applied sciences

MDPI

Communication

Scientific Impact Prediction via Virtual Geography Hawkes Process

Babusurya Ganeshbabu [1], Xin Liu [2], Akiyoshi Matono [2], Kyoung-Sook Kim [2] and Xun Shen [1,*]

[1] Graduate School of Advanced Interdisciplinary Science, Tokyo University of Agriculture and Technology, Koganei 184-8588, Tokyo, Japan
[2] Intelligent Platforms Research Institute, National Institute of Advanced Industrial Science and Technology (AIST), Koto 135-0064, Tokyo, Japan; ks.kim@aist.go.jp (K.K.)
* Correspondence: shen@go.tuat.ac.jp

Abstract

This brief proposes a novel Virtual-Geography Hawkes Process (VG-Hawkes) to model citation dynamics considering academic networks. The VG-Hawkes model incorporates academic relationships between authors as virtual distances by extending the conventional temporal Hawkes process, enabling a more detailed and realistic representation of citation behavior. Validation results on real-world datasets show that the VG-Hawkes model consistently achieves higher log-likelihood scores than temporal Hawkes models and effectively captures citation peaks and distributional patterns. While this study focuses on selected datasets and pairwise interactions, the model is general and readily extensible. Future work includes scaling to broader datasets and incorporating more complex author relationships. The VG-Hawkes model provides a novel and flexible framework for academic network analysis and scientific impact prediction.

Keywords: citation prediction; point process; kernel method; maximum likelihood estimation; expectation maximization algorithm

1. Introduction

1.1. Background

There is growing demand for methods that can identify high-quality recent research within the rapidly growing body of literature [1–3] Citation count is widely used as a fundamental metric for evaluating a researcher's scientific impact. Popular indicators include the impact factor [4], h-index [5], and g-index [6]. The impact factor reflects the average number of citations received per year by articles published in a journal [7], while the h-index is defined as the largest number h, such that a researcher has h papers each cited at least h times. The g-index refines this by identifying the largest g such that the top g articles collectively have at least g^2 citations. These citation-based indices are commonly adopted in academic hiring, promotion, and funding decisions, as they provide quantitative estimates of scientific importance and influence. However, such metrics primarily capture past performance and cannot predict the future impact of a publication. For instance, even papers published in the same journal can diverge widely in their long-term citation counts—from a handful to over a thousand. This variability illustrates that high current metrics do not guarantee sustained influence, while researchers with modest early metrics may later produce groundbreaking work. Consequently, metrics like the impact factor and h-index, which reflect only current evaluation, may not offer fair or forward-looking criteria for assessment.

Check for updates

Academic Editors: Eric Rondeau, Karl Andersson and Ah-Lian Kor

Received: 20 January 2026
Revised: 11 February 2026
Accepted: 12 February 2026
Published: 20 February 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1.2. Research Aim and Objectives

This study aims to develop a structure-aware point-process framework for modeling and analyzing citation dynamics by extending spatiotemporal Hawkes processes with virtual academic proximity that reflects author relationships (VG-Hawkes).

To achieve this aim, we pursue the following objectives: (i) formulate a Hawkes-process intensity in which excitation is modulated by virtual distances derived from academic relationships rather than physical geography; (ii) provide a modeling mechanism that captures both self-excitation and mutual excitation among authors/publications under the same probabilistic framework; and (iii) demonstrate feasibility on real-world citation records by showing improved likelihood-based fit and interpretable intensity patterns compared with temporal Hawkes baselines.

1.3. Contributions

Existing approaches to academic citation modeling often enhance point-process or temporal models by introducing author/paper representations as covariates (e.g., embedding-based features), by imposing network kernels that encode observed graph proximity, or by directly parameterizing relational intensities using pre-defined links or similarity measures. While effective in their respective settings, these designs typically treat the network structure as either an external feature source or a fixed coupling mechanism.

In contrast, our notion of virtual geography provides a unified and explicitly interpretable mechanism to organize heterogeneous academic interactions into a latent "distance" space that directly modulates excitation in a Hawkes-process intensity. Rather than only injecting embeddings as static covariates or prescribing a specific network kernel, virtual geography induces a structured interaction map that governs how events propagate across entities and interaction types, enabling a consistent treatment of citations and other interaction channels under a single temporal point-process framework.

The main contributions are summarized as follows.

C1 **Virtual-geography Hawkes modeling.** We propose a Hawkes-process formulation in which excitation is modulated by a virtual-geography structure, offering an interpretable latent organization of interaction strength beyond fixed graph kernels or purely feature-driven covariates.

C2 **Unified modeling of heterogeneous interactions.** We provide a unified event-driven framework that can incorporate multiple academic interaction channels (e.g., citations and collaboration-related signals) through a shared intensity construction, instead of handling each relation with separate ad hoc models.

C3 **Learning procedure and feasibility demonstration.** We present a practical learning pipeline for the proposed model and provide an initial empirical validation demonstrating feasibility and the added value of the virtual-geography mechanism within the Communication scope.

2. Related Work

2.1. Citation-Impact and Scientometric Models

Forecasting long-term citation impact is valuable for research assessment, funding decisions, and discovery systems that rely on citation-driven relevance signals. Beyond descriptive indices (e.g., impact factor, h-index), several predictive models have been proposed to explain and forecast citation accumulation. For example, Wang et al. [8] introduced a generative model combining a publication's fitness and longevity to support long-horizon citation forecasting, and Sinatra et al. [9] proposed the Q-model to quantify an author's intrinsic scientific ability and its effect on impact trajectories. While influential, these approaches typically do not explicitly model relational excitation mechanisms induced by

academic interactions such as co-authorship ties and citation-network structure, which are central to network-driven propagation of influence.

2.2. Point-Process and Hawkes-Process Modeling of Citation Dynamics

Temporal point processes provide a principled probabilistic framework for modeling event-time data and have been widely used to capture self-excitation, where past events increase the likelihood of future events. In particular, Hawkes processes [10] offer an interpretable mechanism for event clustering through a baseline intensity and an excitation kernel. This paradigm has been adopted in multiple domains and is naturally suited for modeling citation events over time because citations often exhibit bursts and long-tail decay. Likelihood-based learning procedures, including EM-type methods, further make Hawkes models attractive for data-driven inference.

2.3. Network-Aware Extensions and Representation-Based Approaches

A key challenge in citation modeling is to incorporate academic relationships, such as co-authorship, topical similarity, and influence between communities, in a way that goes beyond purely temporal recurrence [11,12]. One direction augments temporal models with network-derived covariates or representation learning (e.g., author/paper embeddings) to capture latent similarity, while another direction introduces relational coupling via graph-based kernels or interaction functions that depend on distances or proximities in an underlying space [13–15]. Relatedly, spatiotemporal Hawkes processes incorporate spatial dependence through physical distance [16,17], which can be effective when geography is meaningful; however, in academic networks, physical proximity is often less relevant than latent scholarly proximity.

2.4. Positioning of VG-Hawkes

Motivated by the above, we propose the Virtual-Geography Hawkes Process (VG-Hawkes), which adapts the spatiotemporal Hawkes principle [17,18] to academic networks by replacing physical distance with virtual academic distance. This design provides an interpretable mechanism to modulate excitation by latent scholarly proximity, enabling a unified event-driven model that captures both temporal clustering and network-structured mutual excitation. Compared with approaches that treat network structure only as external features or fixed graph proximity, VG-Hawkes introduces a distance-based interaction map tailored to academic relationships, while retaining likelihood-based inference and extensibility within the Hawkes framework.

3. Method

3.1. Concept

We conceptualize academic network formation as a hierarchical clustering process, in which researchers gather around thematic topics, form collaborative clusters, and establish new co-authorship ties through scholarly interaction. We hypothesize that the structural evolution of these networks is driven by both self-excitation (e.g., influence from a researcher's own prior work) and mutual excitation (e.g., influence from others' work or external publications). An overview of this conceptual framework is illustrated in the left side of Figure 1. Leveraging the Virtual-Geography Hawkes Process (VG-Hawkes), our model naturally captures the group structure and scale-free characteristics observed in academic networks. Author relationships are quantitatively encoded as distances in a virtual space, enabling explicit and interpretable modeling of both citation and co-authorship dynamics. This approach offers a novel perspective for analyzing the evolving structure of academic networks. While modeling cross-disciplinary influences and interactions between

otherwise unrelated authors remains a subject for future exploration, the proposed method provides a solid foundation for structurally informed citation impact prediction.

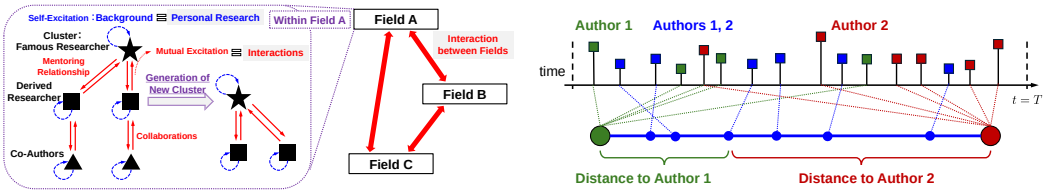


Figure 1. Intuitive explanation of the concept and overview of the proposed method. The (**left**) side: hierarchical structure of academic networks. The (**right**) side: overview of the proposed method.

3.2. Overview of the Proposed Method

This study introduces the virtual network Hawkes process, an extension of the conventional spatiotemporal Hawkes process, designed to model academic relationships between authors as virtual spatial positions. Rather than relying on physical distances, the proposed model defines virtual distances based on citation and co-authorship relationships, enabling it to capture the latent structural characteristics of academic networks. By doing so, the model effectively represents the temporal dynamics of research paper citations with the academic proximity of authors. This concept is illustrated on the right side of Figure 1. The horizontal axis denotes time, and each marker (square) corresponds to an event, interpreted as a paper publication occurring at a particular time point. Events are color-coded by author: green for Author 1, red for Author 2, and blue for publications associated with both Author 1 and Author 2. This visualization reflects how events propagate and cluster over time, conditioned on virtual distances in the latent academic space.

3.3. Hawkes Process

The Hawkes process [10] is a probabilistic point process that accounts for the possibility of one event "triggering" another. This process exhibits self-exciting behavior, where the occurrence of an event temporarily increases the likelihood of subsequent events, resulting in a clustering pattern. Due to this property, the Hawkes process has been widely applied in various domains such as modeling aftershock sequences in seismology and the spread of infectious diseases. The Hawkes process consists mainly of two components: a background rate and an excitation term. Its intensity function is defined as follows:

$$\lambda(t) = \mu + \sum_{t_i < t} g(t - t_i). \tag{1}$$

Here, μ represents the baseline rate of spontaneous events (background events), and $g(t - t_i)$ denotes the triggering kernel that quantifies the likelihood of a direct offspring event generated by a previous event occurring at time t_i. Both μ and g are non-negative functions. Through this excitation mechanism, the Hawkes process enables dynamic prediction of future events based on the history of past events.

While the standard Hawkes process focuses on the temporal characteristics of events, it cannot accommodate additional attributes, such as the magnitude of an earthquake or the severity of an infectious disease case. The marked Hawkes process extends the original model by incorporating additional information (called "marks") associated with each event. The intensity function is extended as follows:

$$\lambda(t, m) = \mu(t, m) + \sum_{k: t_k < t} g(t - t_k, m_k) \tag{2}$$

In many systems, the intensity function depends not only on time but also on spatial location. When the standard Hawkes process is extended to incorporate both temporal and spatial dimensions, the conditional intensity function is expressed as follows:

$$\lambda(t,x) = \mu(t,x) + \sum_{k:t_k<t} g(t-t_k, x-x_k) \tag{3}$$

In this spatio-temporal Hawkes model, the excitation function $g(t-t_k, x-x_k)$ represents spatial dependency based on the physical distance between the location $\mathbf{x}_k$ of a past event and the current location $\mathbf{x}$.

3.4. Virtual Network Hawkes Process

In academic network structures such as citation and co-authorship relationships, physical distance is not necessarily a dominant factor influencing the likelihood of event occurrence. For instance, in citation or collaboration networks, the influence between authors is more strongly determined by academic relevance or the thematic similarity of research topics rather than by geographic proximity. To address this, we formulate the model using a Virtual Network Hawkes Process, in which the notion of "physical distance" is replaced with a "virtual distance." The virtual distance defined in this study quantifies the academic closeness between researchers or papers. It is computed based on factors such as citation relationships and co-authorship connections. A smaller virtual distance indicates a stronger academic connection, while a larger distance represents weaker relevance. This virtual distance is then incorporated into the model as a substitute for physical spatial distance. We formulate a model that captures the citation dynamics of academic papers by integrating the Hawkes process framework with virtual academic distance. Let the virtual position of author i in the Virtual-Geography Hawkes Process (VG-Hawkes) be represented by $\mathbf{x}_{\mathsf{a}}^{(i)} \in \mathbb{R}^n$, $i = 1, \ldots, n+1$, and the position of publication k be denoted by $\mathbf{x}_{\mathsf{p}}^{(k)} \in \mathbb{R}^n$, $k = 1, \ldots, N$. These positions are learned based on prior knowledge and a collected dataset. The publication year of each paper is denoted by t_k, $k = 1, \ldots, N$, as described in Section 3.

Remark 1 (On the construction of virtual distances)**.** *In this Communication, the virtual geography is implemented by treating the author and publication positions* $\{\mathbf{x}_{\mathsf{a}}^{(i)}\}$ *and* $\{\mathbf{x}_{\mathsf{p}}^{(k)}\}$ *as* learnable latent variables *that are estimated from the observed event history under the VG-Hawkes likelihood. Accordingly, the "virtual distance" used in the triggering kernel is the Euclidean distance in this latent space, i.e.,* $\|\mathbf{x}_{\mathsf{p}} - \mathbf{x}_{\mathsf{p}}^{(k)}\|$*, which provides an interpretable notion of academic proximity induced by citation/co-authorship interactions through maximum-likelihood fitting. An explicit deterministic mapping from raw citation/co-authorship relations to virtual coordinates (e.g., via graph embedding or similarity-based construction) can also be integrated into the framework, but is beyond the scope of the present brief and will be investigated in future work.*

Using these positions and publication times, we define a kernel-based intensity model as follows:

$$\lambda_{\mathrm{vgh}}(t, \mathbf{x}_{\mathsf{p}}) := \mu_0 \mu(t, \mathbf{x}_{\mathsf{p}}) + A \sum_{k:t_k<t} g(t-t_k, \mathbf{x}_{\mathsf{p}} - \mathbf{x}_{\mathsf{p}}^{(k)}) \tag{4}$$

Here, $\mu(t, \mathbf{x}_{\mathsf{p}})$ is the background rate, and $g(t-t_k, \mathbf{x}_{\mathsf{p}} - \mathbf{x}_{\mathsf{p}}^{(k)})$ represents the excitation effect. The constants μ_0 and A are relaxation coefficients. To incorporate citation count m into the model, we define:

$$s(m) = \beta \exp(-\beta m), \quad \kappa(m) = A \exp(\alpha m),$$

and extend the model to a marked Hawkes process as follows:

$$\lambda_{\mathsf{vgh}} := s(m)\left[\mu_0\mu(t,\mathbf{x}_\mathsf{p}) + A\sum_{k:t_k<t}\kappa(m_k)g(t-t_k,\mathbf{x}_\mathsf{p}-\mathbf{x}_\mathsf{p}^{(k)},m_k)\right] \tag{5}$$

Here, the probability that event j is a background event (background probability), and the probability that event j was triggered by event i (triggering probability), are given as follows [19]:

$$\varphi_j = \frac{\mu(t_j,x_j)}{\lambda(t_j,x_j)}, \quad \rho_{ij} = \frac{\kappa(m_k)g(t_j-t_i,x_j-x_i)}{\lambda(t_j,x_j)} \tag{6}$$

Moreover, the sum of the background and triggering probabilities is always equal to 1:

$$\varphi_j + \sum_{i=1}^{j-1}\rho_{ij} = 1 \quad \text{for all } j \tag{7}$$

When an event occurs at a given point (t,x), it can be interpreted that φ_j background events are observed at (t,x), and for each $i=1,\ldots,j-1$, event i is responsible for generating ρ_{ij} direct offspring events at (t_j,x_j). Thus, event j can be decomposed into a background event component and descendant events triggered by past events. This formulation enables the use of non-parametric methods to estimate the background rate function $\mu(t,x)$ and the triggering kernel $g(t,x)$. Based on these definitions, the model captures the influence of past publication events over time and space on the current publication intensity.

3.5. Learning Algorithm

We describe the parameter estimation procedure for the model components: $\mu(t,\mathbf{x}_\mathsf{p})$, $g(t,\mathbf{x}_\mathsf{p})$, μ_0, and A, based on observed data. The estimation is performed using an Expectation-Maximization (EM) algorithm, inspired by the work of Zhuang et al. [19], Zhuang [20]. The procedure begins by initializing all parameters and then iteratively performs two steps:

- In the E-step, the algorithm estimates the probability that each observed event was generated by the background component or was triggered by a previous event [16].
- In the M-step, the model parameters are updated to maximize the expected complete-data log-likelihood [21].

This iterative process continues until convergence and allows for principled, data-driven learning of both the background and triggering dynamics underlying citation behavior in the proposed VG-Hawkes model. The proposed algorithm for estimating the virtual network Hawkes process is summarized in Algorithm 1.

Remark 2 (On parameter settings)**.** *The proposed VG-Hawkes model uses kernel bandwidths (e.g., h_t, h_m) and mark-related parameters (e.g., α, β), together with relaxation coefficients (e.g., μ_0, A). In our experiments, these parameters are initialized following standard practice in kernel-based Hawkes estimation and then selected using a simple data-driven tuning protocol on the available event history (details and the exact values used in the experiments are provided in the revised manuscript for reproducibility). A comprehensive sensitivity analysis over these hyperparameters is an important topic, but is outside the scope of this Communication and will be pursued in future work.*

Algorithm 1 EM-based Parameter Estimation for the Virtual Network Hawkes Process

Initialize model components: background intensity $\mu(t, \mathbf{x}_\mathsf{p})$, triggering kernel $g(\cdot)$, mark scaling $\kappa(m)$, and relaxation parameters μ_0, A.
Set tolerance $\varepsilon > 0$ and maximum iterations $I_{\max}$.
Set iteration counter $r \leftarrow 0$ and compute initial log-likelihood $\ell^{(0)}$.
while $r < I_{\max}$ **do**
    E-step: estimate origin of each event.
    For each event j, compute the intensity $\lambda_j := \lambda_{\mathsf{vgh}}(t_j, \mathbf{x}_\mathsf{p}^{(j)}; \boldsymbol{\theta})$.
    Background responsibility: $\varphi_j := \dfrac{\mu_0\, \mu(t_j, \mathbf{x}_\mathsf{p}^{(j)})}{\lambda_j}$.
    Triggering responsibility (for each $i < j$): $\rho_{ij} := \dfrac{A\, \kappa(m_i)\, g(t_j - t_i, \mathbf{x}_\mathsf{p}^{(j)} - \mathbf{x}_\mathsf{p}^{(i)}, m_i)}{\lambda_j}$.
    M-step: update model parameters.
    Update $\mu(\cdot)$ and $g(\cdot)$ using $\{\varphi_j\}$ and $\{\rho_{ij}\}$ (kernel-based updates following [19,20]).
    Update relaxation parameters: $A \leftarrow \dfrac{N - \sum_{j=1}^{N} \varphi_j}{\sum_{j=1}^{N} G_j/\lambda_j}$, $\quad \mu_0 \leftarrow \dfrac{N - A \cdot \sum_{j=1}^{N} G_j/\lambda_j}{\sum_{j=1}^{N} \mu_j/\lambda_j}$,
    where $\mu_j := \mu(t_j, \mathbf{x}_\mathsf{p}^{(j)})$ and $G_j := \sum_{i<j} \kappa(m_i)\, g(t_j - t_i, \mathbf{x}_\mathsf{p}^{(j)} - \mathbf{x}_\mathsf{p}^{(i)}, m_i)$.
    Convergence check: compute the log-likelihood $\ell^{(r+1)} := \log L(\mathcal{D}_N, \boldsymbol{\theta})$ (see Equation (13)).
    if $|\ell^{(r+1)} - \ell^{(r)}| < \varepsilon$ **then**
        break
    $r \leftarrow r + 1$

4. Results

4.1. Dataset

In this study, we evaluate the proposed model using a dataset collected from Google Scholar for each author, without restricting the source to any specific academic journal. During data collection, we focused on the research field of reinforcement learning. This decision was made because these fields are the authors' primary areas of expertise, facilitating an easier understanding of the research content and algorithms presented in the papers. We first identify a seed researcher in a chosen domain and then construct a co-authorship network by recursively tracing co-authorship relationships from that seed researcher. In this process, the nodes represent individual authors, and the edges represent co-authorship links. For each publication, we collect the following attributes as features: publication year, publisher, co-authors, and the number of citations. The dataset uses Mohammad Ghavamzadeh and Alessandro Lazaric as the cluster and derived authors. Tables 1–4 show the collected dataset.

We select two author-centric clusters in reinforcement learning as a controlled proof-of-concept setting: these clusters provide sufficiently long publication histories and dense citation/co-authorship interactions to highlight the benefit of incorporating virtual academic proximity into the Hawkes intensity while keeping the evaluation transparent. We acknowledge that such a selection may emphasize centralized or mentorship-driven structures; we therefore discuss this potential bias and the expected behavior in less centralized or more interdisciplinary settings in Section 4.

Table 1. Dataset for Mohammad Ghavamzadeh: Basic Information.

Title	Publisher (or Conference)	Time	Co-Author
Continuous-time Hierarchical Reinforcement Learning			
Hierarchically Optimal Average Reward Reinforcement Learning	ICML	8 July 2002	
Hierarchical Policy Gradient Algorithms	ICML	2003	
Bayesian Actor-Critic Algorithms	ICML	2007	
Hierarchical Average Reward Reinforcement Learning	JMLR	2007	
Regularized Policy Iteration	NeurIPS	2008	Amir-massoud Farahmand (Offspring)
Natural Actor-Critic Algorithms	Automatica	2009	
Analysis of a Classification-based Policy Iteration Algorithm	ICML	2010	Alessandro Lazaric (Offspring)
Bayesian Multi-task Reinforcement Learning	ICML	2010	Alessandro Lazaric (Offspring)
Finite-sample Analysis of LSTD	ICML	21 July 2010	Alessandro Lazaric (Offspring)
LSTD with Random Projections	NeurIPS	2010	Alessandro Lazaric (Offspring)
Finite-sample Analysis of Lasso-TD	ICML	2011	Alessandro Lazaric (Offspring)
Speedy Q-learning	NeurIPS	2011	
Classification-based Policy Iteration with a Critic	ICML	5 May 2011	
Approximate Modified Policy Iteration	AAAI	2012	
Finite-sample Analysis of Least-squares Policy Iteration	JMLR	1 October 2012	Alessandro Lazaric (Offspring)
Actor-Critic Algorithms for Risk-sensitive MDPs	NeurIPS	2013	Prashanth L.A. (Offspring)
Approximate Dynamic Programming Finally Performs Well in the Game of Tetris	NeurIPS	2013	
Algorithms for CVaR Optimization in MDPs	NeurIPS	2014	Yinlam Chow (Offspring)
High Confidence Policy Improvement	ICML	2015	Philip Thomas (Offspring)
Approximate Modified Policy Iteration and its Application to the Game of Tetris	JMLR	2015	
Maximum Entropy Semi-Supervised Inverse Reinforcement Learning	IJCAI	2015	Alessandro Lazaric (Offspring)
Personalized Ad Recommendation Systems for Life-time Value Optimization with Guarantees	IJCAI	2015	Philip Thomas (Offspring)
Policy Gradient for Coherent Risk Measures	NeurIPS	2015	Aviv Tamar (Co), Yinlam Chow (Offspring)
Proximal Gradient Temporal Difference Learning Algorithms	IJCAI	2016	
Regularized Policy Iteration with Nonparametric Function Spaces	JMLR	1 January 2016	
Analysis of Classification-based Policy Iteration Algorithms	JMLR	2016	Alessandro Lazaric (Offspring)
Safe Policy Improvement by Minimizing Robust Baseline Regret	NeurIPS	2016	Yinlam Chow (Offspring)
Bayesian Policy Gradient and Actor-Critic Algorithms	JMLR	2016	
Improved Learning Complexity in Combinatorial Pure Exploration Bandits	AISTATS	2 May 2016	Alessandro Lazaric (Offspring)
Sequential Decision-making with Coherent Risk	IEEE TAC	2017	Aviv Tamar (Co), Yinlam Chow (Offspring)
Risk-constrained Reinforcement Learning with Percentile Risk Criteria	JMLR	2018	Yinlam Chow (Offspring)
More Robust Doubly Robust Off-policy Evaluation	ICML	2018	Yinlam Chow (Offspring)
A Lyapunov-based Approach to Safe Reinforcement Learning	NeurIPS	2018	Yinlam Chow (Offspring)
Path Consistency Learning in Tsallis Entropy Regularized MDPs	ICML	3 July 2018	Yinlam Chow (Offspring)
Garbage In, Reward Out: Bootstrapping Exploration in Multi-armed Bandits	ICML	2019	
Tight Regret Bounds for Model-based Reinforcement Learning with Greedy Policies	NeurIPS	2019	Yinlam Chow (Offspring)
Risk-sensitive Generative Adversarial Imitation Learning	AISTATS	11 April 2019	Yinlam Chow (Offspring)
Optimizing Over a Restricted Policy Class in MDPs	AISTATS	11 April 2019	
Predictive Coding for Locally-linear Control	ICML	21 November 2020	Yinlam Chow (Offspring)

Table 2. Dataset for Mohammad Ghavamzadeh: Citation Information.

Title	Citations	Citations (1 Year)	Citations (3 Years)	Citations (5 Years)
Continuous-time Hierarchical Reinforcement Learning	55	1	4	14
Hierarchically Optimal Average Reward Reinforcement Learning	21	1	4	7
Hierarchical Policy Gradient Algorithms	79	1	7	17
Bayesian Actor-Critic Algorithms	75	1	8	23
Hierarchical Average Reward Reinforcement Learning	42	8	13	16
Regularized Policy Iteration	163	1	15	45
Natural Actor-Critic Algorithms	1084	5	43	127
Analysis of a Classification-based Policy Iteration Algorithm	88	2	19	41
Bayesian Multi-task Reinforcement Learning	143	1	14	35
Finite-sample Analysis of LSTD	87	3	25	40
LSTD with Random Projections	74	1	14	28
Finite-sample Analysis of Lasso-TD	57	4	23	33
Speedy Q-learning	215	4	11	27
Classification-based Policy Iteration with a Critic	30	9	17	23
Approximate Modified Policy Iteration	62	1	16	30
Finite-sample Analysis of Least-squares Policy Iteration	133	3	10	22
Actor-Critic Algorithms for Risk-sensitive MDPs	310	1	2	17
Approximate Dynamic Programming Finally Performs Well in the Game of Tetris	78	1	13	35
Algorithms for CVaR Optimization in MDPs	370	4	17	52
High Confidence Policy Improvement	220	2	28	76
Approximate Modified Policy Iteration and its Application to the Game of Tetris	154	2	11	37
Maximum Entropy Semi-Supervised Inverse Reinforcement Learning	38	2	10	22
Personalized Ad Recommendation Systems for Life-time Value Optimization with Guarantees	200	3	22	61
Policy Gradient for Coherent Risk Measures	142	2	8	23
Proximal Gradient Temporal Difference Learning Algorithms	32	1	8	16
Regularized Policy Iteration with Nonparametric Function Spaces	125	2	12	41
Analysis of Classification-based Policy Iteration Algorithms	52	1	10	23
Safe Policy Improvement by Minimizing Robust Baseline Regret	161	1	13	55
Bayesian Policy Gradient and Actor-Critic Algorithms	61	1	8	26
Improved Learning Complexity in Combinatorial Pure Exploration Bandits	45	2	14	30
Sequential Decision-making with Coherent Risk	80	6	23	57
Risk-constrained Reinforcement Learning with Percentile Risk Criteria	562	3	43	181
More Robust Doubly Robust Off-policy Evaluation	275	6	96	180
A Lyapunov-based Approach to Safe Reinforcement Learning	579	5	114	355
Path Consistency Learning in Tsallis Entropy Regularized MDPs	46	2	21	31
Garbage In, Reward Out: Bootstrapping Exploration in Multi-armed Bandits	78	6	39	68
Tight Regret Bounds for Model-based Reinforcement Learning with Greedy Policies	78	5	46	70
Risk-sensitive Generative Adversarial Imitation Learning	34	3	20	29
Optimizing Over a Restricted Policy Class in MDPs	12	3	9	11
Predictive Coding for Locally-linear Control	24	2	16	24

Table 3. Dataset: Alessandro Lazaric.

Title	Publisher (or Conference)	Date	Co-Author
Reinforcement learning in continuous action spaces through sequential monte carlo methods	NeurIPS	2007	
Transfer of samples in batch reinforcement learning	ICML	5 July 2008	
Finite-sample analysis of LSTD	ICML	21 June 2010	M. Ghavamzadeh (Cluster)
Analysis of a Classification-based Policy Iteration Algorithm	ICML	2010	M. Ghavamzadeh (Cluster)
Bayesian Multi-task Reinforcement Learning	ICML	2010	M. Ghavamzadeh (Cluster)
LSTD with random projections	NeurIPS	2010	M. Ghavamzadeh (Cluster)
Finite-sample analysis of Lasso-TD	ICML	2011	M. Ghavamzadeh (Cluster)
Transfer from multiple MDPs	NeurIPS	2011	
Risk-aversion in multi-armed bandits	NeurIPS	2012	
Finite-sample analysis of least-squares policy iteration	JMLR	1 October 2012	M. Ghavamzadeh (Cluster)
Sequential transfer in multi-armed bandit with finite set of models	NeurIPS	2013	
Exploiting easy data in online optimization	NeurIPS	2014	
Best-arm identification in linear bandits	NeurIPS	2014	
Sparse multi-task reinforcement learning	NeurIPS	2014	
Maximum entropy semi-supervised inverse reinforcement learning	IJCAI	25 July 2015	M. Ghavamzadeh (Cluster)
Direct policy iteration with demonstrations	IJCAI	25 July 2015	
Analysis of Classification-based Policy Iteration Algorithms	JMLR	2016	M. Ghavamzadeh (Cluster)
Improved Learning Complexity in Combinatorial Pure Exploration Bandits	AISTATS	2 May 2016	M. Ghavamzadeh (Cluster)
Regret minimization in MDPs with options without prior knowledge	NeurIPS	2017	
Near optimal exploration-exploitation in non-communicating Markov decision processes	NeurIPS	2018	
Fighting boredom in recommender systems with linear reinforcement learning	NeurIPS	2018	
Efficient bias-span-constrained exploration-exploitation in reinforcement learning	ICML	3 July 2018	
Improved regret bounds for Thompson sampling in linear quadratic control problems	ICML	3 July 2018	Marc Abeille (Offspring)
A structured prediction approach for generalization in cooperative multi-agent reinforcement learning	NeurIPS	2019	
Exploration bonus for regret minimization in discrete and continuous average reward MDPs	NeurIPS	2019	
Regret bounds for learning state representations in reinforcement learning	NeurIPS	2019	
Limiting extrapolation in linear approximate value iteration	NeurIPS	2019	E. Brunskill (Co-author)
Provably efficient reward-agnostic navigation with linear value iteration	NeurIPS	2020	E. Brunskill (Co-author)
Improved sample complexity for incremental autonomous exploration in MDPs	NeurIPS	2020	Jean Tarbouriech (Offspring)
Learning near optimal policies with low inherent Bellman error	ICML	21 November 2020	E. Brunskill (Co-author)
Frequentist regret bounds for randomized least-squares value iteration	AISTATS	3 June 2020	E. Brunskill (Co-author)
Active model estimation in Markov decision processes	CUAI	2020	Jean Tarbouriech (Offspring)
No-regret exploration in goal-oriented reinforcement learning	ICML	21 November 2020	Jean Tarbouriech (Offspring)
Efficient optimistic exploration in linear-quadratic regulators via Lagrangian relaxation	ICML	21 November 2020	Marc Abeille (Offspring)
Reinforcement learning with prototypical representations	ICML	1 July 2021	Denis Yarats (Offspring)
Mastering visual continuous control: Improved data-augmented reinforcement learning	arXiv	20 July 2021	Denis Yarats (Offspring)
A provably efficient sample collection strategy for reinforcement learning	NeurIPS	6 December 2021	Jean Tarbouriech (Offspring)
Stochastic shortest path: Minimax, parameter-free and towards horizon-free regret	NeurIPS	6 December 2021	Jean Tarbouriech (Offspring)
Reinforcement learning in linear MDPs: Constant regret and representation selection	NeurIPS	6 December 2021	Matteo Papini (Offspring)
Adaptive multi-goal exploration	AISTATS	3 May 2022	Jean Tarbouriech (Offspring)

Table 4. Dataset: Alessandro Lazaric: Citation Information.

Title	Citations	Co-Author	Citations (1 Year)	Citations (3 Years)	Citations (5 Years)
Reinforcement learning in continuous action spaces through sequential monte carlo methods	197		6	26	45
Transfer of samples in batch reinforcement learning	209		4	13	28
Finite-sample analysis of LSTD	87	M. Ghavamzadeh (Cluster)	3	25	40
Analysis of a Classification-based Policy Iteration Algorithm	88	M. Ghavamzadeh (Cluster)	2	19	41
Bayesian Multi-task Reinforcement Learning	143	M. Ghavamzadeh (Cluster)	1	14	35
LSTD with random projections	73	M. Ghavamzadeh (Cluster)	6	17	30
Finite-sample analysis of Lasso-TD	57	M. Ghavamzadeh (Cluster)	4	23	33
Transfer from multiple MDPs	62		4	13	22
Risk-aversion in multi-armed bandits	186		5	17	27
Finite-sample analysis of least-squares policy iteration	133	M. Ghavamzadeh (Cluster)	3	10	22
Sequential transfer in multi-armed bandit with finite set of models	121		1	18	32
Exploiting easy data in online optimization	62		1	19	32
Best-arm identification in linear bandits	212		2	21	45
Sparse multi-task reinforcement learning	83		1	10	22
Maximum entropy semi-supervised inverse reinforcement learning	38	M. Ghavamzadeh (Cluster)	2	10	22
Direct policy iteration with demonstrations	44		2	12	26
Analysis of Classification-based Policy Iteration Algorithms	52	M. Ghavamzadeh (Cluster)	1	10	23
Improved Learning Complexity in Combinatorial Pure Exploration Bandits	45	M. Ghavamzadeh (Cluster)	2	14	30
Regret minimization in MDPs with options without prior knowledge	30		2	9	16
Near optimal exploration-exploitation in non-communicating Markov decision processes	49		1	12	27
Fighting boredom in recommender systems with linear reinforcement learning	53		3	24	50
Efficient bias-span-constrained exploration-exploitation in reinforcement learning	115		1	18	61
Improved regret bounds for Thompson sampling in linear quadratic control problems	106	Marc Abeille (Offspring)	1	45	87
A structured prediction approach for generalization in cooperative multi-agent reinforcement learning	27		1	7	21
Exploration bonus for regret minimization in discrete and continuous average reward MDPs	44		1	24	36
Regret bounds for learning state representations in reinforcement learning	13		2	9	13
Limiting extrapolation in linear approximate value iteration	39	E. Brunskill (Co-author)	10	23	38

Table 4. *Cont.*

Title	Citations	Co-Author	Citations (1 Year)	Citations (3 Years)	Citations (5 Years)
Provably efficient reward-agnostic navigation with linear value iteration	61	E. Brunskill (Co-author)	1	16	47
Improved sample complexity for incremental autonomous exploration in MDPs	17	Jean Tarbouriech (Offspring)	1	10	17
Learning near optimal policies with low inherent Bellman error	234	E. Brunskill (Co-author)	1	67	198
Frequentist regret bounds for randomized least-squares value iteration	146	E. Brunskill (Co-author)	2	49	126
Active model estimation in Markov decision processes	29	Jean Tarbouriech (Offspring)	1	5	23
No-regret exploration in goal-oriented reinforcement learning	43	Jean Tarbouriech (Offspring)	3	29	42
Efficient optimistic exploration in linear-quadratic regulators via Lagrangian relaxation	40	Marc Abeille (Offspring)	1	28	40
Reinforcement learning with prototypical representations	210	Denis Yarats (Offspring)	1	90	209
Mastering visual continuous control: Improved data-augmented reinforcement learning	259	Denis Yarats (Offspring)	1	78	258
A provably efficient sample collection strategy for reinforcement learning	18	Jean Tarbouriech (Offspring)	1	6	18
Stochastic shortest path: Minimax, parameter-free and towards horizon-free regret	34	Jean Tarbouriech (Offspring)	3	28	
Reinforcement learning in linear MDPs: Constant regret and representation selection	19	Matteo Papini (Offspring)	6	19	
Adaptive multi-goal exploration	4	Jean Tarbouriech (Offspring)	1	4	

4.2. Validations

The red lines in Figure 2b,c show the convergence of the log-likelihood for Mohammad Ghavamzadeh and Alessandro Lazaric with two methods. The horizontal axis represents the number of iterations in the learning algorithm, and the vertical axis shows the log-likelihood $\log L(\mathcal{D}_N, \boldsymbol{\theta})$. On the other hand, the red lines in Figures 3a and 4a depict the correspondence between the intensity function. The horizontal axis indicates normalized time (i.e., normalized paper publication date from 0 to 1), and the upper vertical axis shows the intensity function.

Next, we present the results obtained using the proposed Virtual Network Hawkes Process model. This model extends the temporal Hawkes process by incorporating the academic relationships and virtual distances between authors. The blue lines in Figure 2 show the convergence behavior of the log-likelihood. Figure 2a displays the overall dataset (including both Mohammad Ghavamzadeh and Alessandro Lazaric), while Figure 2b,c show individual log-likelihoods: blue for Marked VG Hawkes and red for Marked Temporal Hawkes. The horizontal axis represents the number of learning iterations, and the vertical axis shows the log-likelihood $\log L(\mathcal{D}_N, \boldsymbol{\theta})$, indicating the convergence of the model during training. The blue lines in Figures 3 and 4 illustrate the correspondence between the intensity function and citation counts for Mohammad Ghavamzadeh and Alessandro Lazaric, respectively. In each figure, the horizontal axis represents normalized time (i.e., normalized publication timing), the upper vertical axis indicates the intensity function, and the lower vertical axis shows the normalized number of citations for each paper.

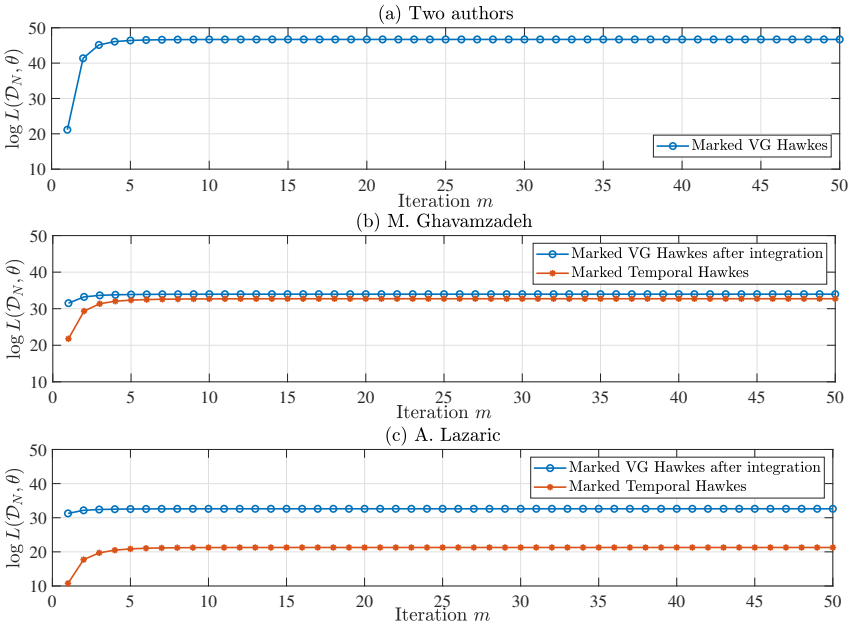


Figure 2. Log-likelihood convergence.

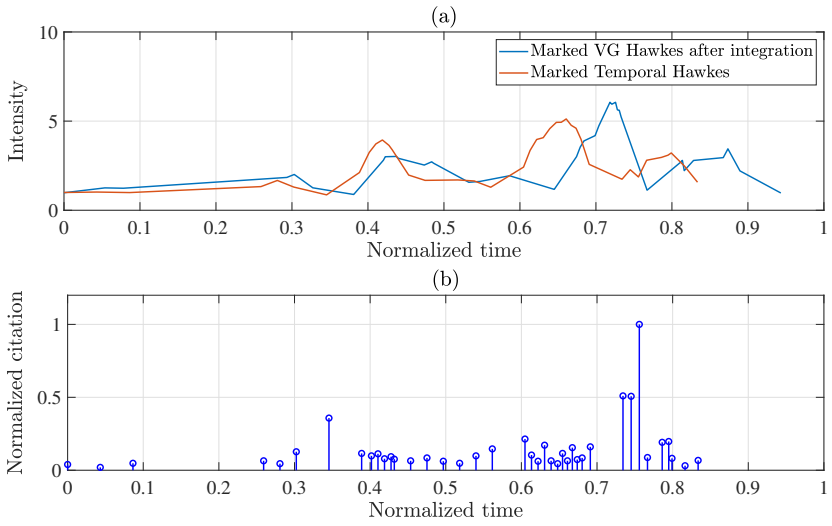


Figure 3. Mohammad Ghavamzadeh: (**a**) Intensity function; (**b**) Citation correspondence.

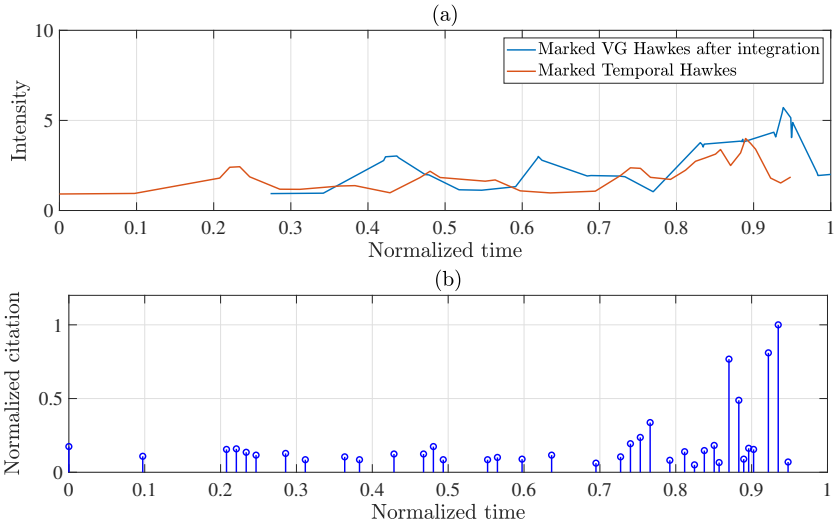


Figure 4. Alessandro Lazaric: (**a**) Intensity function; (**b**) citation correspondence.

The results confirm the effectiveness of the proposed Virtual Network Hawkes Process model in capturing citation dynamics. As shown in Figure 2, the VG-Hawkes model consistently achieves higher log-likelihood values compared to the conventional temporal Hawkes model across all datasets, indicating improved fit and convergence during training. This suggests that incorporating virtual academic distances and author relationships enables the model to better represent the underlying generative structure of citation events. In Figures 3 and 4, the intensity functions generated by the VG-Hawkes model more closely follow the actual citation trends over normalized time, particularly in capturing periods of heightened academic activity. This improvement is attributed to the model's ability to account for mutual excitation effects between authors based on their latent academic proximity, rather than relying solely on temporal recurrence. By embedding structural information from citation and co-authorship networks, the VG-Hawkes framework provides a more accurate and interpretable representation of how scholarly influence propagates over time.

We use log-likelihood as the primary quantitative evaluation criterion because VG-Hawkes is a probabilistic generative point-process model learned via maximum likelihood. Therefore, log-likelihood directly measures goodness-of-fit of the observed event history under the modeled intensity and provides a standard, objective basis for comparing temporal Hawkes and VG-Hawkes under the same inference procedure. Complementary predictive evaluation is also possible by held-out forecasting (e.g., next-event time prediction or citation count prediction over a future window) and reporting prediction errors; we leave such broader predictive benchmarking to future work.

5. Discussion and Limitations

This study presents the Virtual Network Hawkes Process as a novel approach to modeling citation dynamics by integrating latent academic relationships into a spatiotemporal point process framework. The empirical results demonstrate promising performance in both likelihood convergence and citation-intensity correspondence, indicating the model's potential to capture the structural and temporal complexities of academic influence.

While the proposed method shows clear advantages over conventional temporal Hawkes processes, several limitations remain. First, this work does not yet cover the full generality of the proposed framework, such as explicitly modeling cross-disciplinary interactions or dynamically evolving author embeddings. Second, the current implementation relies on predefined datasets and background functions, which may be extended in future work through more adaptive or learning-based formulations.

Nonetheless, this paper represents an important first step toward structure-aware citation modeling. The approach is conceptually general, interpretable, and computationally tractable offering a flexible basis for further development. As a brief report, this work aims not to exhaustively address every technical dimension, but to provide a clear direction for future research. The framework and results presented here are well-positioned to spark deeper exploration into the dynamics of academic networks and influence propagation.

The current evaluation is based on two reinforcement-learning author-centric clusters, which may over-represent hierarchical or mentorship-driven structures and thus bias citation dynamics toward stronger clustering. For less centralized authors, interdisciplinary researchers, or fields with weaker mentorship patterns, we expect the excitation to be more diffuse and the learned virtual distances to reflect weaker, broader interaction channels. Systematic validation across more authors and research domains is an important direction for future work.

The empirical evaluation in this Communication is intended as an initial feasibility study of the proposed VG-Hawkes framework, using two representative author-centric

clusters within a single field to provide a controlled setting and to highlight the effect of incorporating virtual academic proximity. Accordingly, we focus on likelihood-based goodness-of-fit and intensity–citation correspondence as primary validation signals. We acknowledge that a more comprehensive empirical study would further strengthen the evidence, including (i) systematic sensitivity analysis with respect to kernel and mark-related parameters, (ii) predictive evaluation on held-out citation counts using additional metrics (e.g., count prediction error), and (iii) broader coverage across more authors and research domains. These extensions are important and will be pursued in future work.

We do not report confidence intervals for the estimated intensity functions in this Communication. Rigorous uncertainty quantification for point-process intensities typically requires additional inference procedures (e.g., bootstrap-based variability assessment or Bayesian posterior inference), and we leave this extension as future work.

6. Conclusions and Future Work

This Communication proposed a Virtual-Geography Hawkes process (VG-Hawkes) to model citation dynamics by augmenting a temporal Hawkes process with virtual academic proximity extracted from citation and co-authorship relations. By embedding structural information into the excitation mechanism, the resulting intensity captures both temporal recurrence and network-driven mutual excitation. Empirical results on two author-centric clusters suggest improved likelihood convergence and a closer correspondence between learned intensities and observed citation patterns compared with a purely temporal Hawkes baseline.

Several extensions are natural and will be pursued in future work. First, broader validation across additional authors, disciplines, and less centralized collaboration structures is necessary to assess generalization. Second, complementary predictive evaluations (e.g., held-out forecasting of citation counts or next-event prediction) and systematic sensitivity analyses for kernel/mark-related hyperparameters will strengthen the empirical evidence. Third, uncertainty quantification for intensity estimates (e.g., via bootstrap-based variability assessment or Bayesian inference) is an important direction beyond the scope of this brief report.

Author Contributions: Conceptualization, X.S. and X.L.; methodology, X.S., X.L. and B.G.; software, B.G. and X.S.; validation, B.G. and X.S.; formal analysis, B.G. and X.S.; investigation, X.S. and B.G.; resources, X.L., A.M. and K.K.; data curation, X.S.; writing—original draft preparation, B.G. and X.S.; writing—review and editing, B.G., X.L., A.M., K.K. and X.S.; visualization, X.S.; supervision, X.S.; project administration, X.S., X.L., A.M. and K.K.; funding acquisition, X.L., A.M. and K.K. All authors have read and agreed to the published version of the manuscript.

Funding: This paper is based on results obtained from the project, JPNP25006, commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

Data Availability Statement: All the information is available in google scholar and it is summarized in the Tables.

Conflicts of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Seki, Y.; Matsuo, Y. Citation Count Prediction Using Citation Information. In Proceedings of the 25th Annual Conference of the Japanese Society for Artificial Intelligence, Morioka, Japan, 1–3 June 2011.
2. Hirako, J.; Sasano, R.; Takeda, K. CiMaTe: Citation Count Prediction Effectively Leveraging the Main Text. *arXiv* **2024**, arXiv:2410.04404. Available online: https://arxiv.org/html/2410.04404v1 (accessed on 20 January 2026).

3. Mori, J.; Hara, T.; Sakaki, T.; Kajikawa, Y.; Sakata, I. Identifying Central Researchers in Emerging Research Areas via Co-authorship Network Analysis of Large-Scale Academic Data. In Proceedings of the 29th Annual Conference of the Japanese Society for Artificial Intelligence, Hakodate, Japan, 30 May–2 June 2015.
4. Garfield, E. Citation Analysis as a Tool in Journal Evaluation. *Science* **1972**, *178*, 471–479. [CrossRef] [PubMed]
5. Hirsch, J.E. An Index to Quantify an Individual's Scientific Research Output. *Proc. Natl. Acad. Sci. USA* **2005**, *102*, 16569–16572. [CrossRef] [PubMed]
6. Egghe, L. An Improvement of the h-Index: The g-Index. *Scientometrics* **2006**, *69*, 131–152. [CrossRef]
7. Wang, L.; Du, W.; Chen, Z. Multi-Feature-Enhanced Academic Paper Recommendation Model with Knowledge Graph. *Appl. Sci.* **2024**, *14*, 5022. [CrossRef]
8. Wang, D.; Song, C.; Barabási, A.L. Quantifying Long-Term Scientific Impact. *Science* **2013**, *342*, 127–132. [CrossRef] [PubMed]
9. Sinatra, R.; Wang, D.; Deville, P.; Song, C.; Barabási, A.L. Quantifying the Evolution of Individual Scientific Impact. *Science* **2016**, *354*, aaf5239. [CrossRef] [PubMed]
10. Hawkes, A.G. Point Spectra of Some Mutually Exciting Point Processes. *J. R. Stat. Soc. Ser. Stat. Methodol.* **1971**, *33*, 438–443. [CrossRef]
11. Yu, Q.; Long, C.; Lv, Y.; Shao, H.; He, P.; Duan, Z. Predicting Co-Author Relationship in Medical Co-Authorship Networks. *PLoS ONE* **2014**, *9*, e101214. [CrossRef] [PubMed]
12. Zhao, Z.; Liu, W.; Qian, Y.; Nie, L.; Yin, Y.; Zhang, Y. Identifying advisor-advisee relationships from co-author networks via a novel deep model. *Inf. Sci.* **2018**, *466*, 258–269. [CrossRef]
13. Asatani, K.; Mori, J.; Ochi, M.; Sakata, I. Detecting trends in academic research from a citation network using network representation learning. *PLoS ONE* **2018**, *13*, e0197260. [CrossRef] [PubMed]
14. Wang, W.; Tang, T.; Xia, F.; Gong, Z.; Chen, Z.; Liu, H. Collaborative Filtering With Network Representation Learning for Citation Recommendation. *IEEE Trans. Big Data* **2022**, *8*, 1233–1246. [CrossRef]
15. Wang, S.; Gai, K.; Yu, J.; Zhang, Z.; Zhu, L. PraVFed: Practical Heterogeneous Vertical Federated Learning via Representation Learning. *IEEE Trans. Inf. Forensics Secur.* **2025**, *20*, 2693–2705. [CrossRef]
16. Zhuang, J.; Mateu, J. A Semiparametric Spatiotemporal Hawkes-Type Point Process Model with Periodic Background for Crime Data. *J. R. Stat. Soc. Ser. A* **2019**, *182*, 919–942. [CrossRef]
17. Bernabeu, A.; Zhuang, J.; Mateu, J. Spatio-Temporal Hawkes Point Processes: A Review. *J. Agric. Biol. Environ. Stat.* **2025**, *30*, 89–119. [CrossRef]
18. Laub, P.J.; Lee, Y.; Pollett, P.K.; Taimre, T. Hawkes Models and Their Applications. *Annu. Rev. Stat. Its Appl.* **2025**, *12*, 233–258. [CrossRef]
19. Zhuang, J.; Ogata, Y.; Vere-Jones, D. Analyzing Earthquake Clustering Features by Using Stochastic Reconstruction. *J. Geophys. Res.* **2004**, *109*, B05301. [CrossRef]
20. Zhuang, J. Second-Order Residual Analysis of Spatiotemporal Point Processes and Applications in Model Evaluation. *J. R. Stat. Soc. Ser. Stat. Methodol.* **2006**, *68*, 635–653. [CrossRef]
21. Lange, K. *MM Optimization Algorithms*; SIAM: Philadelphia, PA, USA, 2016.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.