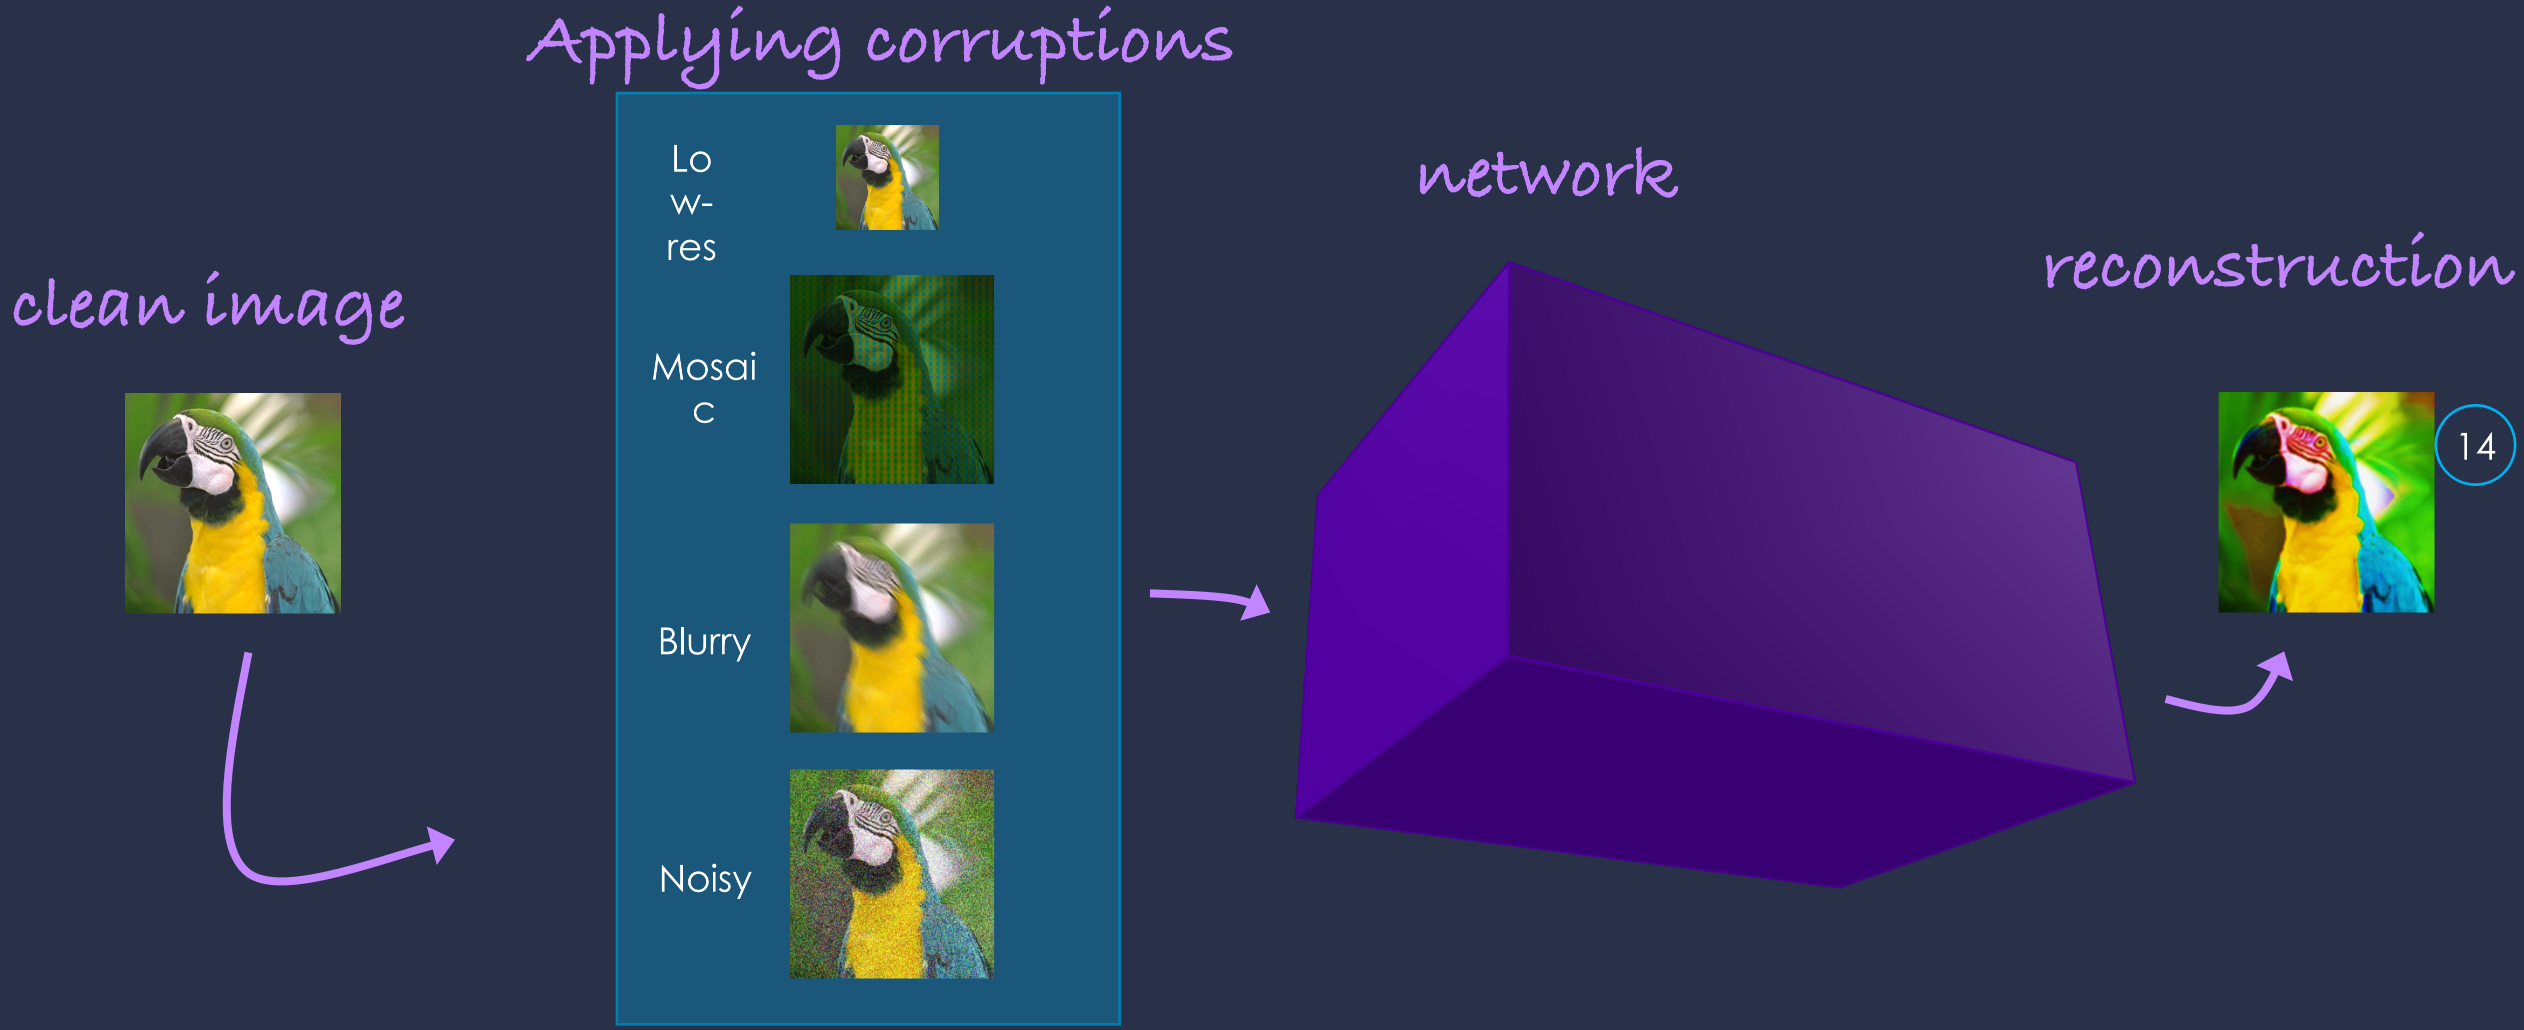

Image restoration: Supervised, End-to-End

Gaussian noise with $\sigma = 50$

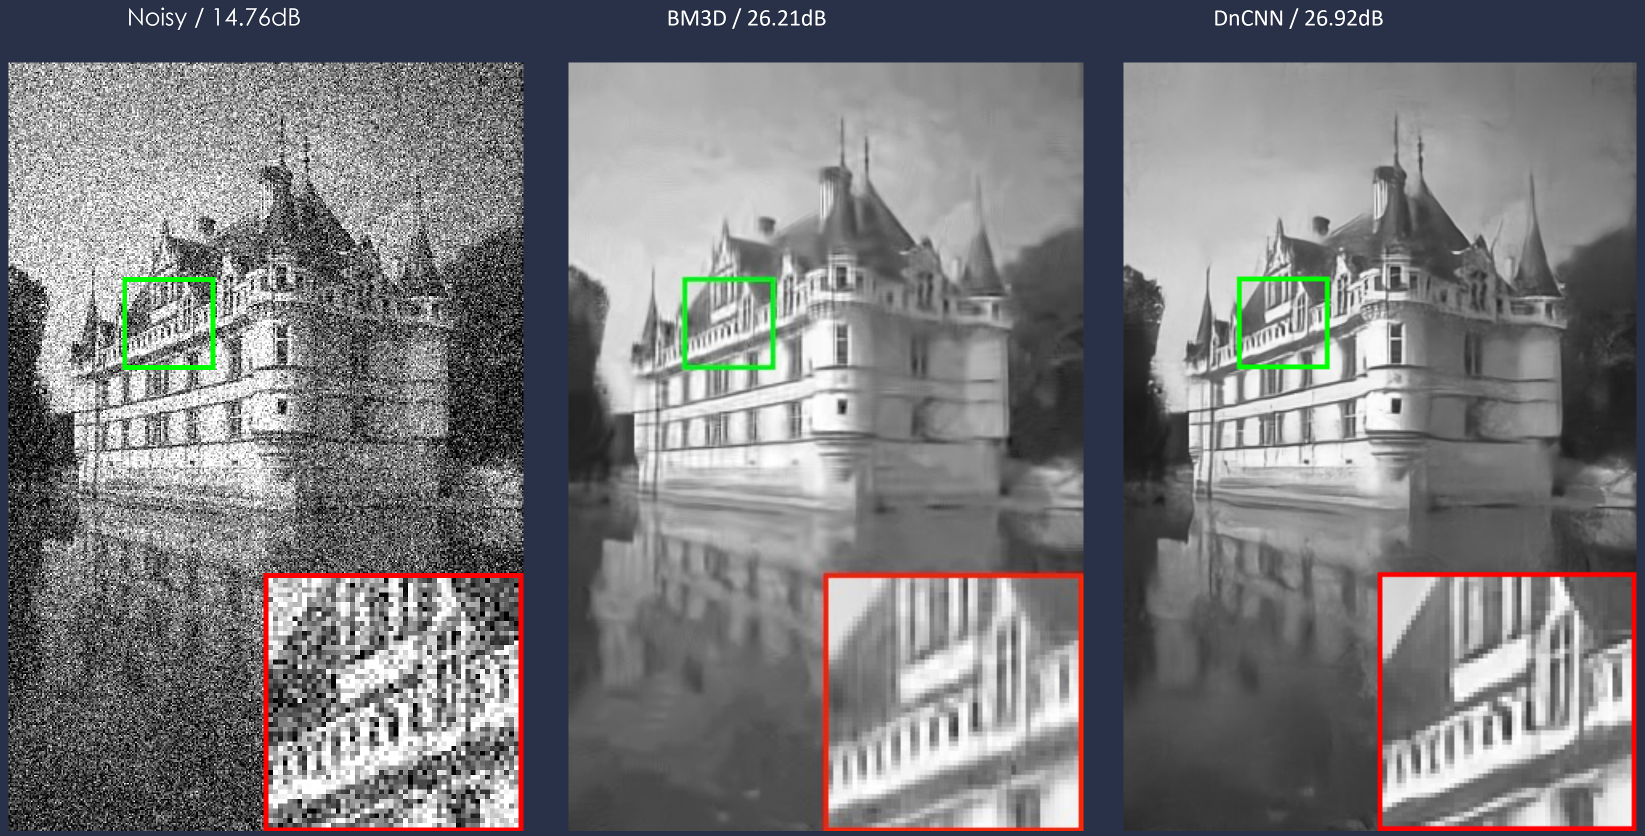

Real camera noise removal

- **Hand-crafted models such as BM3D still outperform CNNs!**

[Plötz & Roth, Benchmarking Denoising Algorithms with Real Photographs, 2017]

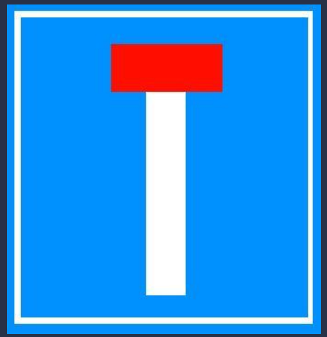

Real camera noise removal

- Hand-crafted models such as BM3D still outperform CNNs!

  [Plötz & Roth, Benchmarking Denoising Algorithms with Real Photographs, 2017]

- **BM3D: Assumption about image sparsity**

- **CNN: Assumption about noise distribution**

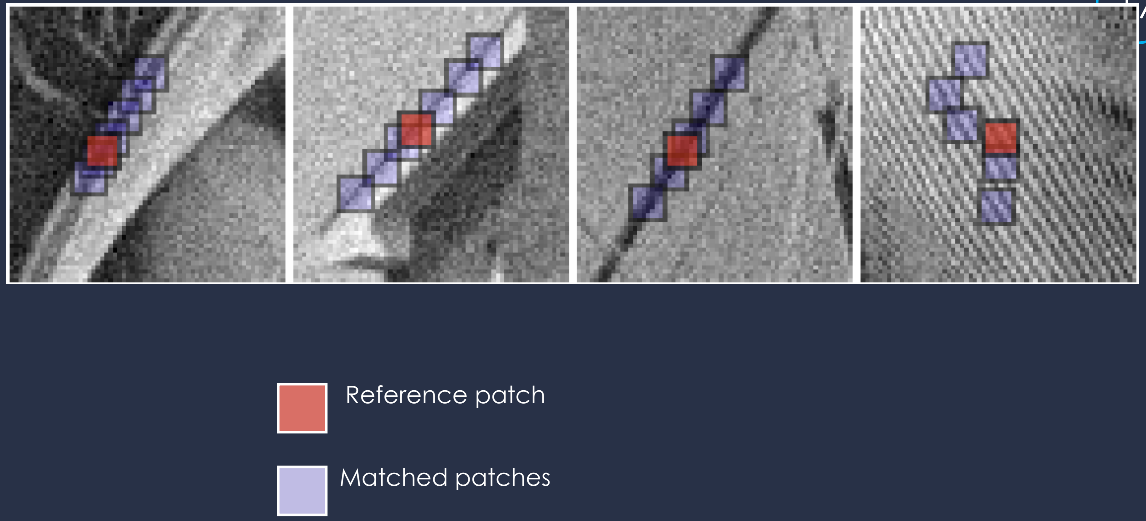

:: csem

Image restoration: Maximum a-Posteriori

- Given y, maximize posterior probability

$$\underset{x}{\operatorname{argmax}}\, p(x|y) = \underset{x}{\operatorname{argmax}}\, p(y|x)p(x) =$$

$$\underset{x}{\operatorname{argmax}} \log p(y|x) + \log p(x)$$

- $\log p(y|x)$: data term (likelihood)
  - How well does the solution explain the observed data
- $\log p(x)$: image prior
  - Probability of the solution

Image restoration: Denoising Autoencoder (DAE)

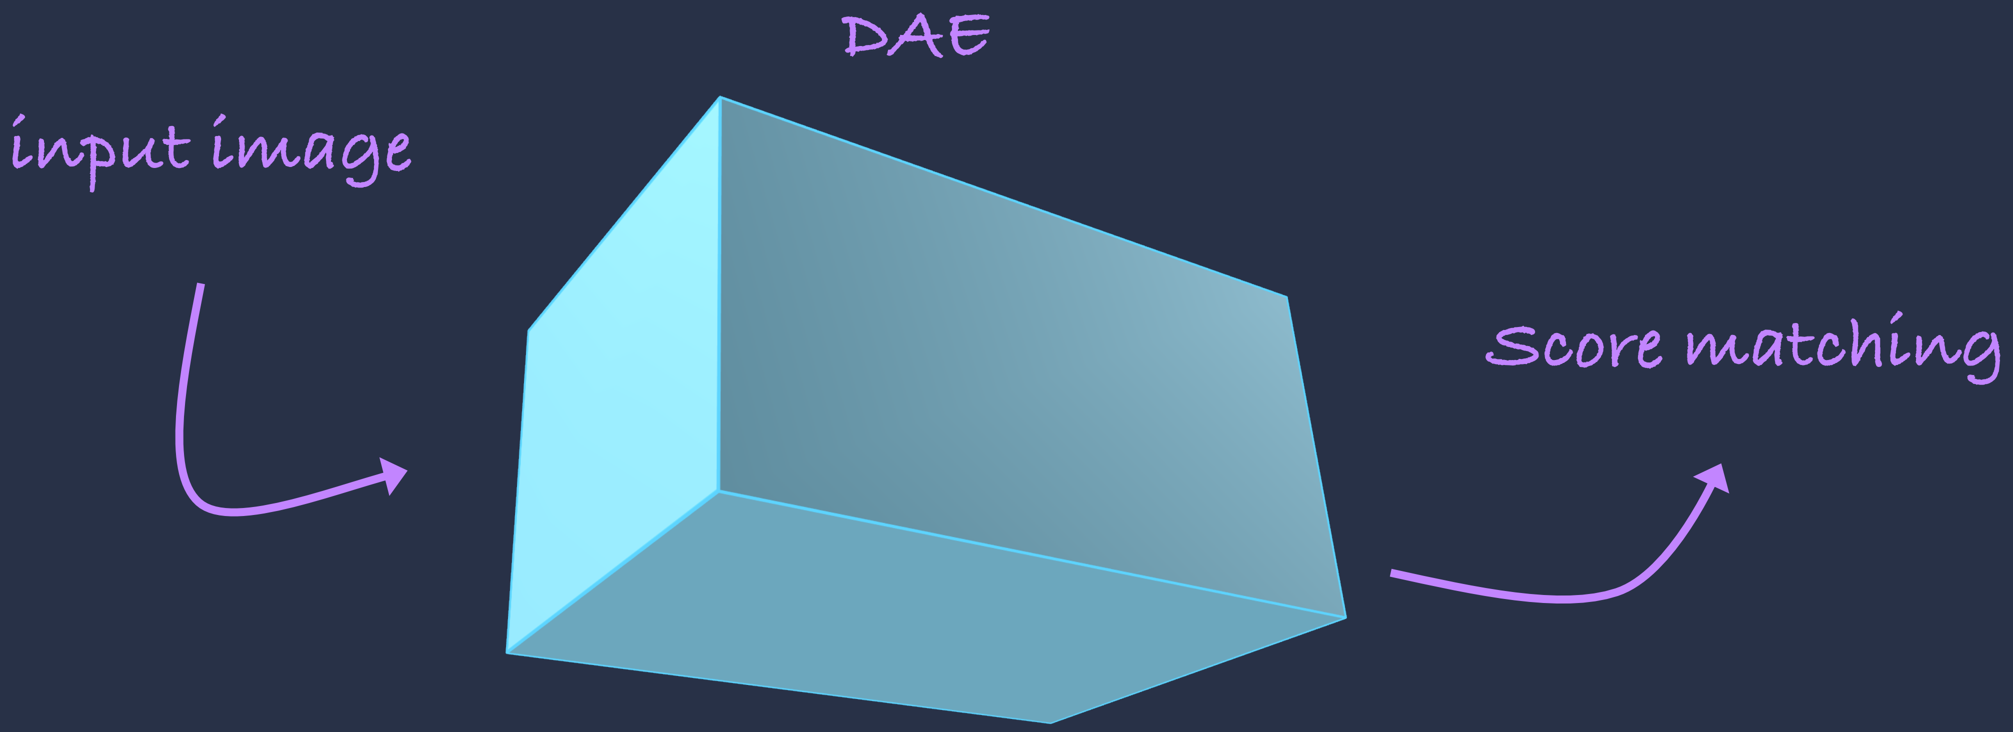

Image restoration using autoencoding priors [Bigdeli and Zwicker, 2017]
Deep mean-shift priors for image restoration [Bigdeli et al., 2018]

Image restoration: Denoising Autoencoders

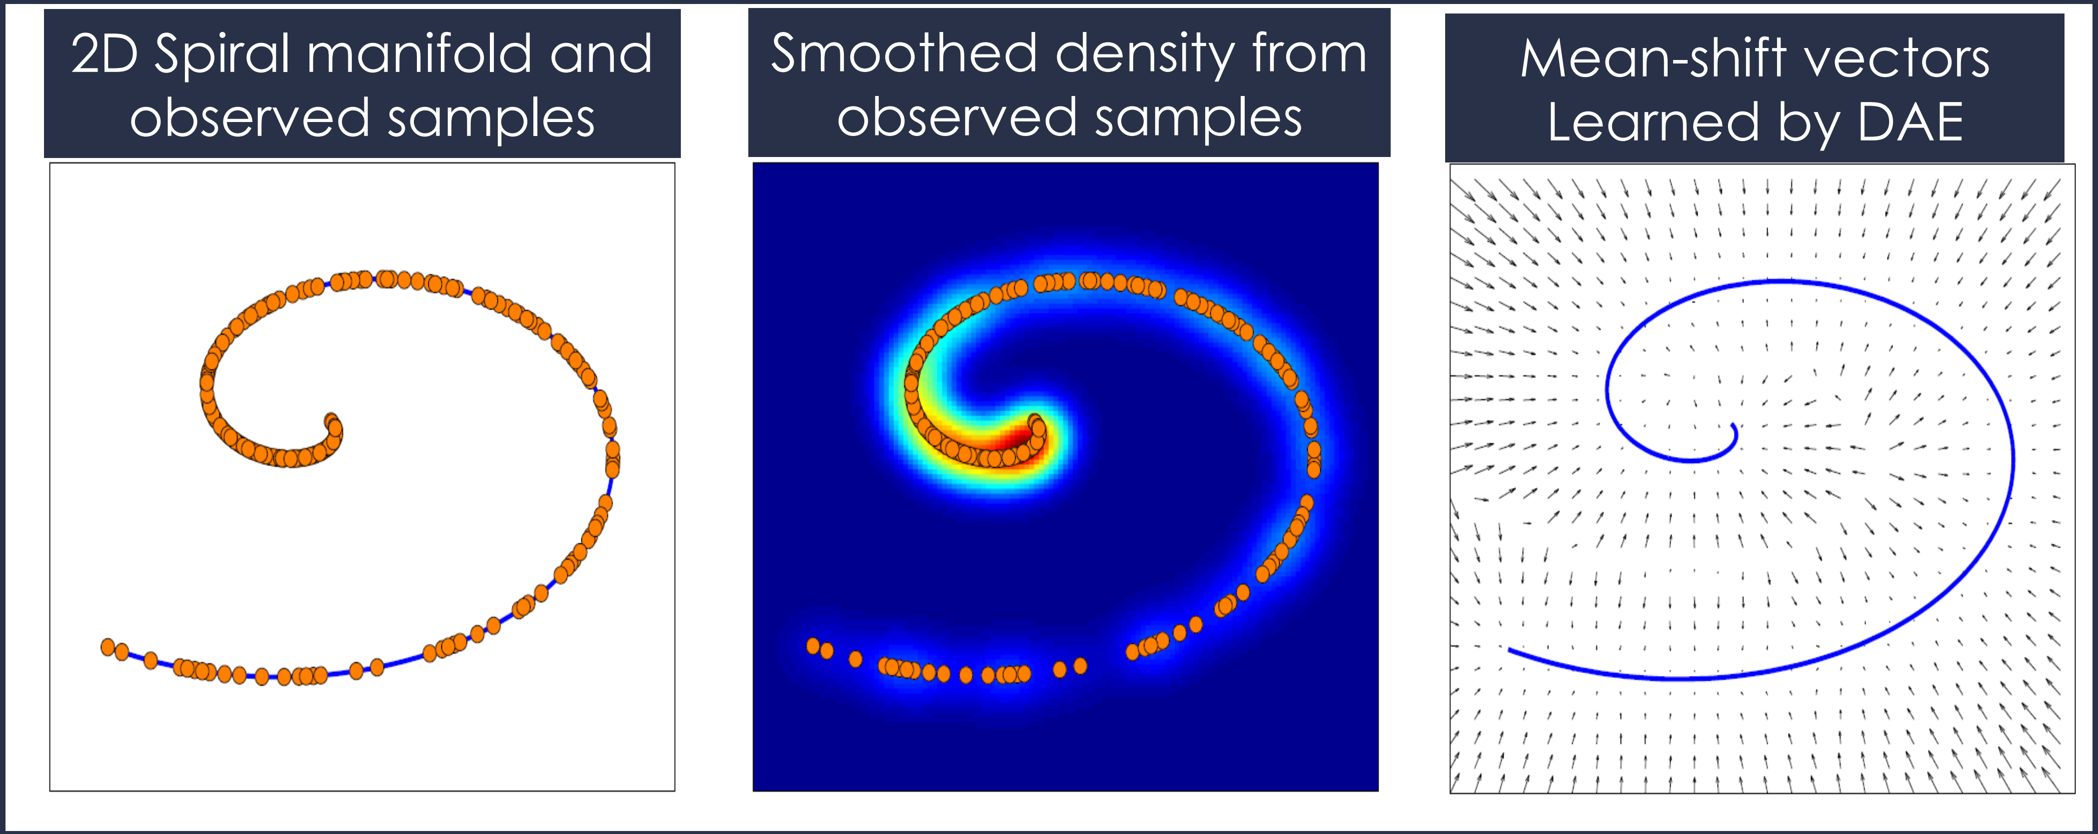

Image restoration using autoencoding priors [Bigdeli and Zwicker, 2017]
Deep mean-shift priors for image restoration [Bigdeli et al., 2018]

Image restoration: Denoising Autoencoder (DAE)

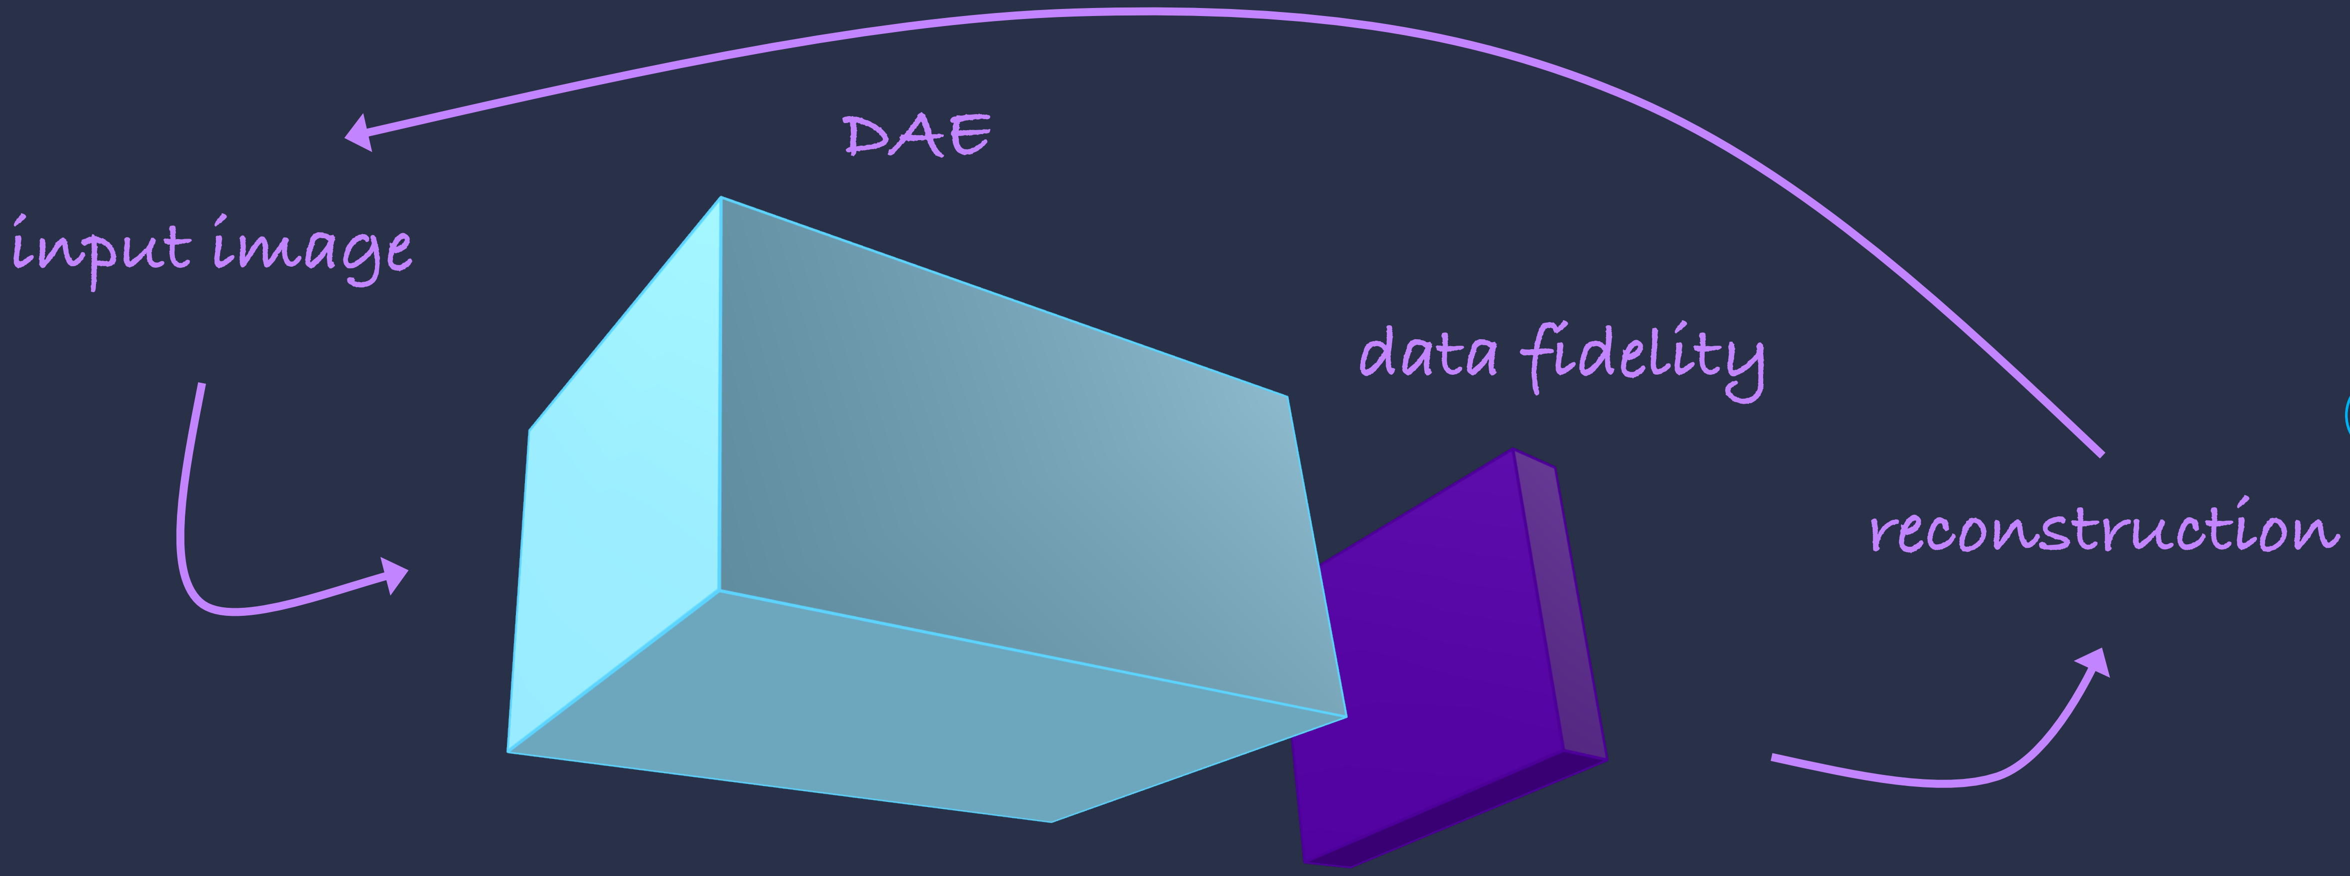

Image restoration: Gradient descent to find MAP

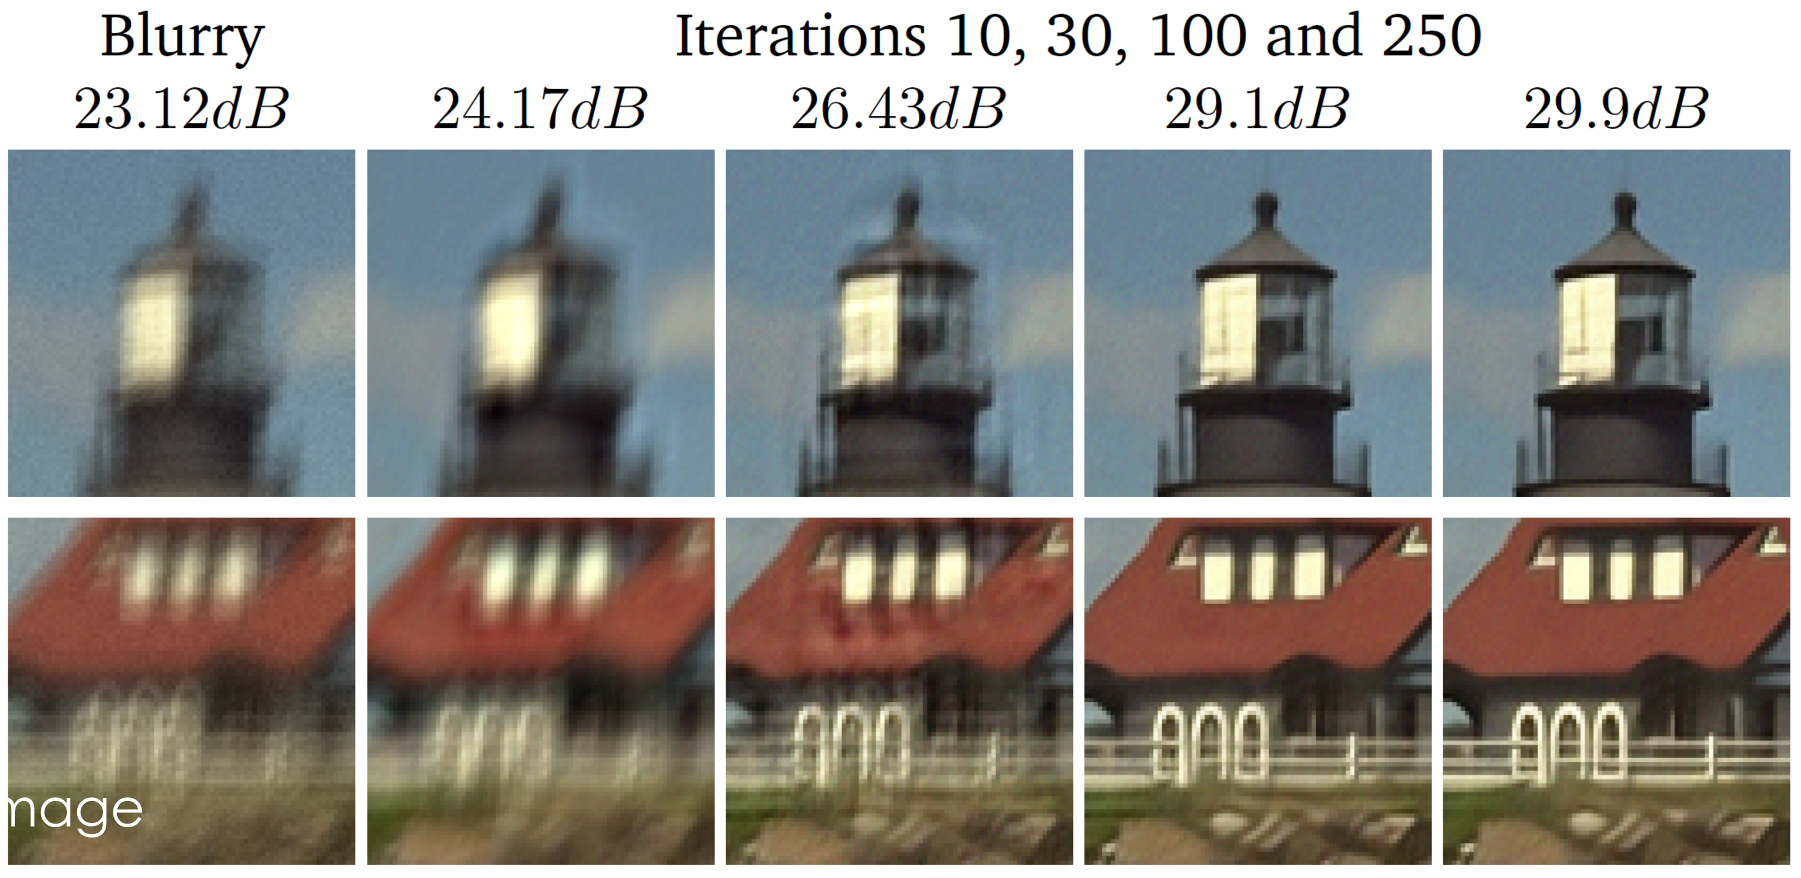

Natural image

Image restoration: Removing blur

Blurry and noisy

Dataset: Kodak Image Suite, 2013

Image restoration: Removing blur

Blurry and noisy

Dataset: Kodak Image Suite, 2013

Image restoration: image up-sampling (super-resolution)

Deep mean-shift priors for image restoration [Bigdeli et al., 2018]
Image Denoising via MAP estimation using Deep Neural Networks [Bigdeli and Süsstrunk, 2019]

4x down-sampled

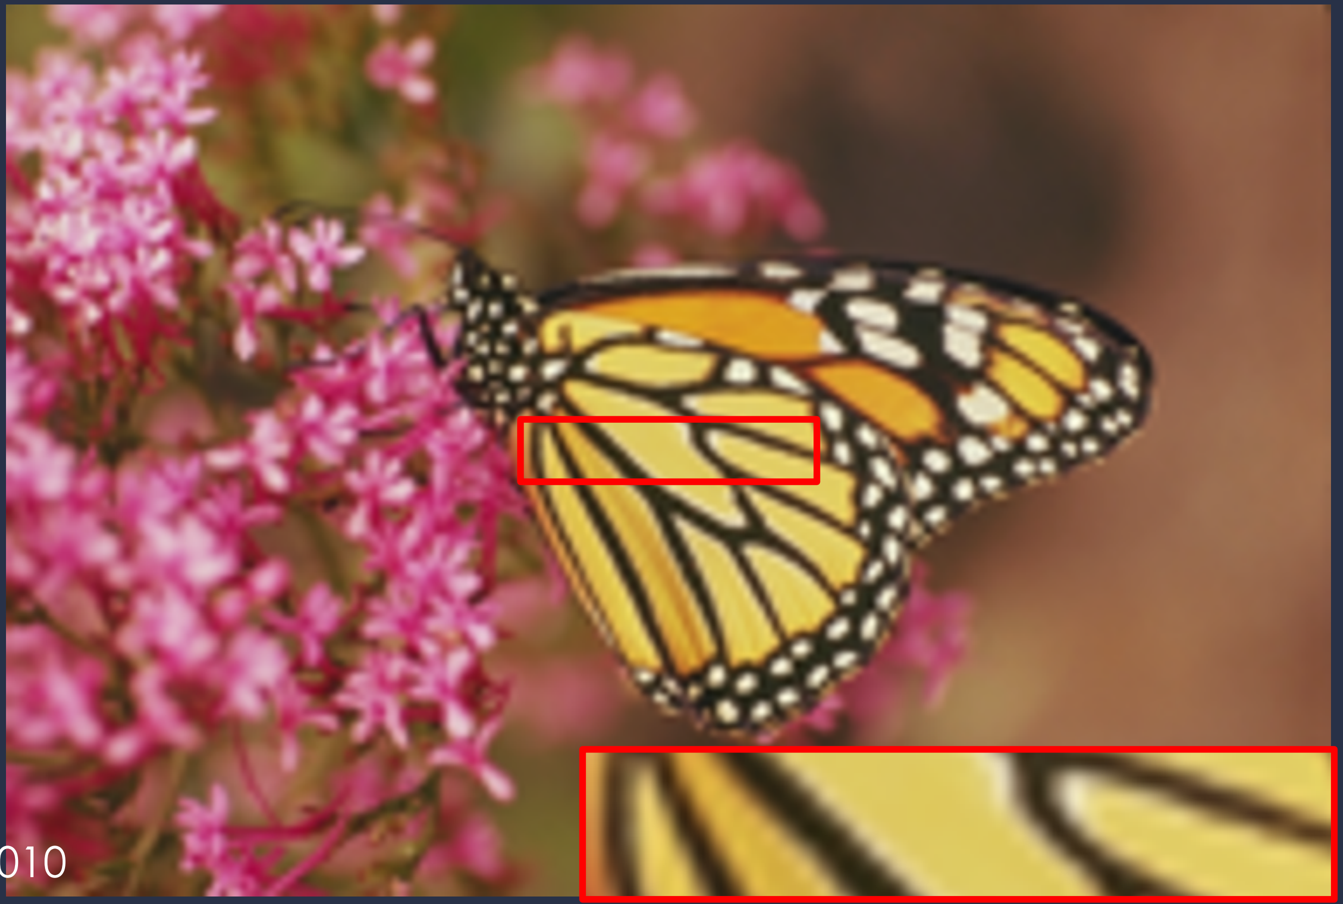

Dataset: Set14, Zeyed et al, 2010

:: csem

Image restoration: image up-sampling (super-resolution)

Deep mean-shift priors for image restoration [Bigdeli et al., 2018]
Image Denoising via MAP estimation using Deep Neural Networks [Bigdeli and Süsstrunk, 2019]

Our solution

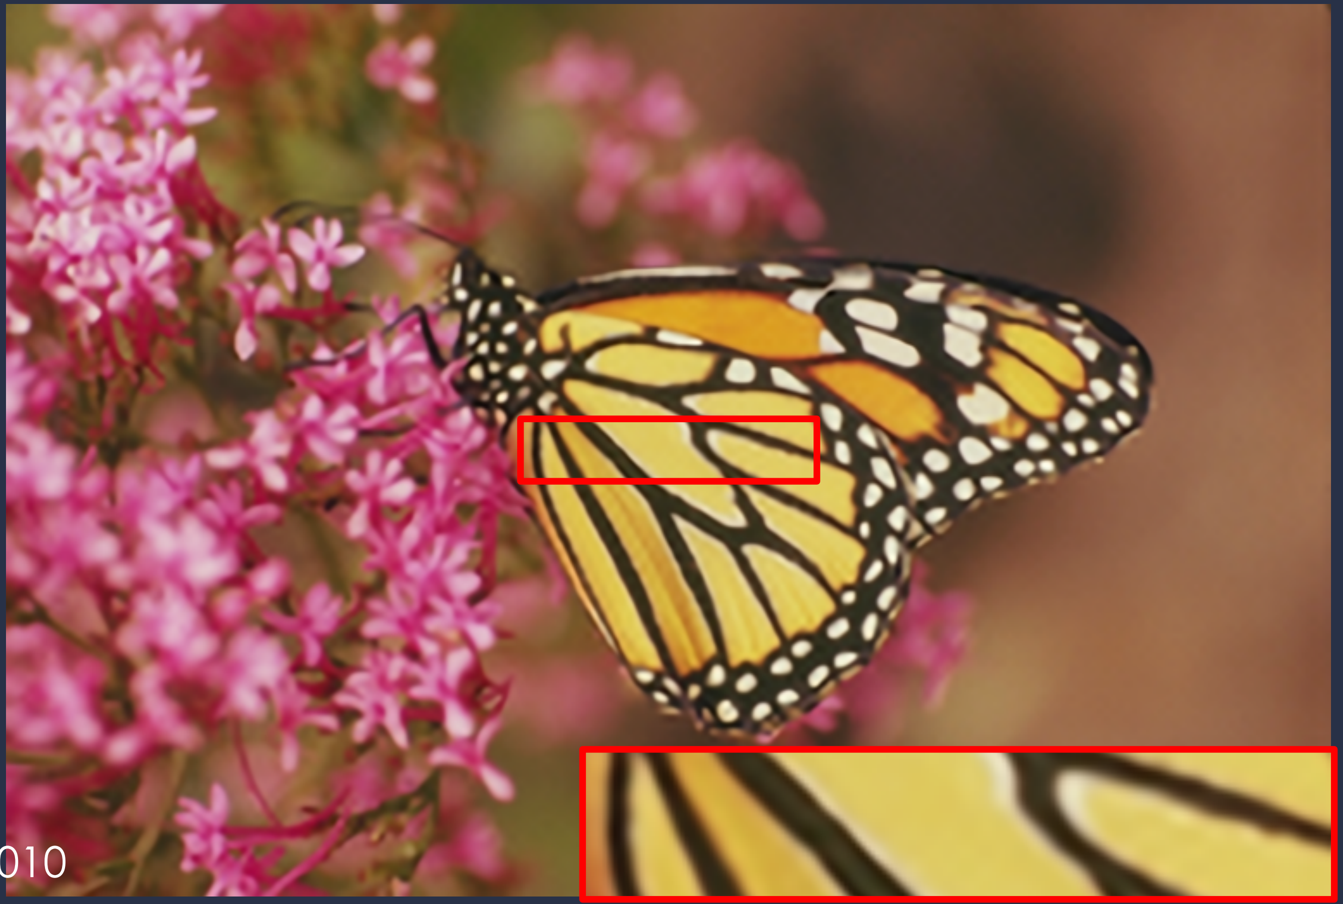

Dataset: Set14, Zeyed et al, 2010

Image restoration: Comparing against end-to-end

Image restoration using autoencoding priors [Bigdeli and Zwicker, 2017]

- DNCNN: trained on ×2,×3,×4

Method	$\times 2$	$\times 3$	$\times 4$	$\times 5$
Bicubic	28.53	25.92	24.44	23.46
SRCNN	30.52	27.48	25.76	24.05
TNRD	30.53	27.60	25.92	24.61
VDSR	30.72	27.81	26.16	24.01
DnCNN-3	30.99	27.93	**26.25**	24.26
IRCNN	30.79	27.68	25.96	24.73
DAEP (Ours)	**31.07**	**27.93**	26.13	**24.88**

Image restoration: Hole-filling

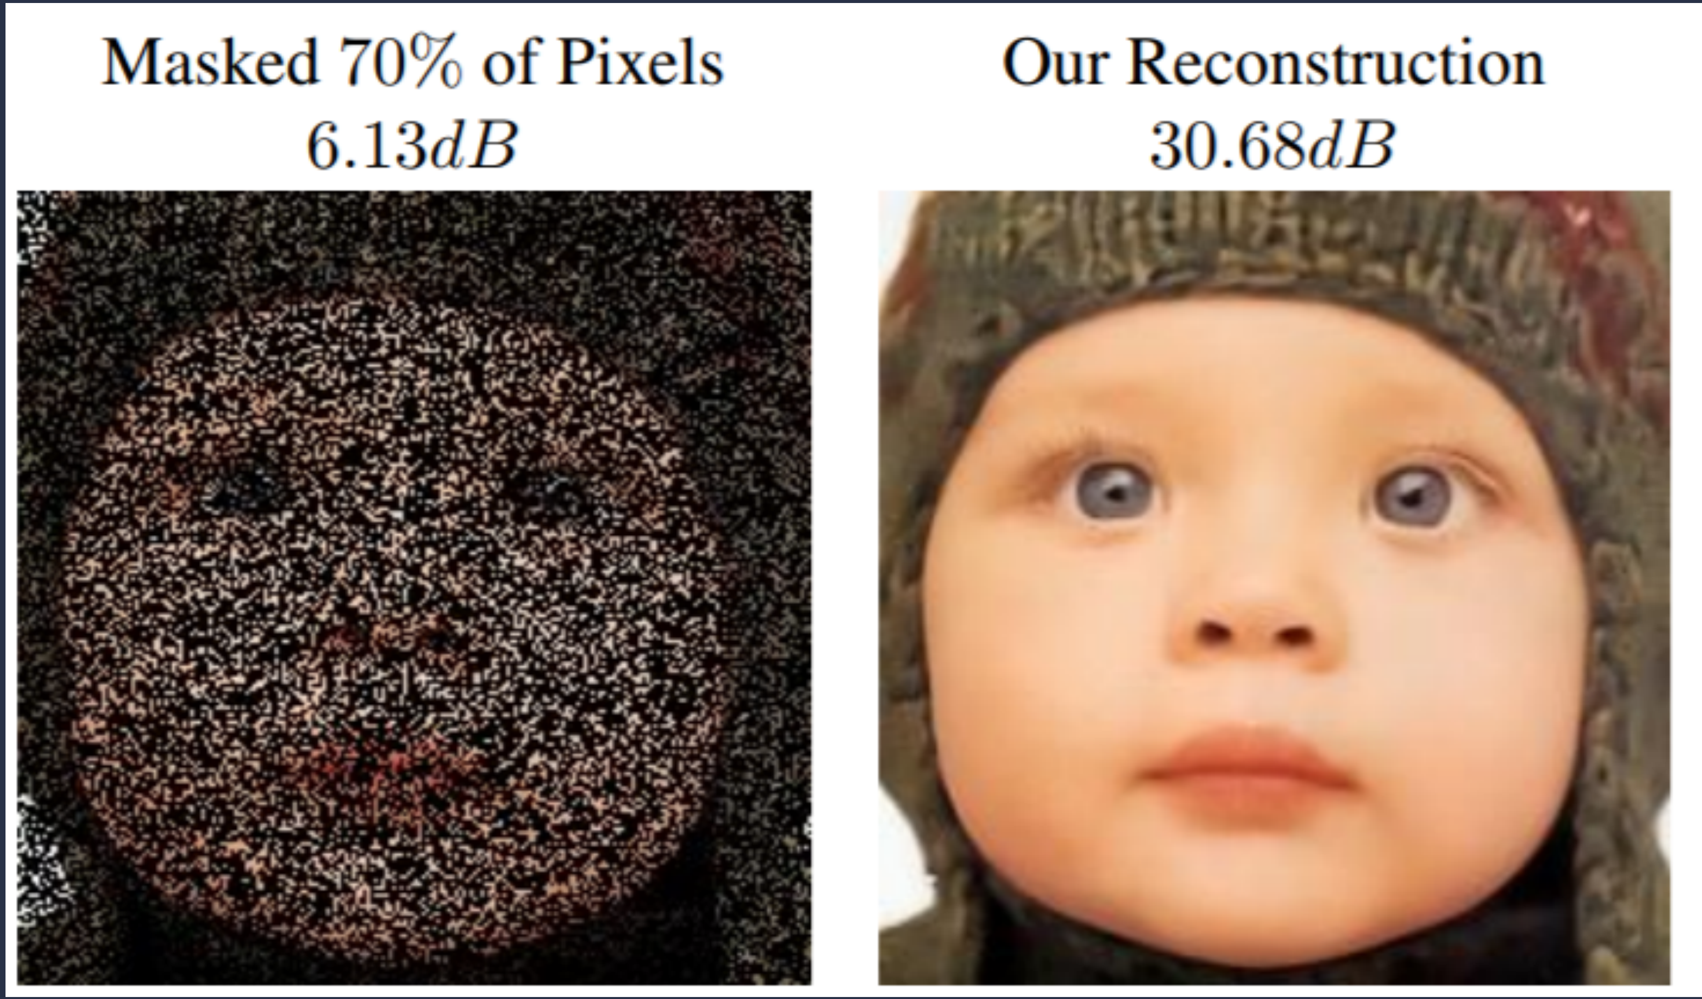

Image restoration using autoencoding priors [Bigdeli and Zwicker, 2017]
Deep mean-shift priors for image restoration [Bigdeli et al., 2018]

Image restoration: Comparing against end-to-end

Deep mean-shift priors for image restoration [Bigdeli et al., 2018]

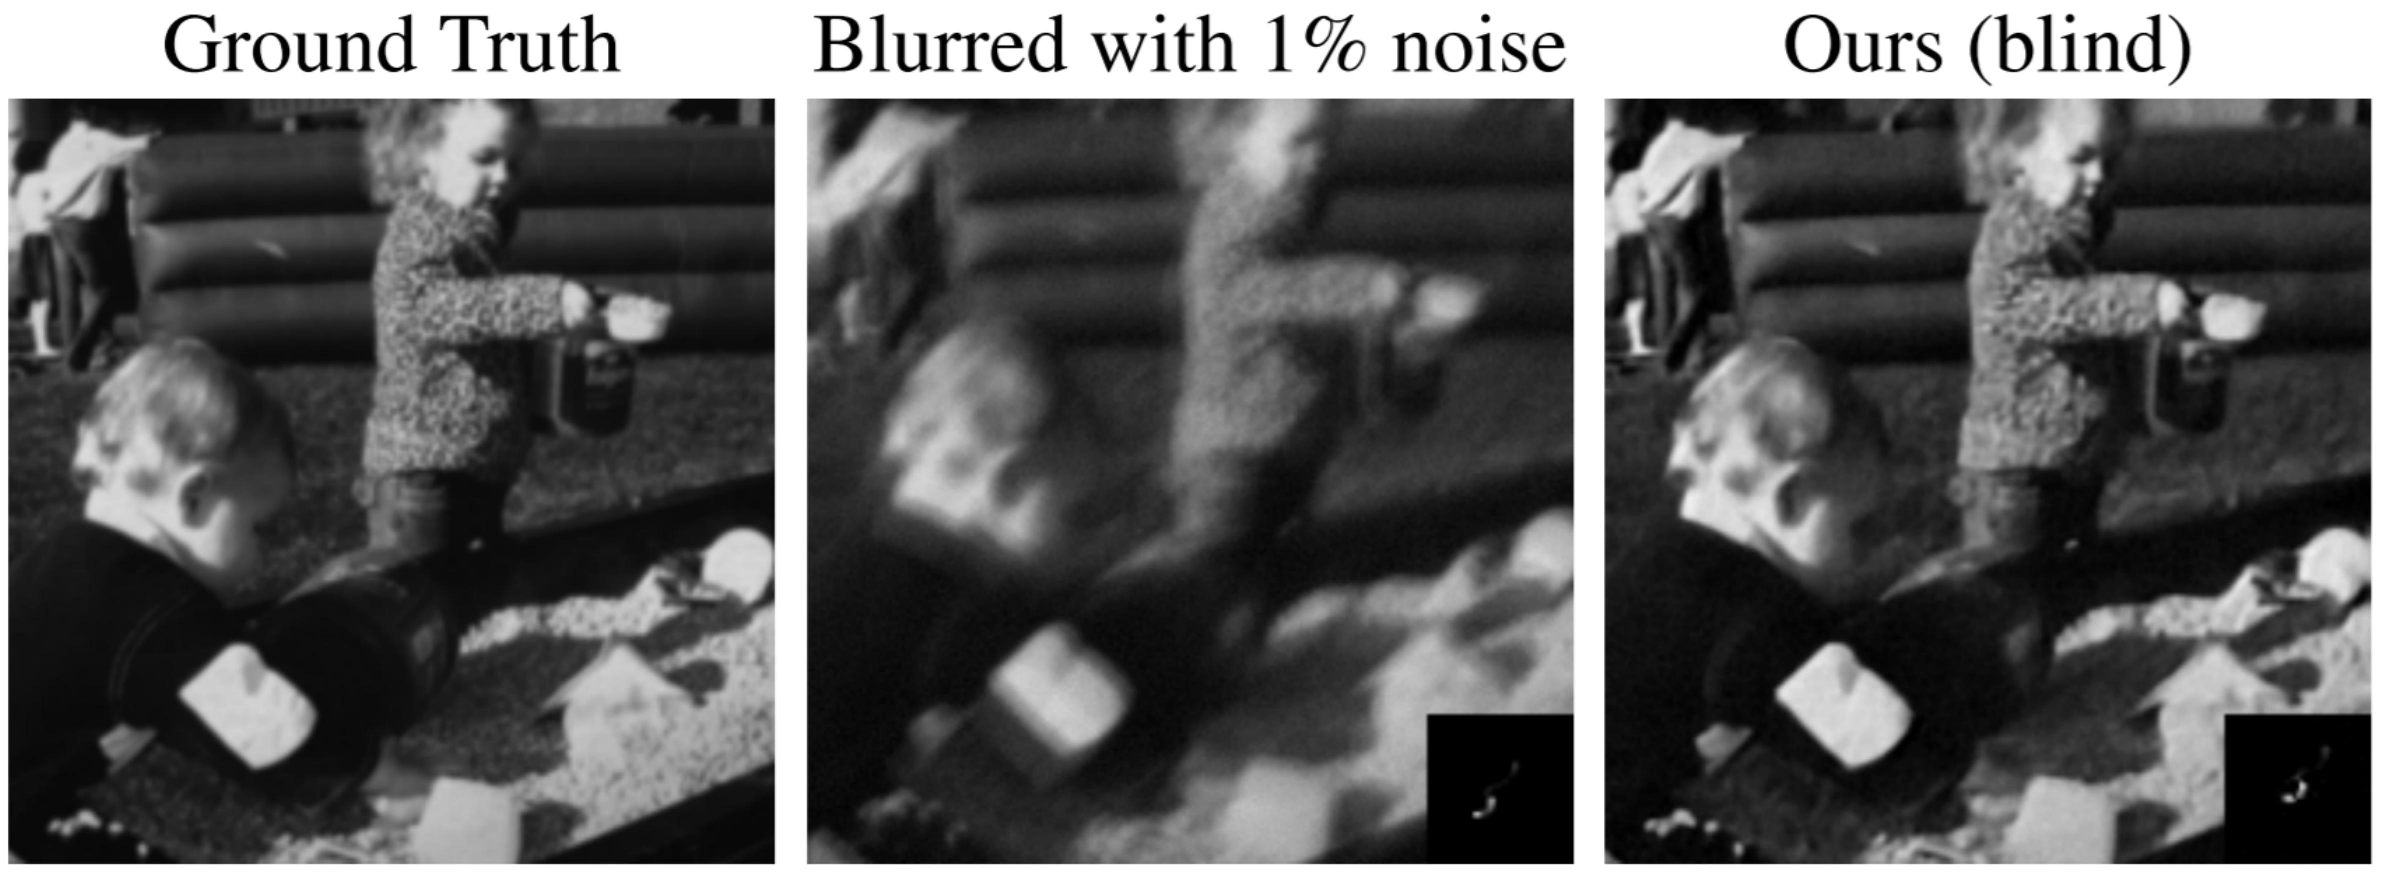

Dataset: Levin et al., 2007

Image restoration: ADMM -> 70x faster

Image Denoising via MAP estimation using Deep Neural Networks [Bigdeli and Süsstrunk, 2019]

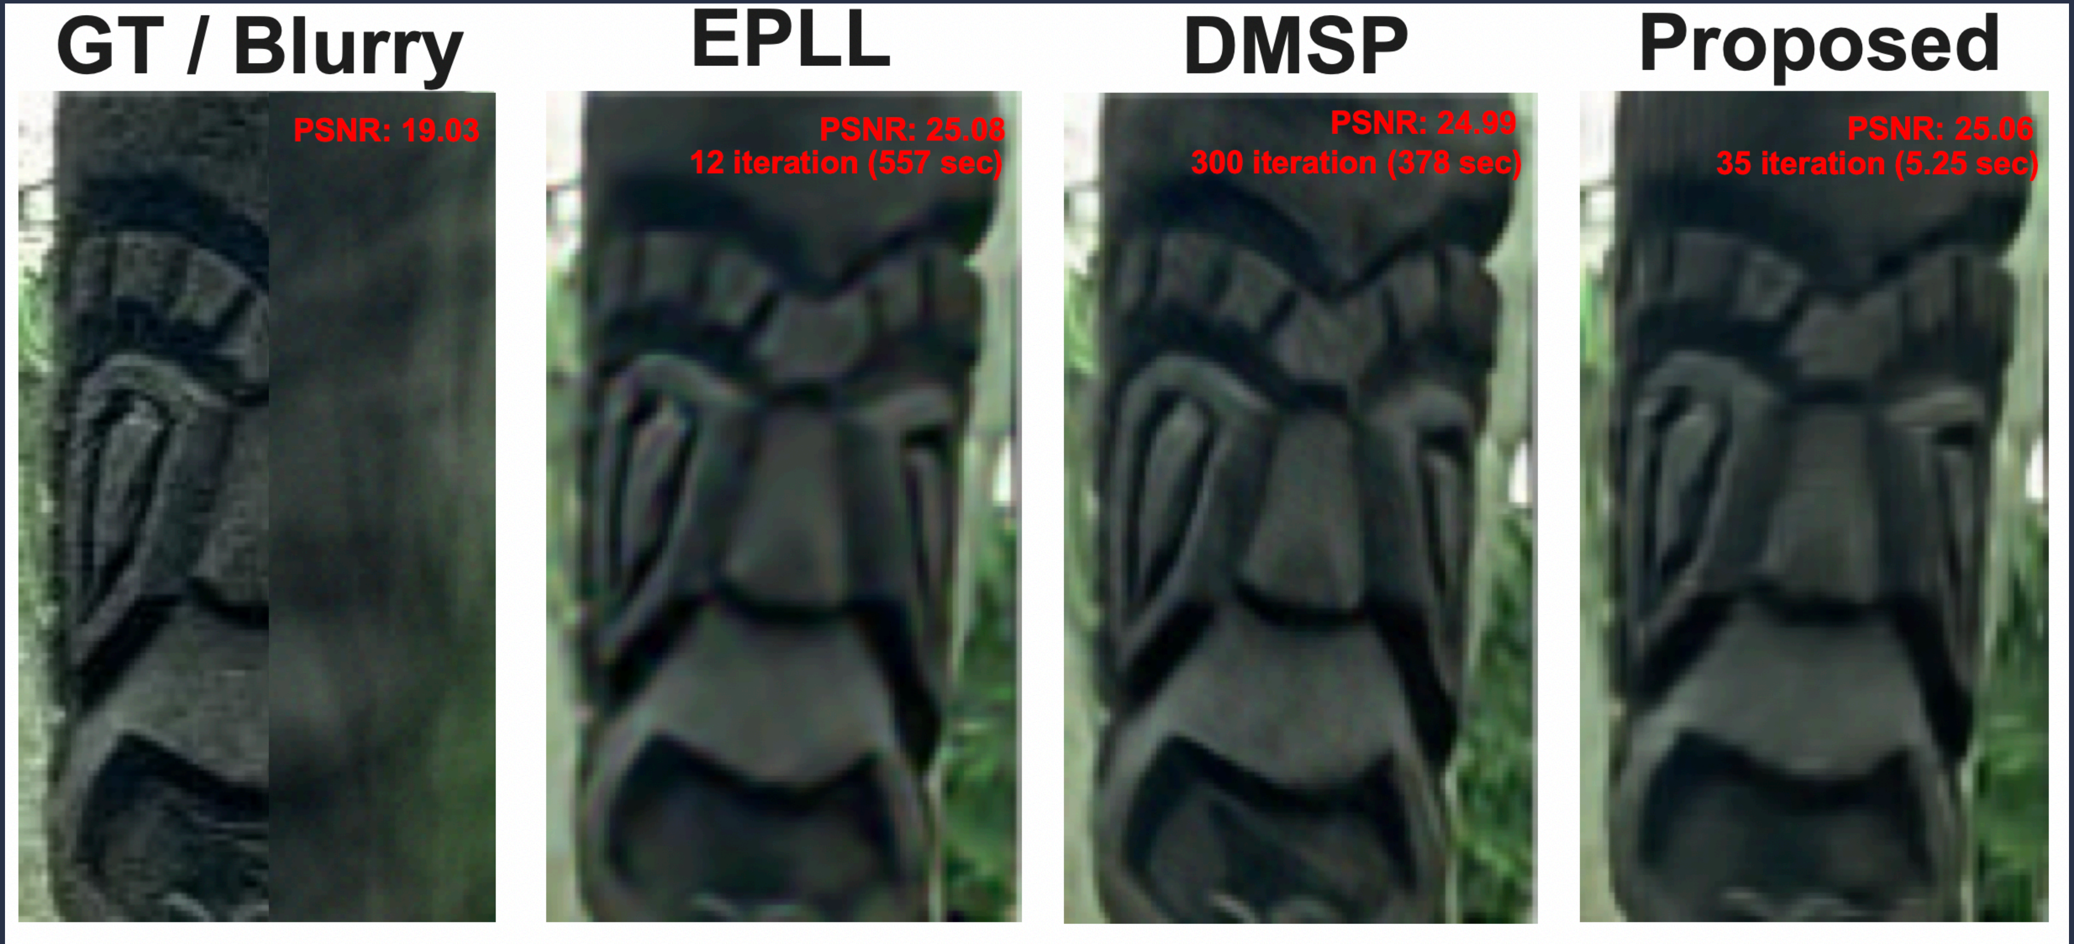

Image restoration: Summary

- Unsupervised training for one network
- Several restoration tasks
- Without degradation parameters at train time
- Optimizing degradation parameters at test time
  - ✓ Can be used to learn the optics error
  - ✓ Can be used for noise modeling

0 ENERGY COMPUTING: Convolution optics as a Neural Network Layer

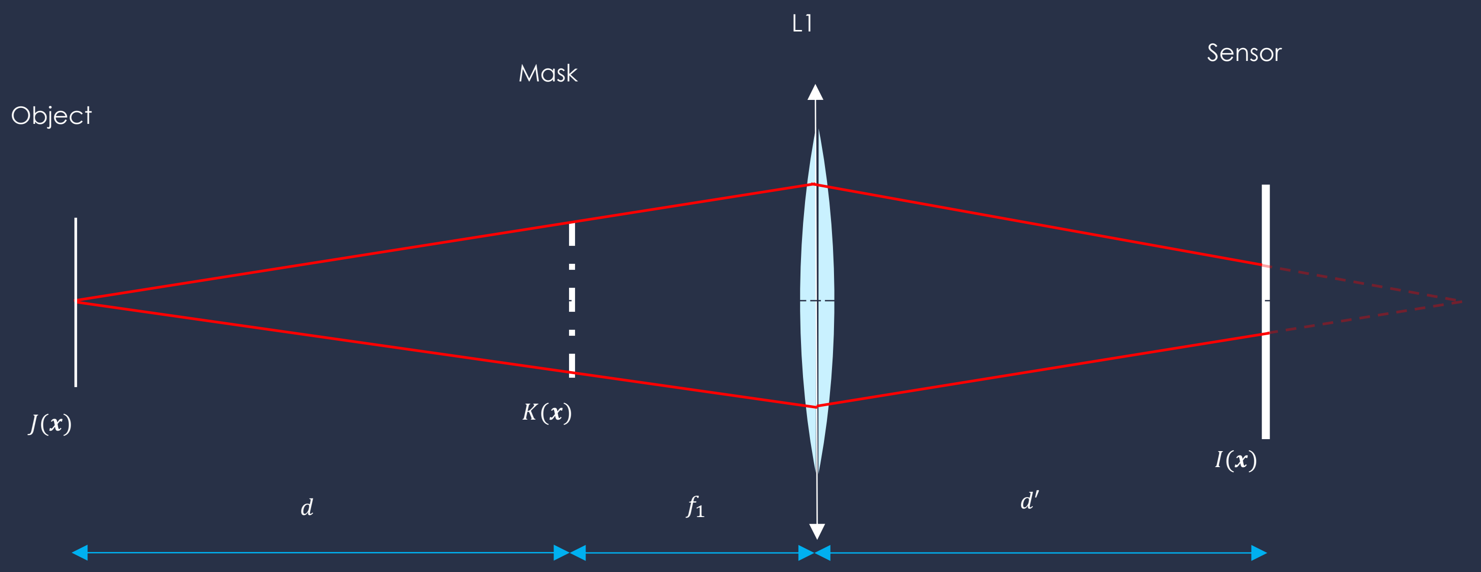

$\alpha = -\frac{f_1}{d}$

$\beta = 1 + \frac{d}{f_1}\left(1 - \frac{d'}{f_1}\right)$

- Field intensity on I(x)

$$\mathrm{I}(\boldsymbol{x}) = \int J\left(\frac{1}{\alpha}\boldsymbol{x} - \overline{\boldsymbol{u}}\right) K\left(-\frac{1}{\beta}\overline{\boldsymbol{u}}\right) w\left(\frac{1}{\beta f_1}\overline{\boldsymbol{u}}\right) d\overline{\boldsymbol{u}} \longrightarrow I(\boldsymbol{x}) = \left(J\left(\frac{1}{\alpha}\cdot\right) * \frac{1}{(\beta f_1)^2} K\left(-\frac{1}{\beta}\cdot\right) w\left(\frac{1}{\beta f_1}\cdot\right)\right)(\boldsymbol{x})$$

Scaled convolution between the field J(x) and K(x)

csem

Convolution measurment

pinhole

cross

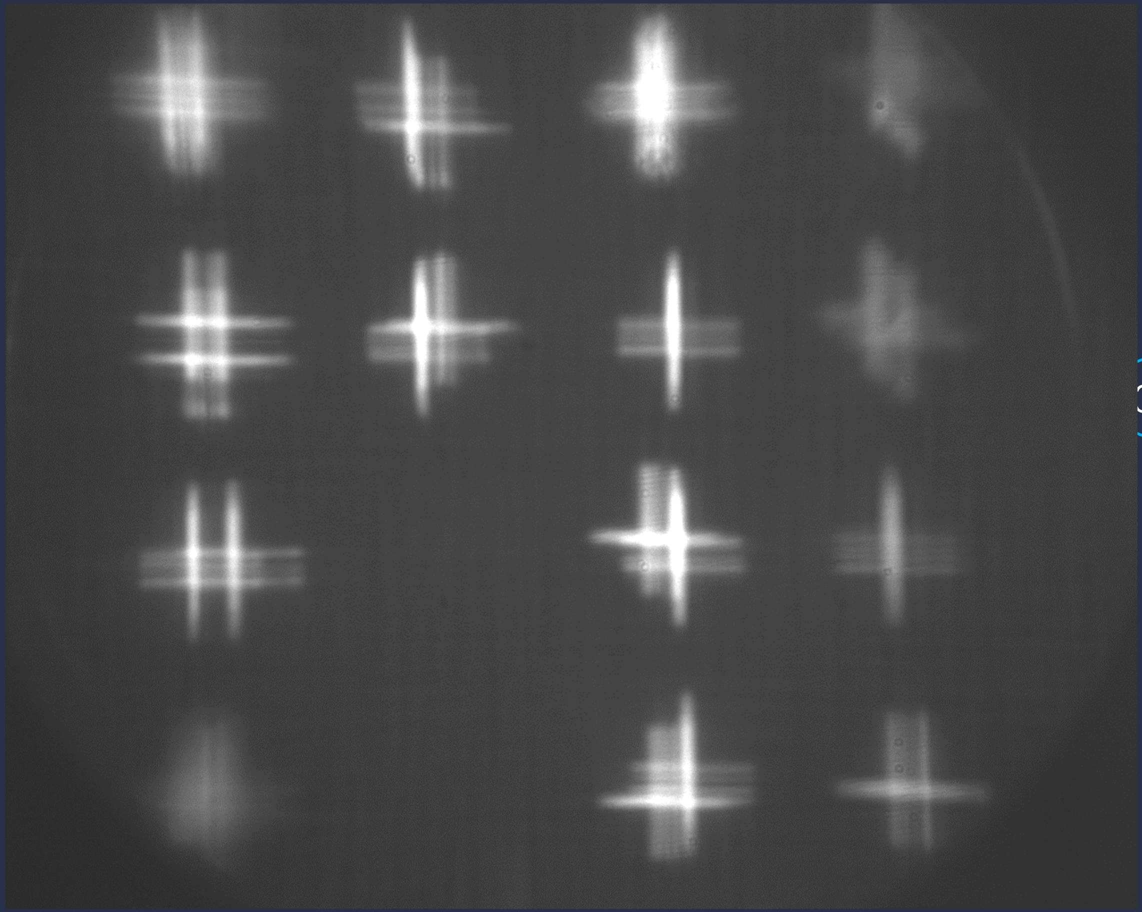

Workflow

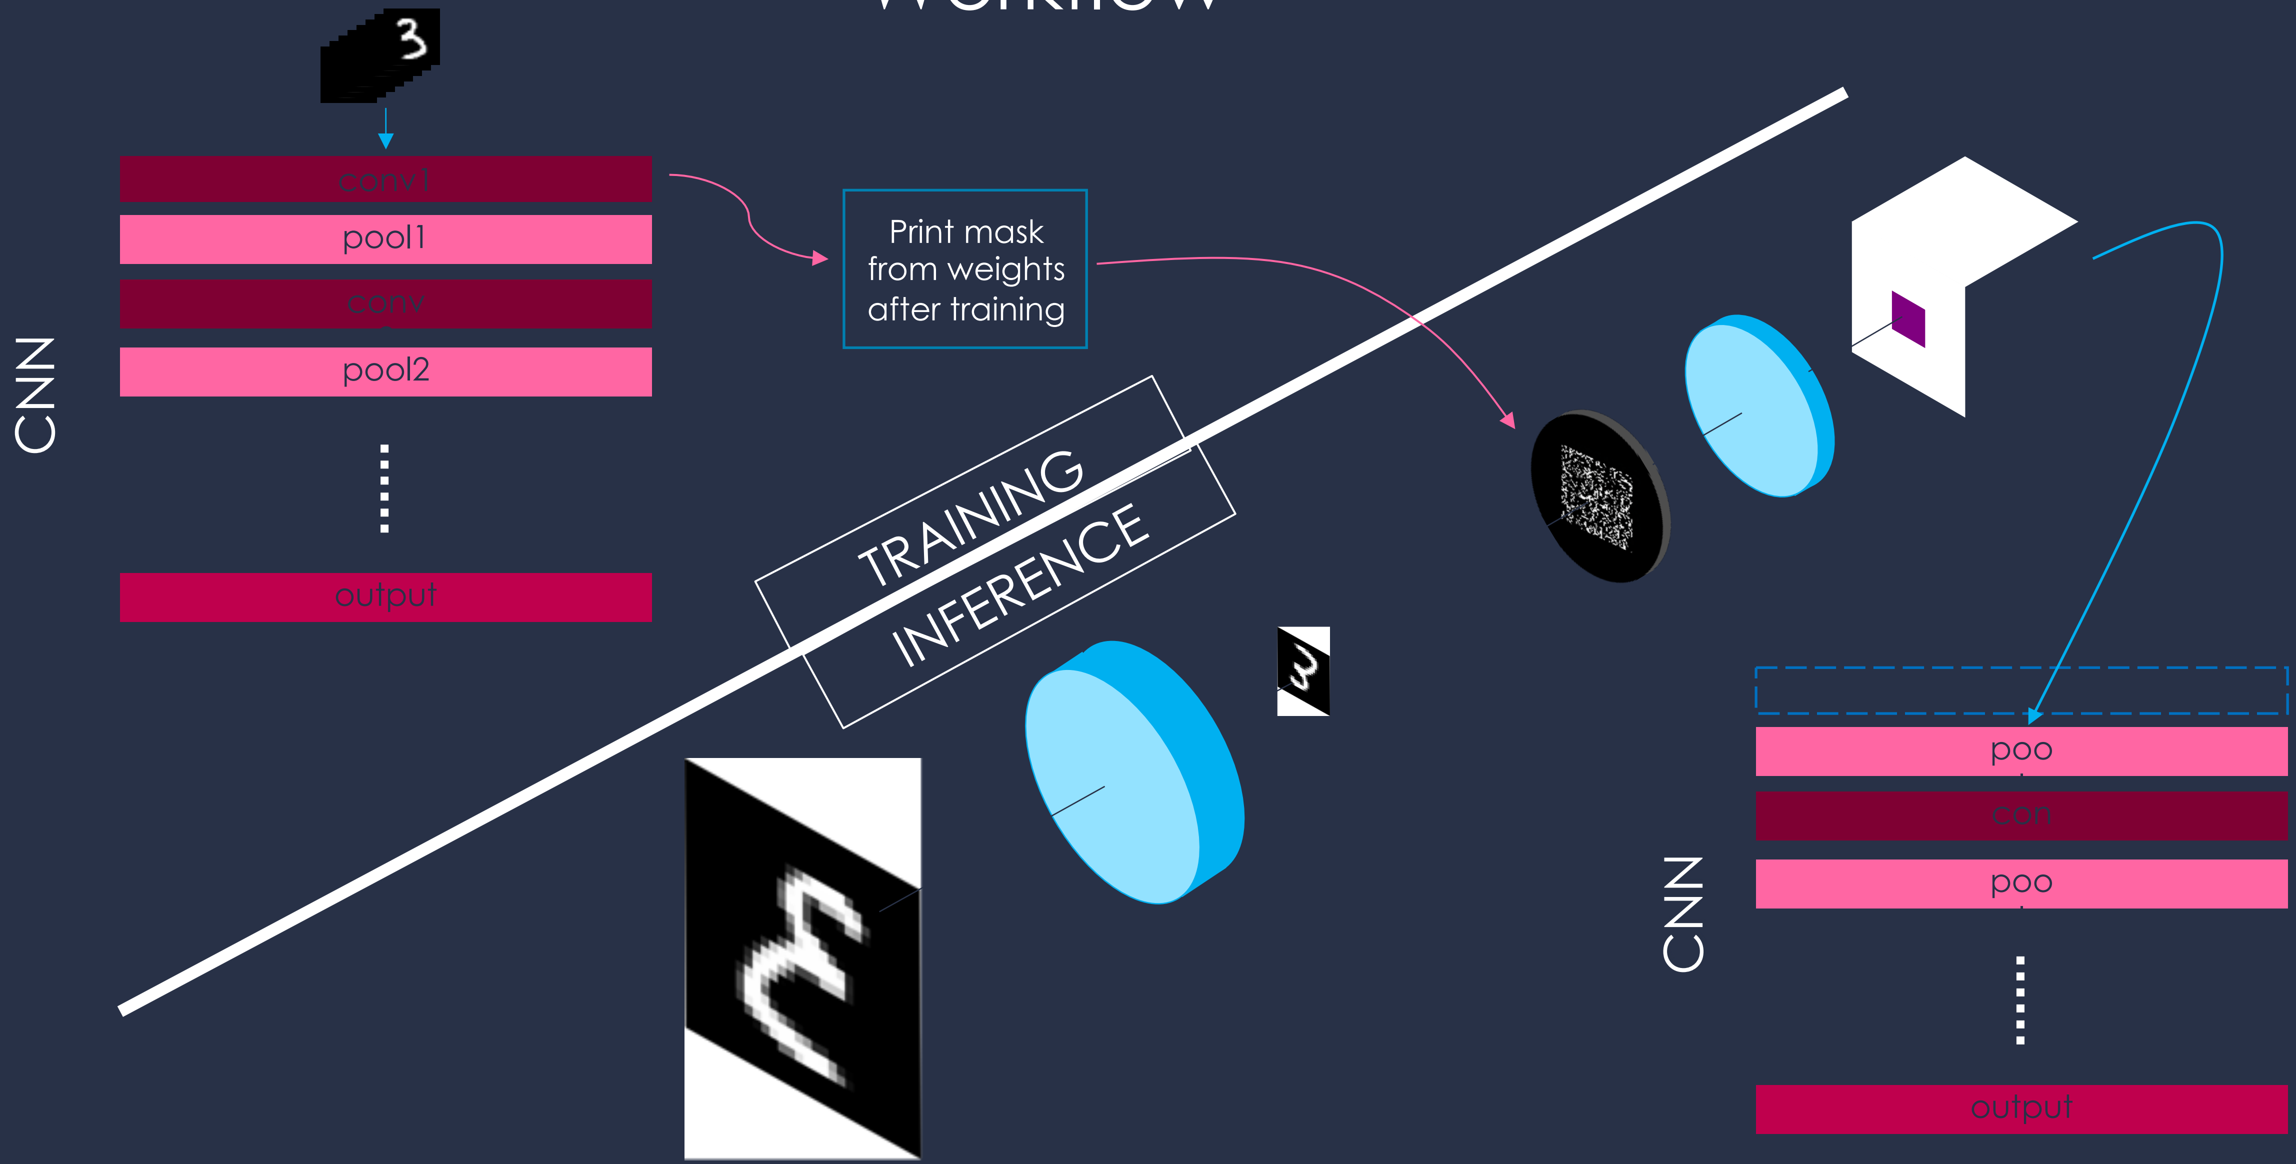

CNN
conv1
pool1
conv
pool2
output
Print mask from weights after training
TRAINING
INFERENCE
CNN
poo
con
poo
output

Digits recognition (MNIST) filters of the 1st layer

Inkjet printing
120 µm
resolution

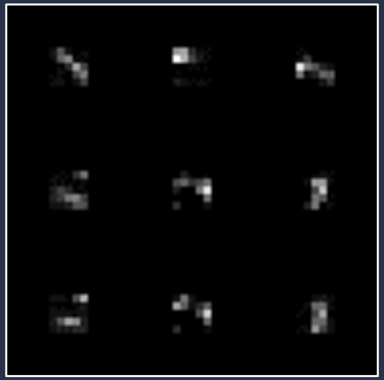

Set-up

mask

Sensor

Image plane

Objective lenses (f = 27 mm)

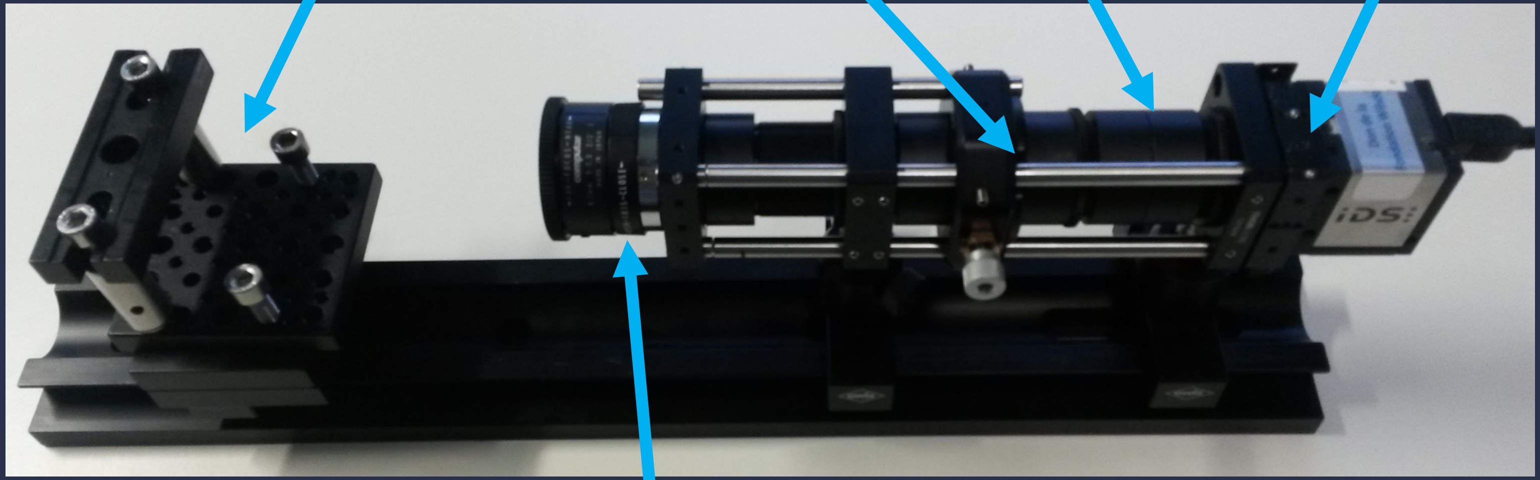

Front lens (f = 8.5 mm)

Printed mask

Output

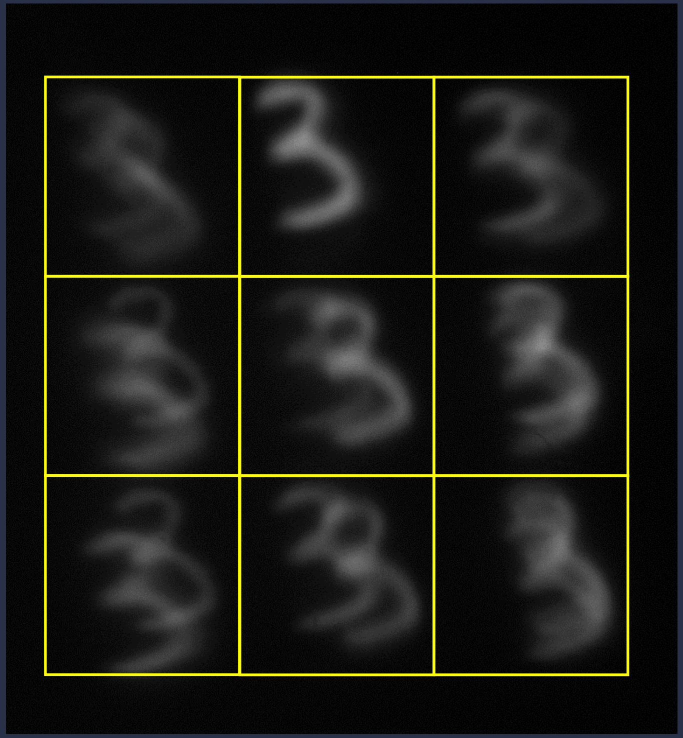

Input

Proof of concept

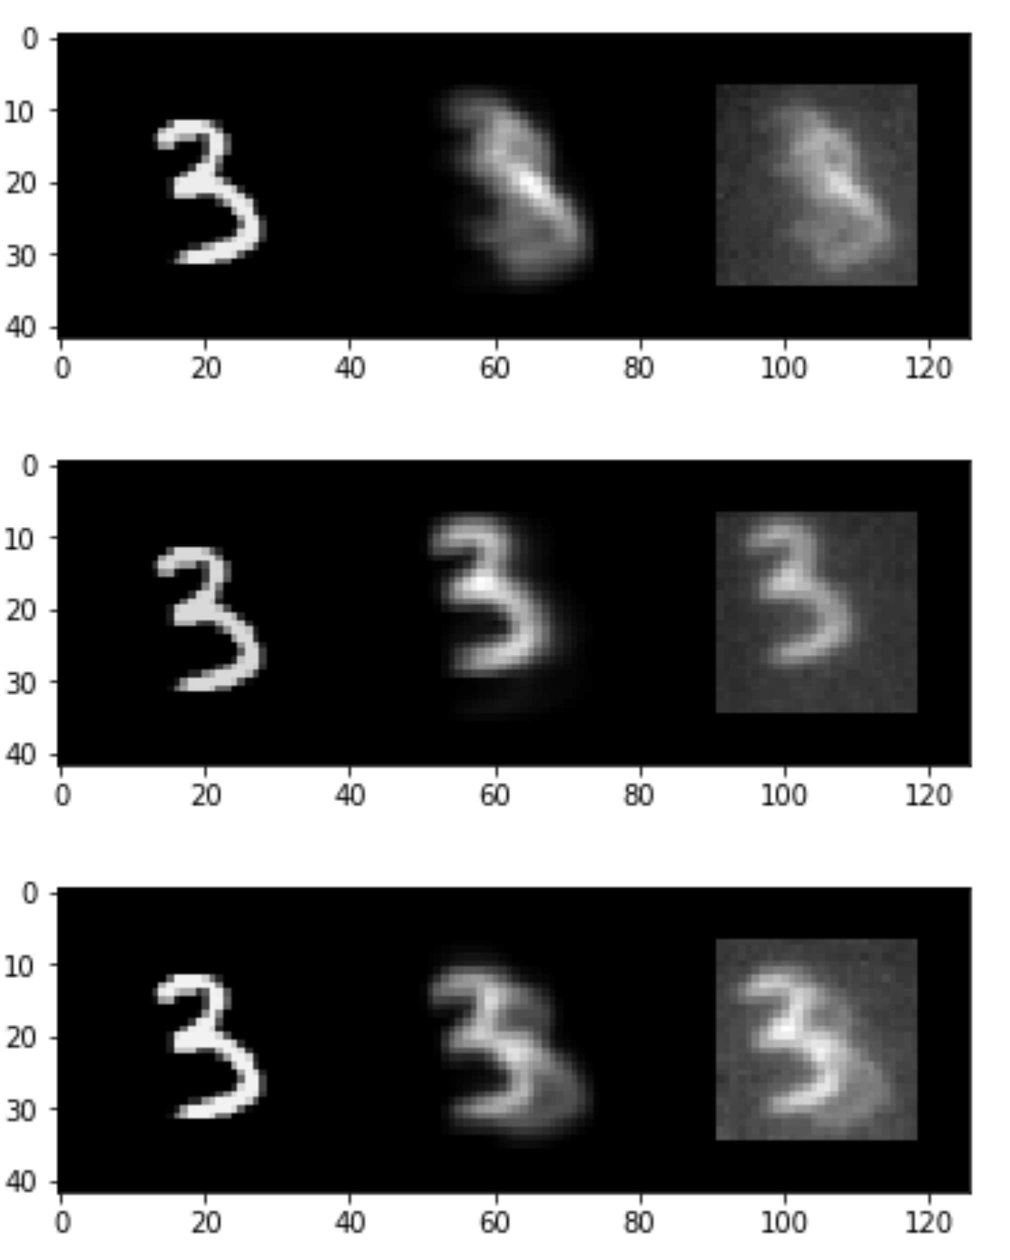

:: csem

Thank you