Blind Deconvolution

Dekel Galor, Arya Raeesi

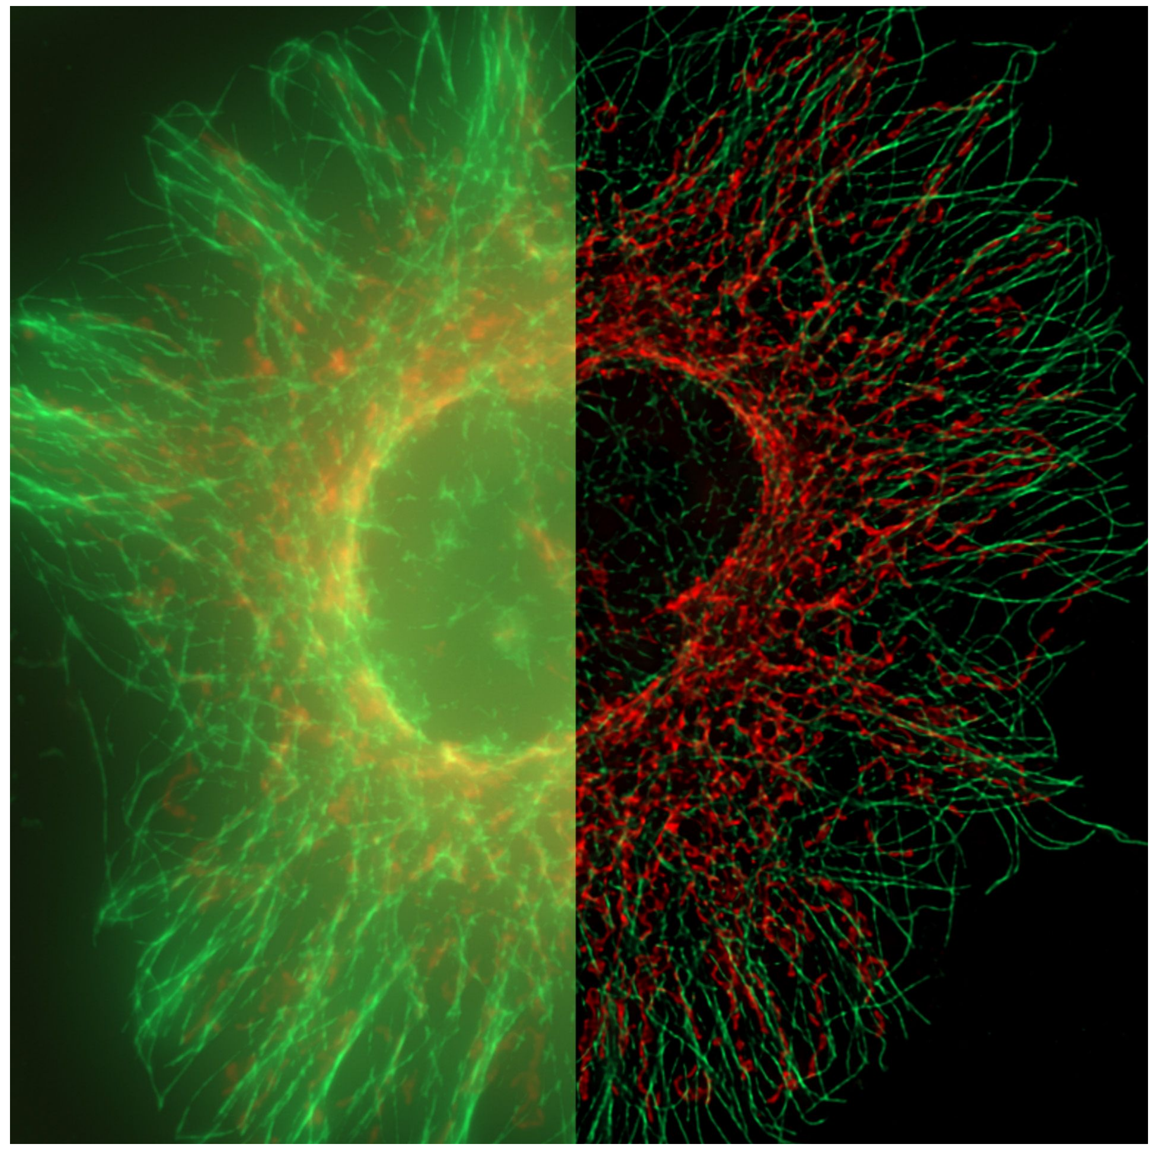

Thursday 12/04/25

Schedule

1. Motivation
2. Theory
3. Experiments

Schedule

1. **Motivation**
2. Theory
3. Experiments

Motivation

Blind Deconvolution

Reconstruction requires precise system identification.

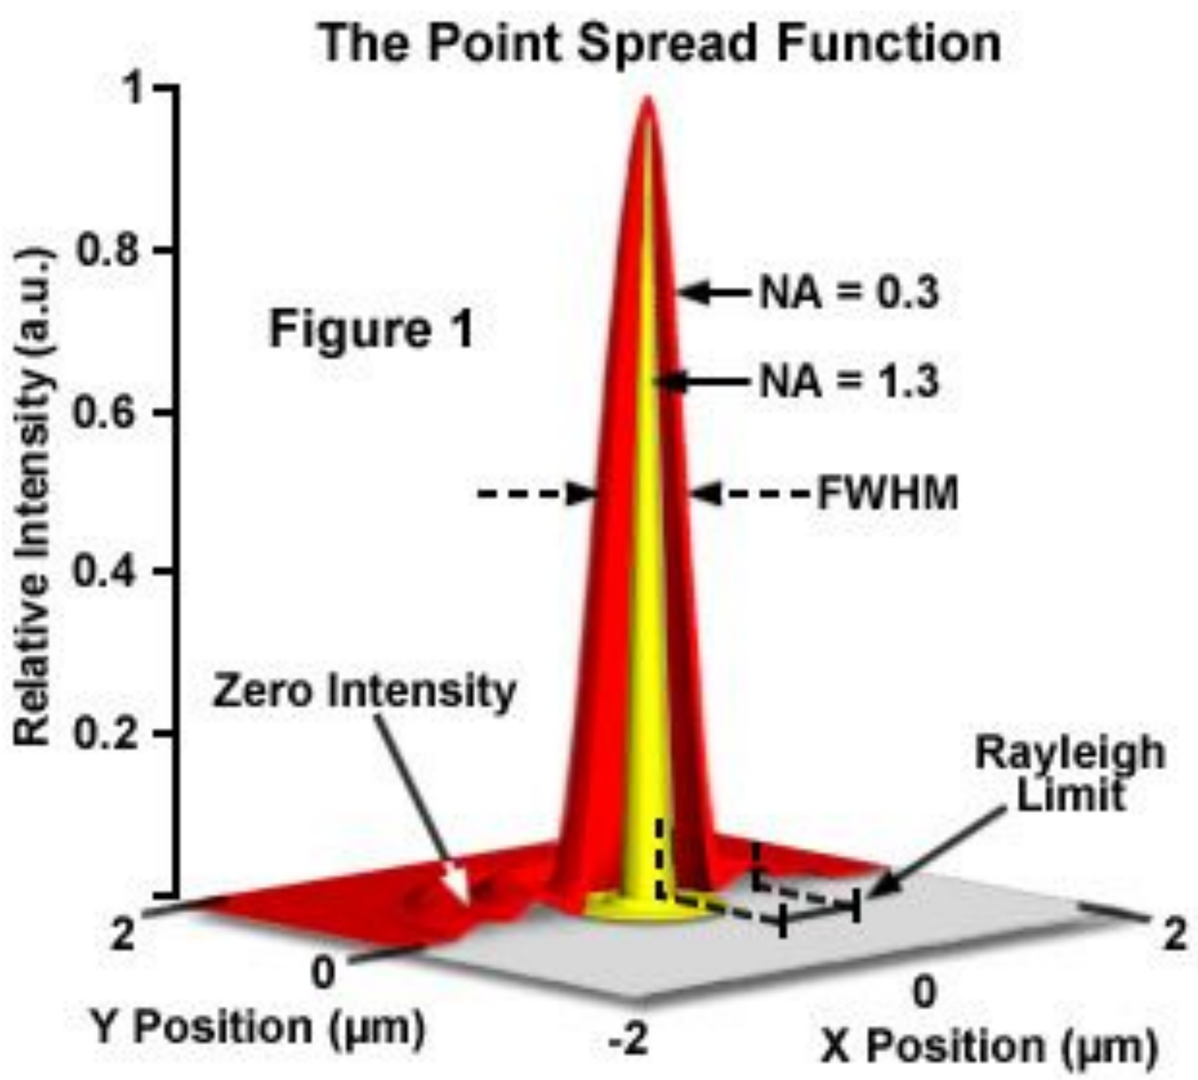

Motivation

Blind Deconvolution

Reconstruction requires precise system identification.

What if we don’t have access to the system?

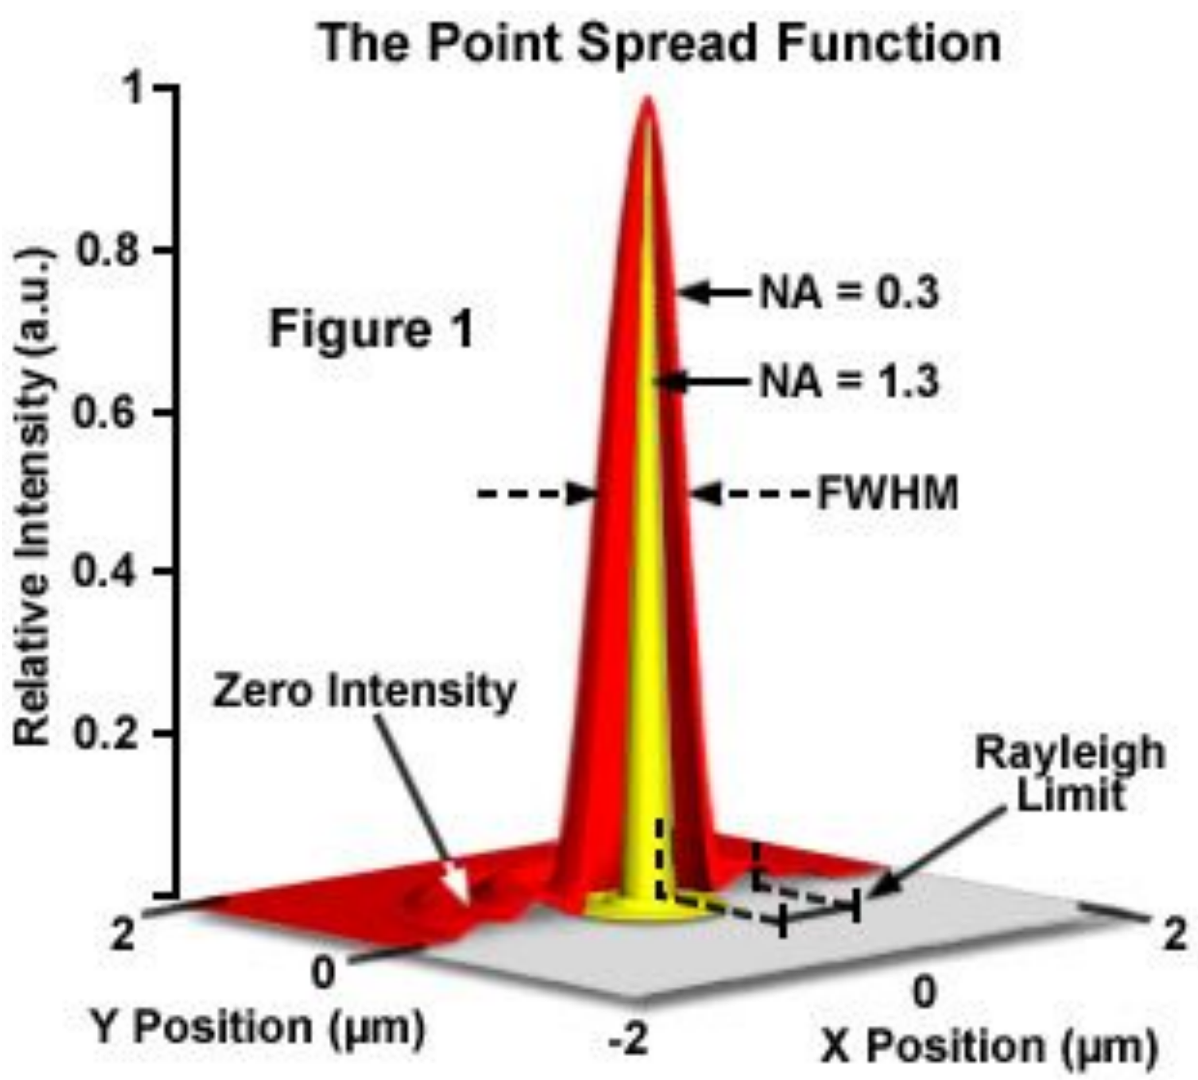

Motivation

Blind Deconvolution

Reconstruction requires precise system identification.

What if we don’t have access to the system?

Imprecise system identification

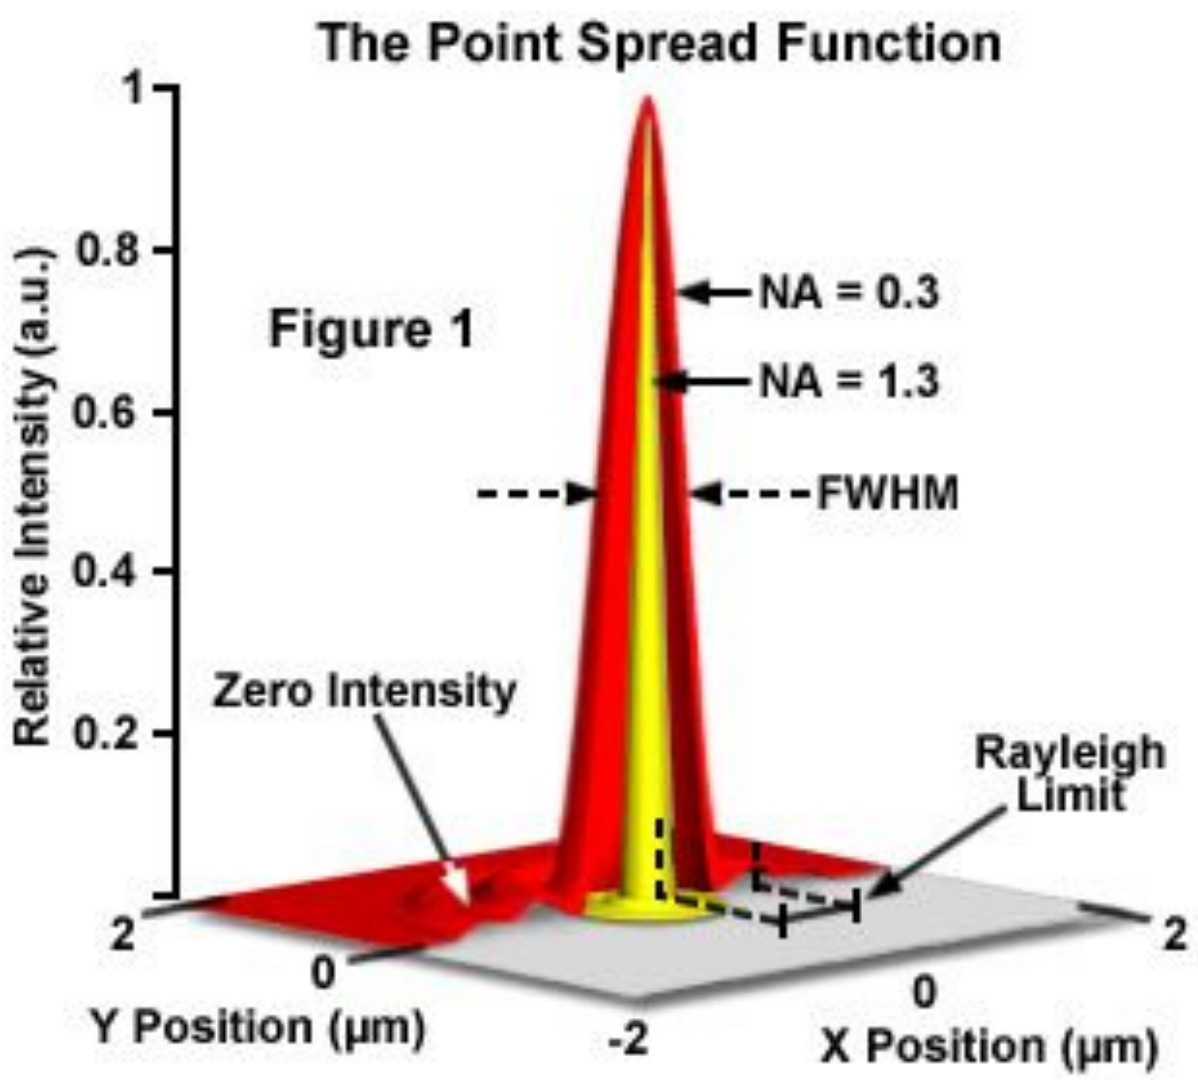

Motivation

Blind Deconvolution

Reconstruction requires precise system identification.

What if we don't have access to the system?

Imprecise system identification

No system identification!

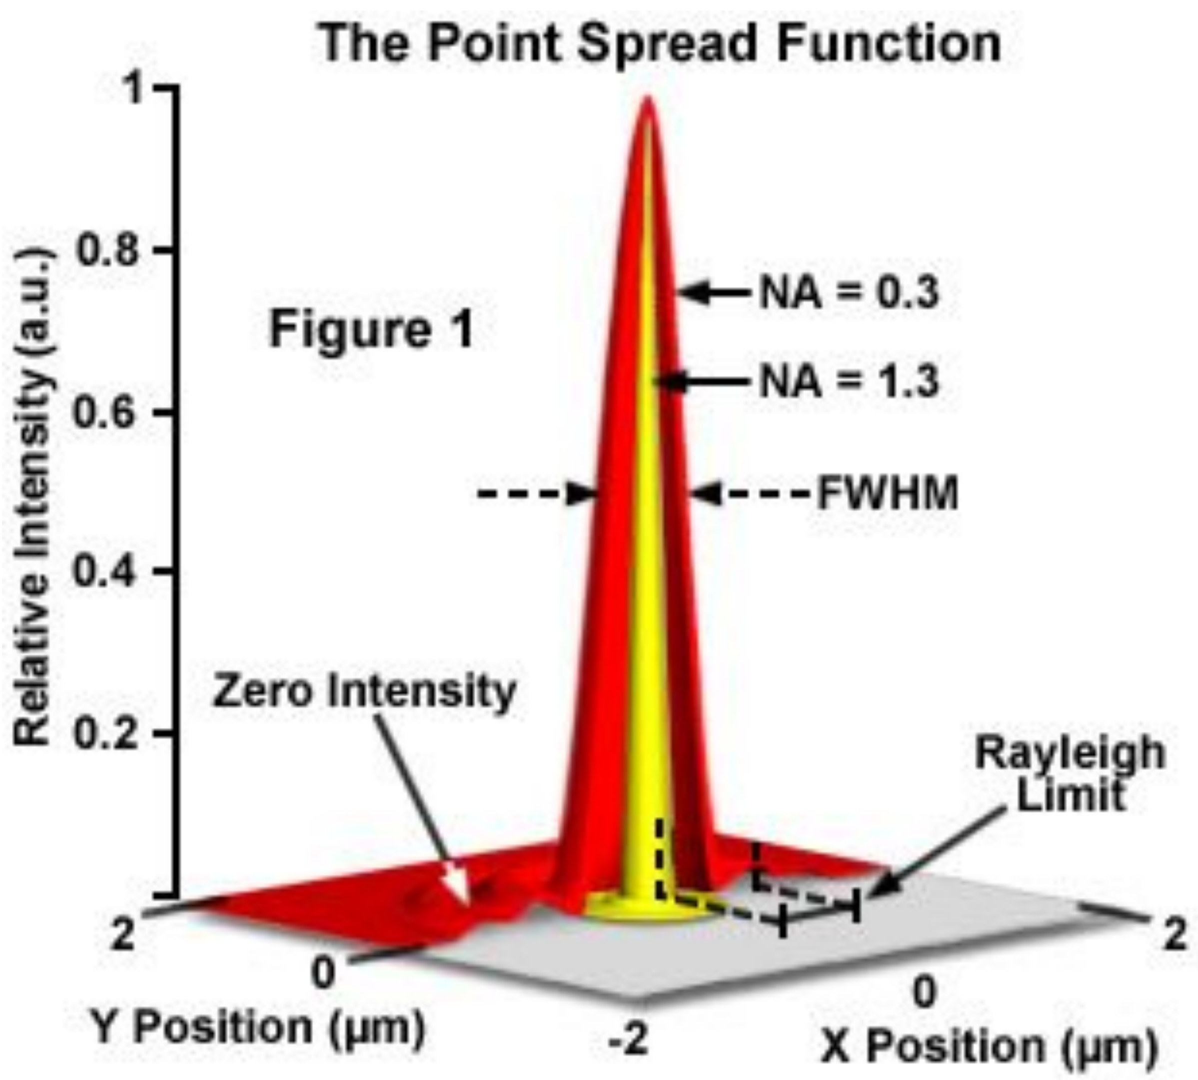

Motivation

Blind Deconvolution

Reconstruction requires precise system identification.

What if we don't have access to the system?

Imprecise system identification

No system identification!

Paradoxical?

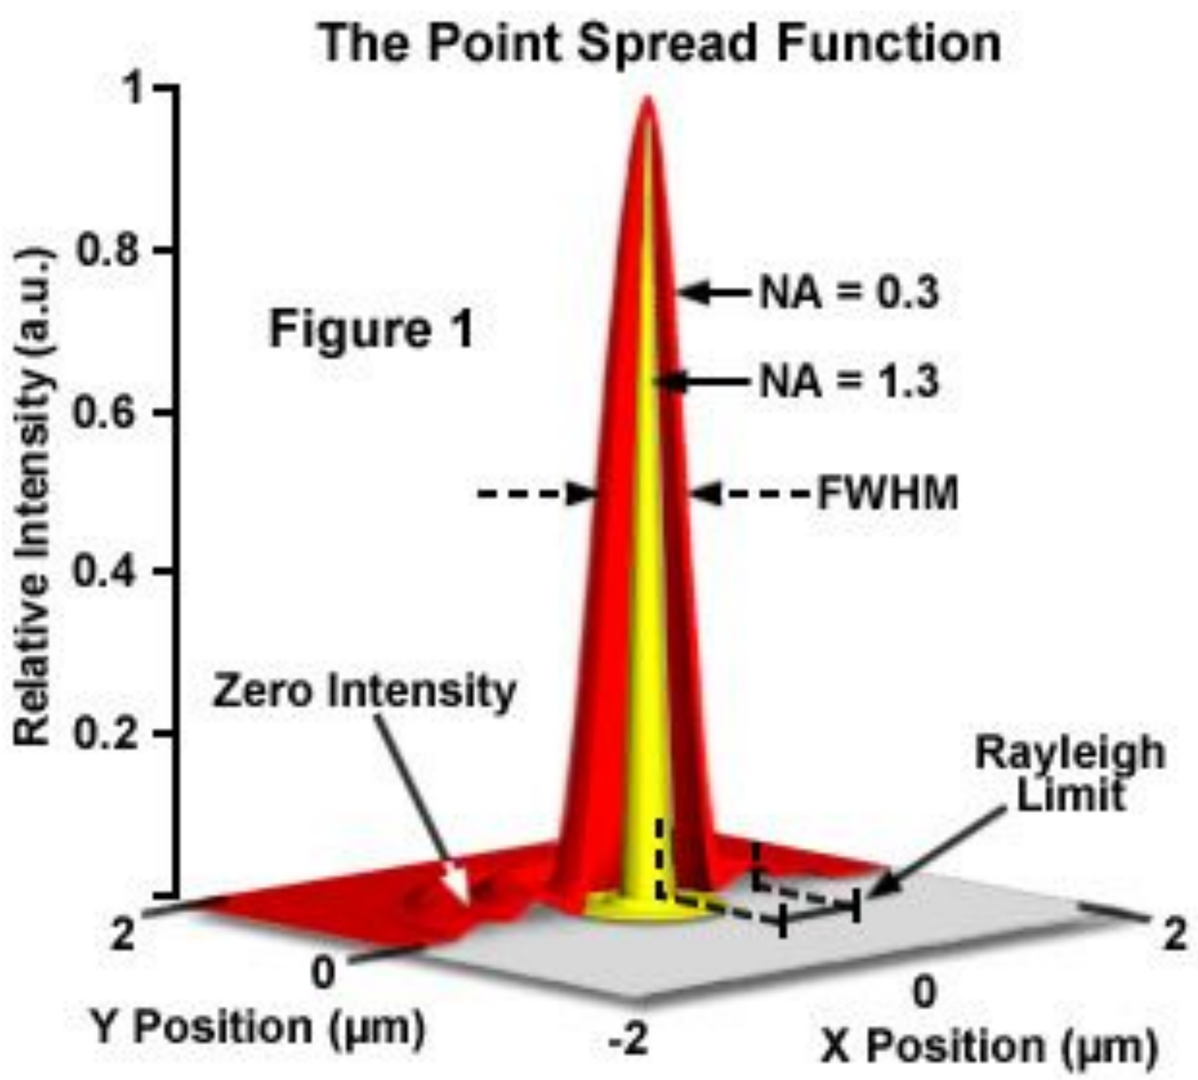

Motivation

Successful examples of blind inverse problems.

Motivation

Successful examples of blind inverse problems.

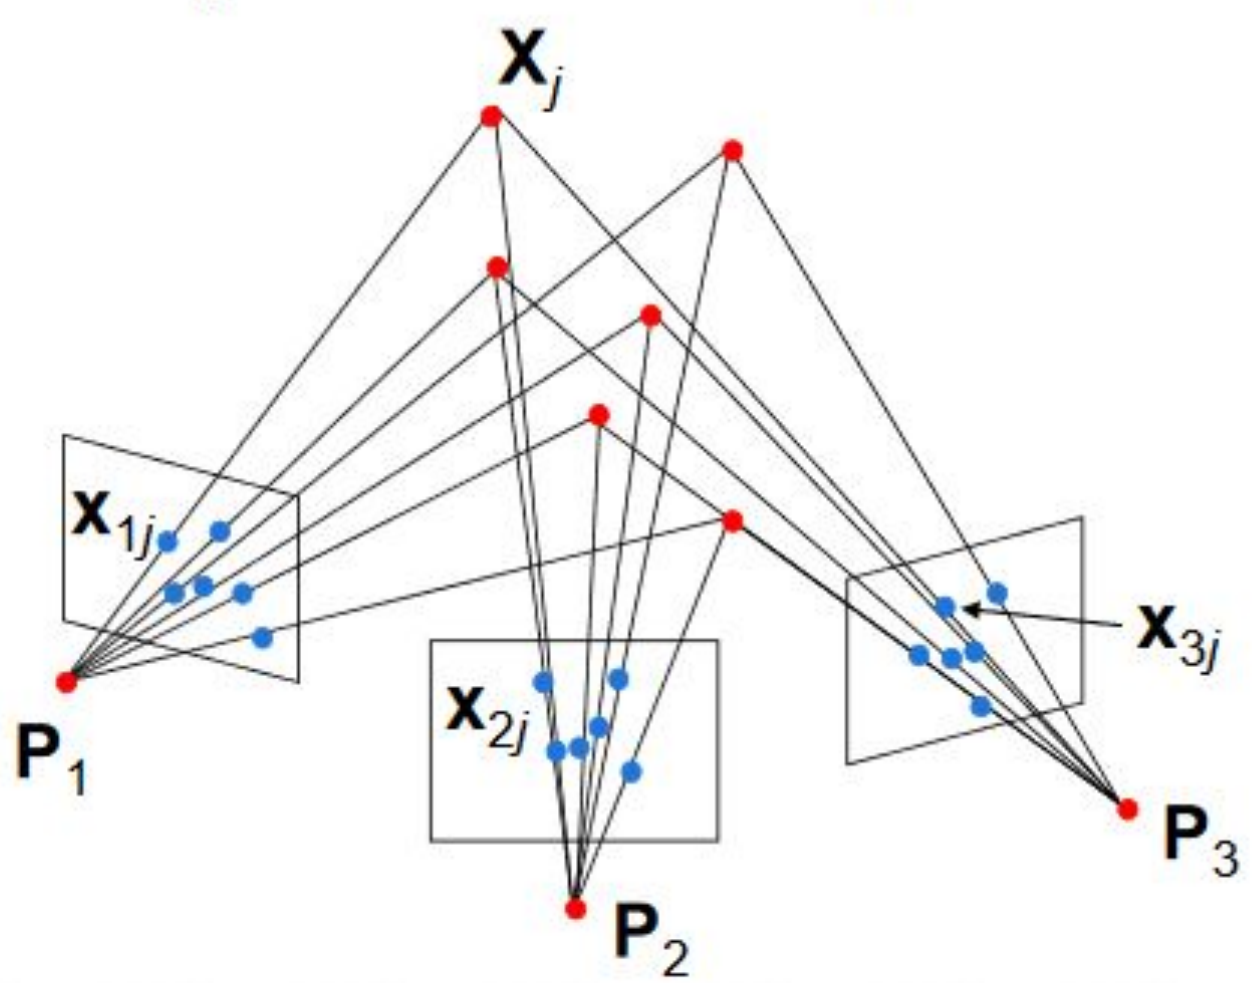

1. Joint pose and geometry estimation (Schönberger 2016)

Motivation

Successful examples of blind inverse problems.

1. Joint pose and geometry estimation (Schönberger 2016)
2. Joint spike timing and waveform estimation (spike sorting)

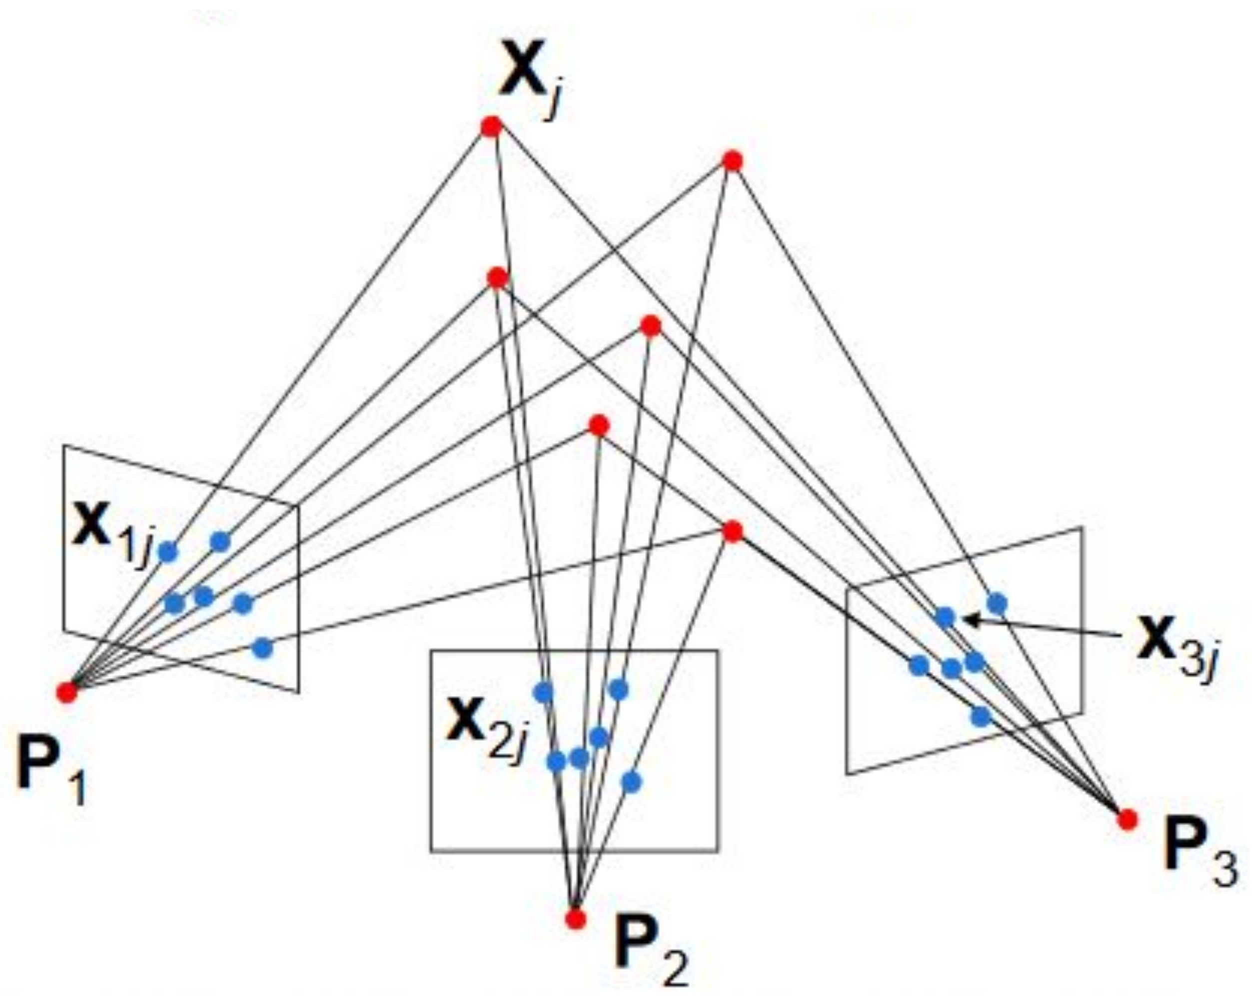

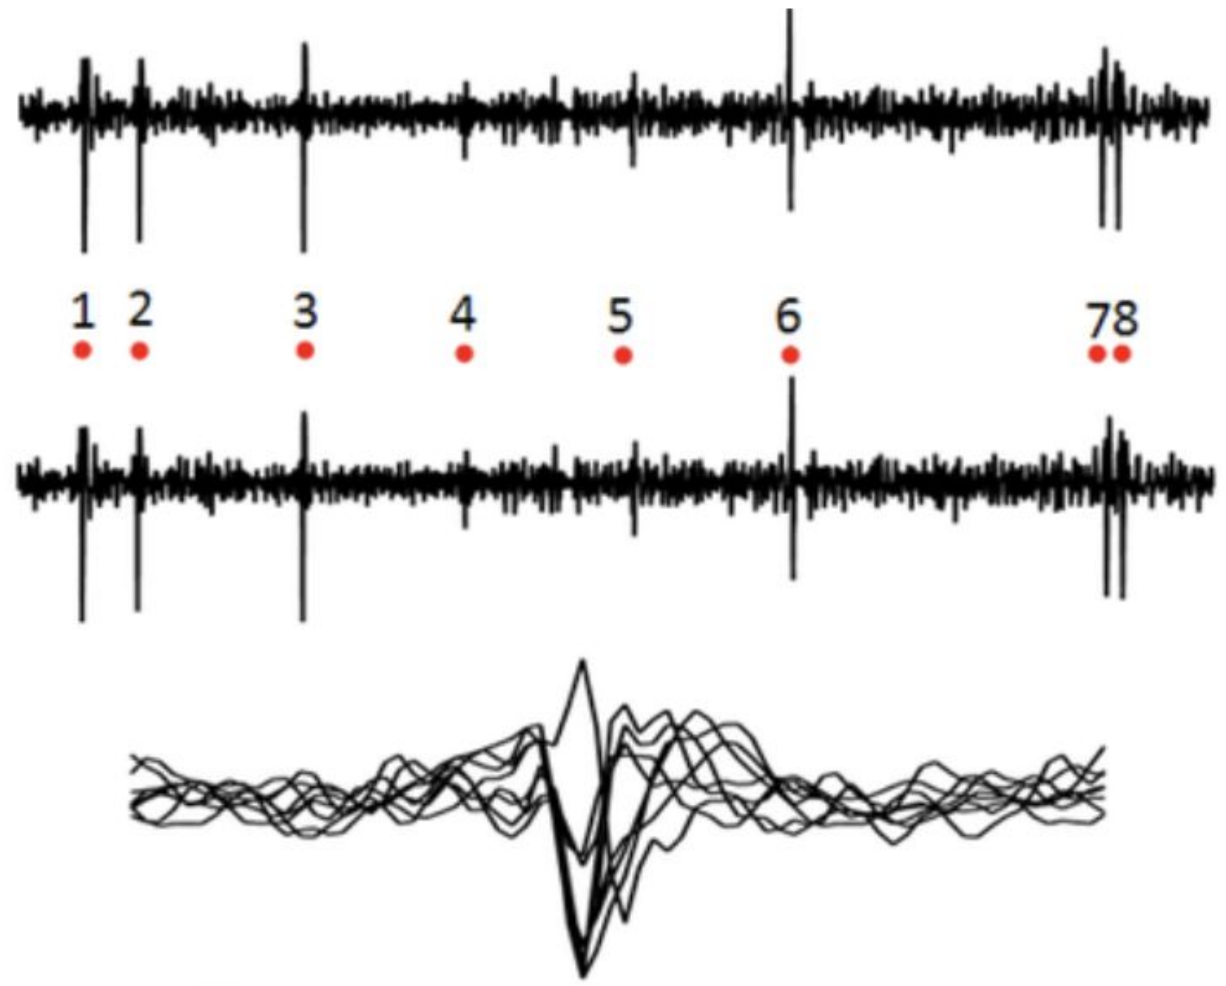

Motivation

Successful examples of blind inverse problems.

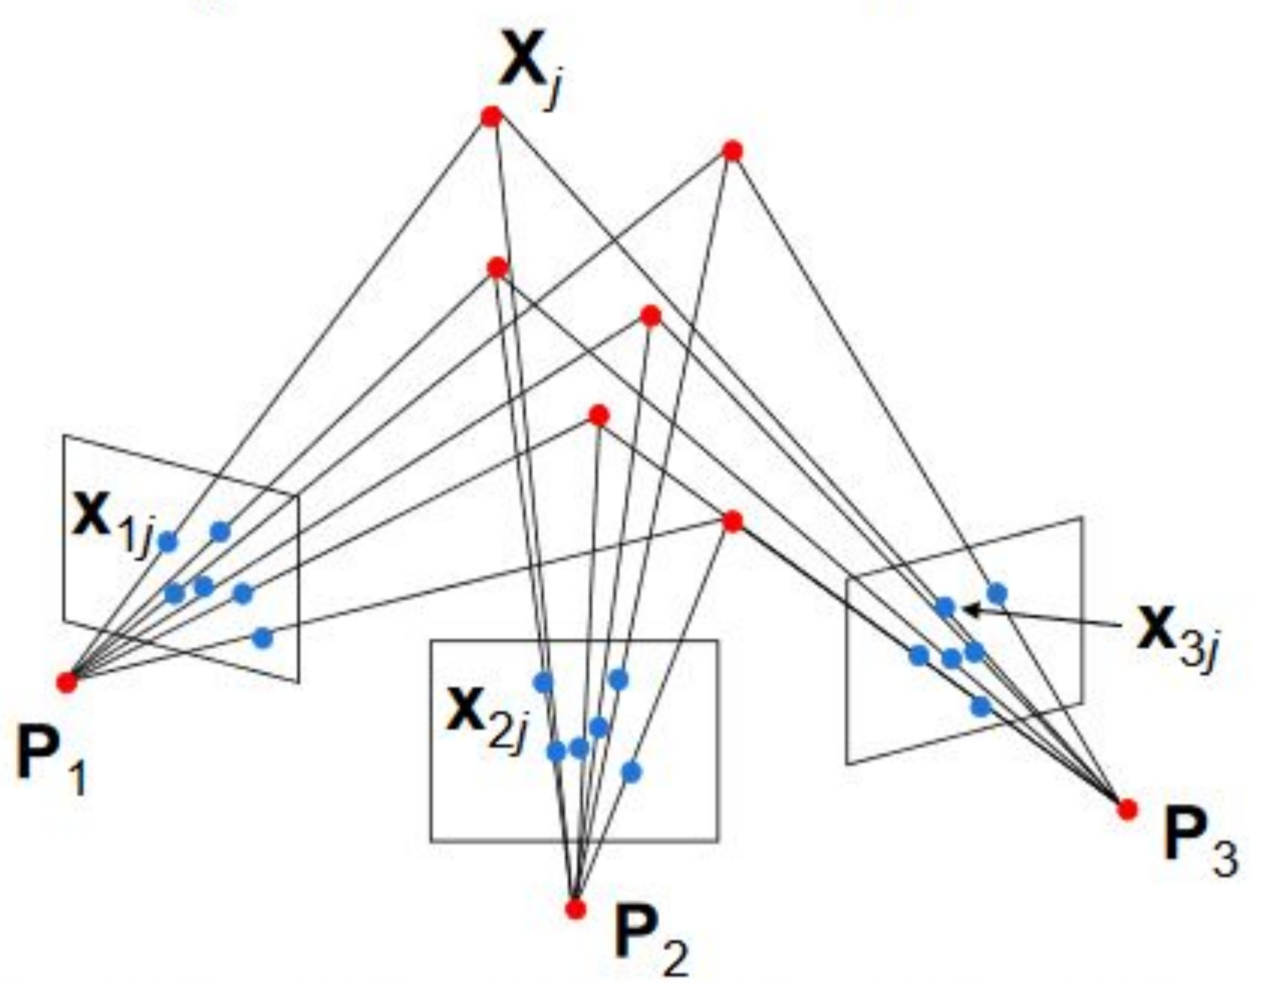

1. Joint pose and geometry estimation (Schönberger 2016)
2. Joint spike timing and waveform estimation (spike sorting)
3. Jointly eye-trace and neural response estimation (Yates 2023)

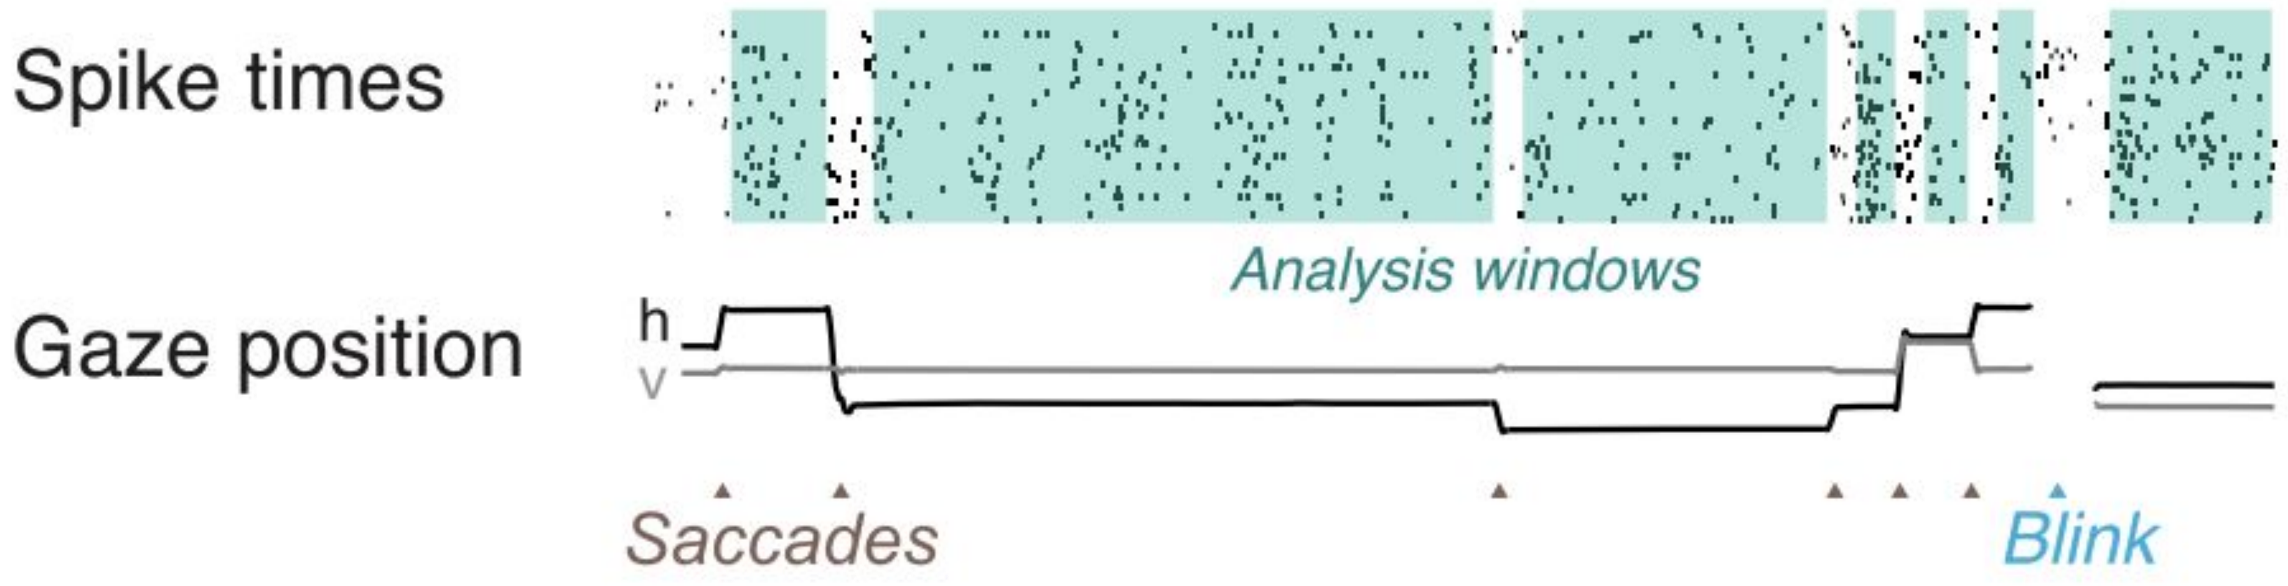

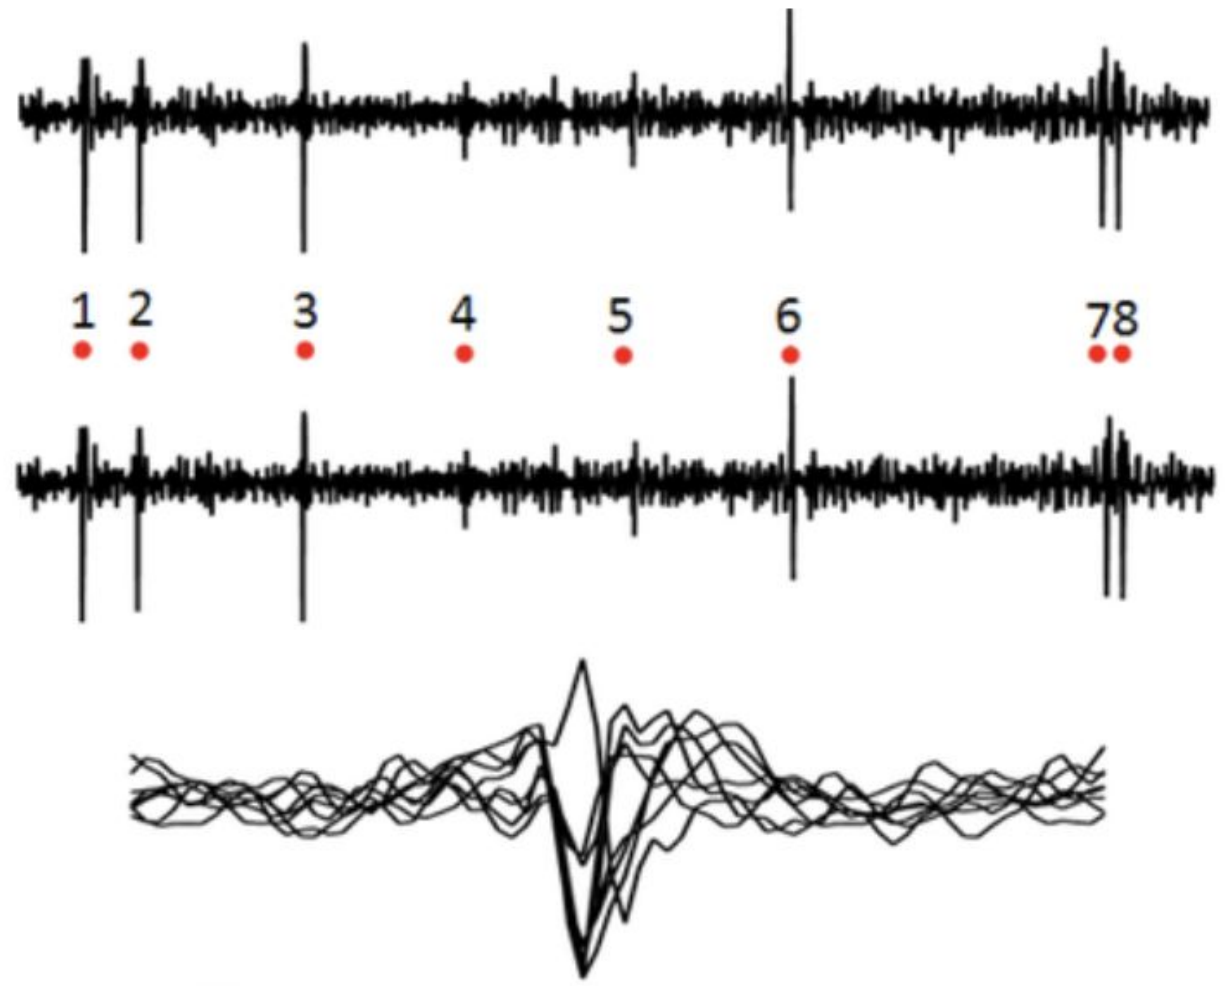

Motivation

Successful examples of blind inverse problems.

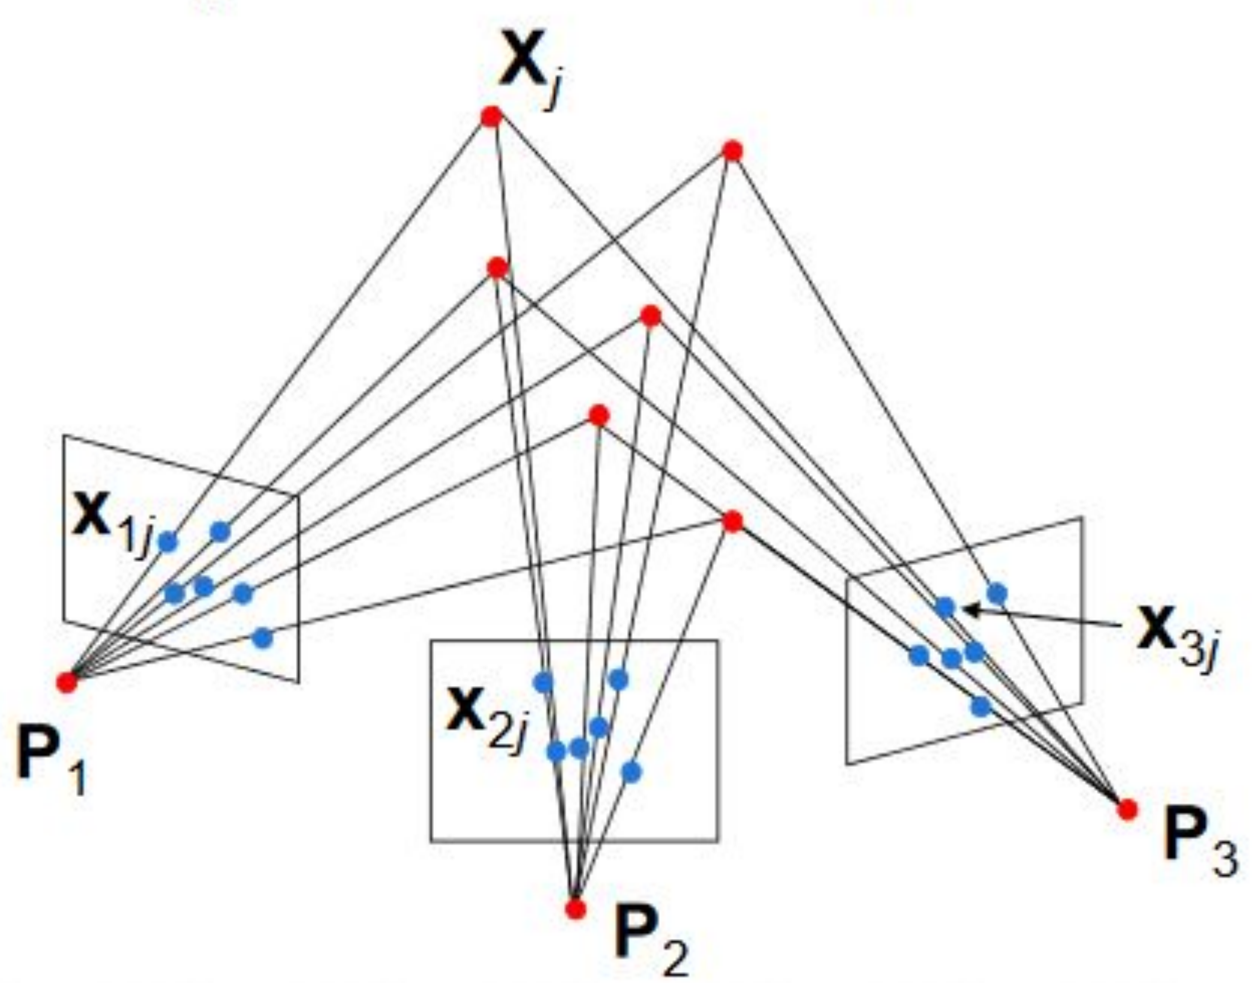

1. Joint pose and geometry estimation (Schönberger 2016)
2. Joint spike timing and waveform estimation (spike sorting)
3. Jointly eye-trace and neural response estimation (Yates 2023)

Blind inverse problems can be solved when prior knowledge constrains the problem sufficiently.

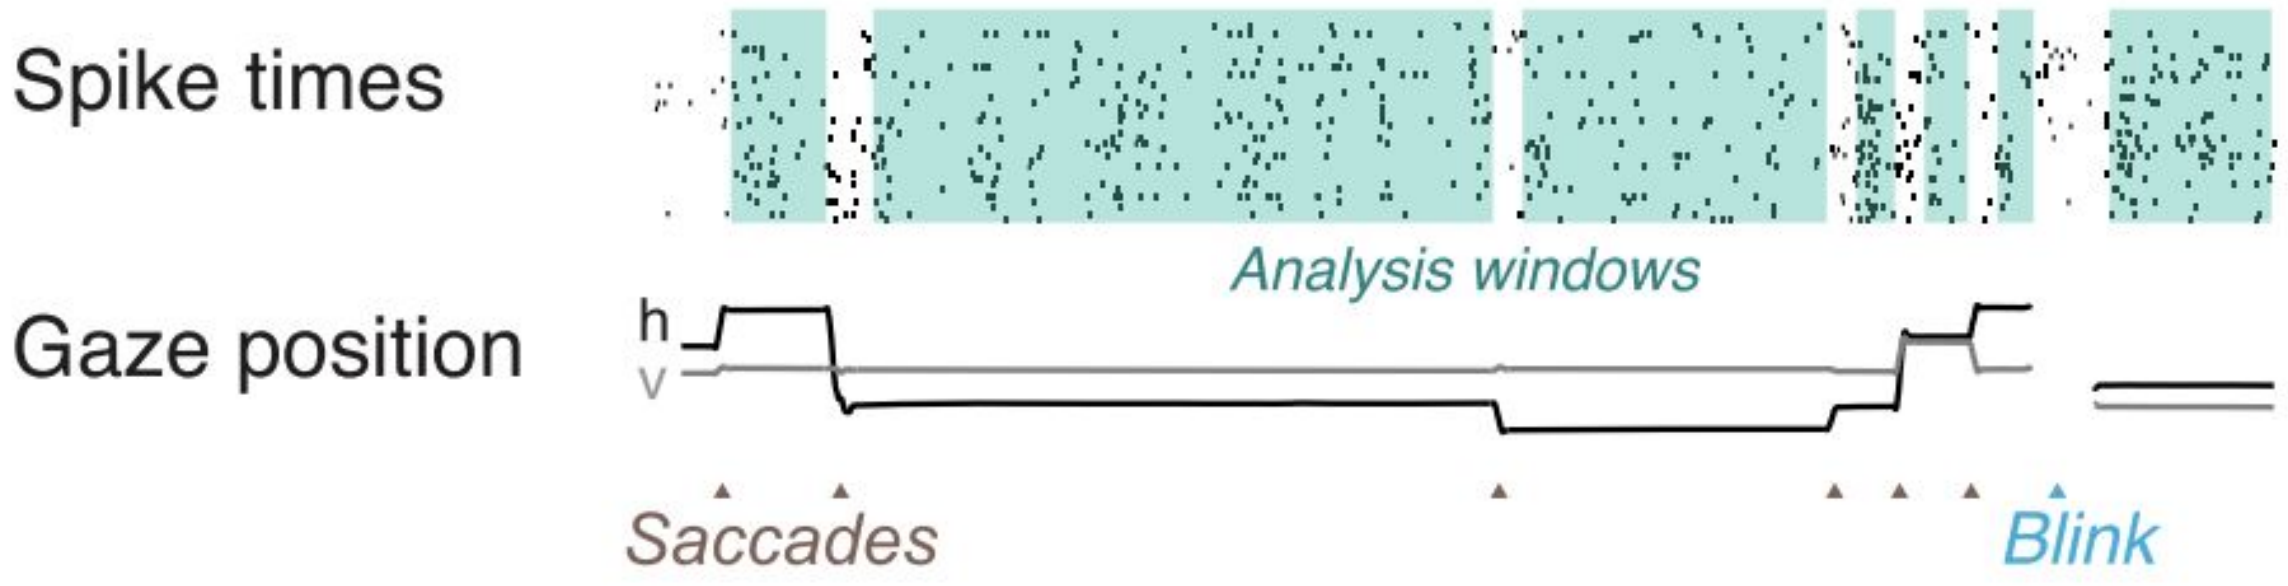

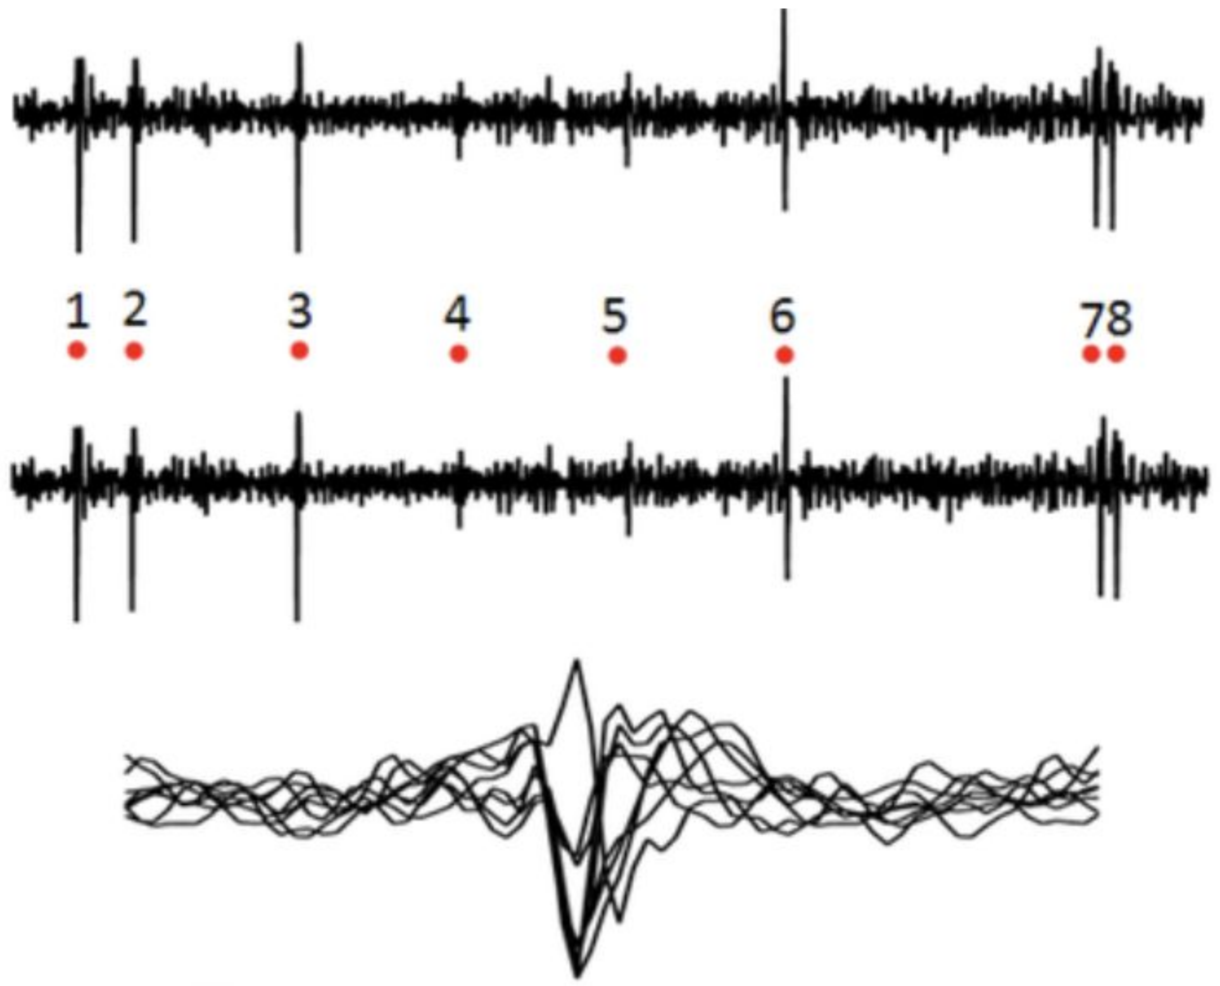

Schedule

1. Motivation
2. **Theory**
3. Experiments

Deconvolution

Make an effort to model

1. Prior - What kinds of objects do I intend to image? $$p(x)$$

2. Likelihood - What is the forward model? $$p(y|x)$$

The Posterior

Posterior
(Reconstruction knowledge)

Likelihood
(forward model)

Prior
(Object knowledge)

$$p\left(x \mid y\right) = \frac{p\left(y \mid x\right) p\left(x\right)}{p\left(y\right)}$$

Evidence
Measurement Knowledge

The Posterior

Likelihood
(forward model)

Prior
(Object knowledge)

Posterior
(Reconstruction knowledge)

$$p(x \mid y) = \frac{p(y \mid x)\, p(x)}{p(y)}$$

Evidence
Measurement Knowledge

constant

The Posterior

Posterior (Reconstruction knowledge) | **Likelihood** (forward model) | **Prior** (Object knowledge)

$$p\left(x \mid y\right) \propto p\left(y \mid x\right) p\left(x\right)$$

Baked-in Assumption

We assume that the kernel k is **known**

$$p\,(y\,|\,x) = x * k + \epsilon$$

Baked-in Assumption

We assume that the kernel k is **known**

$$p\,(y\,|\,x) = x * k + \epsilon$$

What if we don't know k?

Blind Deconvolution

$$p\left(x, k \mid y\right) \propto p\left(y \mid x, k\right) p\left(x, k\right)$$

Posterior
(Object and kernel reconstruction knowledge)

Likelihood
(forward model)

Prior
(Object and kernel knowledge)

Blind Deconvolution

$$p\left(x, k \mid y\right) \propto p\left(y \mid x, k\right) p\left(x, k\right)$$

Posterior
(Object and kernel reconstruction knowledge)

Likelihood
(forward model)

Prior
(Object and kernel knowledge)

Blind Deconvolution

$$p\left(x,k \mid y\right) \propto p\left(y \mid x,k\right) p\left(x,k\right)$$

Posterior
(Object and kernel reconstruction knowledge)

Likelihood
(forward model)

Prior
(Object and kernel knowledge)

Sensor Independence Assumption

$$p\left(x,k\right)=p\left(x\right)p\left(k\right)$$

Object and kernel knowledge are independent from each other

Blind Deconvolution

$$p\left(x,k\,|\,y\right) \propto p\left(y\,|\,x,k\right)p\left(x\right)p\left(k\right)$$

Posterior
(Object and kernel reconstruction knowledge)

Likelihood
(forward model)

Priors
(Object and kernel knowledge)

Structured Priors

$$p(k) = \delta(k - K)$$

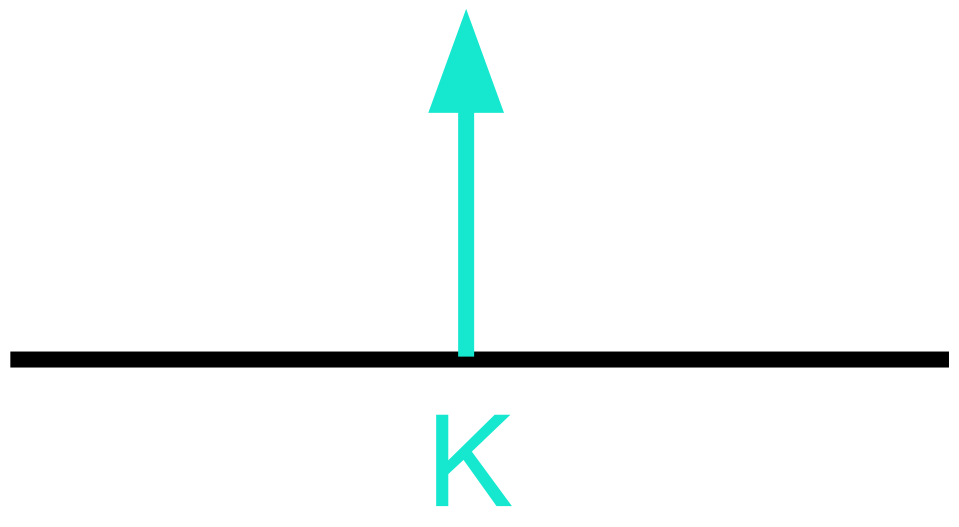

- Regular deconvolution
- Very structured
- Minimum entropy

$$p(k) = \mathcal{U}(a, b)$$

- Completely blind deconvolution
- No structure
- Maximum entropy

Structured Priors

$$p(k) = \delta(k - K)$$

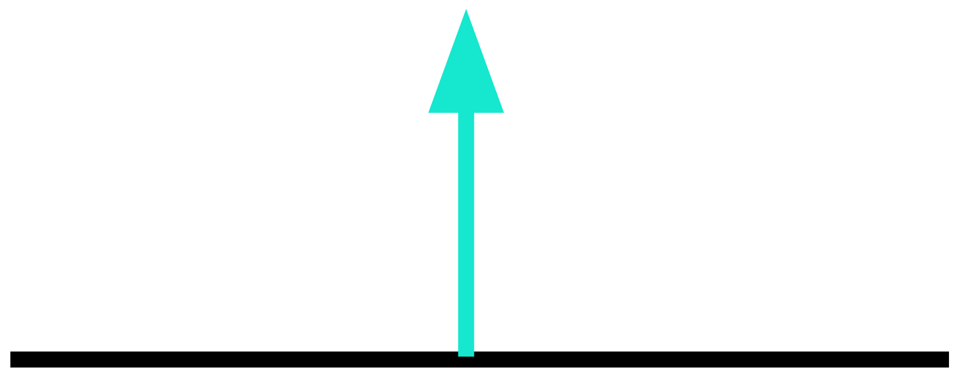

- Regular deconvolution
- Very structured
- Minimum entropy

Impossible in practice!

$$p(k) = \mathcal{U}(a, b)$$

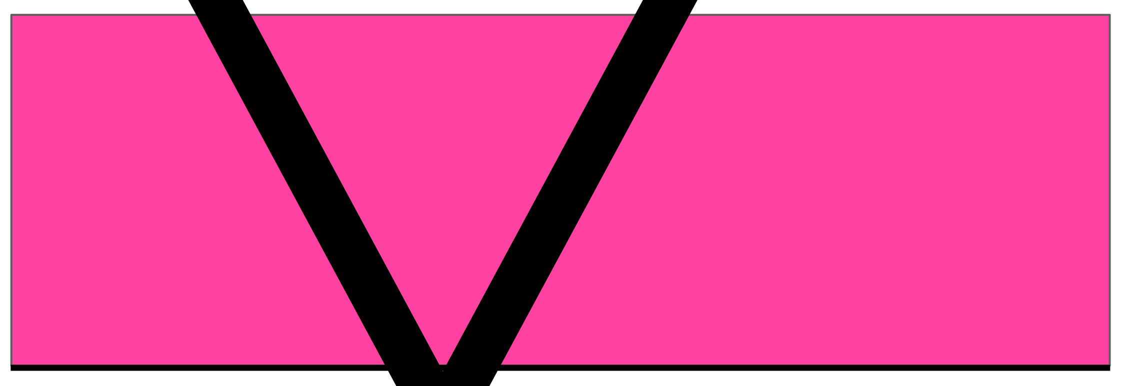

- Completely blind deconvolution
- No structure
- Maximum entropy

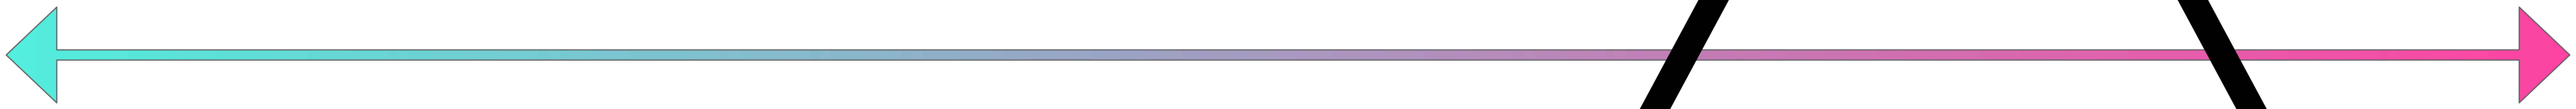

Structured Priors

$$p(k) = \delta(k - K)$$

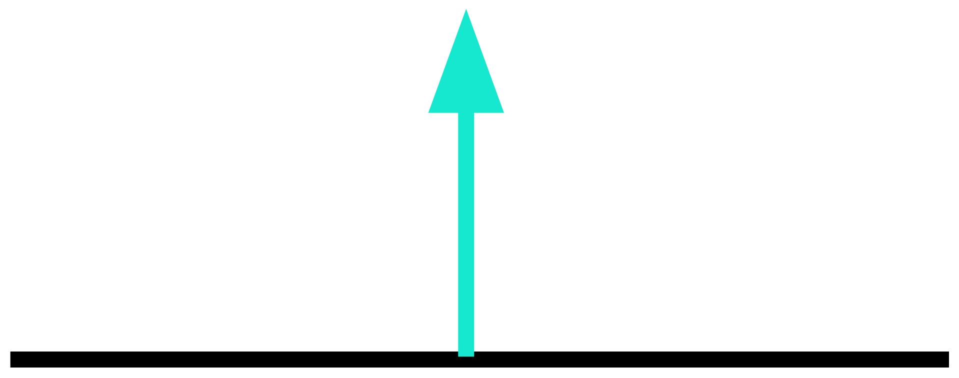

- Regular deconvolution
- Very structured
- Minimum entropy

Impossible in practice!

$$p(k) = \mathcal{U}(a, b)$$

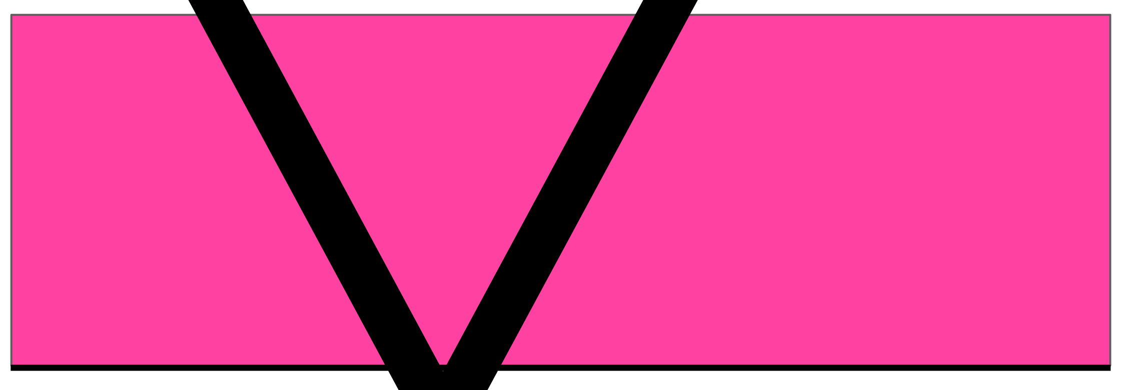

- Completely blind deconvolution
- No structure
- Maximum entropy

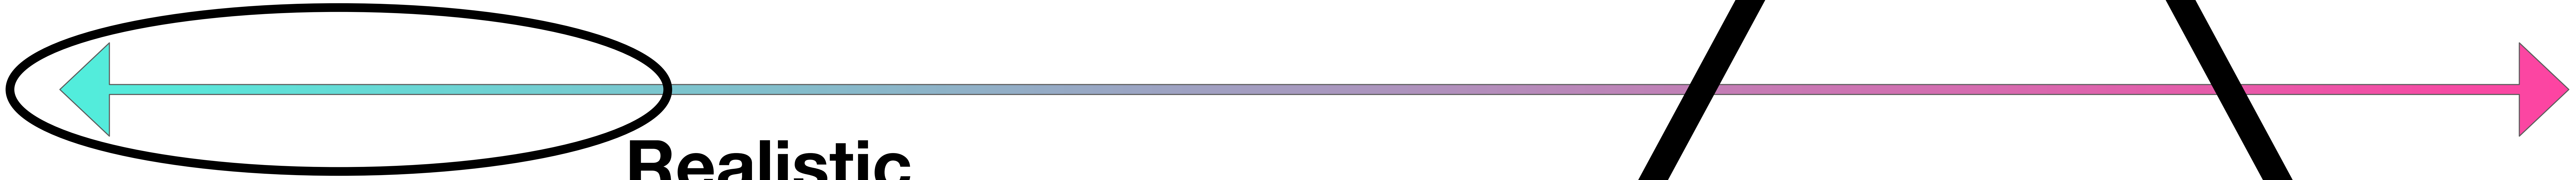

Schedule

1. Motivation
2. Theory
3. **Experiments**

Experiments

- Metrics used
  - PSNR
  - SSIM
  - Kernel Error (KE)
- Experiments
  - Motion kernel – synthetic circle
  - Motion kernel – dog eating grass
  - Gaussian kernel – cat
- Notable findings

Experiments

Metrics – PSNR

- SNR between reconstructed image and true image including mean square error

- PSNR > 30 dB → really good
- 25 dB < PSNR < 30 dB → good
- 20 dB < PSNR < 25 dB → ok
- PSNR < 20 dB → bad

$$\mathrm{PSNR} = 10\log_{10}\left(\frac{L^2}{\mathrm{MSE}}\right)$$

Experiments

Metrics – SSIM

- Measures structural similarity via:
  - Luminance
  - Contrast
  - Structure

$$\mathrm{SSIM}(x, \hat{x}) = \frac{(2\mu_x\mu_{\hat{x}} + C_1)(2\sigma_{x\hat{x}} + C_2)}{(\mu_x^2 + \mu_{\hat{x}}^2 + C_1)(\sigma_x^2 + \sigma_{\hat{x}}^2 + C_2)}$$

- SSIM > 0.9 → really good
- 0.85 < SSIM < 0.9 → good
- 0.75 < SSIM < 0.85 → ok
- SSIM < 0.75 → bad

Experiments

Metrics – Kernel Error (KE)

- How close the estimated PSF is to the real one

$$\mathrm{KE} = \left\|\hat{k} - k_{\mathrm{true}}\right\|_2$$

Experiments

Motion kernel – synthetic circle

- PNSR = 26.83 dB
- SSIM = 0.9642
- KE = 0.9926

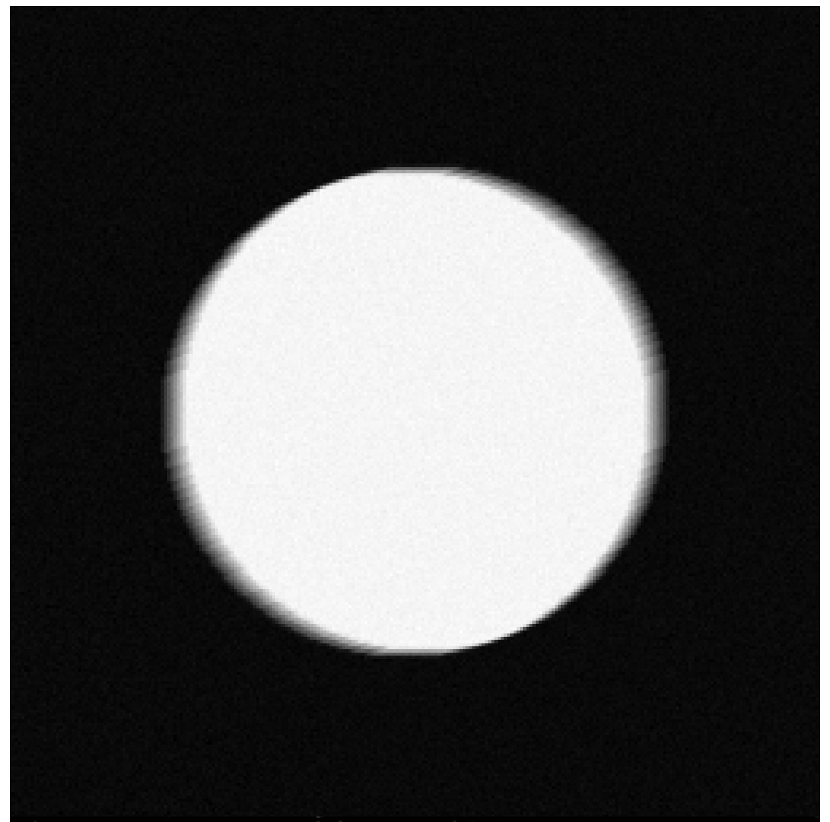

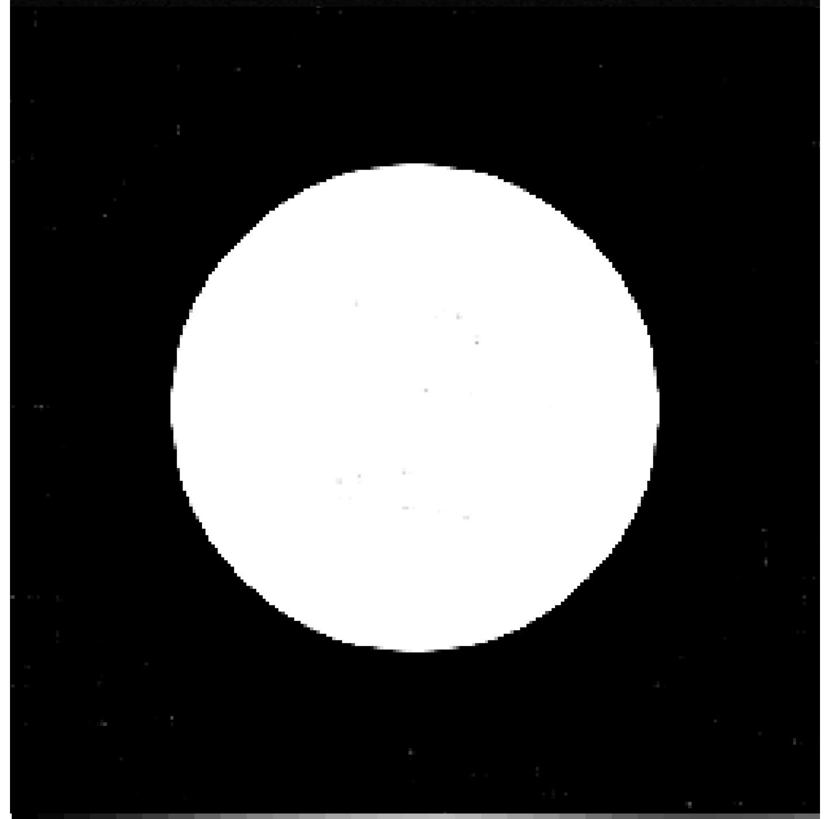

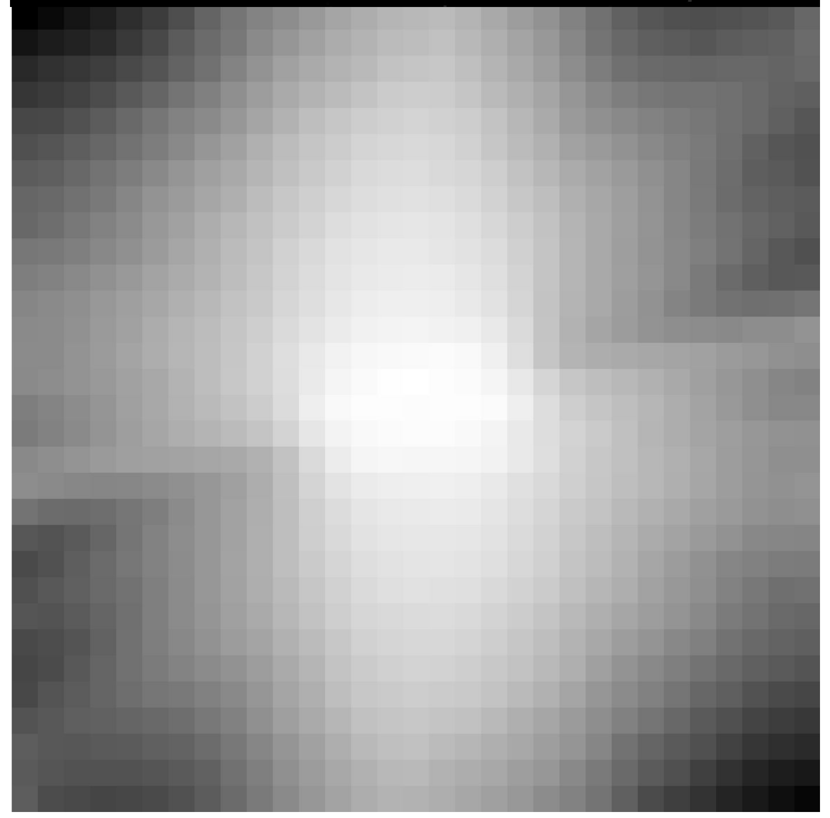

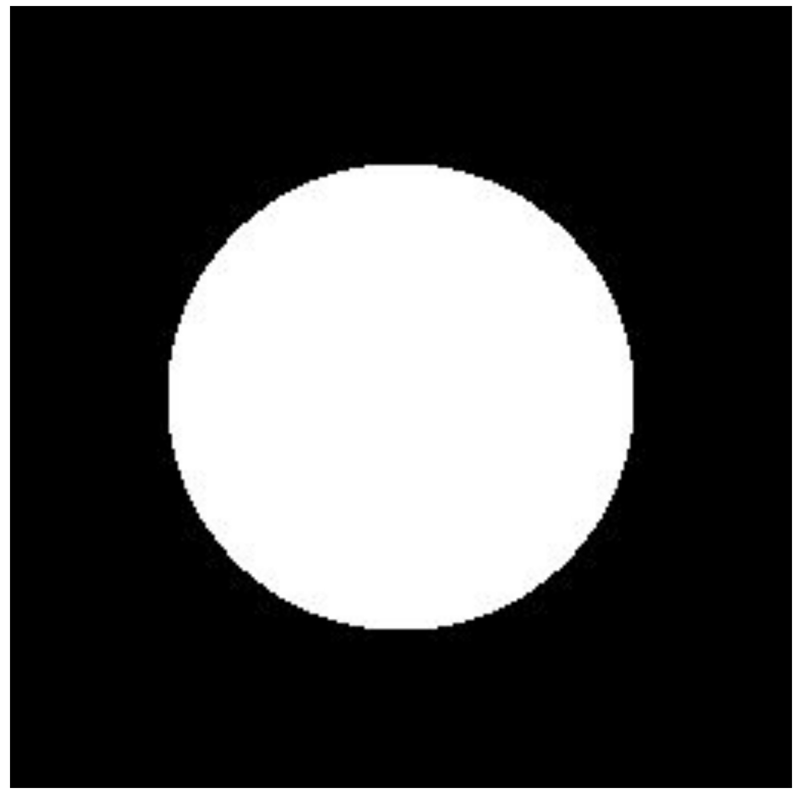

Experiments

Motion kernel – dog eating grass

- PNSR = 22.74 dB
- SSIM = 0.7313
- KE = 3.4408

Measured image

Reconstructed image

Reconstructed PSF

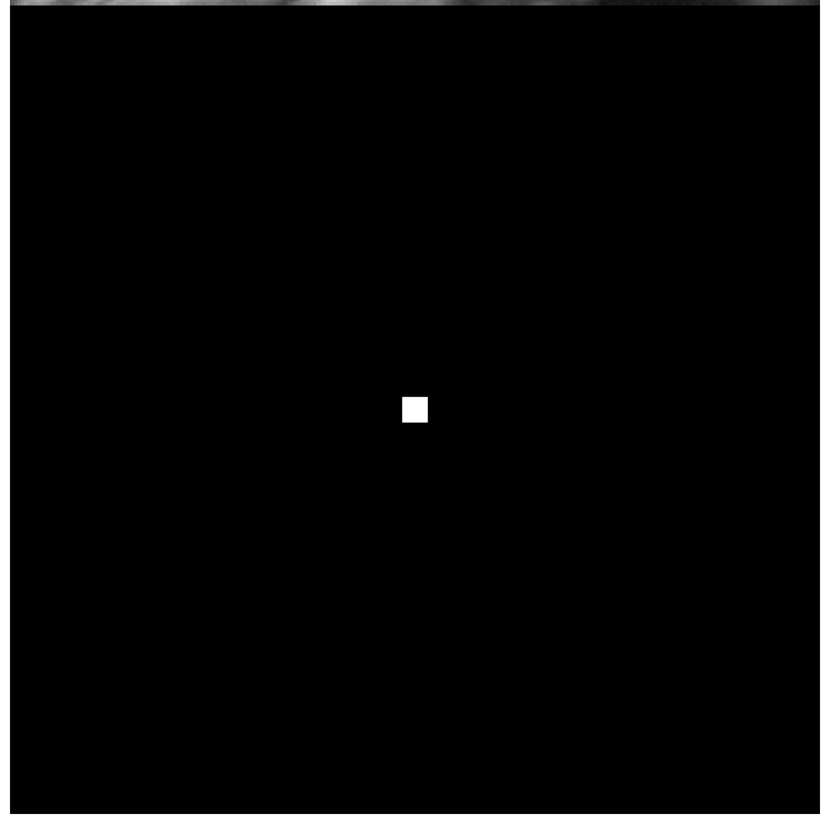

Real image

Experiments

Gaussian kernel – cat

- PNSR = 28.16 dB
- SSIM = 0.7412
- KE = 0.7803

Measured image

Reconstructed image

Reconstructed PSF

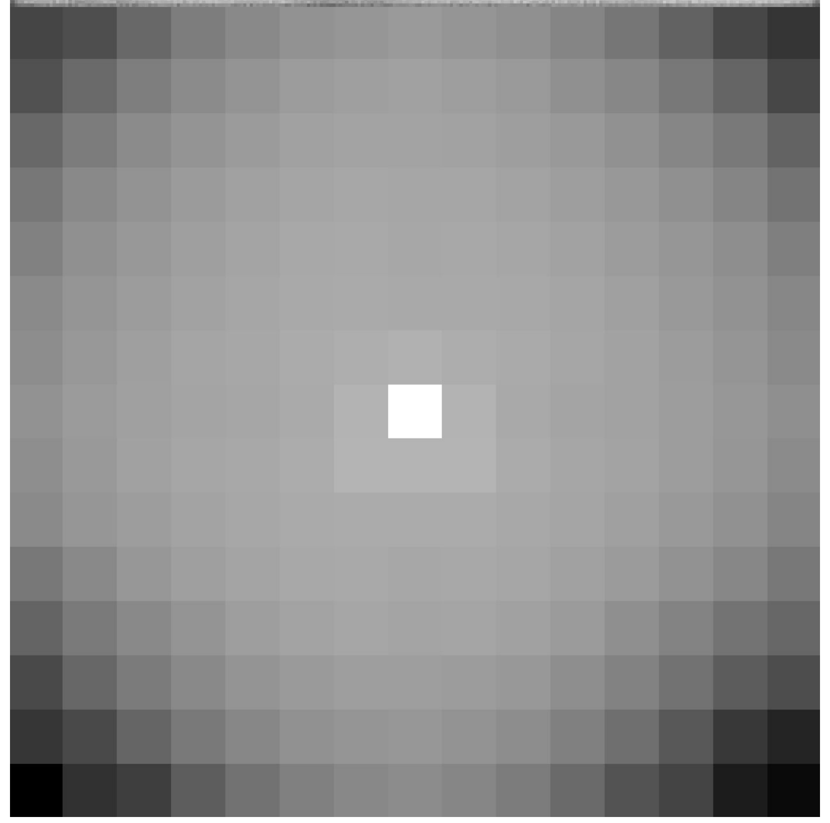

Real image

Experiments

Notable Findings

- *Kernel error is consistently bad*
  - Image priors (diffusion prior) dominate the loss function → kernel error stays high (goes toward identity)
- *PSNR is more stable than the SSIM and kernel error*
  - Diffusion prior tends to hallucinate plausible textures, but not necessarily the correct ones → PSNR is stable, impact on SSIM
  - Diffusion prior helps the image more than the kernel

Blind deconvolution succeeded more or less in reconstructing the image, not the kernel → acts more as regular deconvolution

Experiments

Notable Findings – Solutions

- Increase kernel prior weights → diffusion prior doesn’t impact the loss function as much → model corrects itself more based on the kernel
- Constrain the diffusion prior (lower weight or lower denoising strength) → reduces hallucinated textures → improves SSIM and preserves structural information

Preliminary Results on Video

Curated a Video Dataset

- Developed a video testbed
- Useful for training ML blind deconvolution algorithms
- Spatial and temporal structures that are useful for blind deconvolution
- Limitation: we don't know p(x) for videos

Video Synthesis $x \sim p\left(x\right)$

Shift-Equivariant Gaussian Approximation

- Shift-equivariant gaussian approximation of the first dataset
- Developed synthetic video dataset
- Know the p(x) here

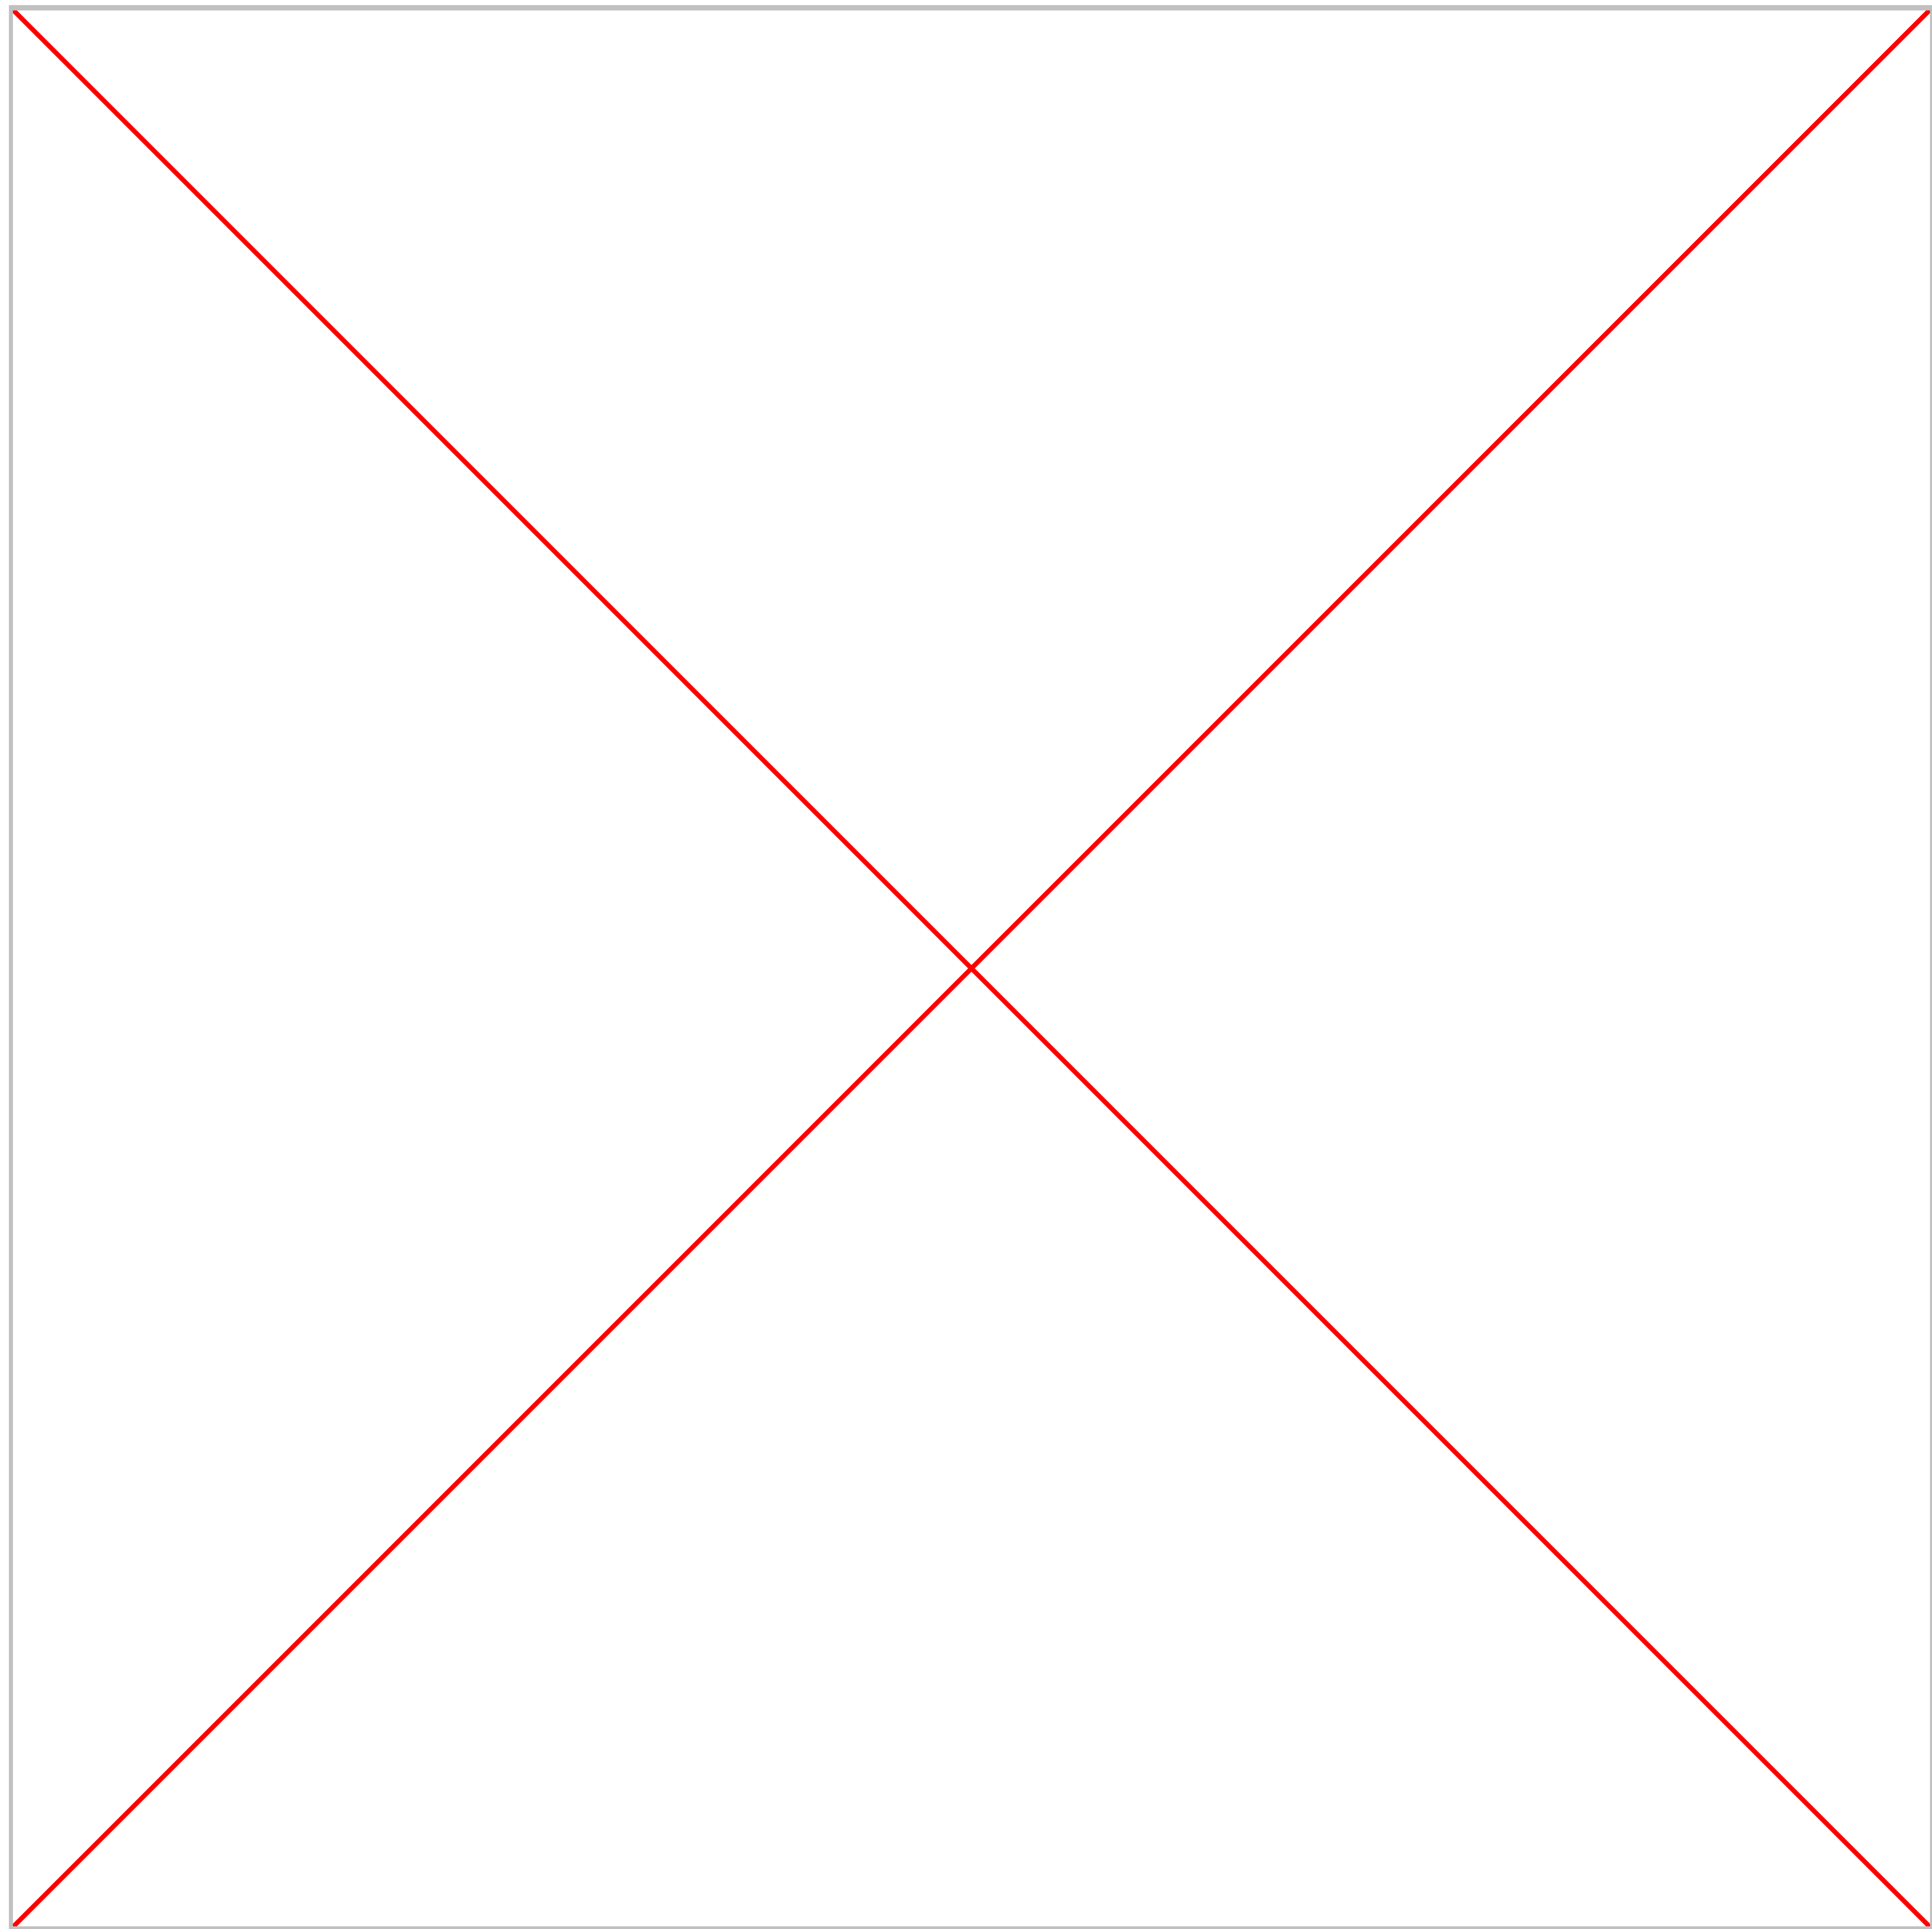

Kernel Synthesis $k \sim p\left(k\right)$

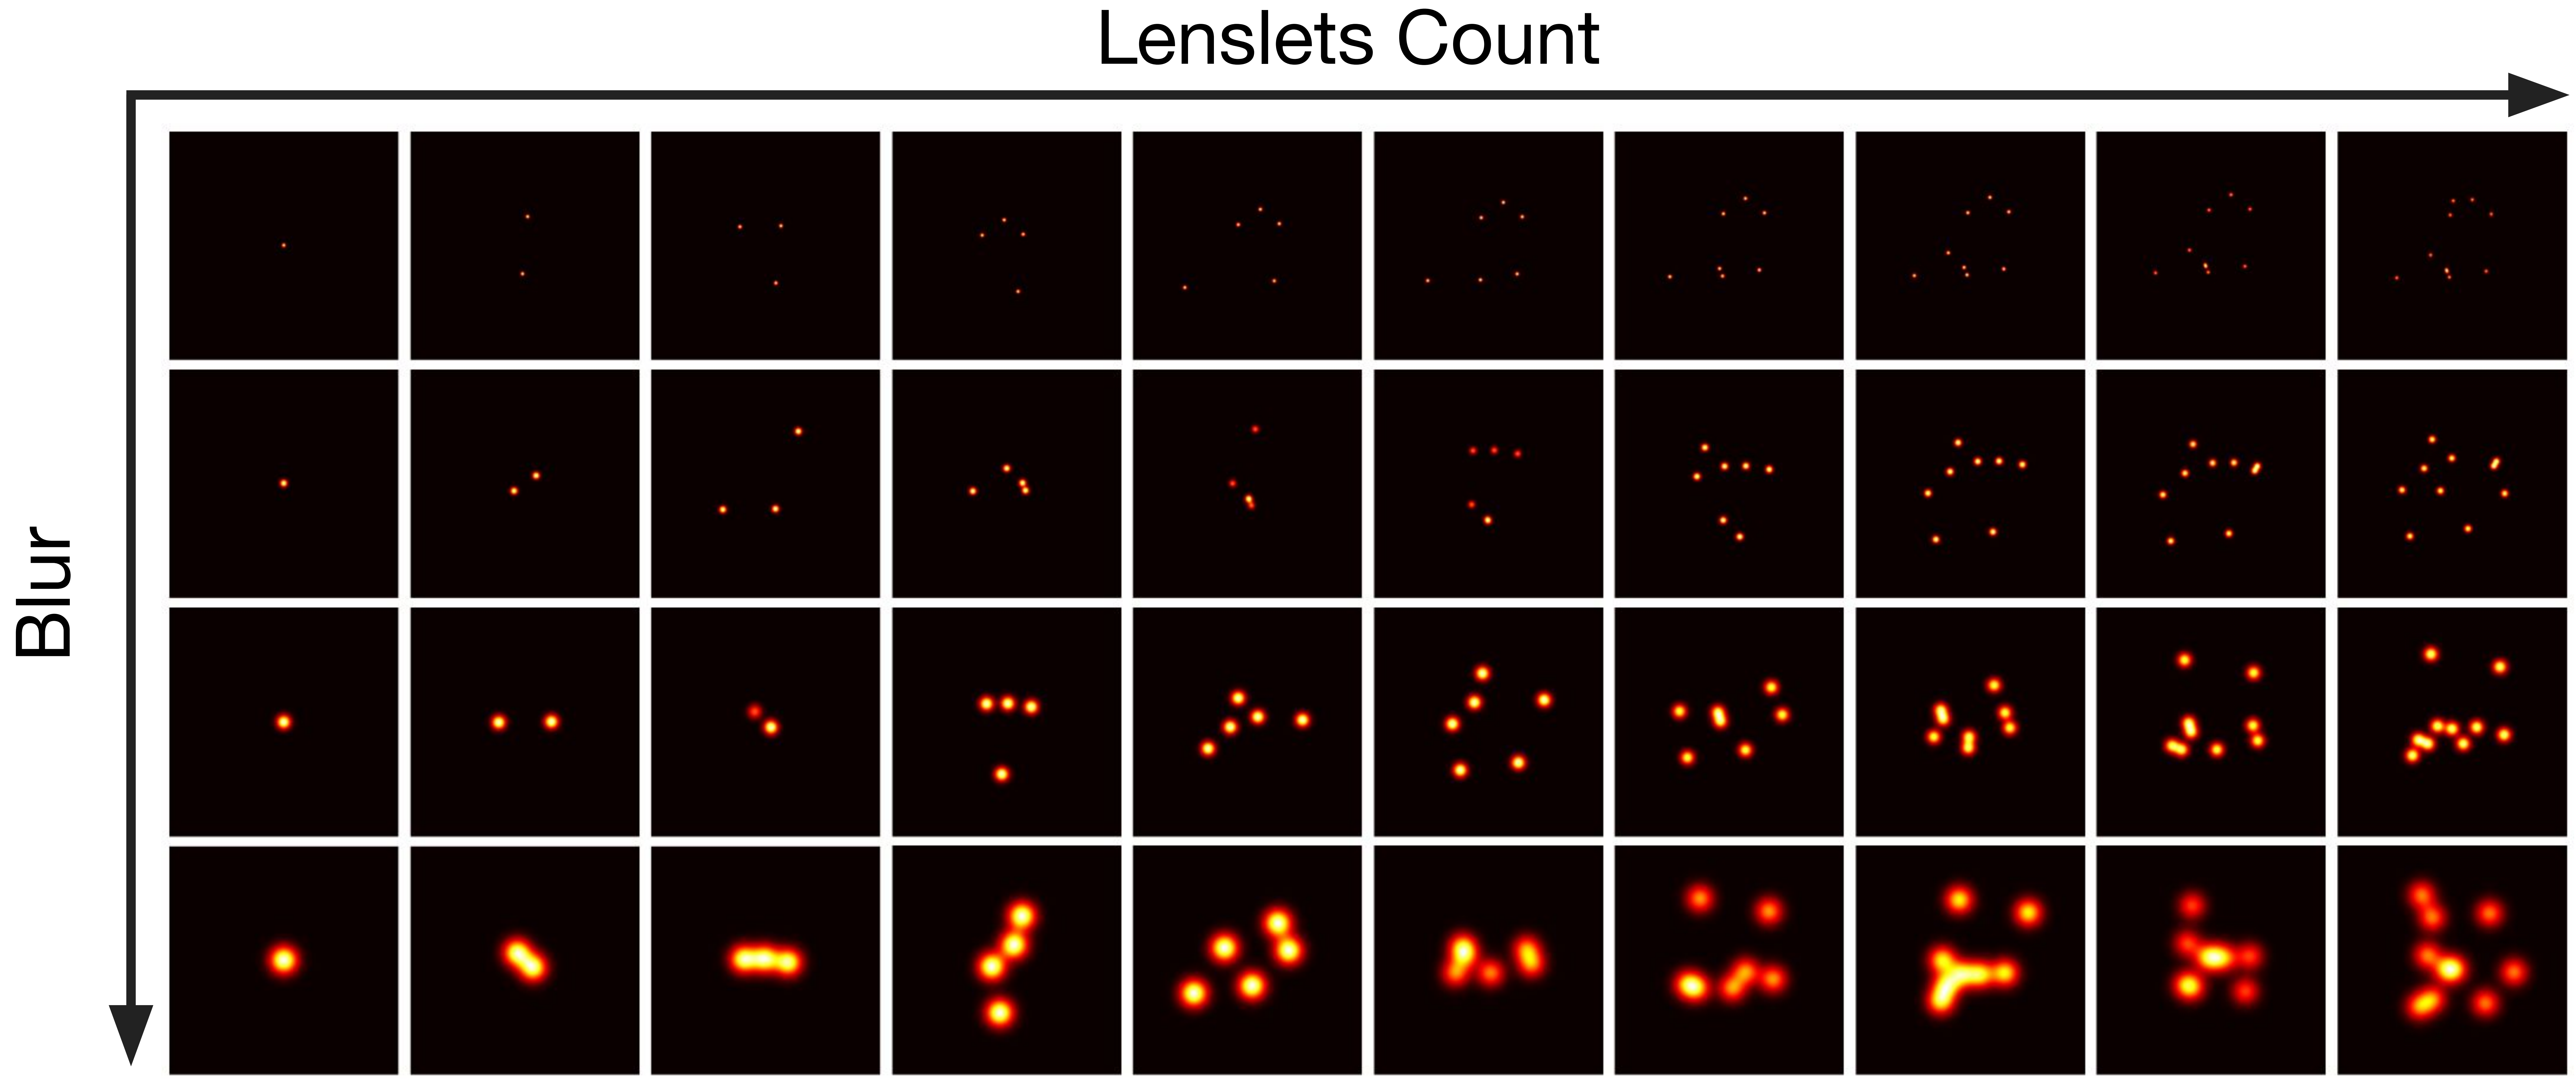

References

Schönberger & Frahm, CVPR 2016

Khairuddin et al., IEEE RAM 2016

Ekanadham et al., PhD Thesis, NYU 2015

Yates et al., Nat. Commun. 2023

Z. Wang et al., IEEE Trans 2004

Ulyanov et al. 2018, CVPR 2018

Temporal Structure

Videos have spatiotemporal structure