

467 A Related Works

468 **Hawkes Process** Hawkes process has long been used to model event sequences (Hawkes, 1971),
469 such as earthquake aftershock sequences (Ogata, 1999), financial transactions (Bauwens and Hautsch,
470 2009), and events on social networks (Fox et al., 2016; Farajtabar et al., 2017). Its variant, mixture
471 of Hawkes processes model, has also been proved effective in many area (Yang and Zha, 2013; Li
472 and Zha, 2013; Xu and Zha, 2017). In most cases, the learning methodology is variational inference
473 or maximum likelihood estimation (Rasmussen, 2013; Zhou et al., 2013; Zhao et al., 2015). Other
474 possible methods includes least-squares-based method (Eichler et al., 2017), Wiener-Hopf-based
475 methods (Bacry et al., 2012), and cumulants-based methods (Achab et al., 2017).

476 Instead of predefine an impact function here, some non-parametric methods use discretization or
477 kernel-estimation when learning models (Reynaud-Bouret et al., 2010; Zhou et al., 2013; Hansen
478 et al., 2015). Those methods usually target small datasets, and do not need a good scalability.
479 Recently, some attempts have been made to further enhance the flexibility of Hawkes processes.
480 The time-dependent Hawkes process (TiDeH) in Kobayashi and Lambiotte (2016) and the neural
481 network-based Hawkes process in Mei and Eisner (2017) learn very flexible Hawkes processes with
482 complicated intensity functions. Those methods usually target very long and multi-dimensional
483 sequences, instead of short sequences.

484 Existing works targeting short sequences is usually in specific cases (Xu et al., 2017a,b), such as the
485 data is censored. However, there is no work targeting general short sequences as we do here.

486 There are lines of research that involves both point processes and graphs. One is using point process to
487 find the latent graph (Blundell et al., 2012; Linderman and Adams, 2014; Tran et al., 2015). Another
488 one is considering the interaction of the nodes as point process and use it to construct a dynamic
489 graph, instead of the event happens on nodes as we consider here (Farajtabar et al., 2016; Zarezhade
490 et al., 2017; Trivedi et al., 2018). These works have vary different aims from our work.

491 **Meta Learning** Meta learning has been studied since last century (Bengio et al., 1990; Chalmers,
492 1991). Some works focus on learning the hyperparameters, such as learning rates or initial conditions
493 (Maclaurin et al., 2015). Some works aim to learn a metric so that a simple K nearest neighbors can
494 perform well under such a metric (Koch et al., 2015; Vinyals et al., 2016; Sung et al., 2018; Snell
495 et al., 2017). Some works design specific deep neural networks so that the information of different
496 tasks are memorized and thus the model can easily generalize to new tasks (Santoro et al., 2016;
497 Munkhdalai and Yu, 2017; Ravi and Larochelle, 2016).

498 Model-Agnostic Meta Learning (MAML) method (Finn et al., 2017) opens another line of research,
499 i.e., it designs an optimization scheme so that the model can fast adapt to new tasks. Reptile (Nichol
500 and Schulman, 2018), a variant of MAML, is proposed to simplify the computation of MAML. None
501 of those works, however, considers the relational information between tasks like our method, which
502 is critical in modeling short sequences.

503 One interesting line of follow-up works of MAML is connecting MAML with Bayesian inference
504 (Finn et al., 2018; Ravi and Beatson, 2018; Grant et al., 2018). Since HARMLESS combines a
505 Bayesian model with MAML, it has the potential to be rewritten into a pure Bayesian model that has
506 better quantification of uncertainty. We left this for future work.

507 B Definition of Operator \mathcal{D}

508 As we mentioned earlier,

$$\min_{\theta} \sum_{\mathcal{T}_i \in \Gamma} \mathcal{F}_{\mathcal{T}_i}(\tilde{\theta}_i) = \sum_{\mathcal{T}_i \in \Gamma} \mathcal{F}_{\mathcal{T}_i}(\theta - \eta \mathcal{D}(\mathcal{F}_{\mathcal{T}_i}, \theta))$$

509 is the loss function for MAML, FOMAML, and Reptile algorithm with different definition of the
510 operator \mathcal{D} .

511 For simplicity, here we define the operator of one gradient step. The cases of few gradient steps can
512 be defined analogously.

513 For MAML, $\mathcal{D}(\mathcal{F}_{\mathcal{T}_i}, \theta)$ is defined as $\nabla_{\theta}(\mathcal{F}_{\mathcal{T}_i}(\theta))$.

514 For First Order MAML (FOMAML), $\mathcal{D}(\mathcal{F}_{\mathcal{T}_i}, \theta)$ is also defined as $\nabla_{\theta}(\mathcal{F}_{\mathcal{T}_i}(\theta))$. The difference is
515 that the output of the operator just a value, not a function of θ , i.e., when we solve the gradient of
516 $\mathcal{F}_{\mathcal{T}_i}(\theta - \eta\mathcal{D}(\mathcal{F}_{\mathcal{T}_i}, \theta))$, the gradient does not back-propagate into $\mathcal{D}(\mathcal{F}_{\mathcal{T}_i}, \theta)$.
517 For Reptile, the algorithm of reptile is as follows [Nichol and Schulman \(2018\)](#).

Algorithm 1 Reptile

while not converged **do**
 Sample task \mathcal{T} with loss $\mathcal{F}_{\mathcal{T}}$;
 $W \leftarrow \text{SGD}(\mathcal{F}_{\mathcal{T}}, \theta, k)$, where k is the number of SGD steps;
 Do the update $\theta \leftarrow \theta - \eta(\theta - W)$;
end while

518 From the algorithm we can see, operator \mathcal{D} is defined as $\mathcal{D}(\mathcal{F}_{\mathcal{T}}, \theta) = \text{SGD}(\mathcal{F}_{\mathcal{T}}, \theta, 1)$. Similar as
519 FOMAML, computing the gradient also does not back-propagate into $\mathcal{D}(\mathcal{F}_{\mathcal{T}_i}, \theta)$.

520 C Derivation of Variational EM

521 **Preparation** After adding latent variable \mathbf{z} , the joint distribution is

$$p(\mathbf{T}, \mathbf{Y}, \mathbf{z}, \mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}, \boldsymbol{\pi}) = p(\mathbf{T}|\mathbf{z})p(\mathbf{Y}|\mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow})p(\mathbf{z}|\boldsymbol{\pi})p(\mathbf{z}_{\leftarrow}|\boldsymbol{\pi})p(\mathbf{z}_{\rightarrow}|\boldsymbol{\pi})p(\boldsymbol{\pi}).$$

522 where

$$\begin{aligned} p(\mathbf{T}|\mathbf{z}) &= \prod_{i=1}^N \prod_{k=1}^K (\mathcal{L}_i(\theta_k - \eta\mathcal{D}(\mathcal{L}_i, \theta_k)))^{z_{i,k}}, \\ p(\mathbf{Y}|\mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}) &= \prod_{i=1}^N \prod_{j=1}^N (z_{i \rightarrow j}^T \mathbf{B} z_{i \leftarrow j})^{Y_{ij}} (1 - z_{i \rightarrow j}^T \mathbf{B} z_{i \leftarrow j})^{1 - Y_{ij}} \\ p(\mathbf{z}|\boldsymbol{\pi}) &= \prod_{i=1}^N \prod_{k=1}^K \pi_{i,k}^{z_{i,k}}, \\ p(\mathbf{z}_{\rightarrow}|\boldsymbol{\pi}) &= \prod_{i=1}^N \prod_{j=1}^N \prod_{k=1}^K \pi_{i,k}^{z_{i \rightarrow j, k}}, \\ p(\mathbf{z}_{\leftarrow}|\boldsymbol{\pi}) &= \prod_{i=1}^N \prod_{j=1}^N \prod_{k=1}^K \pi_{j,k}^{z_{i \leftarrow j, k}}, \\ p(\boldsymbol{\pi}) &= \prod_{i=1}^N \text{Dirichlet}(\pi_i|\alpha) = \prod_{i=1}^N C(\alpha) \prod_{k=1}^K \pi_{i,k}^{\alpha-1}. \end{aligned}$$

523 Note that in this section we represent $z_i, z_{i \rightarrow j}, z_{i \leftarrow j}$ as one-hot vector, while in the main paper we
524 use scalar $z_i = k$ representing the identities.

525 The posterior distribution is defined as

$$p(\mathbf{z}, \mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}, \boldsymbol{\pi}|\mathbf{T}, \mathbf{Y}, \alpha, \boldsymbol{\theta}, B).$$

526 We aim to find a distribution $q(\mathbf{z}, \mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}, \boldsymbol{\pi}) \in \mathcal{Q}$, such that the Kullback-Leibler (KL) divergence
527 between the above posterior distribution and $q(\mathbf{z}, \mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}, \boldsymbol{\pi})$ is minimized. This can be achieved by
528 maximize the Evidence Lower BOund (ELBO),

$$\mathcal{B}(q) = \mathbb{E}_q[\log p(\mathbf{z}, \mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}, \boldsymbol{\pi}, \mathbf{T}, \mathbf{Y})] - \mathbb{E}_q[\log q(\mathbf{z}, \mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}, \boldsymbol{\pi})].$$

529 **Variational family** We adopt the mean-field variational family, i.e.,

$$q(\mathbf{z}, \mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}, \boldsymbol{\pi}) = q_1(\boldsymbol{\pi}) \prod_i q_2(z_i) \prod_j q_3(z_{i \rightarrow j}) q_4(z_{i \leftarrow j}).$$

530 We pick $q_1(\pi_i)$ as PDF of Dirichlet(β), $q_2(z_i)$ as PDF of Categorical(γ_i), $q_3(z_{i \rightarrow j})$ as PDF of
 531 Categorical(ϕ_{ij}), $q_4(z_{i \leftarrow j})$ as PDF of Categorical(ψ_{ij}).

532 **Update for q_1** Again, our goal is to maximize

$$\mathcal{B}(q) = \mathbb{E}_q[\log p(\mathbf{z}, \mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}, \boldsymbol{\pi}, \mathbf{T}, \mathbf{Y})] - \mathbb{E}_q[\log q(\mathbf{z}, \mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}, \boldsymbol{\pi})].$$

533 Now we focus on q_1 , and treat q_2, q_3 and q_4 as given. We want to maximize

$$\begin{aligned} \mathcal{F}_{\boldsymbol{\pi}}(q_1) &= \mathbb{E}_q[\log p(\mathbf{z}, \mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}, \boldsymbol{\pi}, \mathbf{T}, \mathbf{Y})] - \mathbb{E}_q[\log q(\mathbf{z}, \mathbf{z}_{\rightarrow}, \mathbf{z}_{\leftarrow}, \boldsymbol{\pi})] \\ &= \mathbb{E}_q[\log p(\mathbf{T}|\mathbf{z}) + \log p(\mathbf{Y}|\mathbf{z}_{\leftarrow}, \mathbf{z}_{\rightarrow}) + \log p(\mathbf{z}|\boldsymbol{\pi}) + \log p(\mathbf{z}_{\leftarrow}|\boldsymbol{\pi}) + \log p(\mathbf{z}_{\rightarrow}|\boldsymbol{\pi}) + \log p(\boldsymbol{\pi})] \\ &\quad - \mathbb{E}_{q_1}[\log q_1(\boldsymbol{\pi})] + \text{const} \\ &= \mathbb{E}_q[\log p(\mathbf{z}|\boldsymbol{\pi}) + \log p(\mathbf{z}_{\leftarrow}|\boldsymbol{\pi}) + \log p(\mathbf{z}_{\rightarrow}|\boldsymbol{\pi}) + \log p(\boldsymbol{\pi})] - \mathbb{E}_{q_1}[\log q_1(\boldsymbol{\pi})] + \text{const} \\ &= \int q_1(\boldsymbol{\pi}) (\mathbb{E}_{q_2}[\log p(\mathbf{z}|\boldsymbol{\pi}) + \log p(\mathbf{z}_{\leftarrow}|\boldsymbol{\pi}) + \log p(\mathbf{z}_{\rightarrow}|\boldsymbol{\pi}) + \log p(\boldsymbol{\pi})] - \log q_1(\boldsymbol{\pi})) d\boldsymbol{\pi} + \text{const}. \end{aligned}$$

534 Take the derivative,

$$\frac{\delta \mathcal{F}_{\boldsymbol{\pi}}(q_1)}{\delta q_1} = \mathbb{E}_{q_2}[\log p(\mathbf{z}|\boldsymbol{\pi}) + \log p(\mathbf{z}_{\leftarrow}|\boldsymbol{\pi}) + \log p(\mathbf{z}_{\rightarrow}|\boldsymbol{\pi}) + \log p(\boldsymbol{\pi})] - \log q_1(\boldsymbol{\pi}) - 1 = 0.$$

535 Substitute the expressions of the distributions, after some derivation we get the update for β as

$$\beta_{i,k} \leftarrow \alpha_k + \gamma_{i,k} + \sum_{j=1}^N \phi_{ij,k} + \sum_{j=1}^N \psi_{ij,k}. \quad (15)$$

536 **Update for q_2** Similarly, we have

$$\begin{aligned} \mathcal{F}_{\mathbf{z}}(q_2) &= \mathbb{E}_q[\log p(\mathbf{T}|\mathbf{z}) + \log p(\mathbf{z}|\boldsymbol{\pi})] - \mathbb{E}_{q_2}[\log q_2(\mathbf{z})] + \text{const} \\ &= \int q_2(\mathbf{z}) (\mathbb{E}_{q_1}[\log p(\mathbf{T}|\boldsymbol{\theta}, \mathbf{z}) + \log p(\mathbf{z}|\boldsymbol{\pi})] - \log q_2(\mathbf{z})) d\mathbf{z} + \text{const}. \end{aligned}$$

537 Take the derivative,

$$\frac{\delta \mathcal{F}_{\mathbf{z}}(q_2)}{\delta q_2} = \log p(\mathbf{T}|\boldsymbol{\theta}, \mathbf{z}) + \mathbb{E}_{q_1}[\log p(\mathbf{z}|\boldsymbol{\pi})] - \log q_2(\mathbf{z}) - 1 = 0.$$

538 After some derivation, we have

$$\gamma_{i,k} \leftarrow \mathcal{L}_i(\theta_k - \eta \mathcal{D}(\mathcal{L}_i, \theta_k)) \exp \left(f_{\text{dg}}(\beta_{i,k}) - f_{\text{dg}} \left(\sum_{\ell} \beta_{i,\ell} \right) \right), \quad (16)$$

$$\gamma_{i,k} \leftarrow \frac{\gamma_{i,k}}{\sum_{\ell} \gamma_{i,\ell}}, \quad (17)$$

539 where f_{dg} is the digamma function.

540 **Update for q_3 and q_4** The derivation of update for q_3 and q_4 is very similar to the update for q_2 , so
 541 we will not elaborate on that. Readers who are interested might also refer to [Airoldi et al. \(2008\)](#).
 542 The updates are

$$\phi_{ij,k} \leftarrow e^{\mathbb{E}_q[\log \pi_{i,k}]} \prod_{\ell=1}^K \left(B_{k\ell}^{Y_{ij}} (1 - B_{k\ell})^{1-Y_{ij}} \right)^{\psi_{ij,\ell}}, \quad \phi_{ij,k} \leftarrow \frac{\phi_{ij,k}}{\sum_{\ell} \phi_{ij,\ell}}, \quad (18)$$

$$\psi_{ij,\ell} \leftarrow e^{\mathbb{E}_q[\log \pi_{j,\ell}]} \prod_{k=1}^K \left((B_{k\ell})^{Y_{ij}} (1 - B_{k\ell})^{1-Y_{ij}} \right)^{\phi_{ij,k}}, \quad \psi_{ij,k} \leftarrow \frac{\psi_{ij,k}}{\sum_{\ell} \psi_{ij,\ell}}, \quad (19)$$

543 **Update for $\boldsymbol{\theta}$** We update $\boldsymbol{\theta}$ using gradient ascent. We first pick the terms that is relevant to $\boldsymbol{\theta}$,

$$\begin{aligned} \mathcal{F}_{\boldsymbol{\theta}}(\boldsymbol{\theta}) &= \mathbb{E}_q[\log p(\mathbf{T}|\boldsymbol{\theta}, \mathbf{z})] + \text{const} \\ &= \int q_2(\mathbf{z}) [\log p(\mathbf{T}|\boldsymbol{\theta}, \mathbf{z})] d\mathbf{z} + \text{const} \\ &= \sum_{i=1}^N \sum_{k=1}^K \gamma_{i,k} \log \mathcal{L}_i(\theta_k - \eta \mathcal{D}(\mathcal{L}_i, \theta_k)) + \text{const}. \end{aligned}$$

544 So the gradient ascent update is,

$$\theta \leftarrow \theta + \eta_1 \nabla_{\theta} \left(\sum_{i=1}^N \sum_{k=1}^K \gamma_{i,k} \log \mathcal{L}_i(\theta_k - \eta \mathcal{D}(\mathcal{L}_i, \theta_k)) \right). \quad (20)$$

545 **Update for α and B** From [Airoldi et al. \(2008\)](#), we have the update for α and B as follows

$$\alpha_k \leftarrow \alpha_k + \eta_{\alpha} \left(N(f_{\text{dg}}(\sum_{\ell} \alpha_{\ell}) - f_{\text{dg}}(\alpha_k)) + \sum_{i=1}^N (f_{\text{dg}}(\beta_{i,k}) - f_{\text{dg}}(\sum_{\ell} \beta_{i,\ell})) \right), \quad (21)$$

$$B_{k\ell} \leftarrow \frac{\sum_{ij} Y_{ij} \phi_{ij,k} \psi_{ij,\ell}}{\sum_{ij} \phi_{ij,k} \psi_{ij,\ell}}, \quad (22)$$

546 D Derivation of Evaluation Metric

547 In this section, we give more details on the evaluate metrics. Specifically, we show how to compute
 548 the NLL of the test set. Given a sequence $\tau_i = \{\tau_i^{(1)}, \tau_i^{(2)}, \dots, \tau_i^{(M_i)}\}$, we would like to predict the
 549 timestamp of $\tau_i^{(M_i+1)}$. Here, we use the probability of the arrival at time $\tau_i^{(M_i+1)}$ and no arrival in
 550 $[\tau_i^{(M_i)}, \tau_i^{(M_i+1)}]$ given history before $\tau_i^{(M_i)}$ as evaluation metric.

551 Consider a Hawkes process with parameter θ , the probability density is

$$\begin{aligned} \mathcal{P}(\theta) &= \lambda(\tau_i^{(M_i+1)}; \theta, \tau_i) \exp \left(- \int_{\tau_i^{(M_i)}}^{\tau_i^{(M_i+1)}} \lambda(t; \theta, \tau_i) dt \right) \\ &= \left(\mu + \sum_{m=1}^{M_i} \delta \omega e^{-\omega(\tau_i^{(M_i+1)} - \tau_i^{(m)})} \right) \exp \left(-\mu(\tau_i^{(M_i+1)} - \tau_i^{(M_i)}) - \delta(1 - e^{-\omega(\tau_i^{(M_i+1)} - \tau_i^{(M_i)})}) \right). \end{aligned}$$

552 In the generative process, for subject i , we first sample z_i , then use parameter $\tilde{\theta}_{z_i}^{(i)} = \theta_{z_i} - \eta \mathcal{D}(\mathcal{L}_i, \theta_{z_i})$.
 553 The posterior distribution of z_i is $q_2(z_i)$, *i.e.*, Categorical(γ_i). Therefore we have

$$\mathbb{P}(z_i = k) = \gamma_{i,k}.$$

554 So the likelihood of next arrival $\tau_i^{(M_i+1)}$ is

$$\begin{aligned} \tilde{\mathcal{L}}_i &= \sum_{k=1}^K \mathbb{P}(z_i = k) \mathbb{P}(\text{next arrival is } \tau_i^{(M_i+1)} \mid \text{Hawkes model with } \theta_k) \\ &= \sum_{k=1}^K \gamma_{i,k} \mathcal{P}(\tilde{\theta}_k^{(i)}). \end{aligned}$$

555 And then we sum $\tilde{\mathcal{L}}_i$ over every subject.

556 E Detailed Settings of the Experiments

557 Note that we can also adopt a non-informative α instead of updating it in every iteration. After
 558 some trial experiments, we find setting $\alpha = \mathbf{1}_K$ is numerically more stable than updating it in every
 559 iteration. Therefore we adopt $\alpha = \mathbf{1}_K$ in the following experiments.

560 Besides, we find that ν causes nearly no effect to the result when varying from 10^{-10} to 10^{-1} . We
 561 fix it as 10^{-2} .

562 E.1 Synthetic Dataset

563 Both the baselines and our proposed methods are fine tuned. We first perform a coarse grid search to
 564 find hyper-parameters for all methods. The grid search finds learning rate from 1×10^{-7} to 1 for
 565 both inner and outer updates. To perform the multi-split procedure, all hyper-parameters are then
 566 selected in the following range listed in Table 4 and Table 5. For each range, we perform experiment
 567 on three values: the lower one, the upper one, and the middle one. Method *MTL* adopt $\nu_{\text{mtl}} = 0.1$.

Table 4: Learning rates of experiments.

K_0		1	3	6	10
DMHP	lr.	$1 \pm .1 \times 10^{-3}$	$3 \pm .1 \times 10^{-3}$	$6.5 \pm .1 \times 10^{-3}$	$7 \pm .1 \times 10^{-3}$
Two Step	inner lr.	$1 \pm .1 \times 10^{-5}$	$5 \pm .1 \times 10^{-5}$	$5 \pm .1 \times 10^{-5}$	$1 \pm .1 \times 10^{-4}$
	outer lr.	$1 \pm .1 \times 10^{-3}$	$1 \pm .1 \times 10^{-2}$	$1.5 \pm .1 \times 10^{-2}$	$1 \pm .1 \times 10^{-2}$
HARMLESS (MAML)	inner lr.	$5 \pm .1 \times 10^{-5}$	$5 \pm .1 \times 10^{-6}$	$2 \pm .1 \times 10^{-4}$	$7 \pm .1 \times 10^{-5}$
	outer lr.	$6 \pm .1 \times 10^{-4}$	$2 \pm .1 \times 10^{-4}$	$6 \pm .1 \times 10^{-5}$	$4.5 \pm .1 \times 10^{-6}$
HARMLESS (FOMAML)	inner lr.	$5 \pm .1 \times 10^{-4}$	$1 \pm .1 \times 10^{-5}$	$3 \pm .1 \times 10^{-5}$	$1.5 \pm .1 \times 10^{-6}$
	outer lr.	$6 \pm .1 \times 10^{-4}$	$2 \pm .1 \times 10^{-4}$	$6 \pm .1 \times 10^{-5}$	$4.5 \pm .1 \times 10^{-6}$

Table 5: Learning rates of baseline experiments.

Method	Learning Rate
MLE-Sep	$5 \pm .1 \times 10^{-5}$
MLE-Com	$1 \pm .1 \times 10^{-3}$
MTL	$1 \pm .1 \times 10^{-3}$

568 E.2 Real Datasets

569 In this section, we introduce the experimental detail of the real datasets. We run our experiment
570 with same inner and outer learning rate, denoted by η . For simplicity, we also set $\eta = \eta_\alpha = \eta_\theta$,
571 and search over $\{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\} \otimes \{1, 2, 3, 4, 5\}$, where the element-wise product of two
572 sets is defined as $A \otimes B = \{ab | a \in A, b \in B\}$. We search $K \in \{2, 3, 5\}$ and ν_{mtl} in range
573 $\{0.1, 0.01, 0.001\}$. We perform grid search over the hyper-parameters, and obtain the candidate
574 models. Then we perform multi-split procedure.

575 Because StackOverflow dataset is very large, it is too expensive to perform grid search. To accom-
576 modate this, we first split a validation set and a test set, then performing hyper-parameter search by
577 flipping. Each experiment of StackOverflow dataset is run under 5 different settings.

578 In Table 6 we report one of the models that is picked by multi-split procedure. We remark that in
579 most cases, the procedure picks only one model repeatedly.

Table 6: Settings of experiments.

data type	911-Calls	LinkedIn	MathOverflow	StackOverflow
Baseline 1	$\eta = 4 \times 10^{-4}$	$\eta = 1 \times 10^{-3}$	$\eta = 5 \times 10^{-4}$	$\eta = 5 \times 10^{-4}$
Baseline 2	$\eta = 3 \times 10^{-4}$	$\eta = 5 \times 10^{-3}$	$\eta = 1 \times 10^{-3}$	$\eta = 1 \times 10^{-3}$
MTL	$\eta = 3 \times 10^{-5}, \nu_{\text{mtl}} = 0.1$	$\eta = 1 \times 10^{-2}, \nu_{\text{mtl}} = 0.1$	$\eta = 4 \times 10^{-4}, \nu_{\text{mtl}} = 0.1$	$\eta = 5 \times 10^{-4}, \nu_{\text{mtl}} = 0.1$
DMHP	$\eta = 3 \times 10^{-5}, K = 2$	$\eta = 1 \times 10^{-3}, K = 3$	$\eta = 4 \times 10^{-3}, K = 3$	$N \setminus A$
MAML	$\eta = 3 \times 10^{-4}, K = 3$	$\eta = 5 \times 10^{-1}, K = 3$	$\eta = 3 \times 10^{-4}, K = 3$	$\eta = 1 \times 10^{-3}, K = 2$
FOMAML	$\eta = 3 \times 10^{-5}, K = 2$	$\eta = 1 \times 10^{-2}, K = 5$	$\eta = 2 \times 10^{-4}, K = 2$	$\eta = 4 \times 10^{-4}, K = 3$
Reptile	$\eta = 5 \times 10^{-3}, K = 2$	$\eta = 2 \times 10^{-1}, K = 3$	$\eta = 4 \times 10^{-2}, K = 2$	$\eta = 4 \times 10^{-2}, K = 2$

580 E.3 Ablation study

581 In this section we introduce the experimental detail of the ablation study. Specifically, the tuning
582 process of the ablation study is as follows: We start from the same setting as the corresponding real
583 experiment in previous section. For example, experiment *Remove graph (FOMAML)* corresponds to
584 HARMLESS (FOMAML). We first use the same learning rate and K as HARMLESS (FOMAML)
585 to perform experiment. If the experiment runs well, we adopt the experiment result. If the training
586 does not converge, we decrease the learning rate and run again.

587 References

588 ACHAB, M., BACRY, E., GAÏFFAS, S., MASTROMATTEO, I. and MUZY, J.-F. (2017). Uncovering
589 causality from multivariate hawkes integrated cumulants. *The Journal of Machine Learning*
590 *Research*, **18** 6998–7025.

Table 7: Learning rates of experiments of ablation study.

data type	LR
Remove inner heterogeneity ($K = 3$)	0.1
Remove inner heterogeneity ($K = 5$)	0.1
Remove grouping (MAML)	0.1
Remove grouping (FOMAML)	0.01
Remove grouping (Reptile)	0.2
Remove graph (MAML)	0.2
Remove graph (FOMAML)	0.005
Remove graph (Reptile)	0.2

- 591 AIROLDI, E. M., BLEI, D. M., FIENBERG, S. E. and XING, E. P. (2008). Mixed membership
592 stochastic blockmodels. *Journal of machine learning research*, **9** 1981–2014.
- 593 BACRY, E., DAYRI, K. and MUZY, J.-F. (2012). Non-parametric kernel estimation for symmetric
594 hawkes processes. application to high frequency financial data. *The European Physical Journal B*,
595 **85** 157.
- 596 BAUWENS, L. and HAUTSCH, N. (2009). Modelling financial high frequency data using point
597 processes. In *Handbook of financial time series*. Springer, 953–979.
- 598 BENGIO, Y., BENGIO, S. and CLOUTIER, J. (1990). *Learning a synaptic learning rule*. Université
599 de Montréal, Département d’informatique et de recherche ?
- 600 BLEI, D. M., KUCUKELBIR, A. and MCAULIFFE, J. D. (2017). Variational inference: A review for
601 statisticians. *Journal of the American Statistical Association*, **112** 859–877.
- 602 BLUNDELL, C., BECK, J. and HELLER, K. A. (2012). Modelling reciprocating relationships with
603 hawkes processes. In *Advances in Neural Information Processing Systems*.
- 604 BOX, G. E. and TIAO, G. C. (2011). *Bayesian inference in statistical analysis*, vol. 40. John Wiley
605 & Sons.
- 606 CHALMERS, D. J. (1991). The evolution of learning: An experiment in genetic connectionism. In
607 *Connectionist Models*. Elsevier, 81–90.
- 608 CLEEREMANS, A. and MCCLELLAND, J. L. (1991). Learning the structure of event sequences.
609 *Journal of Experimental Psychology: General*, **120** 235.
- 610 EICHLER, M., DAHLHAUS, R. and DUECK, J. (2017). Graphical modeling for multivariate hawkes
611 processes with nonparametric link functions. *Journal of Time Series Analysis*, **38** 225–242.
- 612 FARAJTABAR, M., YANG, J., YE, X., XU, H., TRIVEDI, R., KHALIL, E., LI, S., SONG, L. and
613 ZHA, H. (2017). Fake news mitigation via point process based intervention. In *Proceedings of the*
614 *34th International Conference on Machine Learning-Volume 70*. JMLR. org.
- 615 FARAJTABAR, M., YE, X., HARATI, S., SONG, L. and ZHA, H. (2016). Multistage campaigning in
616 social networks. In *Advances in Neural Information Processing Systems*.
- 617 FINN, C., ABBEEL, P. and LEVINE, S. (2017). Model-agnostic meta-learning for fast adaptation of
618 deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume*
619 *70*. JMLR. org.
- 620 FINN, C., XU, K. and LEVINE, S. (2018). Probabilistic model-agnostic meta-learning. In *Advances*
621 *in Neural Information Processing Systems*.
- 622 FOX, E. W., SHORT, M. B., SCHOENBERG, F. P., CORONGES, K. D. and BERTOZZI, A. L. (2016).
623 Modeling e-mail networks and inferring leadership using self-exciting point processes. *Journal of*
624 *the American Statistical Association*, **111** 564–584.
- 625 GIRVAN, M. and NEWMAN, M. E. (2002). Community structure in social and biological networks.
626 *Proceedings of the national academy of sciences*, **99** 7821–7826.

- 627 GRANT, E., FINN, C., LEVINE, S., DARRELL, T. and GRIFFITHS, T. (2018). Recasting gradient-
628 based meta-learning as hierarchical bayes. *arXiv preprint arXiv:1801.08930*.
- 629 HANSEN, N. R., REYNAUD-BOURET, P., RIVOIRARD, V. ET AL. (2015). Lasso and probabilistic
630 inequalities for multivariate point processes. *Bernoulli*, **21** 83–143.
- 631 HAWKES, A. G. (1971). Spectra of some self-exciting and mutually exciting point processes.
632 *Biometrika*, **58** 83–90.
- 633 HOFFMAN, M. D., BLEI, D. M., WANG, C. and PAISLEY, J. (2013). Stochastic variational inference.
634 *The Journal of Machine Learning Research*, **14** 1303–1347.
- 635 KARRER, B. and NEWMAN, M. E. (2011). Stochastic blockmodels and community structure in
636 networks. *Physical review E*, **83** 016107.
- 637 KOBAYASHI, R. and LAMBIOTTE, R. (2016). Tideh: Time-dependent hawkes process for predicting
638 retweet dynamics. In *Tenth International AAAI Conference on Web and Social Media*.
- 639 KOCH, G., ZEMEL, R. and SALAKHUTDINOV, R. (2015). Siamese neural networks for one-shot
640 image recognition. In *ICML deep learning workshop*, vol. 2.
- 641 LAUB, P. J., TAIMRE, T. and POLLETT, P. K. (2015). Hawkes processes. *arXiv preprint*
642 *arXiv:1507.02822*.
- 643 LI, L. and ZHA, H. (2013). Dyadic event attribution in social networks with mixtures of hawkes
644 processes. In *Proceedings of the 22nd ACM international conference on Information & Knowledge*
645 *Management*. ACM.
- 646 LINDERMAN, S. and ADAMS, R. (2014). Discovering latent network structure in point process data.
647 In *International Conference on Machine Learning*.
- 648 LUO, D., XU, H., ZHEN, Y., NING, X., ZHA, H., YANG, X. and ZHANG, W. (2015). Multi-task
649 multi-dimensional hawkes processes for modeling event sequences. In *Twenty-Fourth International*
650 *Joint Conference on Artificial Intelligence*.
- 651 MACLAURIN, D., DUVENAUD, D. and ADAMS, R. (2015). Gradient-based hyperparameter opti-
652 mization through reversible learning. In *International Conference on Machine Learning*.
- 653 MEI, H. and EISNER, J. M. (2017). The neural hawkes process: A neurally self-modulating
654 multivariate point process. In *Advances in Neural Information Processing Systems*.
- 655 MUNKHDALAI, T. and YU, H. (2017). Meta networks. In *Proceedings of the 34th International*
656 *Conference on Machine Learning-Volume 70*. JMLR. org.
- 657 NICHOL, A., ACHIAM, J. and SCHULMAN, J. (2018). On first-order meta-learning algorithms.
658 *arXiv preprint arXiv:1803.02999*.
- 659 NICHOL, A. and SCHULMAN, J. (2018). Reptile: a scalable metalearning algorithm. *arXiv preprint*
660 *arXiv:1803.02999*.
- 661 OGATA, Y. (1999). Seismicity analysis through point-process modeling: A review. In *Seismicity*
662 *patterns, their statistical significance and physical meaning*. Springer, 471–507.
- 663 PARANJAPE, A., BENSON, A. R. and LESKOVEC, J. (2017). Motifs in temporal networks. In
664 *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. ACM.
- 665 RASMUSSEN, J. G. (2013). Bayesian inference for hawkes processes. *Methodology and Computing*
666 *in Applied Probability*, **15** 623–642.
- 667 RAVI, S. and BEATSON, A. (2018). Amortized bayesian meta-learning.
- 668 RAVI, S. and LAROCHELLE, H. (2016). Optimization as a model for few-shot learning.
- 669 REYNAUD-BOURET, P., SCHBATH, S. ET AL. (2010). Adaptive estimation for hawkes processes;
670 application to genome analysis. *The Annals of Statistics*, **38** 2781–2822.

- 671 ROSS, S. M., KELLY, J. J., SULLIVAN, R. J., PERRY, W. J., MERCER, D., DAVIS, R. M.,
672 WASHBURN, T. D., SAGER, E. V., BOYCE, J. B. and BRISTOW, V. L. (1996). *Stochastic*
673 *processes*, vol. 2. Wiley New York.
- 674 SANTORO, A., BARTUNOV, S., BOTVINICK, M., WIERSTRA, D. and LILICRAP, T. (2016).
675 Meta-learning with memory-augmented neural networks. In *International conference on machine*
676 *learning*.
- 677 SNELL, J., SWERSKY, K. and ZEMEL, R. (2017). Prototypical networks for few-shot learning. In
678 *Advances in Neural Information Processing Systems*.
- 679 SUNG, F., YANG, Y., ZHANG, L., XIANG, T., TORR, P. H. and HOSPEDALES, T. M. (2018).
680 Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE*
681 *Conference on Computer Vision and Pattern Recognition*.
- 682 TRAN, L., FARAJTABAR, M., SONG, L. and ZHA, H. (2015). Netcodec: Community detection from
683 individual activities. In *Proceedings of the 2015 SIAM International Conference on Data Mining*.
684 SIAM.
- 685 TRIVEDI, R., FARAJTABAR, M., BISWAL, P. and ZHA, H. (2018). Dyrep: Learning representations
686 over dynamic graphs.
- 687 VINYALS, O., BLUNDELL, C., LILICRAP, T., WIERSTRA, D. ET AL. (2016). Matching networks
688 for one shot learning. In *Advances in neural information processing systems*.
- 689 XIE, J., KELLEY, S. and SZYMANSKI, B. K. (2013). Overlapping community detection in networks:
690 The state-of-the-art and comparative study. *Acm computing surveys (csur)*, **45** 43.
- 691 XU, H., LUO, D., CHEN, X. and CARIN, L. (2017a). Benefits from superposed hawkes processes.
692 *arXiv preprint arXiv:1710.05115*.
- 693 XU, H., LUO, D. and ZHA, H. (2017b). Learning hawkes processes from short doubly-censored event
694 sequences. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*.
695 JMLR. org.
- 696 XU, H. and ZHA, H. (2017). A dirichlet mixture model of hawkes processes for event sequence
697 clustering. In *Advances in Neural Information Processing Systems*.
- 698 YANG, S.-H. and ZHA, H. (2013). Mixture of mutually exciting processes for viral diffusion. In
699 *International Conference on Machine Learning*.
- 700 ZAREZADE, A., KHODADADI, A., FARAJTABAR, M., RABIEE, H. R. and ZHA, H. (2017). Corre-
701 lated cascades: Compete or cooperate. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- 702 ZHANG, Y. and YANG, Q. (2017). A survey on multi-task learning. *arXiv preprint arXiv:1707.08114*.
- 703 ZHAO, Q., ERDOGDU, M. A., HE, H. Y., RAJARAMAN, A. and LESKOVEC, J. (2015). Seismic: A
704 self-exciting point process model for predicting tweet popularity. In *Proceedings of the 21th ACM*
705 *SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM.
- 706 ZHOU, K., ZHA, H. and SONG, L. (2013). Learning social infectivity in sparse low-rank networks
707 using multi-dimensional hawkes processes. In *Artificial Intelligence and Statistics*.