Fast and Flexible Temporal Point Processes with Triangular Maps

tl;dr

- Represent temporal point process densities as transformations
- Define flexible, tractable and efficient (parallelizable) transformations on sequences as compositions of simple maps
- Reparametrization trick + differentiable relaxation enable learning TPPs with sampling-based losses
- Use TPPs for variational inference in Markov jump processes

What is a temporal point process?

Temporal point process (TPP) is a probability distribution over variablelength event sequences $\mathbf{t} = (t_1, \dots, t_N)$ on an interval [0, T]



How are TPPs usually defined?

Conditional intensity function $\lambda^*(t)$ describes the rate of arrival new events at time t given the history

$$p(t) = \left(\prod_{i}^{N} \lambda^{*}(t_{i})\right) \exp\left(-\int_{0}^{T} \lambda^{*}(u) du\right)$$

State-of-the-art approaches parametrize $\lambda^*(t)$ with autoregressive neural nets

✓ Autoregressive models are flexible X Sequential sampling is extremely slow

Research question

How can we define TPPs that are

- as flexible as autoregressive models?
- where we can both sample & compute likelihood in parallel?





Defining TPPs as transformations

Integrated intensity function (compensator)

$$\Lambda^*(t_i) \coloneqq \Lambda(t_i | t_1, \dots, t_{i-1}) = \int_0^t \lambda^*(u) du$$

Random time change theorem

$$\begin{bmatrix} z_1 \\ z_2 \\ \vdots \\ z_N \end{bmatrix} = \begin{bmatrix} \Lambda(t_1) \\ \Lambda(t_2|t_1) \\ \vdots \\ \Lambda(t_N|t_1, \dots, t_{N-1}) \end{bmatrix} = \mathbf{F}\left(\begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_N \end{bmatrix} \right)$$

The sequence z = F(t) is distributed according to the Standard Poisson Process

The triangular map *F* defines the TPP density

$$p(\mathbf{t}) = \left(\prod_{i=1}^{N} \frac{\partial \Lambda^{*}(t_{i})}{\partial t_{i}}\right) \exp(-\Lambda^{*}(T))$$

Inverse transform sampling for TPPs

- Sample z from the standard Poisson process
- 2. Transform $t = F^{-1}(z)$
- Discard events t_i after T

Main idea

Define TPPs by directly specifying the triangular map *F*

How should we parametrize *F*?

The parametrization of *F* must be

- flexible able to represent different distributions
- fast we can compute F(t) and $F^{-1}(z)$ analytically in O(N) parallel operations

Solution: Define *F* as a composition of easy-to-invert transformations



 \rightarrow New efficient parametrizations for existing TPP models

 \rightarrow TriTPP: A new flexible model where all operations can be done in parallel





Maximum mean discrepancy between model samples and the test set







Learning TPPs with sampling-based losses

Sampling losses are common in variational inference, reinforcement learning

 $\mathcal{L}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{t} \sim p_{\boldsymbol{\theta}}(\boldsymbol{t})}[g(\boldsymbol{t})]$

We show how to optimize such objectives thanks to

1. Reparametrization trick for TPPs that allows computing $\nabla_{\theta} \mathcal{L}$ with Monte Carlo

$$\mathbb{E}_{\boldsymbol{t} \sim p_{\boldsymbol{\theta}}(\boldsymbol{t})}[g(\boldsymbol{t})] = \mathbb{E}_{\boldsymbol{z} \sim \text{SPP}}\left[g\left(\boldsymbol{F}_{\boldsymbol{\theta}}^{-1}(\boldsymbol{z})\right)\right]$$

2. Differentiable relaxation to \mathcal{L} that removes discontinuities w.r.t. $\boldsymbol{\theta}$

- \rightarrow Instead of discarding $t_i > T$, use indicator functions
- \rightarrow Sigmoid relaxes the indicator function

$$\mathbb{I}(t_i < T) \approx \sigma(T - t_i)$$

Scalability analysis



Density estimation

Test set negative log-likelihood

	Hawkes1	Hawkes2	SC	IPP	MRP	RP	PUBG	Reddit-C	Reddit-S	Taxi	Twitter	Yelp1	Yelp2
IPP	1.06	1.03	1.00	0.71	0.70	0.89	-0.06	-1.59	-4.08	-0.68	1.60	0.62	-0.05
RP	0.65	0.08	0.94	0.85	0.68	0.24	0.12	-2.08	-4.00	-0.58	1.20	0.67	-0.02
MRP	0.65	0.07	0.93	0.71	0.36	0.25	-0.83	-2.13	-4.38	-0.68	1.23	0.61	-0.10
Hawkes	0.51	0.06	1.00	0.86	0.98	0.39	0.11	-2.40	-4.19	-0.64	1.04	0.69	0.01
RNN	0.52	-0.03	0.79	0.73	<u>0.37</u>	0.24	-1.96	-2.40	-4.89	-0.66	1.08	0.67	-0.08
TriTPP	0.56	<u>0.00</u>	<u>0.83</u>	0.71	0.35	0.24	-2.41	-2.36	<u>-4.49</u>	-0.67	<u>1.06</u>	0.64	<u>-0.09</u>

	Hawkes1	Hawkes2	SC	IPP	MRP	RP	PUBG	Reddit-C	Reddit-S	Taxi	Twitter	Yelp1	Yelp2
IPP	0.08	0.09	0.58	0.02	0.15	0.07	0.01	0.10	0.21	0.10	0.16	0.15	0.16
RP	0.06	0.06	1.13	0.34	1.24	0.01	0.46	0.07	0.18	0.57	0.14	0.16	0.23
MRP	0.05	0.06	0.50	0.02	<u>0.11</u>	0.02	<u>0.12</u>	0.09	0.20	<u>0.09</u>	0.13	<u>0.13</u>	<u>0.16</u>
Hawkes	<u>0.02</u>	0.04	0.58	0.36	0.65	0.05	0.16	0.04	0.35	0.20	0.20	0.20	0.32
RNN	0.01	0.02	0.19	0.09	0.17	0.01	0.23	0.04	0.09	0.13	0.08	0.19	0.18
TriTPP	0.03	<u>0.03</u>	<u>0.23</u>	0.02	0.08	0.01	0.16	0.07	<u>0.16</u>	0.08	0.08	0.12	0.14

Variational inference in Markov jump processes

- Goal: Infer the posterior distribution over the latent state in a continuoustime discrete-state system
- Approximate posterior over the jump times modeled with TriTPP
- ELBO can be optimized thanks to fast sampling + reparametrization trick + differentiable relaxation

