




list



Mixed Poisson processes, diffusions and inverse problems


Thibaut Pellerin, ETICS 2026

Supervized by Alain Durmus (CMAP), Elinor Berger (CEA-List), Eric Barat (CEA-List)




1. (Poisson) inverse problems

Positron Emission Tomography




- An image $x \in \mathbb{R}^d$, $d = 256 \times 256$, from the distribution of PET brain images π

Positron Emission Tomography




- An image $x \in \mathbb{R}^d$, $d = 256 \times 256$, from the distribution of PET brain images π
- The Radon transform of x , $\mathcal{A}(x) \in \mathbb{R}^{d_y}$, $d_y = 180 \times 180$

Positron Emission Tomography



- An image $x \in \mathbb{R}^d$, $d = 256 \times 256$, from the distribution of PET brain images π
- The Radon transform of x , $\mathcal{A}(x) \in \mathbb{R}^{d_y}$, $d_y = 180 \times 180$
- The Poisson observations $y \sim \text{Pois}(\mathcal{A}(x))$ measured by the detectors

Positron Emission Tomography



- An image $x \in \mathbb{R}^d$, $d = 256 \times 256$, from the distribution of PET brain images π
- The Radon transform of x , $\mathcal{A}(x) \in \mathbb{R}^{d_y}$, $d_y = 180 \times 180$
- The Poisson observations $y \sim \text{Pois}(\mathcal{A}(x))$ measured by the detectors
- A naive reconstruction

Bayesian inverse problems

Let π denote a probability measure on $\mathbb{R}_+^d$, and $\mathcal{A} : \mathbb{R}_+^d \rightarrow \mathbb{R}_+^{d_y}$. If $X \sim \pi$, consider

$$\boxed{Y} = \text{Pois}(\boxed{\mathcal{A}(X)}) .$$

observation forward operator

π is the prior distribution, and we want the posterior distribution of X given $Y = y$,

$$\boxed{\pi^y(dx)} \propto \boxed{g(y|x)} \boxed{\pi(dx)} \propto \prod_{i=1}^{d_y} e^{-\mathcal{A}(x)_i} \mathcal{A}(x)_i^{y_i} \pi(dx)$$

posterior likelihood prior

Poisson denoising, $\mathcal{A} = a\mathbf{I}_d$




Figure 1: Ground truth x (McAlpine et al., 2016), observation $y \sim \text{Pois}(\mathcal{A}(x))$

Blur, $\mathcal{A} =$ convolution with Gaussian kernel




Figure 2: Ground truth x (McAlpine et al., 2016), observation $y \sim \text{Pois}(\mathcal{A}(x))$

PET, $\mathcal{A} = \text{Radon operator}$





Figure 3: Ground truth x (Aubert-Broche et al., 2006), observation $y \sim \text{Pois}(\mathcal{A}(x))$, and filtered-back projection

Posterior sampling



The diagram shows the formula for the posterior distribution: $\pi^y(dx) \propto g(y|x) \pi(dx)$. The term $\pi^y(dx)$ is enclosed in a pink box, with a red arrow pointing to it from the label "posterior" below. The term $g(y|x)$ is enclosed in an orange box, with an orange arrow pointing to it from the label "likelihood" above. The term $\pi(dx)$ is enclosed in a blue box, with a blue arrow pointing to it from the label "prior" below.

$$\pi^y(dx) \propto g(y|x) \pi(dx)$$

posterior

likelihood

prior

- Sampling from π^y is possible with MCMC if we have an analytic prior π (e.g., regularization term)
- Recent techniques use a diffusion model trained on π to sample from π^y instead
→ **data-driven prior**
- Most of the litterature consider Gaussian inverse problems: $Y = \mathcal{A}(X) + b\varepsilon$

Gaussian diffusion models



- Goal: obtain new samples from $X \sim \pi$ given a dataset

"forward" noising process

$X_t = tX + dB_t \Leftrightarrow dX_t = D_t^\pi(X_t)dt + dW_t$

"backward" generative process

- The Gaussian denoiser $D_t^\pi(X_t) = \mathbb{E}[X|X_t]$ can be learned as the minimizer of $f \mapsto \mathbb{E}[\|X - f(t, X_t)\|_2^2]$, using the forward equation
- Then, new π samples are obtained by simulating the backward equation

Diffusion models + Gaussian inverse problems

Now we want to sample from π^y , the posterior of the Gaussian inverse problem $Y = \mathcal{A}(X) + b\varepsilon$, with $X \sim \pi$.

- If we know the *conditional denoiser* D^{π^y} , we're good, but we don't want to train it for a specific inverse problem
- For general, and especially linear $\mathcal{A}$, there are a lot of ways to estimate D^{π^y} using a pre-trained D^π : *diffusion-based posterior sampling methods* (Chung et al., 2023; Song et al., 2023; Zhang et al., 2025)
- **Base case:** if $\mathcal{A} = aI_d$, then D^{π^y} is explicit !
- For Poisson inverse problems, this mostly does not work



2. Mixed Poisson processes as generative models

Counting processes

A **counting process** $\mathbf{Z} = (Z_t)_{t \geq 0}$ is a càdlàg process with unit jumps, and $Z_0 = 0$ a.s.

Ex: Poisson process $\mathbf{N}^\alpha$, with independent increments $N_t^\alpha - N_s^\alpha \sim \text{Pois}(\alpha(t - s))$.




Figure 4: Sample from Poisson process with unit rate.

Counting processes

A **counting process** $\mathbf{Z} = (Z_t)_{t \geq 0}$ is a càdlàg process with unit jumps, and $Z_0 = 0$ a.s.

Ex: Poisson process $\mathbf{N}^\alpha$, with independent increments $N_t^\alpha - N_s^\alpha \sim \text{Pois}(\alpha(t - s))$.




Figure 4: Sample from Poisson process with unit rate.

Compensator and intensity

The compensator of $\mathbf{Z}$, denoted $\mathbf{\Lambda}$, is a non-decreasing process such that $\mathbf{Z} - \mathbf{\Lambda}$ is a $\mathbf{F}$ -local martingale. For $\mathbf{N}^\alpha$, this is just $t \mapsto \alpha t$ since $\mathbb{E}[N_t^\alpha - N_s^\alpha | \mathcal{F}_s] = \alpha(t - s)$.




Figure 5: Sample from Poisson process with unit rate, and its compensator.

Compensator and intensity

The compensator of $\mathbf{Z}$, denoted $\mathbf{\Lambda}$, is a non-decreasing process such that $\mathbf{Z} - \mathbf{\Lambda}$ is a $\mathbf{F}$ -local martingale. For $\mathbf{N}^\alpha$, this is just $t \mapsto \alpha t$ since $\mathbb{E}[N_t^\alpha - N_s^\alpha | \mathcal{F}_s] = \alpha(t - s)$.




Figure 5: Sample from Poisson process with unit rate, and its compensator.

Compensator and intensity



The intensity λ is the "derivative" of the compensator: $\Lambda_t = \int_0^t \lambda_s ds$. For $\mathbf{N}^\alpha$, this is just $t \mapsto \alpha$. In general,

$$\mathbb{P}(\underbrace{Z_{t+h} - Z_t = 1}_{\text{a jump occurs between } t \text{ and } t+h} \mid \underbrace{\mathcal{F}_t}_{\text{given the process up to time } t \text{ is observed}}) = h \underbrace{\lambda_t}_{\text{intensity at time } t} + o(h).$$

If we know the intensity, we can more or less simulate the counting process $\mathbf{Z}$.

Mixed Poisson processes (Grandell, 1997)

What if $X \sim \pi$ is a random variable on $\mathbb{R}_+$, and conditionally on X , $\mathbf{Z}$ is a Poisson process with rate X ? $\mathbf{Z}$ is called a mixed Poisson process with (random) rate X .

- In particular, $Z_t|X \sim \text{Pois}(tX)$
 $\Rightarrow \mathbb{E}[Z_t] = \mathbb{E}[\mathbb{E}[Z_t|X]] = t\mathbb{E}[X]$ and $\mathbb{V}(Z_t) = tX$
- Almost surely, $\lim_{t \rightarrow \infty} Z_t/t = X$
- If we knew the intensity of $\mathbf{Z}$, we could get new sample from π !

Mixed Poisson processes (Grandell, 1997)

Theorem

If $\mathbf{Z}$ is the MPP with rate X , its intensity is $t \mapsto \mathbb{E}[X|\mathcal{F}_{t-}]$.




Figure 6: Mixed Poisson process, with rate X Gamma distributed, and its compensator.

Mixed Poisson processes

Theorem

If $\mathbf{Z}$ is the MPP with rate X , its intensity is $t \mapsto K_t^\pi(Z_{t-}) = \mathbb{E}[X|Z_{t-}]$.




Figure 7: Mixed Poisson process, with rate X Gamma distributed, and its compensator.

Generative modeling with MPPs




- We extend the previous construction to $X \sim \mathbb{R}_+^d$. Conditionally on X , $\mathbf{Z}$ is a *multivariate* Poisson process with rate X , and we define



$$[K_t^\pi(z_t)]_i = \mathbb{E}[X_i | Z_t = z_t], \quad 1 \leq i \leq d$$

- We learn the denoiser $\tilde{K}^\pi \approx K^\pi$ by minimizing $\int_0^T \mathbb{E}[\mathcal{D}(X, \tilde{K}_t^\pi(Z_t))] dt$
 - draw $t, x \sim \pi, z_t \sim \text{Pois}(tx)$
 - compute $\mathcal{D}(x, \tilde{K}_t^\pi(z_t))$
- Then, we simulate the MPP using the learned denoiser/intensity



Generative modeling with MPPs





Generative modeling with MPPs





Generative modeling with MPPs





Generative modeling with MPPs





Generative modeling with MPPs



Generative modeling with MPPs



Generative modeling with MPPs




Generative modeling with MPPs



Generative modeling with MPPs



Generative modeling with MPPs





3. Back to Poisson inverse problems

Posterior sampling with MPPs



Goal: sample from π^y , conditional distribution of X given $Y \sim \text{Pois}(\mathcal{A}(X))$, using the learned MPP with random rate $X \sim \pi$.

$$\pi^y(dx) \propto \underbrace{g(y|x)}_{\text{likelihood}} \underbrace{\pi(dx)}_{\text{prior}}$$

posterior

- We learned K^π , we can simulate the associated MPP and sample from π
- If we knew the conditional denoiser K^{π^y} , we could sample from the posterior π^y !

Poisson denoising





Theorem


When $\mathcal{A} = a\mathbf{I}_d$, i.e, $Y \sim \text{Pois}(aX)$, the conditional denoiser is

$$(t, z) \mapsto K_t^{\pi^y}(z) = K_{t+a}^{\pi}(z + y)$$




Poisson denoising







Poisson denoising







Poisson denoising




Poisson denoising




Poisson denoising




Poisson denoising




Poisson denoising




Poisson denoising




Poisson denoising



Poisson denoising



General forward operators




The diagram illustrates the relationship between the posterior, likelihood, and prior distributions. It features three colored boxes: a pink box for the posterior $\pi^y(dx)$, an orange box for the likelihood $g(y|x)$, and a blue box for the prior $\pi(dx)$. An orange arrow labeled "likelihood" points from the orange box to the pink box. A blue arrow labeled "prior" points from the blue box to the pink box. A red arrow labeled "posterior" points from the pink box down to the text below.

$$\pi^y(dx) \propto g(y|x) \pi(dx)$$



- Define $g_t(y|z_t) = \int g(y|x)\pi_t(dx|z_t)$, where $\pi_t(dx|z_t)$ is the distribution of X given $Z_t = z_t$: for large t , $\pi_t(dx|z_t) \approx \delta_{K_t^\pi(z_t)}(dx)$. We have

$$\begin{aligned} K_t^{\pi^y}(z_t)_i &= K_t^\pi(z_t)_i \frac{g_t(y|K_t^\pi(z_t + e_i))}{g_t(y|K_t^\pi(z_t))} \\ &\approx K_t^\pi(z_t)_i (1 + \partial_i \log g(y|K_t^\pi(z_t))) \end{aligned}$$




PET






PET







PET







PET






PET







PET







PET







PET



PET



PET




Conclusion



- Multivariate counting processes, and in particular mixed Poisson processes, are a natural framework to design "Poissonian diffusion models"
- For a simple Poisson inverse problem $Y \sim \text{Pois}(aX)$, we can use a pre-trained MPP with rate π (the prior) to simulate the MPP with rate π^Y (the posterior)
- For more general Poisson inverse problems, we can adapt several diffusion-based posterior sampling algorithms to this framework



list



Merci !

CEA SACLAY- LIST/DIN
Département Instrumentation Numérique
91191 Gif-sur-Yvette Cedex
thibaut.pellerin@cea.fr

References (1/2)

-  Aubert-Broche, B., Griffin, M., Pike, G. B., Evans, A. C., & Collins, D. L. (2006). **Twenty new digital brain phantoms for creation of validation image data bases.** *IEEE transactions on medical imaging*, 25(11), 1410–1416.
<https://doi.org/10.1109/TMI.2006.883453>
-  Chung, H., Kim, J., Mccann, M. T., Klasky, M. L., & Ye, J. C. (2023). **Diffusion posterior sampling for general noisy inverse problems.** *The Eleventh International Conference on Learning Representations*.
<https://openreview.net/forum?id=OnD9zGAGT0k>
-  Grandell, J. (1997). **Mixed Poisson Processes.** Chapman; Hall/CRC.
<https://doi.org/10.1201/9781003059950>

References (2/2)

-  McAlpine, S., Helly, J. C., Schaller, M., Trayford, J. W., Qu, Y., Furlong, M., Bower, R. G., Crain, R. A., Schaye, J., Theuns, T., Vecchia, C. D., Frenk, C. S., McCarthy, I. G., Jenkins, A., Rosas-Guevara, Y., White, S. D. M., Baes, M., Camps, P., & Lemson, G. (2016). **The eagle simulations of galaxy formation: Public release of halo and galaxy catalogues.** *Astronomy and Computing*, 15, 72–89. <https://doi.org/https://doi.org/10.1016/j.ascom.2016.02.004>
-  Song, J., Vahdat, A., Mardani, M., & Kautz, J. (2023). **Pseudoinverse-guided diffusion models for inverse problems.** *International Conference on Learning Representations*. https://openreview.net/forum?id=9_gsMA8MRKQ
-  Zhang, B., Chu, W., Berner, J., Meng, C., Anandkumar, A., & Song, Y. (2025). **Improving diffusion inverse problem solving with decoupled noise annealing.** *Proceedings of the Computer Vision and Pattern Recognition Conference*, 20895–20905.