RIFT: RECTIFIED LOG-INTENSITY FLOWS FOR TEMPORAL POINT PROCESSES

Anonymous authors
Paper under double-blind review

ABSTRACT

Generative models of temporal point processes (TPPs) either generate events one at a time, which accumulates errors and is slow for long horizons, or operate directly on event times, which gives up the intensity function that defines a point process. We propose RIFT (Rectified log-Intensity Flows for Temporal point processes), which generates a whole event sequence by transporting a homogeneous Poisson intensity to the data intensity along a straight path in *log-intensity* space. Working in log-intensity space keeps every intermediate intensity positive, and each transition between consecutive steps is realized exactly by Poisson thinning and superposition. As a result, the proposed RIFT never has to move individual event times. We use a transformer to predict the log-intensity velocity at any query time given the current event set. RIFT is trained by maximizing a Poisson likelihood of the observed event sequence under the intensity obtained by integrating the velocity. Furthermore, a straightness penalty is added which enforces the velocity to be constant along the path, and requires a few steps for the denoising process (RIFT samples sequences with $K = 5$ steps).On a benchmark of 13 synthetic and real-world datasets RIFT consistently to be the top two performers in terms of the the kernel density and forecasting tasks, which shows great potential of the proposed RIFT.

1 INTRODUCTION

Temporal point processes (TPPs) provide a principled probabilistic framework for modeling the occurrence of random events over continuous time. Their capacity to capture complex temporal dependencies has led to applications in neuroscience (Linderman, 2016; Apostolopoulou et al., 2019), healthcare (Yan et al., 2019; Enguehard et al., 2020), finance (Hawkes, 2018; Fan et al., 2023), and ride-hailing demand forecasting (Du et al., 2016; Okawa et al., 2019).

Autoregressive neural TPPs (Shchur et al., 2020a; Zuo et al., 2020; Shchur et al., 2020b; Zhou et al., 2020; Lin et al., 2022) generate one event at a time, conditioning each new event on a representation of the history. Sampling a sequence therefore needs as many sequential network evaluations as there are events, and errors in early events propagate to later ones.

A recent line of work instead generates the whole sequence at once. Add-Thin (Lüdke et al., 2023) defines a diffusion process that adds and thins events and learns its posterior with a classifier for kept events and an intensity model for the removed ones, however, it usually needs $K = 100$ denoising steps. EventFlow (Kerrigan et al., 2025) applies flow matching to event times, obtaining strong forecasting results with 25 function evaluations. EdiTPP (Lüdke et al., 2026) learns an edit-based flow of insertions, deletions and substitutions and samples in milliseconds. Point Set Diffusion (Lüdke et al., 2025) diffuses point sets for conditional generation and imputation. These methods either move individual event times or learn discrete edit rates; none of them transports the *intensity function* itself, although the intensity function is the object that defines a TPP and its likelihood.

We propose RIFT (Rectified log-Intensity Flows), a flow on log-intensity functions inspired by rectified flow (Liu et al., 2023; Lee et al., 2024). RIFT moves from a homogeneous Poisson intensity $\lambda_1(t) = \lambda_1$ to the ground-truth data intensity function $\lambda_0(t)$ along the straight path $\log \lambda_r(t) = r \log \lambda_1 + (1 - r) \log \lambda_0(t)$, where $r \in [0, 1]$ is the transportation step. Three properties follow. (i) Every intermediate intensity is positive since the intensity is in log space. (ii) A step between two intensities is an *exact* point-process operation: thinning existing events where the intensity decreases and superposing new events where it increases, so events are never moved and the number of events

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

is determined by the intensity function. (iii) The ground-truth velocity $\log \lambda_1 - \log \lambda_0(t)$ does not depend on r, so the flow is straight and fewer Euler steps are required to generate sequences.

Our contributions are:

- **A flow on log-intensities for TPPs.** We define the log-linear intensity path, show how each discretised transition is realized exactly by the `Scaled-Sampling` operation (thinning and superposition Function 1), and derive a sampler that tracks the intensity exactly along the discretised path (Algorithm 2).
- **Likelihood-based training without the true intensity.** The velocity field $\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)$ is trained by maximizing a Poisson log-likelihood of the observed event sequence under the intensity function derived from the integrated velocity, with a kernel estimate used only to build noised training sequences. We make the underlying modelling assumption explicit and characterize which processes this family can and cannot represent (Section 4.6).
- **Few-step generation with an intensity-aware architecture.** A transformer encodes the current event set once and reads out the velocity at arbitrary query times by cross-attention, so the cost of a step is $\mathcal{O}(N^2 + QN)$ for N events and Q query times. RIFT samples with $K = 5$ steps, and we report its sampling time and its quality as a function of K.

2 Related Work

Autoregressive neural TPPs. Autoregressive TPPs parameterise each event's conditional distribution given the history with recurrent or attention-based encoders (Zuo et al., 2020; Shchur et al., 2020a; Zhou et al., 2020; Lin et al., 2022) or neural intensity integrals (Omi et al., 2019). TriTPP (Shchur et al., 2020b) parallelises sampling with triangular maps. Sampling remains sequential in the number of events for most models.

Whole-sequence generative TPPs. Add-Thin (Lüdke et al., 2023) is a diffusion model on event sets whose forward process superposes Poisson noise and thins data events; the reverse process is learned with a classifier of kept events and an intensity model of the remaining ones, and uses $K = 100$ steps. Point Set Diffusion (Lüdke et al., 2025) extends set diffusion to conditional generation and arbitrary-mask imputation. EventFlow (Kerrigan et al., 2025) applies flow matching to the joint distribution of event times for a sampled number of events, reporting results with 25, 10 and 1 function evaluations. EdiTPP (Lüdke et al., 2026) learns an edit flow (insertions, deletions, substitutions) between noise and data sequences. GD (Lin et al., 2022) applies Gaussian diffusion to inter-event times within an autoregressive model. RIFT differs from all of these in what is transported: the intensity function, along a straight path in log space, with exact thinning/superposition transitions and a likelihood-based objective.

Rectified flows and flow matching. Rectified flow and flow matching (Liu et al., 2023; Lipman et al., 2023; Albergo et al., 2023) learn velocity fields along straight interpolation paths and allow few-step sampling; practical choices such as time-step sampling and weight averaging matter for quality (Esser et al., 2024). RIFT carries the straight-path idea over to log-intensities of point processes, where the regression target is unavailable and a likelihood objective replaces the velocity-matching loss.

3 Methodology

Setting. Let $\mathbf{e}_0 = \{t_1, \ldots, t_n\} \subset [0, T]$ be an observed event sequence. We model each sequence as a Poisson process given a sequence-specific intensity $\lambda_0(t)$, i.e., $\mathbf{e}_0 \sim \mathrm{PP}(\lambda_0(t))$. Appendices H and G motivate transporting the intensity and relate RIFT to Gaussian flow matching, which analogue the point-process as $\mathcal{N}(\mathbf{0}, \mathbf{I})$ in flow matching, with closed-form sampling and likelihood, and independent increments without dependence structure. The velocity network is conditioned on $\log(\lambda_1/\bar{\lambda}_1)$, with $\bar{\lambda}_1$ the mean training rate. Every step of the flow remains an exact thinning/superposition of a Poisson process given λ_1. The lengths of $\mathbf{e}_0$ and $\mathbf{e}_1$ generally differ. $\boldsymbol{\phi}$ denotes the parameters of RIFT.

Function 1 `Scaled-Sampling` ($\overline{\mathbf{e}}$, $\overline{\lambda}(t)$, $\overline{\omega}(t)$)

Input: Current event sequence $\overline{\mathbf{e}}$, current intensity function $\overline{\lambda}(t)$, scaling factor $\overline{\omega}(t) > 0$
Output: New event sequence $\mathbf{e}_{\text{new}} \sim \text{PP}(\overline{\lambda}(t)\overline{\omega}(t))$

$\mathbf{e}_{\text{new}} = \emptyset$
For every $t_i \in \overline{\mathbf{e}}$: append t_i to $\mathbf{e}_{\text{new}}$ with probability $\min(\overline{\omega}(t_i), 1)$ ▷ thinning
$\widetilde{\mathbf{e}} \sim \text{PP}\big(\max(\overline{\omega}(t) - 1, 0)\,\overline{\lambda}(t)\big)$ by rejection sampling; $\mathbf{e}_{\text{new}} = \mathbf{e}_{\text{new}} \cup \widetilde{\mathbf{e}}$ ▷ superposition
return $\mathbf{e}_{\text{new}}$

Algorithm 1 Forward process $\{\mathbf{e}_{r_k}\}_{k=1}^K$ for one training sequence

Input: Event sequence $\mathbf{e}_{r_0} = \mathbf{e}_0$, base rate λ_1, grid $r_0 = 0 < r_1 < \ldots < r_K = 1$
Output: $\{\mathbf{e}_{r_k}\}_{k=1}^K$

for $k = 1, \ldots, K$ **do**
  $\mathbf{e}_{r_k} \sim$ `Scaled-Sampling`$\big(\mathbf{e}_{r_{k-1}},\ \lambda_1^{r_{k-1}}\widehat{\lambda}_0^{1-r_{k-1}}(t),\ \lambda_1^{r_k - r_{k-1}}\widehat{\lambda}_0^{r_{k-1}-r_k}(t)\big)$
end for
return $\{\mathbf{e}_{r_k}\}_{k=1}^K$

3.1 Rectified log-Intensity Flows (RIFT)

Naive approach. A direct analogue of rectified flow would regress a velocity of intensity functions onto $\lambda_1 - \lambda_0(t)$. However, such an approach would cause two problems: (1) since the learned velocity might be negative, the inferred intensity function might have negative values which is incorrect in practice; (2) $\lambda_0(t)$ is not observed.

Log-intensity path. RIFT instead defines the velocity in log-intensity space,

$$\mathrm{d}\log\lambda_r(t)/\mathrm{d}r = \log\lambda_1 - \log\lambda_0(t), \quad \forall\, r \in (0,1), \tag{1}$$

The resulting path is the log-linear interpolation

$$\lambda_r(t) = \lambda_1^r \lambda_0^{1-r}(t) = \lambda_0(t)\,(\lambda_1/\lambda_0(t))^r\,, \tag{2}$$

which is positive for every r and equals λ_1 at $r = 1$ (left panel of Figure 1).

Scaled-Sampling. Moving along the path multiplies the intensity by a positive factor. `Scaled-Sampling` (Function 1) realises such a multiplication on an event set: given events $\overline{\mathbf{e}}$ with intensity $\overline{\lambda}(t)$ and a factor $\overline{\omega}(t) > 0$, it keeps each event $t_i \in \overline{\mathbf{e}}$ with probability $\min(\overline{\omega}(t_i), 1)$ and superposes new events from $\text{PP}\big(\max(\overline{\omega}(t) - 1, 0)\,\overline{\lambda}(t)\big)$, drawn by rejection sampling.

Remark 3.1 (Correctness of `Scaled-Sampling`). If $\overline{\mathbf{e}} \sim \text{PP}(\overline{\lambda}(t))$, then `Scaled-Sampling`$(\overline{\mathbf{e}}, \overline{\lambda}(t), \overline{\omega}(t)) \sim \text{PP}(\overline{\lambda}(t)\overline{\omega}(t))$, by the thinning and superposition theorems for Poisson processes (Daley & Vere-Jones, 2007). The proof is in Appendix A. The statement is exact for the backward (denoising) process, whose input is Poisson by construction; for the forward process it holds to the extent that the observed $\mathbf{e}_0$ behaves like a sample of $\text{PP}(\widehat{\lambda}_0(t))$ (see below).

3.2 Forward process ($\mathbf{e}_0 \to \mathbf{e}_r$)

Since $\lambda_0(t)$ is unknown, the forward process uses a kernel estimate built from the sequence itself,

$$\widehat{\lambda}_0(t) = \sum_{t_i \in \mathbf{e}_0} \kappa_h(t - t_i), \qquad \kappa_h(u) = \tfrac{1}{h}\phi(u/h), \tag{3}$$

with a Gaussian kernel ϕ; note that $\int_0^T \widehat{\lambda}_0 \approx n$, so the number of expected events is n. The bandwidth h is a per-dataset hyperparameter (Appendix K; sensitivity in Section 4.5). Noised sequences are produced by the forward process on the grid $0 = r_0 < r_1 < \cdots < r_K = 1$ (Algorithm 1). In each transportation step r_k, $\mathbf{e}_{r_k}$ is obtained from $\mathbf{e}_{r_{k-1}}$ by `Scaled-Sampling` with $\overline{\lambda} = \widehat{\lambda}_{r_{k-1}}$ and $\omega_{r_k}(t) = \frac{\widehat{\lambda}_{r_k}(t)}{\widehat{\lambda}_{r_{k-1}}(t)} = \left(\frac{\lambda_1}{\widehat{\lambda}_0(t)}\right)^{r_k - r_{k-1}}$.

Algorithm 2 Backward denoising process

Input: Trained velocity $\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)$, grid $r_0 = 0 < r_1 < \ldots < r_K = 1$, base rate λ_1
Output: Generated event sequence $\mathbf{e}_0$

$n \sim$ training counts; $\lambda_1 = (m + s(n-m))/T$; $\mathbf{e}_{r_K} \sim \mathrm{HPP}(\lambda_1)$; $\lambda_{r_K}(t) = \lambda_1$
for $k = K, \ldots, 1$ **do**
  $\omega_{r_{k-1}}(t) = \exp\left(\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t)\,(r_{k-1} - r_k)\right)$ ▷ velocity held constant for $r \in [r_{k-1}, r_k]$
  $\mathbf{e}_{r_{k-1}} = \texttt{Scaled-Sampling}(\mathbf{e}_{r_k}, \lambda_{r_k}(t), \omega_{r_{k-1}}(t))$
  $\lambda_{r_{k-1}}(t) = \lambda_{r_k}(t) \cdot \omega_{r_{k-1}}(t)$ ▷ exact intensity tracking
end for
return $\mathbf{e}_{r_0}$

It is noted that the kernel estimate is only used to generate noised inputs. The velocity network never sees $\widehat{\lambda}_0(t)$ and is trained on the observed event times (Section 3.4). Appendix F compares this forward process with one driven by the true generating intensity on synthetic data. The chain can instead be started from $\widehat{\mathbf{e}}_0 \sim \mathrm{PP}(\widehat{\lambda}_0)$, which makes every $\mathbf{e}_{r_k}$ exactly $\mathrm{PP}(\widehat{\lambda}_{r_k})$ by Remark 3.1; we found this to be worse (Section 4.5), because the resampled inputs no longer carry information about interactions between the observed events.

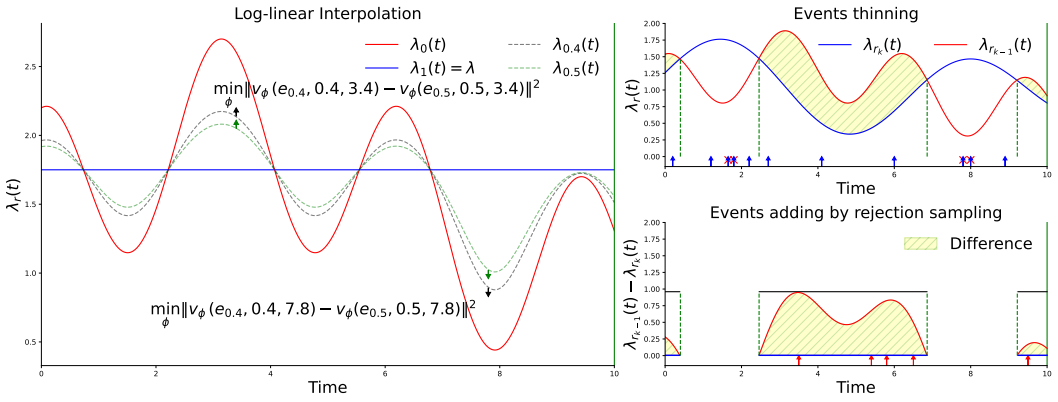


Figure 1: Log-linear interpolation between the data intensity $\lambda_0(t)$ (red) and the homogeneous Poisson intensity λ_1 (blue). Dashed lines show $\lambda_{0.4}(t)$ and $\lambda_{0.5}(t)$. **Left panel:** The red solid line represents the ground-truth intensity function $\lambda_0(t)$ and the blue solid line is the constant base intensity $\lambda_1(t) = \lambda_1$. The black and green dashed lines show the noised intensity functions $\lambda_{0.4}(t)$ and $\lambda_{0.5}(t)$ at transition steps $r = 0.4$ and $r = 0.5$, respectively. These intermediate intensities are obtained via the log-linear interpolation in Eq. (2). Minimizing the straightness loss (Eq. (6)) equalizes the velocities $\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{0.4}, 0.4, t)$ and $\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{0.5}, 0.5, t)$ at the same query time t, encouraging the trajectories to be straight in log-intensity space. **Right panels** illustrate thinning (red $\times$) and superposition (red $\uparrow$) in `Scaled-Sampling`. The blue solid line is the intensity $\lambda_{r_k}(t)$ and the red solid line is $\lambda_{r_{k-1}}(t)$. Blue arrows ($\uparrow$) mark the events in $\mathbf{e}_{r_k}$, red arrows ($\uparrow$) mark newly generated events in $\mathbf{e}_{r_{k-1}}$, and red crosses ($\times$) indicate thinned events. For each event $t_i \in \mathbf{e}_{r_k}$: if $\lambda_{r_{k-1}}(t_i) < \lambda_{r_k}(t_i)$, the event is retained with probability $\lambda_{r_{k-1}}(t_i)/\lambda_{r_k}(t_i)$ (thinning); otherwise it is always kept. In regions where $\lambda_{r_{k-1}}(t) > \lambda_{r_k}(t)$, new events are generated via rejection sampling from a Poisson process with intensity $\max(0, \lambda_{r_{k-1}}(t) - \lambda_{r_k}(t))$. The final sequence $\mathbf{e}_{\text{new}}$ is the union of the retained events from $\mathbf{e}_{r_k}$ and the newly generated events.

3.3 Backward Denoising Process ($\mathbf{e}_1 \to \mathbf{e}_0$)

Velocity network. A network $\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)$ is used to approximate $\mathrm{d}\log\lambda_r(t)/\mathrm{d}r$ at query time t given the current event set $\mathbf{e}_r$ and step r. Although the ideal velocity is independent of r, the input $\mathbf{e}_r$ carries different amounts of information about λ_0 at different r, so the network is conditioned on r.

Discretisation. We use Euler method to approximate the log-intensity function, which assumes the velocity function is constant on each segment $[r_{k-1}, r_k]$. The log-intensity after K steps is:

$$\log \lambda_0(t) \approx \log \lambda_1 + \sum_{k=1}^{K} \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t)\,(r_{k-1} - r_k), \tag{4}$$

As a result, the factor function between consecutive steps is $\omega_{r_{k-1}}(t) = \lambda_{r_{k-1}}(t)/\lambda_{r_k}(t) = \exp\big(\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t)(r_{k-1} - r_k)\big)$. If the learned velocity is constant in r, Eq. (4) is exact for any K.

Sampling. Starting from $\mathbf{e}_1 \sim \mathrm{PP}(\lambda_1)$, Algorithm 2 applies `Scaled-Sampling` K times and tracks $\lambda_{r_k}(t)$ exactly by accumulating $\log \omega_{r_k}(t)$ on a fine time grid. Since the input of every step is Poisson given the previous steps, each step is exact by Remark 3.1; the only approximation is the learned velocity.

Architecture. The velocity network is a diffusion-transformer-style model (Peebles & Xie, 2023). Event times in $\mathbf{e}_r$ are embedded with sinusoidal features followed by an MLP and encoded once by self-attention blocks whose layer norms are modulated by an embedding of r (adaLN). Each query time t is embedded with the same time embedding, added to a pooled summary of the encoded events, and passed through adaLN cross-attention blocks that attend to the encoded events; a small MLP head outputs $\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)$. Encoding once and reading out by cross-attention makes a step cost $\mathcal{O}(N^2 + QN)$ for N events and Q queries, and all sequences of a batch are processed in one padded call. Details are in Appendix K.

3.4 Training RIFT

The velocity-matching loss $\mathbb{E}\|\mathbf{v}_{\boldsymbol{\phi}} - (\log \lambda_1 - \log \lambda_0(t))\|^2$ is unavailable because $\lambda_0(t)$ is unknown. We have used kernel estimation $\widehat{\lambda}_0(t)$ to replace $\lambda_0(t)$ and explored such performance in Appendix J, with discussions in Appendix I.

RIFT trains on the observed event sequence with a likelihood objective and a straightness regularizer.

Likelihood loss. Plugging the discretised intensity of Eq. (4) into the Poisson log-likelihood $\sum_i \log \lambda(t_i) - \int_0^T \lambda(t)\,\mathrm{d}t$ of the observed sequence gives

$$\mathcal{L}_{\text{ll}}(\boldsymbol{\phi}) = n \log \lambda_1 + \sum_{i=1}^{n} \sum_{k=1}^{K} \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t_i)\,\Delta r_k - \lambda_1 \int_0^T \exp\Big(\sum_{k=1}^{K} \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t)\,\Delta r_k\Big)\mathrm{d}t, \tag{5}$$

with $\Delta r_k = r_{k-1} - r_k$ and $\{\mathbf{e}_{r_k}\}$ from Algorithm 1. The integral is estimated with stratified random quadrature: one uniformly drawn point in each of M equal cells, redrawn at every training step and shared across k. $\mathcal{L}_{\text{ll}}$ uses the sequence's own forward chain and is therefore a training surrogate, not a held-out log-likelihood.

Straightness loss. To keep trajectories close to the straight path, consecutive velocities at the same query time are encouraged to agree:

$$\mathcal{L}_{\text{straight}}(\boldsymbol{\phi}) = \frac{1}{K} \sum_{k=1}^{K} \mathbb{E}_t\Big[\big(\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t) - \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_{k-1}}, r_{k-1}, t)\big)^2\Big]. \tag{6}$$

The pairs $(\mathbf{e}_{r_{k-1}}, \mathbf{e}_{r_k})$ are on the same forward process trajectory from the same data sequence (Algorithm 1; $\mathbf{e}_{r_0} = \mathbf{e}_0$ for $k = 1$), evaluated at the same quadrature times t as $\mathcal{L}_{\text{ll}}$. Since this chain is stochastic and not a trajectory of the learned ODE, $\mathcal{L}_{\text{straight}}$ is a *pathwise-consistency* regulariser favouring a velocity that is constant along the path (straight in log-intensity), not a guarantee that the learned flow is straight.

Objective. The training objective is

$$\mathcal{L}(\boldsymbol{\phi}) = \alpha_{\text{straight}}\,\mathcal{L}_{\text{straight}}(\boldsymbol{\phi}) - \alpha_{\text{ll}}\,\mathcal{L}_{\text{ll}}(\boldsymbol{\phi}), \tag{7}$$

with $\alpha_{\text{straight}}, \alpha_{\text{ll}} \geq 0$.

Training cost and simulation. It is noted that the training does not simulate the learned ODE. All the event sequences $\{\mathbf{e}_{r_k}\}_k$ are generated through one realization of the forward process, which

is independent of ϕ. With $\lambda_0(t)$ unobserved, the likelihood constrains only $\sum_k \mathbf{v}_k \Delta r_k$, so each sequence needs all K velocity evaluations.

Conditional generation for forecasting. For forecasting, the events before a window are encoded by a recurrent history encoder whose output is added to the step embedding in every adaLN block, and RIFT generates the events inside the window (rescaled to $[0, T]$). For forecasting we use a single global reference rate: the sequence-level rate is drawn independently of the history and degrades forecasts.

4 EXPERIMENTS

We evaluate RIFT on (1) *density estimation*: unconditional generation of whole sequences, compared with held-out test sequences; and (2) *forecasting*: generation of the events in a future window given the history.

4.1 EXPERIMENTAL SETUP

Datasets. We use the 13-dataset benchmark of (Shchur et al., 2020b; Lüdke et al., 2023), also used by EventFlow and EdiTPP. Six synthetic datasets follow (Omi et al., 2019): two Hawkes processes (H1, H2) (Hawkes, 1971), a self-correcting process (SC) (Isham & Westcott, 1979), an inhomogeneous Poisson process (IPP), a stationary renewal process (RP) and a non-stationary renewal process (MRP) (Cox, 1972; Sevast'yanov, 1975). Seven real-world datasets are PUBG, Reddit comments (RedditC) and submissions (RedditS), Taxi, Twitter, and two Yelp check-in datasets (Yelp1, Yelp2) (Shchur et al., 2020b). Sequences are split 60/20/20 into train/validation/test with the split seed of (Lüdke et al., 2023). Forecasting uses the seven real-world datasets.

Baselines. Regarding baseline methods, we use autoregressive methods which include IF-RNN (Shchur et al., 2020a), IF-TF (the intensity-free decoder with the Transformer encoder of (Zuo et al., 2020)) and TriTPP (Shchur et al., 2020b); GD (Lin et al., 2022), as well as whole-sequence models which include Add-Thin (Lüdke et al., 2023) ($K = 100$), EventFlow (Kerrigan et al., 2025) (25 function evaluations) and EdiTPP (Lüdke et al., 2026). All baseline numbers are taken from the cited papers; we did not re-run the baselines. IF-RNN, IF-TF, TriTPP, GD and Add-Thin are reported by Lüdke et al. (2023) (MMD on 4,000 samples), EventFlow by Kerrigan et al. (2025) (MMD on 1,000 samples), and EdiTPP by Lüdke et al. (2026) (4,000 samples, checkpoints selected on the Wasserstein distance of inter-event times); each is a mean over five training seeds on the same splits and with the same MMD. For forecasting we use the baseline rows of Kerrigan et al. (2025), who evaluate all of them with the same windows and the same MARE.

Metrics. *Density estimation.* Following (Shchur et al., 2020b; Lüdke et al., 2023), we report the maximum mean discrepancy (MMD) between 1,000 generated sequences and the test sequences. Times are rescaled to $[0, 1]$; for sequences $\gamma = \{t_\gamma^1, \dots, t_\gamma^n\}$ and $\eta = \{t_\eta^1, \dots, t_\eta^m\}$ with $n \le m$ the distance is $d(\gamma, \eta) = \sum_{k=1}^{n} |t_\gamma^k - t_\eta^k| + \sum_{k=n+1}^{m} (1 - t_\eta^k)$ (Xiao et al., 2017), the kernel is $k(\gamma, \eta) = \exp(-d(\gamma, \eta)/(2\sigma^2))$ with σ the median of all pairwise distances (Gretton et al., 2012). We also report the Wasserstein distance between the distributions of sequence lengths (W_1, normalised by the maximum training length). *Forecasting.* For each test sequence we draw 50 windows of length ΔT (4 time units; 5 for PUBG) with the history before the window as input, and report the mean absolute relative error of the event count, $\mathrm{MARE} = |\hat{N} - N| / \max(N, 1)$ (Kerrigan et al., 2025), and the forecast Wasserstein distance of (Lüdke et al., 2023). For RIFT we train five models with different seeds, draw 1,000 sequences from each, and report the mean $\pm$ standard deviation of the metric over the five seeds; baseline values are the seed-averaged means reported in their sources (Section 4.1, *Baselines*).

Implementation. RIFT uses $K = 5$ steps for training and sampling, a moment-matched sequence-level reference rate (Section 3), a fixed kernel bandwidth, $M = 100$ stratified quadrature points and a bounded velocity head $v = c \tanh(\tilde{v}/c)$. We train with AdamW (learning rate 10^{-4}, batch size 32, gradient-norm clipping), keep an exponential moving average (EMA) of the weights for evaluation, and select the checkpoint by validation MMD computed on 500 samples every 5 epochs.

Table 1: Density estimation: MMD ($\times 10^{-1}$; lower is better). RIFT: mean$_{\pm\text{std}}$ over five training seeds, 1,000 samples each. **Bold**: best; underline: second best. RIFT uses $K = 5$ steps, Add-Thin $K = 100$, EventFlow 25 function evaluations. The last row is the MMD between training and test sequences of the same dataset, an estimate of the irreducible sampling error. Baselines are copied from their sources (means over five seeds; standard deviations in Appendix Table 9): IF-RNN to Add-Thin from Lüdke et al. (2023) (4,000 samples, one decimal), EventFlow from Kerrigan et al. (2025) (1,000 samples) and ‡EdiTPP from Lüdke et al. (2026) (4,000 samples, checkpoints selected on a different validation metric); differences within one decimal or one standard deviation are not significant. **RIFT** uses the moment-matched sequence-level reference rate (Section 3); "RIFT, global λ_1" is the same model with a single reference rate.

Method	*Synthetic*						*Real-world*						
	H1	H2	SC	IPP	RP	MRP	PUBG	RedditC	RedditS	Taxi	Twitter	Yelp1	Yelp2
IF-RNN	0.2	**0.1**	0.8	0.5	**0.1**	**0.3**	0.4	0.1	0.2	0.4	0.3	0.7	**0.3**
IF-TF	0.3	0.4	1.9	1.0	0.2	1.9	0.6	0.5	0.9	0.9	0.8	1.2	1.4
TriTPP	0.3	0.4	2.3	0.4	0.2	0.5	0.6	0.9	1.2	0.7	0.4	0.6	0.6
GD	0.6	0.6	1.3	0.8	0.5	1.4	1.1	0.3	0.3	1.0	1.5	1.2	1.0
Add-Thin	0.2	0.2	1.9	0.3	0.2	1.0	0.3	0.1	0.2	0.4	0.4	0.8	0.4
EventFlow	0.19	0.22	0.83	0.38	0.17	0.42	0.15	**0.07**	0.07	0.35	0.49	0.66	**0.30**
EdiTPP‡	**0.11**	0.12	**0.77**	**0.17**	**0.10**	0.35	**0.14**	0.08	0.24	**0.31**	**0.13**	**0.37**	0.40
RIFT, global λ_1	0.85	0.89	1.83	0.48	0.78	0.97	0.48	0.36	0.82	0.38	1.36	0.58	0.51
RIFT	$0.14_{\pm.05}$	$0.12_{\pm.03}$	$0.80_{\pm.07}$	$0.27_{\pm.08}$	$\mathbf{0.10}_{\pm.03}$	$0.32_{\pm.03}$	$0.15_{\pm.05}$	$\mathbf{0.07}_{\pm.02}$	$\mathbf{0.06}_{\pm.01}$	$0.35_{\pm.04}$	$0.28_{\pm.07}$	$0.55_{\pm.18}$	$0.38_{\pm.08}$
Data (train vs. test)	0.14	0.12	0.62	0.22	0.13	0.35	0.23	0.04	0.04	0.35	0.29	0.62	0.25

Table 2: Forecasting: MARE of the event count in 50 random windows per test sequence (lower is better), mean$_{\pm\text{std}}$ over five seeds. Baselines as reported by Kerrigan et al. (2025) with the same windows and metric (GD is the diffusion model of Lin et al. (2022), NHP the neural Hawkes process). RIFT uses a single global reference rate for forecasting.

	PUBG	RedditC	RedditS	Taxi	Twitter	Yelp1	Yelp2
Average seq. length	76.5	295.7	1129.0	98.4	14.9	30.5	55.2
IF-RNN	$1.05_{\pm.14}$	$1.69_{\pm.39}$	$0.79_{\pm.20}$	$0.60_{\pm.11}$	$0.88_{\pm.08}$	$0.76_{\pm.07}$	$0.76_{\pm.05}$
NHP	$1.02_{\pm.08}$	$0.95_{\pm.01}$	$1.00_{\pm.00}$	$0.67_{\pm.11}$	$2.48_{\pm.40}$	$0.80_{\pm.22}$	$1.07_{\pm.34}$
GD	$1.95_{\pm.48}$	$1.28_{\pm.09}$	$1.12_{\pm.56}$	$0.49_{\pm.07}$	$0.66_{\pm.04}$	$0.65_{\pm.07}$	$0.72_{\pm.07}$
Add-Thin	$0.43_{\pm.01}$	$0.99_{\pm.10}$	$0.38_{\pm.01}$	$0.33_{\pm.02}$	$0.60_{\pm.02}$	$0.42_{\pm.01}$	$0.46_{\pm.03}$
EventFlow	$0.40_{\pm.01}$	$0.70_{\pm.07}$	$\mathbf{0.16}_{\pm.01}$	$\mathbf{0.28}_{\pm.01}$	$0.46_{\pm.03}$	$0.56_{\pm.06}$	$0.50_{\pm.02}$
EdiTPP	$\mathbf{0.35}_{\pm.04}$	$\mathbf{0.36}_{\pm.16}$	$0.28_{\pm.01}$	$1.23_{\pm.04}$	$\mathbf{0.23}_{\pm.30}$	$0.35_{\pm.01}$	$0.90_{\pm.04}$
RIFT	$0.39_{\pm.04}$	$0.47_{\pm.16}$	$0.20_{\pm.01}$	$0.44_{\pm.04}$	$0.24_{\pm.30}$	$\mathbf{0.28}_{\pm.01}$	$\mathbf{0.43}_{\pm.04}$

The full hyperparameter list is in Appendix K. All experiments run on a single NVIDIA V100 GPU. Density-estimation runs train for at most 1,500 epochs or 1.5–1.7 hours, forecasting runs for at most 8.5 hours (Appendix K).

4.2 Density Estimation

Table 1 reports MMD on the 13 datasets. We use a moment-matching method to determine the reference rate for each sequence. Compared with a single global rate, it performas much better in MMD, such as from 0.85 to 0.14 on H1, 0.89 to 0.12 on H2, 0.78 to 0.10 on RP and 1.36 to 0.28 on Twitter. In general, RIFT's MMD scores are always the top two among all the comparison methods, including RIFT perform the best in RP, RedditC, RedditS. In the datasets where RIFT is not the best performer, its MMD values are either close to the top one or to the data (train vs. test).

4.3 Forecasting

Table 2 reports forecasting results. RIFT is the best in Yelp1 and Yelp2, and is the second best on PUBG, RedditC, RedditS and Twitter. In Taxi, RIFT's performance is quite comparable to the top two performers. In RIFT, history is used only through the conditioning embedding; the generative process inside the window is the same as for density estimation.

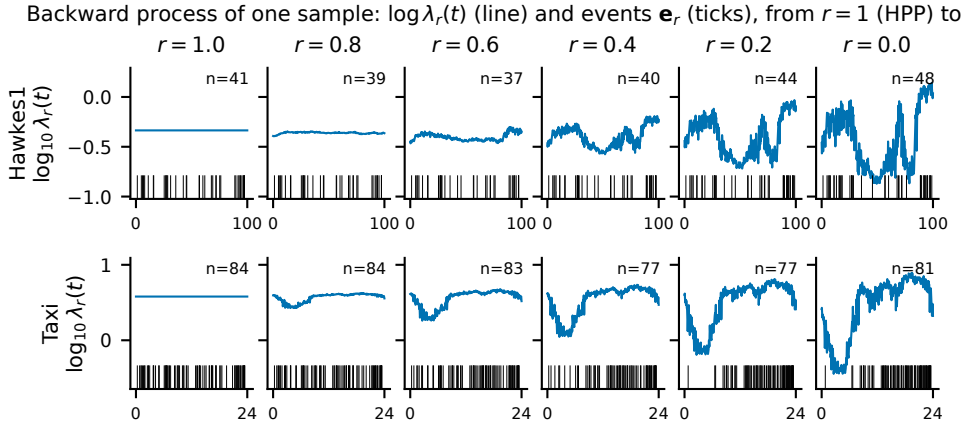


Figure 2: The backward process for one generated Hawkes1 and one Taxi sequence: log-intensity $\log_{10} \lambda_r(t)$ (line) and events $\mathbf{e}_r$ (ticks) after each of the $K = 5$ steps, from the homogeneous reference ($r = 1$) to the sample ($r = 0$); n is the number of events. For both processes the intensity profile (the daily pattern of Taxi, the bursts of Hawkes1) emerges gradually along the path, with each step thinning and adding a few events.

4.4 KDE FUNCTIONS VS GROUND-TRUTH INTENSITY FUNCTIONS

Driving the forward process with the kernel estimate instead of the true intensity makes no significant difference (Table 3). As we can see, there is no clear differences of performs between the one using kernel estimate and the ground-truth intensity function. The kernel estimate is thus a reliable substitute for the unknown intensity in the forward process.

Table 3: MMD ($\times 10^{-1}$, mean $\pm$ std over five seeds) with the forward process driven by the kernel estimate $\widehat{\lambda}_0(t)$ or by the ground-truth intensity function $\lambda_0(t)$.

Intensity	H1	H2	SC	IPP	RP	MRP
$\widehat{\lambda}_0(t)$	0.14 ± 0.05	0.12 ± 0.03	0.80 ± 0.07	0.27 ± 0.08	0.10 ± 0.03	0.32 ± 0.03
$\lambda_0(t)$	0.16 ± 0.06	0.12 ± 0.04	0.80 ± 0.04	0.31 ± 0.06	0.10 ± 0.04	0.30 ± 0.05

4.5 ABLATION STUDY

Figure 2 shows the backward process for one generated Hawkes1 and one Taxi sequence. Table 4 compares reference processes. Moment matching gives the best H1 result among the variants that keep Taxi and SC unchanged; using $\lambda_1 = n/T$ directly is worse on H1 and, as discussed above, on Poisson-count data. Self-exclusion (below) is the only change that substantially improves SC (1.09), but it hurts PUBG and Twitter in our runs, so we do not use it by default; count features combined with velocity matching improve H1 further (0.18) at the cost of Taxi and Yelp. The stratified quadrature with a bounded velocity is necessary for stable training, five sampling steps suffice ($K = 1$ is far worse, $K = 10, 20$ do not help), the straightness loss acts only as a mild regulariser, and query-time conditioning is essential on non-stationary data (Taxi: $0.39 \rightarrow 2.20$ without it).

4.6 ANALYSIS

What limits RIFT? Reference floors. To separate optimisation error from what the model family can represent, we compare against simple reference generators without any training (Appendix Table 5): a homogeneous Poisson process at the training rate, uniform event times with lengths drawn from the training set, and a Cox process that resamples a random training sequence's kernel intensity, $\mathrm{PP}(\widehat{\lambda}_0)$. The last one is exactly the distribution the forward chain targets at $r = 0$, so it indicates the

Table 4: Ablation of the reference process on H1, SC and Taxi: MMD ($\times 10^{-1}$), mean $\pm$ std, recipe of Table 8. †fewer than three seeds.

Variant	H1	SC	Taxi
Reference process (five seeds)			
RIFT: moment-matched sequence-level λ_1	$\mathbf{0.14 \pm 0.05}$	$\mathbf{0.80} \pm 0.07$	$\mathbf{0.35} \pm 0.04$
sequence-level $\lambda_1 = n/T$ (no moment matching)	0.33 ± 0.10	1.87 ± 0.02	0.41 ± 0.05
global λ_1	0.85 ± 0.09	1.83 ± 0.05	0.38 ± 0.03
$\lambda_1 = n/T$ + self-exclusion (three seeds)	0.25 ± 0.03	1.09 ± 0.22	0.38 ± 0.05
$\lambda_1 = n/T$ + count features + velocity matching†	0.18 ± 0.01	1.94 ± 0.06	0.50 ± 0.02

best MMD a perfectly trained RIFT with a Poisson-resampled chain could reach. On H1, RP, Taxi and Yelp1 it is within the train-vs.-test floor, so the family is expressive enough on these datasets and the remaining gap is an optimisation problem. On SC it stays at 0.19 for every bandwidth, the same as a homogeneous Poisson process: the self-correcting process is under-dispersed (count variance/mean ≈ 0.01), which no Cox process can reproduce. Capturing it requires the velocity to depend on the other events (the data-sourced chain), not only on a smoothed intensity.

Why a sequence-level reference rate. Figure 5 shows the effect of the reference process on the count distribution. With a single global λ_1 every sample starts with $\mathrm{Poisson}(\lambda_1 T)$ events and five near-exact thinning/superposition steps leave the count almost Poisson, so the MMD, which is dominated by length differences, stays near the "uniform times, training lengths" floor. Drawing the reference count from the moment-matched training counts gives each sample the right total mass from the start, so the flow only has to place the events, and the generated sequences become bursty like the data (Appendix Figures 2 and 3).

Sampling efficiency. With batched evaluation, generating 100 sequences with $K = 5$ takes 0.24 s on a V100 for H1, Taxi, Twitter and PUBG alike (one training step: 0.42 s for a batch of 32); including evaluation overheads, sampling the 1,000 test sequences takes 4–10 ms per sequence across all 13 datasets. Batching the velocity network over sequences made training $15\times$ and sampling $29\times$ faster without changing the outputs. Since we did not re-run the baselines, we do not report a wall-clock comparison on the same hardware.

Conclusion

We introduced RIFT, a flow on log-intensity functions for TPPs whose transitions are exact thinning/superposition steps and which samples with $K = 5$ steps. With a moment-matched, sequence-level reference rate, the proposed RIFT consistently to be the top two performers among all the datasets in both kernel density and forecasting tasks. **Limitations.** RIFT's forward process relies on a kernel estimate. Its training requires one full forward process which might be less effective than a simulation-free velocity-matching objective (Appendix J), and it is restricted to unmarked, univariate sequences.

AI use statement

We used generative AI tools (Claude) to help debug the training, to orchestrate experiments on a compute cluster, and to check its consistency with the implementation. All AI-assisted code was tested (unit tests for estimator correctness and batched/per-sequence equivalence) and all reported numbers come from logged runs. We reviewed all AI-assisted work and take responsibility for the final content of this paper.

Reproducibility statement

The method is specified in Section 3 and Algorithms 2–1; implementation details and all hyperparameters are in Appendix K (Table 8); datasets, splits and metric definitions are in Section 4.1.

Anonymised code, configuration files and job scripts that reproduce every table are provided in the supplementary material.

REFERENCES

Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic Interpolants: A Unifying Framework for Flows and Diffusions. *arXiv preprint arXiv:2303.08797*, 2023.

Ifigeneia Apostolopoulou, Scott Linderman, Kyle Miller, and Artur Dubrawski. Mutually regressive point processes. *Advances in Neural Information Processing Systems*, 32, 2019.

David R Cox. The statistical analysis of dependencies in point processes. In Peter A W Lewis (ed.), *Stochastic Point Processes: Statistical Analysis, Theory, and Applications*, pp. 55–66. Wiley, New York, 1972.

Daryl J Daley and David Vere-Jones. *An Introduction to the Theory of Point Processes: Volume II: General Theory and Structure*. Springer Science & Business Media, 2007.

Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomez-Rodriguez, and Le Song. Recurrent marked temporal point processes: Embedding event history to vector. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1555–1564, 2016.

Joseph Enguehard, Dan Busbridge, Adam Bozson, Claire Woodcock, and Nils Hammerla. Neural temporal point processes for modelling electronic health records. In *Machine Learning for Health*, pp. 85–113. PMLR, 2020.

Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=FPnUhsQJ5B.

Lina Fan, Hao Yang, Jia Zhai, and Xiaotao Zhang. Forecasting stock volatility during the stock market crash period: The role of Hawkes process. *Finance Research Letters*, pp. 103839, 2023.

Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(1):723–773, 2012.

Alan G Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58 (1):83–90, 1971.

Alan G Hawkes. Hawkes processes and their applications to finance: a review. *Quantitative Finance*, 18(2):193–198, 2018.

Valerie Isham and Mark Westcott. A self-correcting point process. *Stochastic processes and their applications*, 8(3):335–347, 1979.

Gavin Kerrigan, Kai Nelson, and Padhraic Smyth. EventFlow: Forecasting temporal point processes with flow matching. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2025.

Sangyun Lee, Zinan Lin, and Giulia Fanti. Improving the training of rectified flows. *Advances in neural information processing systems*, 37:63082–63109, 2024.

Haitao Lin, Lirong Wu, Guojiang Zhao, Liu Pai, and Stan Z. Li. Exploring generative neural temporal point process. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL https://openreview.net/forum?id=NPfS5N3jbL.

Scott Warren Linderman. *Bayesian Methods for Discovering Structure in Neural Spike Trains*. PhD thesis, 2016.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow Matching for Generative Modeling. *ICLR*, 2023.

Xingchao Liu, Chengyue Gong, and qiang liu. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow. *ICLR*, 2023.

David Lüdke, Marin Biloš, Oleksandr Shchur, Marten Lienen, and Stephan Günnemann. Add and thin: Diffusion for temporal point processes. *Advances in Neural Information Processing Systems*, 36:56784–56801, 2023.

David Lüdke, Enric Rabasseda Raventós, Marcel Kollovieh, and Stephan Günnemann. Unlocking point processes through point set diffusion. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=4anfpHj0wf.

David Lüdke, Marten Lienen, Marcel Kollovieh, and Stephan Günnemann. Edit-based flow matching for temporal point processes. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=FNf9IV1P2L.

Maya Okawa, Tomoharu Iwata, Takeshi Kurashima, Yusuke Tanaka, Hiroyuki Toda, and Naonori Ueda. Deep mixture point processes: Spatio-temporal event prediction with rich contextual information. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 373–383, 2019.

Takahiro Omi, Kazuyuki Aihara, et al. Fully neural network based model for general temporal point processes. *Advances in neural information processing systems*, 32, 2019.

William S Peebles and Saining Xie. Scalable diffusion models with transformers. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4195–4205, 2023.

BA Sevast'yanov. Renewal theory. *Journal of Soviet Mathematics*, 4(3):281–302, 1975.

Oleksandr Shchur, Marin BiloÅ¡, and Stephan GÃ¼nnemann. Intensity-free learning of temporal point processes. In *International Conference on Learning Representations*, 2020a. URL https://openreview.net/forum?id=HygOjhEYDH.

Oleksandr Shchur, Nicholas Gao, Marin Biloš, and Stephan Günnemann. Fast and flexible temporal point processes with triangular maps. *Advances in neural information processing systems*, 33: 73–84, 2020b.

Shuai Xiao, Mehrdad Farajtabar, Xiaojing Ye, Junchi Yan, Le Song, and Hongyuan Zha. Wasserstein learning of deep generative point process models. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Junchi Yan, Hongteng Xu, and Liangda Li. Modeling and applications for temporal point processes. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 3227–3228, 2019.

Feng Zhou, Zhidong Li, Xuhui Fan, Yang Wang, Arcot Sowmya, and Fang Chen. Efficient inference for nonparametric hawkes processes using auxiliary latent variables. *Journal of Machine Learning Research*, 21(241):1–31, 2020.

Simiao Zuo, Haoming Jiang, Zichong Li, Tuo Zhao, and Hongyuan Zha. Transformer hawkes process. In *International conference on machine learning*, pp. 11692–11702. PMLR, 2020.

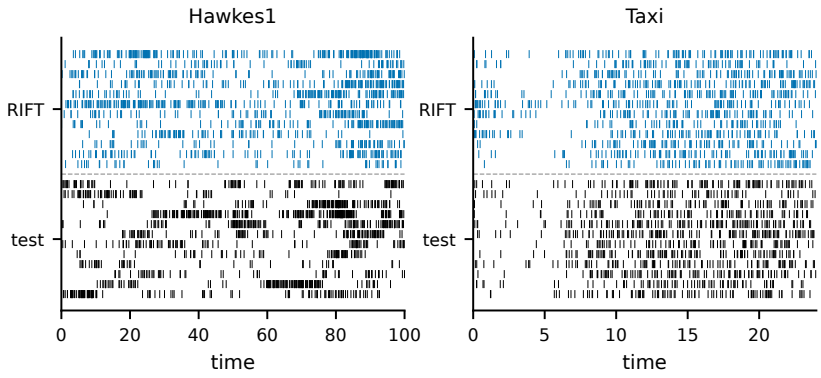


Figure 3: Twelve test sequences (black) and twelve RIFT samples (blue) on Hawkes1 and Taxi. RIFT reproduces the daily pattern of Taxi and the bursty, over-dispersed sequences of Hawkes1.

Table 5: Reference floors (MMD, not scaled; 1,000 samples; no training). "Cox-KDE" resamples $\mathrm{PP}(\widehat{\lambda}_0)$ of a random training sequence with the best of $\sigma^2 \in \{0.05, 0.2, 0.5, 1.5, 5\}$.

	H1	SC	RP	Taxi	Yelp1
Train vs. test	0.014	0.062	0.013	0.035	0.062
HPP(λ_1)	0.243	0.187	0.212	0.349	0.145
Uniform times, training lengths	0.029	0.206	0.028	0.221	0.108
Cox-KDE (best σ^2)	0.010	0.186	0.012	0.038	0.069

A Detailed description of Scaled-sampling

Given a current event sequence $\overline{\mathbf{e}}$ with intensity $\overline{\lambda}(t)$ and an intensity scaling function $\overline{\omega}(t) = \lambda_{\text{new}}(t)/\overline{\lambda}(t)$, the function `Scaled-Sampling`$(\overline{\mathbf{e}}, \overline{\lambda}(t), \overline{\omega}(t))$ produces a new event sequence $\mathbf{e}_{\text{new}}$ distributed as $\mathrm{TPP}(\overline{\lambda}_{\text{new}}(t))$ where $\overline{\lambda}_{\text{new}}(t) = \overline{\lambda}(t) \cdot \overline{\omega}(t)$. The procedure is implemented through two complementary operations:

Event thinning. Each event $t_i \in \overline{\mathbf{e}}$ is retained in the output with probability $\min(\overline{\omega}(t_i), 1)$. When $\overline{\omega}(t_i) \geq 1$, the event is always kept; when $\overline{\omega}(t_i) < 1$, events are randomly discarded at the ratio $\overline{\omega}(t_i)$, reducing the local intensity to $\overline{\lambda}(t) \cdot \overline{\omega}(t)$.

Event addition. In regions where $\overline{\omega}(t) > 1$, the target intensity exceeds the original, requiring new events. These are generated via rejection sampling: candidate events $\widetilde{\mathbf{e}}$ are drawn from a homogeneous Poisson process with intensity equal to the maximum excess intensity $\max_t \big((\overline{\omega}(t) - 1)\overline{\lambda}(t)\big)$, then each candidate event at time t is accepted with probability $(\overline{\omega}(t) - 1)\overline{\lambda}(t)/M^*$ where M^* is the upper bound used in rejection sampling. The final sequence $\mathbf{e}_{\text{new}}$ is the concatenation of the retained events from $\overline{\mathbf{e}}$ and the accepted new events $\widetilde{\mathbf{e}}$.

Correctness. By the superposition and thinning theorems for Poisson processes (Daley & Vere-Jones, 2007), thinning $\overline{\mathbf{e}}$ with retention probability $\min(\overline{\omega}(t_i), 1)$ yields a TPP with intensity $\min(\overline{\omega}(t), 1) \cdot \overline{\lambda}(t)$, and adding independent events from $\mathrm{TPP}(\max((\overline{\omega}(t) - 1)\overline{\lambda}(t), 0))$ recovers the remaining intensity. Their superposition therefore follows $\mathrm{PP}(\overline{\lambda}(t) \cdot \overline{\omega}(t))$. These theorems require the input $\overline{\mathbf{e}}$ to be a Poisson process with intensity $\overline{\lambda}$. This holds exactly in the backward process, whose initial sequence is $\mathrm{HPP}(\lambda_1)$ and whose every step is a thinning/superposition of a Poisson process. In the forward process the input is the observed sequence $\mathbf{e}_0$, which is not in general a sample of $\mathrm{PP}(\widehat{\lambda}_0)$; the noised sequences are then approximate samples of $\mathrm{PP}(\widehat{\lambda}_r)$, and exact ones if the chain is started from $\widehat{\mathbf{e}}_0 \sim \mathrm{PP}(\widehat{\lambda}_0)$.

B Other basic settings

A single global rate would make the number of events of every sample start from the same Poisson distribution, and a few steps of thinning and superposition cannot turn it into the strongly over-dispersed counts of, e.g., Hawkes data (count variance/mean ≈ 20). We therefore use a *sequence-level* reference rate: for a training sequence with n events, $\lambda_1 = \tilde{n}/T$ with $\tilde{n} = m + s\,(n - m)$, where m and v are the mean and variance of the training counts and $s = \sqrt{\max(v - m, 0)/v}$; at generation time, n is drawn from the empirical distribution of training counts and $\mathbf{e}_1 \sim \mathrm{HPP}(\tilde{n}/T)$. The reference is thus a mixed Poisson process whose count variance, $s^2 v + m$, equals the variance v of the data whenever the data are over-dispersed ($v > m$); for Poisson-like or under-dispersed counts $s = 0$ and the global rate is recovered (using $\lambda_1 = n/T$ directly would add the Poisson variance m on top of v, doubling the count variance of Poisson-count data) (EventFlow similarly samples the number of events first (Kerrigan et al., 2025), and EdiTPP offers a noise distribution with empirical counts (Lüdke et al., 2026)).

C Integration approximation details

We discretize the transition interval $[0, 1]$ into K segments $\{[r_{k-1}, r_k]\}_{k=1}^{K}$ where $r_0 = 0$ and $r_K = 1$. Within each segment, the velocity field is assumed constant and evaluated at the right endpoint r_k. Under this piecewise-constant assumption, the ODE integral is approximated by a Riemann sum:

$$\int_1^0 \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)\,\mathrm{d}r \approx \sum_{k=1}^{K} \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t)\,(r_{k-1} - r_k). \tag{8}$$

The ground-truth intensity function is recovered as

$$\lambda_0(t) \approx \lambda_1 \cdot \exp\left(\sum_{k=1}^{K} \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t)\,(r_{k-1} - r_k)\right). \tag{9}$$

The intermediate intensity at step r_k is approximated by $\lambda_{r_k}(t) \approx \lambda_1 \cdot \exp(\sum_{j=k+1}^{K} \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_j}, r_j, t)\,(r_{j-1} - r_j))$, from which the intensity ratio between consecutive steps follows as

$$\omega_{r_{k-1}}(t) = \frac{\lambda_{r_{k-1}}(t)}{\lambda_{r_k}(t)} = \exp(\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t)\,(r_{k-1} - r_k)). \tag{10}$$

The approximation quality depends on two factors: (1) the learned network $\mathbf{v}_{\boldsymbol{\phi}}$ accurately estimating the ground-truth velocity at each step, and (2) the discretisation error decreasing as K increases. The first depends on training quality; the second is controlled by the step granularity and is mitigated by the straightness loss (Section 3.4).

Alternative straightness penalty. An alternative approach is to directly penalize the velocity derivative $\min_{\boldsymbol{\phi}} \|\partial \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)/\partial r\|^2$. However, this is difficult to implement because $\mathbf{e}_r$ is generated discretely and its gradient with respect to r is non-trivial to compute. The pairwise penalty in Eq. (6) avoids this issue by comparing velocities at discrete steps.

D Preliminary knowledge

Temporal Point Processes (TPPs) Temporal point processes (TPPs) provide a probabilistic framework for modeling random events distributed over continuous time (Daley & Vere-Jones, 2007). Let $\mathcal{N}(t)$ denote the counting process that represents the total number of events occurring up to time t. The event occurrence rate is characterized by the (conditional) intensity function $\lambda(t)$, defined as $\lambda(t) = \lim_{\Delta t \to 0^+} \frac{\mathbb{E}[\mathcal{N}(t+\Delta t) - \mathcal{N}(t)]}{\Delta t}$.. By construction, $\lambda(t) \geq 0$ and quantifies the instantaneous likelihood of an event occurring at time t, where larger values indicate higher event arrival rates.

Given a sequence of n observed events $\{t_1, \ldots, t_n\}$ with $t_i \in [0, T]$ and a specified intensity function $\lambda(t)$, the log-likelihood of the event sequence under a TPP model is given by

$$\mathcal{L}(\{t_1, \ldots, t_n\} \mid \lambda(t)) = \sum_{i=1}^{n} \log \lambda(t_i) - \int_0^T \lambda(t)\,\mathrm{d}t. \tag{11}$$

Rectified flows Rectified flow (Lipman et al., 2023; Liu et al., 2023; Albergo et al., 2023), also referred to as flow matching or stochastic interpolation, models the transformation between two distributions $p_0(\cdot)$ and $p_1(\cdot)$ via Ordinary Differential Equations (ODEs). Specifically, given a clean sample $\mathbf{x}_0 \sim p_0(\mathbf{x}_0)$ and a noise sample $\mathbf{x}_1 \sim p_1(\mathbf{x}_1)$, rectified flow defines a linear interpolation trajectory $\mathbf{x}_r = (1-r)\mathbf{x}_0 + r\mathbf{x}_1, \quad r \in [0,1]$.. The model is trained by learning a velocity field $\mathbf{v}_{\boldsymbol{\theta}}(\mathbf{x}_r, r)$ that matches the constant velocity along this path, with the objective

$$\mathcal{L}_{\text{rf}}(\boldsymbol{\theta}) = \int \mathbb{E}_{\mathbf{x}_0,\mathbf{x}_1}\left[\left\|(\mathbf{x}_1 - \mathbf{x}_0) - \mathbf{v}_{\boldsymbol{\theta}}(\mathbf{x}_r, r)\right\|^2\right] \mathrm{d}r, \tag{12}$$

where $\mathbf{v}_{\boldsymbol{\theta}}(\mathbf{x}_r, r)$ denotes the learned velocity field.

For generation, rectified flow starts from a noise sample $\mathbf{x}_1$ and solves the ODE $\mathrm{d}\mathbf{x}_r = \mathbf{v}_{\boldsymbol{\theta}}(\mathbf{x}_r, r)\mathrm{d}r, \quad r \in [0,1]$, which yields the recovered clean sample $\widehat{\mathbf{x}}_0 = \mathbf{x}_1 + \int_1^0 \mathbf{v}_{\boldsymbol{\theta}}(\mathbf{x}_r, r)\mathrm{d}r$.

By minimizing $\mathcal{L}_{\text{rf}}(\boldsymbol{\theta})$, the learned velocity field $\mathbf{v}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)$ is encouraged to approximate the straight-line transport from $\mathbf{x}_1$ to $\mathbf{x}_0$. As a result, rectified flow can achieve high-quality sample generation using significantly fewer integration steps than conventional diffusion models, whose denoising trajectories are typically curved and require many steps.

E KERNEL DENSITY ESTIMATION FOR THE FORWARD PROCESS

The forward process uses the kernel intensity estimate $\widehat{\lambda}_0(t) = \sum_{i=1}^n \kappa_h(t - t_i)$ with a Gaussian kernel $\kappa_h(u) = \frac{1}{h}\phi(\frac{u}{h})$ (Eq. (3)). Because $\int \kappa_h = 1$, $\int_0^T \widehat{\lambda}_0(t)\,\mathrm{d}t \approx n$, so $\widehat{\lambda}_0$ has the scale of an intensity and the forward chain ends at the rate λ_1. The normalised estimate $\frac{1}{n}\sum_i \kappa_h$ is a probability density; used as an intensity it makes the chain end at a rate of order $\lambda_1 T$ instead of λ_1 (we measured roughly twice the data count at $r = 1$). The bandwidth h is fixed per dataset. Rules of thumb for density estimation of the event *times* (e.g. Silverman's $h = 1.06\,\hat{\sigma} n^{-1/5}$) are not appropriate here: $\hat{\sigma}$ measures the spread of all event times over $[0, T]$, which gives $h \approx 12$ on $T = 100$ and an almost constant $\widehat{\lambda}_0$. The velocity network is trained on the observed event times; the kernel estimate only shapes its noised inputs. The rejection sampling procedure for superposition is given in Function 2.

F KERNEL ESTIMATE VERSUS TRUE INTENSITY IN THE FORWARD PROCESS

The forward process uses a kernel estimate $\widehat{\lambda}_0$ of each sequence's intensity because the generating intensity is unknown on real data. To test whether this substitution matters, we train RIFT on synthetic data whose generating intensity is known and compare two forward processes that differ *only* in the intensity that defines the path $\lambda_r = \lambda_0^{1-r}\lambda_1^r$: the kernel estimate $\widehat{\lambda}_0$ (Eq. (3)) and the true intensity λ_0 of the sequence. We use three processes: Hawkes1, whose true intensity is the conditional intensity $\lambda_0(t) = 0.2 + 0.8\sum_{t_i<t} e^{-(t-t_i)}$ given the sequence's own past; the self-correcting process, $\lambda_0(t) = \exp(t - N(t^-))$; and an inhomogeneous Poisson process generated for this experiment, $\lambda_0(t) = 1 + 0.8\sin(2\pi t/25)$ on $[0, 100]$ (1,000 sequences, 99.7 events on average). The parameters of the first two match the benchmark generators (a log-likelihood check on the training data prefers $\alpha = 0.8$ over 0.6 and 0.9 for Hawkes1). Everything else (architecture, objective, sampler, sequence-level reference rate, five seeds) is identical, and evaluation uses the same MMD protocol as Table 1.

G FLOW MATCHING VERSUS RIFT

Table 6 places each component of RIFT next to its counterpart in Gaussian flow matching (FM; Appendix D), and Figure 4 illustrates the forward processes of both.

From a discrete sequence to a continuous intensity. The forward process has three stages. (i) *Smoothing*: the observed point set $\mathbf{e}_0$ is turned into a continuous, positive function $\widehat{\lambda}_0(t) = \sum_i \kappa_h(t - t_i)$ (Eq. (3)), whose integral is approximately the number of events. This plays the role that Gaussian noise plays in FM, but acts along the time axis of a single sequence rather than on the

Table 6: Gaussian flow matching versus RIFT, component by component.

	Flow matching	**RIFT**
Data space	vectors $\mathbf{x}_0 \in \mathbb{R}^d$	event sets $\mathbf{e}_0 \subset [0, T]$ of variable size; transported object: $\log \lambda_r(t)$
Prior / reference	$\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$	$\mathbf{e}_1 \sim \mathrm{HPP}(\lambda_1)$
Interpolation path	$\mathbf{x}_r = (1-r)\mathbf{x}_0 + r\,\boldsymbol{\epsilon}$	$\log \lambda_r = (1-r)\log \widehat{\lambda}_0 + r \log \lambda_1$ (Eq. (2))
Drawing an intermediate sample	add Gaussian noise to $\mathbf{x}_0$, directly at any r	thin / superpose the events of $\mathbf{e}_{r_{k-1}}$ (`Scaled-Sampling`); event times are never perturbed
Smoothing of the data	convolution of p_0 with Gaussian noise in sample space	kernel estimate $\widehat{\lambda}_0 = \sum_i \kappa_h(\cdot - t_i)$ turns the point set into a continuous intensity
Training target	known conditional velocity $\boldsymbol{\epsilon} - \mathbf{x}_0$; regression at one random r	unknown (λ_0 unobserved); Poisson likelihood of the observed events through all K steps (Eq. (5))
Sampler	ODE solver (e.g. Euler) on $\mathbf{x}_r$	K exact Poisson steps (thinning / superposition) with exact intensity tracking (Algorithm 2)

coordinates of a sample. (ii) *Path*: the log-intensity is interpolated linearly from $\log \widehat{\lambda}_0$ to the constant $\log \lambda_1$, so peaks are flattened and troughs filled geometrically while λ_r stays positive. (iii) *Events*: an event set for λ_{r_k} is obtained from $\mathbf{e}_{r_{k-1}}$ by thinning where $\lambda_{r_k} < \lambda_{r_{k-1}}$ and superposing new events where $\lambda_{r_k} > \lambda_{r_{k-1}}$. The sequence is thus never represented by a continuous object except through $\widehat{\lambda}_0$; every intermediate state is again a finite event set whose length changes along the path. The backward process reverses stage (iii) with the learned velocity and needs neither the kernel estimate nor stage (i).

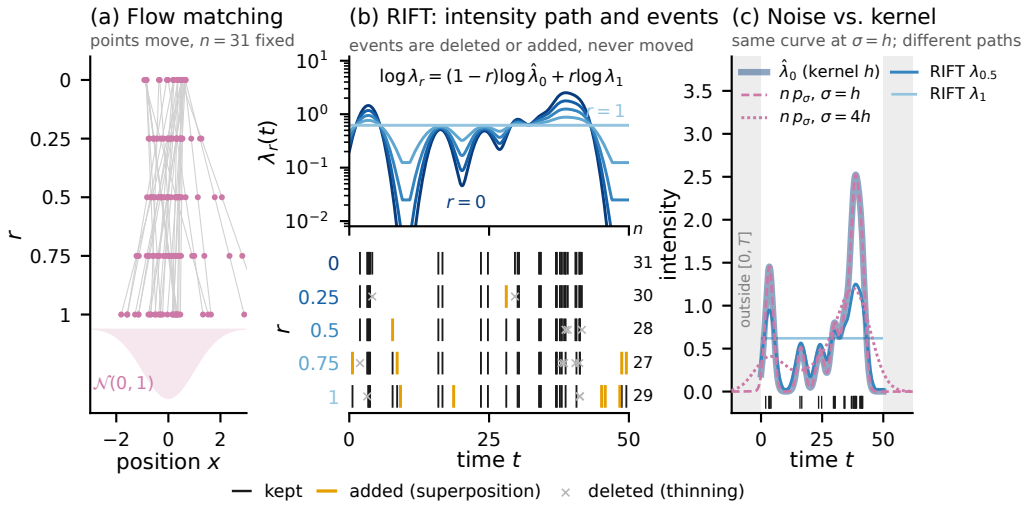


Figure 4: Forward processes of flow matching and RIFT on a 1-D point set. (a) Gaussian FM interpolation $\mathbf{x}_r = (1-r)\mathbf{x}_0 + r\,\boldsymbol{\epsilon}$: every point moves continuously towards noise. (b) RIFT: log-intensity path from the kernel estimate $\widehat{\lambda}_0$ to the homogeneous rate λ_1, with the events thinned and added at $r \in \{0, 0.25, 0.5, 0.75, 1\}$. (c) Gaussian noise smoothing (perturbing the event times) versus kernel smoothing (turning the event set into an intensity).

H WHY TRANSPORT THE INTENSITY?

Event sequences are the observed objects, so modelling event times directly is the more immediate choice. Transporting the intensity instead has the following advantages.

- *Exact likelihood as training signal.* Because the model outputs an intensity, the training loss is the exact Poisson log-likelihood of the observed events (Eq. (5)) rather than a surrogate loss on event positions, and the compensator $\int_0^T \lambda$ penalises mass placed where there are no events.
- *Exact transitions and variable length.* Every intermediate state is a Poisson process with explicit intensity $\lambda_r(t)$, so each step is an exact thinning/superposition (Remark 3.1); events are inserted and deleted rather than moved, and the number of events is handled natively without a fixed-length representation or a separate length model.
- *Interpretability.* The intensity $\lambda_{r_k}(t)$ is available at every step of generation (Figure 2).
- *Global, non-autoregressive, few steps.* The velocity is a function of the whole event set and all query times at once, so there is no accumulation of error over events, and the straight path in log-intensity allows sampling with $K = 5$ steps.

These advantages matter most for non-stationary or bursty data, where a global $\lambda(t)$ over $[0, T]$ carries most of the structure (removing the query-time input degrades the non-stationary Taxi data sharply, Appendix Table **??**), and when few sampling steps are required.

Bottlenecks. (i) *Count under-dispersion.* Each step is Poisson given the past steps, so the model is a Cox process and cannot represent counts that are less variable than Poisson; the self-correcting process (count variance/mean ≈ 0.01, Figure 5) is out of reach for this reason (Section 4.6). (ii) *Reliance on the kernel estimate.* The forward path is defined by $\widehat{\lambda}_0$ and its bandwidth; on synthetic data, replacing it by the true intensity makes no significant difference (Appendix F), but the bandwidth remains a hyperparameter. (iii) *Optimisation through the K-step chain.* The likelihood constrains only the sum of the K velocities along one chain (Appendix I); the reference floors of Section 4.6 indicate that most of the remaining gap on several datasets is due to optimisation rather than to the model family. Promising directions are velocities that depend explicitly on the other events (self-exclusion improves the self-correcting process, Table 4), non-Poisson transitions for under-dispersed data, and additional per-step supervision.

I TRAINING COST AND SIMULATION-FREE TRAINING

In flow matching, training is simulation-free in two senses: the intermediate sample $\mathbf{x}_r$ is available in closed form, and the regression target $\boldsymbol{\epsilon} - \mathbf{x}_0$ is known, so a single network evaluation at a random r gives an unbiased loss. RIFT keeps the first property and loses the second.

What is kept. The forward chain $\mathbf{e}_{r_0}, \ldots, \mathbf{e}_{r_K}$ (Algorithm 1) is sampled by thinning and superposition from the kernel-driven path. It does not depend on $\boldsymbol{\phi}$, so the learned ODE is never simulated during training, no gradient passes through a sampler, and the chain costs $\mathcal{O}(K)$ cheap Poisson operations per sequence without any network evaluation.

What is lost. The target velocity $\log \lambda_1 - \log \lambda_0(t)$ is unknown because λ_0 is not observed. The likelihood of Eq. (5) evaluates only the end-point intensity $\lambda_1 \exp(\sum_k \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t)\Delta r_k)$, so it constrains only the sum $\sum_k \mathbf{v}_k \Delta r_k$ and requires all K velocity evaluations on the same sequence; the straightness term adds the evaluation at r_0. A training step therefore costs $K + 1$ network passes per sequence ($K = 5$), each at $n + M$ query times, instead of one pass at one random r. Batching over sequences and encoding each $\mathbf{e}_{r_k}$ once (Section 3.3) keeps this affordable; timings are reported in Section 4.6.

A simulation-free alternative. The kernel path does provide a per-step target, $\log \lambda_1 - \log \widehat{\lambda}_0(t)$, and regressing $\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t)$ onto it gives an FM-style objective with a known target. It needs no likelihood, but it ties the learned velocity to the kernel estimate $\widehat{\lambda}_0$, whereas the likelihood objective only uses $\widehat{\lambda}_0$ to generate noised inputs. Appendix J compares the two with the same network, reference process and forward chain: they perform comparably, each ahead on some datasets. The

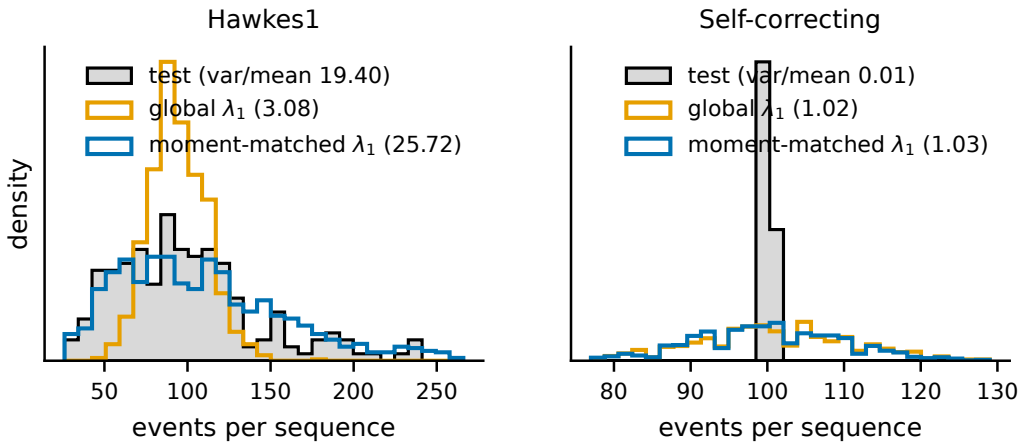


Figure 5: Distribution of the number of events per sequence (1,000 samples; legend: count variance/mean). A global reference rate yields nearly Poisson counts (Hawkes1: 3.1 vs. 19.4 in the data); the moment-matched sequence-level reference rate recovers the over-dispersion (25.7). The under-dispersed self-correcting counts (0.01) are out of reach for both, since moment matching falls back to the global rate when the data are not over-dispersed.

practical difference is therefore not accuracy but what is optimised: RIFT maximises the exact TPP likelihood of the observed events, which is available without any regression target, while the velocity-matching variant needs the kernel estimate to define its target.

J Flow matching on the kernel path

Setup. We train the same velocity network with the same moment-matched reference process and the same kernel-driven forward chain, but replace the objective by pure velocity matching onto the kernel path, $\frac{1}{K}\sum_k \mathbb{E}_t\big[(\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t) - (\log\lambda_1 - \log\max(\widehat{\lambda}_0(t), 10^{-3}\lambda_1)))^2\big]$, with t ranging over the data events and the $M = 100$ stratified quadrature points, and $\alpha_{\text{ll}} = \alpha_{\text{straight}} = 0$. The training budget, checkpoint selection (validation MMD), sampler ($K = 5$) and evaluation are identical to RIFT. Forecasting uses the global reference rate for both models. We ran three seeds for the baseline (RIFT: five).

Results. Table 7 shows that the RIFT consistently outperform the flow matching objective.

K Implementation details

There are two integration terms in Eq. (5): $\int_1^0 \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)\mathrm{d}r$ and $\int_0^T \exp\left(\int_1^0 \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)\mathrm{d}r\right)\mathrm{d}t$.

For the first term, since the trajectory with $\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)$ as velocity is enforced to be straight, the Euler ODE solver would conveniently approximate this integration well. That is, we divide the transition range into $1 = r_K > r_{K-1} > \ldots > r_0 = 0$ and approximate the integration as $\int_1^0 \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)\mathrm{d}r = \sum_k \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t)(r_{k-1} - r_k)$. As a result, $\lambda_0(t)$ can be obtained as $\lambda_0(t) \approx \lambda_1 \exp\left(\sum_k \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t)(r_{k-1} - r_k)\right)$.

A list of event sequences $\{\mathbf{e}_{r_k}\}_k$ are required in such integration approximations. The detail procedure is written in Algorithm 1. We sequentially generate the noised event sequences from $\mathbf{e}_0$ to $\mathbf{e}_1$ as r_k increases.

Table 7: Flow matching (FM) on the kernel path versus RIFT's likelihood objective, same network, reference process and forward chain. Density: MMD ($\times 10^{-1}$); forecasting: MARE and forecast Wasserstein distance (W). Mean $\pm$ standard deviation over seeds (FM: three; RIFT: five for density, three for forecasting). Lower is better; the better of the two is in bold when the gap exceeds the larger standard deviation.

Density (MMD)			*Forecasting (global reference rate)*				
Dataset	FM	RIFT	Dataset	MARE, FM	MARE, RIFT	W, FM	W, RIFT
H1	0.22 ± 0.01	0.14 ± 0.05	PUBG	$\mathbf{0.55} \pm 0.04$	0.61 ± 0.03	2.54 ± 0.07	2.54 ± 0.05
H2	0.23 ± 0.04	0.12 ± 0.03	Taxi	0.37 ± 0.01	0.39 ± 0.01	3.83 ± 0.04	$\mathbf{3.73} \pm 0.05$
SC	1.88 ± 0.01	0.80 ± 0.07	Twitter	$\mathbf{0.83} \pm 0.09$	1.04 ± 0.08	$\mathbf{2.66} \pm 0.08$	2.84 ± 0.10
IPP	0.39 ± 0.08	0.27 ± 0.18	Yelp1	0.55 ± 0.03	0.62 ± 0.07	1.82 ± 0.02	1.86 ± 0.08
RP	0.20 ± 0.02	0.10 ± 0.03	Yelp2	0.58 ± 0.08	0.61 ± 0.05	2.85 ± 0.12	2.81 ± 0.04
MRP	0.97 ± 0.01	0.32 ± 0.03					
PUBG	$\mathbf{0.48} \pm 0.06$	0.15 ± 0.05					
RedditC	0.06 ± 0.01	0.07 ± 0.02					
RedditS	0.09 ± 0.00	$\mathbf{0.06} \pm 0.01$					
Taxi	0.41 ± 0.01	0.35 ± 0.04					
Twitter	$\mathbf{0.39} \pm 0.01$	0.28 ± 0.07					
Yelp1	0.88 ± 0.06	$\mathbf{0.55} \pm 0.18$					
Yelp2	0.53 ± 0.01	$\mathbf{0.38} \pm 0.08$					

The second integral is estimated by stratified random quadrature. In summary, Eq. (5) is approximated as:

$$\mathcal{L}_{\text{ll}} \approx n \log \lambda_1 + \sum_{i,k} \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t_i)(r_{k-1} - r_k) - \frac{\lambda_1 T}{M} \sum_{m=1}^{M} \exp\Big(\sum_{k} \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t_m)(r_{k-1} - r_k)\Big) \tag{13}$$

where the first term $n \log \lambda_1$ does not depend on $\boldsymbol{\phi}$ and is dropped in the implementation, and $t_m = (m - 1 + u_m)\, T/M$ with $u_m \sim \mathrm{Uniform}[0, 1]$ redrawn at every training step and shared across k, so that the quadrature is an unbiased estimate of the compensator. The velocity head is bounded as $\mathbf{v} = c \tanh(\tilde{\mathbf{v}}/c)$, which keeps $\exp(\cdot)$ finite without clipping the loss (a clipped exponent removes the gradient of the compensator and let training diverge in our experiments).

Training budget. All density-estimation runs in Table 1 train for at most 1,500 epochs or 1.5 hours (1.7 hours for the Reddit datasets). Training the final model for up to 3.5 hours on the six datasets where it did not improve over the global reference rate changes MMD by at most 0.04 (Taxi 0.40 ± 0.01, Yelp1 0.62 ± 0.02, Yelp2 0.34 ± 0.05, PUBG 0.59 ± 0.17, SC 1.83 ± 0.07, MRP 0.99 ± 0.03), mainly reducing the variance across seeds, so these gaps are not due to under-training.

Table 8: Hyperparameters of RIFT (shared by all datasets unless stated).

Steps K (training and sampling)	5, uniform grid $r_k = k/K$
Reference rate λ_1	$(m + s(n - m))/T$, $s = \sqrt{\max(v - m, 0)/v}$; n drawn from the training counts at generation
Kernel bandwidth	fixed $\sigma^2 = h^2 = 1.5$ on the $T = 100$ time scale
Quadrature	$M = 100$ stratified random points
Velocity bound c	50
Loss weights	$\alpha_{\text{ll}} = 1$, $\alpha_{\text{straight}} = 1$
Architecture	hidden size 128, 2 self-attention + 2 cross-attention adaLN blocks, 2 heads
Time / step embedding	sinusoidal of $1000\, t/T$ and $1000\, r$, followed by a 2-layer MLP
Optimiser	AdamW, learning rate 10^{-4}, weight decay 10^{-4}, batch size 32, grad-norm clip 10
EMA	decay 0.999, started after 2,000 steps; evaluation uses EMA weights
Model selection	validation MMD on 500 samples every 5 epochs, patience 20 validations, min. 100 epochs
Sampler	exact intensity tracking on a 1,000-point grid (Algorithm 2)
Hardware	one NVIDIA V100 (32 GB)

Table 9: Baseline standard deviations of the density-estimation MMD ($\times 10^{-1}$) over five seeds, as reported in the sources of Table 1 (IF-RNN to Add-Thin: Lüdke et al. (2023); EventFlow: Kerrigan et al. (2025); EdiTPP: Lüdke et al. (2026)), next to RIFT.

	H1	H2	SC	IPP	RP	MRP	PUBG	RedditC	RedditS	Taxi	Twitter	Yelp1	Yelp2
IF-RNN	0.03	0.02	0.53	0.09	0.01	0.05	0.05	0.02	0.03	0.01	0.03	0.05	0.01
IF-TF	0.11	0.17	0.06	0.34	0.07	0.48	0.14	0.25	0.60	0.14	0.20	0.26	0.48
TriTPP	0.02	0.01	0.03	0.03	0.02	0.04	0.01	0.02	0.03	0.07	0.02	0.05	0.04
GD	0.04	0.02	0.04	0.02	0.02	0.08	0.23	0.01	0.01	0.02	0.11	0.10	0.01
Add-Thin	0.04	0.02	0.13	0.06	0.01	0.30	0.15	0.05	0.01	0.06	0.06	0.10	0.05
EventFlow	0.02	0.01	0.04	0.12	0.03	0.05	0.02	0.01	0.01	0.01	0.07	0.12	0.05
EdiTPP	0.02	0.01	0.04	0.03	0.02	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.03
RIFT	0.05	0.03	0.07	0.18	0.03	0.03	0.05	0.02	0.01	0.04	0.07	0.18	0.08

Algorithm 3 Training process of RIFT

Input: Ground-truth event sequence $\mathbf{e}_0$, base intensity value λ_1
Output: Trained velocities $\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)$

while not converged **do**
  Given $\mathbf{e}_0, \lambda_1$, use Algorithm 1 to generate the forward chain $\{\mathbf{e}_{r_k}\}_{k=1}^{K}$ along the path of Eq. (2) with $\widehat{\lambda}_0$
  Calculate likelihood loss $\mathcal{L}_{\text{ll}}$ based on Eq. (13)
  Calculate straightness loss $\mathcal{L}_{\text{straight}}$ based on Eq. (6), with pairs $(\mathbf{e}_{r_{k-1}}, \mathbf{e}_{r_k})$ from the same chain
  Use gradient descent to update $\boldsymbol{\phi}$ with $\nabla_{\boldsymbol{\phi}}\mathcal{L}$ of Eq. (7)
end while
return $\boldsymbol{\phi}$

Function 2 Backward process of RIFT with exact intensity tracking (implementation of Algorithm 2)

Input: Trained velocity $\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t)$, grid $1 = r_K > r_{K-1} > \ldots > r_0 = 0$, base rate λ_1, time grid $\{g_j\}_{j=1}^{G}$ on $[0, T]$
Output: Generated event sequence $\mathbf{e}_0$

$\mathbf{e}_{r_K} \sim \text{HPP}(\lambda_1)$; $\ell_j = \log \lambda_1$ for all j ▷ $\ell_j \approx \log \lambda_{r_k}(g_j)$
for $k = K, \ldots, 1$ **do**
  $\Delta = r_{k-1} - r_k$; $a_j = \mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, g_j)\,\Delta$ for all j ▷ $\log \omega$ on the grid
  $\mathbf{e}_{r_{k-1}} = \{t_i \in \mathbf{e}_{r_k} : u_i < \exp(\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, t_i)\Delta)\}$, $u_i \sim \text{Uniform}[0,1]$ ▷ thinning
  $M^* = 1.2 \max_j (e^{a_j} - 1)_+ e^{\ell_j}$; $C \sim \text{Poisson}(TM^*)$; $s_1, \ldots, s_C \sim \text{Uniform}[0, T]$
  for $c = 1, \ldots, C$ **do** ▷ superposition by rejection sampling
    $\ell(s_c)$ = linear interpolation of $\{\ell_j\}$; $\omega_c = \exp(\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_{r_k}, r_k, s_c)\Delta)$
    add s_c to $\mathbf{e}_{r_{k-1}}$ with probability $(\omega_c - 1)_+ \, e^{\ell(s_c)}/M^*$
  end for
  $\ell_j \leftarrow \ell_j + a_j$ for all j ▷ $\lambda_{r_{k-1}} = \lambda_{r_k}\,\omega_{r_{k-1}}$
end for
return $\mathbf{e}_{r_0}$

Algorithm 4 Velocity network $\mathbf{v}_{\boldsymbol{\phi}}(\mathbf{e}_r, r, t_{1:Q})$

Input: Event set $\mathbf{e}_r = \{t^1, \ldots, t^N\}$, step r, query times $t_1, \ldots, t_Q$; $\text{TE}(\cdot)$ time embedding, $\text{StE}(\cdot)$ step embedding
Output: Velocities $v_1, \ldots, v_Q$

$\mathbf{c} = \text{StE}(1000\, r)$ (+ history embedding for forecasting) ▷ adaLN conditioning
$\mathbf{h}_i = \text{Linear}(\text{TE}(1000\, t^i/T))$ for $i = 1, \ldots, N$
$\mathbf{h}_{1:N} \leftarrow \text{SelfAttnBlock}_\ell(\mathbf{h}_{1:N}; \mathbf{c})$ for $\ell = 1, \ldots, L$ ▷ encode $\mathbf{e}_r$ once, $\mathcal{O}(N^2)$
$\mathbf{s} = \frac{1}{N}\sum_i \mathbf{h}_i$ ▷ sequence summary
$\mathbf{q}_j = \text{TE}(1000\, t_j/T) + \mathbf{s}$ for $j = 1, \ldots, Q$
$\mathbf{q}_{1:Q} \leftarrow \text{CrossAttnBlock}_\ell(\mathbf{q}_{1:Q}, \mathbf{h}_{1:N}; \mathbf{c})$ for $\ell = 1, \ldots, L$ ▷ read out, $\mathcal{O}(QN)$
$v_j = c \tanh(\text{MLP}(\mathbf{q}_j)/c)$
return $v_{1:Q}$