Inference and Sampling of Point Processes from Diffusion Excursions

Ali Hasan [1] Yu Chen [2] Yuting Ng [1] Mohamed Abdelghani [3] Anderson Schneider [2] Vahid Tarokh [1]

[1]Duke University [2]Morgan Stanley [3]University of Alberta

Problem Statement

- Point processes model the distribution of discrete events.
- In many cases, events occur due to an unobserved *continuous process*.
- Understanding the distribution of latent continuous processes may provide additional insight on the behavior of the point process.

Diffusions, Excursions, and First Hitting Times

An *Itô diffusion* is given by the equation

$$\mathrm{d}Z_t = \mu(Z_t, t)\mathrm{d}t + \sigma(Z_t, t)\mathrm{d}W_t$$

An *excursion* is a continuous path of Z_t constrained to be positive or negative.

We can reconstruct a sample path of Z_t by concatenating excursions.

Another common way of using diffusions for representing point processes is through computing the first hitting time of a boundary. However, first hitting time models do not result in a single, continuous path.

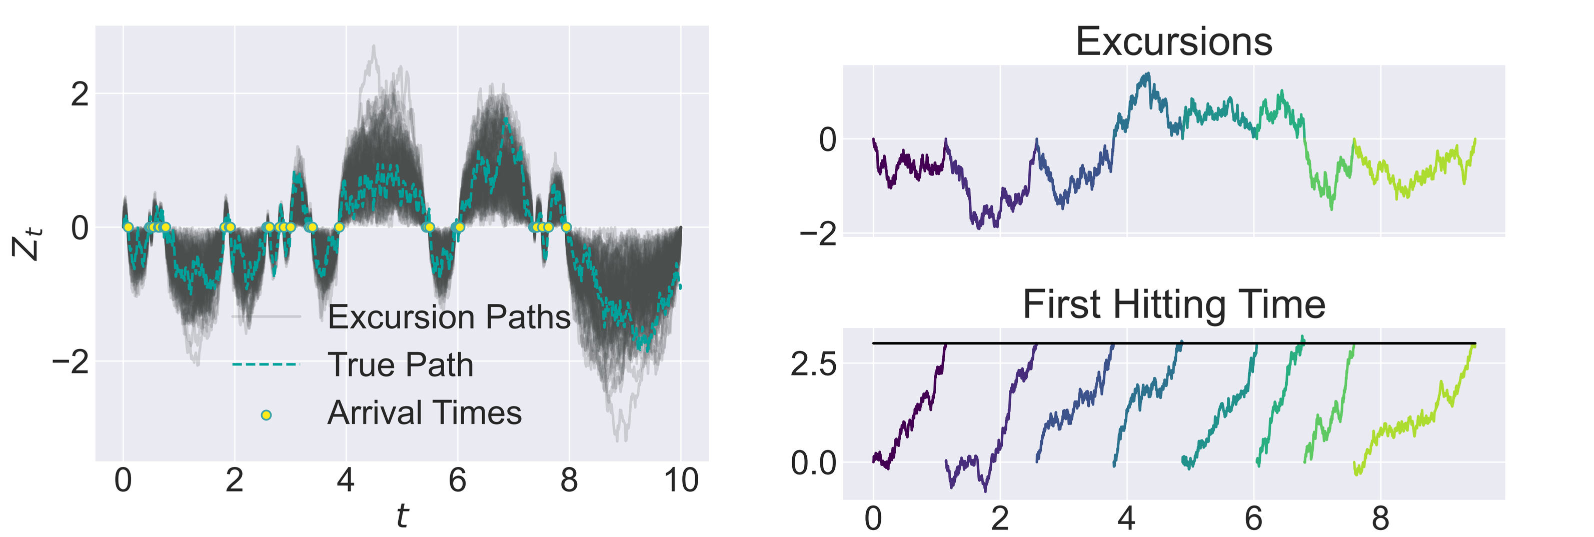

Figure 1. (Left) Decomposition of arrival times into distributions of Brownian excursions. Dashed green line represents the true unobserved signal and solid black lines represent possible excursions between observations of arrival times indicated by yellow circle markers. (Right) A comparison between excursions and first hitting times. Excursions lead to a continuous sample path while first hitting times have discontinuities.

Example Applications

- **Neuroscience:** Recover a continuous signal that generates neuronal spike trains.
- **Health Care:** Estimate a patient's health based on visits to a clinic.
- **Finance:** Infer a hidden price that governs an asset.
- **Social Science:** Recover the propagation of information within a social structure based on social media posts.

All assume the existence of a hidden, continuous process that models and generates point process observations.

Motivating Question

Given discrete observations of a temporal point process, can we model the arrivals of points as excursions of a diffusion?

Change-of-Measure Estimator

To estimate the parameters μ and σ, we consider the likelihood of observing a sequence of excursion lengths $\{\tau_1, \dots, \tau_n\}$. We formalize this in the following proposition.

Diffusion Excursion Density Let Z_t satisfy an SDE with drift μ such that Z_t is recurrent at zero. Then the density of the excursion lengths τ of Z_t is given by:

$$p_Z(\tau) = p_e(\tau;\delta)\,\mathbb{E}_{\mathbb{Q}^\delta_\updownarrow}\left[\exp\left(\int_0^\tau \mu(e_t, t;\theta)\mathrm{d}e_t - \frac{1}{2}\int_0^\tau \mu^2(e_t,t;\theta)\mathrm{d}t\right)\right].$$

$\mathbb{Q}^\delta_\updownarrow$ is the measure of excursion paths of length τ with minimum height δ.

Multi-Dimensional Point Processes

There are two ways to extend to multi-dimensional processes: 1) consider different types of excursions (e.g. positive or negative) and 2) consider a multivariate diffusion.

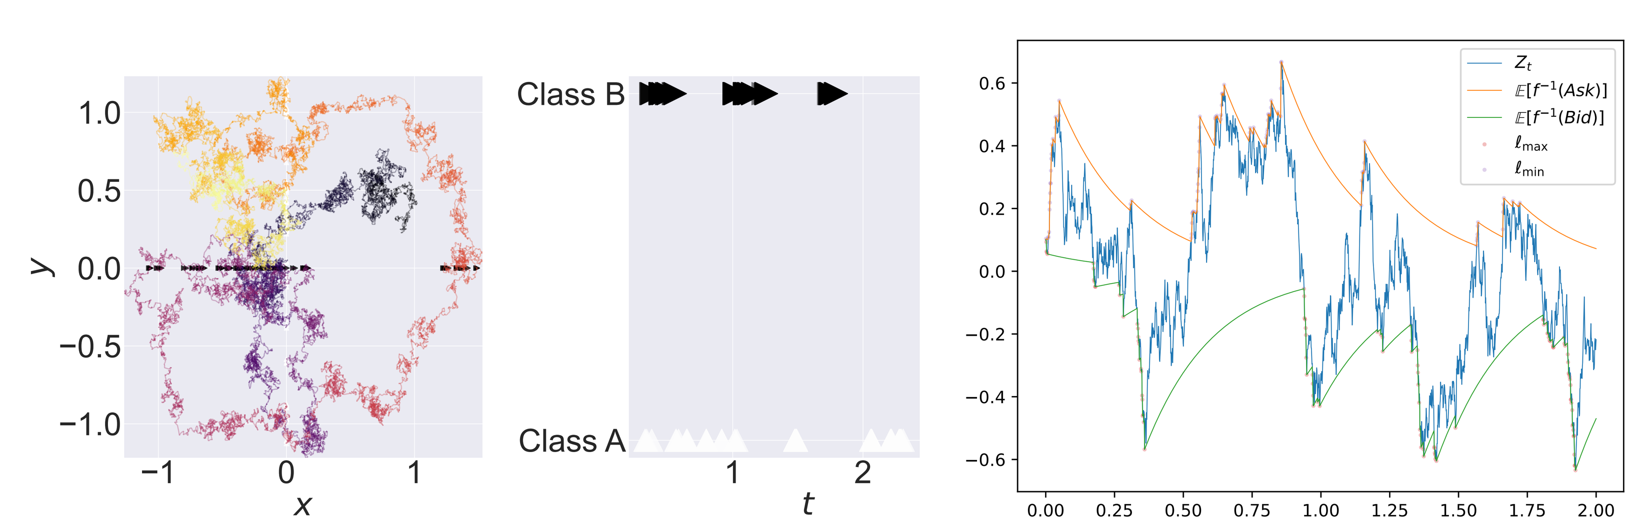

Figure 2. (Left) 2-dimensional latent diffusion crosses the axes producing points in the point process. (Center) Arrival times of the point process generated by the diffusion process on the left. White up triangles represent excursions from the y-axis and black left triangles represent excursions from the x-axis. (Right) An example of a continuous process that generates two classes of point processes in a financial market (bids in green and asks in orange).

Single Path For a single path, we can consider multiple boundaries that correspond to arrivals of different types of points. The estimator then changes to take into account the label for each of the arrivals. In Figure 2 on the right, the continuous process corresponds to a hidden "fair price" that the market is acting upon. Bids are generated with the fair price experiences an excursion below the boundary and and asks are generated as excursions above the boundary.

Multivariate Diffusion The diffusion then becomes multidimensional and the likelihood of the joint arrival is given by

$$p_Z(\tau^{(1)},\dots,\tau^{(d)}) = \prod_{k=1}^{d} p_e(\tau^{(k)};\delta)\mathbb{E}_{\mathbb{Q}_\updownarrow}\Bigg[\exp\Bigg(\int_0^{\left(\bigvee_{k=1}^d \tau^{(k)}\right)\wedge T} \mu(\mathbf{e}_t, t;\theta)\mathrm{d}\mathbf{e}_t - \frac{1}{2}\int_0^{\left(\bigvee_{k=1}^d \tau^{(k)}\right)\wedge T} \mu^\dagger\mu(\mathbf{e}_t,t;\theta)\mathrm{d}t\Bigg)\Bigg].$$

Reconstructing a Stimulus from Neuron Spikes

As an example application, we consider the problem of reconstructing a latent stimulus from spike train observations.

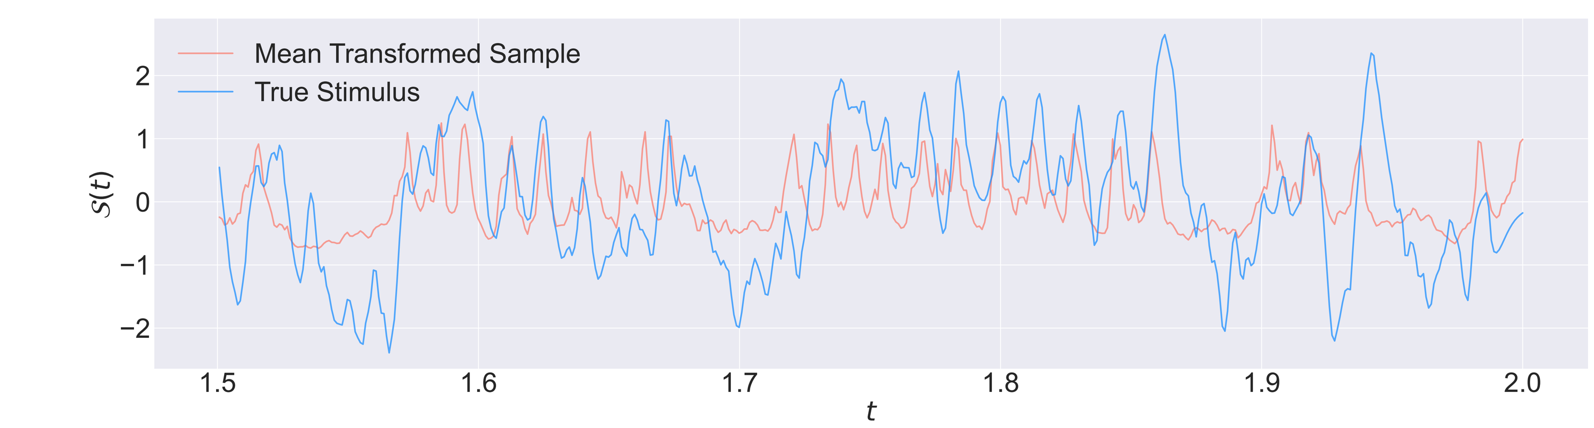

Figure 3. Average of learned sample paths (red) on $t \in [1.5, 2]$ compared with the true stimulus (blue). Spikes of the learned stimulus align well with the true stimulus.

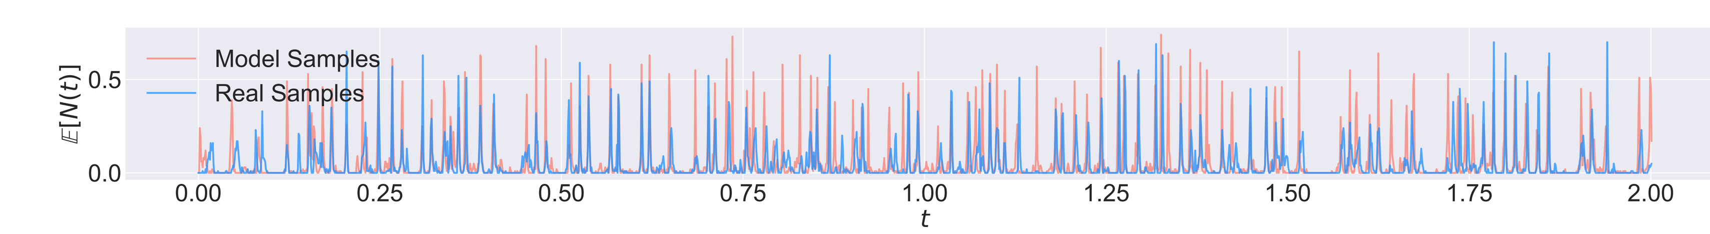

Figure 4. Comparison of the histogram of observed arrival times (blue) with the histogram of generated arrival times (red). Sampled spikes align well with true spikes.

Approximating Renewal Processes

In addition to the theoretical proof of expressiveness, we empirically validate our method on the approximation of different renewal processes from observations.

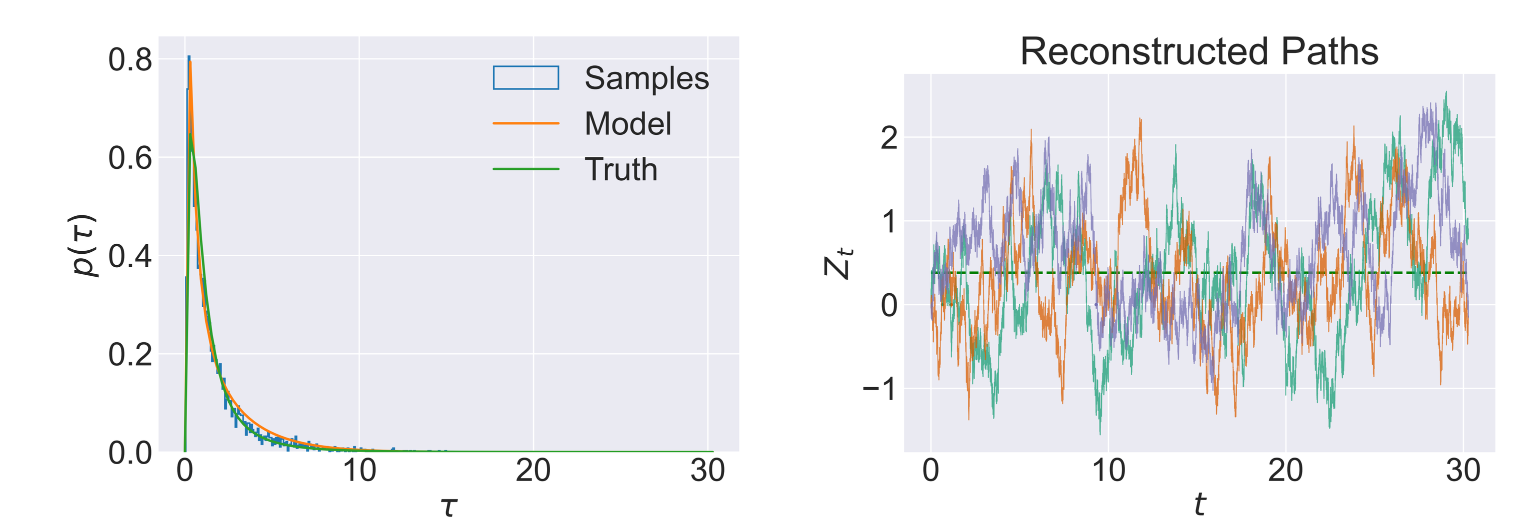

Figure 5. Example of estimated log-normal renewal process with samples generated from learned diffusion. (Left) Histogram of samples (blue) compared with the true density (green) and estimated density (orange). (Right) Learned sample paths with excursion lengths corresponding to the histogram. The dashed green line corresponds to minimum height δ.

Future Directions

1. Extension to higher dimensional processes such as spatial-temporal point processes.
2. Identifiability theory to determine under what conditions we can recover the parameters of the diffusion.

Conference on Uncertainty in Artificial Intelligence 2023 — UAI-2023

ali.hasan@duke.edu