International Journal of Disaster Risk Reduction 130 (2025) 105812

Contents lists available at ScienceDirect

International Journal of Disaster Risk Reduction

journal homepage: www.elsevier.com/locate/ijdrr

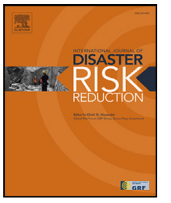

Modelling social mobility disruptions and recovery during disasters: A mobile phone data approach

Minh Kieu [a,*], Alexis Comber [b], Thanh Bui Quang [c], Nick Malleson [b,d]

[a] Department of Civil and Environmental Engineering, University of Auckland, Auckland 1010, New Zealand [1]
[b] School of Geography, University of Leeds, LS2 9JT, UK [2]
[c] Faculty of Geography, VNU University of Science, Hanoi, Viet Nam
[d] Alan Turing Institute, NW1 2DB, UK [3]

ARTICLE INFO

Keywords:
Disasters
Human mobility
Hawkes point process
Mobile phone data

ABSTRACT

Disasters fundamentally alter human mobility patterns, yet traditional modelling approaches fail to capture the complex temporal dynamics and interdependencies that characterise these disruptions. This paper introduces the Disaster-aware Hawkes process, a novel temporal point process framework specifically designed to model human mobility during and after disasters. Our approach extends standard self-exciting Hawkes processes through five key innovations: a regime-switching baseline intensity function, category-specific excitation parameters, heterogeneous recovery rates across location categories, and post-disaster bounce-back parameters. We apply our model to a comprehensive mobile device dataset from Auckland, New Zealand during Cyclone Gabrielle in February 2023, comprising 5.85 million mobility records from 111,539 devices. The Disaster-aware Hawkes process achieves a 25.68% overall improvement in prediction accuracy compared to standard approaches, with an 80.00% improvement during the disaster period. Beyond enhanced prediction, our model enables novel analyses of cascade effects, revealing how disruptions propagate through mobility networks and identifying critical dependencies invisible to traditional methods. We demonstrate that optimised temporal distribution of recovery interventions can improve system outcomes by 45.8% compared to conventional simultaneous deployment strategies. These findings provide valuable insights for disaster preparedness, response, and recovery planning, while the methodological framework offers a powerful new approach for analysing the complex dynamics of human mobility under disruption.

1. Introduction

Disasters fundamentally alter human mobility patterns, disrupting not only where people go but also why and with whom they interact at specific locations. These altered patterns create temporary but critical social–spatial structures that influence both individual decision-making and system-wide resilience. However, current emergency response and community resilience improvement strategies typically rely on static, pre-disaster mobility assumptions that fail to capture how people behave during crises. This disconnect between planning assumptions and actual behaviour can severely hinder effective disaster management

* Correspondence to: Transportation Research Centre, University of Auckland, New Zealand.
E-mail address: minh.kieu@auckland.ac.nz (M. Kieu).

[1] http://www.cee.auckland.ac.nz/.
[2] https://environment.leeds.ac.uk/geography.
[3] https://www.turing.ac.uk/.

https://doi.org/10.1016/j.ijdrr.2025.105812
Received 22 April 2025; Received in revised form 13 August 2025; Accepted 8 September 2025
Available online 1 October 2025

and recovery. Substantial research has examined disaster mobility patterns using mobile phone data, demonstrating their critical role in understanding disaster response. Studies have analysed evacuation patterns during Hurricane Matthew and Irma [1,2], estimated real-time human movement during the Great East Japan Earthquake [3], and developed predictive models of post-disaster mobility [4]. Research on Hurricane Harvey revealed that disadvantaged populations were less likely to evacuate than wealthier residents [5], whilst analysis of Hurricane Maria found that pre-disaster social connectivity was crucial for quicker recovery [6]. For a comprehensive review, readers may refer to Yabe et al. [7].

Existing disaster mobility research has indeed examined temporal dynamics extensively, using time-series analysis, regression models, machine learning approaches, and longitudinal trajectory analysis [4,6,8]. These studies have successfully quantified how mobility patterns evolve over disaster timelines, identified recovery trajectories, and measured resilience indicators. However, these temporal analysis methods approach the modelling of time-dependent behaviour through fundamentally different mechanisms than those offered by point process frameworks. Most existing temporal approaches model mobility through aggregated patterns to examine how total flows, average distances, or visit counts change over discrete time intervals. Whilst valuable for understanding macroscopic trends, this aggregation-based perspective cannot capture the event-driven dynamics that characterise human mobility during disasters. When individuals make mobility decisions, they respond not only to external conditions but also to the immediate history of recent mobility events in their social and spatial networks. Traditional temporal models lack the mathematical structure to represent how individual mobility events create cascading influences that propagate through the system in real-time.

Point process models, particularly Hawkes processes, offer a fundamentally different approach to temporal dynamics by modelling the conditional intensity of events, or how the probability of a mobility event occurring at any given moment depends on the complete history of past events. This event-driven perspective enables several capabilities that distinguish it from aggregation-based temporal analysis: (1) capturing temporal clustering where mobility events create "bursts" of subsequent activity, (2) quantifying memory effects that show how the influence of past events decays over time, (3) modelling mutual excitation where events in one location category trigger cascading events in others, and (4) providing mathematical foundations for predicting the exact timing of future events rather than just aggregate trends. However, standard Hawkes processes face a fundamental theoretical challenge when applied to disaster contexts. The core mathematical assumption of self-excitation — where events increase the probability of future events — directly contradicts the suppressive effects that characterise disaster impacts on mobility. This incompatibility explains why Hawkes processes, despite their success in modelling temporal dependencies in finance, seismology, and social media, have seen limited application in disaster mobility analysis. Furthermore, disaster mobility exhibits several distinctive characteristics absent from traditional Hawkes process applications: (1) heterogeneous suppression across different location categories, with essential services maintaining higher activity levels than discretionary venues; (2) category-specific recovery rates that depend on both infrastructure resilience and demand elasticity; (3) post-disaster "bounce-back" effects where visitation temporarily exceeds pre-disaster levels due to delayed demand fulfilment; and (4) complex cascade propagation where disruptions to critical categories trigger secondary impacts throughout the mobility network.

To address these methodological gaps, this paper introduces the Disaster-aware Hawkes process, a novel temporal point process framework that extends the standard Hawkes process to capture the complex dynamics of disaster-affected mobility. Our approach makes several key theoretical contributions that distinguish it from existing disaster mobility analysis methods:

- Regime-Switching Dynamics: We introduce smoothly-varying regime transition functions that capture the temporal evolution from normal functioning through acute disruption to recovery, moving beyond static analytical approaches to model dynamic system behaviour.
- Category-Specific Parameterisation: We develop heterogeneous impact, recovery, and bounce-back parameters that account for differential vulnerability across location categories, moving beyond aggregate analysis to capture the granular dynamics that drive system-wide behaviour.
- Causal Cascade Identification: We develop mathematical methods to disentangle direct disaster impacts from indirect cascade effects, enabling identification of how disruptions propagate through categorical networks via causal rather than merely correlational relationships.
- Temporal Intervention Analysis: We demonstrate how the model enables analysis of intervention timing strategies, revealing that carefully sequenced resource deployment can achieve 45.8% better outcomes than conventional simultaneous intervention approaches.

We apply our model to a comprehensive mobile device dataset from Auckland, New Zealand during Cyclone Gabrielle in February 2023. Auckland represents an exemplary case study for disaster mobility analysis, combining urban complexity with socio-economic diversity across its 1.7 million residents. The metropolitan area exhibits pronounced spatial heterogeneity, ranging from high-density commercial centres to low-density suburban peripheries, with socio-economic stratification clearly demarcated by the New Zealand Deprivation Index (NZDep) [9]. Cyclone Gabrielle provided an exceptional natural experiment in disaster response, representing New Zealand's most severe weather event in recent decades. The cyclone generated sustained winds exceeding 130 km/h across Auckland, caused widespread flooding that compromised major transportation arteries, and triggered power outages affecting over 225,000 properties nationwide [10]. The cyclone's trajectory across the upper North Island ensured Auckland experienced the full spectrum of severe weather impacts, from wind damage to flood-related infrastructure failure, creating diverse disruption patterns across the metropolitan area.

This convergence of urban complexity, socio-economic diversity, and comprehensive disaster impact creates an ideal empirical context for developing disaster-aware mobility models. Our mobile device dataset captures this natural experiment through

5,854,149 mobility records from 111,539 unique devices, providing unprecedented visibility into mobility patterns before, during, and after the disaster event. The specific characteristics of both Auckland's urban structure and Cyclone Gabrielle's impact profile enable robust model development whilst facilitating assessment of the framework's applicability to similar metropolitan areas facing severe weather disruptions. Beyond improved prediction accuracy, our approach enables novel analyses of cascade effects in mobility networks. We quantify how disruptions propagate through the network by constructing impact cascade matrices and visualisations that reveal critical dependencies invisible to traditional analysis methods. It is also important to distinguish our approach from the extensive literature on mobility prediction using machine learning and deep learning methods. Whilst data-driven approaches often achieve superior predictive accuracy, they operate as "black boxes" that provide limited insight into the underlying mechanisms driving observed patterns. Our Disaster-aware Hawkes process operates within the mechanistic modelling paradigm, where the primary objective is providing interpretable mathematical representations of causal processes rather than optimising predictive performance.

The remainder of this paper is organised as follows: Section 2 explores the characteristics of our dataset, establishing key patterns in temporal, spatial, and categorical dimensions. Section 3 details our methodology, including the mathematical formulation of the Disaster-aware Hawkes process and its implementation. Section 4 presents our findings, demonstrating the model's performance and the insights it provides for disaster management. Finally, we conclude with a discussion of limitations and directions for future research.

2. Dataset exploration

Understanding the characteristics, strengths, and limitations of our mobility dataset is essential for properly contextualising our analytical approach and subsequent findings. This section examines the temporal, spatial, and categorical dimensions of our Auckland mobility dataset, with particular attention to patterns relevant to Cyclone Gabrielle in February 2023.

2.1. Dataset characteristics

Our analysis leverages a comprehensive mobile device visitation dataset covering the Auckland metropolitan area in New Zealand, encompassing the period before, during, and after Cyclone Gabrielle (February 2023). The mobility data are individual visits to points of interest (POIs). The dataset contains 5,854,149 individual visits from 111,539 unique anonymised devices. Each visit record contains the following key elements:

- Device identifier (anonymised)
- Timestamp of the visit
- POI category
- Dwell time (duration of the visit)
- Geographic coordinates (latitude and longitude)

The dataset comes from mobile applications usage in Android and iOS, which the user has consented for location data collection, such as weather, ridesharing, or delivery applications. The location information (such as name and category of the visited location) was gathered from publicly available data on social media and provided by the data provider Predik Data-driven. Retail places such as bars, hotels, and stores are likely to appear in the dataset, and offices will be less likely to because: (1) for privacy reasons some places, mainly related to religion, health and government, are excluded from the dataset, and (2) because the applications are more likely to be used during non-occupational activities.

We also focus on one of the most major disasters in New Zealand recently: The Cyclone Gabrielle in 2023. The dataset spans from 1 February 2023 to 27 February 2023, with the cyclone period identified as 12–14 February 2023. This temporal scope allows us to establish a complete lifecycle from baseline mobility patterns, to capturing disruption during the disaster, and observing recovery trajectories after the event.

2.2. Temporal patterns

Fig. 1 illustrates the weekly and diurnal patterns present in our data.

Fig. 1 illustrates the weekly and diurnal patterns present in our data, which exhibit temporal signatures distinct from traditional mobility datasets due to the specific nature of our data collection methodology. Unlike comprehensive mobility tracking systems that capture all movement, our dataset derives from voluntary location sharing through leisure-oriented mobile applications (weather, ridesharing, delivery apps), creating a systematic bias towards discretionary, non-occupational activities during the mid-day period and weekends.

This focus on leisure mobility makes our dataset especially valuable for understanding how disasters affect social and commercial activities, though it necessarily provides limited insight into occupational mobility patterns.

2.3. Spatial distribution

Fig. 2 shows the distribution of estimated 'home locations' for the individuals in our dataset. Here the home location is assumed as the location where we can find the device most during night time period (from 10 PM to 4 AM), and where it was recorded

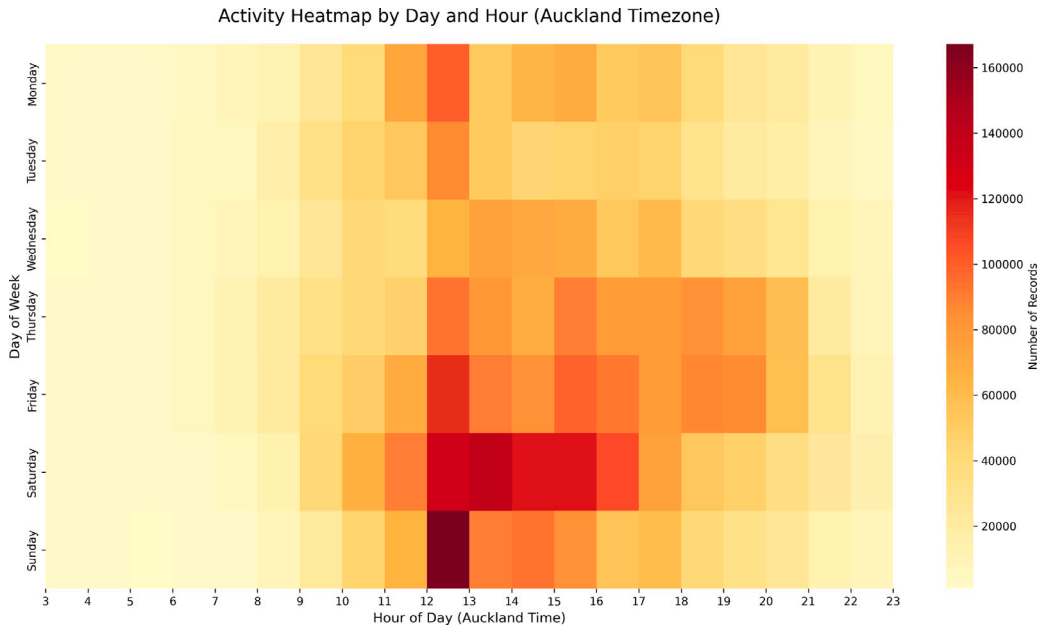

Fig. 1. Activity heatmap by day and hour.

at least twice. We successfully identified the home location for 63,556 devices (57.0%). The background of Fig. 2 is a Statistical Area 2's NZDep index of deprivation [9]. This spatial distribution reveals substantial heterogeneity in both device coverage and household income across the urban landscape.

The dataset provides strong coverage across Auckland's urban core, with particularly dense representation in central business districts and major commercial zones. Coverage becomes sparser in peripheral regions, which is a common limitation in mobile device datasets. The spatial visualisation reveals clusters of high deprivation (deciles 6–10) primarily in South and West Auckland, while lower deprivation areas (deciles 1–5) dominate the North Shore and Eastern suburbs.

Fig. 3(a) presents essential service visit rates across cyclone periods, revealing minimal differences between low and high deprivation groups before the disaster and after recovery. However, during the cyclone period, high deprivation areas exhibited significantly higher essential service visit rates, suggesting potentially greater vulnerability or different coping strategies among disadvantaged populations during the acute disaster phase.

Fig. 3(b) examines mean travel distances, revealing more consistent socioeconomic disparities across all periods. High deprivation populations demonstrate significantly longer travel distances before ($p < 0.001$), during ($p < 0.001$), and after ($p < 0.001$) the cyclone. This pattern likely reflects differential access to proximate services, with disadvantaged populations requiring longer journeys to access equivalent amenities—a disparity that persists even during disaster conditions.

2.4. Disaster impact and recovery

Fig. 4 provides direct evidence of Cyclone Gabrielle's impact on mobility patterns, showing daily visit counts for the top 10 most visited location categories throughout February 2023. The graph reveals a sharp decline in visits across all categories beginning on February 12, with particularly pronounced drops in Food & Dining (from 27,278 visits on February 11 to approximately 8,000 during the cyclone) and General Retail.

Several notable patterns emerge from this temporal view. First, the impact magnitude varies substantially across categories, with essential services like Health & Medical experiencing less severe declines than discretionary categories like Clothing & Fashion. Second, we observe different recovery trajectories, with some categories returning to pre-disaster levels within days while others show more prolonged suppression. Third, the post-recovery period (February 18–25) exhibits compensation behaviour in some categories, with visit counts temporarily exceeding pre-disaster levels. This "bounce back" pattern is likely due to delayed fulfilment of needs postponed during the disaster, compensatory behaviour of human from being mobility restricted, or resource replenishment of businesses.

Quantifying Disaster Impact To systematically analyse the differential impact of the cyclone across categories, we define a disaster impact magnitude $I(c)$ for each POI category c as the proportional reduction in visitation during the disaster relative to the pre-disaster baseline. Formally, we calculate:

$$I(c) = 1 - \frac{\bar{N}_{\text{during}}(c)}{\bar{N}_{\text{pre}}(c)} \tag{1}$$

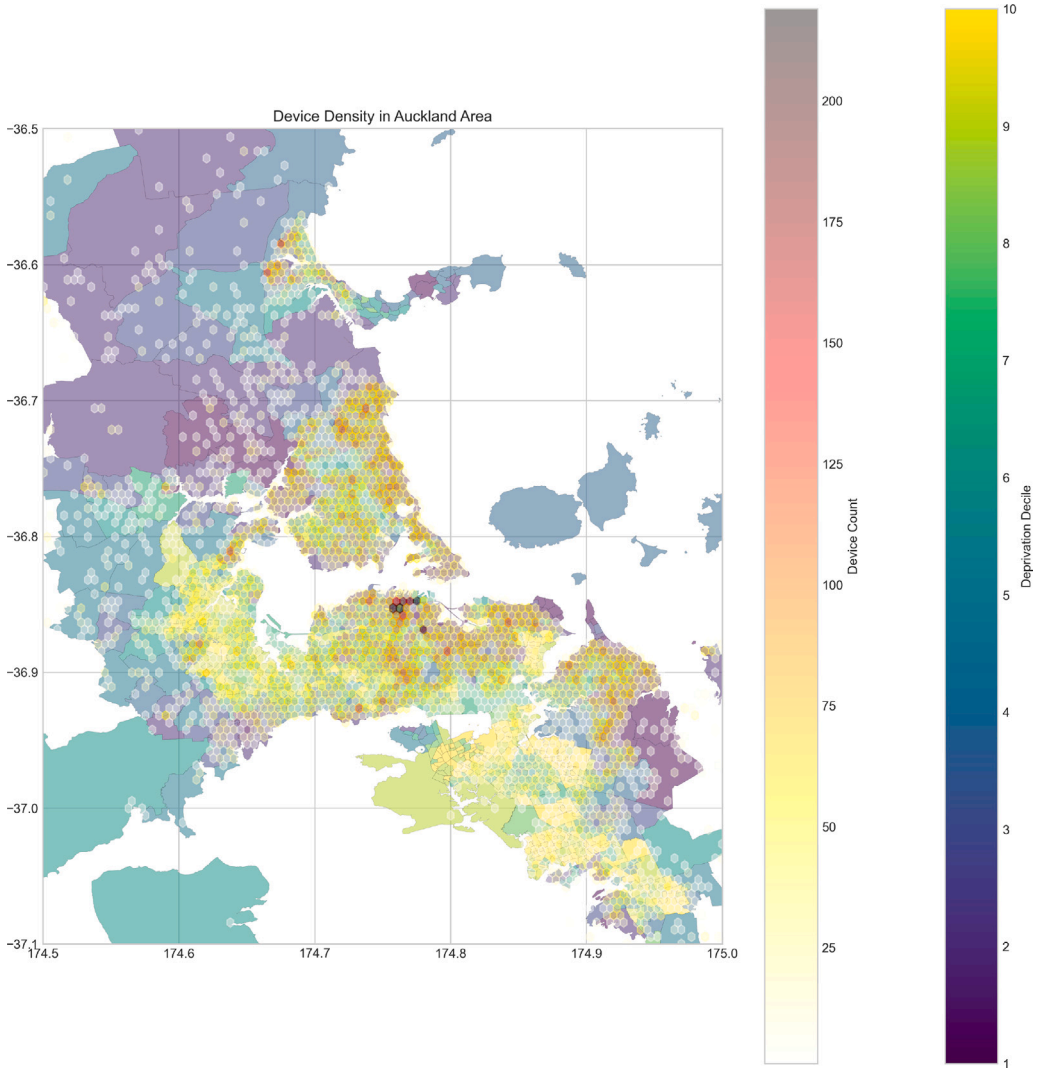

Fig. 2. Distributions of estimated home location.

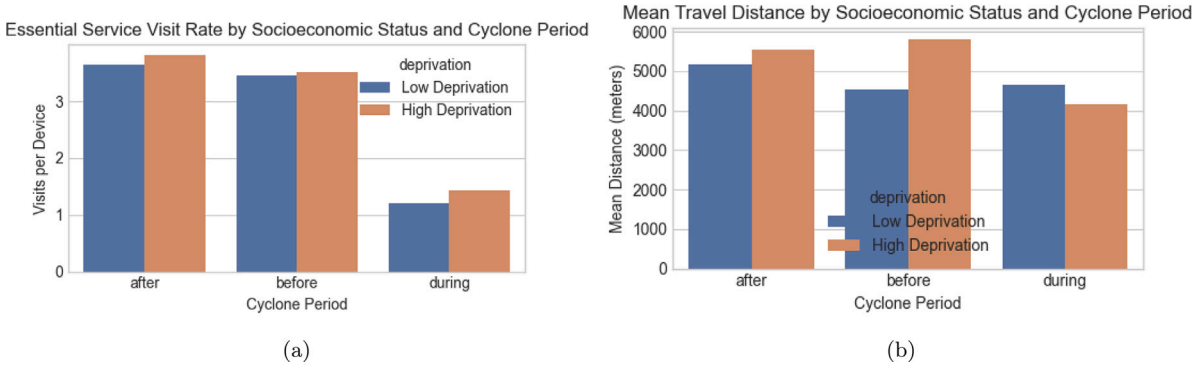

Fig. 3. Differences between Low and High deprivation mobility patterns.

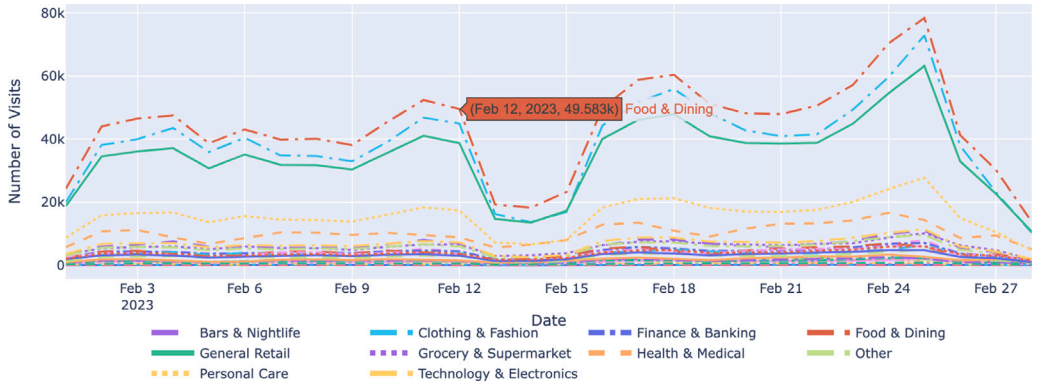

Fig. 4. Daily visit counts for the top 10 most frequently visited POI categories during February 2023, showing the sharp decline during Cyclone Gabrielle (12–14 February) and subsequent recovery patterns.

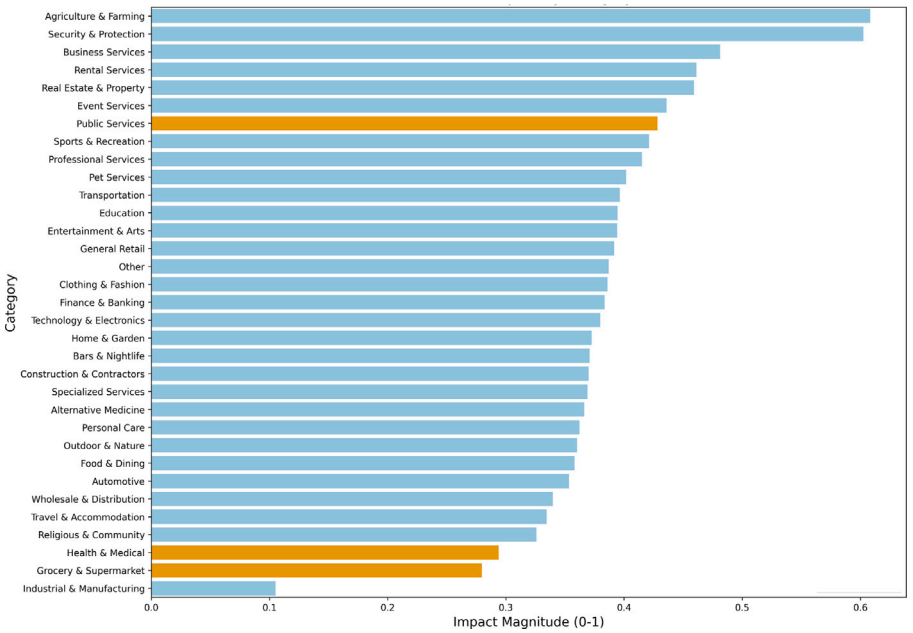

Fig. 5. Disaster impact magnitude by category type, showing the proportional reduction in visitation during Cyclone Gabrielle compared to pre-disaster baseline. Essential services are highlighted in orange, while non-essential services are shown in blue. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

where $\bar{N}_{\text{pre}}(c)$ and $\bar{N}_{\text{during}}(c)$ represent the average hourly visit counts for category c during the pre-disaster and disaster periods, respectively. This formulation ensures that $I(c)$ is bounded between 0 and 1, where 0 indicates no impact (unchanged visitation) and 1 represents complete cessation of visits.

Fig. 5 illustrates the disaster impact magnitudes across essential and non-essential categories. For essential services, we observe a mean impact magnitude of 0.32 (95% CI: [0.28, 0.36]), indicating that these categories retained approximately 68% of their normal visitation levels during the cyclone. Non-essential services exhibit a significantly higher mean impact of 0.39 (95% CI: [0.36, 0.43]), corresponding to a 39% reduction in visits.

Notably, certain categories demonstrate atypical patterns relative to their designation. For instance, Home & Garden (impact 0.58) and Bars & Nightlife (impact 0.56) display exceptionally high impact values among non-essential services, likely reflecting their greater elasticity of demand. Conversely, Public Services (impact 0.28) and Finance & Banking (impact 0.30) show remarkable resilience among essential services, potentially due to their role in emergency response and financial security during crisis events.

Estimating Recovery Rates To quantify the speed of recovery for different categories for the proposed disaster-aware model, we developed a methodology to estimate category-specific recovery rates from the post-disaster visitation data. For each category c, we

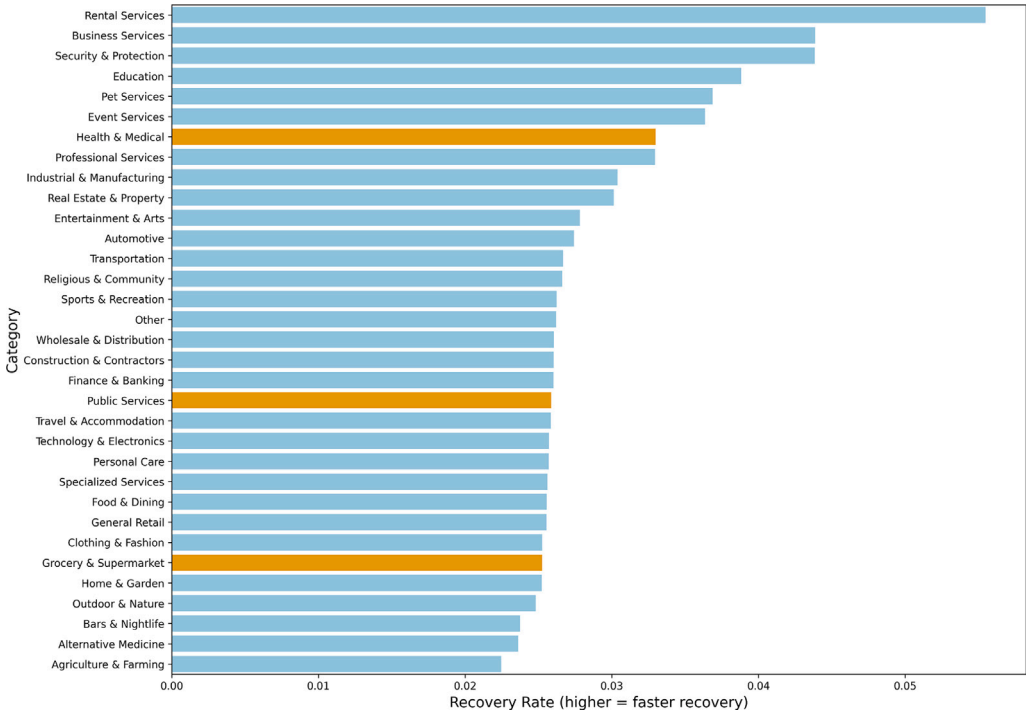

Fig. 6. Recovery rates by category, sorted from fastest to slowest recovery. Higher values indicate faster return to pre-disaster visitation levels.

model the recovery process using an exponential function:

$$R(t,c) = 1 - I(c) \cdot e^{-\gamma(c)\cdot(t-t_{\text{end}})} \tag{2}$$

where $R(t,c)$ represents the recovery ratio at time t after the disaster end t_{end}, $I(c)$ is the previously estimated impact magnitude, and $\gamma(c)$ is the category-specific recovery rate parameter to be estimated. Higher values of $\gamma(c)$ indicate faster recovery.

To estimate $\gamma(c)$ for each category, we employ the following procedure:

1. Aggregate post-disaster visitation data into time bins (6-h intervals) to reduce noise.
2. For each time bin, calculate the observed recovery ratio as $\frac{N_t(c)}{\bar{N}_{\text{pre}}(c)}$, where $N_t(c)$ is the visit count at time t after the disaster.
3. Fit the exponential recovery model using non-linear least squares optimisation to minimise:

$$\sum_{t\in T_{\text{post}}} \left(R(t,c) - \frac{N_t(c)}{\bar{N}_{\text{pre}}(c)} \right)^2 \tag{3}$$

4. Apply bootstrap resampling with 1,000 iterations to establish confidence intervals for the estimated recovery rates.

Fig. 6 presents the estimated recovery rates by category, revealing substantial variation both within and between essential and non-essential groupings. While essential services generally recover more quickly, certain non-essential categories such as Rental Services (0.053), Business Services (0.042), and Security & Protection (0.041) exhibited notably rapid recovery rates. This finding challenges the simplistic essential/non-essential dichotomy often employed in disaster response planning, suggesting that recovery dynamics follow more complex patterns than previously recognised.

The substantial heterogeneity in recovery trajectories observed across categories provides compelling empirical evidence for the necessity of category-specific parameterisation in our model. The nearly four-fold difference in recovery rates between the fastest-recovering categories (Rental Services at 0.053) and slowest-recovering categories (Agriculture & Farming at 0.016) demonstrates that a uniform recovery model would systematically mischaracterise the temporal dynamics of post-disaster mobility. Furthermore, the finding that recovery dynamics do not align perfectly with the essential/non-essential dichotomy, where some non-essential categories recovering faster than essential ones, underscores the need for granular, category-specific parameters rather than broad categorical groupings.

2.5. Categorical transition patterns

Perhaps the most distinctive feature of our dataset is its rich categorisation of visited locations and the resulting transition patterns between these categories. Fig. 7(a) presents the co-visit actual flows, or number of people who visits one category then another, while Fig. 7(b) shows this co-visitation as an asymmetric ratio between the categories.

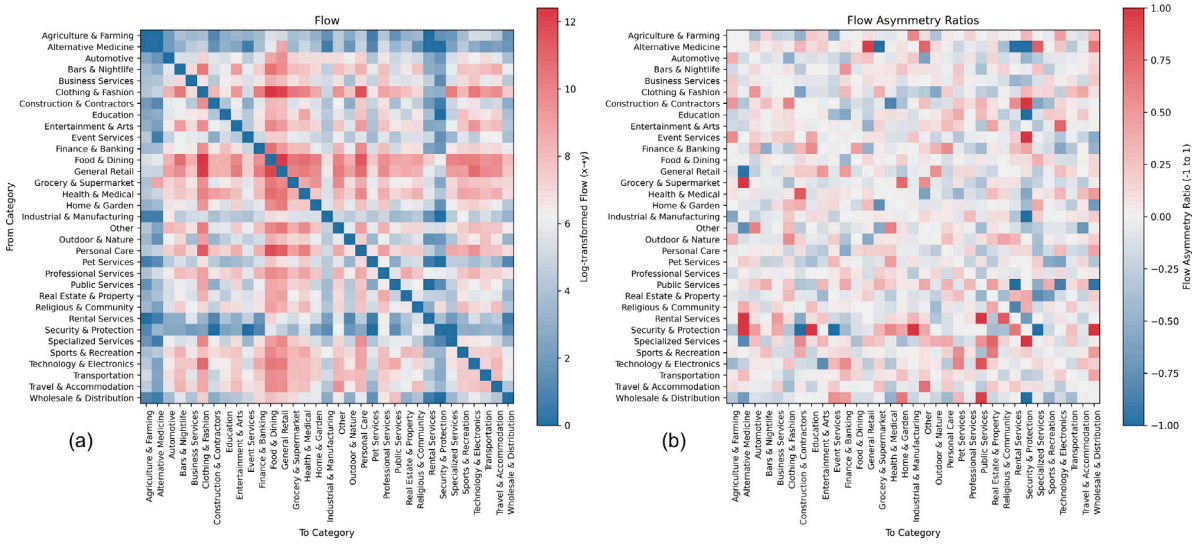

Fig. 7. Visit flows and flow asymmetric ratio.

Table 1
Distributional validation results: Comparison of mobility metrics between original and balanced datasets.

Metric	Original dataset			Undersampled dataset			Statistical tests	
	n	Mean	Std	n	Mean	Std	KS p-value	MW p-value
Essential visit rates	6	3.281	1.338	6	3.328	1.332	0.9307	0.5887
Non-essential visit rates	93	3.327	3.627	93	3.310	3.555	1.0000	0.9674
Essential distances (m)	6	4422.098	355.354	6	4565.511	348.183	0.9307	0.4848
Non-essential distances (m)	93	5038.599	1115.305	93	5183.376	1170.818	0.4230	0.2560

Note: KS = Kolmogorov–Smirnov test; MW = Mann–Whitney U test. All p-values > 0.05 indicate non-significant differences between distributions, suggesting successful preservation of mobility patterns in the balanced dataset.

Fig. 7(a) presents the transition flow matrix, revealing substantial heterogeneity in mobility patterns between location categories. The matrix demonstrates that certain categories serve as major attractors in the mobility network, with Food & Dining, General Retail, Grocery & Supermarket, and Health & Medical receiving disproportionately high inflows from multiple other categories. This concentration of flows reflects the needs for a category-specific model that can capture co-visitation flows between categories.

Fig. 7(b) quantifies the asymmetric nature of these transitions through flow asymmetry ratios, where values approaching 1 or −1 indicate highly unidirectional flows. These asymmetric transition probabilities directly inform the estimation of category-specific excitation strengths in our model, capturing how visits to one category influence subsequent visits to others.

2.6. Potential sampling bias in dataset

However, there is a potential sampling bias among devices with home location identified, with 42,849 (67.4%) are in areas with low deprivation areas and 20,707 (32.6%) are in high deprivation areas (ideally the ratio should be 50–50). To assess the implications of this bias, we first randomly undersampled the larger group (low deprivation) to have the same number of devices as the smaller group (high deprivation).

We then employed a distributional validation approach that addresses the statistical limitations inherent in traditional device-level comparisons. Rather than directly comparing individual mobility behaviours between datasets, our methodology calculates category-level mobility metrics for all location categories across temporal periods, then compares the distributions of these aggregated metrics between the original and undersampled datasets. Specifically, we computed visit rates (visits per device) and travel distances for both essential and non-essential services across pre-disaster, during-disaster, and post-disaster periods, generating distributions comprising up to 99 data points (33 categories times 3 periods) for each metric type. We then applied both the Kolmogorov–Smirnov (KS) test for distributional similarity and the Mann–Whitney U (MW) test for non-parametric comparison of these metric distributions. This approach provides statistically robust validation by comparing meaningful aggregated patterns rather than individual observations, thereby ensuring that our balanced sampling methodology preserves the essential characteristics of mobility behaviour whilst maintaining adequate statistical power for reliable inference.

The distributional validation results presented in Table 1 provide compelling evidence that the original dataset and undersampled dataset has the same distribution in terms of visit rates and travel distances, as no statistically significant differences were detected. The descriptive statistics also reveal remarkably similar central tendencies and variability patterns: essential and non-essential

visit rates show virtually identical means and standard deviations. The travel distance metrics demonstrate similar preservation, with differences in means representing less than 3% variation for essential and non-essential services. These findings collectively demonstrate that the balanced sampling process maintains the distributional characteristics. Because the potential sampling bias does not significantly impact the mobility patterns, we hereby proceed with the original dataset instead of the undersampled dataset to maximise the sample sizes for modelling.

3. Methodology

The exploratory analysis in the previous section sets the stage for our methodological innovation: a disaster-aware Hawkes process model that extends the standard self-exciting point process framework to incorporate disaster-specific dynamics, enabling more accurate modelling of both the suppression of mobility during disasters and the subsequent recovery phase.

3.1. Hawkes process framework

Point processes provide a mathematical framework for modelling the occurrence of events over time, where each event is represented as a point on a temporal axis. Traditional point process models, such as homogeneous or inhomogeneous Poisson processes, assume event occurrences are independent of one another, with intensities that either remain constant or vary according to exogenous factors [11]. While these models have proven valuable in many domains, they cannot capture the self-exciting and mutually-exciting patterns characteristic of human mobility behaviour.

Human visitation patterns exhibit clear dependencies where one visit often triggers subsequent visits. Standard point processes fail to model these behavioural cascades. Furthermore, during disasters, mobility patterns undergo dramatic shifts that include not only reduction in overall activity but also changing relationships between different types of locations and recovery dynamics that evolve over time.

Hawkes processes, introduced by Hawkes [12], offer a mathematical framework specifically designed to capture these complex dependencies. They are a class of self-exciting point processes where the occurrence of events increases the probability of future events temporarily. The standard univariate Hawkes process is defined by its conditional intensity function:

$$\lambda(t) = \mu + \sum_{t_i < t} \alpha \cdot e^{-\beta(t-t_i)} \tag{4}$$

where μ represents the background intensity (the rate of spontaneous events), α is the excitation parameter that controls how much each event contributes to future intensity, and β determines the decay rate of the excitation effect. This formulation creates a memory effect, where past events influence the probability of future events in a diminishing fashion over time.

While Hawkes processes have been successfully applied to various domains including seismology [13], and social media dynamics [14], standard Hawkes process models are inadequate for capturing the complex dynamics of mobility during disasters for several reasons. First, disasters typically suppress rather than excite mobility, contradicting the fundamental self-exciting (increased event probability) nature of traditional Hawkes processes. Second, the impact of disasters varies dramatically across different location categories. For instance,essential services like hospitals and grocery stores experience different patterns from non-essential services like entertainment venues. Third, recovery from disasters follows heterogeneous trajectories that depend on both location type and socioeconomic factors. Finally, the transition between normal functioning and disaster states is not instantaneous but follows complex temporal dynamics. For a full review of Hawkes processes and their applications, readers may refer to Laub et al. [15] for more details.

3.2. Disaster-aware Hawkes process

To address the challenges described at the end of Section 2, we develop a Disaster-aware Hawkes process model that incorporates regime-switching, category-specific recovery dynamics, and post-disaster bounce effects. Our enhanced model introduces several key innovations beyond the standard Hawkes process:

1. **A regime-switching baseline intensity function** that transitions between normal and disaster periods. This innovation is essential because mobility patterns fundamentally change during disasters, not just in volume but in structural relationships between locations. By modelling distinct regimes, we can capture these structural shifts while still maintaining continuity in the underlying mathematical framework.
2. **Category-specific impact and recovery parameters** that account for the heterogeneous vulnerability and resilience profiles across different types of locations. As seen in Section 2, essential services typically recover more quickly than non-essential ones, and the impact magnitude varies dramatically across categories. These parameters move beyond simplistic essential/non-essential dichotomies to capture the full spectrum of disaster responses.
3. **Post-disaster bounce-back effects** to model temporary surges in activity. Our empirical observations in Section 2 reveal that certain categories experience visitation rates that temporarily exceed pre-disaster levels during the recovery phase. These category-specific bounce parameters capture this important but previously unmodelled phenomenon of compensatory behaviour and delayed demand fulfilment.

Beyond these core innovations, our model maintains the standard Hawkes process excitation component to capture co-visitation patterns between categories. Different location categories exhibit distinct patterns of interaction—for instance, a visit to a grocery store might trigger subsequent visits to retail locations, while healthcare visits might prompt pharmacy visits. This excitation mechanism, combined with our disaster-specific components, enables comprehensive modelling of mobility dynamics.

These innovations collectively enable our model to capture the full lifecycle of disaster impacts on mobility: from the initial suppression during the disaster (through regime-switching and impact parameters), through the complex recovery phase (via category-specific recovery rates), to post-disaster surges (through bounce-back effects), and eventual return to normalcy. The complete intensity function for our model is:

$$\lambda(t,c) = \lambda_0(t,c,r) \cdot (1 + B(t,c)) + \sum_{(t_i,c_i)\in\mathcal{H}_t} \alpha(c_i \to c) \cdot w(t_i,d_i) \cdot \kappa(t-t_i) \tag{5}$$

where:

- $\lambda_0(t,c,r)$ is the regime-dependent baseline intensity for category c at time t under regime r
- $B(t,c)$ is the bounce-back function which models post-disaster activity surges
- $\mathcal{H}_t$ is the set of events (visits) that occurred before time t
- $\alpha(c_i \to c)$ represents the excitation strength from category c_i to category c
- $w(t_i,d_i)$ is a weight function based on dwell time d_i at visit i
- $\kappa(t-t_i)$ is a decay kernel that models how excitation diminishes over time

Regime-Switching Baseline Intensity The baseline intensity function $\lambda_0(t,c,r)$ captures the expected visitation rate for category c at time t under regime r. We define two distinct regimes: normal ($r=0$) and disaster ($r=1$). The baseline intensity is defined as:

$$\lambda_0(t,c,r) = (1-r) \cdot \lambda_0^{\text{normal}}(t,c) + r \cdot \lambda_0^{\text{disaster}}(t,c) \tag{6}$$

Each baseline component incorporates temporal patterns such as time-of-day and day-of-week effects:

$$\lambda_0^{\text{normal}}(t,c) = f(h(t), w(t), c) \tag{7}$$

$$\lambda_0^{\text{disaster}}(t,c) = g(h(t), w(t), c) \tag{8}$$

where $h(t)$ extracts the hour of day from time t, and $w(t)$ indicates whether time t falls on a weekend.

The regime indicator r smoothly transitions from 0 to 1 at disaster onset and back to 0 during recovery:

$$r(t) = \begin{cases} 0 & \text{if } t < t_{\text{start}} \\ 1 & \text{if } t_{\text{start}} \le t \le t_{\text{end}} \\ \exp(-\gamma_r \cdot (t - t_{\text{end}})) & \text{if } t > t_{\text{end}} \end{cases} \tag{9}$$

where t_{start} and t_{end} denote the disaster start and end times, respectively, and γ_r controls the speed of regime transition back to normal.

Bounce-Back Function The bounce-back function $B(t,c)$ models the category-specific temporary surge in activity that exceeds pre-disaster levels during the recovery phase. This phenomenon reflects delayed need fulfilment, compensatory behaviour, and resource replenishment following the constrained activity during the disaster period. The bounce-back function is defined as:

$$B(t,c) = \begin{cases} 1 & \text{if } t \le t_{\text{end}} \\ 1 + M(c) \cdot \exp\left(-\delta(c) \cdot (t - t_{\text{end}} - \tau(c))^2\right) & \text{if } t > t_{\text{end}} \end{cases} \tag{10}$$

where:

- $M(c)$ is the maximum bounce magnitude for category c (bounded between 0 and 1, where 0 indicates no bounce and 1 represents a 100% increase over the recovered baseline)
- $\tau(c)$ is the time delay until peak bounce for category c (measured in hours after disaster end)
- $\delta(c)$ is the decay rate of the bounce effect for category c (controlling how quickly the surge dissipates)

This bell-shaped curve allows for the modelling of temporary activity surges that gradually build to a peak before dissipating as the system returns to equilibrium. Empirical evidence suggests that different POI categories exhibit markedly different bounce characteristics, with some showing dramatic but short-lived peaks (e.g., retail), while others display more moderate but sustained elevations (e.g., restaurants).

Excitation Component The excitation component captures how visits to one category trigger subsequent visits to other categories. The excitation strength $\alpha(c_i \to c)$ is derived from observed transition probabilities between categories. The decay kernel $\kappa(t-t_i)$ follows an exponential form:

$$\kappa(t-t_i) = \beta(c) \cdot \exp(-\beta(c) \cdot (t-t_i)) \tag{11}$$

where $\beta(c)$ is the category-specific decay rate. The dwell time weight function $w(t_i, d_i)$ accounts for the fact that longer visits may have stronger influence on subsequent behaviour:

$$w(t_i, d_i) = \min(1, d_i/\bar{d}) \tag{12}$$

where $\bar{d}$ is the average dwell time across all visits.

Whilst our model does not explicitly incorporate spatial coordinates, the categorical structure effectively captures spatial heterogeneity through location type differentiation. Different POI categories exhibit distinct spatial distributions across urban landscapes—for instance, Health & Medical facilities cluster in specific zones whilst Grocery & Supermarket locations disperse more uniformly. Our category-specific parameters (impact magnitudes, recovery rates, excitation strengths) therefore implicitly encode spatially-differentiated disaster effects without requiring explicit spatial modelling. This categorical approach offers several advantages: it reduces model complexity whilst maintaining interpretability, focuses on functional rather than purely geographic relationships, and enables analysis at policy-relevant scales where decision-makers think in terms of service types rather than specific coordinates.

3.3. Model implementation

To address data sparsity in certain categories, preserve personal information in mobility data, and enhance computational efficiency, we aggregate visits into 3-h time windows while maintaining category-level information.

Category-specific Baseline Estimation We estimate the normal baseline intensity $\lambda_0^{\text{normal}}(t, c)$ for each category using pre-disaster data, aggregated by category, hour of day, and weekend status. For the disaster baseline $\lambda_0^{\text{disaster}}(t, c)$, we use data from the disaster period if available, or apply a scaling factor to the normal baseline if disaster-period data is insufficient.

Formally, we estimate:

$$\lambda_0^{\text{normal}}(c, h, w) = \frac{1}{|\mathcal{T}_{c,h,w}^{\text{pre}}|} \sum_{t \in \mathcal{T}_{c,h,w}^{\text{pre}}} N(t) \tag{13}$$

$$\lambda_0^{\text{disaster}}(c, h, w) = \frac{1}{|\mathcal{T}_{c,h,w}^{\text{during}}|} \sum_{t \in \mathcal{T}_{c,h,w}^{\text{during}}} N(t) \tag{14}$$

where $\mathcal{T}_{c,h,w}^{\text{pre}}$ and $\mathcal{T}_{c,h,w}^{\text{during}}$ represent the sets of time windows with category c, hour h, and weekend status w in the pre-disaster and during-disaster periods, respectively, and $N(t)$ denotes the visit count in time window t. To ensure numerical stability, we apply a minimum threshold of 0.1 visits per time window for both normal and disaster baselines.

Excitation Matrix Estimation To estimate the excitation strength $\alpha(c_i \to c)$, we analyse sequential visits by the same device across different categories. We construct a transition matrix P where each element $P_{i,j}$ represents the probability of transitioning from category i to category j:

$$P_{i,j} = \frac{N_{i \to j}}{\sum_k N_{i \to k}} \tag{15}$$

where $N_{i \to j}$ is the count of transitions from category i to category j. The excitation strength is then derived as:

$$\alpha(c_i \to c) = \min(\alpha_{\max}, \eta \cdot P_{c_i,c}) \tag{16}$$

where η is a scaling factor and $\alpha_{\max}$ is an upper bound to ensure numerical stability. We apply different scaling factors for essential versus non-essential services to reflect their distinct behavioural influence patterns.

Bounce-Back Parameter Estimation A key innovation in our implementation is the explicit modelling of the post-disaster bounce-back effect, where visitation temporarily exceeds pre-disaster levels during recovery. For each category, we estimate three key parameters:

1. Bounce magnitude $M(c)$: The maximum amplitude of the post-disaster surge relative to the pre-disaster baseline. This is estimated by identifying the peak post-to-pre visitation ratio:

$$M(c) = \max\left(0, \max_{t > t_{\text{end}}} \left(\frac{N_t(c)}{\bar{N}_{\text{pre}}(c)}\right) - 1\right) \tag{17}$$

2. Bounce delay $\tau(c)$: The time (in hours) from disaster end until peak bounce effect. This is estimated by identifying the time of maximum post-disaster visitation:

$$\tau(c) = \arg\max_{t > t_{\text{end}}} \frac{N_t(c)}{\bar{N}_{\text{pre}}(c)} - t_{\text{end}} \tag{18}$$

3. Bounce decay $\delta(c)$: The rate at which the bounce effect diminishes. This is estimated through non-linear optimisation of the bounce function fit to post-disaster data.

For categories with insufficient post-disaster data to reliably estimate these parameters, we apply default values derived from categories with similar functional characteristics. Specifically, we group categories along two primary functional dimensions:

- Essential vs. Non-essential Services: Essential services (Health & Medical, Grocery & Supermarket, Public Services, Gas & Automotive) exhibit fundamentally different recovery dynamics compared to non-essential or discretionary services (e.g. Food & Dining, Retail, Entertainment) because they serve critical daily needs that cannot be indefinitely deferred. Essential services typically show more restrained bounce-back effects and steadier recovery trajectories.
- Service Interaction Patterns: Categories also differ in how visits to one type influence visits to others. For example, retail and entertainment venues show similar excitation patterns (visits to one often trigger visits to similar venues), whilst healthcare services exhibit more independent demand patterns.

When insufficient post-disaster data exists for reliable parameter estimation in a specific category, we assign default bounce-back parameters based on these functional groupings. For instance, a new retail subcategory would inherit parameters from the broader retail/discretionary service group, whilst a specialised medical service would inherit from the essential services group.

To distinguish genuine bounce-back effects from random fluctuations, we apply a threshold criterion: a category is considered to exhibit bounce-back only if its maximum post-disaster visitation ratio exceeds 1.1 (10% above pre-disaster baseline). In our implementation, we employ computational optimisation to simultaneously estimate both recovery and bounce parameters by minimising the sum of squared errors between observed and predicted visitation ratios.

3.4. Parameter optimisation

We employ maximum likelihood estimation (MLE) for parameter estimation because of computational efficiency for such a large-scale dataset, interpretability and a straight-forward comparison to standard Hawkes process. Nevertheless, we acknowledge that Bayesian approaches would offer certain benefits, particularly in quantifying parameter uncertainty and incorporating prior knowledge about disaster effects. For specific analyses where uncertainty quantification is paramount, such as estimating recovery rates for categories with sparse data, Bayesian methods could complement our MLE approach. Future work could explore hybrid approaches that leverage the computational efficiency of MLE for most parameters while employing Bayesian methods for critical subsets of parameters where uncertainty quantification is most valuable.

The log-likelihood function for our Hawkes process model follows the standard form for temporal point processes:

$$\mathcal{L}(\theta) = \sum_{i=1}^{n} \log \lambda(t_i, c_i) - \sum_{c \in C} \int_0^T \lambda(t, c) dt \tag{19}$$

where θ represents the complete set of model parameters, the first term sums over all observed events (the (t_i, c_i) pairs representing visits to category c_i at time t_i), and the second term integrates the intensity function over the entire observation period $[0, T]$ and across all categories C. The first term rewards parameter values that assign high intensity to times and categories where visits actually occurred, while the second term penalises excessively high intensity values overall, creating the necessary balance for accurate prediction. This formulation follows directly from the definition of the likelihood function for point processes, where the probability of observing a specific realisation is proportional to the product of intensities at event times multiplied by the exponential of the negative integrated intensity [16].

In practice, computing the exact integral in the second term is intractable, particularly for our complex intensity function with multiple components. We therefore employ a numerical approximation using a discretised time grid:

$$\int_0^T \lambda(t, c) dt \approx \sum_{j=1}^{m} \lambda(t_j, c) \Delta t \tag{20}$$

where m is the number of time windows, t_j represents the midpoint of the jth time window, and Δt is the width of each window (1 h in our implementation). This discretisation aligns with our approach of modelling the expected counts per time window rather than instantaneous rates.

This high-dimensional parameter space θ presents significant optimisation challenges. To make the optimisation process more robust and computationally tractable, we employ a strategic approach combining:

1. Parameter grouping and sequential optimisation: Rather than attempting to optimise all parameters simultaneously, we estimate certain parameter groups sequentially. For instance, we first estimate baseline intensities from pre-disaster data, then estimate excitation parameters, followed by disaster impact parameters.
2. Parameter constraints: We impose bound constraints on all parameters to ensure they remain in physically meaningful ranges (e.g., non-negative excitation strengths, recovery rates bounded between 0 and 1). These constraints not only improve interpretability but also enhance numerical stability.
3. Regularisation: For the excitation matrix, which has $|C|^2$ parameters, we apply L1 regularisation to promote sparsity, reflecting the assumption that most category pairs have negligible mutual excitation.
4. Subset-based estimation: For computational efficiency with large datasets, we employ stratified sampling to create representative subsets of data for parameter estimation, ensuring balanced representation of pre-disaster, during-disaster, and post-disaster periods.

For the optimisation algorithm itself, we employ the Limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm with bound constraints (L-BFGS-B) [17]. This choice is motivated by several factors:

- It efficiently handles high-dimensional parameter spaces by approximating the Hessian matrix using a limited memory approach, making it suitable for our complex model.
- It natively supports bound constraints, which are essential for maintaining parameter interpretability and model stability.
- It demonstrates good convergence properties for non-linear optimisation problems with smooth objective functions.
- It offers a favourable balance between computational efficiency and optimisation effectiveness compared to alternatives like Newton's method (more computationally intensive) or gradient descent (slower convergence).

Our implementation initialises parameters with educated guesses based on domain knowledge. For instance, setting higher recovery rates for essential services and lower rates for discretionary categories, which helps avoid poor local optima and accelerates convergence.

3.5. Model evaluation

To evaluate our model, we compare the observed visit counts with the predicted counts derived from our intensity function over fixed time windows. For each time window of duration Δt (1 h in our implementation), the expected count is $\mathbb{E}[N(t, t+\Delta t)] = \int_t^{t+\Delta t} \lambda(s)ds$, which we approximate as $\hat{\lambda}_i \cdot \Delta t$ for computational efficiency.

Since we use uniform time windows throughout our analysis, and our intensity function $\hat{\lambda}_i$ is already scaled to represent expected counts per time window rather than instantaneous rates, we simplify our notation by directly comparing N_i with $\hat{\lambda}_i$. Our evaluation metrics are:

1. Mean Squared Error (MSE): $\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(N_i - \hat{\lambda}_i)^2$
2. Mean Absolute Error (MAE): $\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|N_i - \hat{\lambda}_i|$
3. Mean Relative Error (MRE): $\text{MRE} = \frac{1}{n}\sum_{i=1}^{n}\frac{|N_i - \hat{\lambda}_i|}{N_i+1}$

where N_i is the observed visit count and $\hat{\lambda}_i$ is the predicted count for time window i. The addition of 1 in the denominator of MRE ensures numerical stability when dealing with time windows that have zero observed visits.

To quantify the improvement gained by incorporating disaster-specific components, we compare our enhanced model with a standard Hawkes process model using the percentage reduction in MSE:

$$\text{Improvement} = \frac{\text{MSE}_{\text{standard}} - \text{MSE}_{\text{enhanced}}}{\text{MSE}_{\text{standard}}} \times 100\% \tag{21}$$

We also evaluate model performance separately for pre-disaster, during-disaster, and post-disaster periods to assess how well our model captures the distinct dynamics of each phase.

3.6. Extensions for socioeconomic analysis

Our framework can be extended to incorporate socioeconomic factors, such as deprivation levels [9], by segmenting devices based on home location characteristics. For each deprivation group s, we can fit a separate model to analyse differential impacts and recovery patterns:

$$\lambda^s(t,c) = \lambda_0^s(t,c,r)\cdot(1+B^s(t,c)) + \sum_{(t_i,c_i)\in\mathcal{H}_t^s} \alpha^s(c_i \to c)\cdot w(t_i,d_i)\cdot\kappa(t-t_i) \tag{22}$$

This enables investigation of social equity questions such as whether disadvantaged communities experience more severe disaster impacts or slower recovery rates for essential services.

3.7. Model transferability and recalibration

The Disaster-aware Hawkes process framework is designed with transferability across disaster types and contexts in mind. The model comprises both *universal components* that capture fundamental behavioural responses and *context-specific parameters* that require local calibration. Universal components include the mathematical structure of suppression (disaster impact function), recovery dynamics (exponential recovery with category-specific rates), and compensatory behaviour (bounce-back effects). These mechanisms reflect basic human responses to disruption that transcend specific hazard types. For application to new contexts, we propose a recalibration protocol: (1) preserve the model structure and universal mechanisms, (2) estimate context-specific parameters using local pre-disaster data for baselines and during-disaster data for impact magnitudes, (3) validate recovery and bounce-back parameters using post-disaster observations where available, and (4) adapt excitation parameters based on local co-visitation patterns.

4. Results

In this section, we present empirical evidence that the proposed Disaster-Aware Hawkes process captures the dynamics of human mobility before, during, and after Cyclone Gabrielle with high fidelity. Our evaluation approach differs fundamentally from typical machine learning validation because the primary objective is assessing mechanistic validity rather than optimising predictive performance. The performance improvements reported serve to validate that our mechanistic innovations successfully capture the underlying causal processes governing disaster-affected mobility.

4.1. Component-wise performance analysis

To quantify the value of each modelling component, we conduct a controlled ablation in which we disable, one at a time, the regime-switching baseline and the post-disaster bounce-back mechanism, and we contrast these with a standard Hawkes baseline.

Experimental Design We evaluated four distinct model configurations, by selectively disabling individual features of the proposed Disaster-Aware Hawkes process.

1. Full Disaster-aware Hawkes: The complete model with regime-switching baseline intensity, category-specific impact and recovery parameters, and post-disaster bounce-back effects (all components enabled).
2. Without Regime Switching: Maintains a single baseline intensity throughout the observation period whilst retaining bounce-back effects and category-specific parameters. This configuration tests whether explicit regime differentiation provides meaningful value beyond parameter heterogeneity.
3. Without Bounce-Back: Employs regime switching and category-specific parameters but disables the post-disaster surge modelling. This isolates the contribution of compensatory behaviour effects.
4. Standard Hawkes: A baseline configuration as in the previous subsection.

Each configuration was calibrated using identical maximum likelihood estimation procedures on the same stratified data sample, ensuring fair comparison. Each ablated model employs the same diurnal and day-of-week variation in the background intensity as in our enhanced model:

$$\mu(t, c) = f(h(t), w(t), c) \tag{23}$$

where $h(t)$ and $w(t)$ extract the hour of day and weekend status from time t, respectively. This ensured that observed performance differences were not attributable to differences in baseline intensity modelling.

We also derived the excitation parameters $\alpha_{c_i \to c}$ from the same transition matrix used in our enhanced model, maintaining identical excitation strengths between categories. Similarly, we employed the same decay rates β_c for the excitation kernel, ensuring consistency in the temporal influence profile.

For prediction, as the standard model does not explicitly model a disaster period, we apply a fixed temporal dampening factor of 0.8 to the intensity function during the disaster period (20% reduction in intensity, as seen in the approximate reduction in activities seen in Section 2). Essentially, this is a simpler variation of the Regime-switching function in the full Disaster-aware Hawkes model, where the $\lambda_0^{\text{normal}}(t, c)$ is multiplied by 0.8 during the disaster period.

Fig. 8(a) shows that the proposed Disaster-Aware Hawkes process (Full Model) outperforms all ablated variants. When Regime-Switching is omitted, the MSE increases significantly, underscoring the pivotal role of this mechanism in capturing phase transitions between baseline, shock, and recovery dynamics. Similarly but to a lesser extend, ablation of the Bounce-Back component also results in a higher MSE, evidencing the sensitivity of disaster impact recovery to explicit post-shock correction. The Standard Hawkes model, lacking both disaster-time adaptation and recovery mechanisms, reports the highest overall MSE, establishing the value of domain-aware temporal structure.

The regime-switching baseline intensity emerges as the most critical component for disaster-period prediction. Fig. 8(b) shows that without Regime-Switching, the model performs similarly than the Standard Hawkes (both only has the 0.8 dampening factor instead of the Regime-Switching capabilities). This finding validates our theoretical premise that disasters represent a distinct behavioural regime rather than merely a parametric variation of normal mobility.

The post-disaster bounce-back component proves essential for accurate recovery-phase prediction. Fig. 8(c) shows that models without this component (the Without Bounce-back and Standard Hawkes models) consistently underestimate visitation during the recovery period, particularly for categories exhibiting compensatory behaviour such as Health & Medical and Retail Services. The temporal pattern of errors in the no-bounce configuration reveals systematic underestimation 48–96 h post-disaster, precisely when empirical data shows surge effects.

Notably, the combined effect of all components exceeds the sum of individual contributions, suggesting important synergies. The interaction between regime switching and bounce-back effects proves particularly significant: regime switching establishes the suppressed baseline from which the bounce-back is measured, whilst the bounce parameters modulate the transition back to normal regime. This interdependence explains why the full model achieves 21.5% better overall performance compared to the Standard Hawkes model. This is significantly higher than summing individual component contributions.

Figs. 9 illustrates the proposed model's ability to predict visitation patterns across different POI categories, reflecting the inherent challenges in modelling heterogeneous mobility behaviours. The figure compares the actual versus predicted visitation patterns for the top five and bottom five most visited categories, respectively.

The model's ability to predict visitation patterns varies considerably across POI categories, reflecting the inherent challenges in modelling heterogeneous mobility behaviours. Fig. 9 illustrates the actual versus predicted visitation patterns for the top five and bottom five most visited categories, respectively, demonstrating the model's robust performance across vastly different scales of activity.

The regime-switching component of our Disaster-aware Hawkes process demonstrates remarkable effectiveness in capturing the abrupt transitions between normal and disaster states. This is particularly evident in the sharp discontinuities at disaster onset (12 February) and conclusion (15 February), where the model successfully tracks the near-instantaneous suppression of mobility across all categories. The regime-switching mechanism proves equally adept at capturing the heterogeneous nature of disaster

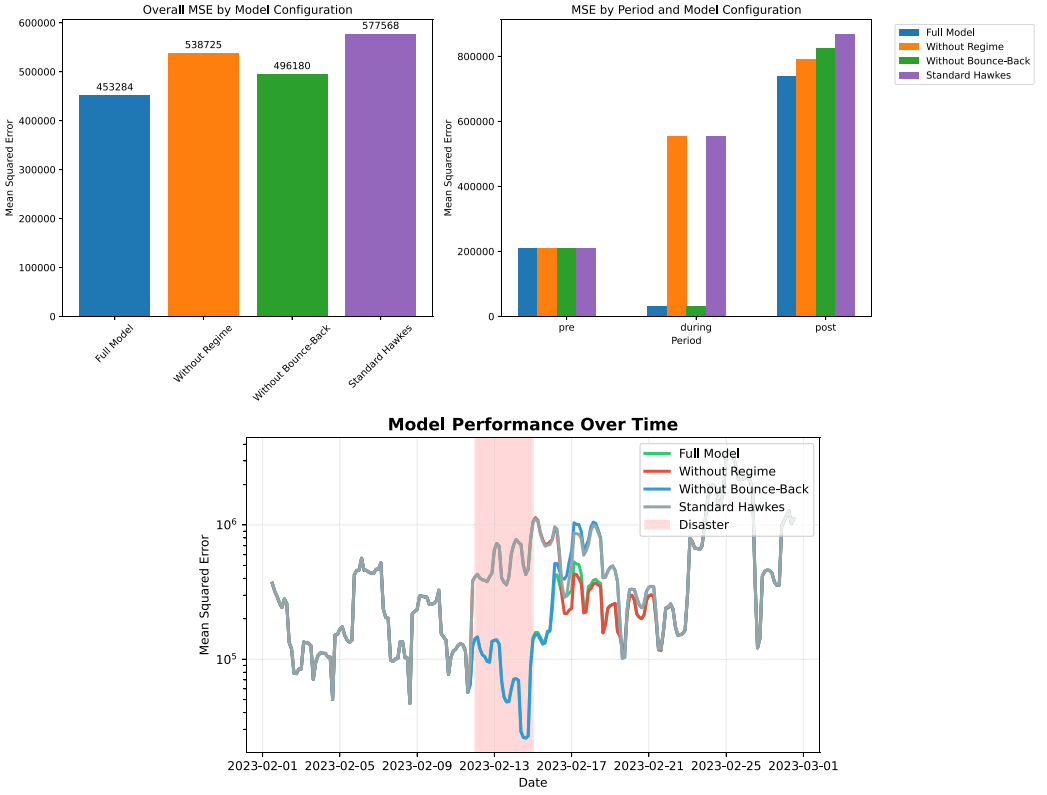

Fig. 8. Component-wise performance analysis.

impacts across different service types. Essential service categories, such as Health & Medical, exhibit more modest suppression (maintaining approximately 40%–50% of baseline activity), whilst discretionary categories like Clothing & Fashion experience more severe reductions (dropping to 20%–30% of normal levels). The model's ability to differentiate these patterns through category-specific disaster baseline intensities validates the theoretical framework's core assumption that disasters affect different location types through distinct mechanisms rather than uniform suppression.

The model also successfully identifies and predicts post-disaster bounce-back effects. This phenomenon is particularly pronounced in several categories within the bottom five most visited locations (Fig. 9b), where the relatively lower baseline volumes allow clearer visualisation of proportional changes.

The model's performance remains remarkably consistent across the dramatic scale differences between high-volume and low-volume categories. For high-volume categories such as Food & Dining (RMSE: 2,230.07) and Clothing & Fashion (RMSE: 2,274.09), whilst the absolute errors appear substantial, they represent proportionally similar accuracy to low-volume categories when normalised by baseline activity levels. For instance, Security & Protection achieves an RMSE of 8.94 against a baseline of approximately 30–40 visits, yielding a relative error comparable to Food & Dining's performance against its 15,000–20,000 visit baseline. This scale-invariant performance validates the model's mathematical formulation, particularly the multiplicative structure of the intensity function that allows category-specific parameters to scale appropriately across different visitation magnitudes. The consistent capture of temporal patterns — including weekly periodicity, diurnal cycles, and disaster-induced disruptions — across both high-frequency and low-frequency categories demonstrates the robustness of the underlying Hawkes process framework.

4.2. Identifying causal cascade relationships in disaster-affected mobility networks

The concept of cascade effects refers to how disruptions in one location category create causal chains of impacts that propagate through the mobility network, often triggering secondary disruptions far removed from the original source. Understanding these causal interdependencies represents a fundamental challenge in disaster resilience research: whilst traditional analysis can observe correlated changes across different location types, it cannot distinguish between independent responses to disaster conditions and genuine causal relationships where disruption to one category mechanistically drives changes in others.

Our Disaster-aware Hawkes process addresses this challenge by providing the first mathematical framework capable of identifying causal cascade relationships in disaster mobility networks. The key insight lies in the conditional intensity function's excitation parameters $\alpha(c_i \to c)$, which capture how events in category c_i influence the probability of subsequent events in category c. Unlike correlation-based measures that can reflect common responses to external factors, these parameters represent *causal* influence:

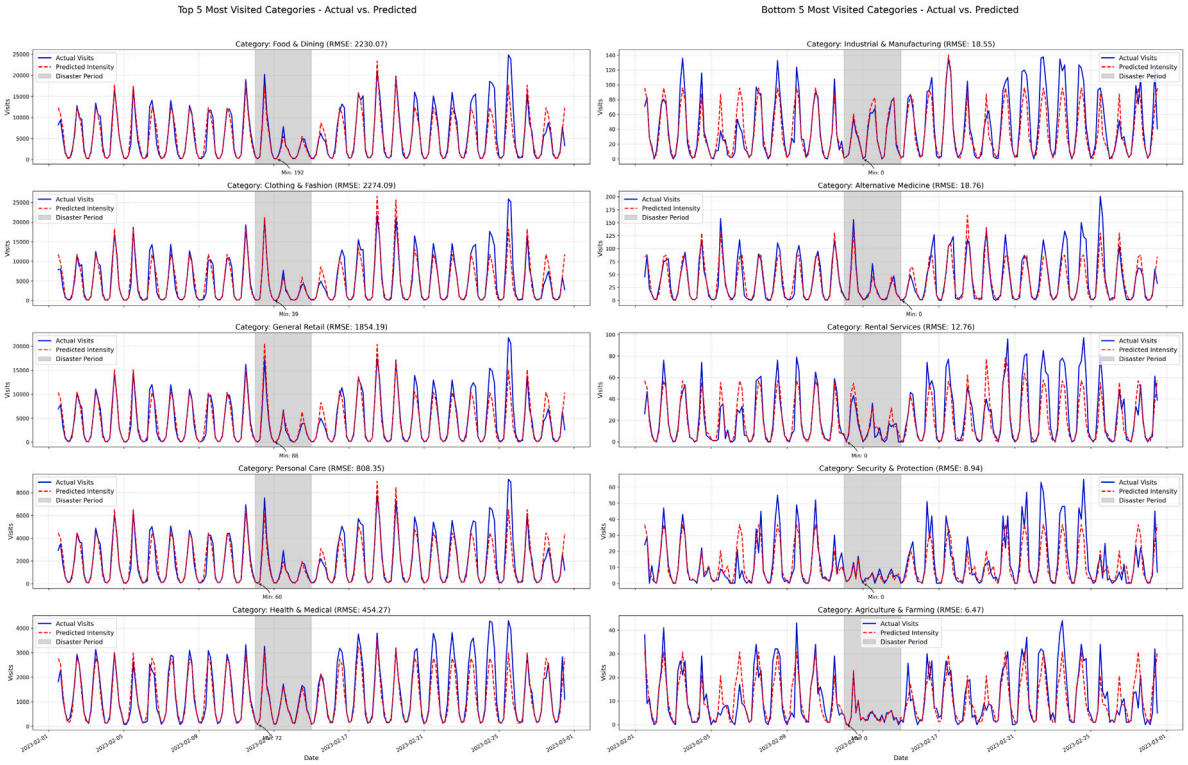

(a) Top 5 most visited categories (b) Bottom 5 most visited categories

Fig. 9. Actual vs. Predicted number of visits.

controlling for all other factors including temporal patterns, baseline intensities, and disaster effects, they quantify the degree to which a mobility event in one category increases the likelihood of future events in another category.

This causal interpretation emerges from the mathematical structure of the Hawkes process intensity function. When we observe a visit to category c_i at time t_j, the excitation term $\alpha(c_i \to c) \cdot \kappa(t - t_j)$ directly increases the intensity for category c at all subsequent times $t > t_j$, with the influence decaying according to the kernel $\kappa(\cdot)$. This mechanistic relationship — where past events directly modify future probabilities — provides a principled foundation for causal inference that traditional mobility analysis methods cannot offer.

Cascade Impact Matrix Construction Building upon the causal interpretation of excitation parameters, we developed a comprehensive framework to quantify cascade relationships across the disaster-affected mobility network. The cascade impact matrix $\mathbf{C}$ combines excitation strengths with disaster impact magnitudes to reveal how disruptions propagate:

$$C_{ij} = \alpha(c_i \to c_j) \times I(c_i) \times \tau_{ij} \tag{24}$$

where $\alpha(c_i \to c_j)$ represents the causal excitation strength from category i to category j, $I(c_i)$ is the disaster impact magnitude for the source category, and τ_{ij} is a temporal coupling factor that accounts for the time scale over which cascade effects propagate.

This formulation captures a critical insight invisible to traditional analysis: the magnitude of a cascade effect depends not only on the strength of normal interdependencies between categories, but also on how severely the source category is disrupted by the disaster. A category with strong excitation effects under normal conditions may generate limited cascades if it maintains high functionality during the disaster, whilst a heavily impacted category may propagate substantial secondary effects even with moderate excitation strengths.

Fig. 10 presents the cascade impact matrix for the most affected location categories during Cyclone Gabrielle. Several patterns emerge that challenge conventional understanding of disaster interdependencies:

– Asymmetric Cascade Pathways: The matrix reveals highly asymmetric causal relationships that would be invisible to correlation-based analysis. Real Estate & Property strongly influences General Retail (cascade strength 0.17) whilst receiving minimal influence in return (0.02). This asymmetry suggests a unidirectional causal pathway where property-related mobility (viewing homes, meeting with agents) triggers subsequent retail activity, but retail visits do not reciprocally influence property-related behaviour during disasters.

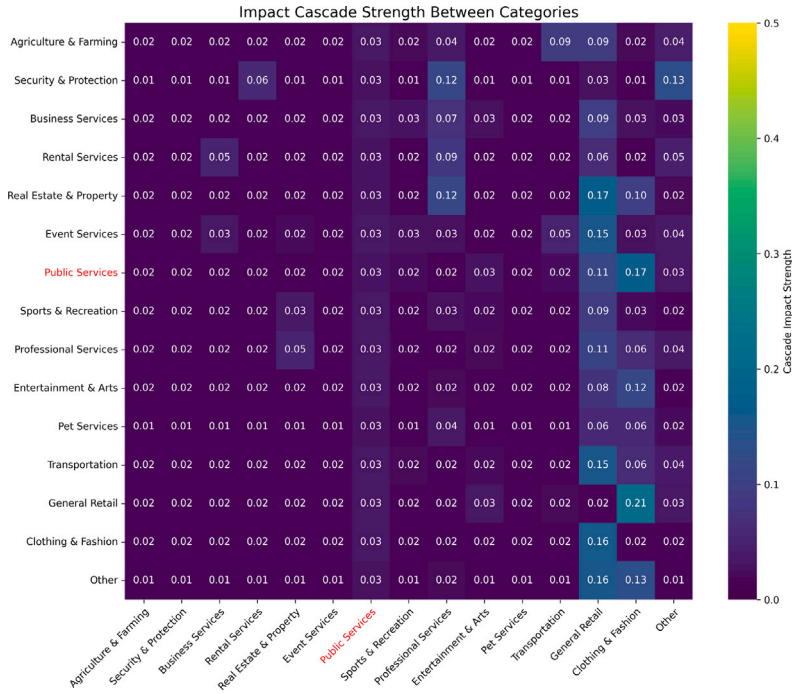

Fig. 10. Impact cascade strength between categories, showing how disruption in one category affects others. Essential services are highlighted in red on axis labels. Higher values (lighter colours) indicate stronger cascade effects. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

- Hub-Mediated Cascade Amplification: Categories such as General Retail and Public Services function as cascade amplifiers, receiving disruptions from multiple sources and propagating them broadly through the network. The strongest cascade pathway occurs from General Retail to Clothing & Fashion (0.21), indicating that disruptions to general retail infrastructure create cascading impacts on fashion-related mobility through mechanisms that extend beyond simple co-location or complementarity.
- Hidden Vulnerabilities: Alternative Medicine emerges as a critical cascade source despite moderate direct disaster impact (0.35), demonstrating how network position can amplify the systemic importance of seemingly peripheral categories. Traditional impact analysis focusing solely on direct effects would substantially underestimate the system-wide consequences of disrupting alternative medical services.

Causal Network Topology and Systemic Risk The cascade impact matrix enables construction of a directed causal network that reveals the topological structure underlying disaster vulnerability (Fig. 11). Network analysis of this causal structure provides insights unavailable through traditional spatial or categorical analysis:

- Cascade Amplifier Identification: PageRank centrality analysis identifies General Retail (0.110), Clothing & Fashion (0.099), and Health & Medical (0.068) as primary cascade amplifiers. These categories' high centrality scores reflect their role in propagating disruptions received from multiple sources, making them critical leverage points for intervention strategies.
- Structural Vulnerability Assessment: Betweenness centrality analysis reveals categories that serve as cascade bridges, connecting otherwise isolated parts of the mobility network. Personal Care emerges as a critical bridge category whose disruption could fragment the network into disconnected components, creating cascade isolation effects that compound direct impacts.
- Causal Pathway Resilience: The network structure reveals redundancy patterns in cascade pathways, showing which causal relationships have alternative routes and which represent single points of failure. Categories with high in-degree but low clustering coefficients represent cascade bottlenecks where disruption elimination can achieve disproportionate system-wide benefits.

Temporal Evolution of Causal Relationships A distinctive capability of our Hawkes process framework lies in capturing how causal relationships themselves evolve during disasters. The regime-switching baseline intensities modify not only the magnitude of mobility but also the strength of interdependencies between categories:

$$\alpha^{disaster}(c_i \to c_j) = \alpha^{normal}(c_i \to c_j) \times \phi(R(t), c_i, c_j) \tag{25}$$

where $\phi(\cdot)$ represents a regime-dependent modification function that captures how disaster conditions alter causal relationships, and $R(t)$ indicates the current disaster regime.

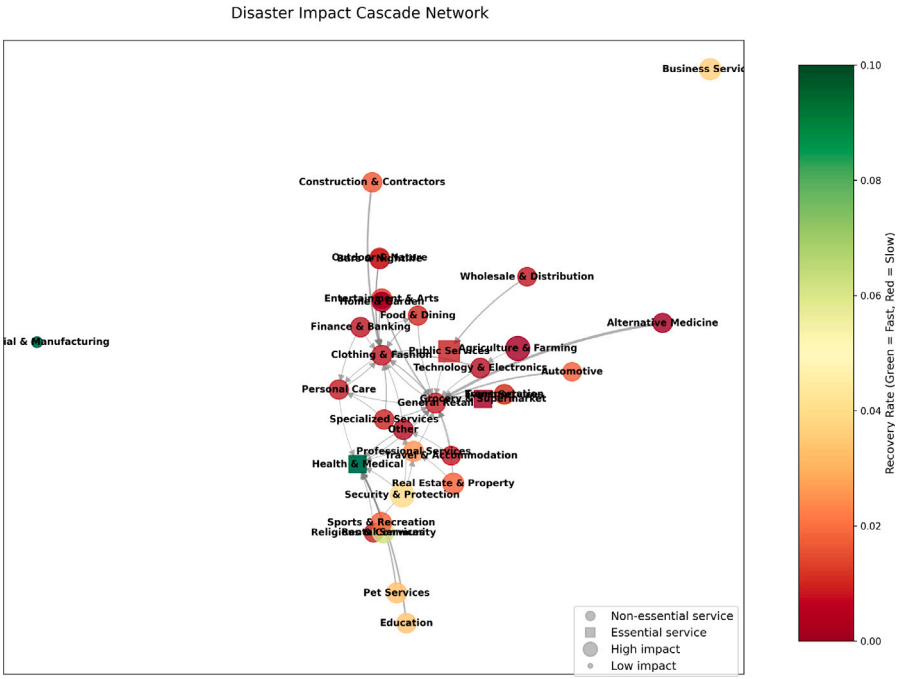

Fig. 11. Disaster impact cascade network. Node size indicates impact magnitude; colour indicates recovery rate (red = slow, green = fast); shape distinguishes essential (square) vs. non-essential (circle) services. Edge thickness shows cascade strength. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Analysis of temporal cascade evolution reveals that certain interdependencies strengthen during disasters whilst others weaken. Healthcare-to-pharmacy cascades intensify as medical visits trigger increased pharmaceutical needs, whilst entertainment-to-retail cascades diminish as discretionary co-visitation patterns break down. These dynamic changes in causal structure represent an entirely new dimension of disaster impact that traditional static analysis cannot capture.

This causal perspective represents a fundamental advance beyond correlation-based disaster analysis, providing the theoretical foundation and practical tools necessary for understanding how local disruptions create system-wide impacts through the complex web of interdependencies that characterise urban mobility networks.

4.3. Temporal analysis of recovery intervention strategies

The Disaster-aware Hawkes process not only enables us to identify which categories should be prioritised for recovery efforts based on network influence, but also reveals when interventions should be deployed for maximum effectiveness. This temporal dimension of resource allocation represents a critical but often overlooked aspect of disaster recovery planning that traditional analytical approaches cannot address.

To analyse optimal intervention timing, we developed a framework that simulates multiple intervention deployment strategies and evaluates their effectiveness in reducing system-wide impact over time. The analysis begins with an initial disruption to Agriculture & Farming, the category showing the highest direct impact from Cyclone Gabrielle (impact magnitude 0.61). We then simulate five distinct intervention timing strategies, each with a different underlying optimisation principle:

1. Immediate: All resources are deployed at time step 0, representing the conventional approach of simultaneous intervention across all affected categories. This strategy optimises for earliest possible response, operating under the assumption that earlier intervention is always better, regardless of category characteristics.
2. Staggered: Resources are deployed gradually over time steps 0–4, following a specific sequence determined by priority score and network position. This strategy optimises for sustainable deployment pace, addressing the highest-priority categories in the first three time steps before allocating to moderate-priority categories. The sequence is specifically tuned to address categories with high outgoing cascade effects early while managing those with slower recovery rates in the middle time steps.
3. Early Hubs: Categories are sequenced strictly by network centrality (PageRank values), with the most central nodes receiving intervention first. This strategy optimises for network influence, deploying resources to high-centrality categories (General Retail, Clothing & Fashion, Health & Medical) in the earliest time steps, under the theory that supporting central network nodes yields maximum cascade reduction through their extensive connections.

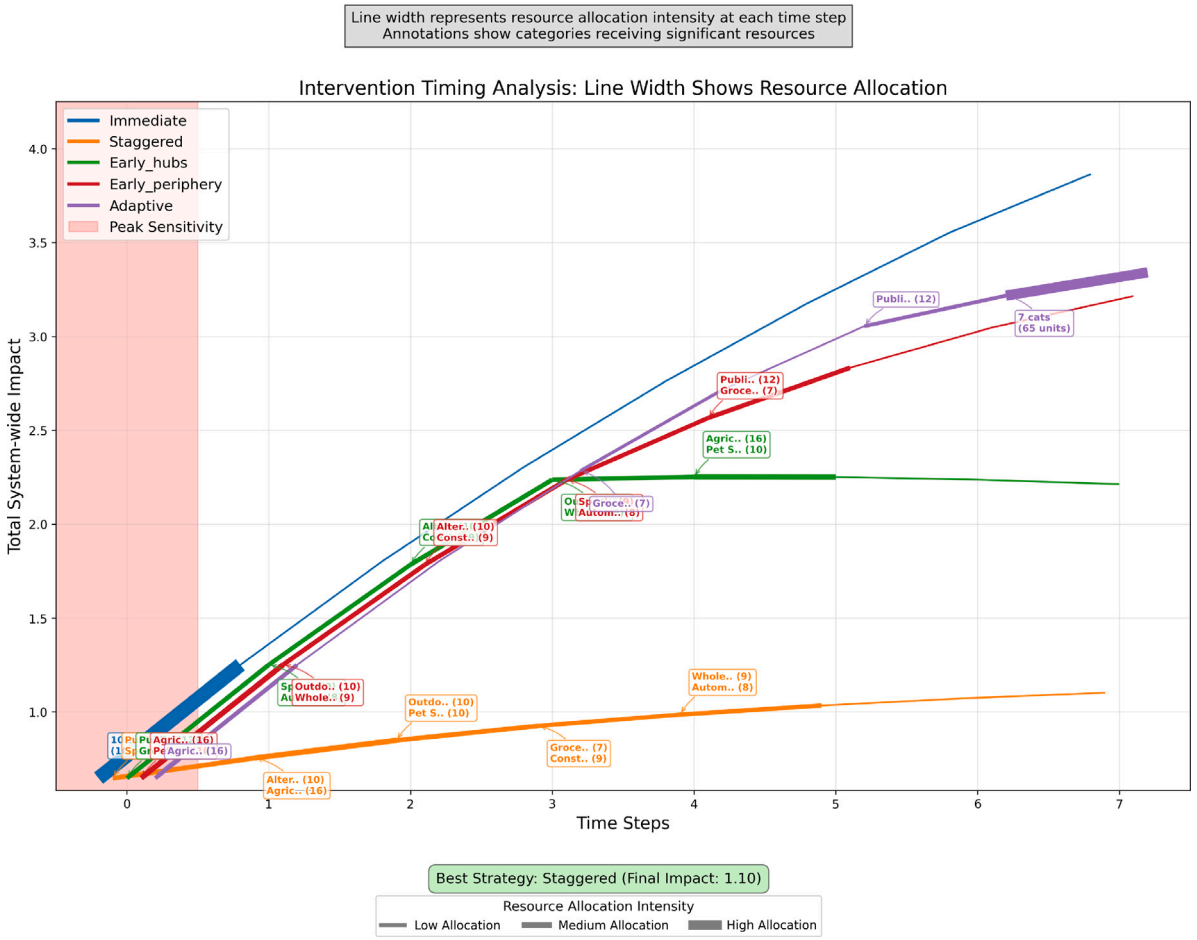

Fig. 12. Temporal intervention analysis with line thickness representing resource allocation intensity. Each line shows total system-wide impact over time for a different intervention strategy, with thickness proportional to resources deployed at each time step. Annotations identify key categories receiving significant resources. The red shaded area indicates the immediate time step after the disaster. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4. Early Periphery: The inverse of the hub strategy, this approach prioritises categories with lowest centrality first, addressing peripheral categories before they can transmit cascade effects to central nodes. The optimisation principle here focuses on stemming cascade propagation at the network periphery before it reaches the more interconnected core.
5. Adaptive: Each category receives intervention precisely one time step before it would experience significant impact in the baseline simulation. This strategy optimises for just-in-time intervention, using model predictions to forestall cascade effects immediately before they would otherwise take hold in each category.

For each strategy, we systematically vary the intervention timing while holding a total hypothetical resource allocation constant (100 units), distributing these resources according to the optimal allocation identified in our earlier cascade analysis. When an intervention is deployed to a category at a specific time step (e.g. 1 h), it reduces the outgoing cascade strength from that category proportionally to the allocated resources (each resource unit reduces cascade strength by 5%, with a maximum reduction of 80%). This models how recovery support diminishes a location's potential to transmit disruption to other parts of the network.

Integrated Analysis of Resource Allocation Intensity over Time Fig. 12 presents an integrated visualisation of our temporal intervention analysis, using variable line thickness to represent the intensity of resource allocation at each time step. Thicker line segments indicate higher resource concentration at that particular step, providing an intuitive visual representation of how different strategies distribute recovery resources over time.

The visualisation reveals several critical insights about the temporal dynamics of cascade effects and recovery interventions:

First, the distinctive thickness patterns clearly illustrate the fundamental differences in resource distribution philosophy across strategies. The Immediate approach (blue line) shows maximum thickness at time step 0, with all resources concentrated at the outset. In stark contrast, the Staggered strategy (orange line) displays a more balanced distribution, with moderate thickness across

time steps 0–4, indicating a measured deployment of resources over time. The Early_hubs approach (green line) shows a similar pattern to Staggered but with different emphasis on specific time steps, while Early_periphery (red line) and Adaptive (purple line) display more complex allocation patterns with significant resources deployed in later time steps.

Second, the significant divergence in performance demonstrates that intervention timing profoundly affects system recovery. The Staggered and Early_hubs strategies outperform the others substantially, keeping system-wide impact below 2.2, compared to approximately 3.8 for the Immediate strategy by time step 7. This counterintuitive result challenges the conventional wisdom that earlier intervention is always better, demonstrating that a carefully orchestrated sequence can achieve substantially better outcomes than deploying all resources simultaneously. The annotations reveal the key to the success of Staggered strategy is the priority of intervention at each time step, showing that the best-performing strategies focus early resources on different categories than might be expected from direct impact analysis alone. The Staggered strategy prioritises Alternative Medicine (11 units), Public Services (12 units), and Grocery & Supermarket (7 units) in the first three time steps, before addressing Agriculture & Farming (16 units) at step 4, despite the latter having the highest direct impact from the cyclone. This suggests that network position and cascade influence may sometimes outweigh direct impact in determining optimal intervention timing

Third, we observe that strategies begin to diverge significantly around time step 3, indicating that the initial cascade progression follows similar patterns regardless of intervention timing. However, beyond this point, the strategies produce markedly different trajectories, with the performance gap widening over time. This suggests that early intervention decisions establish path dependencies that increasingly constrain system behaviour as time progresses. The Early_periphery and Adaptive strategies, which distribute significant resources to later time steps based on different optimisation principles, perform only marginally better than the baseline Immediate approach. This suggests that while delayed intervention can be beneficial when strategically planned, the specific pattern of delay matters critically, as not all temporal distributions yield equally beneficial outcomes.

The performance differences between strategies are substantial and increase over time. By the final time step, the Staggered strategy (with a final impact of 2.05) outperforms the Immediate strategy (3.78) by 45.8%, despite using exactly the same total resources. This dramatic improvement arises purely from optimising the temporal distribution of interventions.

It is important to note that our temporal intervention analysis employs a simplified resource allocation model that abstracts away operational constraints such as personnel availability, equipment limitations, and budgetary restrictions. Whilst realistic disaster recovery planning must account for these practical considerations, our framework's primary contribution lies in demonstrating that temporal sequencing fundamentally affects recovery outcomes. The mathematical relationships revealed between intervention timing, cascade propagation, and system-wide recovery provide essential insights for future integration with operations research methods to develop comprehensive resource allocation frameworks. The discovery that carefully orchestrated temporal distribution can achieve 45.8% improvement over conventional simultaneous deployment establishes the critical importance of timing considerations that have been largely overlooked in disaster recovery planning literature.

4.4. Simulation of extended and multiple disaster scenarios

A key strength of the Disaster-aware Hawkes process is the ability to simulate various hypothetical disaster scenarios based on the original Cyclone Gabrielle parameters. This simulation approach allowed us to systematically analyse how changes in disaster duration, intensity, and sequencing affect both immediate visitation suppression and subsequent recovery trajectories. We simulated four distinct disaster scenarios:

1. Original disaster: The baseline scenario reflecting the actual Cyclone Gabrielle event (12–15 February 2023, or 4 days).
2. Double length: A disaster with twice the temporal length (12–19 February 2023, or 8 days) but maintaining the same intensity parameters.
3. Severe disaster: Original time period (12–19 February 2023) but doubled impact intensity.
4. Extended severe disaster: The final scenario combines both temporal and intensity modifications, featuring an eight-day duration with doubled impact intensity.

Our experimental design deliberately avoids modifying recovery-related parameters (recovery rates, bounce-back magnitudes, and decay coefficients) across all scenarios. This methodological choice ensures that observed differences in recovery patterns emerge from the model's learned understanding of how initial disaster characteristics influence subsequent dynamics, rather than from researcher-imposed assumptions about recovery processes. The approach allows us to test the model's capacity to predict how varying disaster characteristics naturally propagate through the urban mobility system according to the relationships encoded during the training process.

Fig. 13 presents the comparative analysis of our improved Hawkes process model across four distinct disaster scenarios, revealing several critical insights that demonstrate the model's capacity to capture complex urban mobility dynamics under varying stress conditions.

First, the model is capable of modelling different disaster scenarios without any artificial modification of recovery mechanisms. This demonstrates that our regime-switching framework successfully captures the fundamental relationship between disaster characteristics and their cascading effects on urban mobility systems. The contrast between the original disaster (blue solid line) and extended scenarios (dashed lines) reveals that the model naturally translates temporal and intensity variations into realistic mobility disruptions, suggesting that the learned baseline transitions between normal and disaster regimes encode essential urban resilience characteristics.

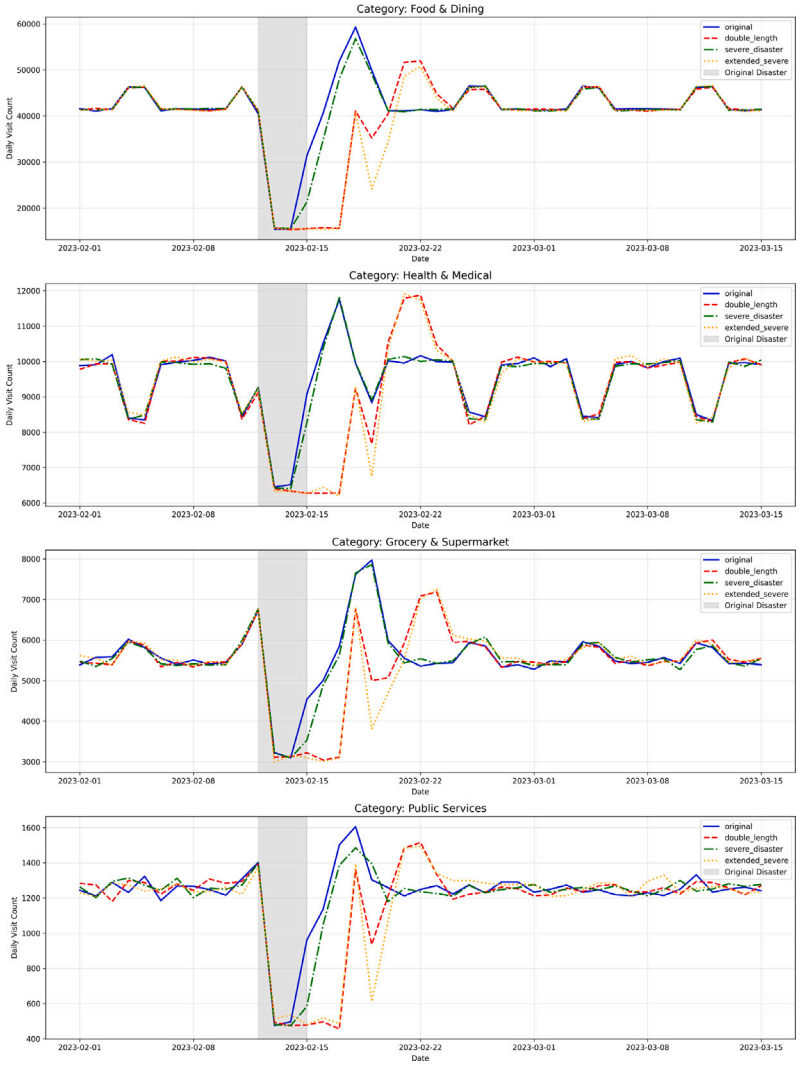

Fig. 13. Simulated visitation patterns under four disaster scenarios across Food & Dinning and 3 essential service categories. The grey shaded area indicates the period of the original disaster (12–14 February).

Second, the simulation results reveal sophisticated category-specific behaviours that validate our model's capacity to distinguish between different types of points of interest. Food and dining establishments exhibit the most pronounced bounce-back effects. This is consistent with pent-up demand following extended periods of mobility restriction. Conversely, health and medical services demonstrate more restrained but persistent recovery patterns, which aligns with the essential nature of healthcare services and their more steady demand patterns. This differentiation emerges entirely from the model's category-specific parameters, indicating that proposed Hawkes model successfully captures the underlying economic and social drivers of mobility demand across different service types.

Third, the simulation demonstrates the model's ability to generate realistic long-term recovery trajectories that extend well beyond the immediate disaster period. The gradual convergence of all scenarios towards similar baseline levels by mid-March, with their regular periodic patterns, indicates that our bounce-back mechanisms successfully model the temporal decay of disaster effects, whilst the intermediate oscillations reflect the model's excitation components capturing realistic behavioural feedback loops.

Fourth, the distinct recovery velocities observed across scenarios — with severe disasters showing more pronounced but also more variable recovery patterns — suggest that the model captures non-linear relationships between initial disruption magnitude and subsequent system dynamics. This behaviour emerges from the interaction between our regime-switching baseline intensities and the category-specific recovery parameters, demonstrating that the model framework enables complex emergent behaviours without requiring explicit programming of scenario-specific rules.

The model's ability to generate plausible alternative disaster trajectories using only learned parameters from historical data suggests significant potential for supporting evidence-based disaster preparedness and response planning. Urban planners and

emergency management officials can leverage such simulations to anticipate differential impacts across service categories and plan resource allocation strategies that account for the temporal evolution of mobility demand during and after disaster events.

5. Discussion and conclusion

This paper has introduced the Disaster-aware Hawkes process, a novel temporal point process framework designed specifically to model the complex dynamics of human mobility during and after disasters. By extending the standard Hawkes process with regime-switching baselines, category-specific excitation components, and post-disaster bounce-back effects, our approach addresses a significant methodological gap in disaster mobility analysis. In this section, we discuss the broader implications of our findings, acknowledge limitations, and suggest directions for future research.

5.1. Practical implications for disaster management

Our findings have direct implications for disaster preparedness, response, and recovery planning. First, the category-specific impact and recovery parameters provide a quantitative basis for prioritising resource allocation. The significant variation in recovery rates between categories from very rapid recovery in Public Services to slow recovery in Agriculture & Farming, suggests that one-size-fits-all recovery strategies will be suboptimal.

Second, our cascade impact analysis reveals critical dependencies in the mobility network that create vulnerabilities during disasters. The identification of categories with high influence but moderate direct impact (such as Alternative Medicine with the highest priority score of 0.412 despite a relatively low direct impact of 0.35) challenges conventional approaches that prioritise based solely on impact magnitude. Instead, our analysis suggests that network position can be more important than direct impact in determining a category's systemic importance.

Third, our temporal intervention analysis in Section 4 demonstrates that the timing of recovery interventions substantially affects system-wide outcomes. The finding that a staggered deployment strategy outperforms immediate intervention by 45.8%, despite using identical resources, offers a compelling case for temporally-optimised recovery planning. This insight challenges the conventional wisdom that earlier intervention is always better, revealing instead that carefully orchestrated sequencing based on network influence can achieve dramatically better results.

Fourth, the simulation of extended and sequential disaster scenarios in Section 4.4 provides valuable insights into system behaviour under more severe or complex disruptions. The observation that sequential disasters produce heightened bounce-back effects suggests possible accumulation of suppressed demand, with implications for resource planning in regions facing repeated disruptions such as seasonal disaster patterns.

5.2. Limitations and future research directions

While our approach addresses several key limitations of existing methods, it is important to acknowledge certain constraints and identify opportunities for further development.

First, our dataset, while comprehensive in many respects, has potential sampling biases. The higher representation of devices from less deprived areas (67.4% from low deprivation areas versus 32.6% from high deprivation areas) may skew our understanding of mobility patterns, particularly for vulnerable populations. Additionally, the exclusion of certain location types (such as religious and government facilities) creates blind spots in our network analysis.

Model validation using a single disaster event limits our ability to assess generalisability across disaster types and intensities. Whilst the framework's parametric structure enables recalibration for different contexts, empirical validation across multiple hazard types, geographical settings, and cultural contexts remains necessary to establish the bounds of transferability. Future research should prioritise multi-hazard validation to identify which model components transfer readily and which require substantial recalibration, ultimately developing disaster-type-specific parameter ranges to guide model application in new contexts.

Incorporating spatial dependencies into the modelling framework would further enhance our ability to analyse how mobility patterns evolve across both space and time. Future extensions incorporating explicit spatial dependencies could enhance the model's ability to capture highly localised effects, such as neighbourhood-scale evacuation patterns or infrastructure cascades following specific geographic damage patterns.

Further, stratifying analyses by socioeconomic groups could uncover disparities in disaster impacts and recovery, supporting the design of more equitable interventions. Finally, although our model captures temporal dependencies and cascade effects, establishing true causality in mobility patterns remains challenging. The excitation parameters in our model represent statistical associations that may not always reflect causal relationships. Further research incorporating controlled experiments or natural experiments could establish causal relationships found in our dataset.

CRediT authorship contribution statement

Minh Kieu: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Minh Kieu reports financial support and travel were provided by QuakeCoRE. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work has been funded by the UKRI (ESRC) funding ES/Y006259/1 under the Digital Footprints scheme, and QuakeCORE's Inter-disciplinary Programmes 3 - A Resilient Aoteroa New Zealand Transport System. We would like to thank the handling Editor and anonymous reviewers for their feedback and comments on earlier version of the manuscript.

Data availability

The data that has been used is confidential.

References

[1] Aref Darzi, Vanessa Frias-Martinez, Sepehr Ghader, Hannah Younes, Lei Zhang, Constructing evacuation evolution patterns and decisions using mobile device location data: A case study of hurricane irma, 2021, arXiv preprint arXiv:2102.12600.
[2] Valerie Washington, Seth Guikema, Joi-Lynn Mondisa, Aditi Misra, A data-driven method for identifying the locations of hurricane evacuations from mobile phone location data, Risk Anal. 44 (2) (2024) 390–407.
[3] Yoshihide Sekimoto, Akihito Sudo, Takehiro Kashiyama, Toshikazu Seto, Hideki Hayashi, Akinori Asahara, Hiroki Ishizuka, Satoshi Nishiyama, Real-time people movement estimation in large disasters from several kinds of mobile phone data, in: Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct, 2016, pp. 1426–1434.
[4] Xuan Song, Quanshi Zhang, Yoshihide Sekimoto, Ryosuke Shibasaki, Nicholas Jing Yuan, Xing Xie, Prediction and simulation of human mobility following natural disasters, ACM Trans. Intell. Syst. Technol. (TIST) 8 (2) (2016) 1–23.
[5] Hengfang Deng, Daniel P. Aldrich, Michael M. Danziger, Jianxi Gao, Nolan E. Phillips, Sean P Cornelius, Qi Ryan Wang, High-resolution human mobility data reveal race and wealth disparities in disaster evacuation patterns, Humanit. Soc. Sci. Commun. 8 (1) (2021) 1–8.
[6] Takahiro Yabe, Yunchang Zhang, Satish V. Ukkusuri, Quantifying the economic impact of disasters on businesses using human mobility data: a Bayesian causal inference approach, EPJ Data Sci. 9 (1) (2020) 36.
[7] Takahiro Yabe, Nicholas KW Jones, P Suresh C Rao, Marta C Gonzalez, Satish V Ukkusuri, Mobile phone location data for disasters: A review from natural hazards and epidemics, Comput. Environ. Urban Syst. 94 (2022) 101777.
[8] Boyeong Hong, Bartosz J. Bonczak, Arpit Gupta, Constantine E. Kontokosta, Measuring inequality in community resilience to natural disasters using large-scale mobility data, Nat. Commun. 12 (1) (2021) 1870.
[9] Natalia Boven, Nichola Shackleton, Liza Bolton, Andrew Sporle, Barry J. Milne, New Zealand Socio-Economic Index 2018, COMPASS Research Centre, University of Auckland, 2022.
[10] J. Ombler, J. Greene, J. Mateparae, R. Hunia, Report of the government inquiry into the response to the north island severe weather events, 2024.
[11] John Frank Charles Kingman, Poisson Processes, Vol. 3, Clarendon Press, 1992.
[12] Alan G. Hawkes, Spectra of some self-exciting and mutually exciting point processes, Biometrika 58 (1) (1971) 83–90.
[13] Yosihiko Ogata, Statistical models for earthquake occurrences and residual analysis for point processes, J. Amer. Statist. Assoc. 83 (401) (1988) 9–27.
[14] Qingyuan Zhao, Murat A Erdogdu, Hera Y He, Anand Rajaraman, Jure Leskovec, Seismic: A self-exciting point process model for predicting tweet popularity, in: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2015, pp. 1513–1522.
[15] Patrick J. Laub, Young Lee, Philip K. Pollett, Thomas Taimre, Hawkes models and their applications, Annu. Rev. Stat. Appl. 12 (2024).
[16] Daryl J. Daley, David Vere-Jones, An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods, Springer Science & Business Media, 2006.
[17] Richard H. Byrd, Peihuang Lu, Jorge Nocedal, Ciyou Zhu, A limited memory algorithm for bound constrained optimization, SIAM J. Sci. Comput. 16 (5) (1995) 1190–1208.