Optimised object decection: a completely non-unified approach

Edward Rosten

COMPUTER VISION CONSULTING
computervisionconsulting.com

Detecting stuff

- ‘Traditional’ object detection
  Identifiable ‘object’, in this case small objects.

- Superresolution fluorescence microscopy
  Objects are now just featureless dots—harder than it sounds.

- Corner detection
  Objects are now...? What are corners anyway?

Object detection

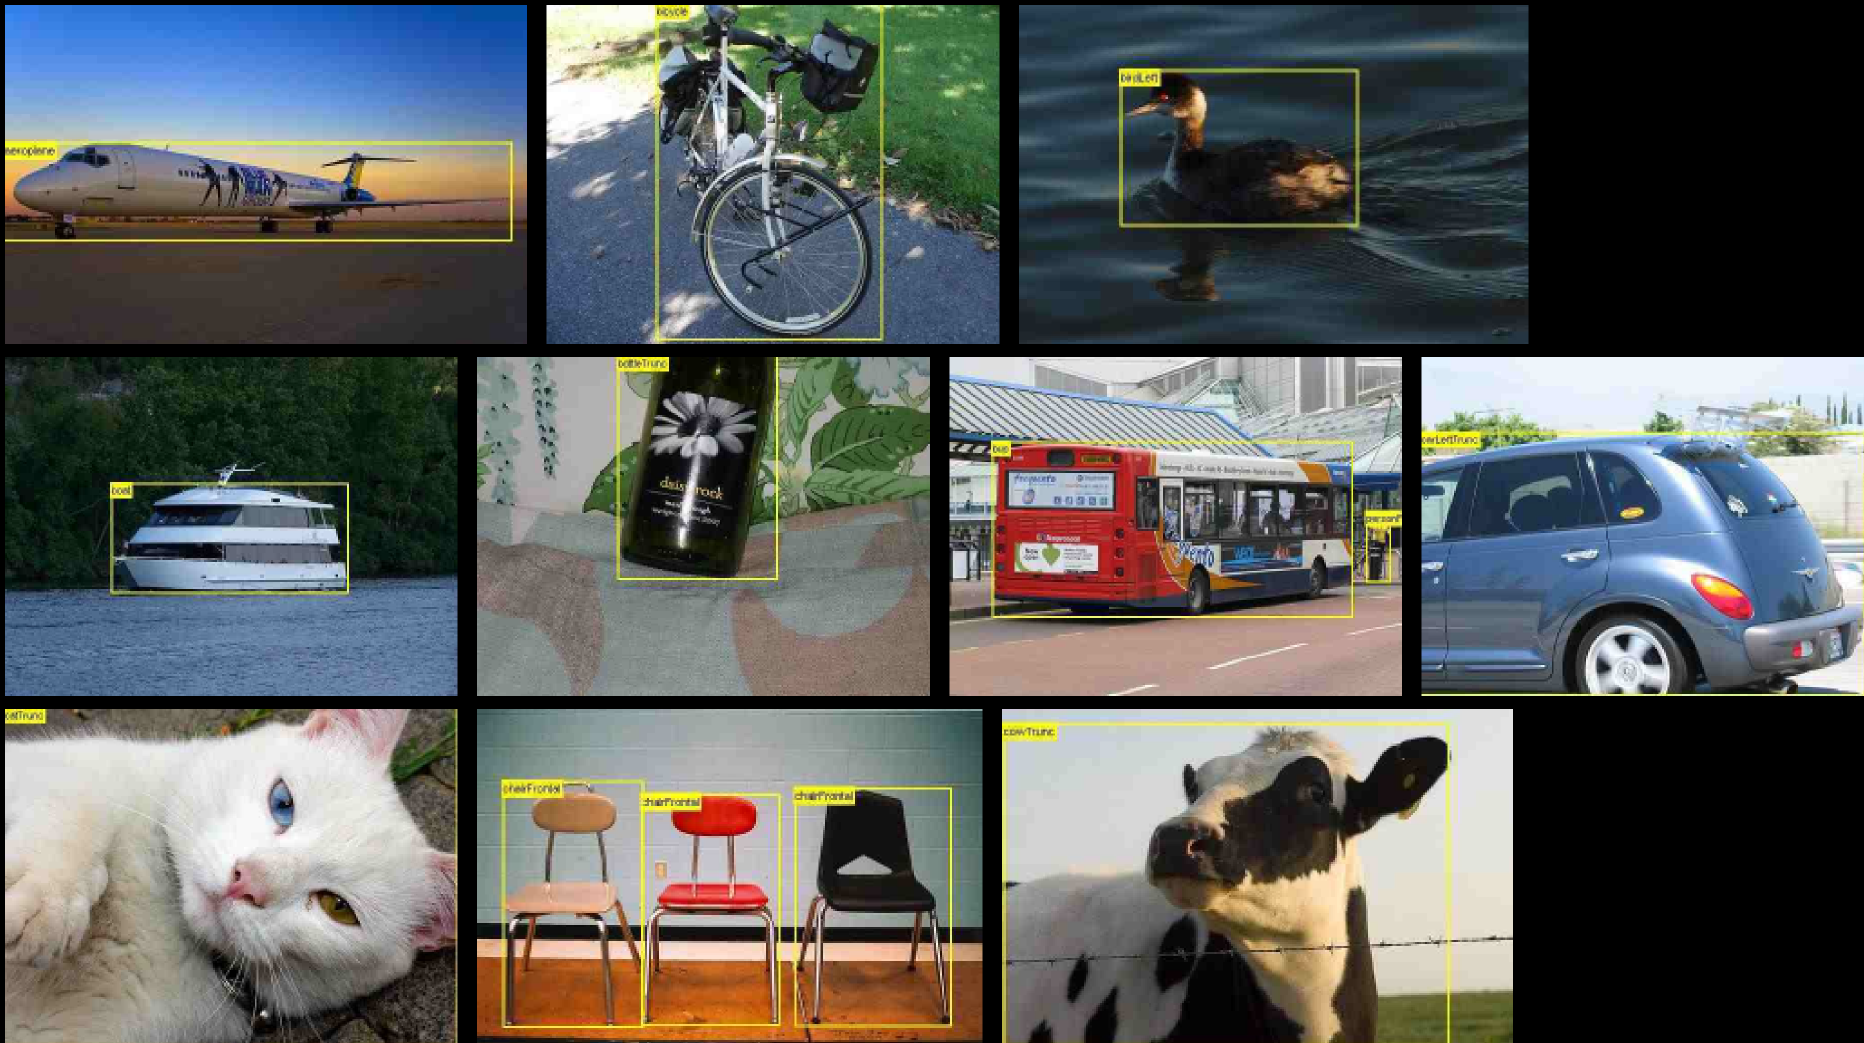

Everingham, Van Gool, Williams, Winn, Zisserman. The PASCAL Visual Object Classes Challenge 2011.

Object detection

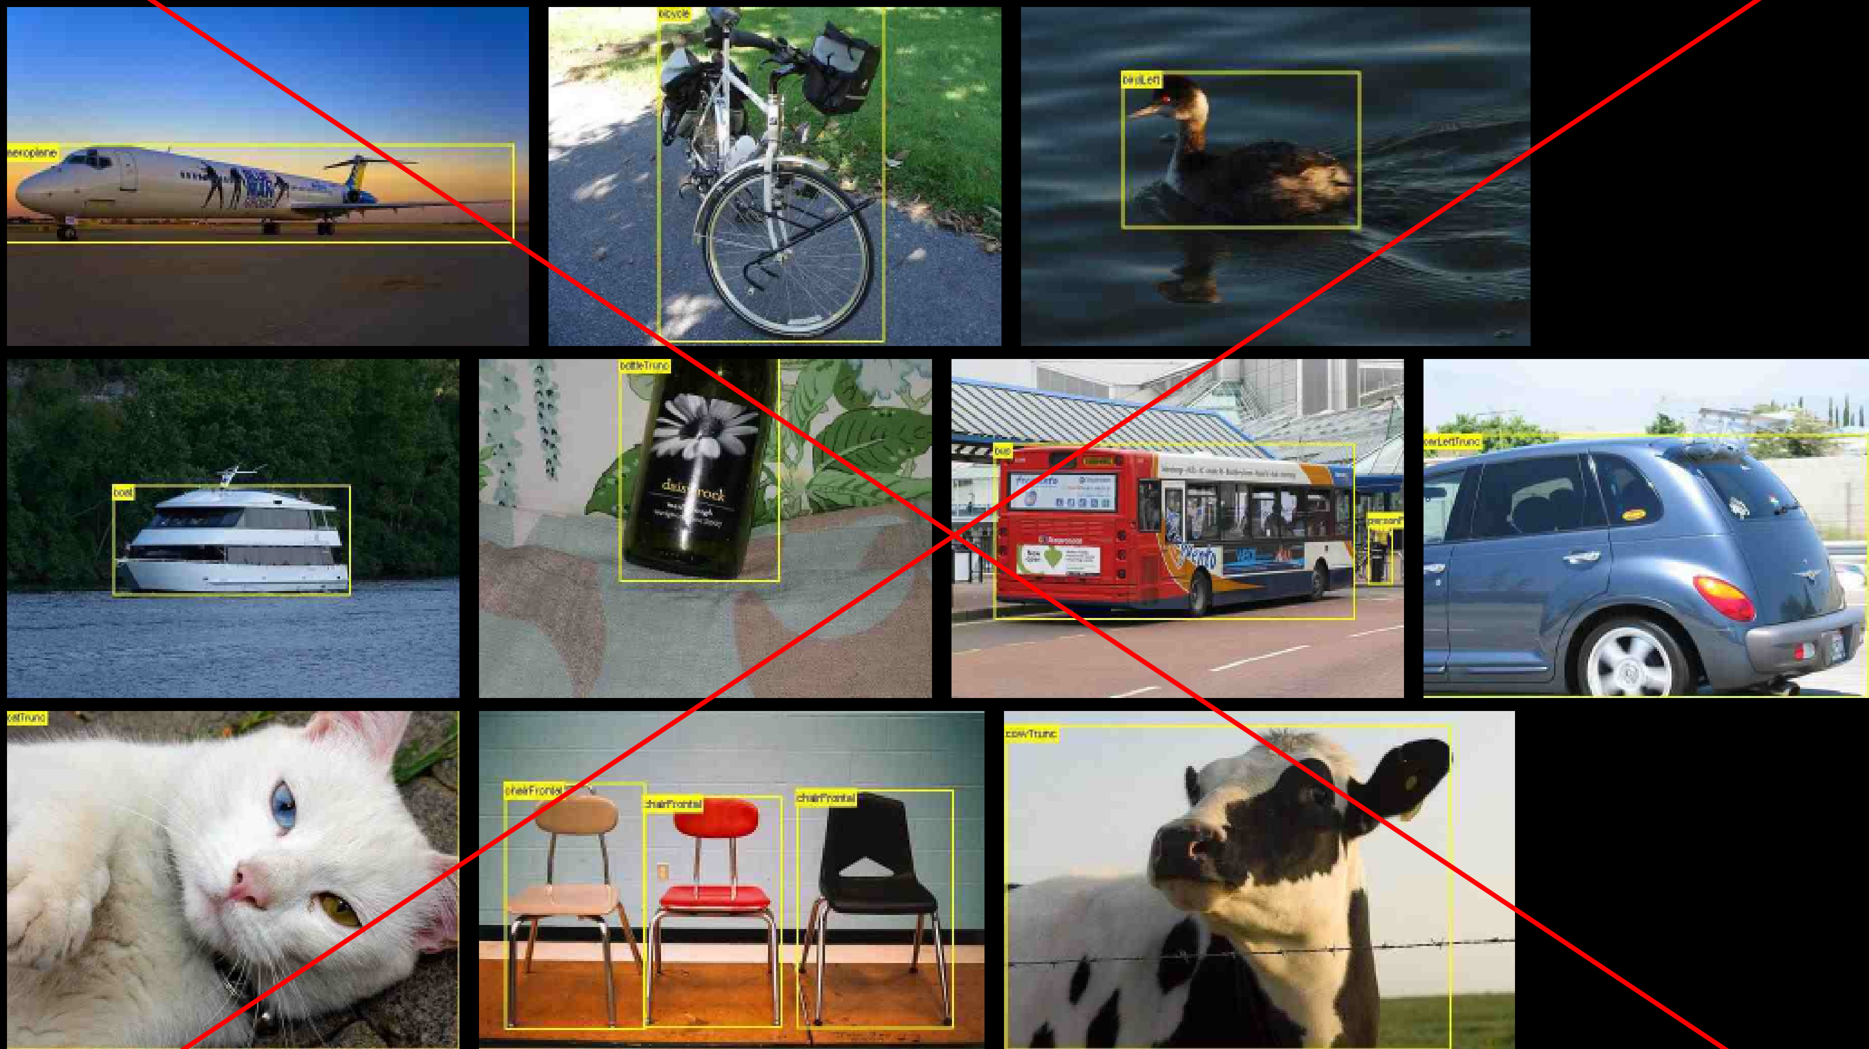

Everingham, Van Gool, Williams, Winn, Zisserman. The PASCAL Visual Object Classes Challenge 2011.

Object detection

Target detection

Traffic analysis

Eads, Rosten, Helmbold. BMVC 2009

Rosten, Loveland, Hickman. arXiv:0912.1310

Object detection: difficulties

Which ones are cars?

- Problem is unstructured

  $$\text{Image} \rightarrow \{(x_1, y_1), (x_2, y_2), \cdots\}$$

- Number of objects unknown *a priori*
- Not a fixed set of labels

What is object detection?

- Target detection:
  - Within boundary
- Tracking
  - Nearby, but with unique assignment

What is object detection?

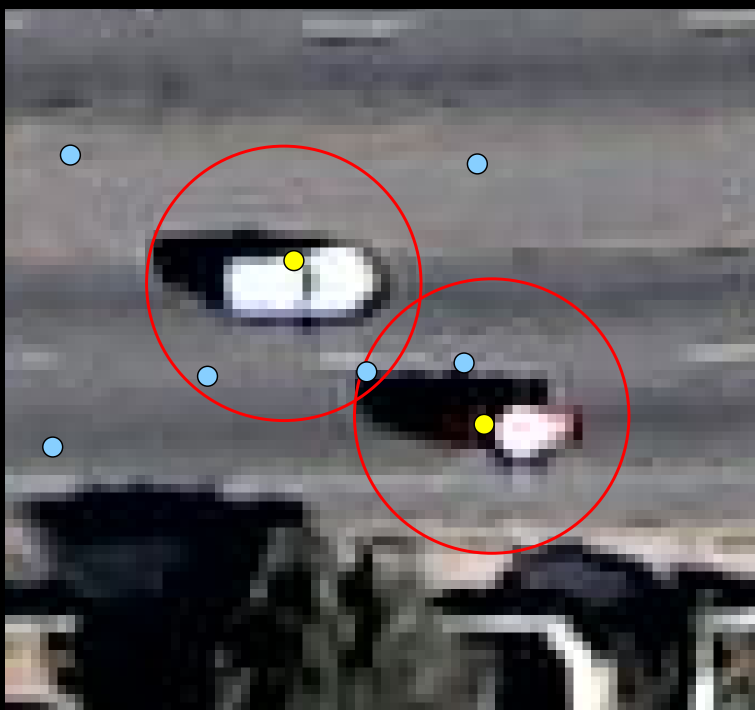

- Target detection:
  - Within boundary
- Tracking
  - Nearby, but with unique assignment

An optimized bag of hacks

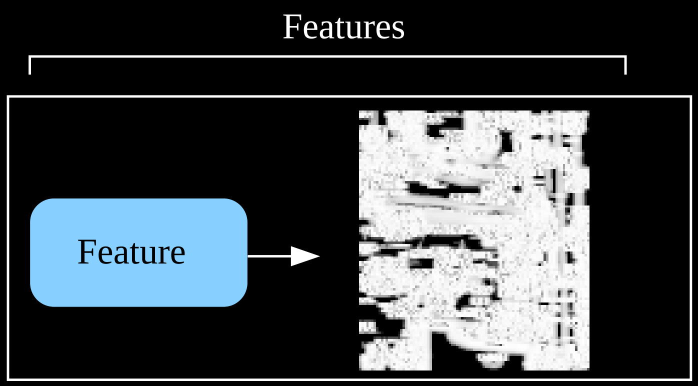

An optimized bag of hacks

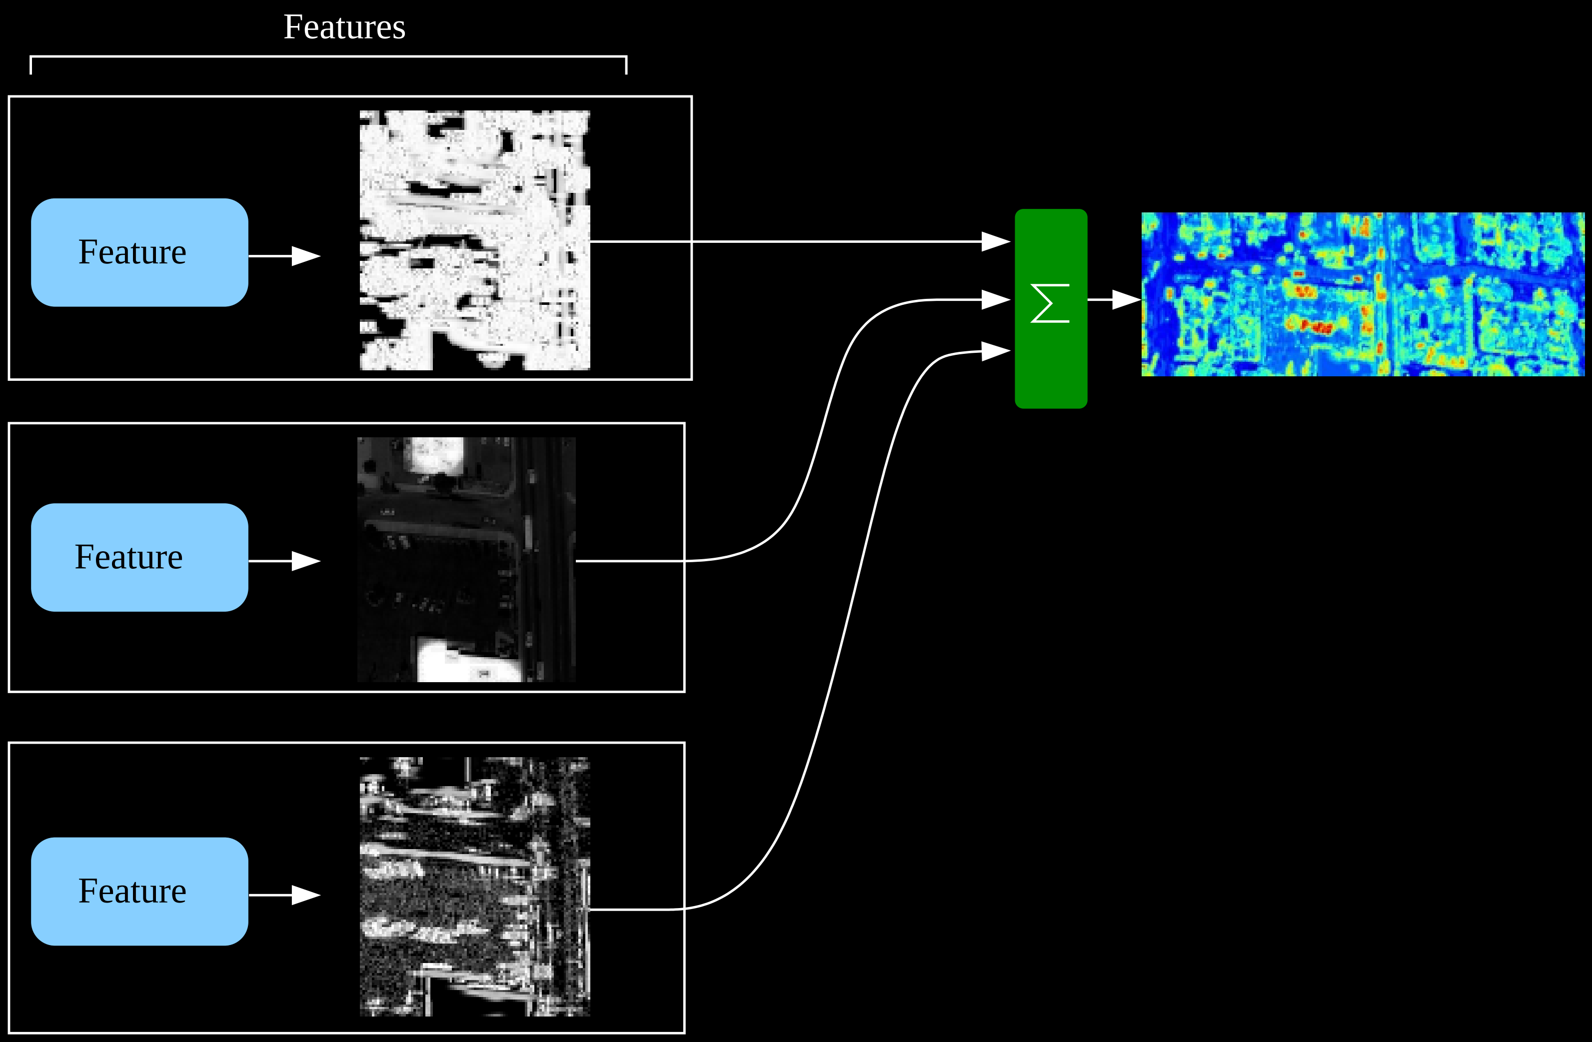

Eads, Rosten, Helmbold. BMVC 2009

An optimized bag of hacks

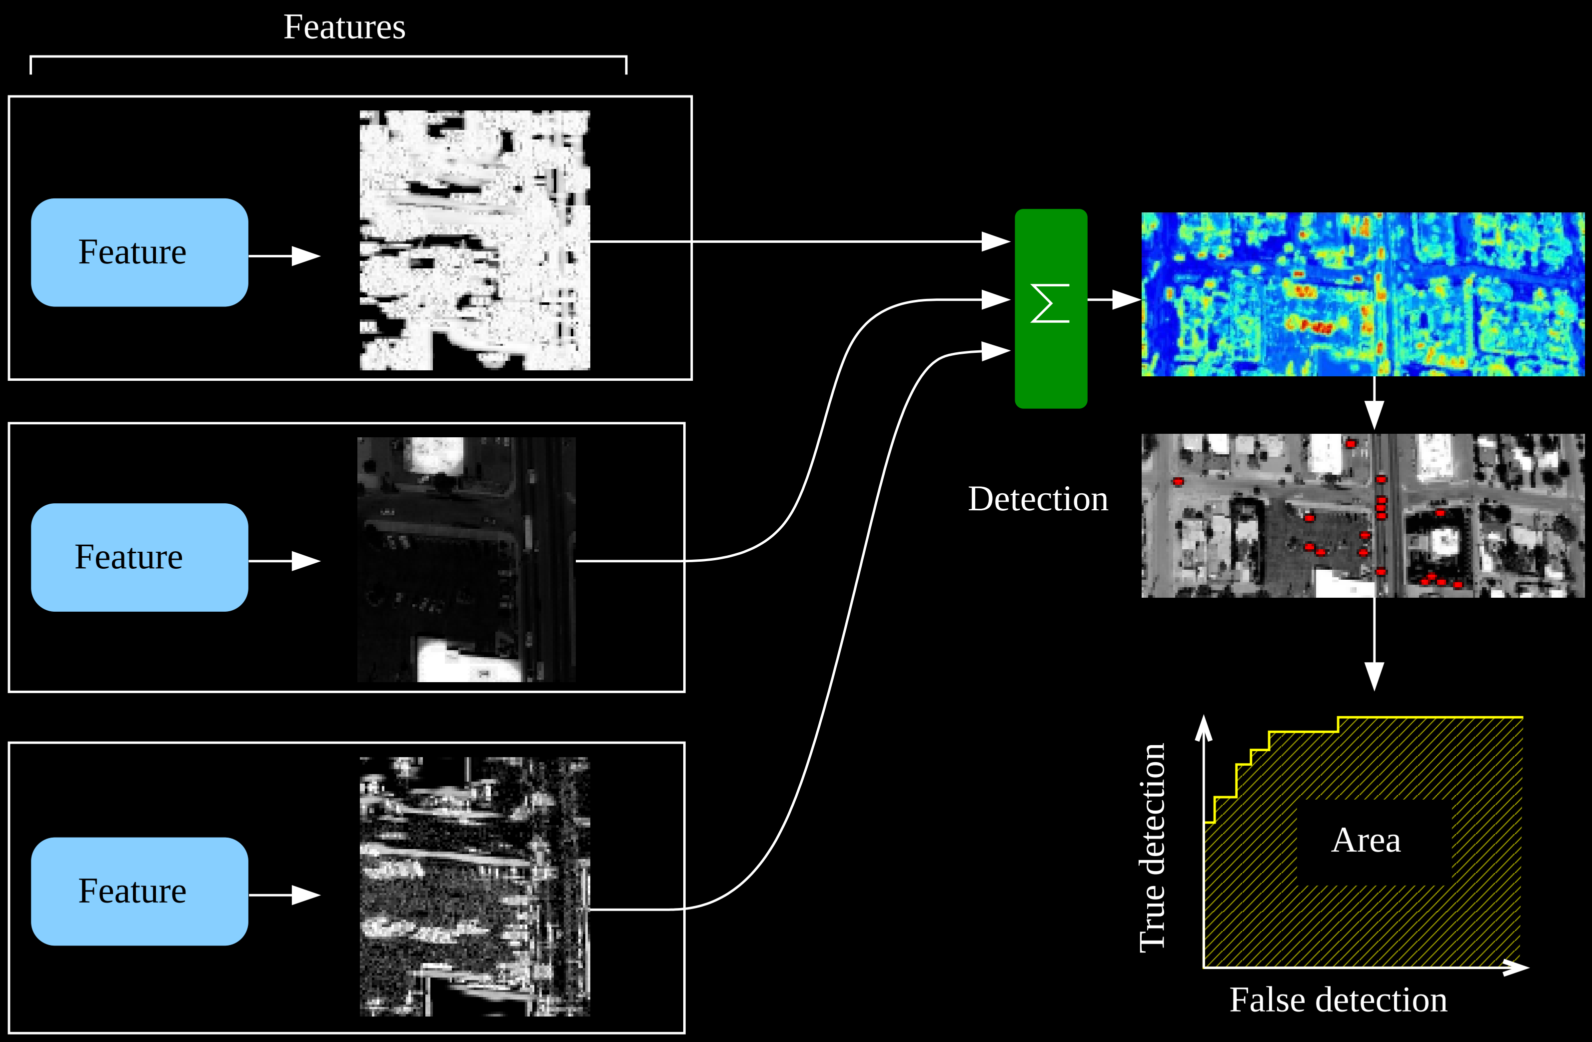

Eads, Rosten, Helmbold. BMVC 2009

An optimized bag of hacks

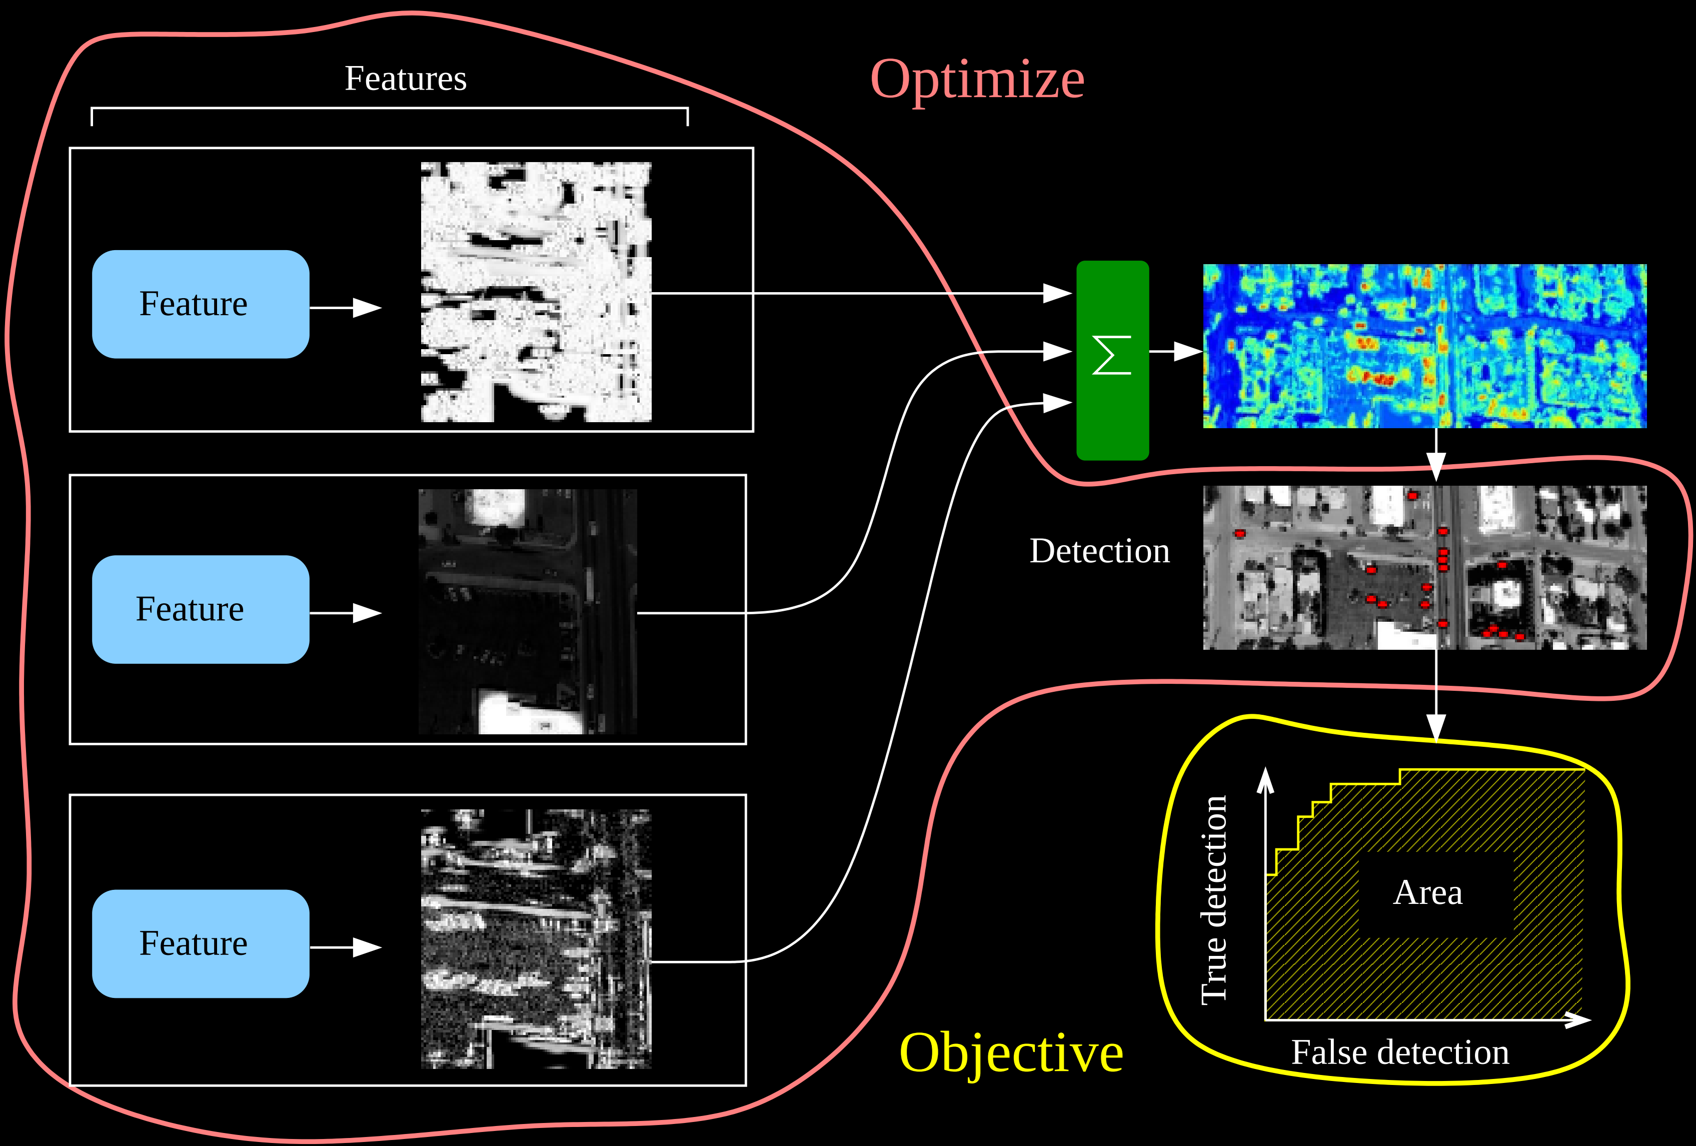

Eads, Rosten, Helmbold. BMVC 2009

Spatial boosting

Training:

- Markup: Image, and list of (x, y) pairs.
- Thresholdable weak detectors

Runtime:

- Image $\rightarrow$ list of (x, y) pairs.

Goal:

- Find weighted sum of weak detectors

Cheat:

- Actual detection stage not included

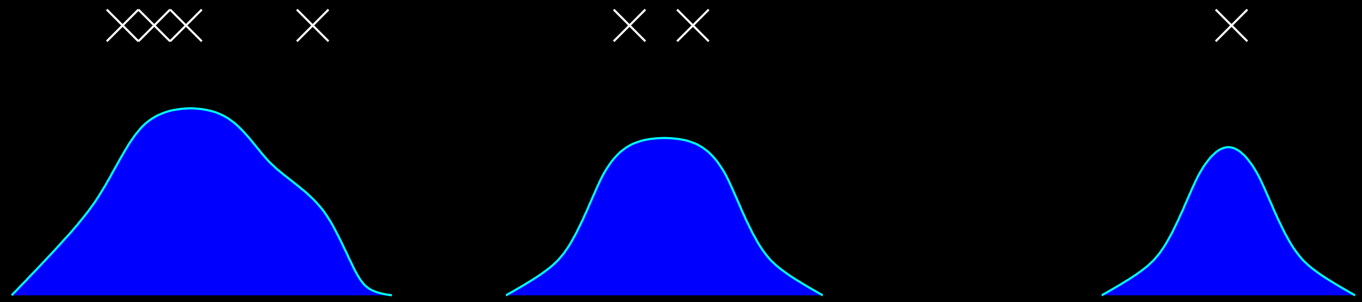

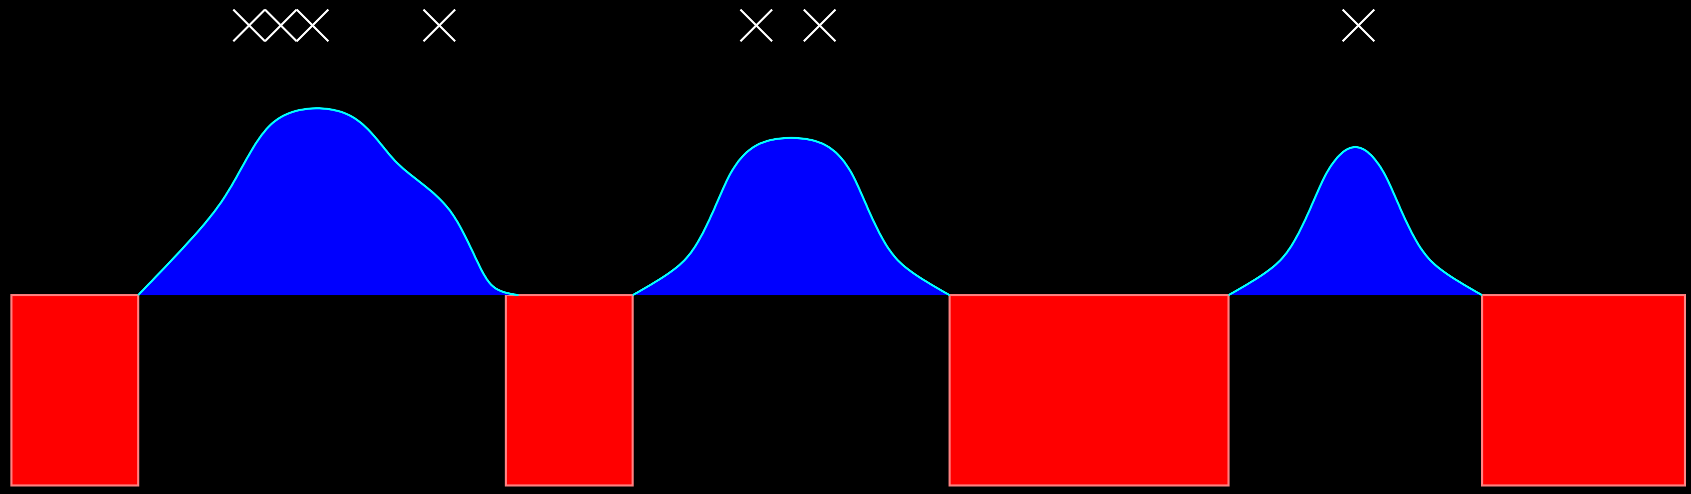

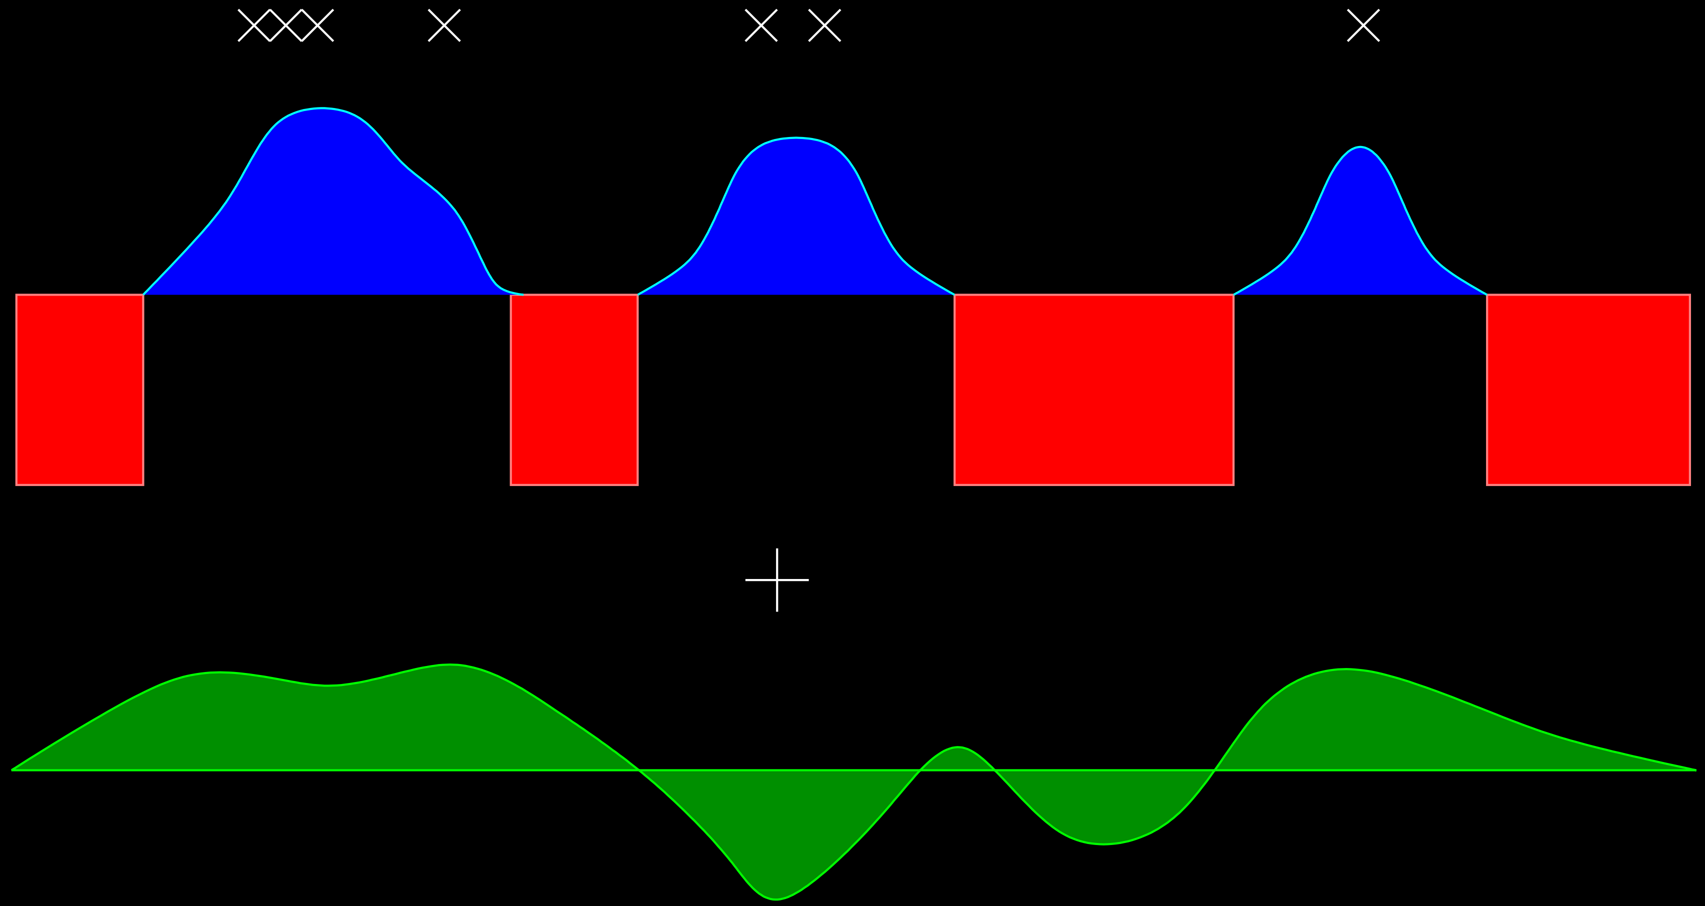

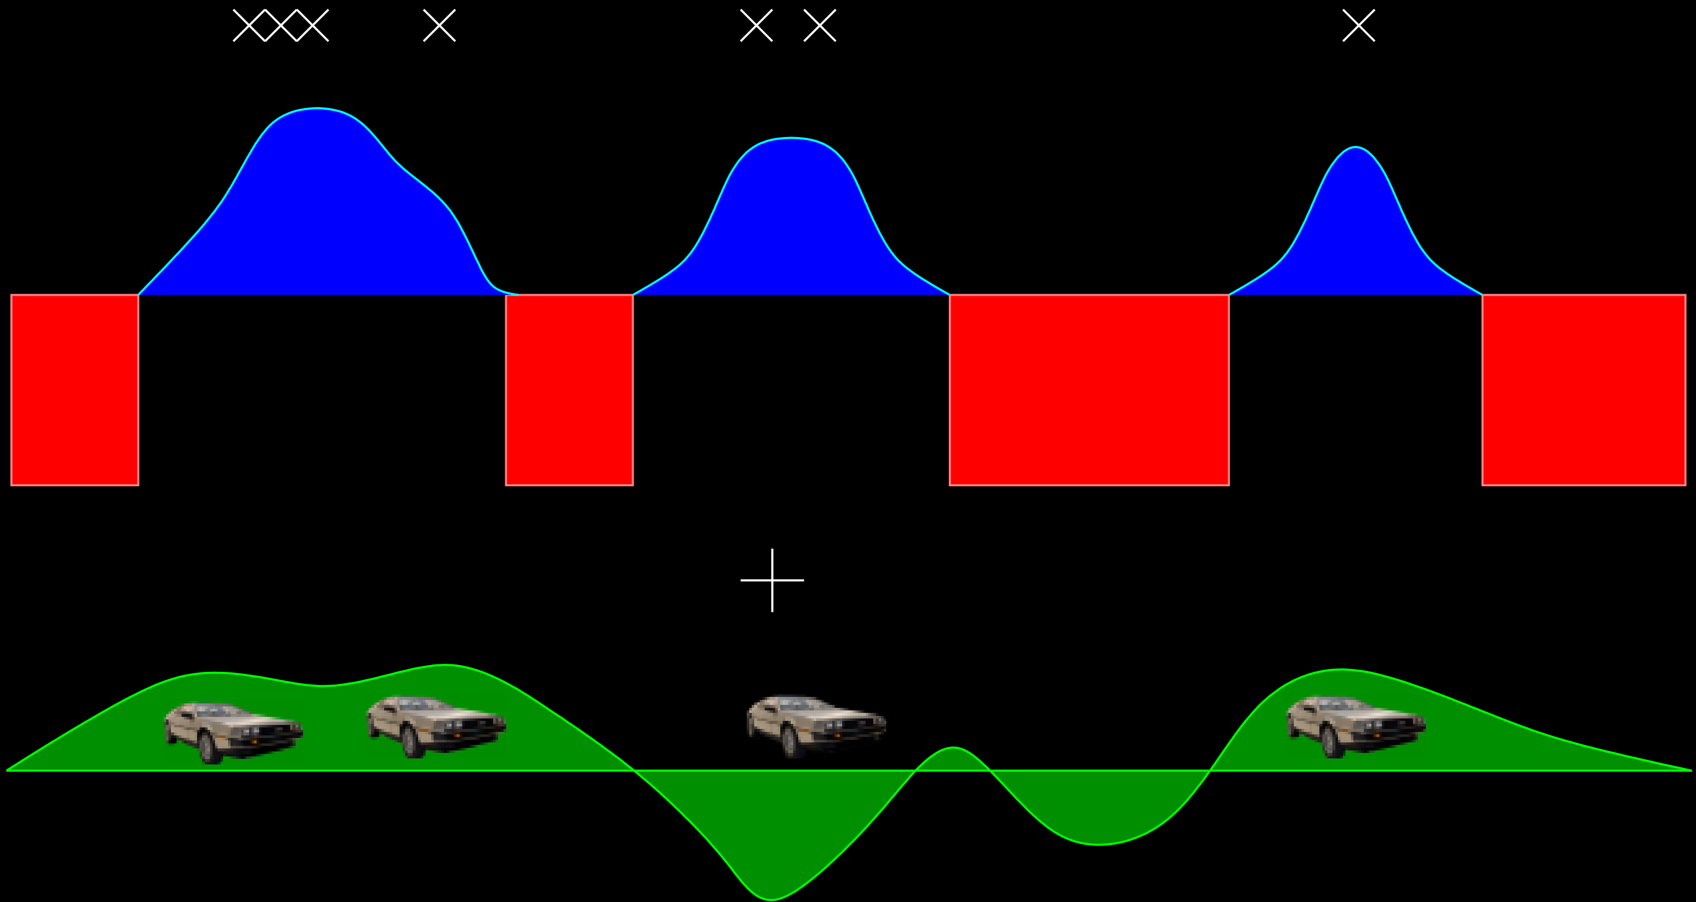

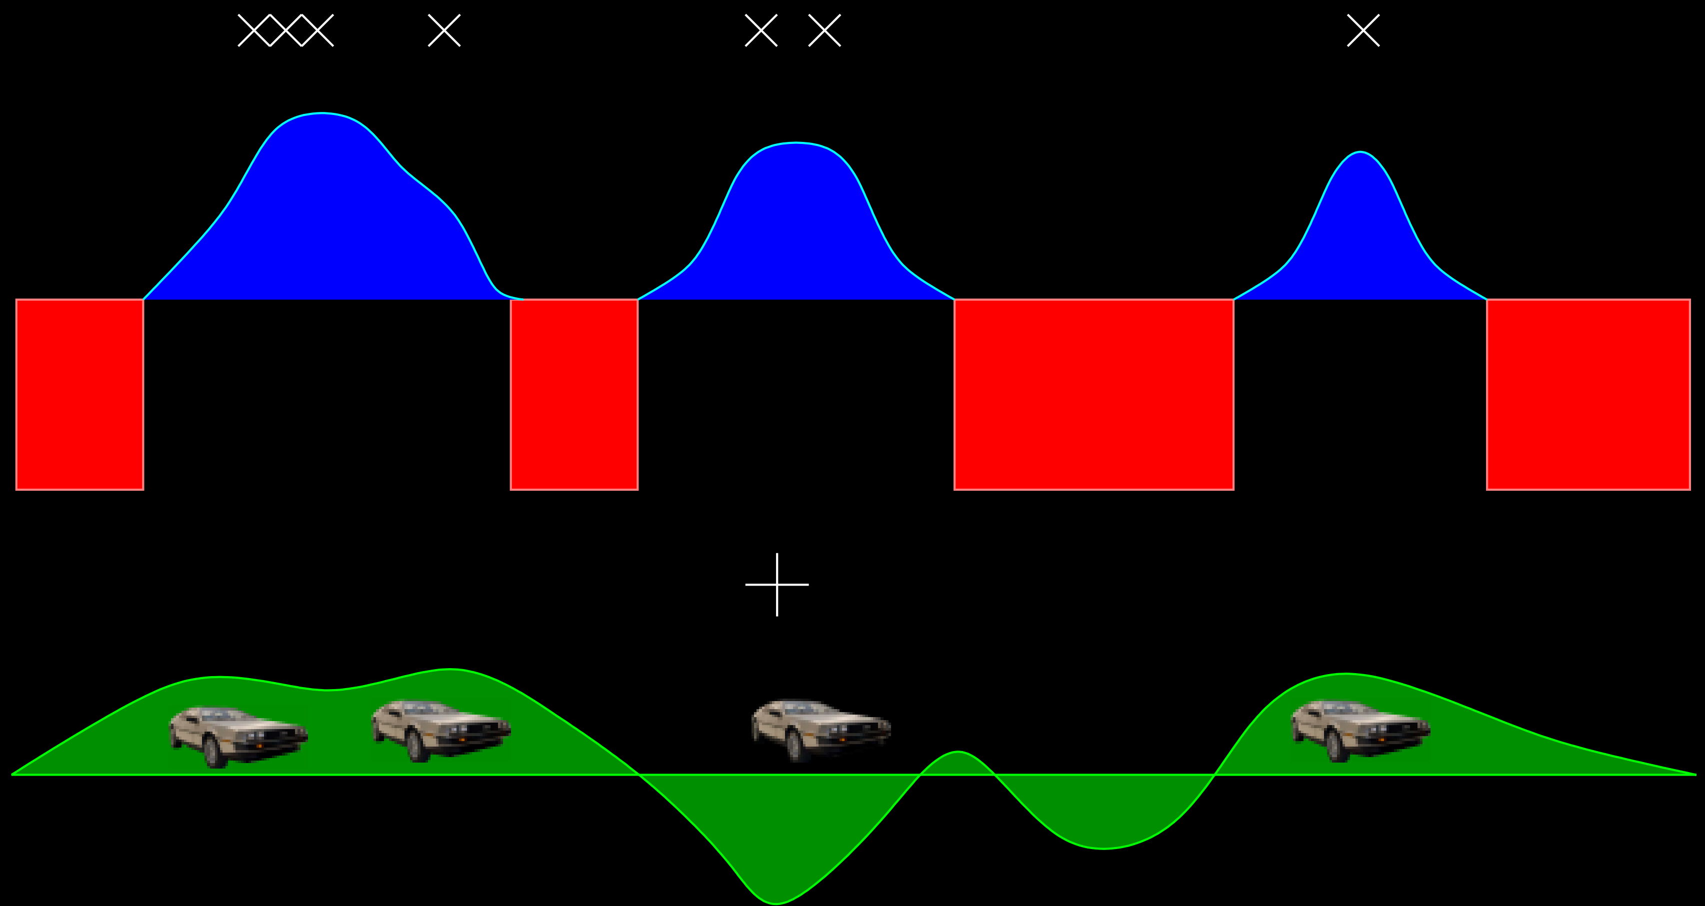

Incrementally (near) optimal

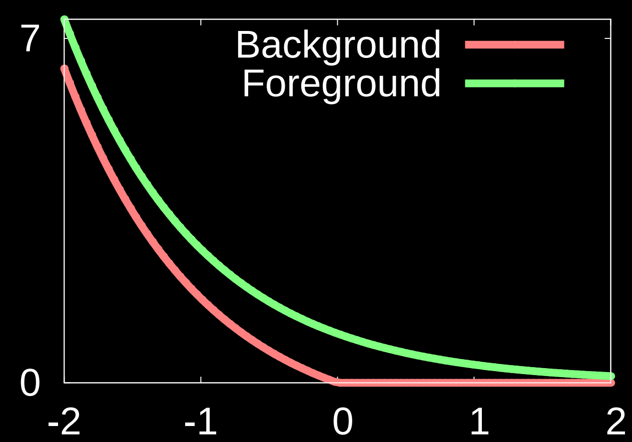

Does it work?

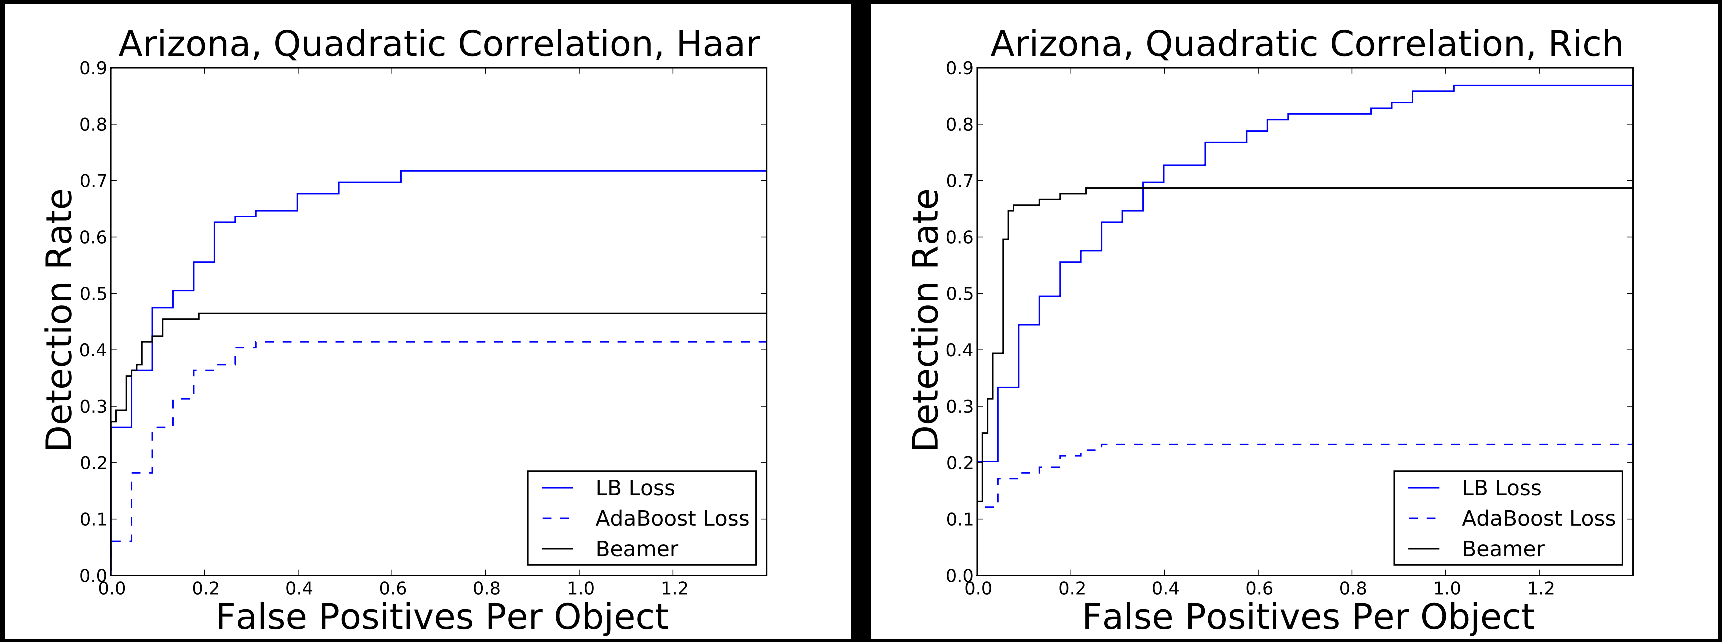

- Treating foreground and background differently makes a big difference.
- No ad-hoc data weighting.

Eads, Rosten, Helmbold. arXiv:1309.1080 2013

Fluorescence microscopy in one slide

Fluorescence microscopy in one slide

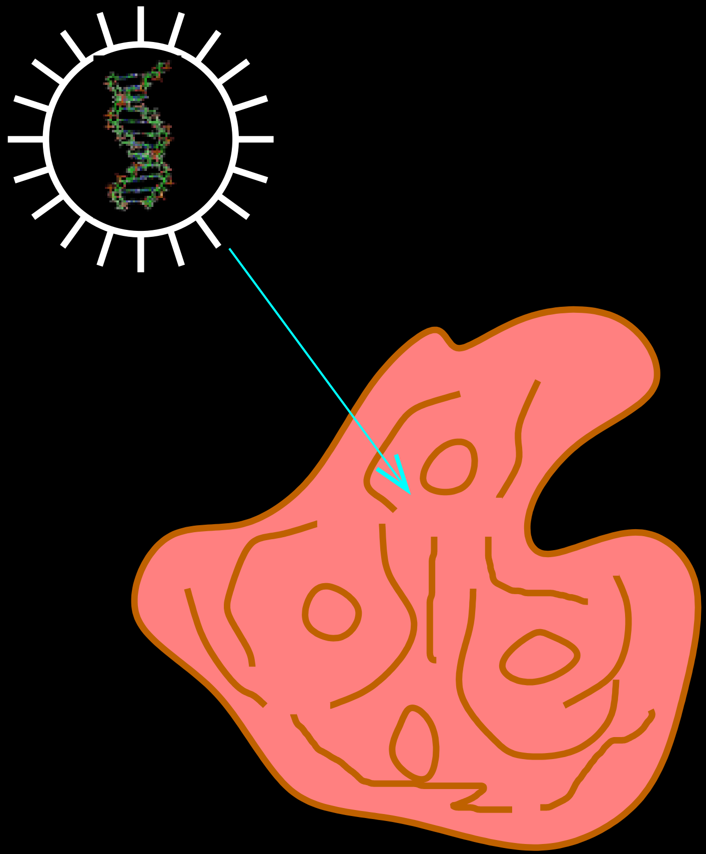

Fluorescence microscopy in one slide

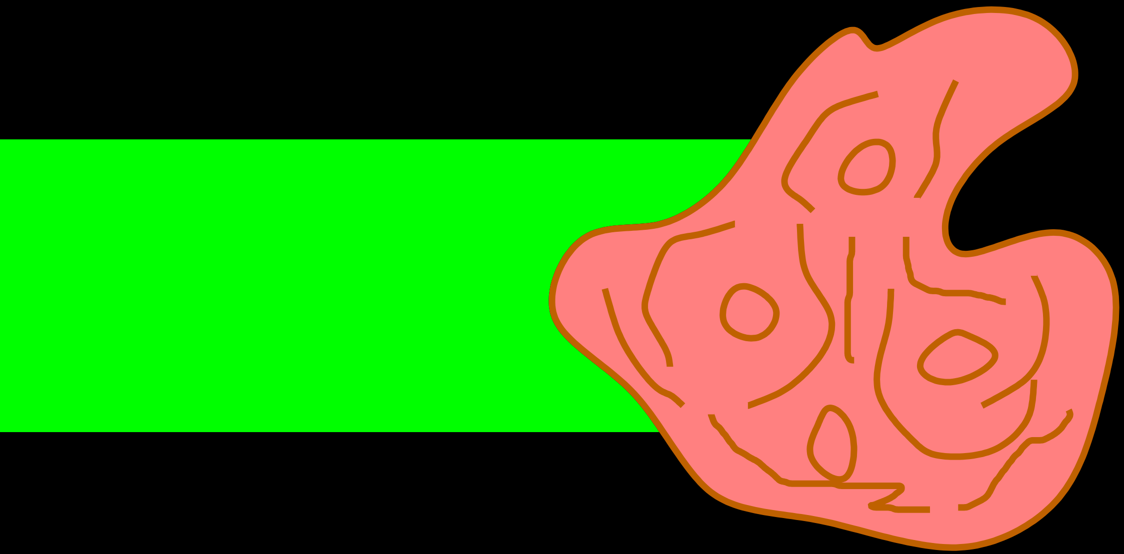

Fluorescence microscopy in one slide

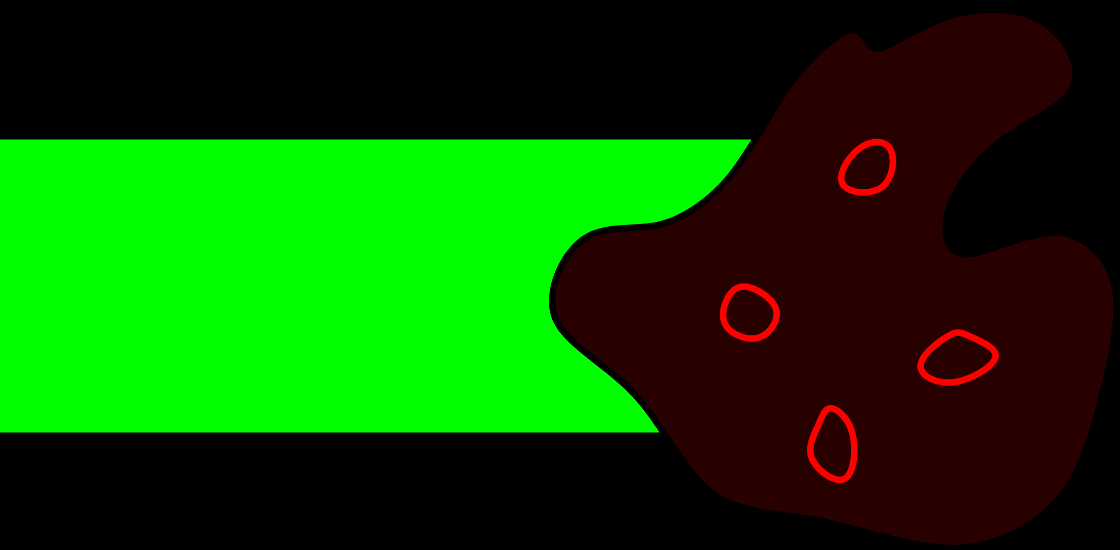

Fluorescence microscopy in one slide

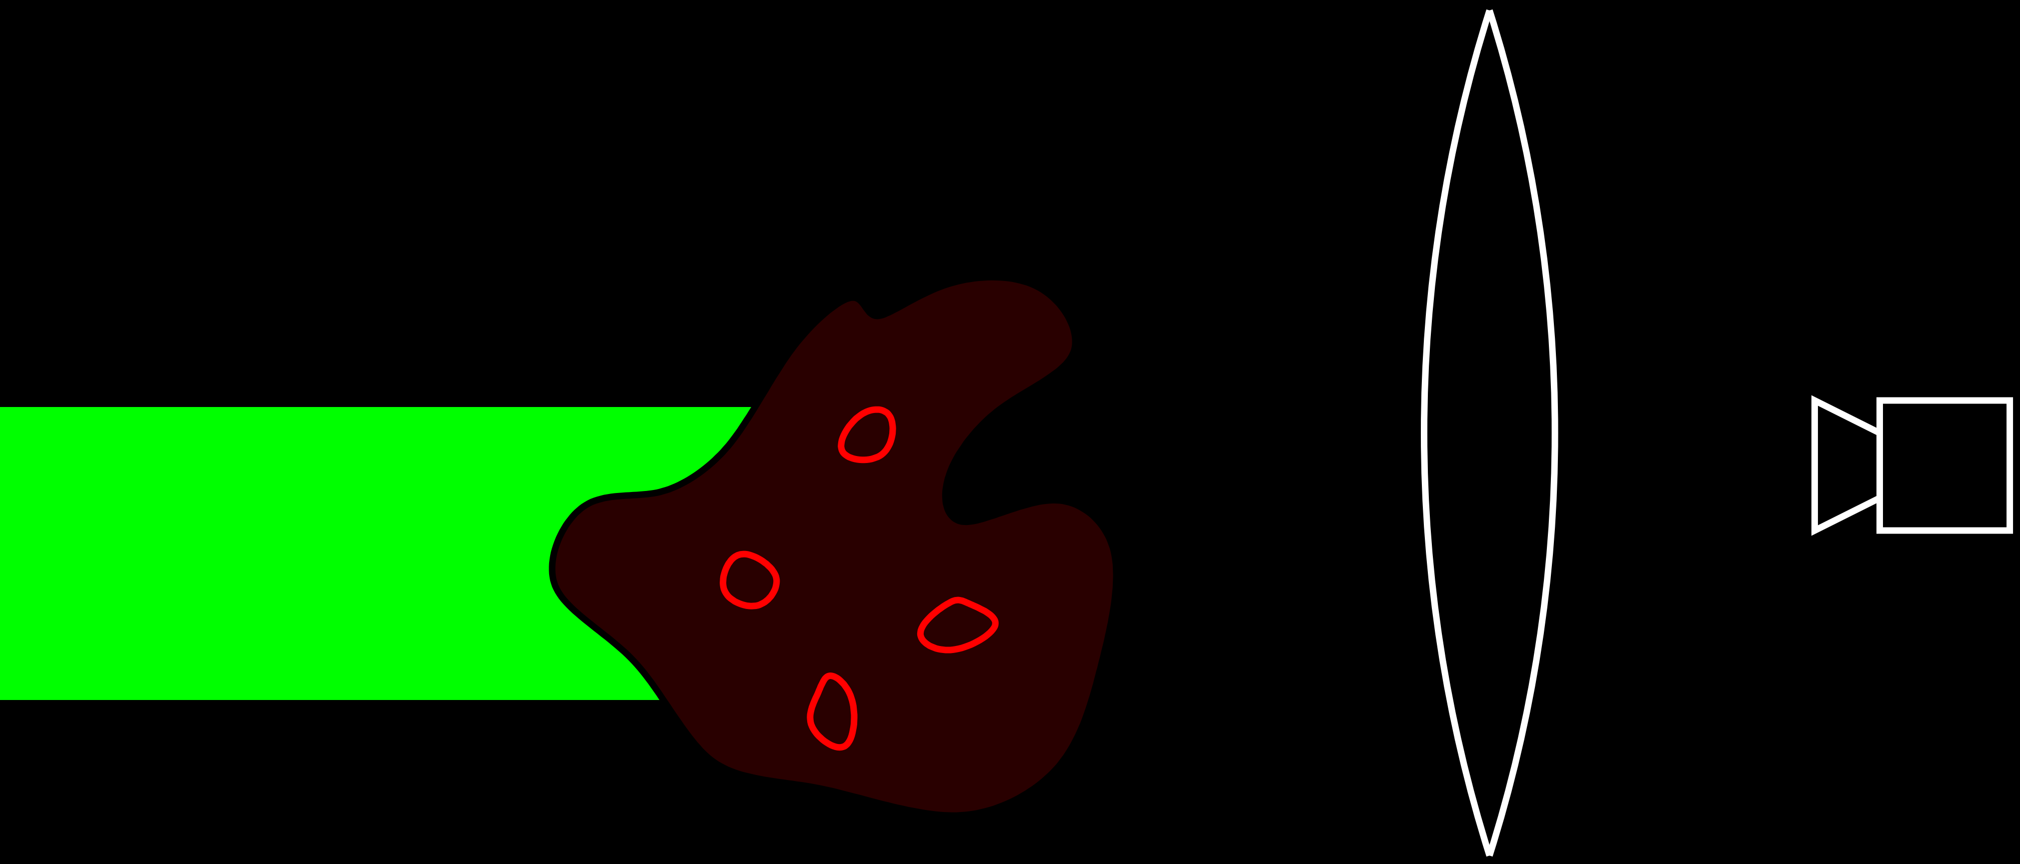

Abbé Limit: Wavelength limits resolution (250 nm)

Fluorescence microscopy in one slide

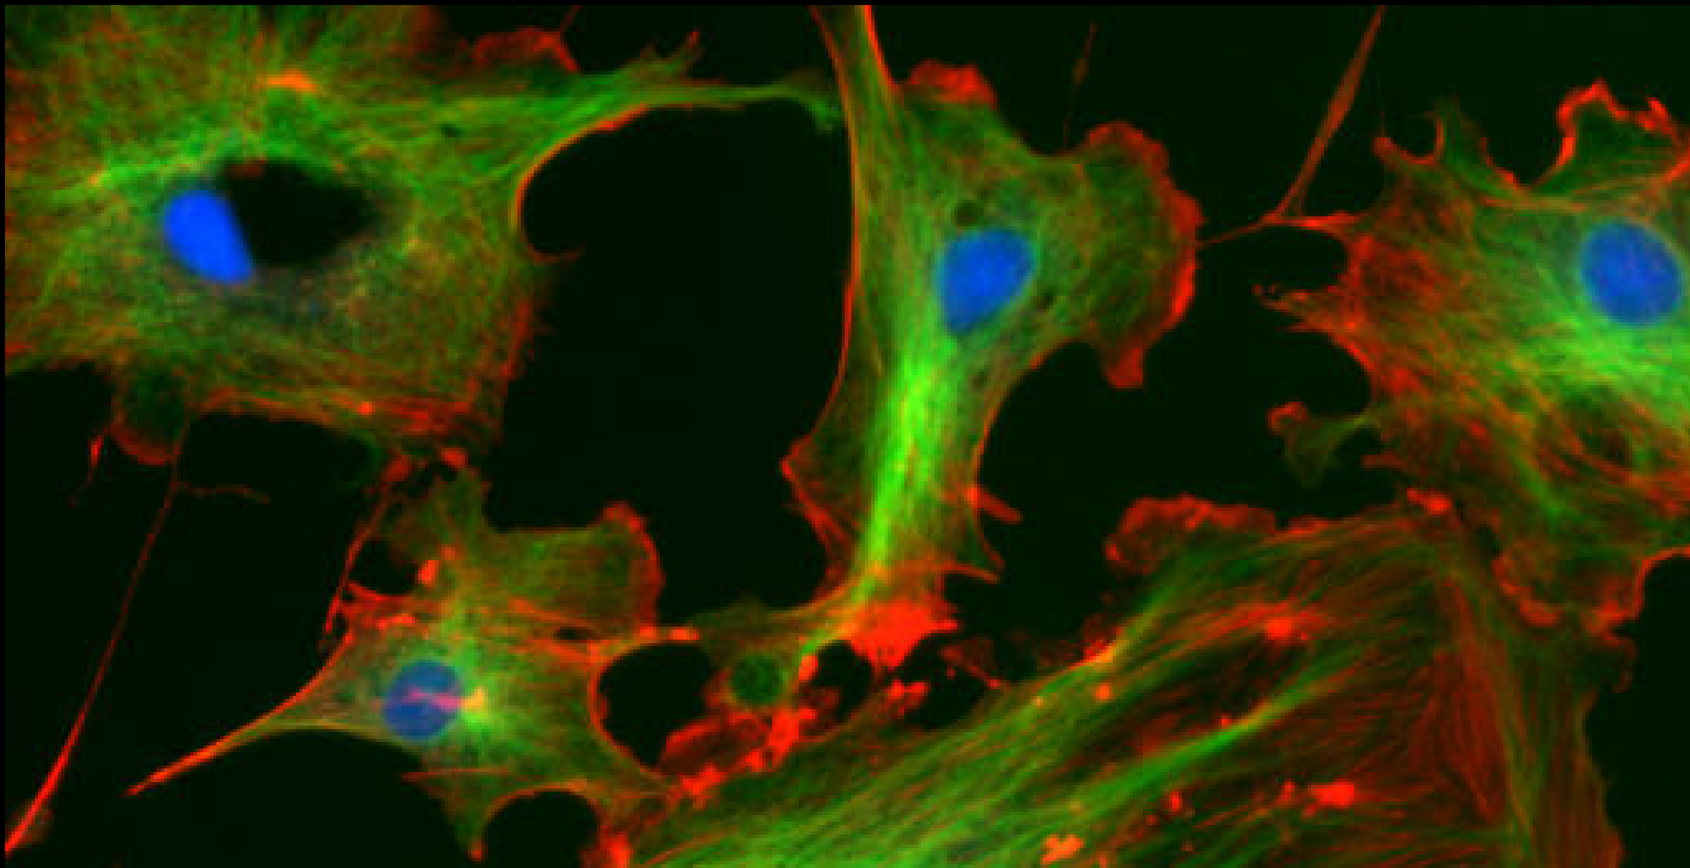

J Benutzer, Wikipedia, 2005

Abbé Limit: Wavelength limits resolution (250 nm)

Localisation microscopy

- Separate fluorophore emission in time
- Localise individual spots

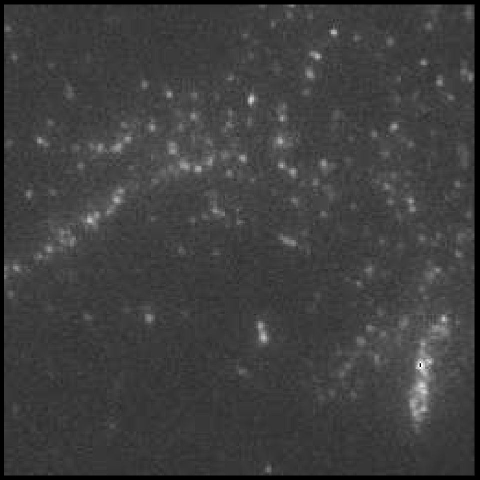

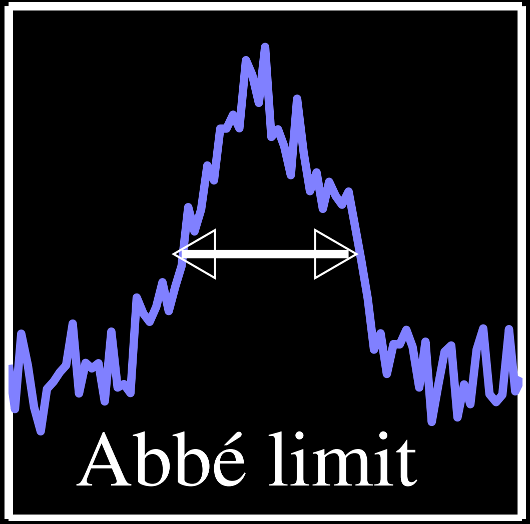

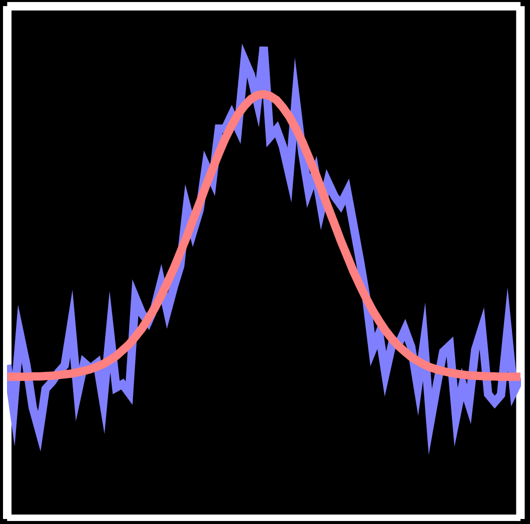

- Frame by frame: threshold and fit Gaussians

Betzig, Patterson, Sougrat, Lindwasser, Olenych, Bonifacino, Davidson, Lippincott-Schwartz, Hess. Science 2006

Localisation microscopy

Frames: 4

Localisation microscopy

Frames: 16

Localisation microscopy

Frames: 64

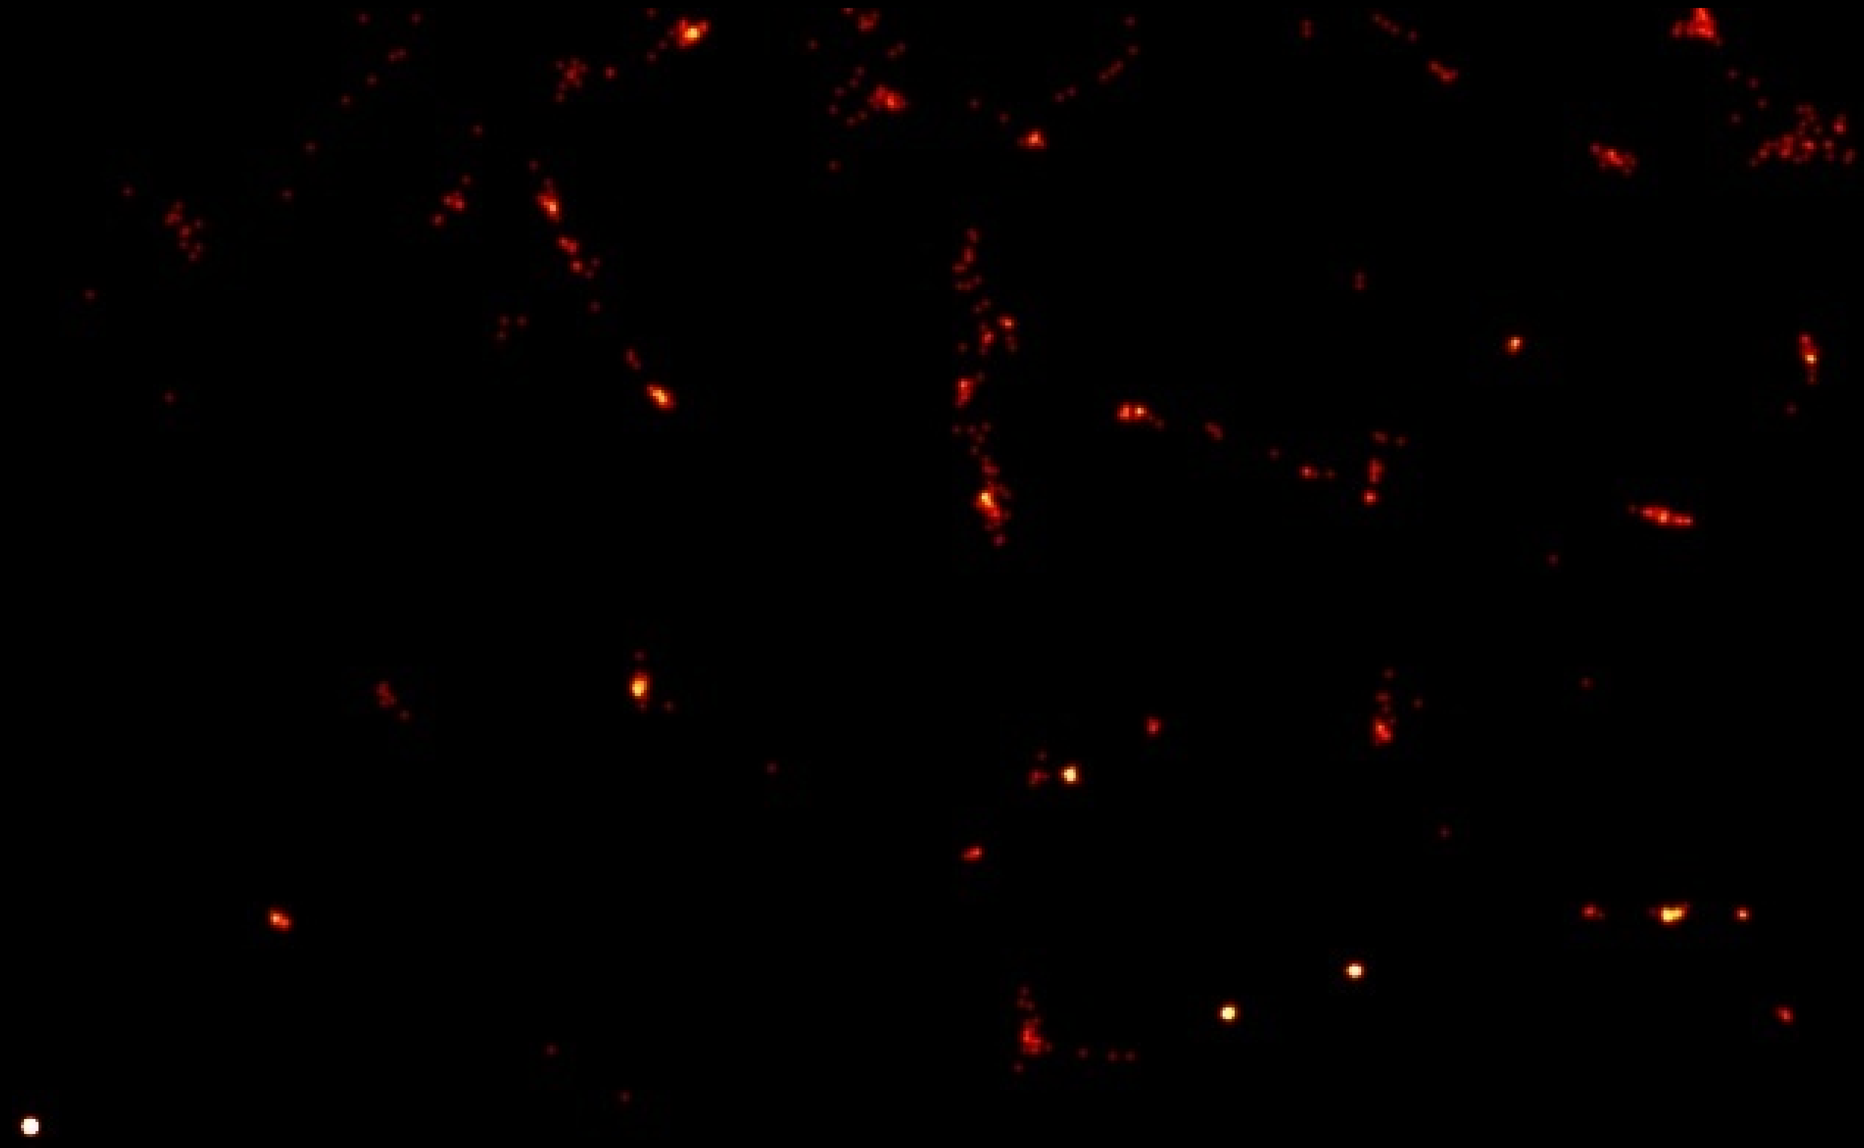

Localisation microscopy

Frames: 256

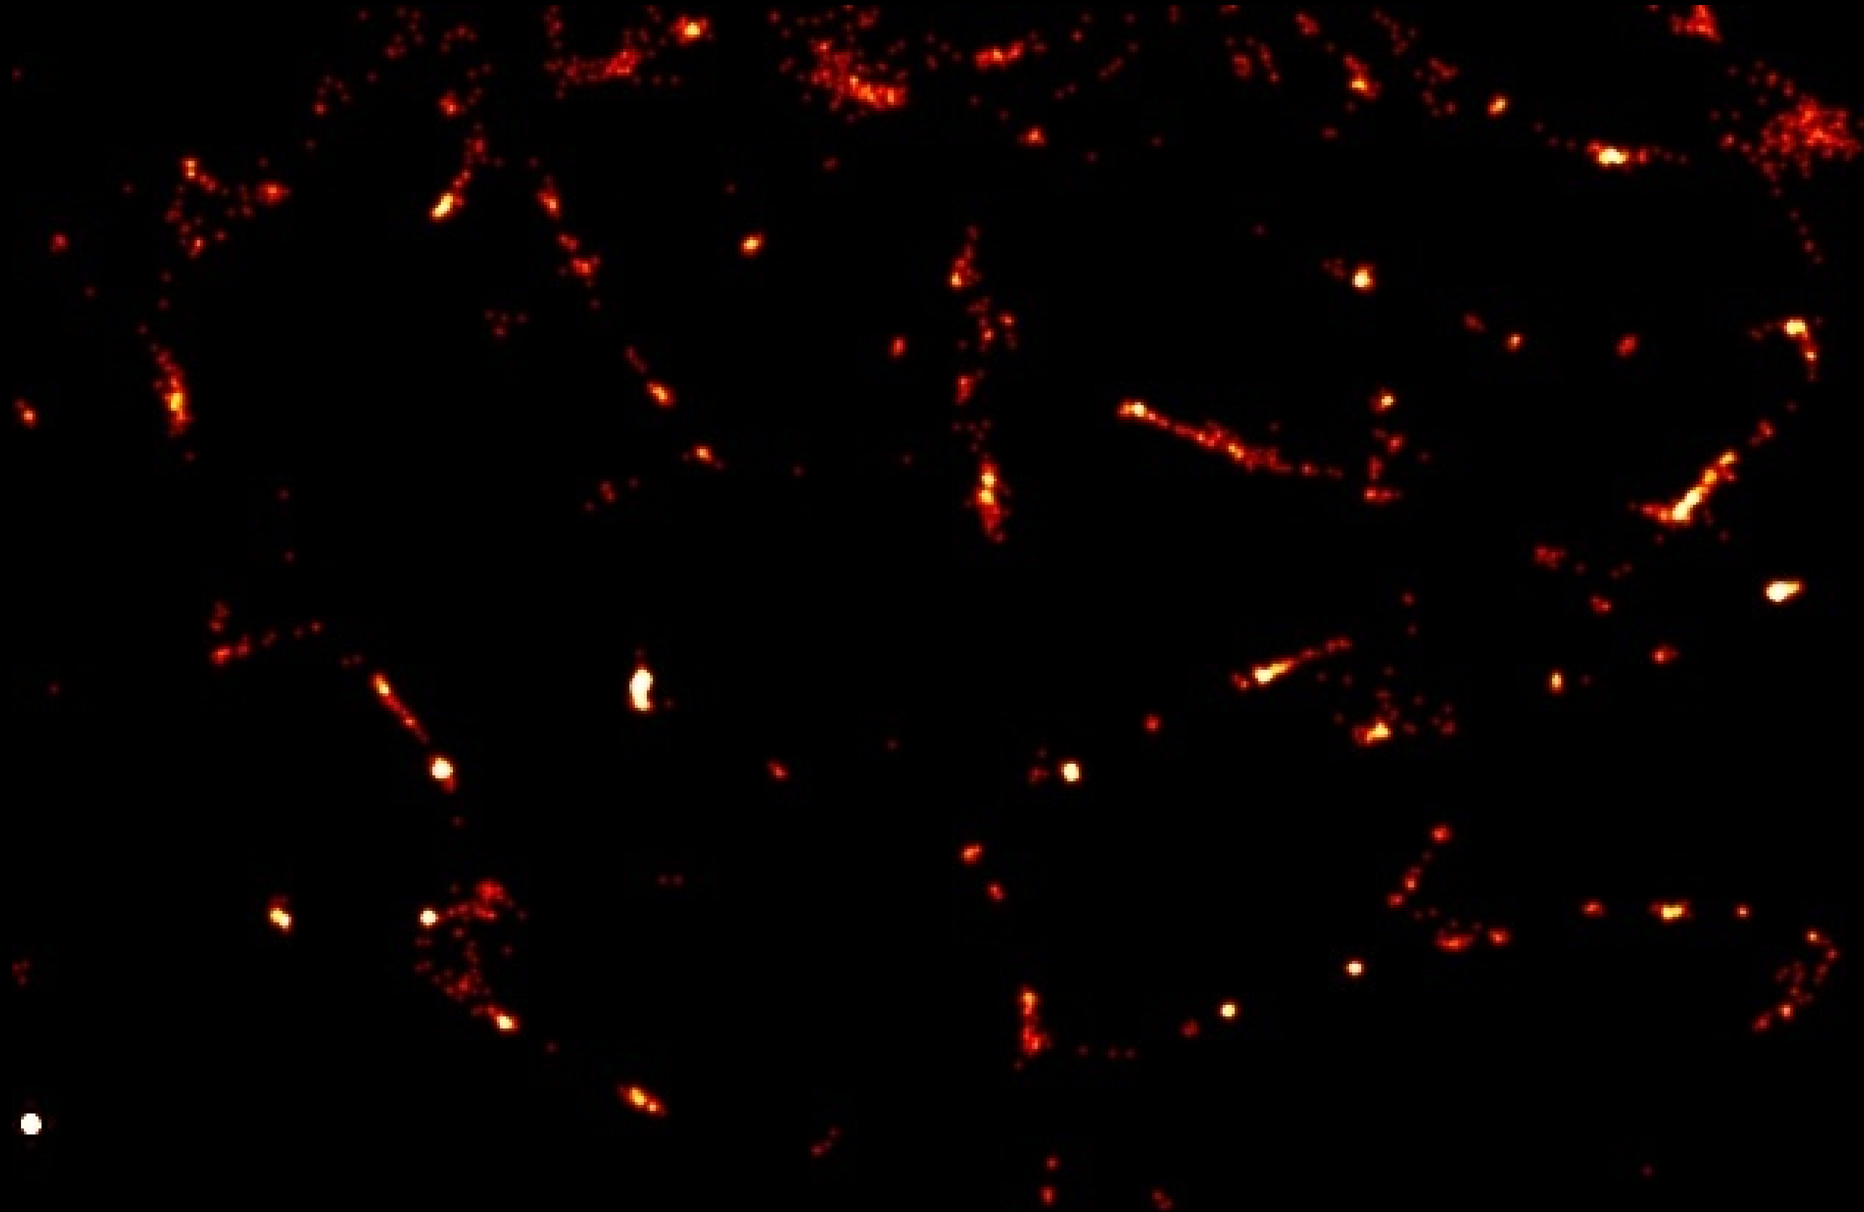

Localisation microscopy

Frames: 1024

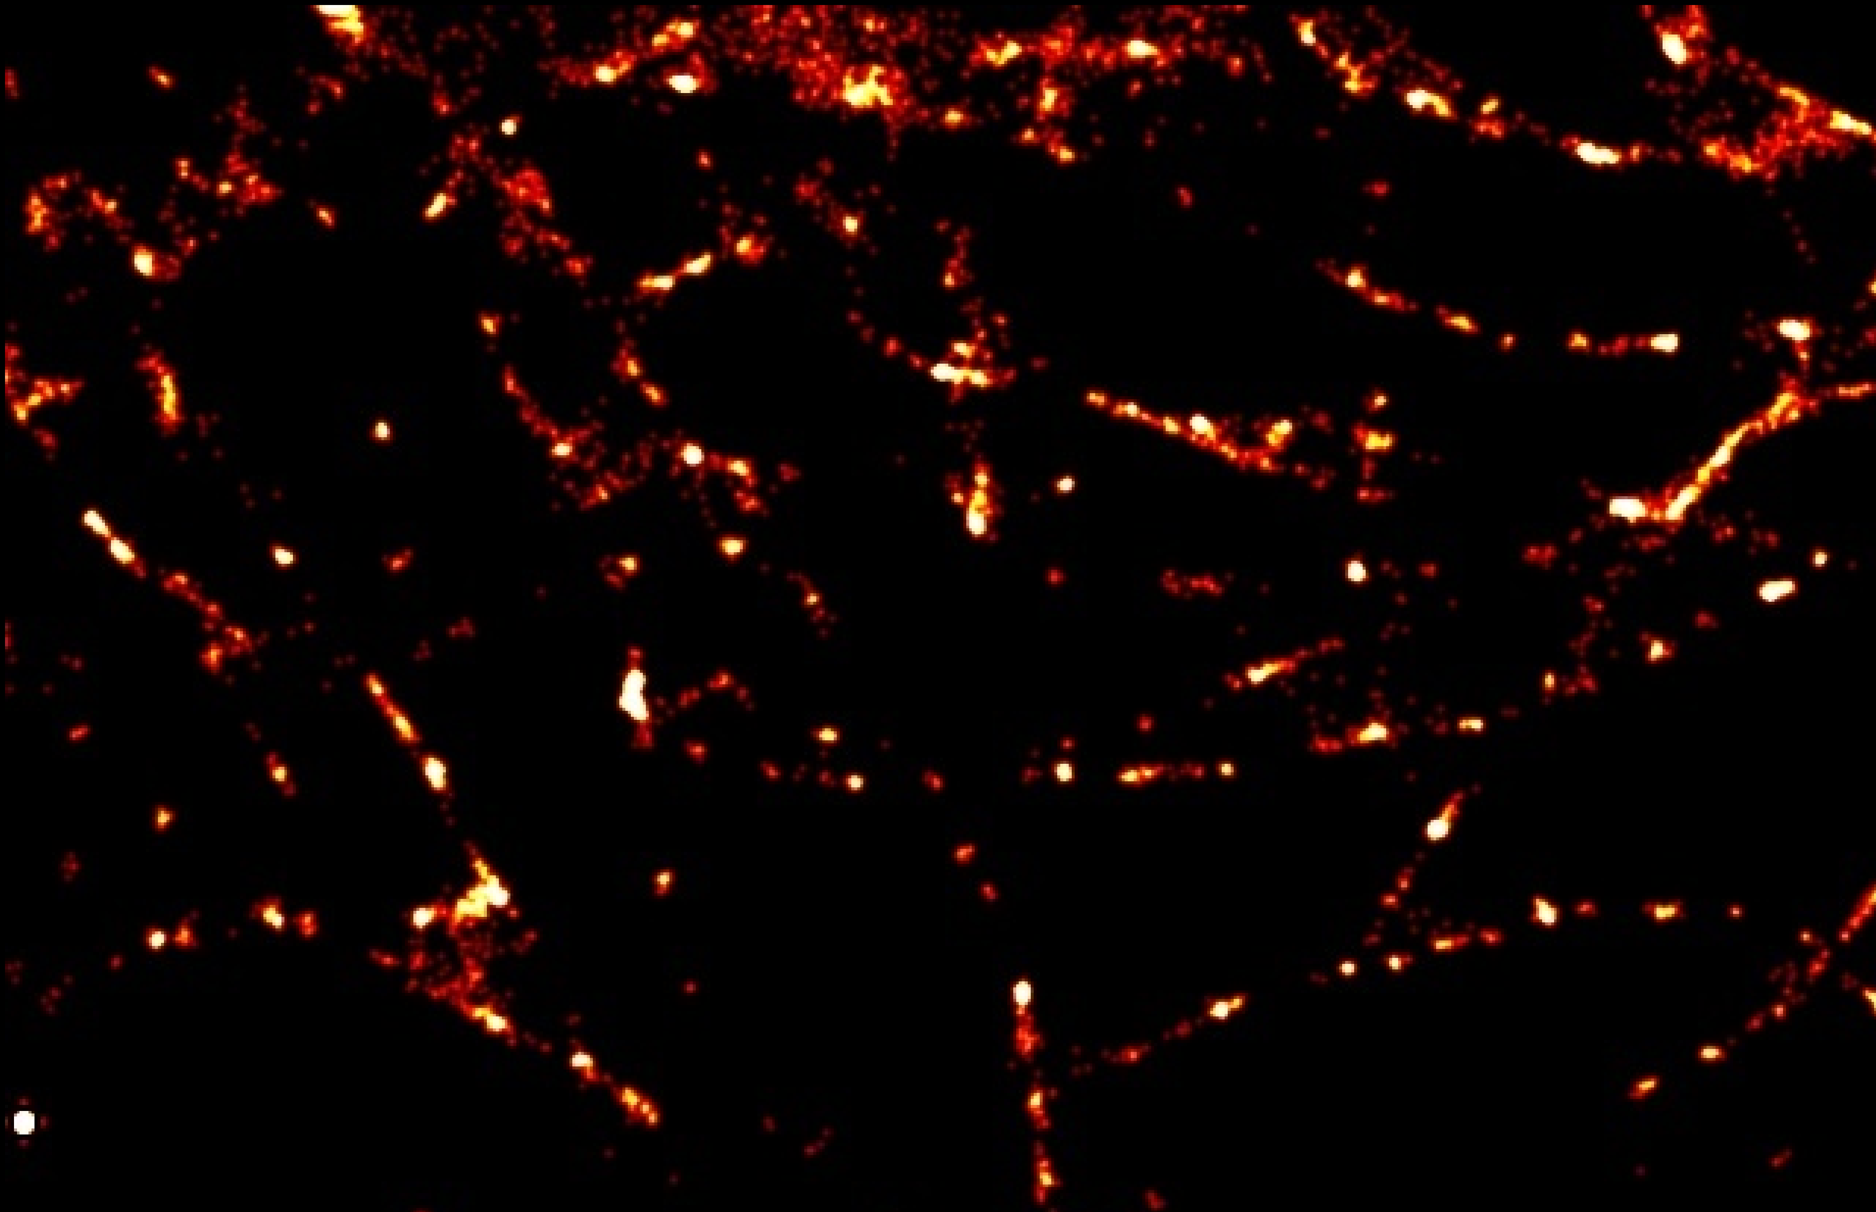

Localisation microscopy

Frames: 4096

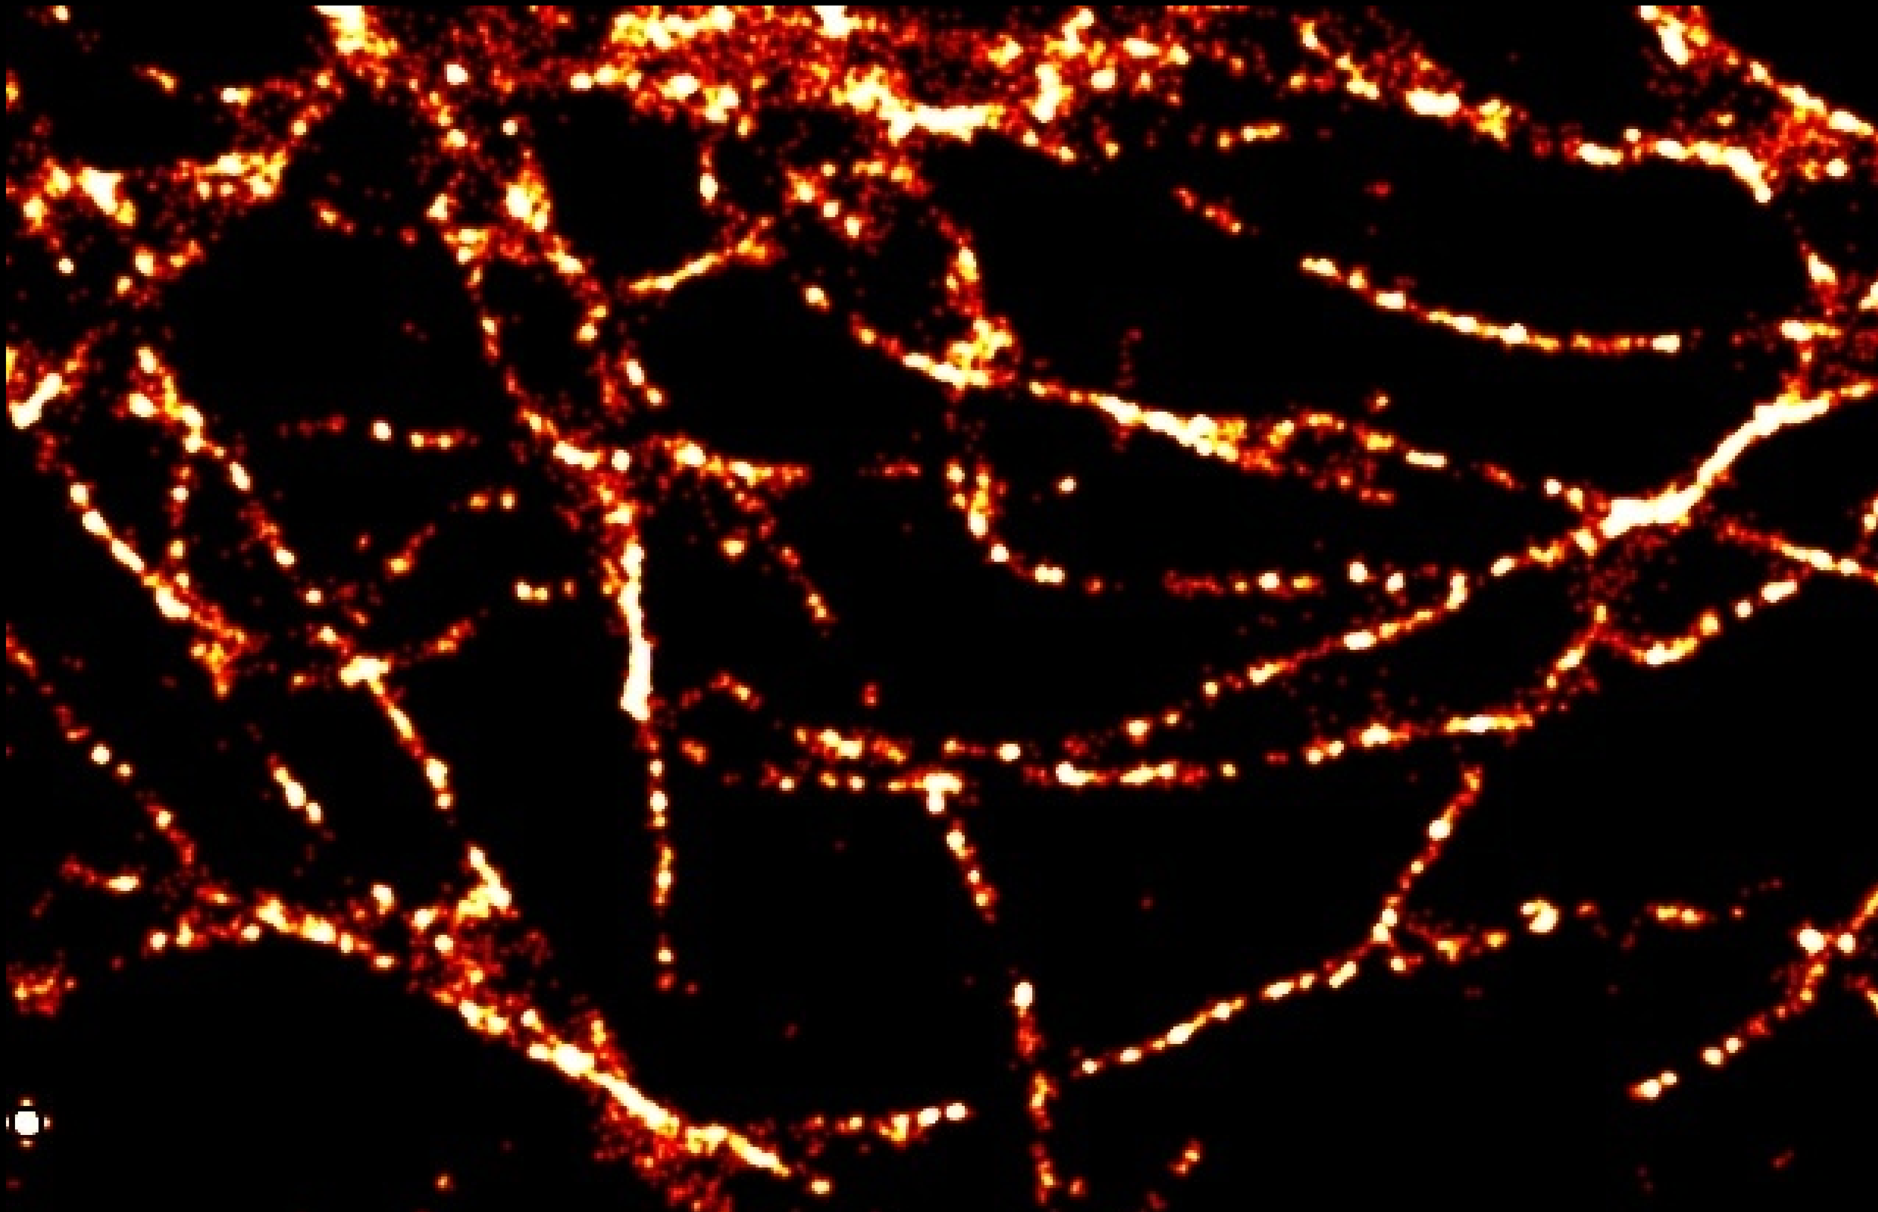

More realistic data

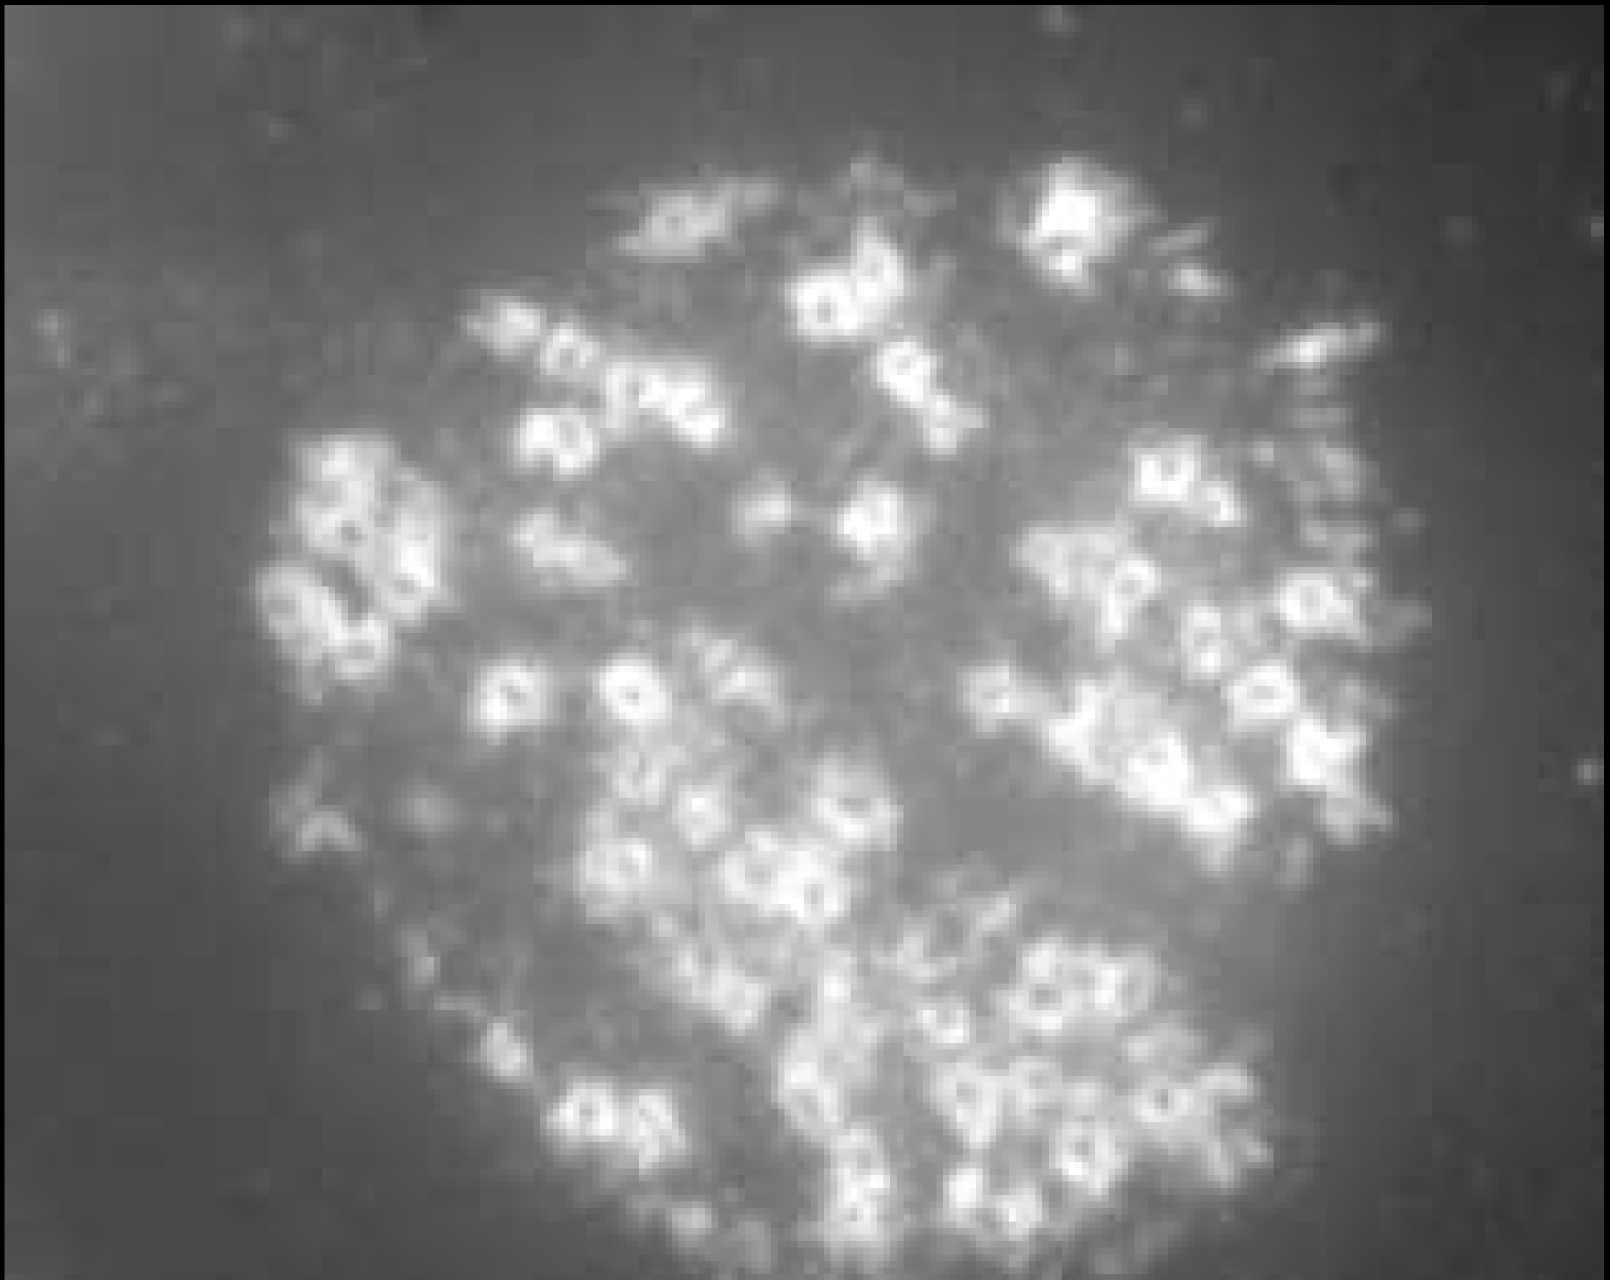

Bayesian analysis of Blinking and Bleaching (3B)

Model entire data set in one go:

- Statistical temporal model of fluorophores
- Resolution — improved by multiple fluorophore reappearances
- Acquisition time — use data from overlapping fluorophores: far fewer images needed
- Fluorophore choice — suitable for answering biologically useful questions: improved data analysis needed to cope with worse data.

Cox, Rosten, Monypenny, Jovanovic-Talisman, Burnette, Lippincott-Schwartz, Jones, Heintzmann. Nature Methods 2012.

Bayesian analysis

- Generative model for images

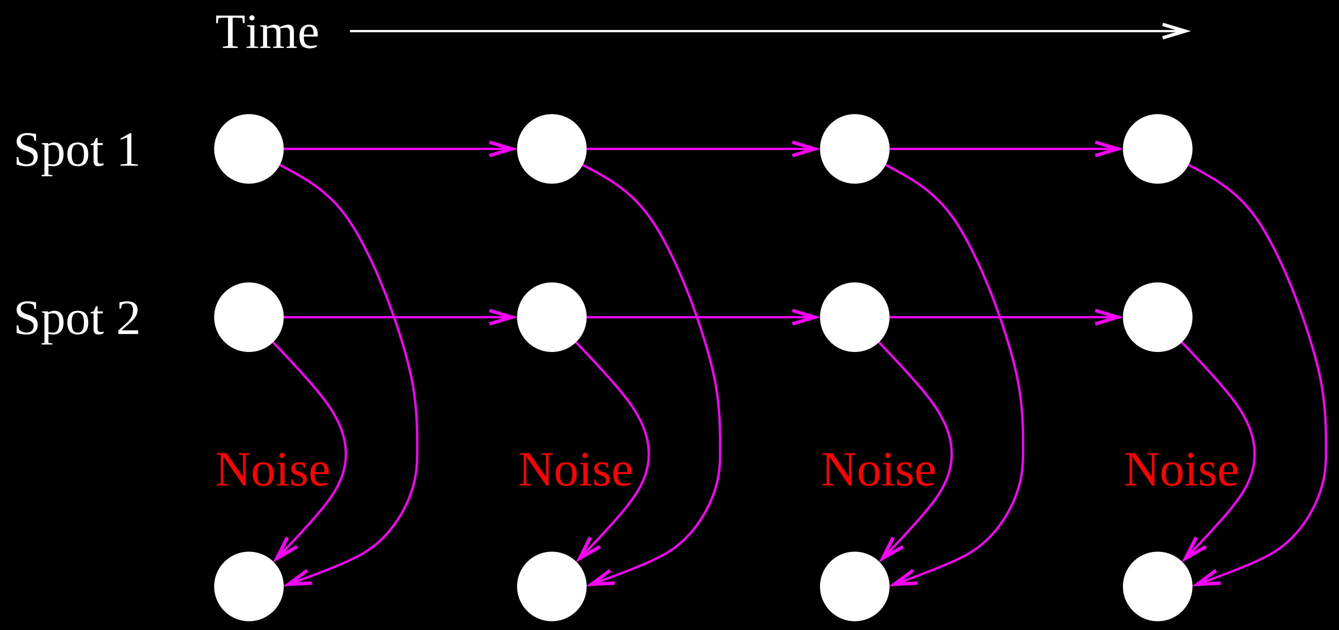

- Model: size / position / brightness / sequence
- Integrate over all state sequences (MCMC)
- Max likelihood estimate and error bars for pos/size/brightness (difficult optimization)

Bayesian analysis

- Generative model for images

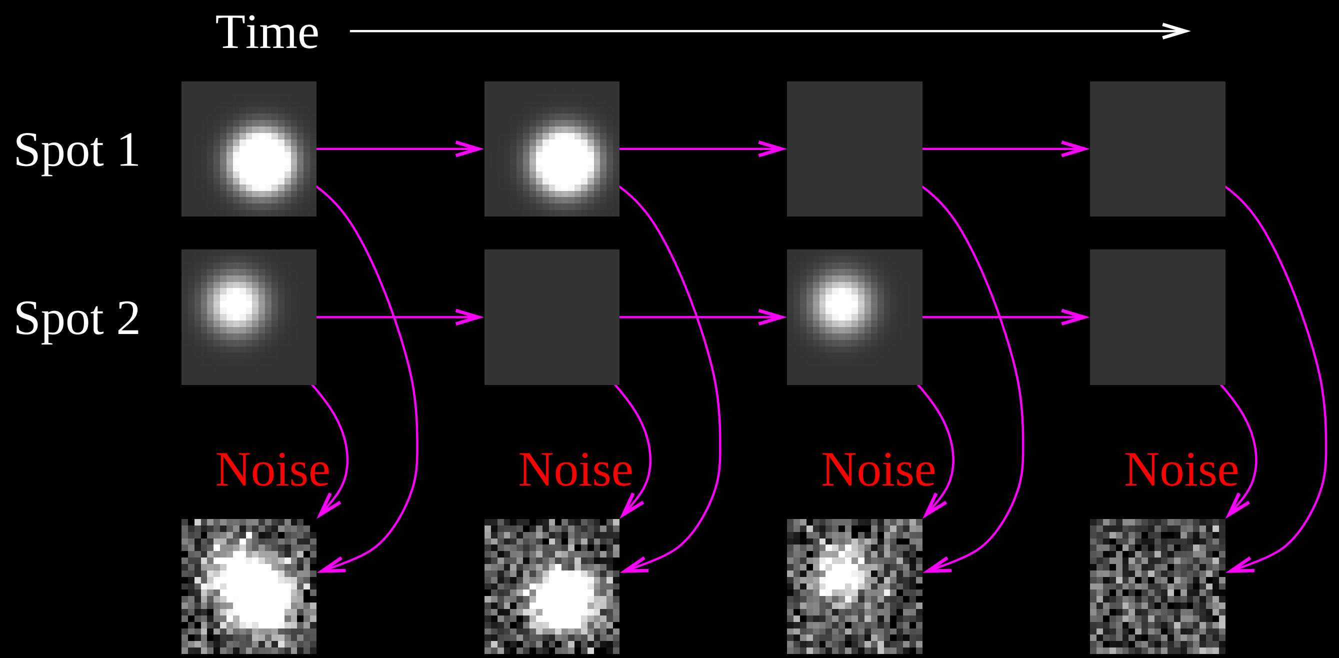

- Model: size / position / brightness / sequence
- Integrate over all state sequences (MCMC)
- Max likelihood estimate and error bars for pos/size/brightness (difficult optimization)

Results: first, some biology (Podosomes)

- Cytoskeletal structures at the cell membrane
- Actin core, surrounded by a ring of other proteins including vinculin and talin

Confocal STED

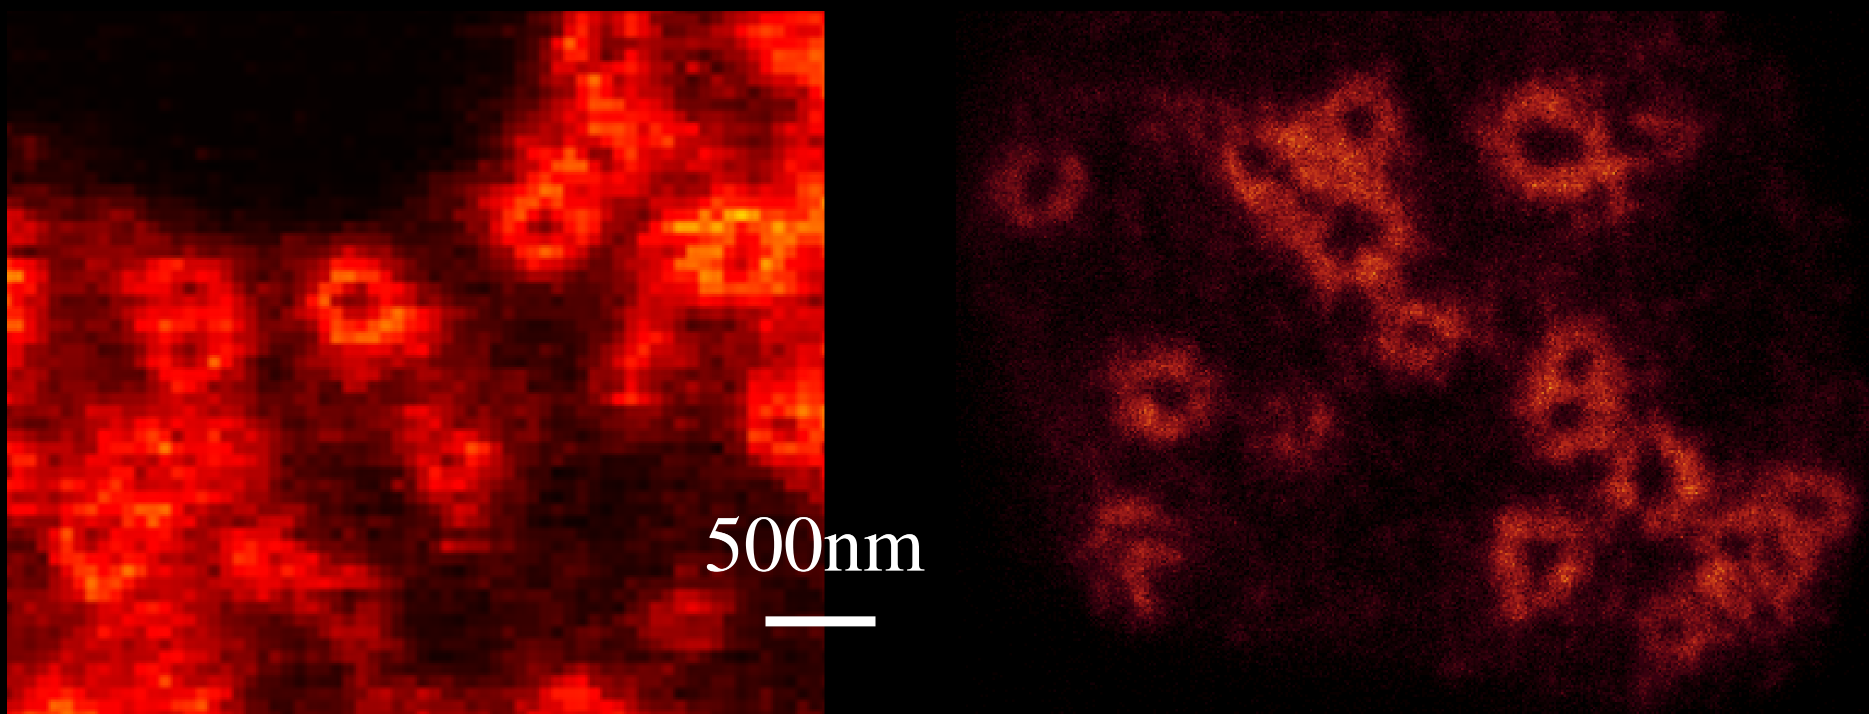

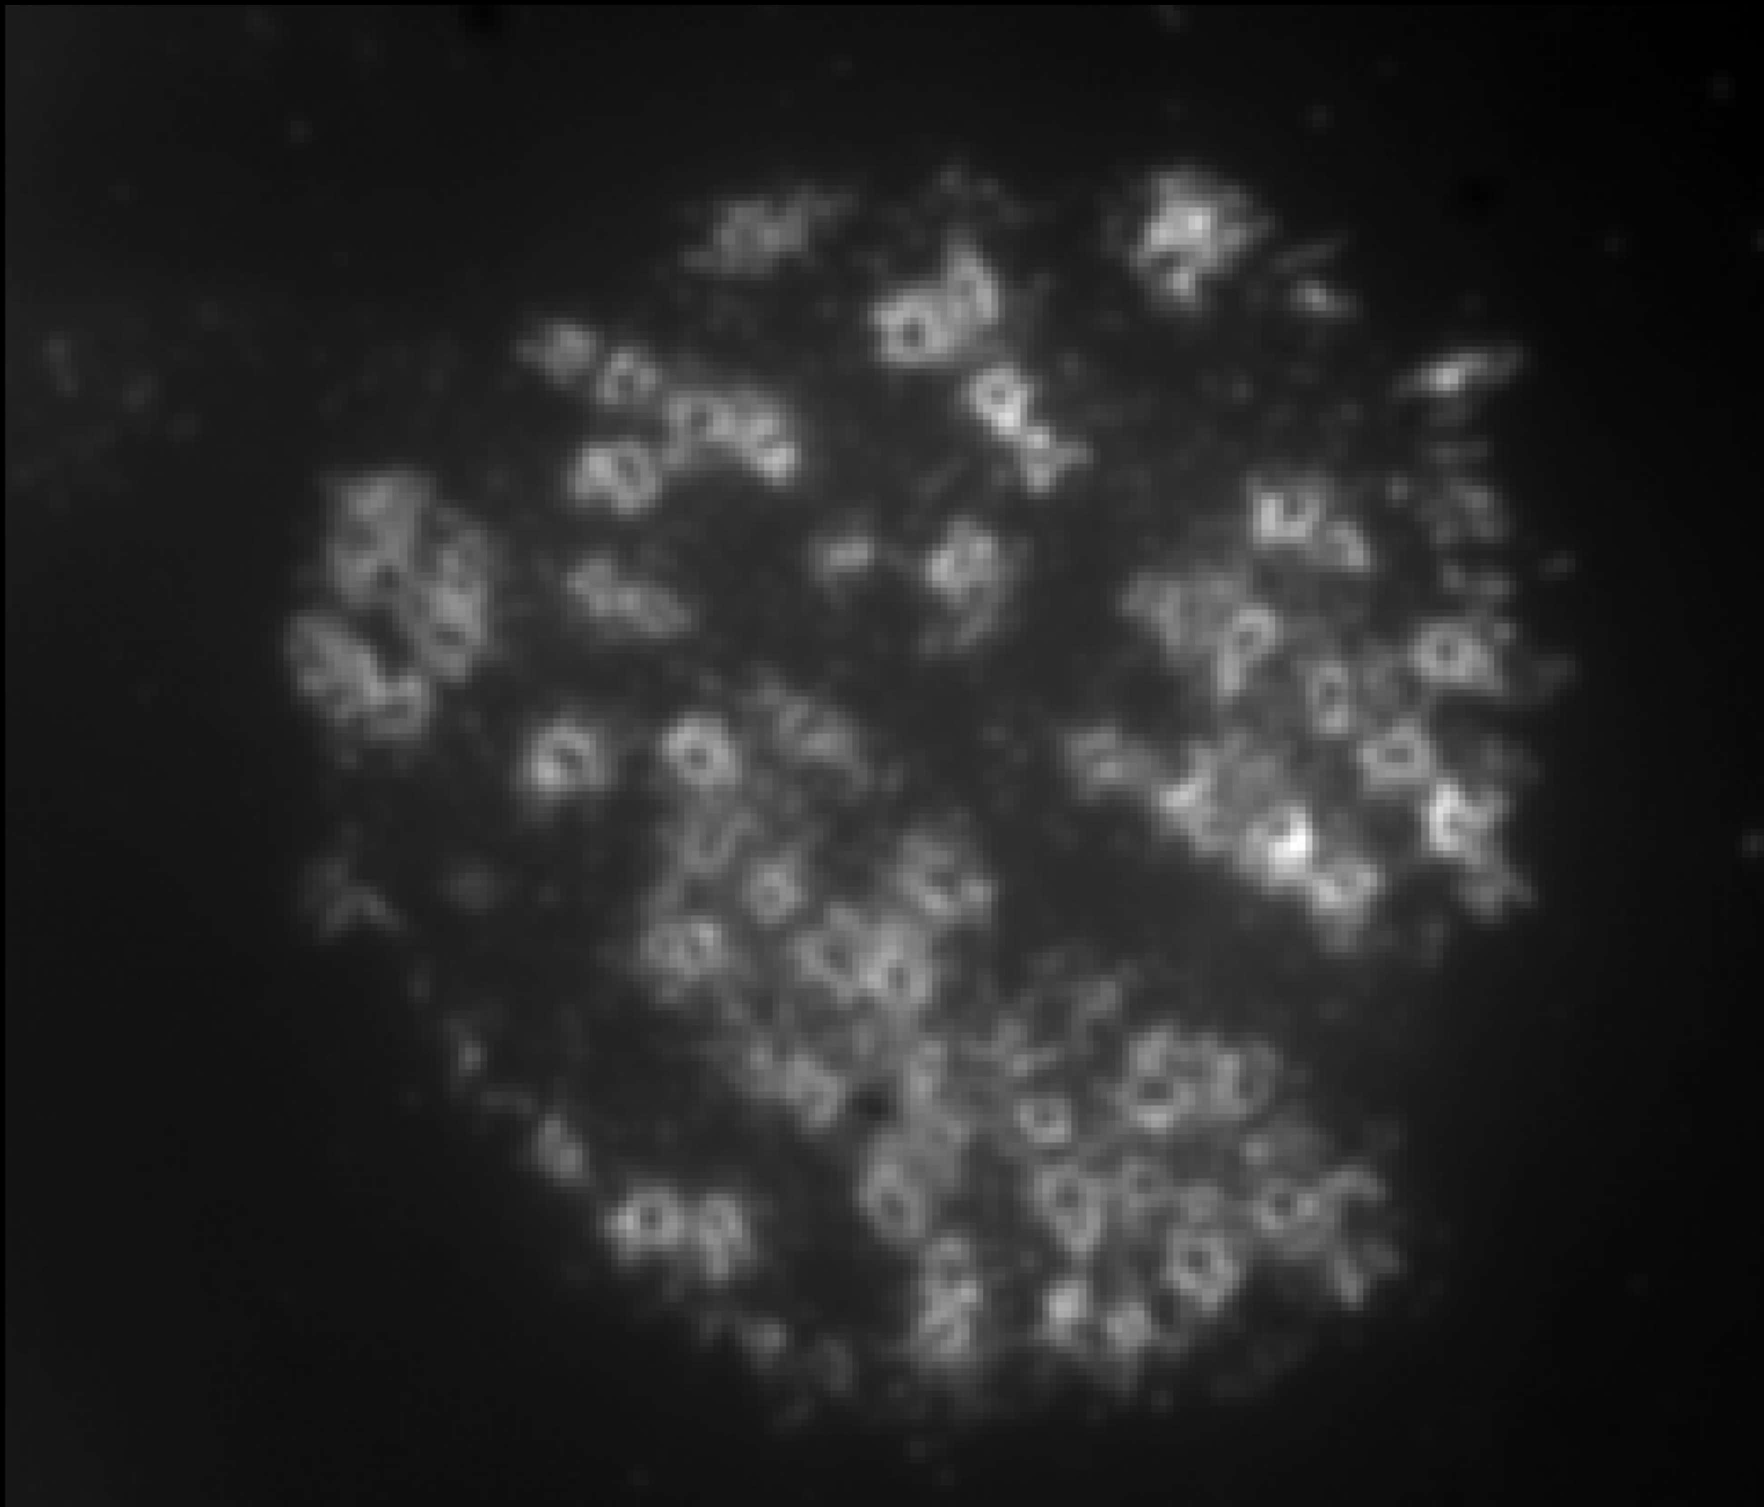

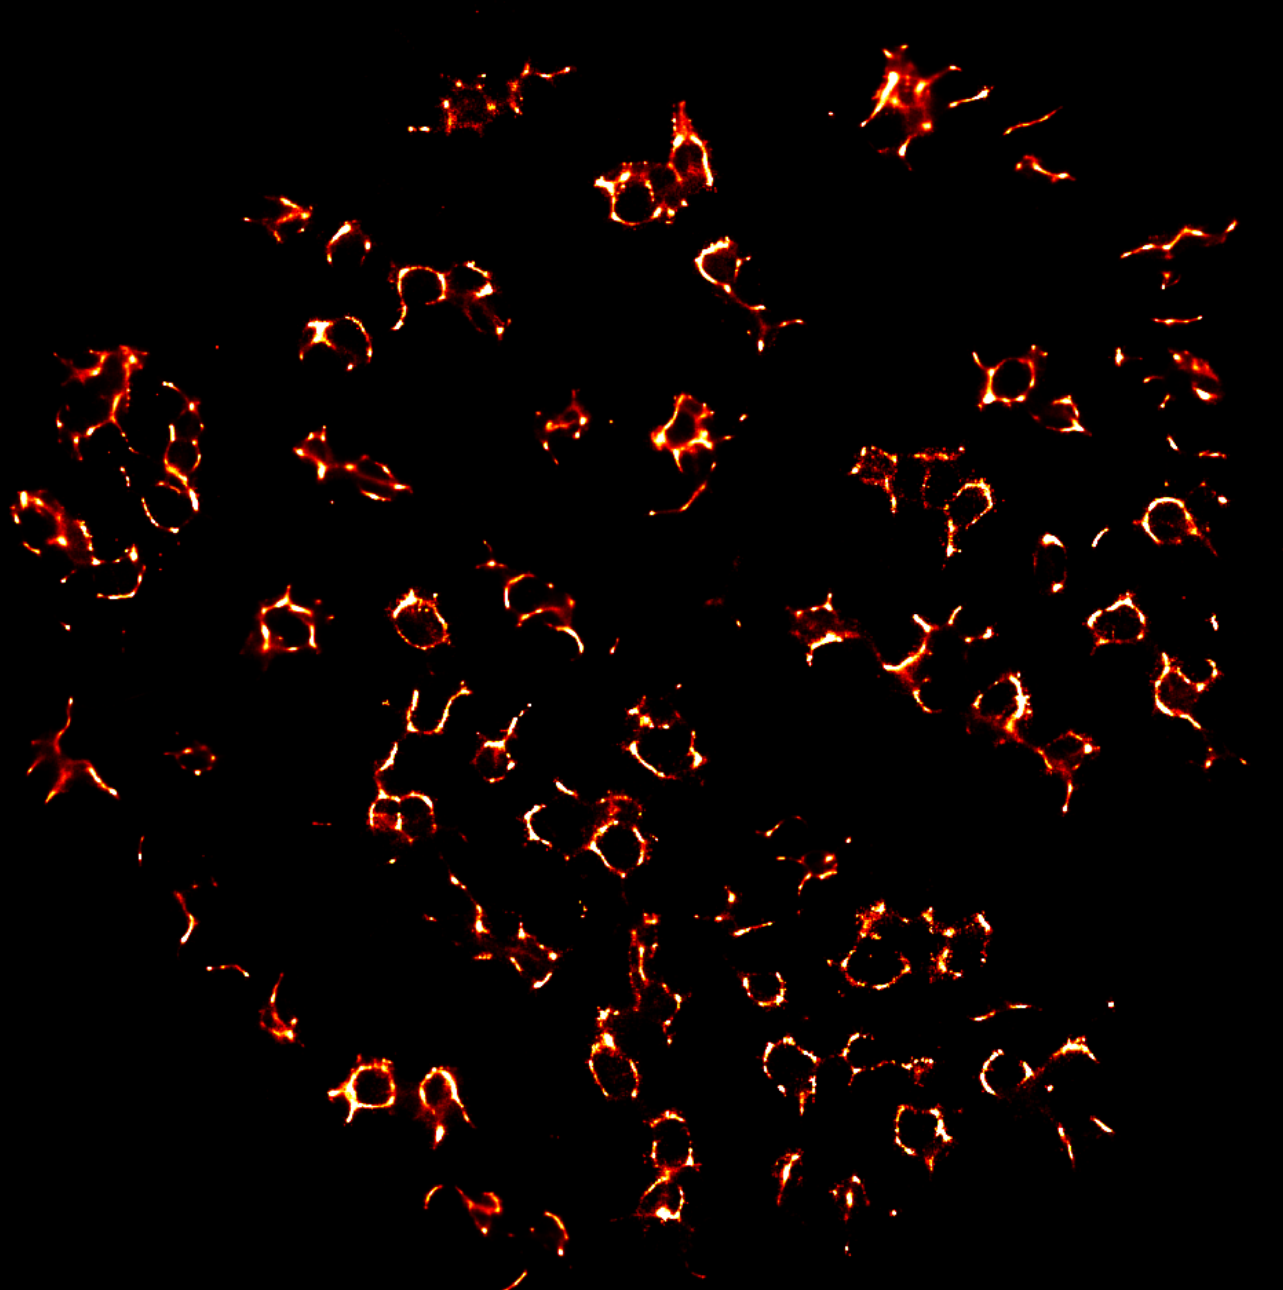

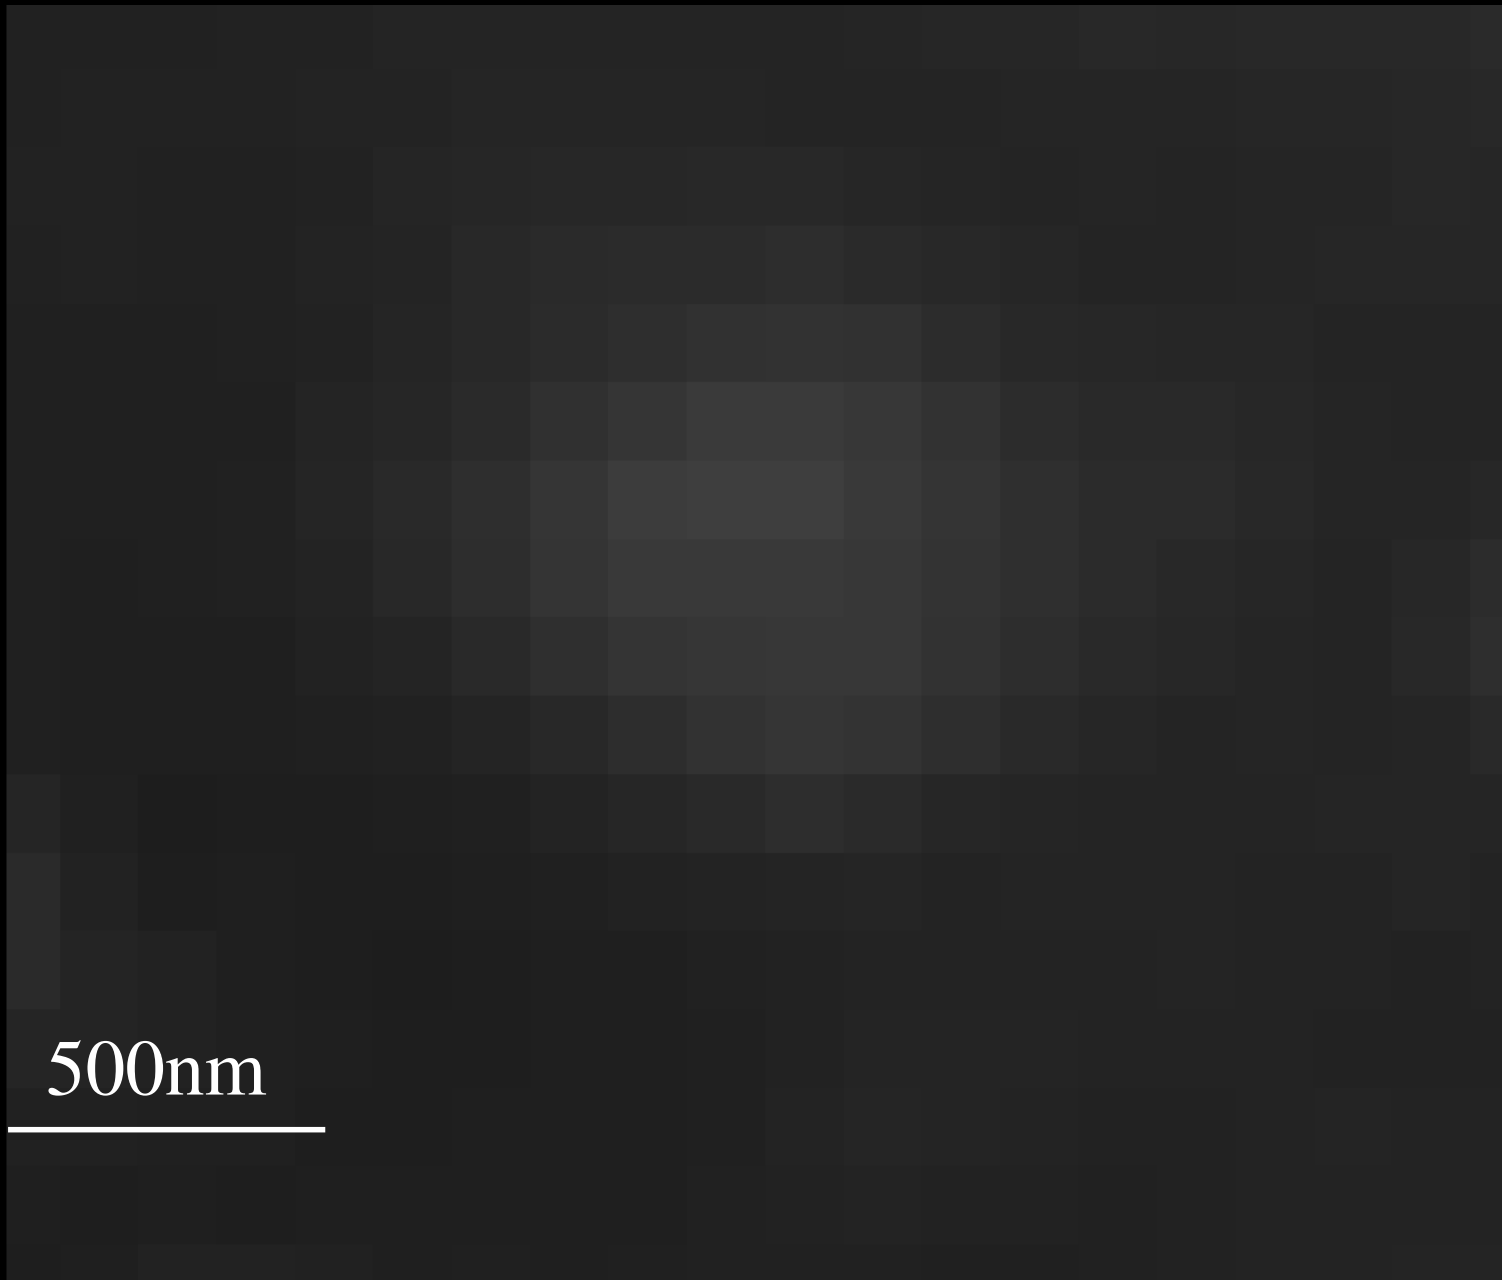
500nm

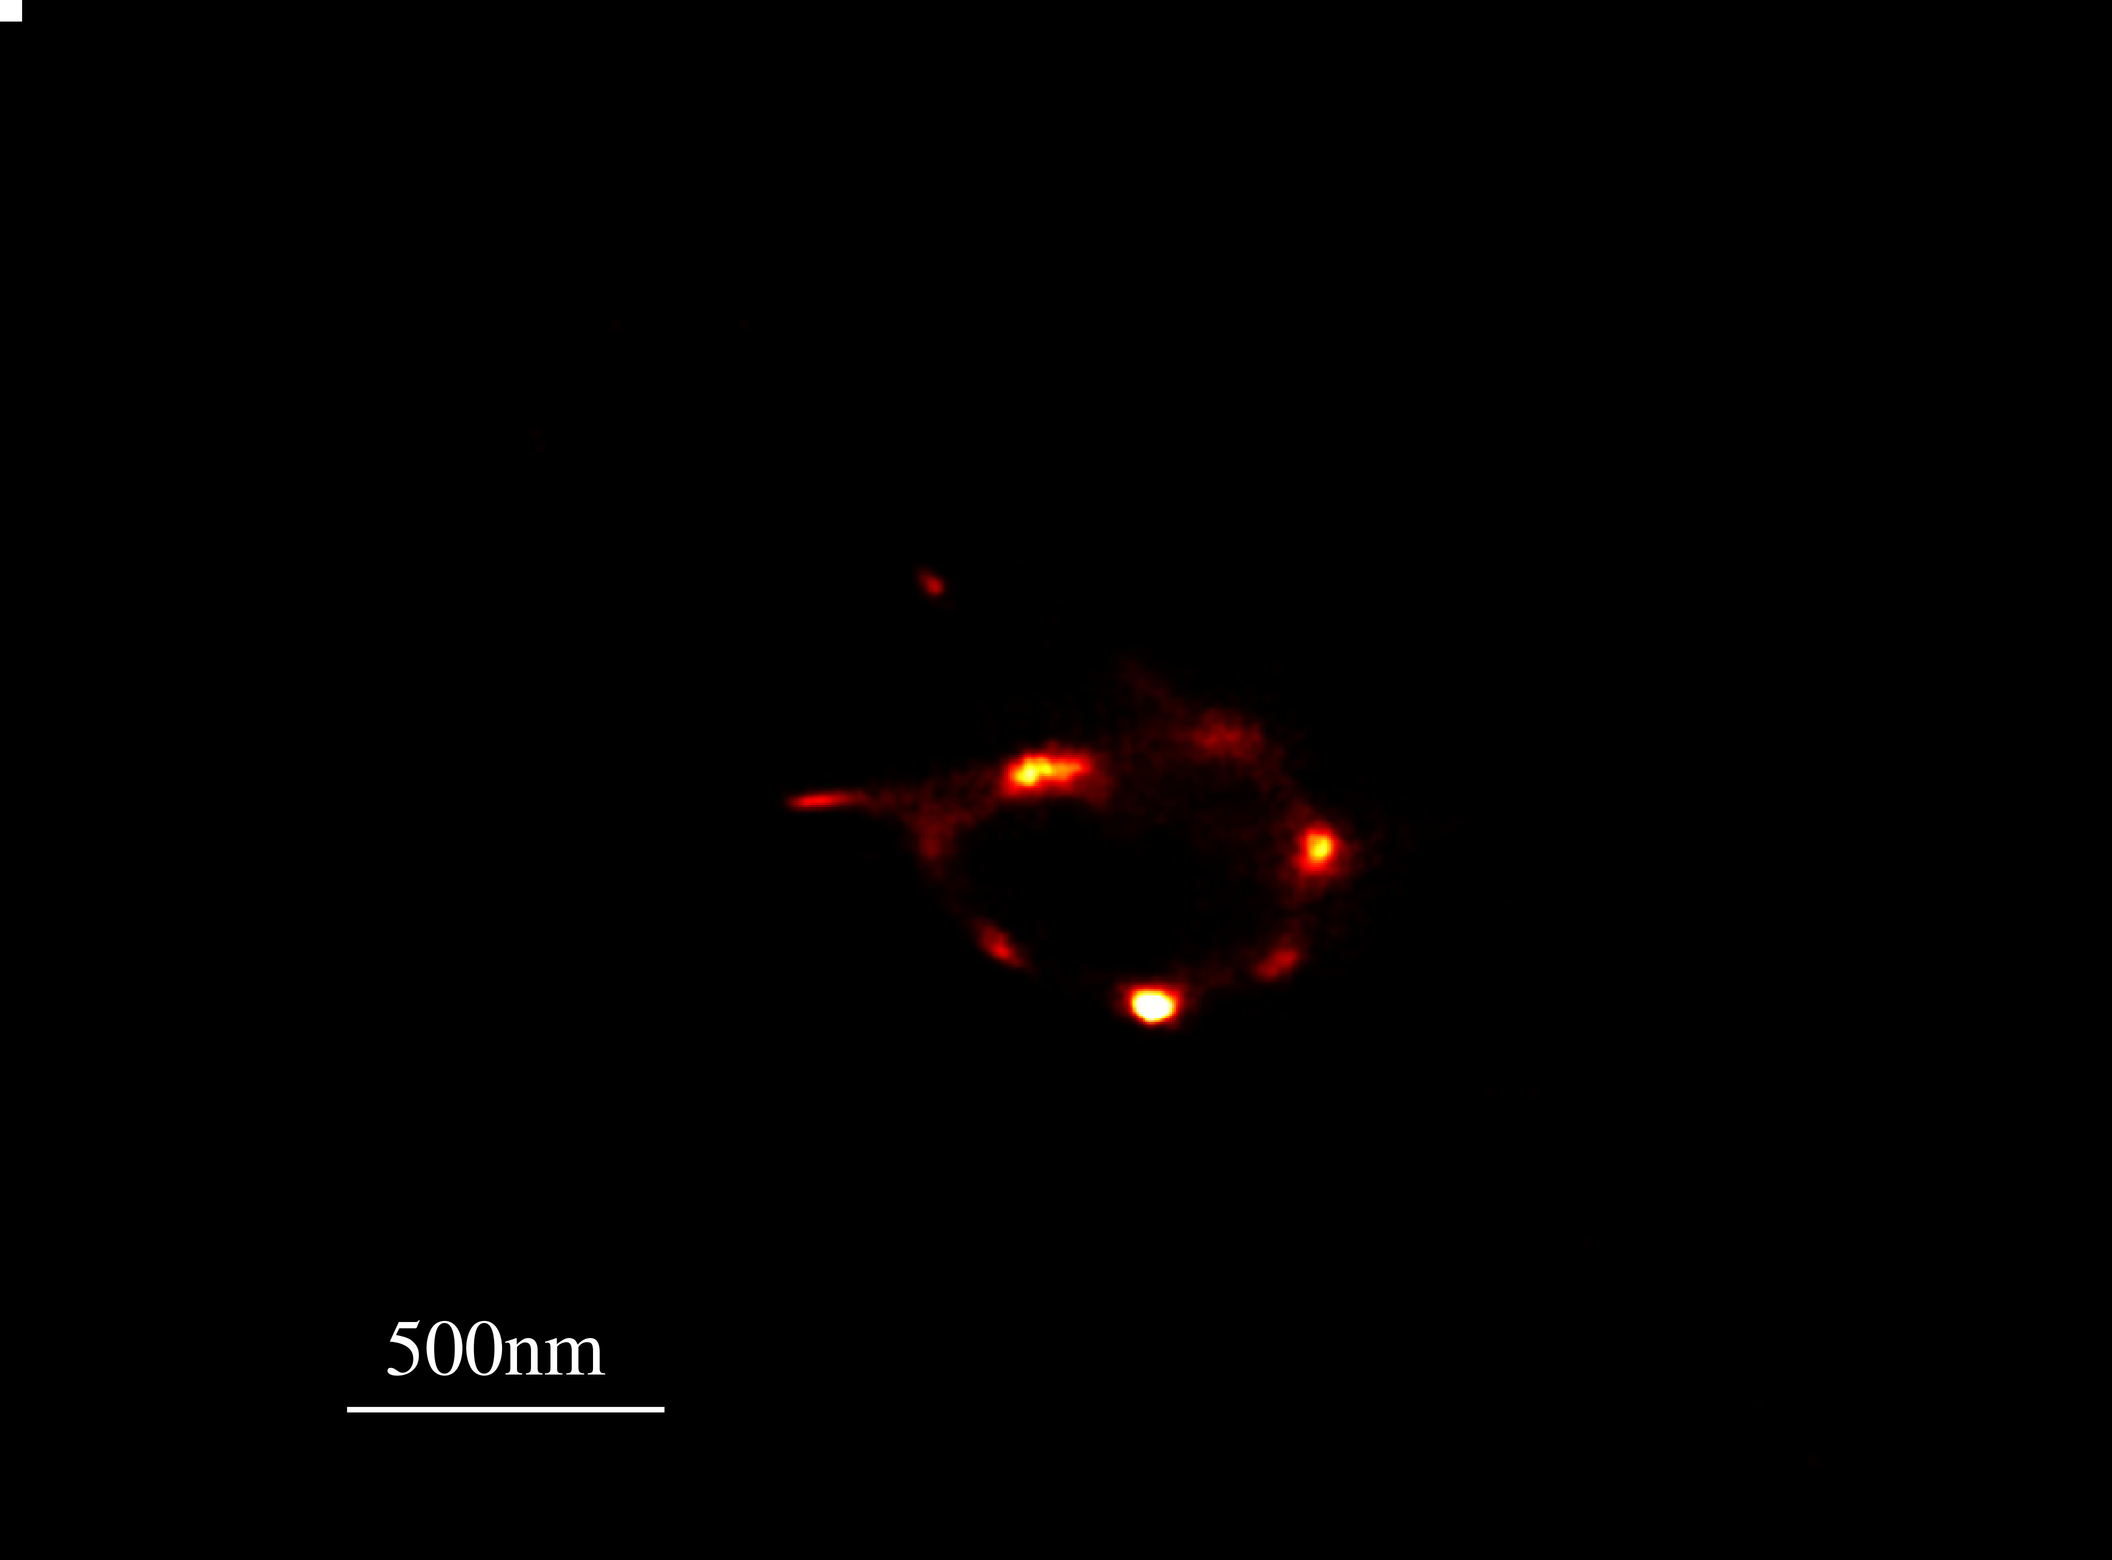
500nm

Live cells: truly awful data

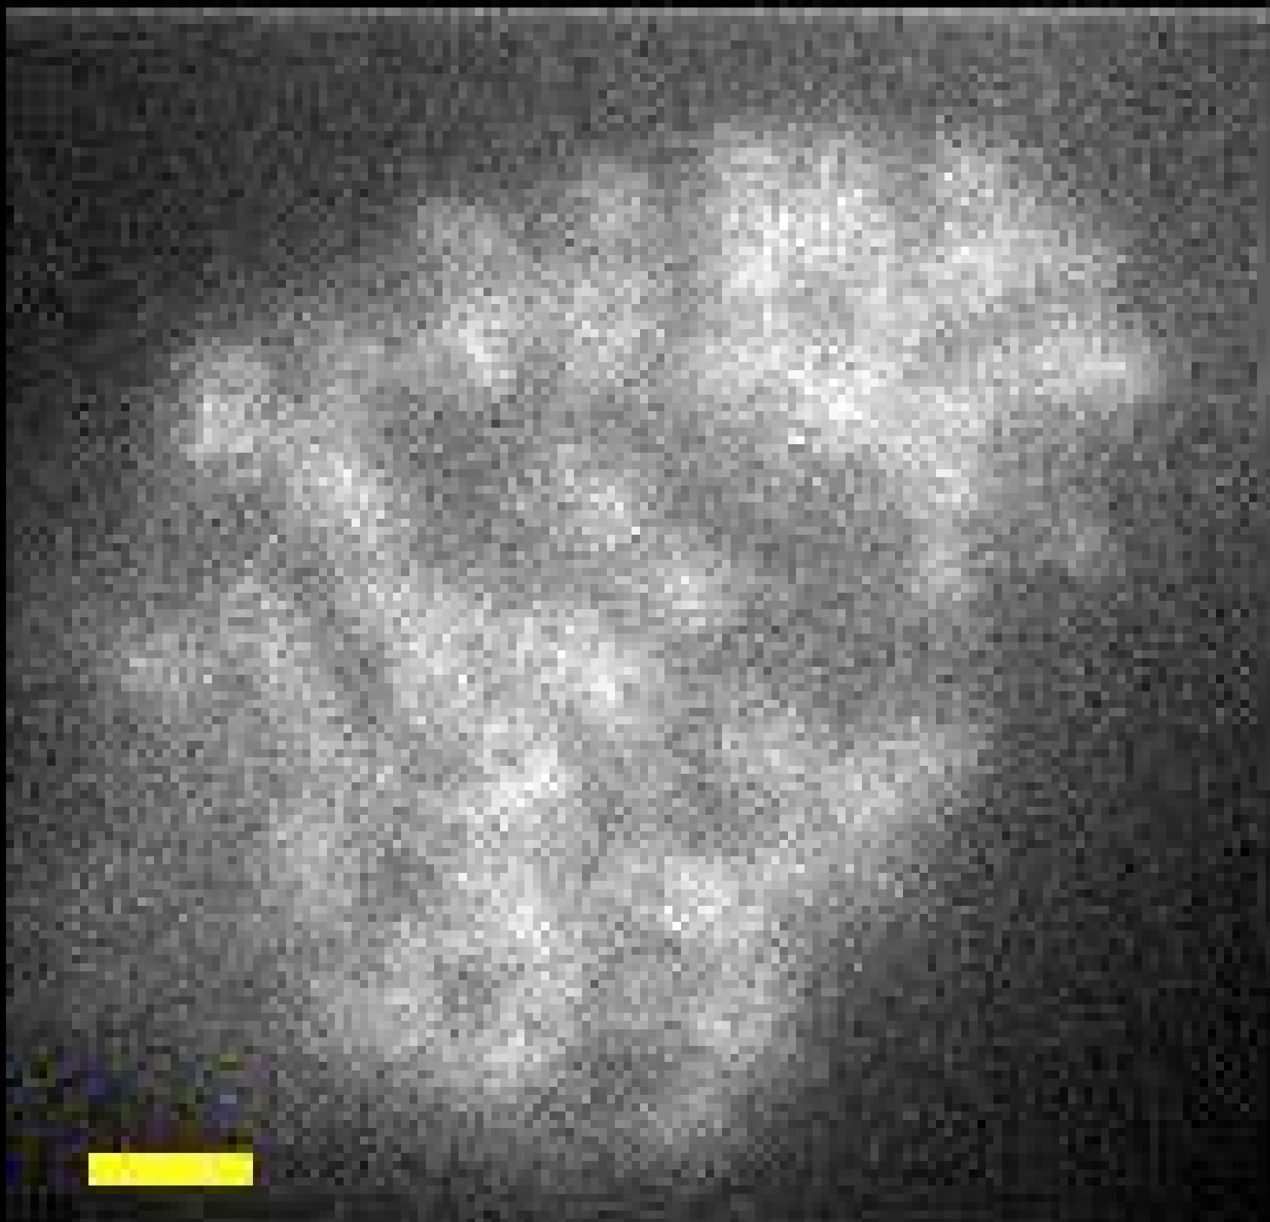

Live cell podosomes

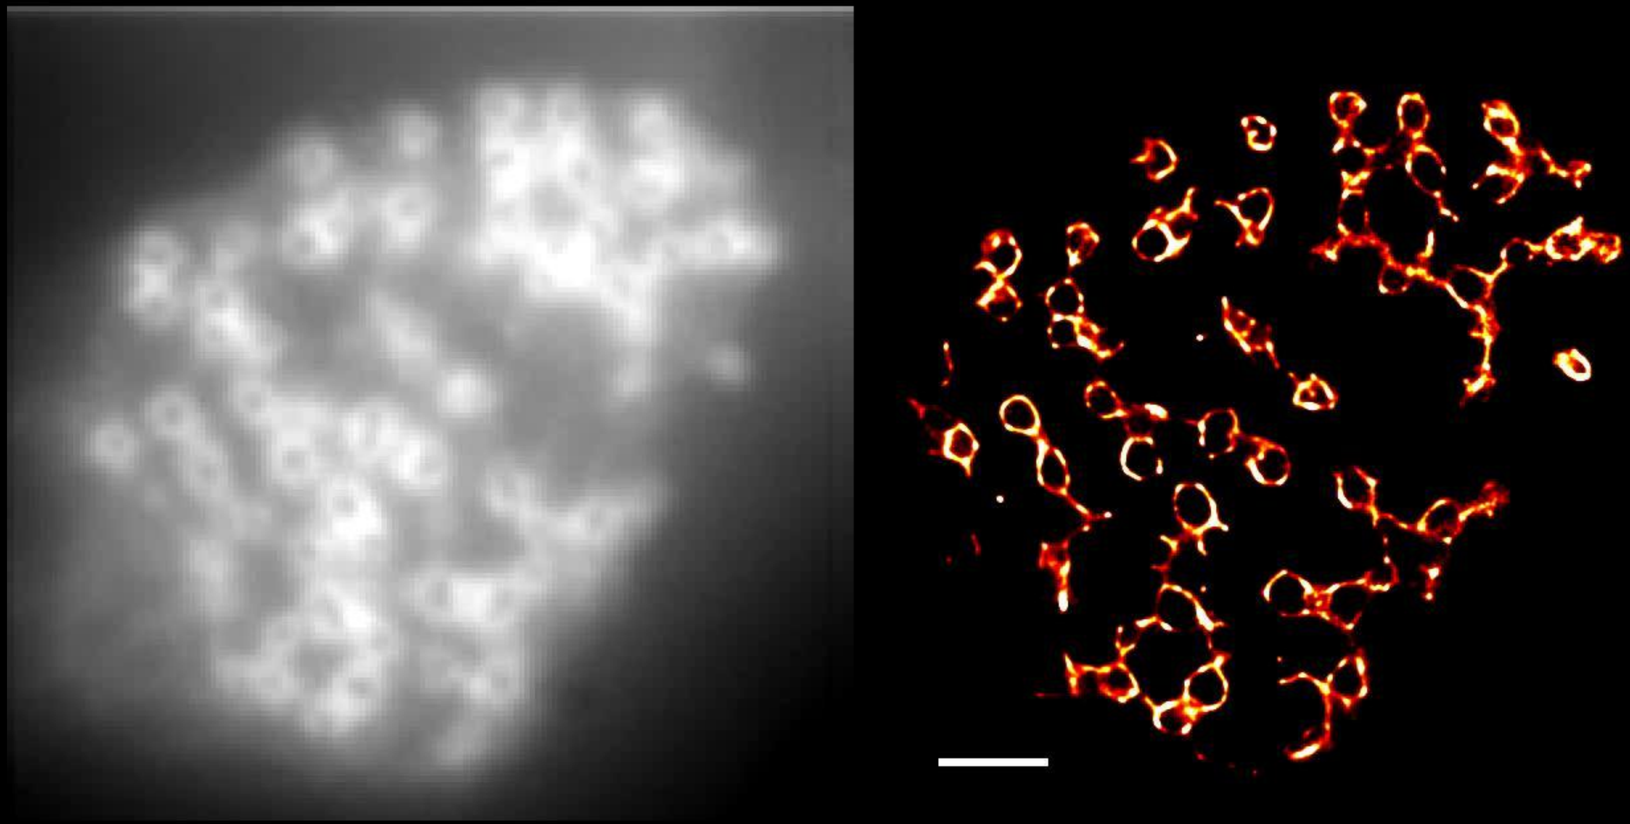

What is corner detection?

Useful for:

- 2D tracking, 3D tracking, live 3D reconstruction, object recognition, etc.

- Visually 'salient' features.
- Localized in 2D.
- Sparse.
- High 'information' content.
- Repeatable between images.

FAST-ER: Enhanced Repeatability

- Define feature detector as:

  A collection of FAST-like tests which detects points with a high repeatability.

- To evaluate repeatability:
  1. Detect features in all frames.
  2. Compute repeatability.

- That is hard to optimize!

  Optimize tree using simulated-annealing.
  Brute force and ignorance?

Rosten, Porter, Drummond. PAMI 2010

FAST: features from accelerated segment test

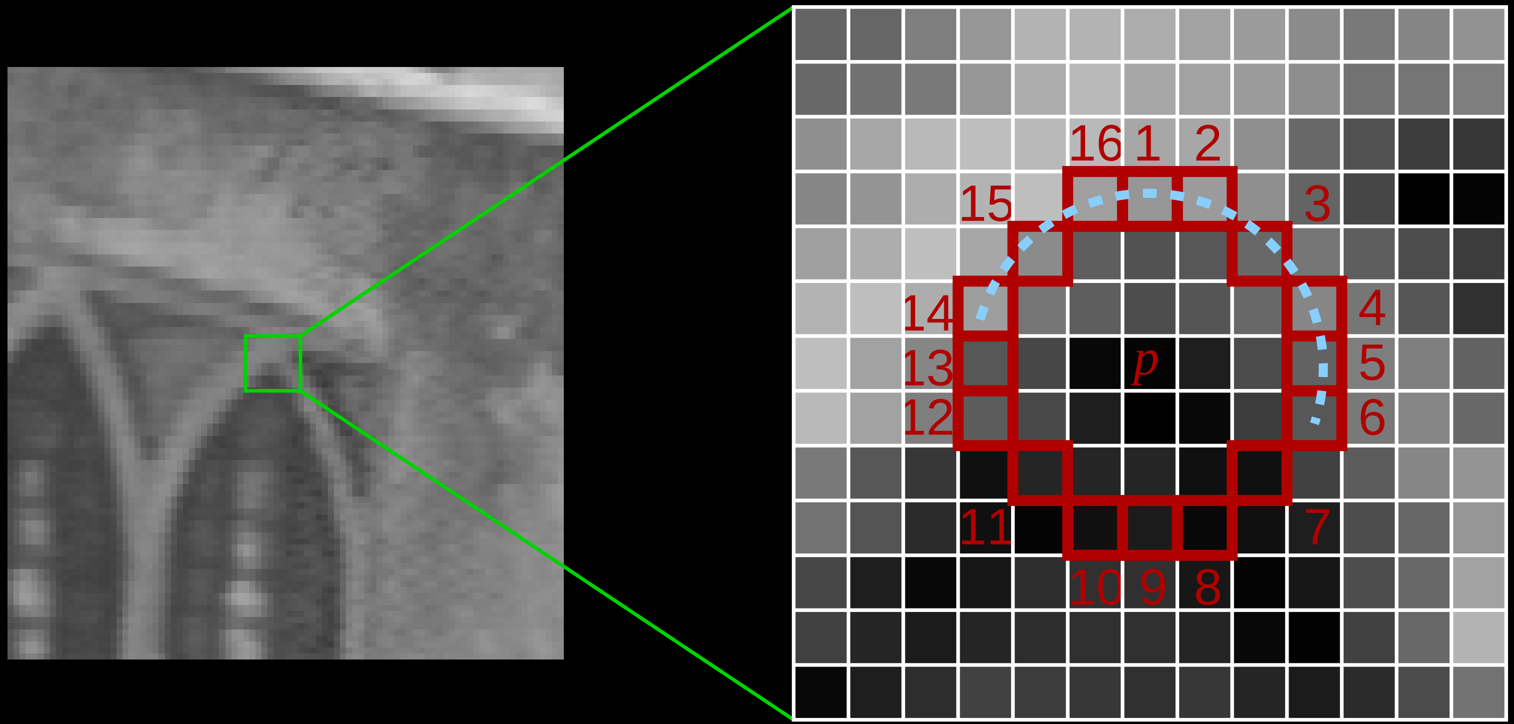

Contiguous arc of 9 or more pixels:

- All much brighter than p (brighter than $p + t$).

or

- All much darker than p (darker than $p - t$).

Rosten, Drummond. ECCV 2006, ICCV 2005

FAST-ER: Enhanced Repeatability

- Define feature detector as:

  A collection of FAST-like tests which detects points with a high repeatability.

- To evaluate repeatability:
  1. Detect features in all frames.
     Slow
  2. Compute repeatability.

- That is hard to optimize!
  - Optimize tree using simulated-annealing.
    Many iterations required. Use machine code!
  - Brute force and ignorance?

Results: datasets

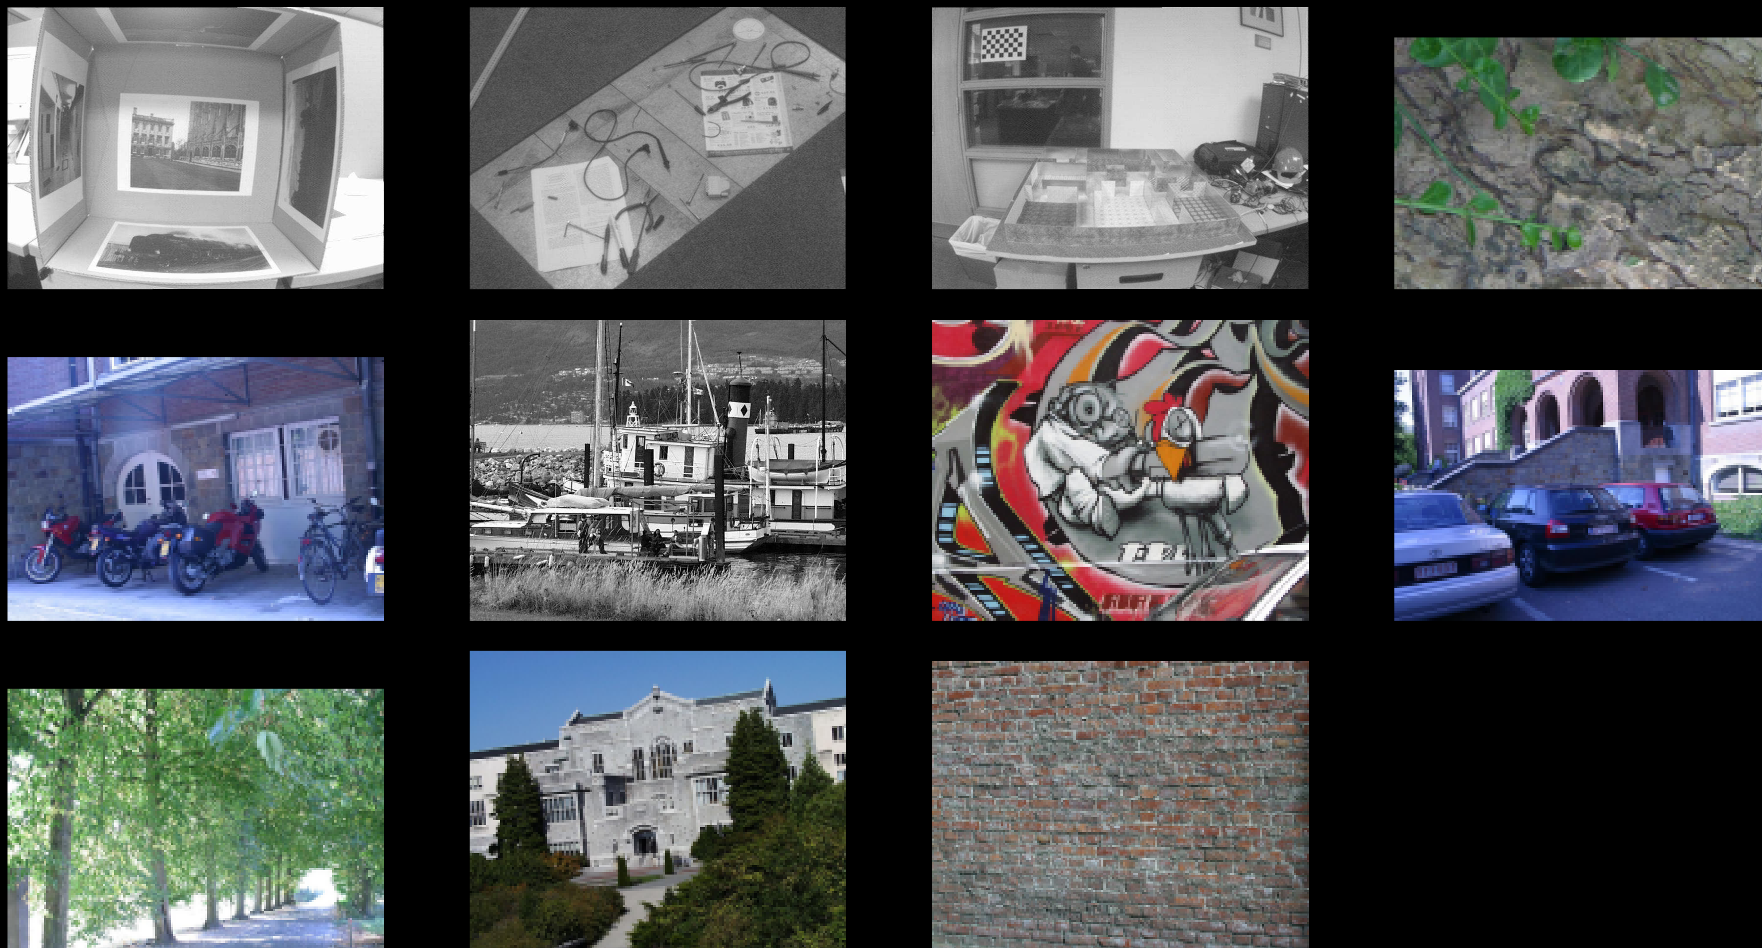

Cambridge Repeatability Dataset, Oxford Affine Covariant Regions Dataset

Aggregate results

Detector	AUR
FAST-ER	1313.6
FAST-9	1304.57
DoG	1275.59
Shi & Tomasi	1219.08
Harris	1195.2
Harris-Laplace	1153.13
FAST-12	1121.53
SUSAN	1116.79
Random	271.73

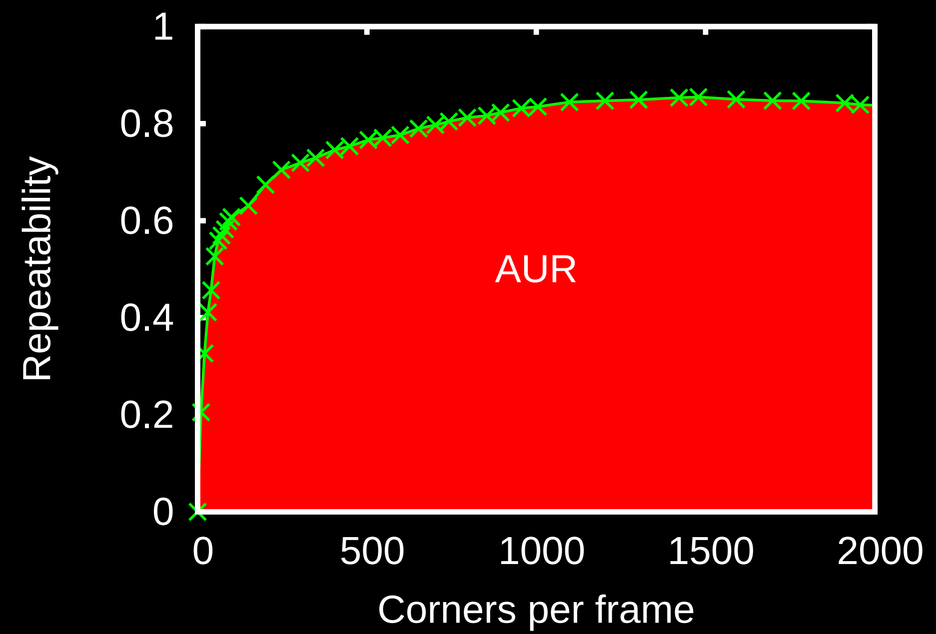

Conclusions

- Detection is meaningless without defining the task
- Task defines the objective function
  - Objects
    - ⋆ Cost based on performance in tracking system. (sort of)
  - Fluorophores
    - ⋆ Cost based reconstructing image from detections
  - Corners
    - ⋆ No real definition, but they need to be repeatable.
  - These may not be easy to optimize (might have to cheat)
- Optimization works (of course)
- Are they the objective functions the best choices?

Acknowledgements

- Susan Cox (KCL)
- Damian Eads (UCSC, LANL, Cambridge, now wise.io, Inc)
- Tijana Jovanovic-Talisman, Dylan Burnette, Jennifer Lippincott-Schwartz (NIH)
- James Monypenny, Gareth Jones (KCL)
- Finale Doshi (Cambridge, MIT)
- Reid Porter, Andy Fraser (LANL)
- Tom Drummond (Cambridge, now Monash University)
- Rainer Heintzmann (University of Jena, KCL)
- David Helmbold (UCSC)

Correlative

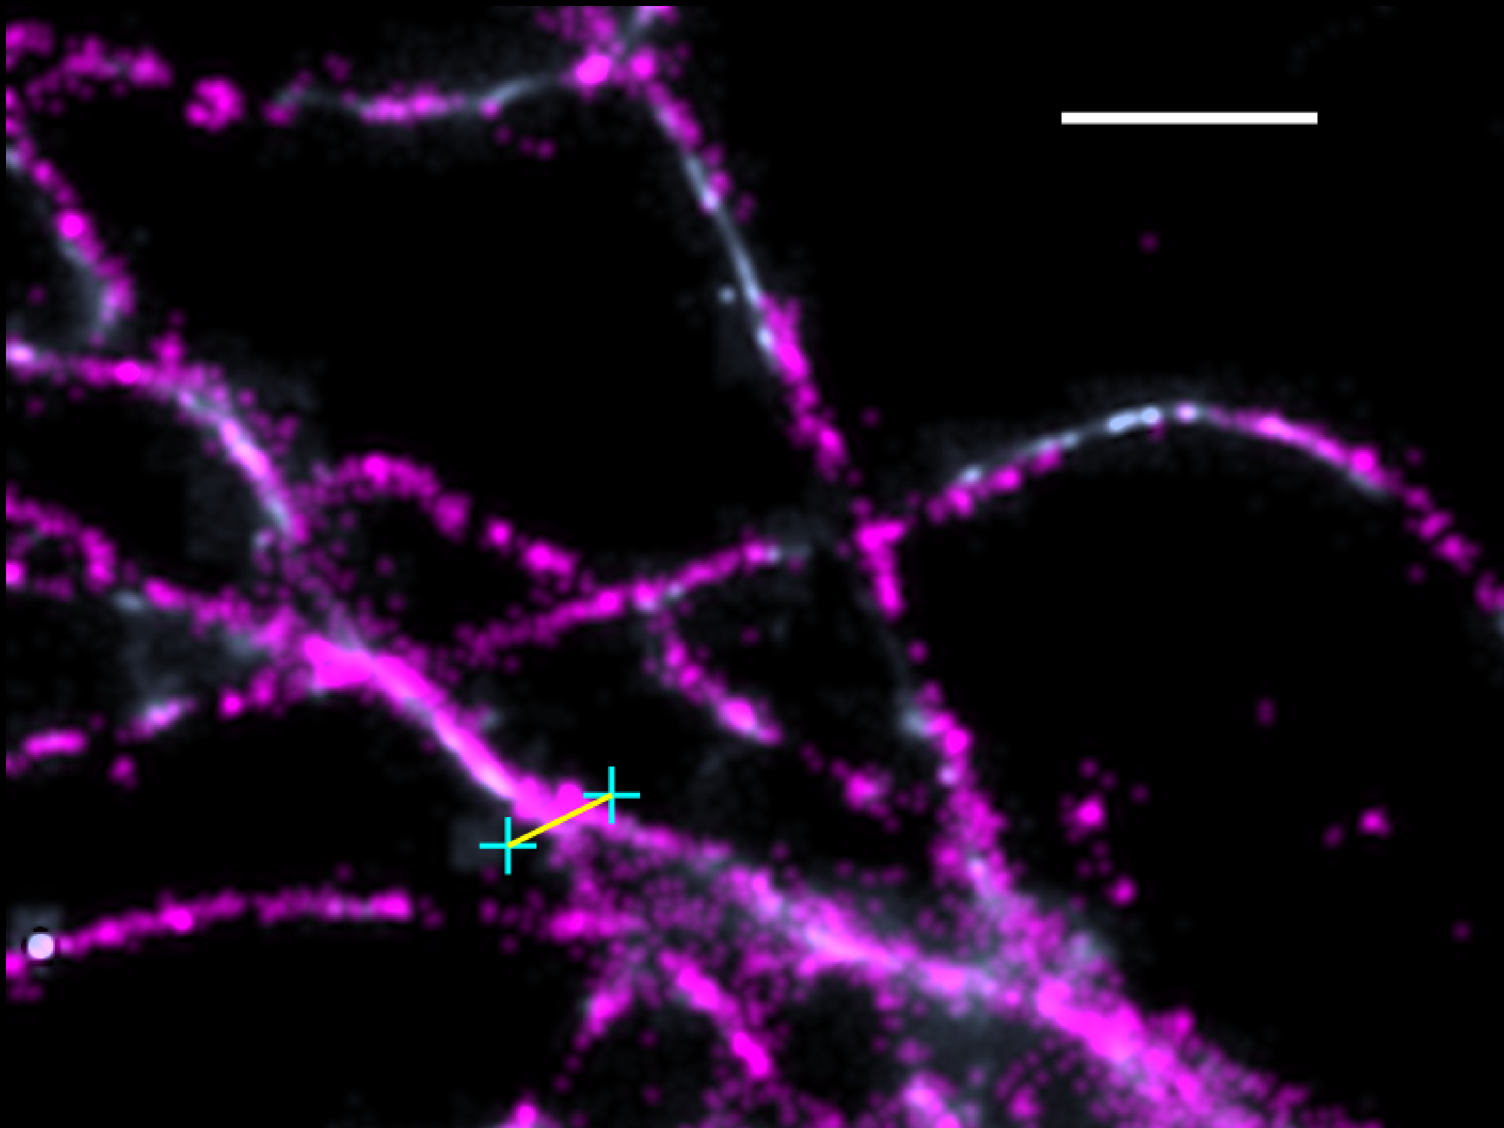

Live cell podosomes

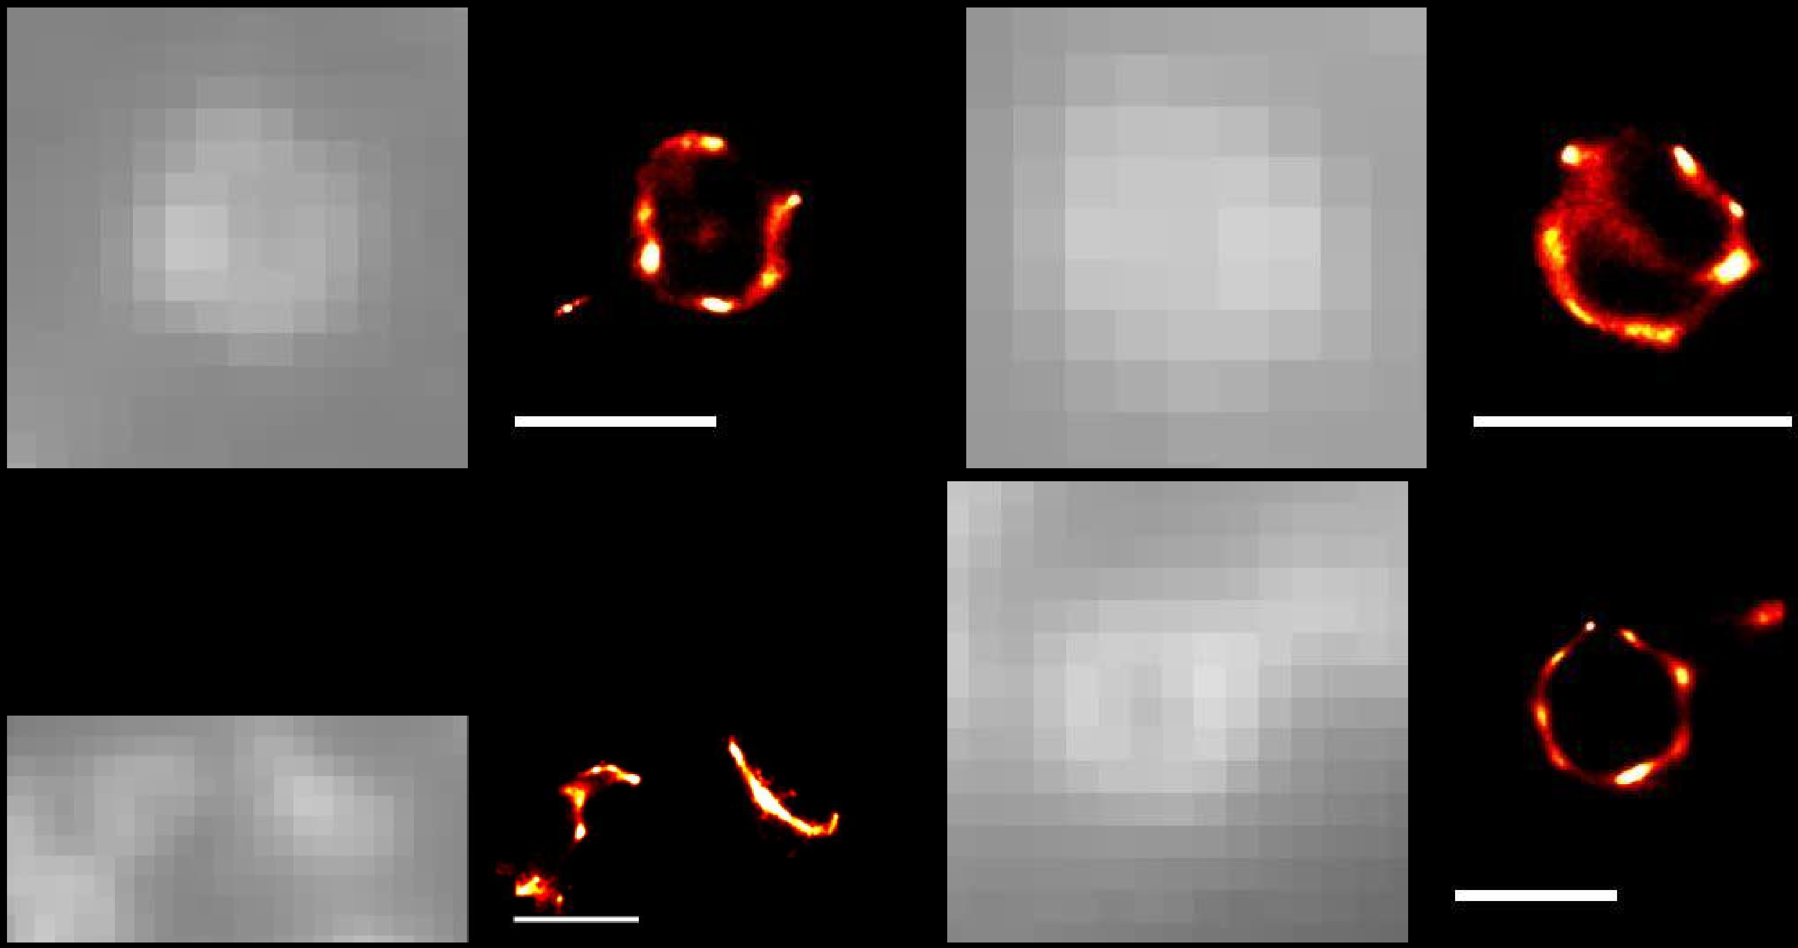

Results: Perspective (box) dataset

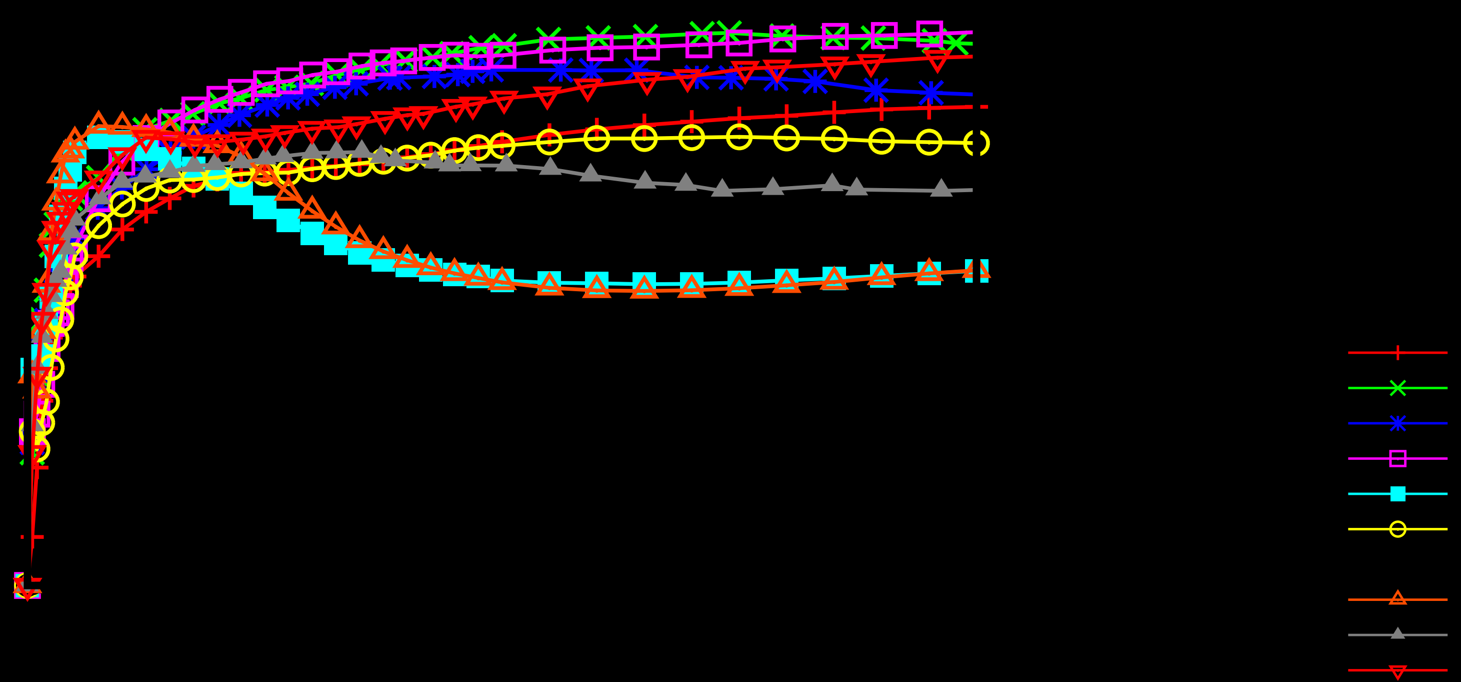

Results: Geometric dataset

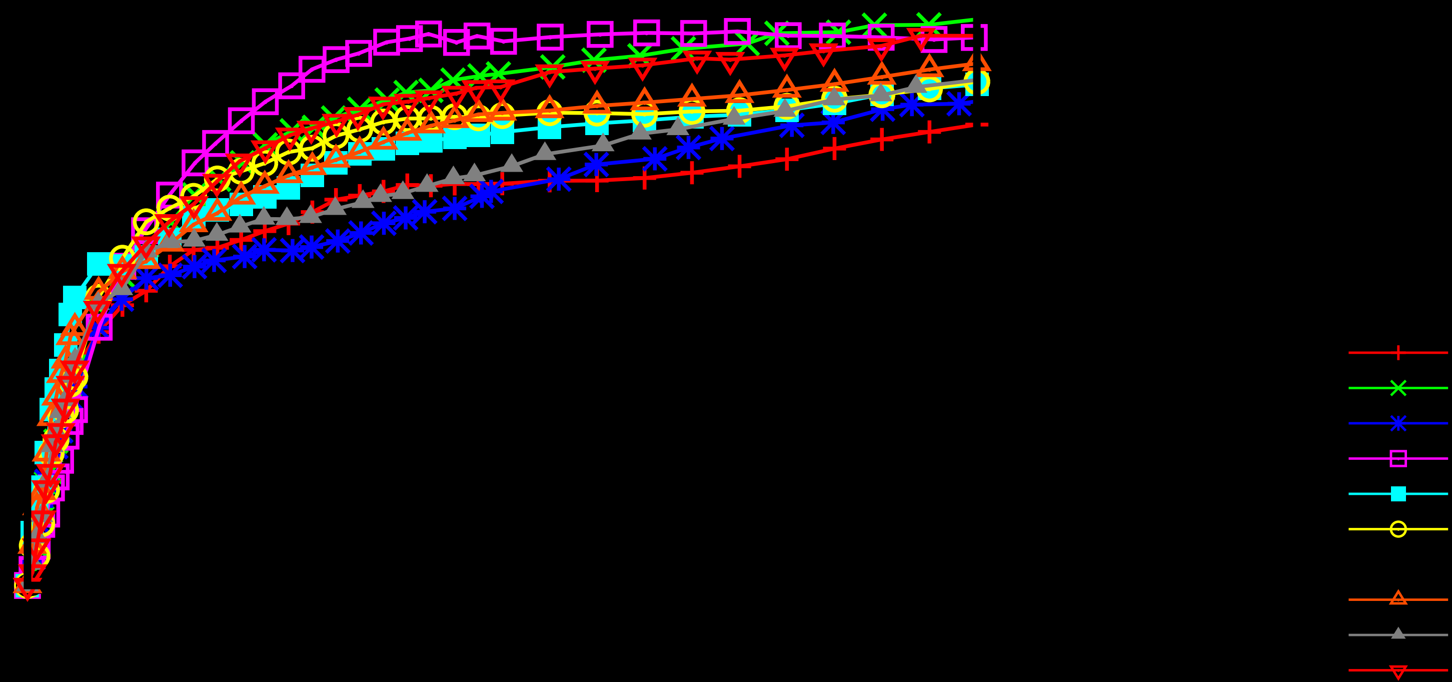

Results: Bas-relief dataset

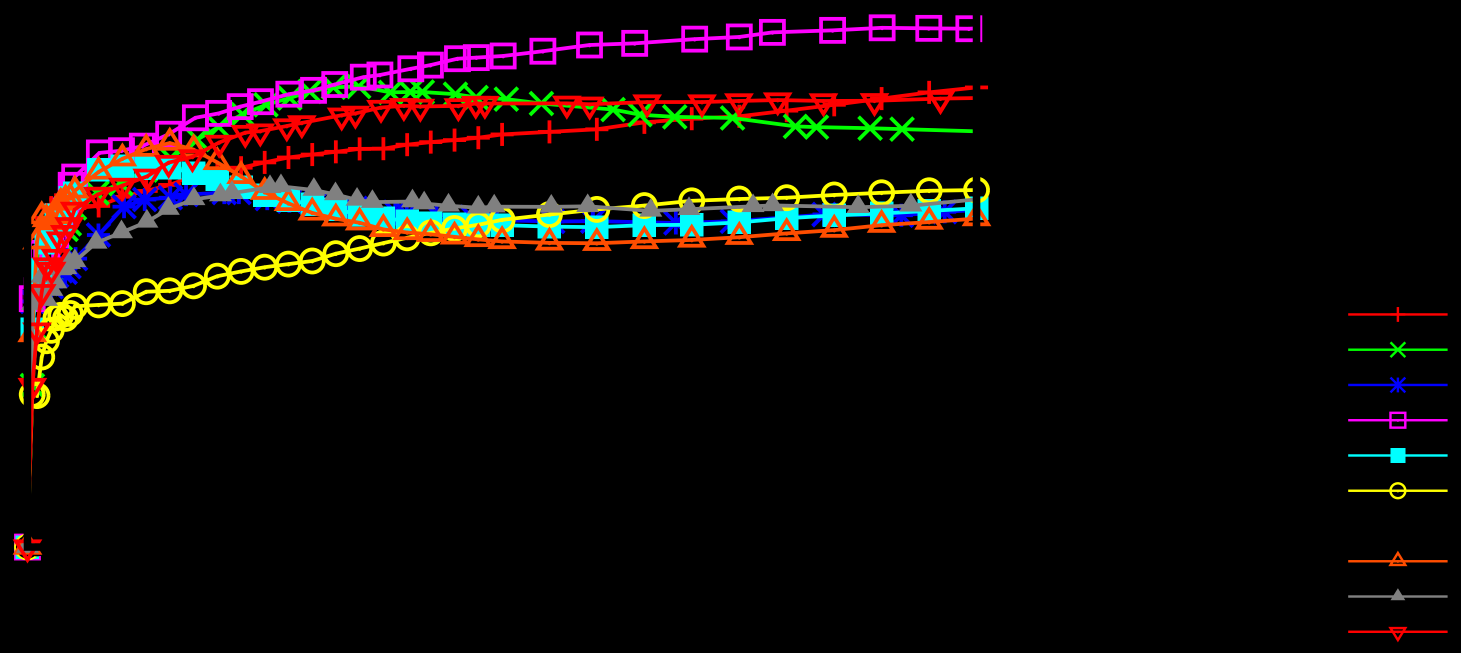

Results: Scale and rotation (bark) dataset

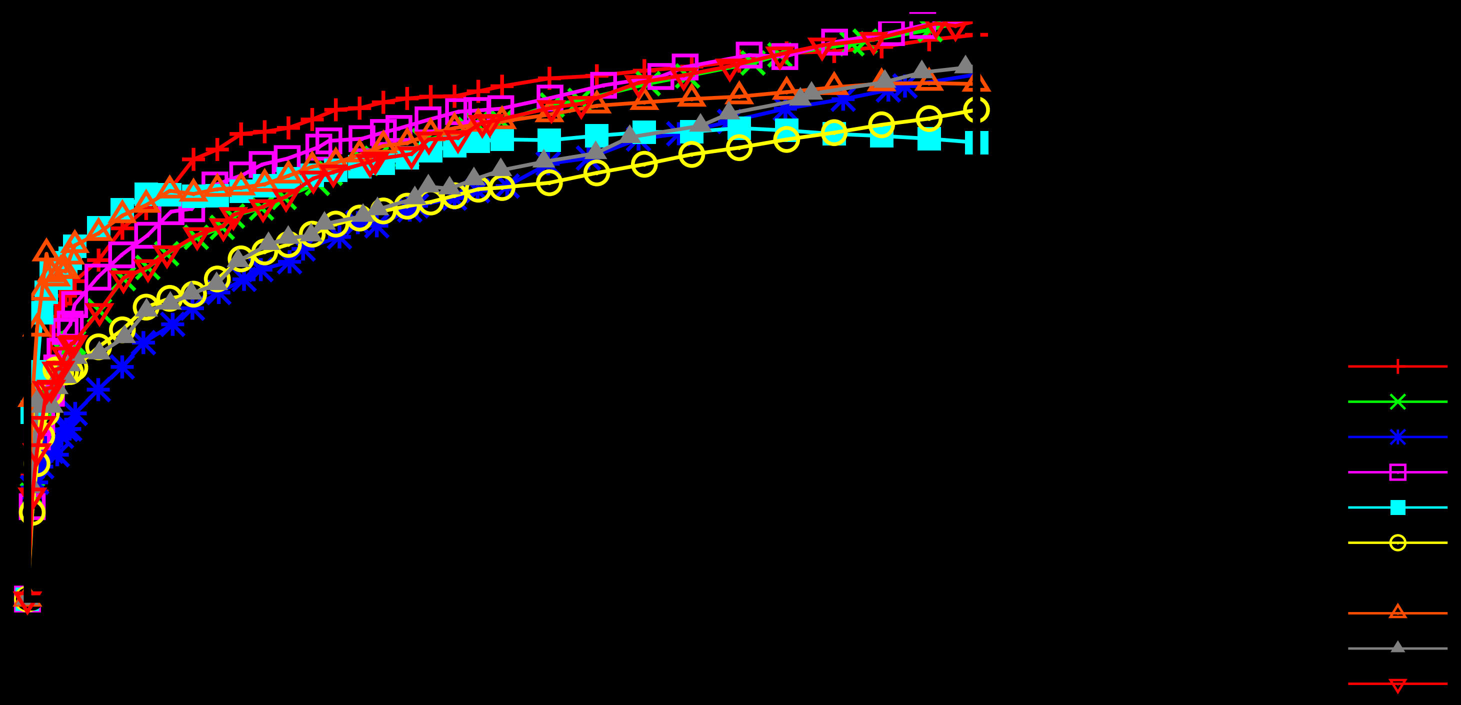

Results: Blur (bikes) dataset

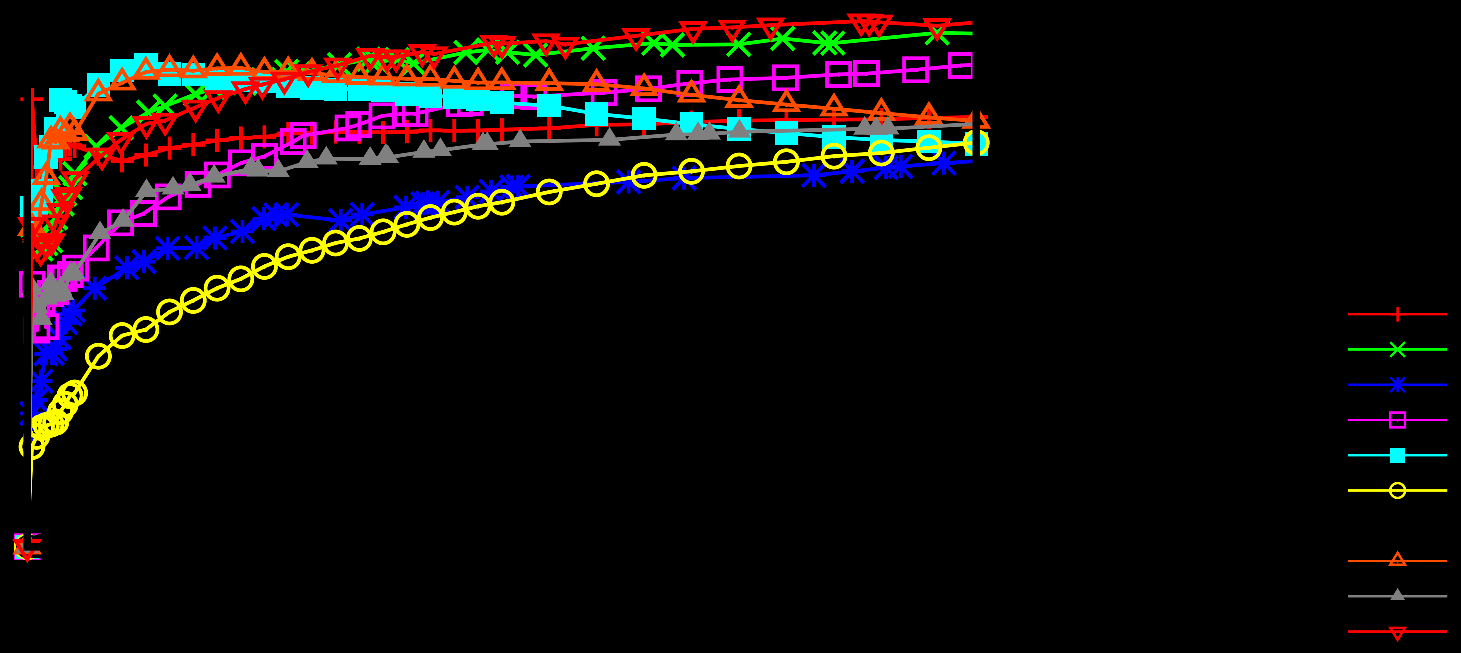

Results: Scale and rotation (boat) dataset

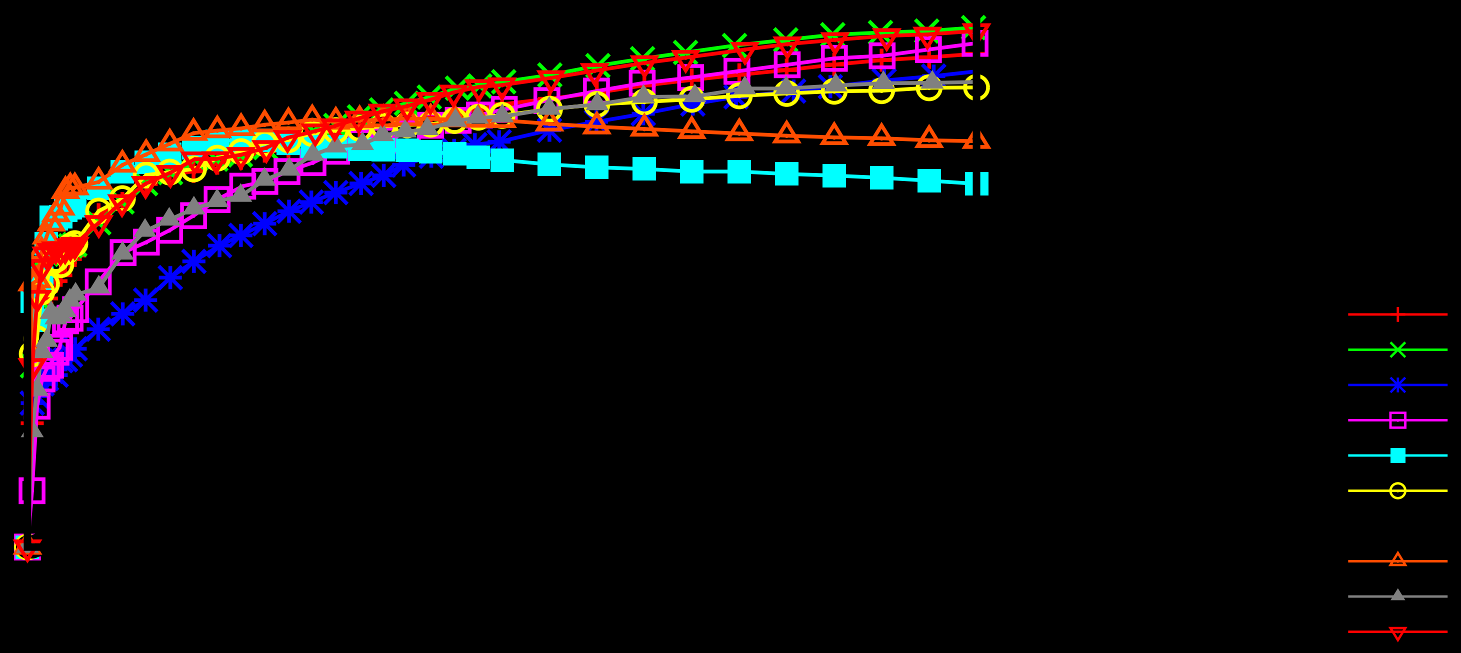

Results: Perspective (graffiti) dataset

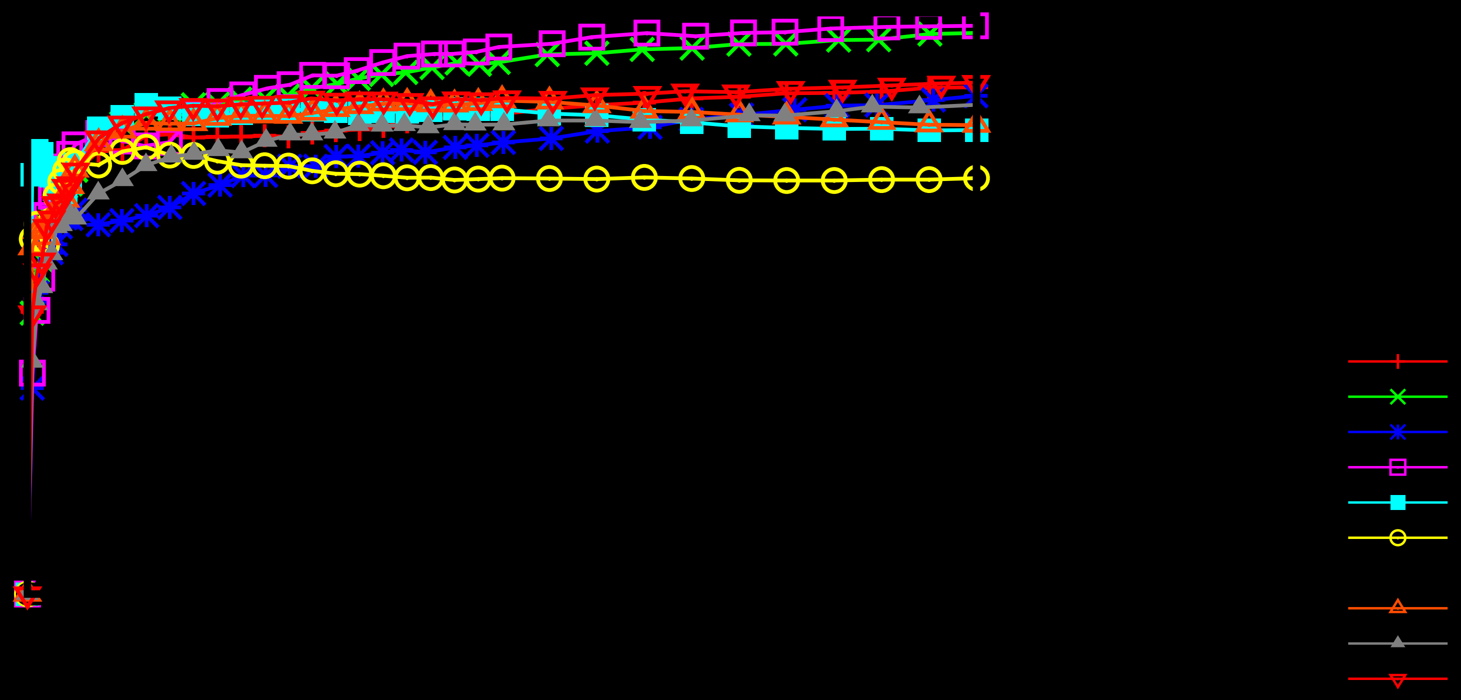

Results: Lighting dataset

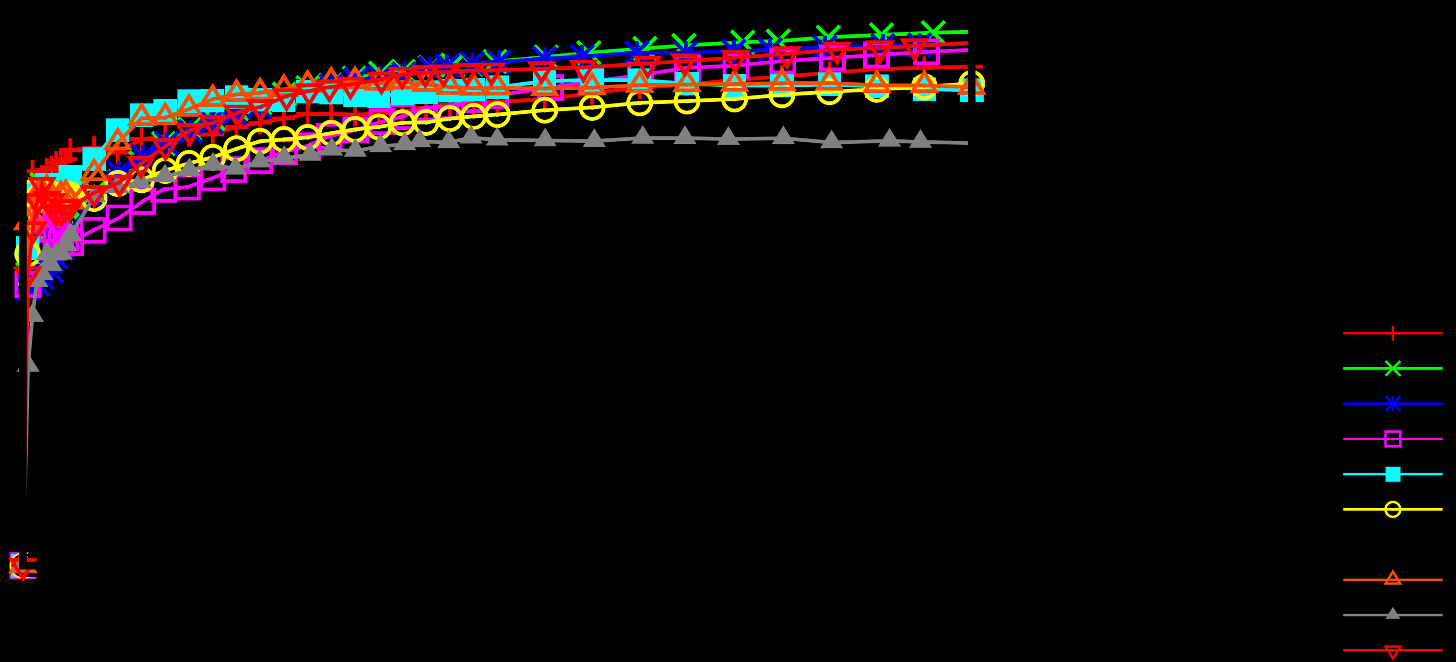

Results: Blur (trees) dataset

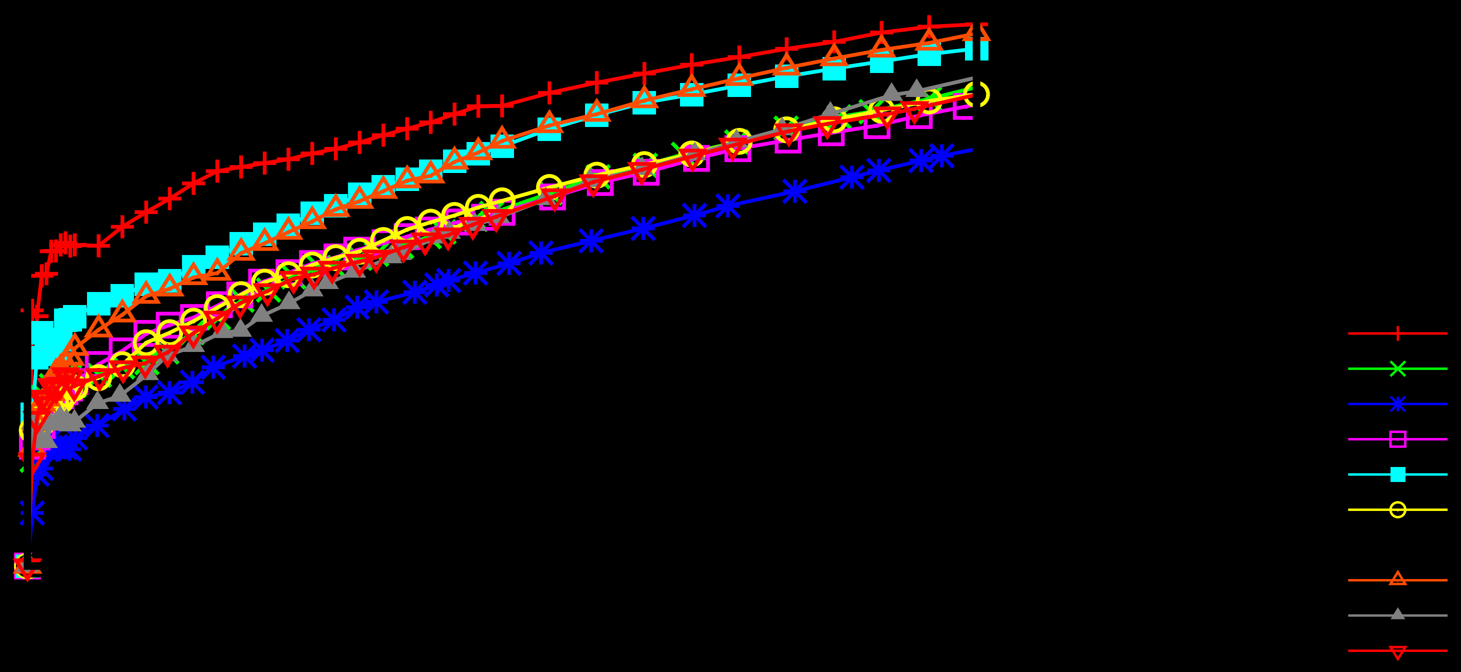

Results: JPEG compression dataset

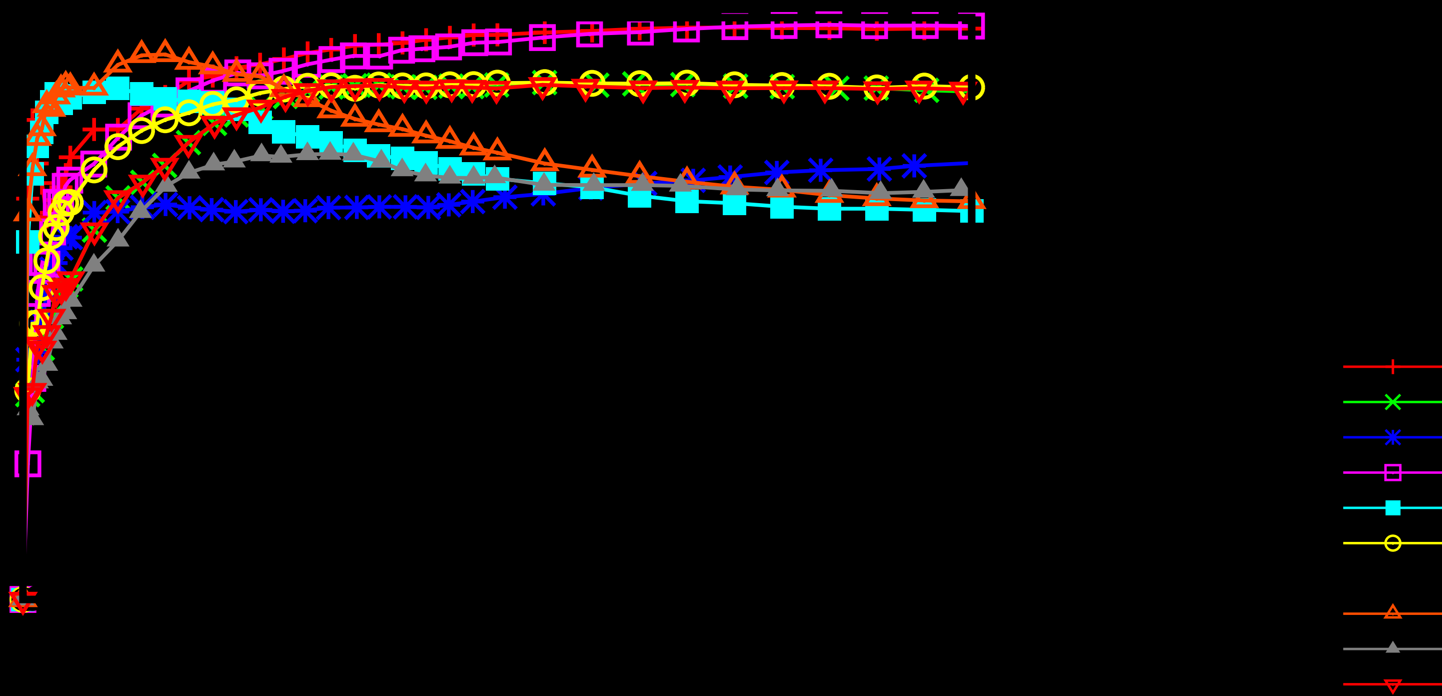

Results: Perspective (wall) dataset