STORM imaging

- microscopy, the diffraction limit
- "super-resolution": techniques to break this barrier.
- masks: with telescopes, you can control them
- masks: in our microscopy setting, we cannot control them, or even observe them... we assume that they change gradually.

The Graphical Model

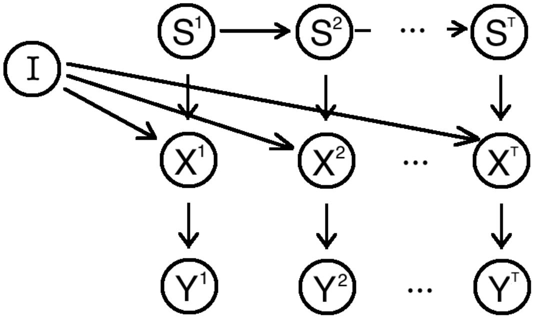

- We are interested in the original image I.
- for each pixel (m, n), $Y^t_{mn} \sim Poi(\gamma[AX^t]_{mn})$
- A is the Gaussian blur matrix. γ is the exposure time of each picture. Y^t_{mn} is the number of photons observed at that pixel and time.
- the masks S^t follow a Markov chain: $S^{t+1} \sim Q(S^t)$

Inference ideas

- **direct inference**, i.e. maximizing $P(Y|I)$ would require integrating out all S^t, i.e. TU^2 parameters.
- for each t, let $\hat{X}^t = \underset{X^t}{argmax} \log P(Y^t|X^t) - \lambda ||X^t||_2^2$
- $\hat{X}^t$ is a noisy realization of X^t:
  $E[\hat{X}^t] = X^t$, $Var(\hat{X}^t) = A^{-1}X^t$.
  ... and since X^{t+1} depends on X^t, maybe we can treat this as an HMM with Poisson emissions.
- but perhaps the simplest thing to do is realize that most pixels are unmasked at some point, so one hack is: $\hat{I}_{mn} = \underset{t}{max} \hat{X}_{mn}$.
  ... or better: take a high empirical quantile (e.g. 95th quantile).

Simulation: 3x3

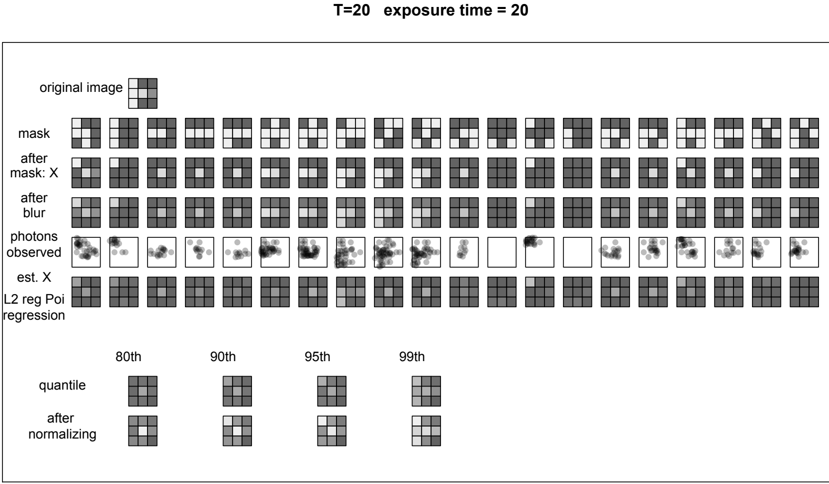

Simulation: 5x5

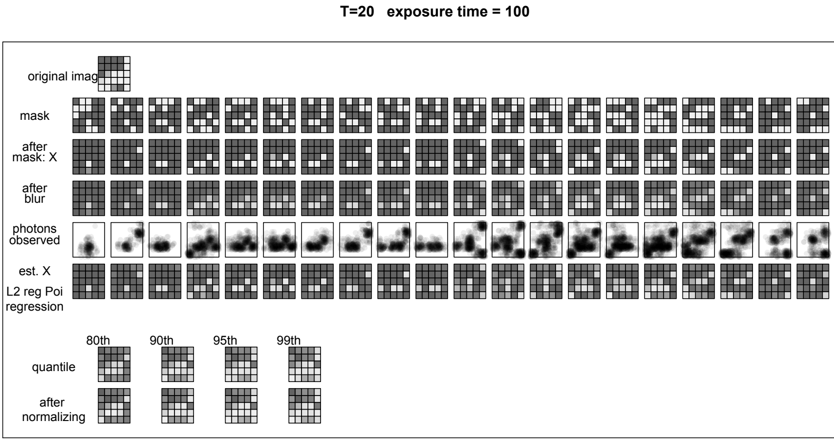

Future work

- assume sparsity of I, or the spatial derivatives of I.
- take advantage of the knowledge that adjacent masks are similar

Thanks to Liam Paninski and Eftychios Pnevmatikakis for discussions.