Comparative study of transformer robustness for multiple particle tracking without clutter

Piyush MISHRA
I2M & Inst. Fresnel

Sup: Philippe ROUDOT

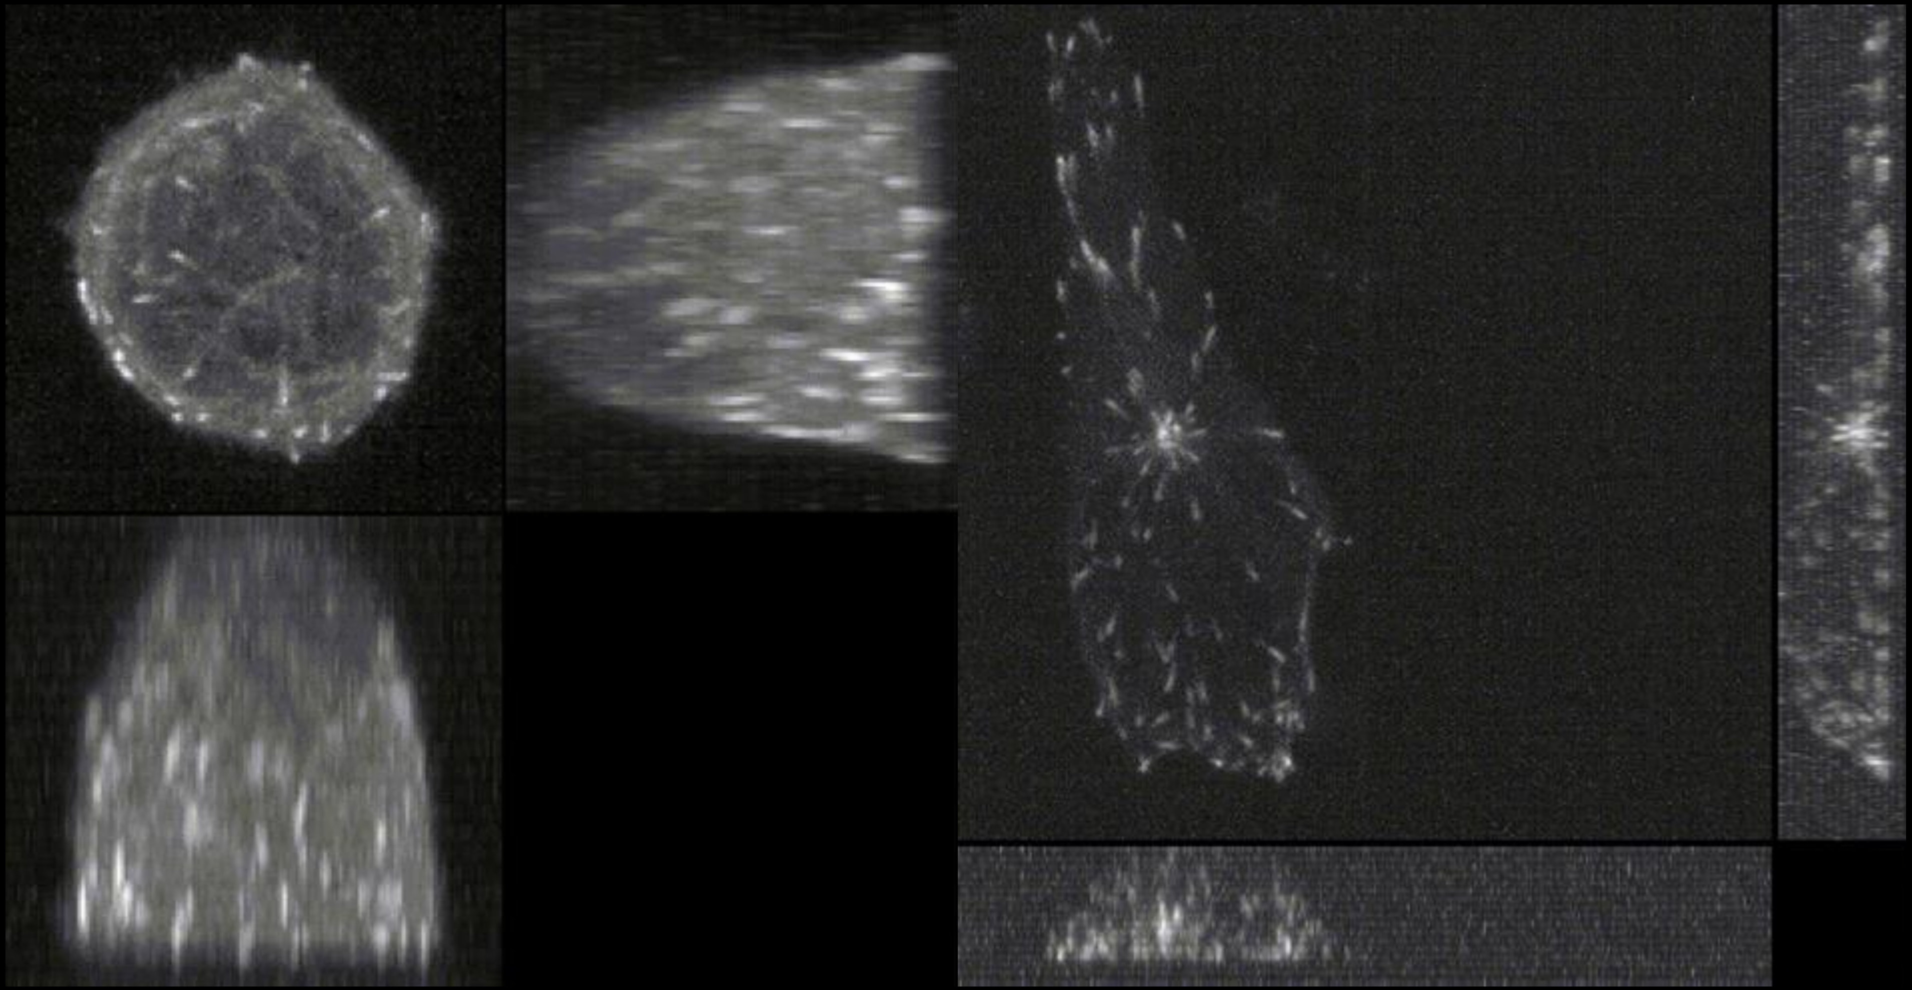

G.D.C Ortega,
ETH Zurich, 2023

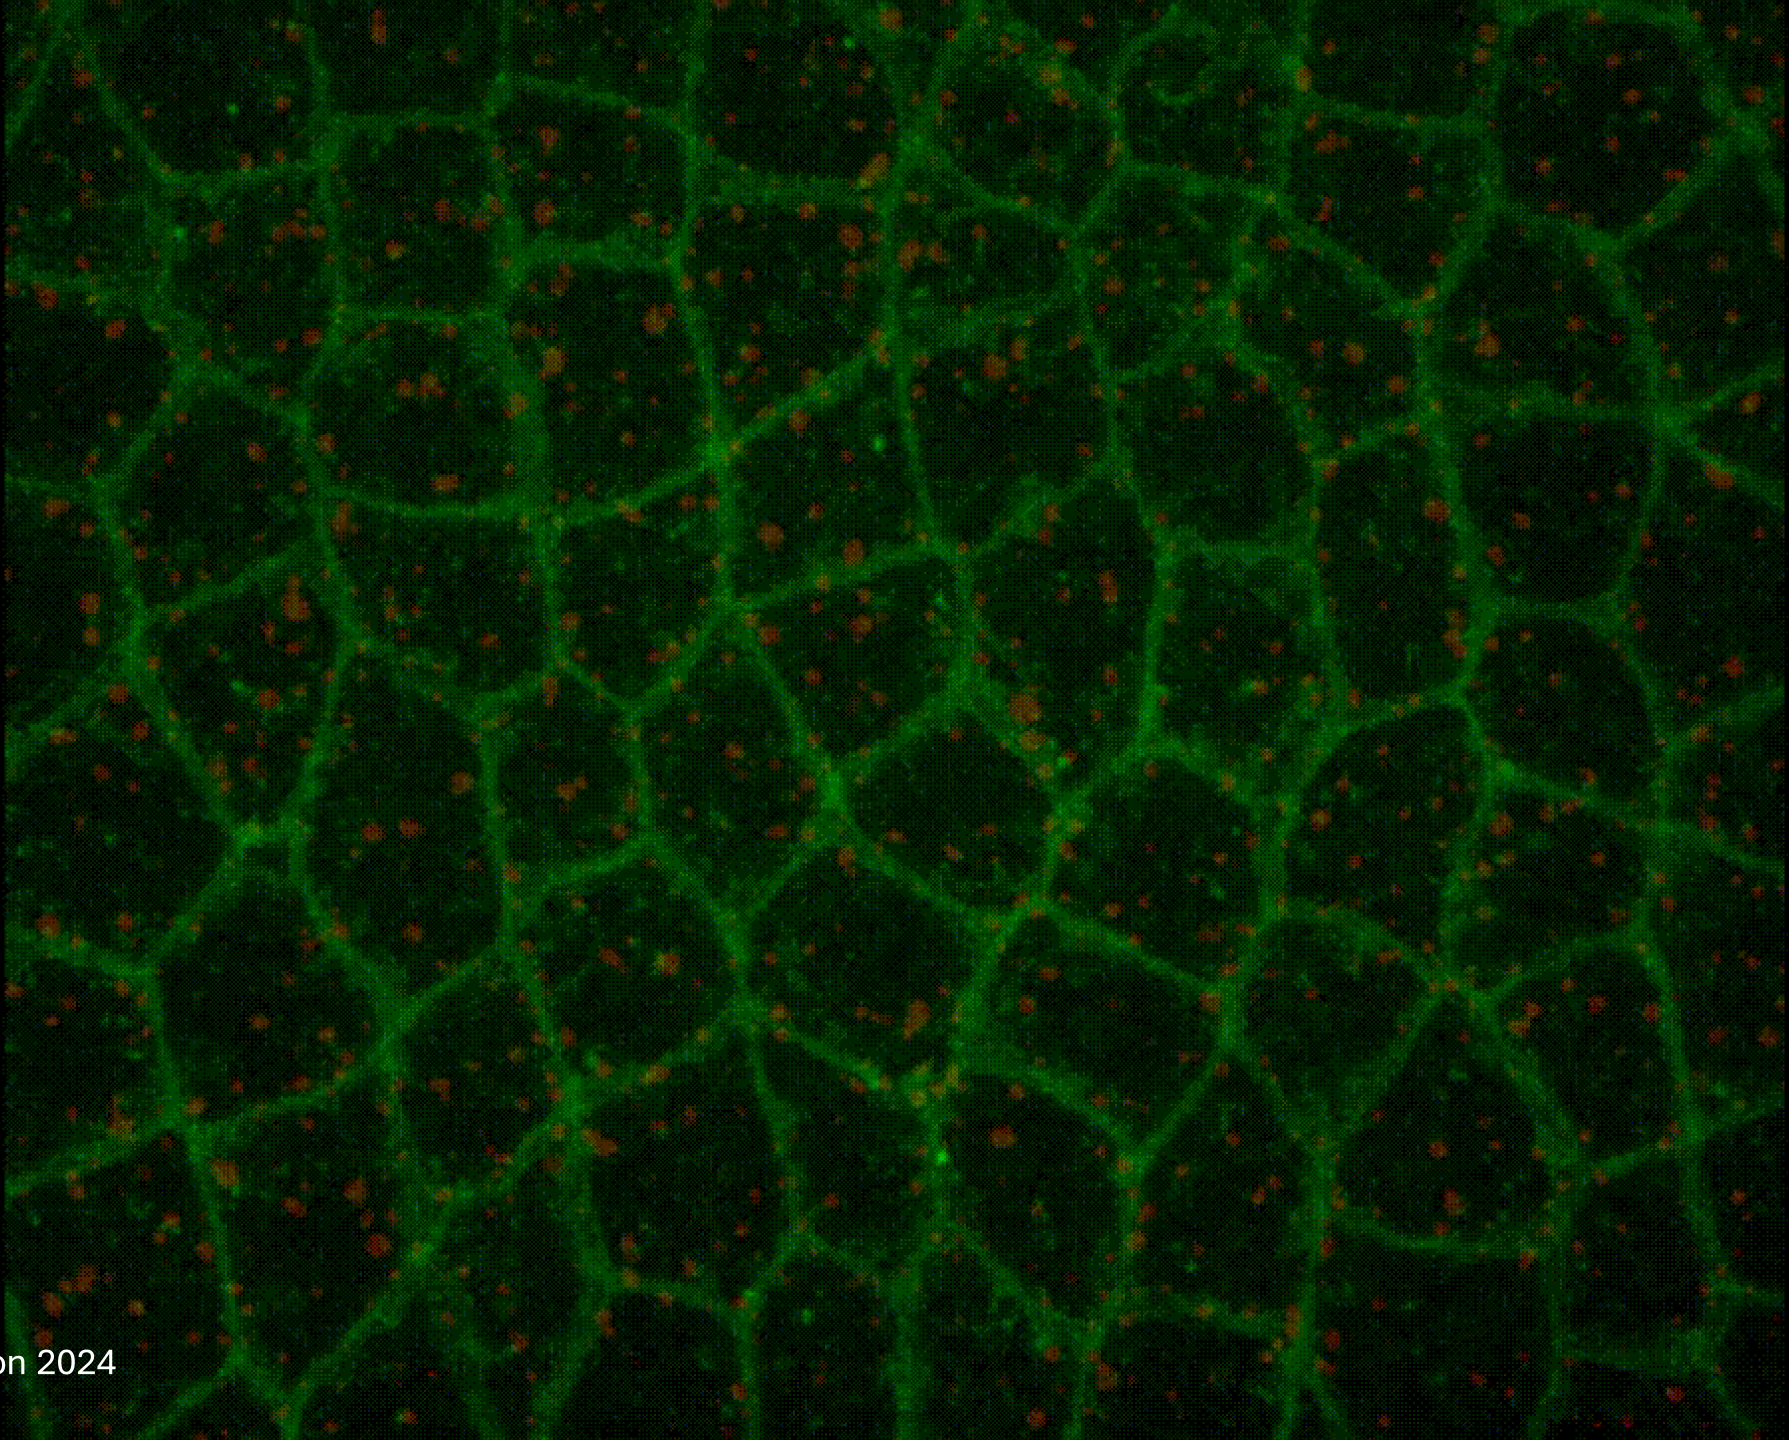

Team Endotrack,
Centuri Hackathon 2024

Mapping measurements to states is an inverse problem of data-association

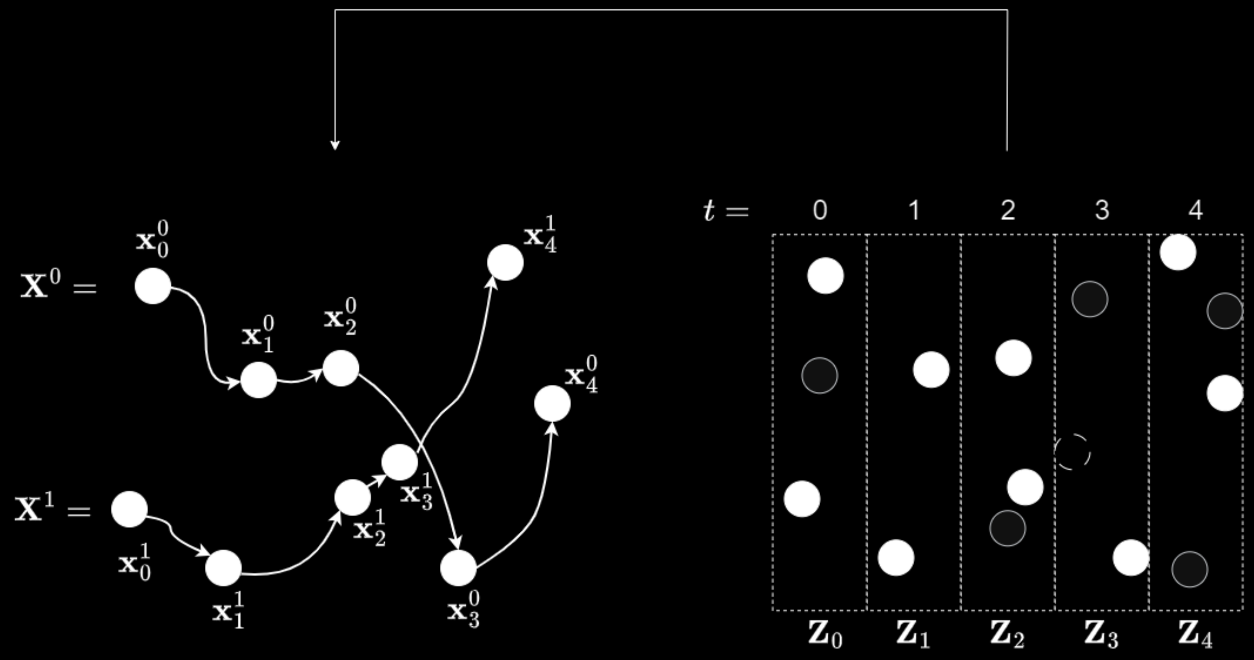

$\mathbf{X} = \{\mathbf{X}^{\mathrm{P}}\}_{p=0:N-1}$

states

$\mathbf{Z} = \{\mathbf{Z}_t\}_{t=0:T-1}$

measurements

$$\mathbf{Z} = \mathbf{\Lambda} \cdot \mathbf{X} + \epsilon$$

Data-association is a combinatorially hard problem

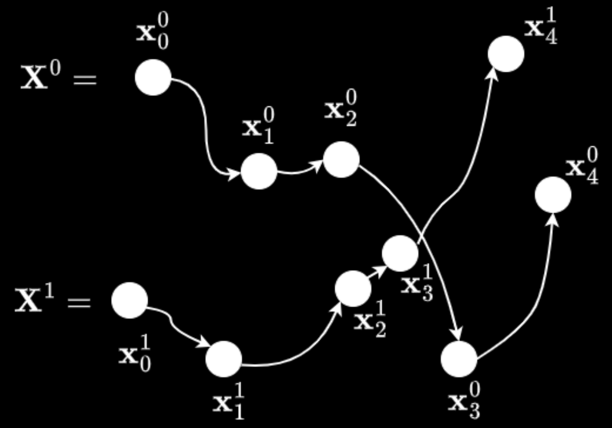

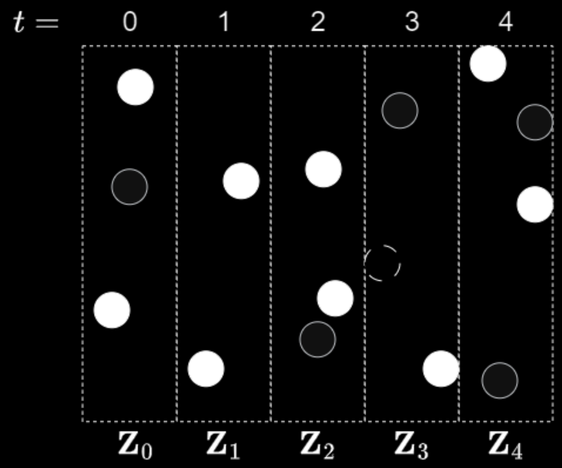

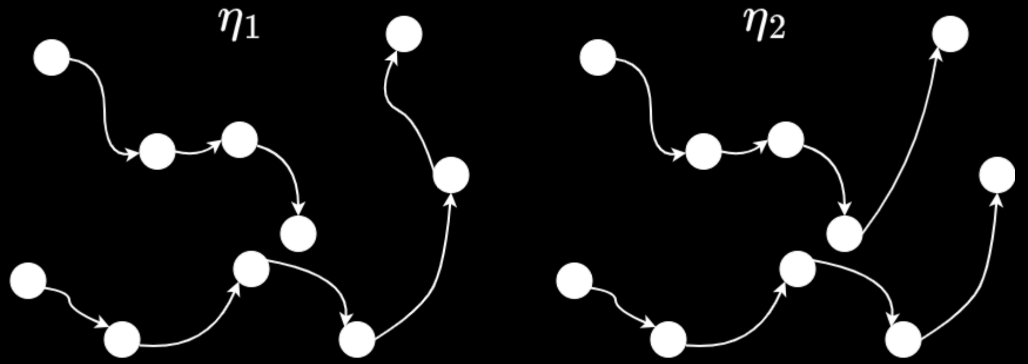

$|\eta_i| \propto N!^t$

$\mathbf{X} = \{\mathbf{X}^{\mathrm{P}}\}_{p=0:N-1}$

states

$\mathbf{Z} = \{\mathbf{Z}_t\}_{t=0:T-1}$

measurements

hypotheses

$$\mathbf{Z} = \mathbf{\Lambda} \cdot \mathbf{X} + \epsilon$$

Conventional methods use an iterative estimator as a suboptimal solution

$$p(\mathbf{X}_t|\mathbf{Z}_{1:t}) = \underbrace{p(\mathbf{Z}_t|\mathbf{X}_t)}_{\text{association}} \int \underbrace{p(\mathbf{X}_t|\mathbf{X}_{t-1})}_{\text{prediction}} \underbrace{p(\mathbf{X}_{t-1}|\mathbf{Z}_{1:t-1})}_{\text{a priori}} d\mathbf{X}$$

Conventional methods must prematurely prune hypotheses based on priors

$$p(\mathbf{X}_t|\mathbf{Z}_{1:t}) = \underbrace{p(\mathbf{Z}_t|\mathbf{X}_t)}_{\text{association}} \int \underbrace{p(\mathbf{X}_t|\mathbf{X}_{t-1})}_{\text{prediction}} \underbrace{p(\mathbf{X}_{t-1}|\mathbf{Z}_{1:t-1})}_{\text{a priori}} d\mathbf{X}$$

$$= \sum_{\boldsymbol{\eta}_p^t \in \mathbf{H}_t'} p(\mathbf{Z}_t|\mathbf{X}_t, \boldsymbol{\eta}_p^t) p(\boldsymbol{\eta}_p^t|\mathbf{X}_t) \int p(\mathbf{X}_t|\mathbf{X}_{t-1}) p(\mathbf{X}_{t-1}|\mathbf{Z}_{1:t-1}) d\mathbf{X}$$

Attention can be used to make decisions on both states & hypotheses

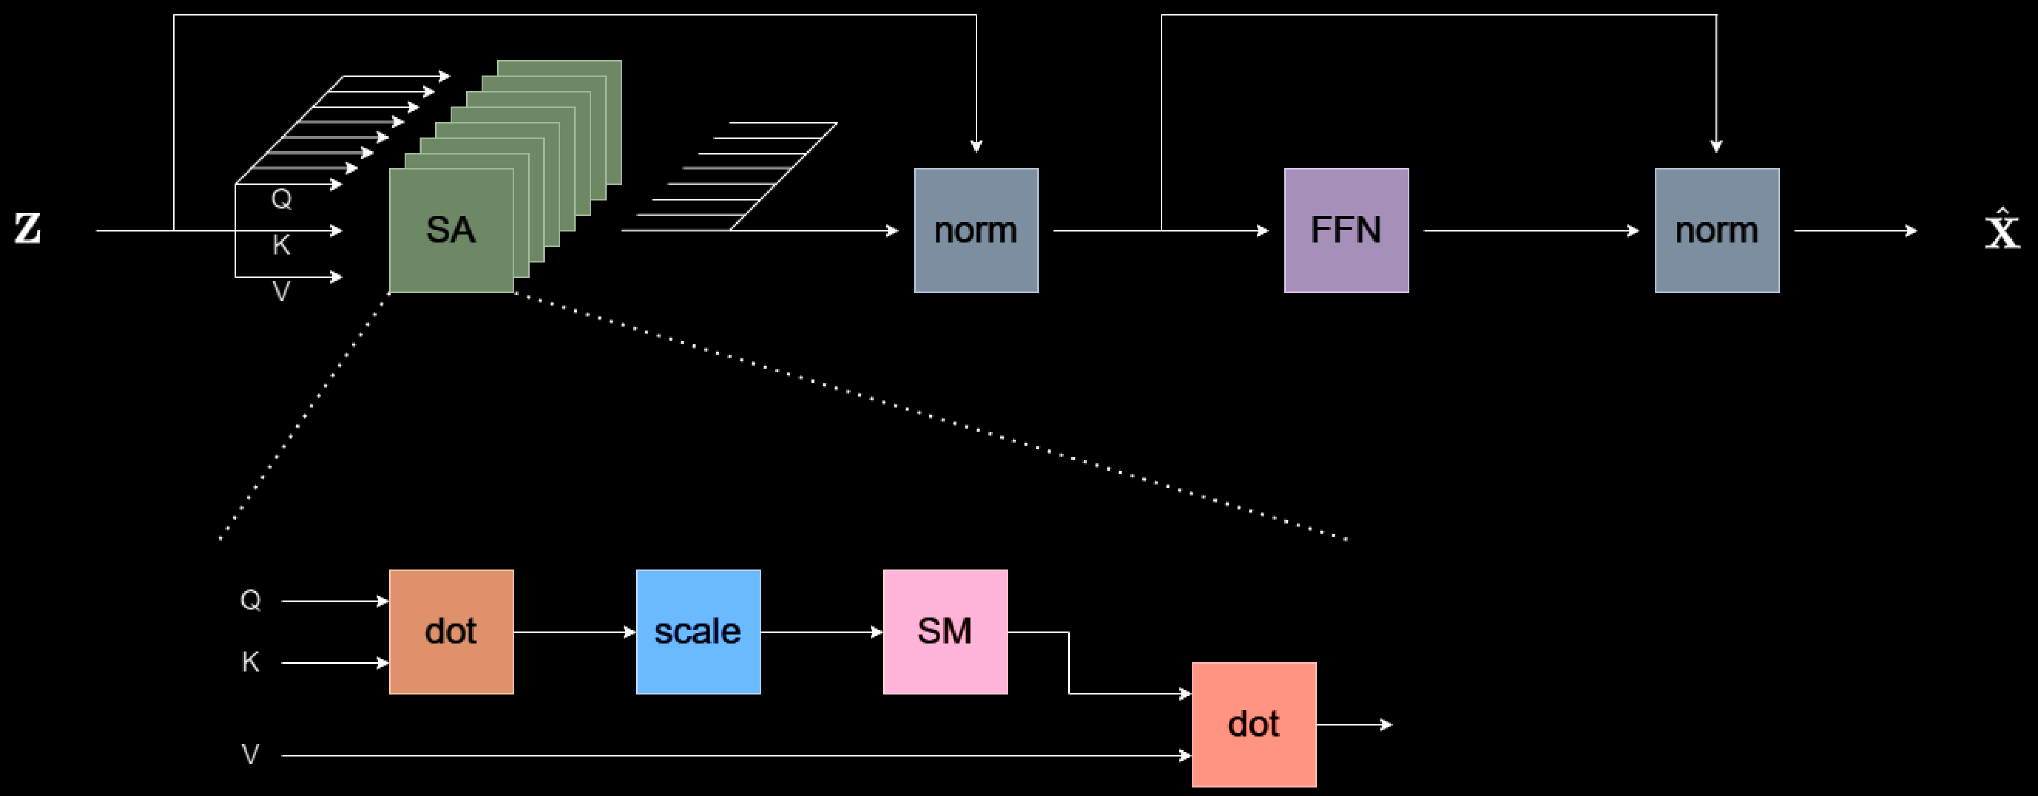

$$\hat{\mathbf{X}} = \mathrm{concat}\left(\mathrm{softmax}\left(\frac{\mathbf{W}_{q,m}\mathbf{Z}_m(\mathbf{W}_{k,m}\mathbf{Z}_m)^T}{\sqrt{d_k}}\right)\mathbf{W}_{v,m}\mathbf{Z}_m, \forall m \leq n_h\right) \cdot \mathbf{W}_l$$

A simple experimental setup for proof-of-concept

$$y^{t,p} = y^{t-1,p} + \varepsilon^{t,p} + \delta^{p}$$

drift

randomness
with process
noise

$$z^{t,p} = y^{t,p} + \omega^{p}$$

measurement noise

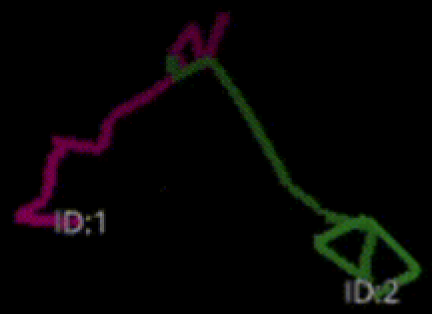

Mishra, Roudot, 2024

Similarity between ground-truth and prediction is given by Jaccard coefficient

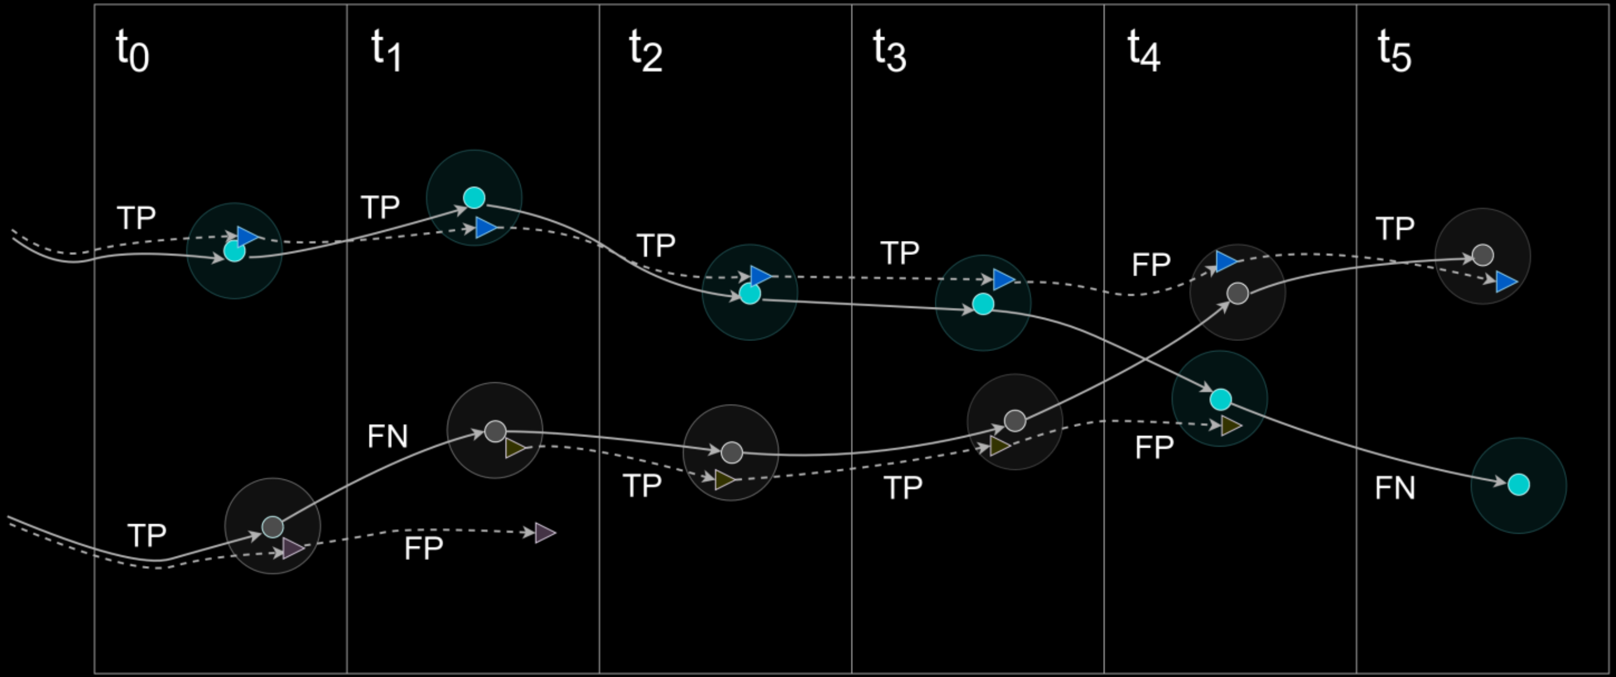

Illustration by Laura Neschen

Attention is robust to increasing noise in long sequences

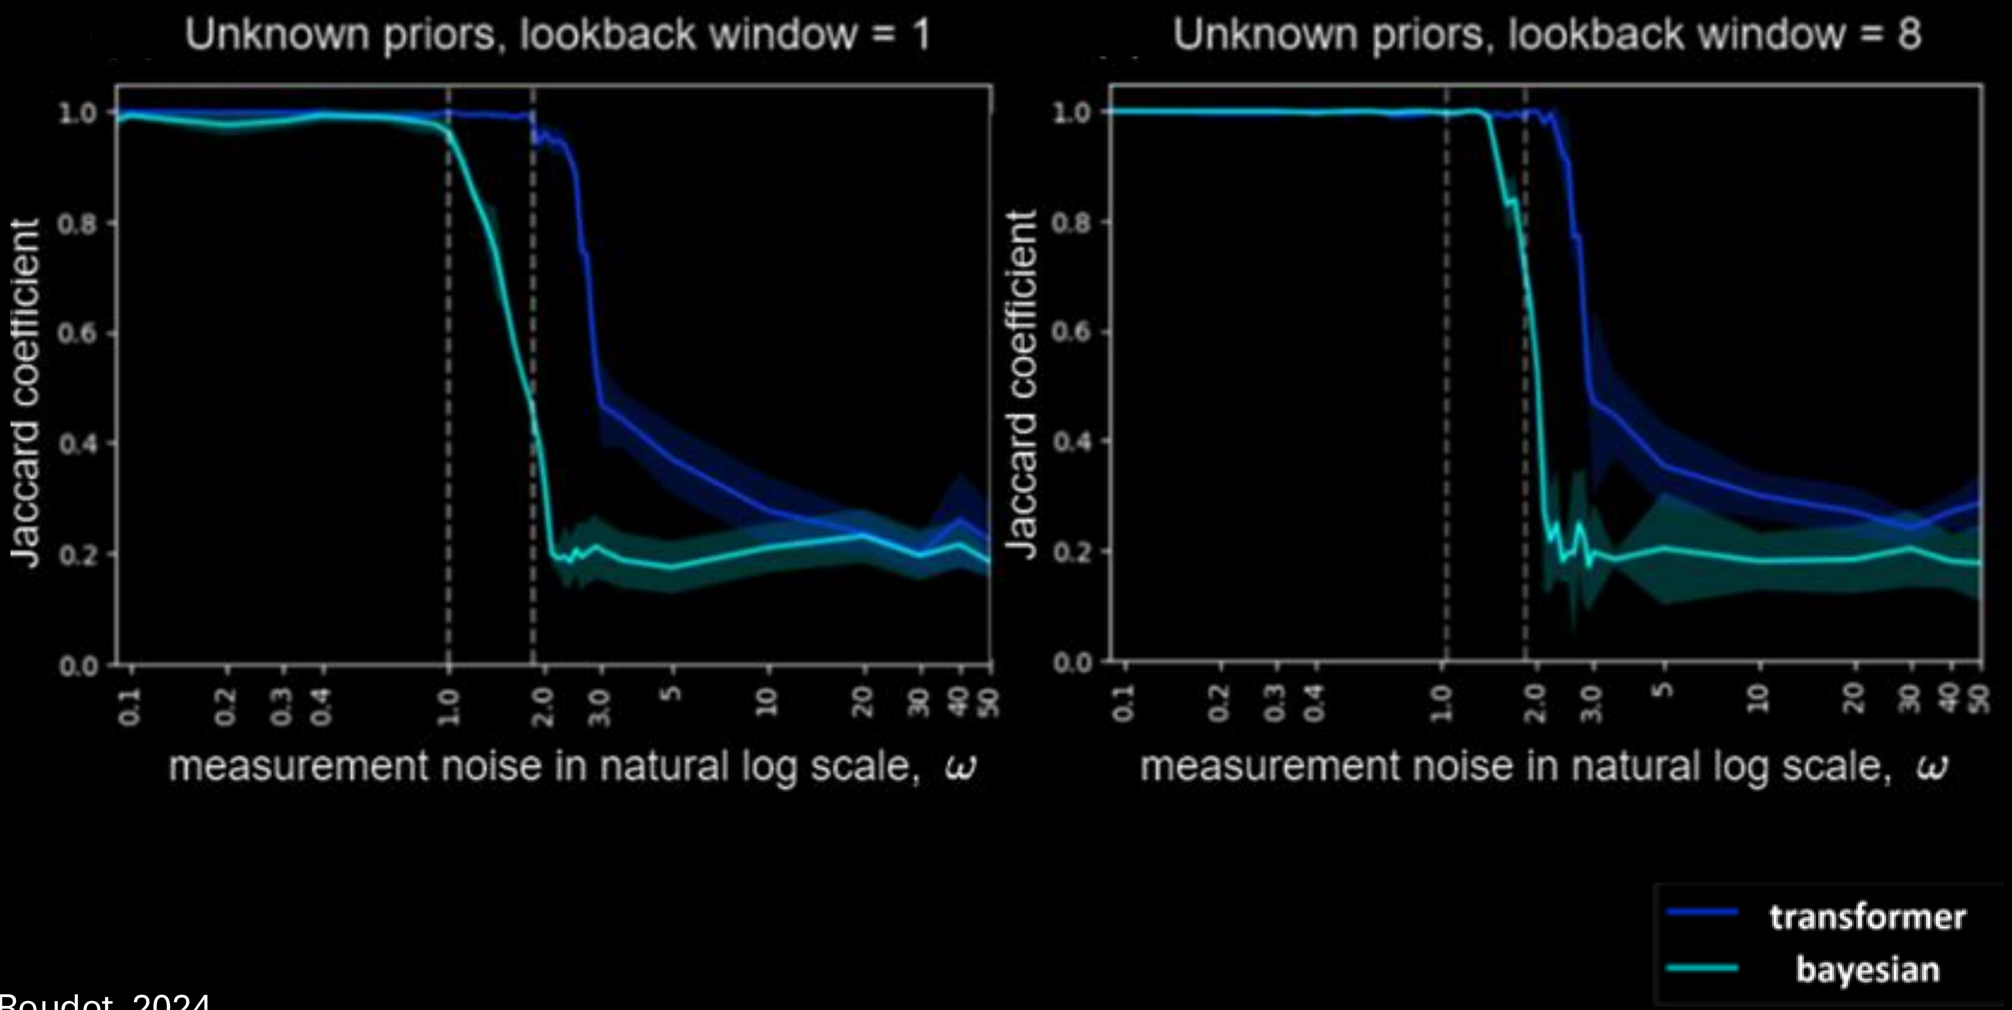

Mishra, Roudot, 2024

When Bayesian filtering is optimal, attention is suboptimal

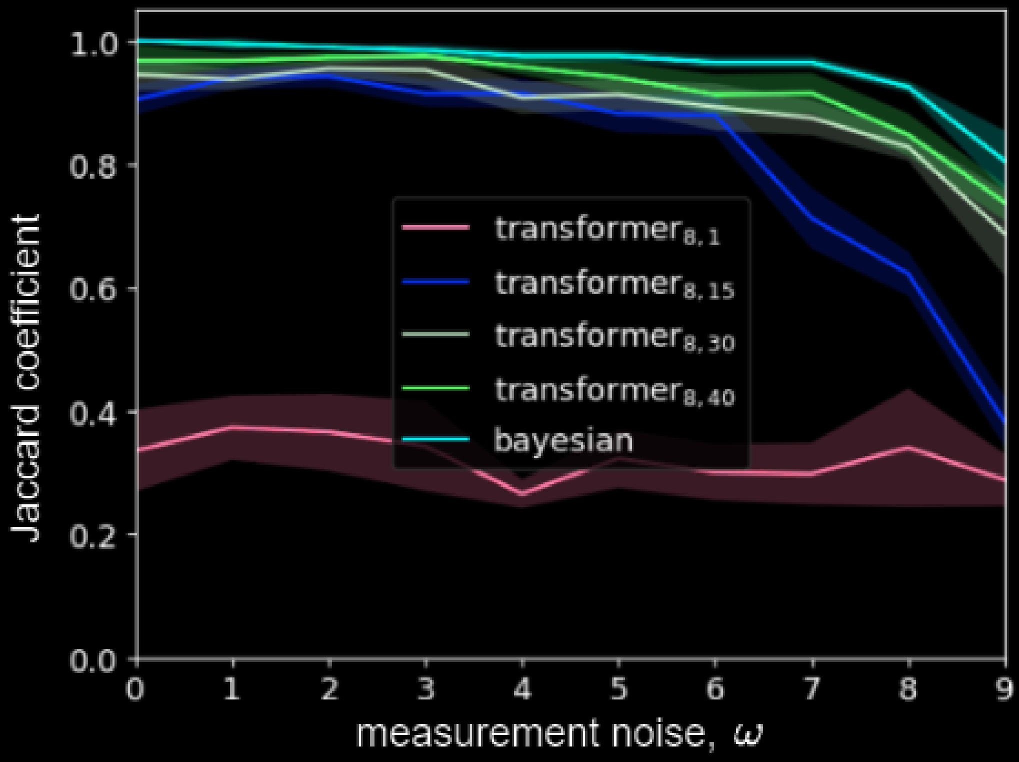

Mishra, Roudot, 2024

When Bayesian filtering is optimal, attention remains suboptimal

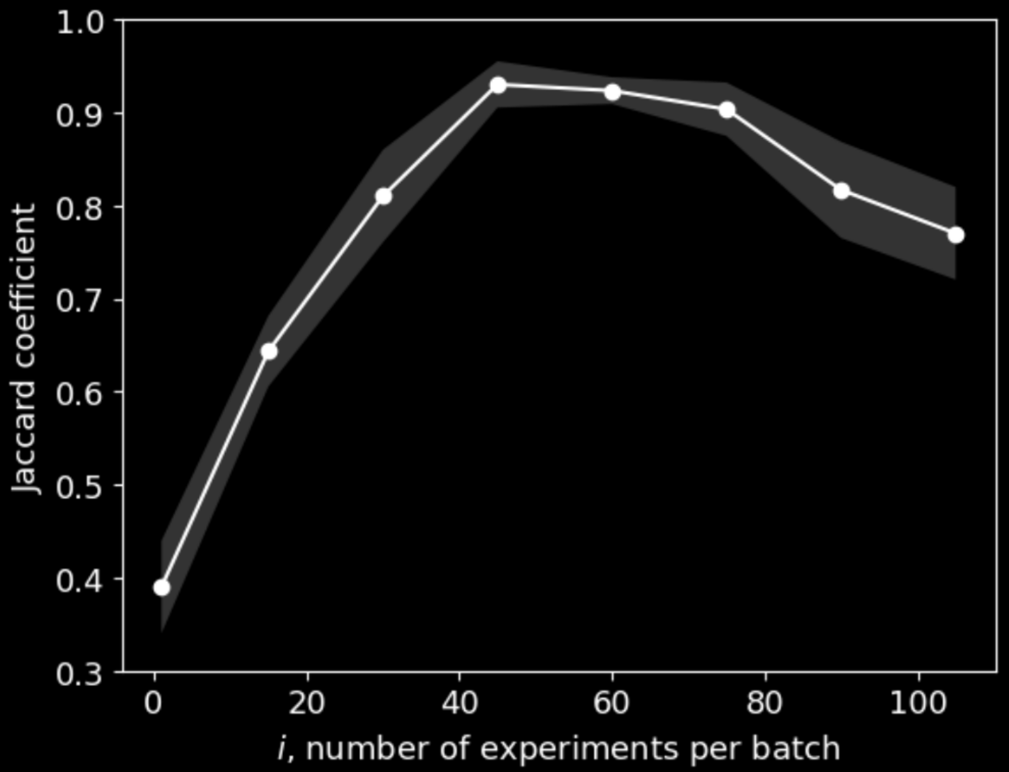

Mishra, Roudot, 2024

Attention is more efficient when increasing the lookback window

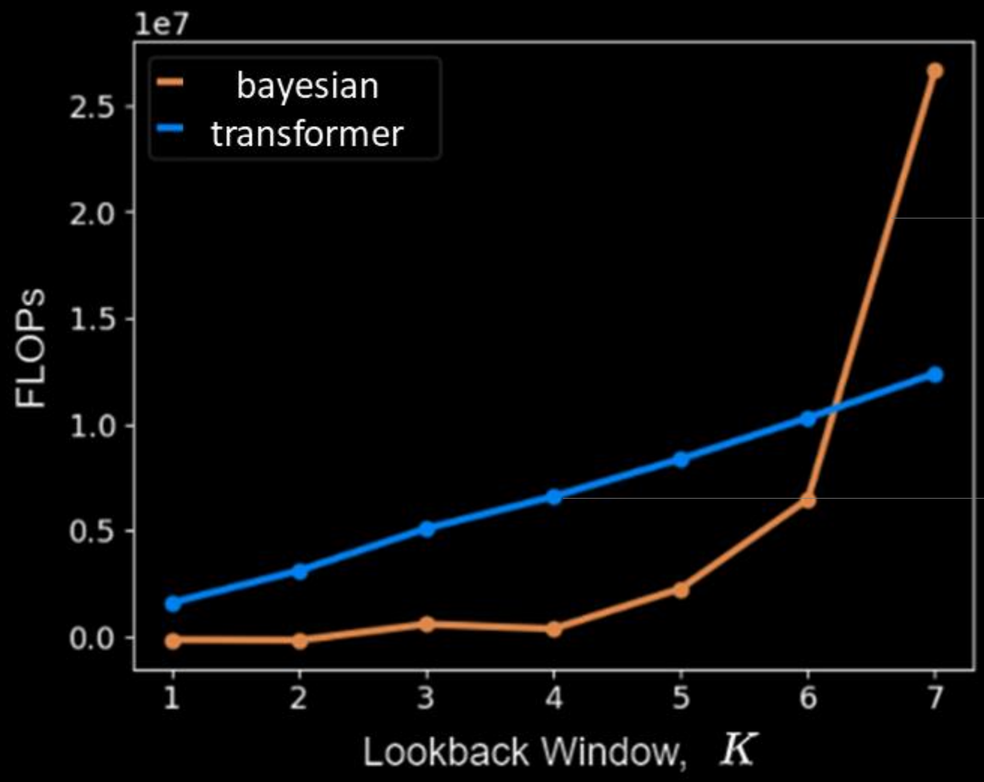

$O(N!^K)$

Reid, 1979

$O(K)$

Mishra, Roudot, 2024

Attention is robust to increasing sequence length

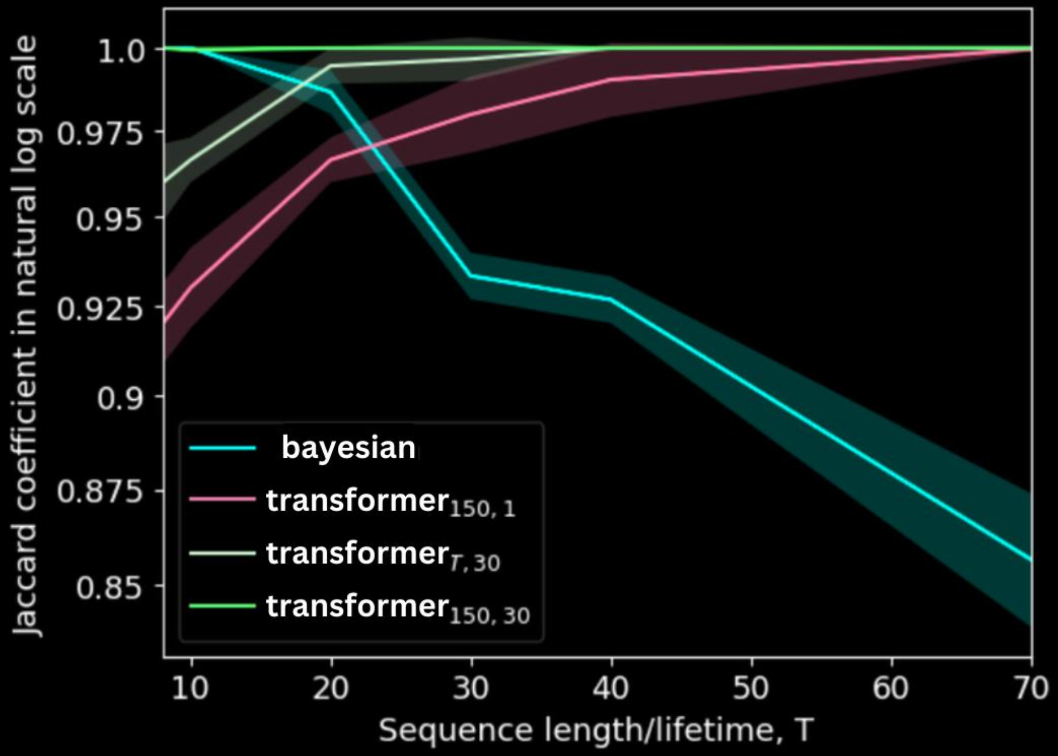

Mishra, Roudot, 2024

Ongoing work: A frugal tracking strategy that uses attention to build global priors

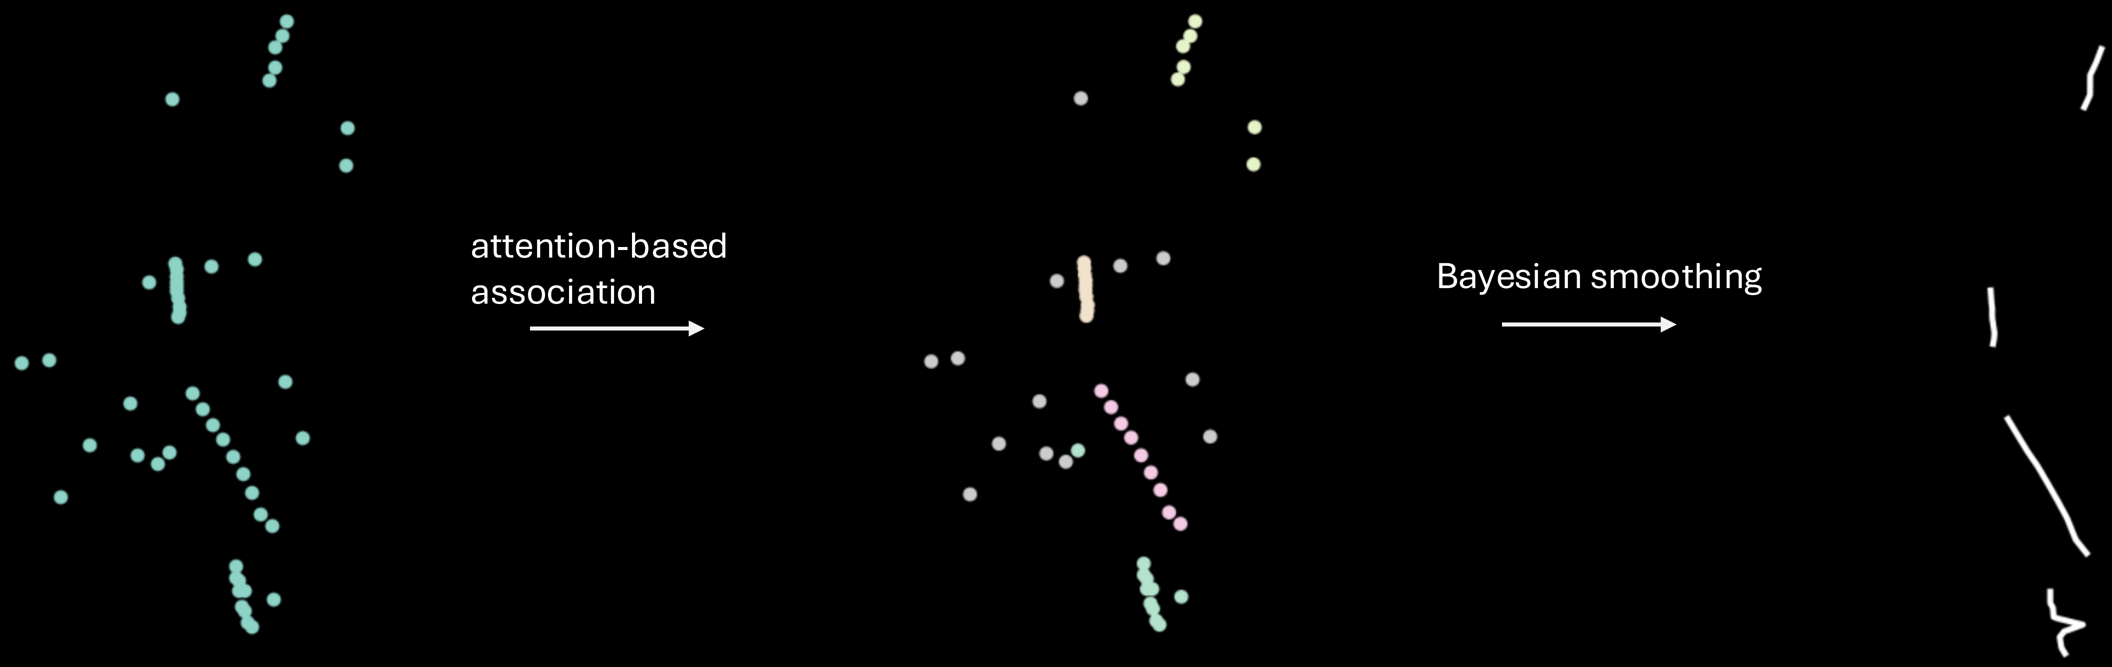

Applications & preliminary results: Stand on a moving cell

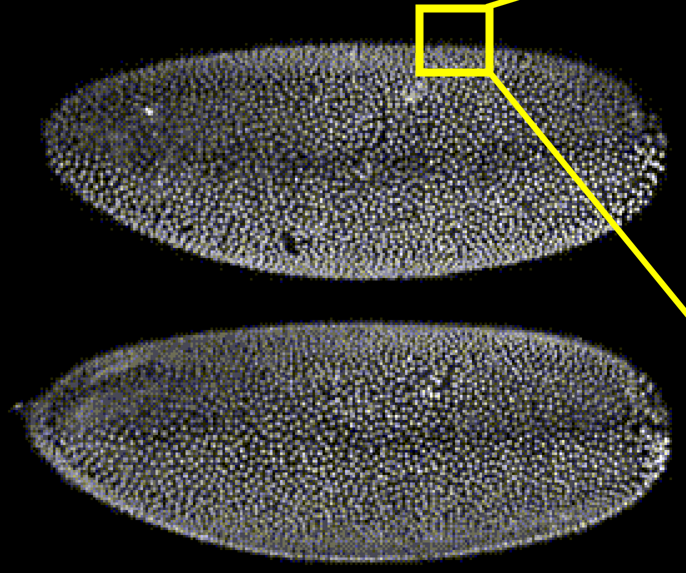

Microscopy images
of fruitfly embryo,
C. Collinet, IBDM

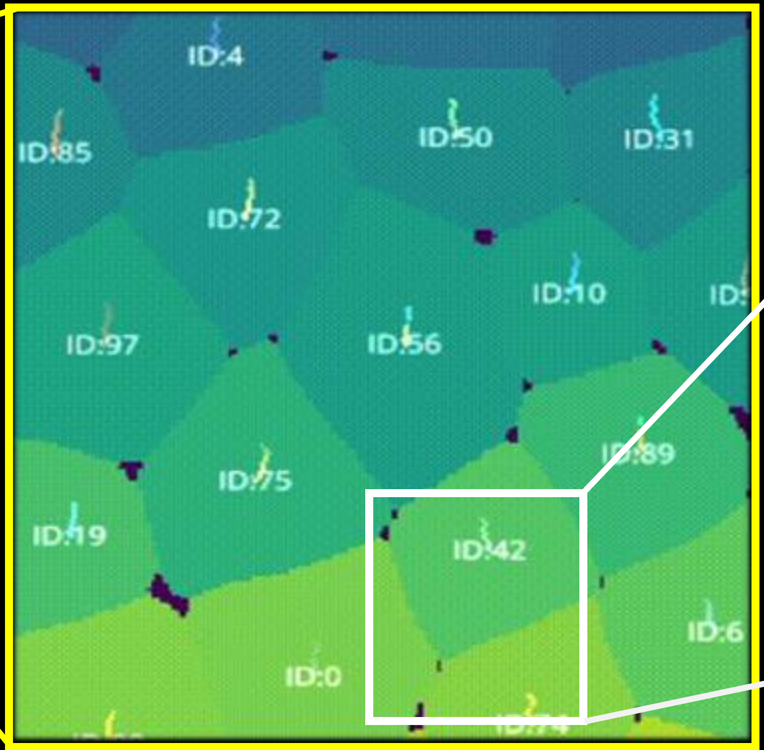

Tracking cells using
the Bayesian-Attention
hybrid strategy

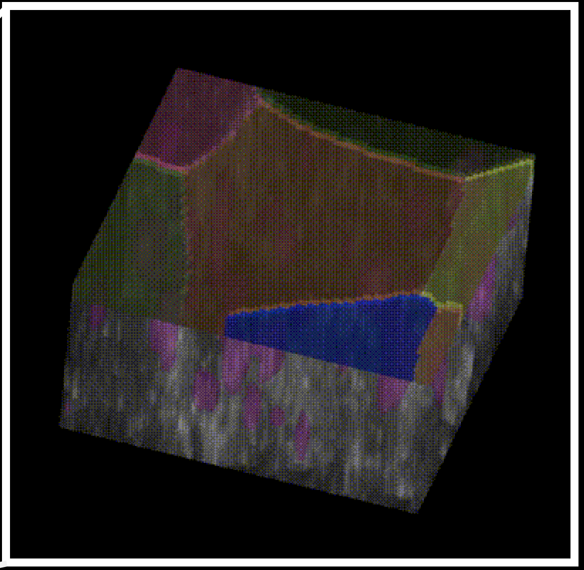

Stabilised region of interest
Team Endotrack
Centuri Hackathon, 2024

- why track?

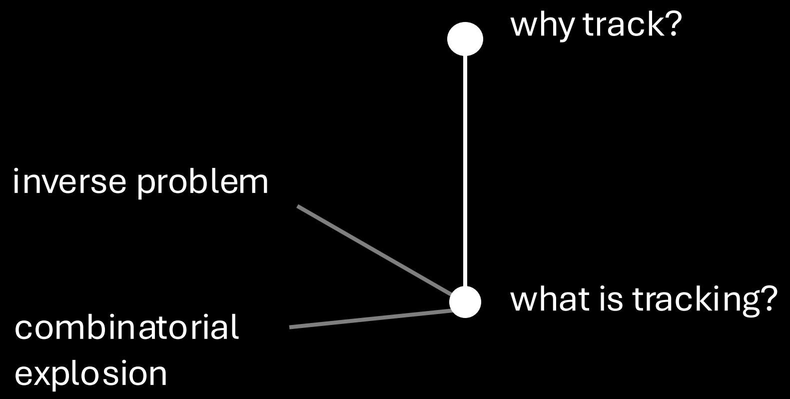
why track?
inverse problem
what is tracking?
combinatorial explosion

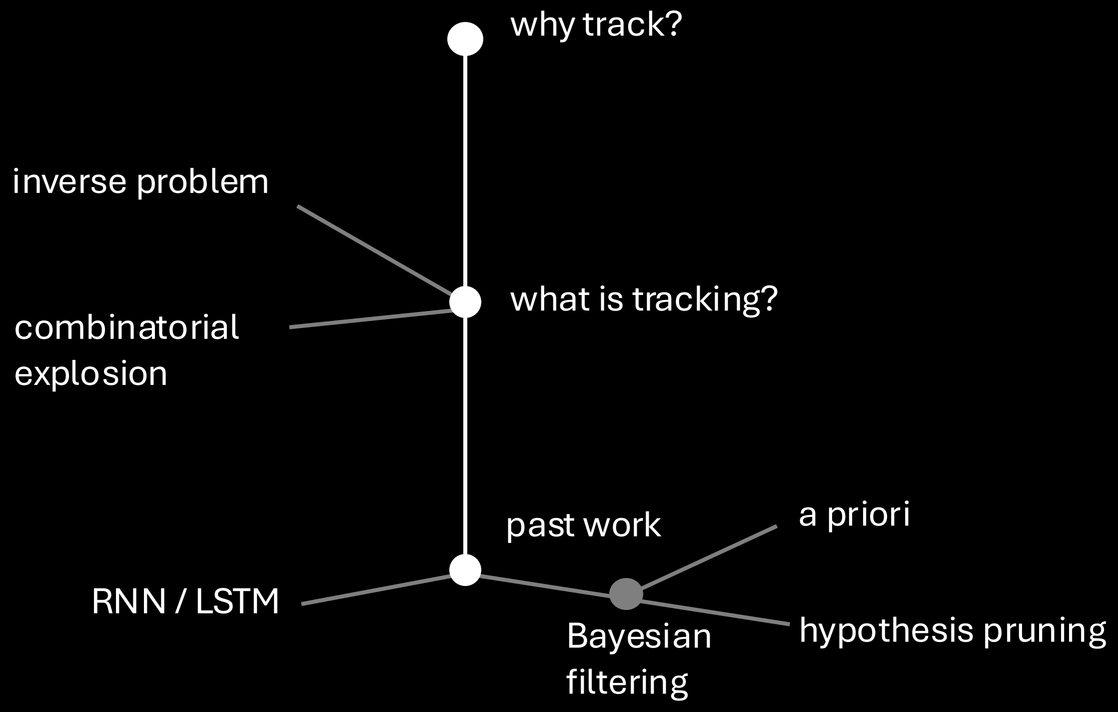
why track?
inverse problem
what is tracking?
combinatorial explosion
past work
a priori
RNN / LSTM
hypothesis pruning
Bayesian filtering

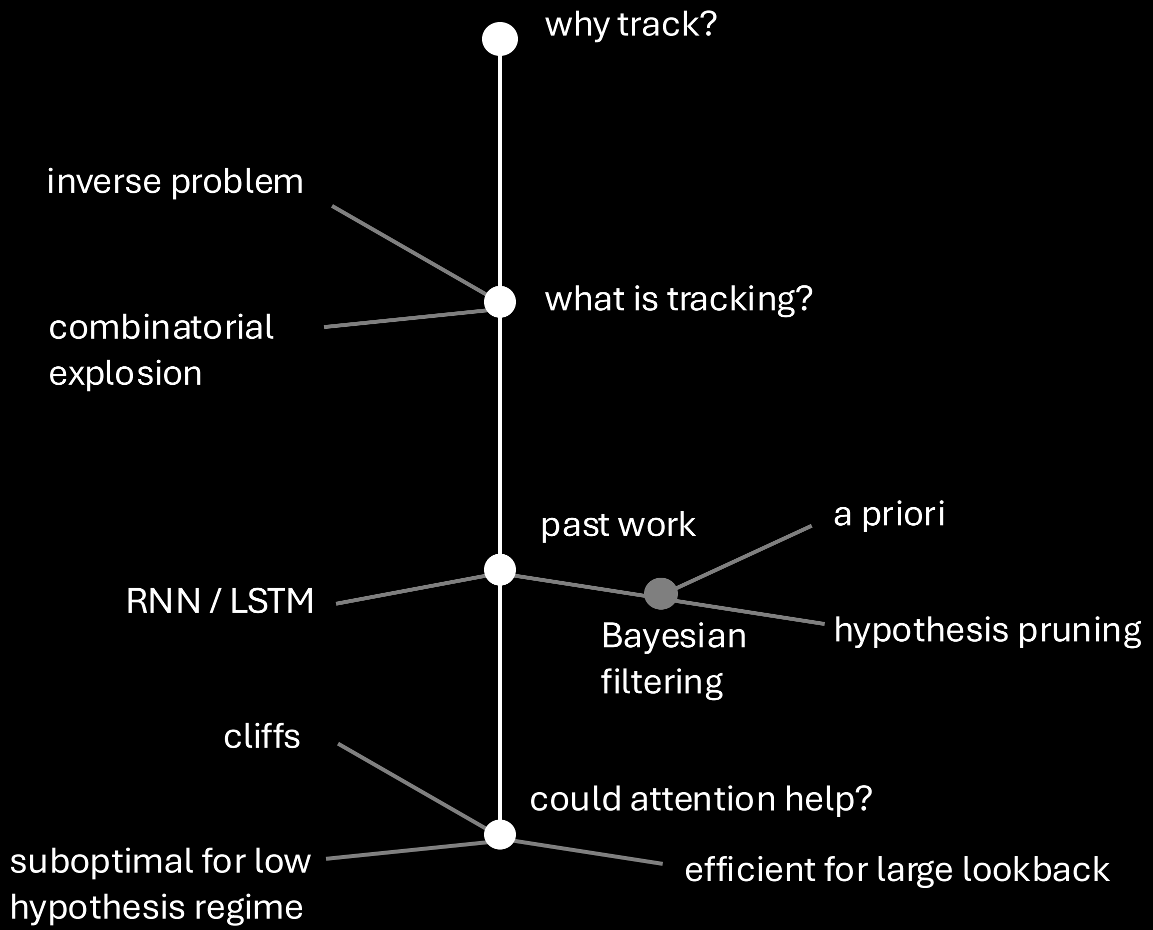
why track?
inverse problem
what is tracking?
combinatorial explosion
past work
a priori
RNN / LSTM
hypothesis pruning
Bayesian filtering
cliffs
could attention help?
suboptimal for low hypothesis regime
efficient for large lookback

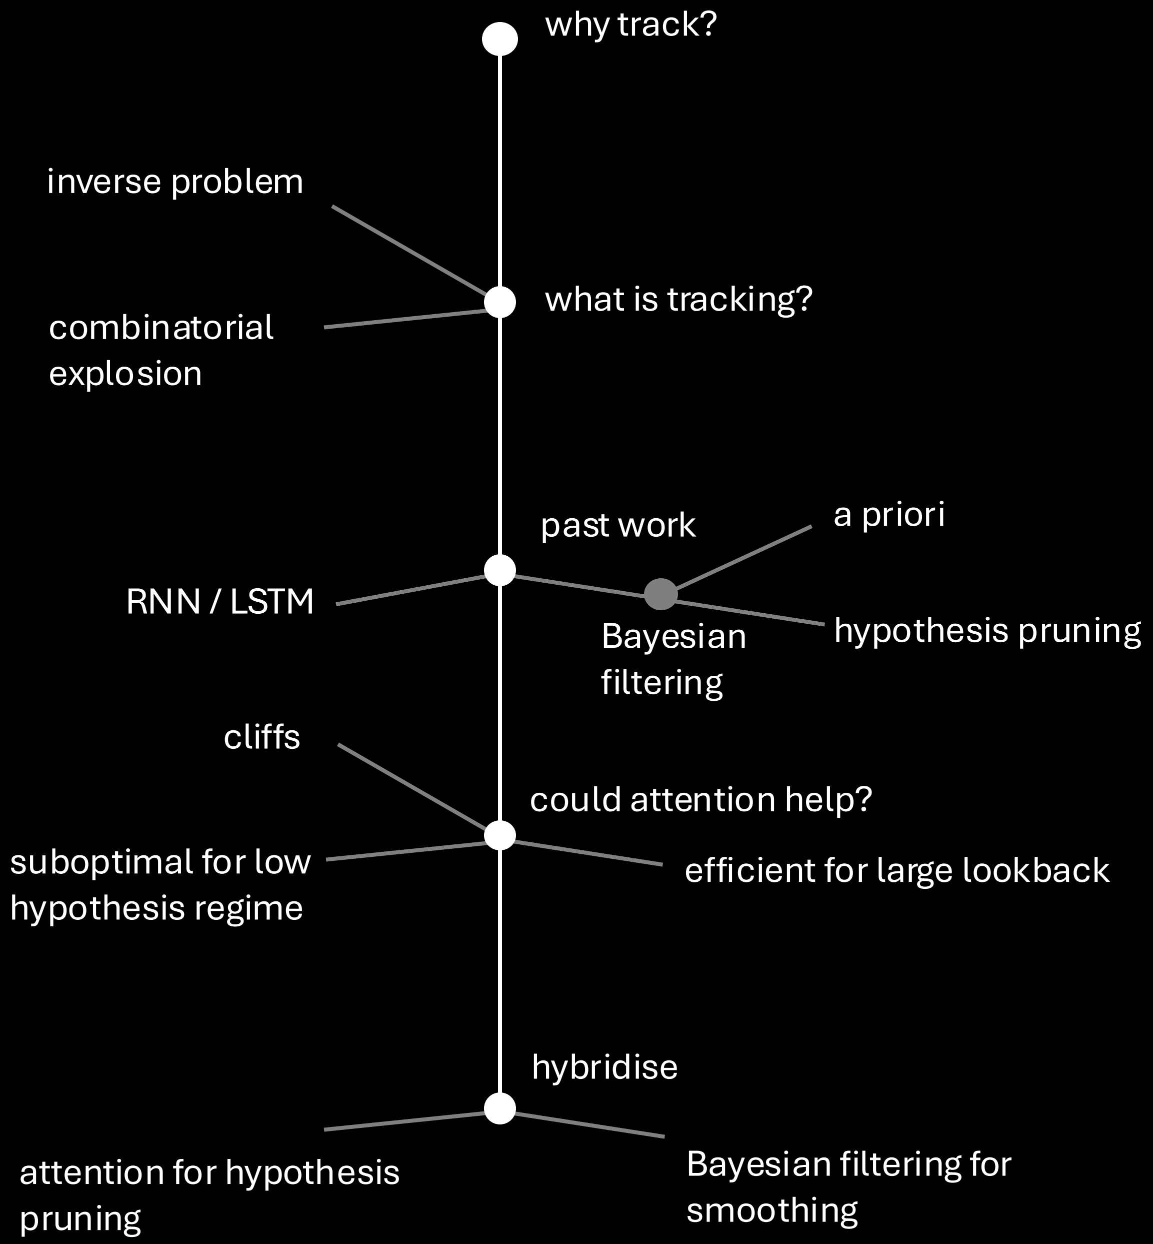
why track?
inverse problem
what is tracking?
combinatorial
explosion
past work
a priori
RNN / LSTM
Bayesian
filtering
hypothesis pruning
cliffs
could attention help?
suboptimal for low
hypothesis regime
efficient for large lookback
hybridise
attention for hypothesis
pruning
Bayesian filtering for
smoothing

Thank you for your attention ☺

Mishra, Roudot, 2024

piyushmishra12.github.io