BMC Medical Informatics and Decision Making

RESEARCH

Open Access

From data to decisions: a modular platform for modelling and simulation of infectious disease diffusion in networks

Francesco Branda[1*†], Annamaria Defilippo[2], Ugo Lomoio[2], Barbara Puccio[2], Massimo Ciccozzi[1], Fabio Scarpa[3], Pierangelo Veltri[4] and Pietro Hiram Guzzi[2*†]

Abstract

Accurately modelling diffusion dynamics in complex networks is essential for improving medical outcomes, guiding pandemic preparedness, and optimizing resource allocation in public health. However, existing approaches often face a trade-off between predictive performance and model interpretability, limiting their utility for clinical decision-making and strategic planning. This study presents a modular computational methodology that integrates classical compartmental models with graph neural networks (GNNs) and explainable artificial intelligence (XAI) to simulate, analyse, and interpret the spread of contagion across heterogeneous network topologies. The approach captures both structural and temporal dimensions of diffusion processes, enabling granular insights into transmission pathways. Simulations are applied to critical public health scenarios, including the identification of super-spreaders and the assessment of targeted containment strategies. By combining mechanistic models with data-driven learning and explainability techniques, the methodology supports outcome forecasting, scenario comparison, and the interpretation of network-based risk factors. Results demonstrate the ability to predict diffusion trajectories with high accuracy while preserving transparency in decision-relevant variables. The approach is intended as a generalizable tool to support medical modelling and simulation with applications ranging from epidemic control to personalized risk assessment and cost-effective intervention planning.

Keywords Network simulation, Diffusion modelling, Graph theory, Explainable artificial intelligence, Stochastic processes

Introduction

Understanding, predicting, and controlling the spread of infectious diseases remains a critical challenge in epidemiology, public health policy, and computational modelling [1–4]. This is particularly relevant for highly contagious diseases such as measles, whose resurgence, despite the availability of a highly effective and safe vaccine, highlights the fragility of immunization programs and the need for improved decision-support tools [5, 6]. A growing body of evidence points to declining vaccination coverage worldwide, driven by a combination of misinformation, logistical barriers, vaccine hesitancy,

†Francesco Branda and Pietro Hiram Guzzi contributed equally to this work.

*Correspondence:
Francesco Branda
f.branda@unicampus.it
Pietro Hiram Guzzi
hguzzi@unicz.it
[1]Unit of Medical Statistics and Molecular Epidemiology, Università Campus Bio-Medico di Roma, Via Álvaro del Portillo, 21, 00128 Rome, RM, Italy
[2]Department of Surgical and Medical Sciences, Magna Graecia University of Catanzaro, Viale Europa, 88100 Catanzaro, CZ, Italy
[3]Department of Biomedical Sciences, University of Sassari, Viale San Pietro, 43, 07100 Sassari, SS, Italy
[4]DIMES, University of Calabria, via P. Bucci, 87036 Rende, CS, Italy

and the disruption of routine immunization services during the COVID-19 pandemic [7, 8]. As a result, measles, once considered close to eradication in many countries, has reemerged as a global threat, necessitating innovative approaches for outbreak prevention and control.

In this landscape, computational models have become indispensable tools for informing public health responses. Classical compartmental models, such as Susceptible-Infectious-Recovered (SIR) and Susceptible-Exposed-Infectious-Recovered (SEIR), have long been used to describe infectious disease dynamics at the population level [9]. However, these models often assume homogeneous mixing, ignoring the complex and heterogeneous nature of real-world social interactions that govern the spread of pathogens. As a response to this limitation, network-based approaches have gained prominence. These methods explicitly incorporate the structure of human interactions, modelling how diseases propagate through social, spatial, or behavioural connections [10, 11]. Such models can simulate not only global outbreak trajectories but also localized spread patterns, superspreader dynamics, and the differential impact of interventions across population subgroups [12–14].

The application of network-based modelling to measles is particularly compelling. Measles is characterized by a basic reproduction number (R_0) between 12 and 18, among the highest of any known infectious disease, which means that even small immunity gaps can trigger large-scale outbreaks [15]. Given this high transmissibility, accurate modelling of contact patterns and vaccination coverage is essential for risk forecasting and policy planning. Computational frameworks that integrate data from multiple sources, such as epidemiological databases, demographic surveys, and mobility records, enable the construction of context-specific contact networks. These networks provide the substrate on which simulated outbreaks can unfold, allowing researchers to explore alternative scenarios, assess the efficacy of control strategies, and identify critical nodes (e.g., superspreaders) that disproportionately influence disease propagation [6, 16, 17].

To support these needs, this work introduces a modular and extensible computational framework for modelling diffusion dynamics across complex networks. The proposed system unifies classical epidemic modelling, network simulation, and machine learning into a cohesive architecture designed to support scenario analysis and strategic intervention planning. Central to the framework is a pipeline that transforms epidemiological and demographic data into network representations, simulates disease transmission using compartmental models adapted to network structure, and applies Graph Neural Networks (GNNs) to analyse node-level diffusion behaviour. XAI techniques are used to interpret the learnt models, enabling insight into structural determinants of spread and supporting the identification of intervention targets, such as high-risk communities or superspreading individuals.

In the context of measles, the framework is applied to explore the interplay between vaccination coverage, network topology, and outbreak dynamics. Simulation experiments test the effectiveness of different containment strategies, such as blanket immunization, targeted vaccination, and isolation of key nodes, under realistic contact patterns. The integration of GNNs and XAI supports the classification of node risk and the explanation of diffusion trajectories, thereby enhancing both the predictive accuracy and interpretability of the model output. This hybrid approach provides a versatile platform for disease modelling that is not only scientifically rigorous but also accessible to public health professionals and decision-makers.

Overall, the framework offers a novel methodological contribution to the modelling and simulation landscape in biomedical research. By bridging traditional epidemiological theory with modern network science and machine learning, it enables deeper understanding of diffusion phenomena and practical planning tools for public health. The case study on measles exemplifies its potential impact: facilitating the design of proactive vaccination policies, identifying structural vulnerabilities in population networks, and supporting resource allocation in outbreak preparedness. As infectious disease threats continue to evolve, such integrative tools will play an increasingly role in guiding timely and effective public health interventions.

The paper is organized as follows: the *Background and Related Work* section provides an overview of existing approaches to epidemic diffusion modelling, and introduces the proposed computational framework. The *System Architecture* section describes the proposed method and its application in epidemiological contexts. The *Results* section presents the experimental evaluation performed on different network topologies. The *Discussion* analyses the implications of the findings for epidemic control strategies and methodological development. Finally, the *Conclusion* summarizes the contributions of the work and outlines possible directions for future research.

Background and related work

Control of epidemic spreading through simulation

The control of infectious disease outbreaks in highly connected and heterogeneous populations requires simulation frameworks that go beyond descriptive epidemiology and embrace predictive, explainable, and adaptive strategies. Traditional models such as SIR or SEIR offer analytical insights into aggregate-level dynamics, but fall short

when applied to structured populations where individual behaviour, network position, and stochastic interactions significantly affect diffusion patterns. Network-based simulation environments enable a finer-grained approach by incorporating real-world contact patterns, demographic data, and behavioural assumptions, thereby allowing the generation of realistic outbreak scenarios under various intervention hypotheses. Mathematical models have been employed to study the dynamics of epidemic transmission and provide a quantitative representation of disease spread. Among them, compartmental models represent a widely used theoretical framework for simulating how an infectious agent diffuses through a population. In these models, the population is subdivided into distinct classes (or compartments), and the homogeneous mixing hypothesis is assumed, meaning that each individual has the same probability of coming into contact with any other. Classic examples include the Susceptible-Infectious-Susceptible (SIS) model, which considers two groups, and the Susceptible-Infectious-Recovered (SIR) model [18], which introduced a three-compartment structure. Over time, extensions such as the Susceptible-Infectious-Recovered-Vaccinated (SIRV) model have been developed to incorporate additional states such as vaccination. Despite their relevance, a major limitation of classical compartmental models lies in their reliance on the homogeneous mixing assumption, which oversimplifies real-world interaction patterns by assuming that individuals mix uniformly within the population. Network-based simulations provide a realistic representation of heterogeneous social interactions, capturing variations in contact frequency and community structures [19]. These approaches integrate a diffusion process, typically from mathematical models, with a contact structure described by temporal networks. In this framework, nodes represent individuals and edges denote interactions through which transmission may occur [20, 21]. Each node has a state (e.g., susceptible, infected, recovered), allowing epidemic dynamics to reflect both the structure of interpersonal contacts and the probabilistic transitions defined by the model. Progression between states is determined by epidemiological parameters and effective contacts, which directly link the spread of infection to the network topology [22, 23].

Graph-based learning and model explainability

GNNs have become a powerful framework for learning from graph-structured data. Through iterative message-passing mechanisms, they enable each node to aggregate information from its neighbours, thereby capturing both local and global structural patterns in the graph [24]. Applying this framework to the epidemiological domain, GNNs can capture complex diffusion dynamics and provide realistic predictions of infection spread or intervention outcomes. Explainable AI (XAI) techniques have been increasingly integrated into graph-based learning, providing interpretability of nodes, connections, and substructures that most influence model predictions [25]. In epidemiology, this transparency helps uncover key transmission pathways, understand the role of highly connected individuals or "superspreaders," and assess the vulnerability of specific communities.

Methodology overview

The presented framework exploits a modular pipeline, that combines network reconstruction from epidemiological databases, modelling of spread through compartmental or stochastic processes and structural analysis using GNNs. Crucially, interpretability is maintained through the integration of XAI methods, which help to uncover the role of specific individuals, connections or substructures in promoting infection. This allows not only a retrospective analysis of the spread, but also a prospective evaluation of targeted interventions. By applying this framework to the case of measles, an archetypal vaccine-preventable and highly transmissible disease, we demonstrate its ability to inform decisions ranging from early diagnosis to the design of tailored public health strategies.

In the process of identifying and immunizing high-risk individuals, such as super-spreaders, through explainable GNN analysis integrated within the proposed simulation framework (Fig. 1). Initially, a contact network is generated and associated with a dynamic SIRVD (Susceptible-Infected-Recovered-Vaccinated-Deceased) model to simulate disease spread across the population. The learned dynamics are then analysed by a GNN, which encodes both structural and temporal features of the diffusion process. Through XAI techniques, the system highlights the most influential nodes, often corresponding to individuals with high centrality, clustering, or temporal reach. Once identified, these critical individuals are selectively vaccinated, and subsequent simulations show the effectiveness of targeted immunization in disrupting transmission pathways. The final network visualization demonstrates a significant reduction in active infections, underscoring the role of tailored interventions in achieving infection containment with minimal resource allocation.

In the case of non-pharmacological interventions, such as quarantine or travel restrictions, the framework simulates the targeted removal of connections in the network, an action that can greatly reduce the spread of infection (Fig. 2). The system begins with SIRVD-based simulations on a contact network, which are used to train a GNN. XAI techniques then identify key transmission pathways (high-risk connections) within the network. Strategic removal or temporary deactivation of selected

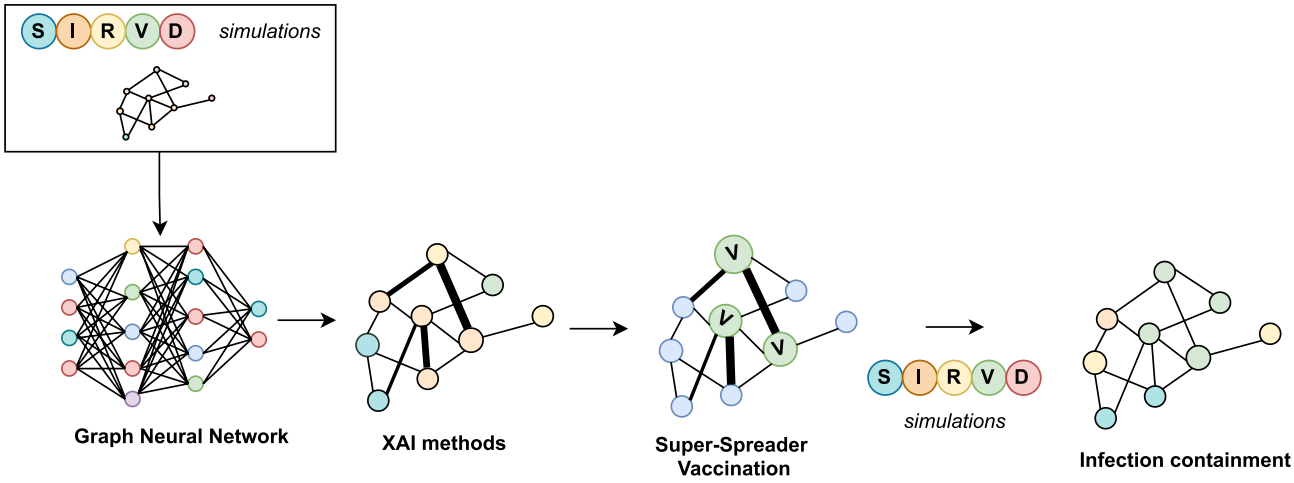

Fig. 1 Infection containment scheme using GNN analysis followed by XAI methods (super-spreader vaccination). After simulating graph structures using the SIRVD model diffusion, explainable GNN analysis is employed to select the most important edges (highlighted with thicker edges), based on which super-spreader nodes are identified and vaccinated (nodes labelled 'V' for 'vaccinated node'). Finally, the SIRVD simulation is applied to verify if the containment driven by the XAI method produces the expected infection containment

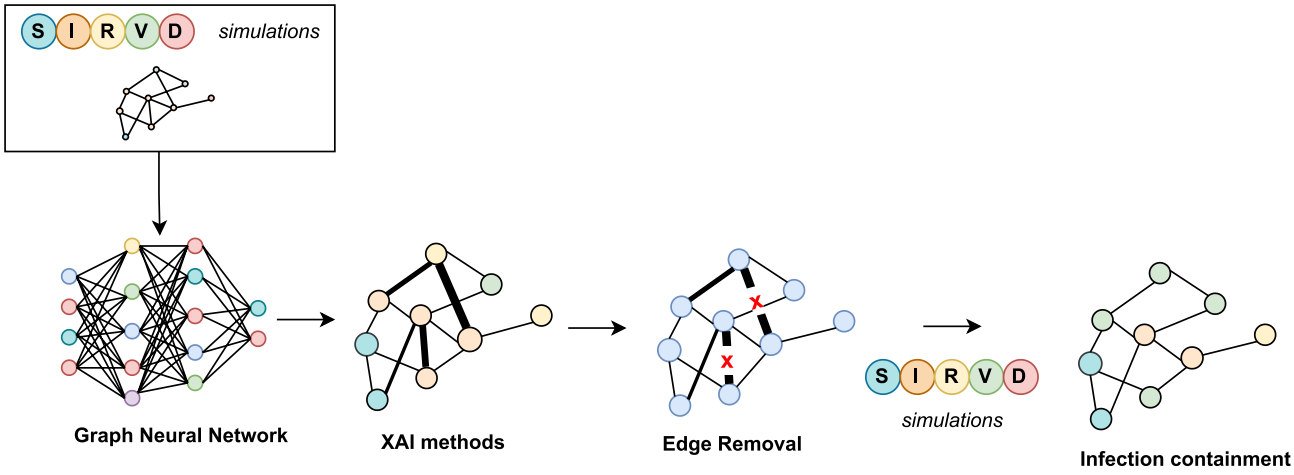

Fig. 2 Infection containment scheme using GNN analysis followed by XAI methods (edge-removal). After simulating graph structures using the SIRVD model diffusion, explainable GNN analysis is employed to select the most important edges (highlighted with thicker edges). These edges are then removed to simulate the targeted infection containment. A follow-up simulation demonstrates reduced spread and effective infection control, validating the impact of network-based intervention strategies

edges simulates non-pharmacological interventions such as quarantine, social distancing, or travel restrictions. Follow-up simulations show effective infection containment, demonstrating how edge-based interventions can suppress outbreak propagation when guided by data-driven insights.

Figure 3 shows the complete simulation-informed intervention pipeline designed to support data-driven epidemic control through the integration of network simulation, GNNs, and XAI. The upper section of the diagram depicts a sequence of SIRVD simulations modelling disease spread across a dynamic contact network over discrete time steps (t_1 to t_n). Each node represents an individual, and edges denote potential transmission routes. As the simulation progresses, nodes transition through epidemiological states (Susceptible, Infected, Recovered, Vaccinated, Deceased), reflecting how an outbreak unfolds over time. A specific simulation snapshot, selected at a critical point (e.g., high infection rate), is extracted for further analysis. The lower section of the diagram illustrates the GNN analysis, followed by an explanation of it with XAI-guided interventions.

System architecture

The pipeline described in Fig. 4 summarized the modular flow we implemented for the construction of an epidemiological database and its transformation into a computational infrastructure for the simulation, starting with the systematic collection of epidemiological bulletins released by governmental authorities and

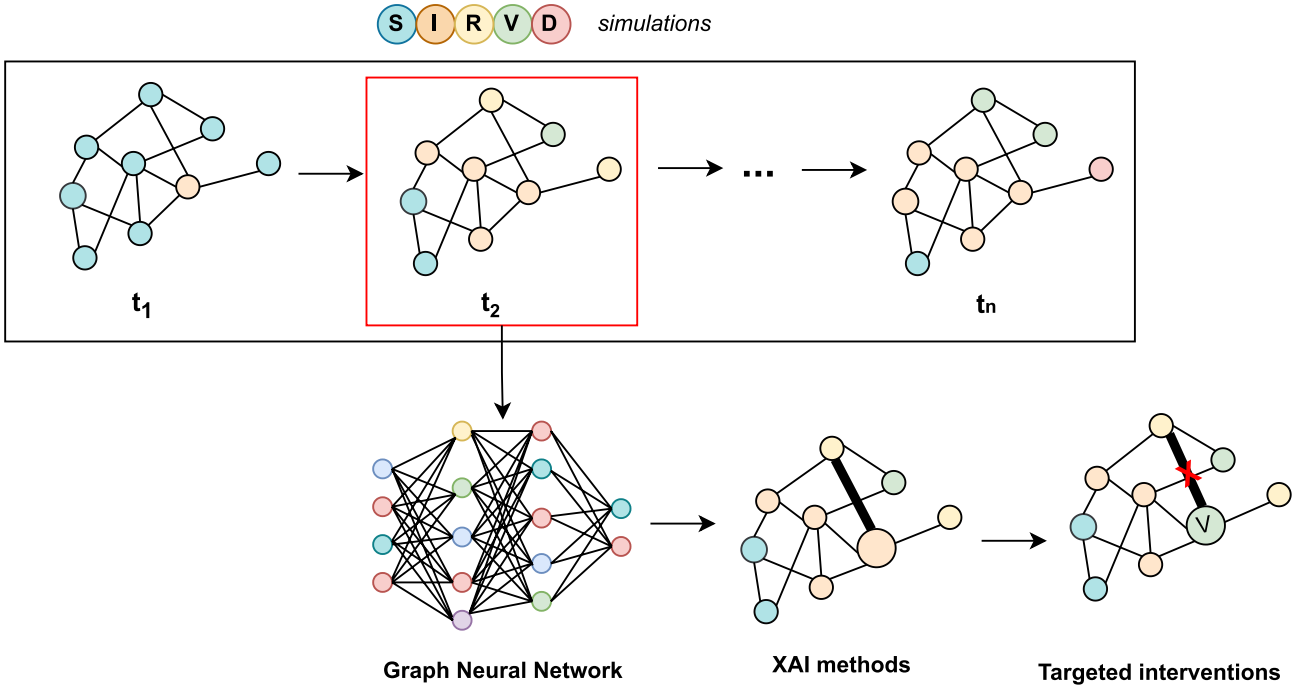

Fig. 3 General framework for simulation-informed intervention planning using GNNs and XAI. The top row illustrates the temporal progression of an epidemic over a contact network, simulated using a compartmental SIRVD model. Nodes change state over discrete time steps, reflecting infection dynamics. A representative snapshot (here t_2) is selected as input for analysis. The corresponding contact graph is processed by a GNN, which learns structural and temporal patterns of disease propagation. XAI methods are then used to interpret the GNN's predictions, identifying key nodes or edges critical to diffusion. This enables the design of targeted interventions, such as selective immunization or contact restriction. The final output is an adapted contact network incorporating the selected intervention strategy, which can then be re-simulated to assess containment outcomes. This iterative process integrates predictive modelling, explainability, and simulation to support dynamic, data-driven public health planning

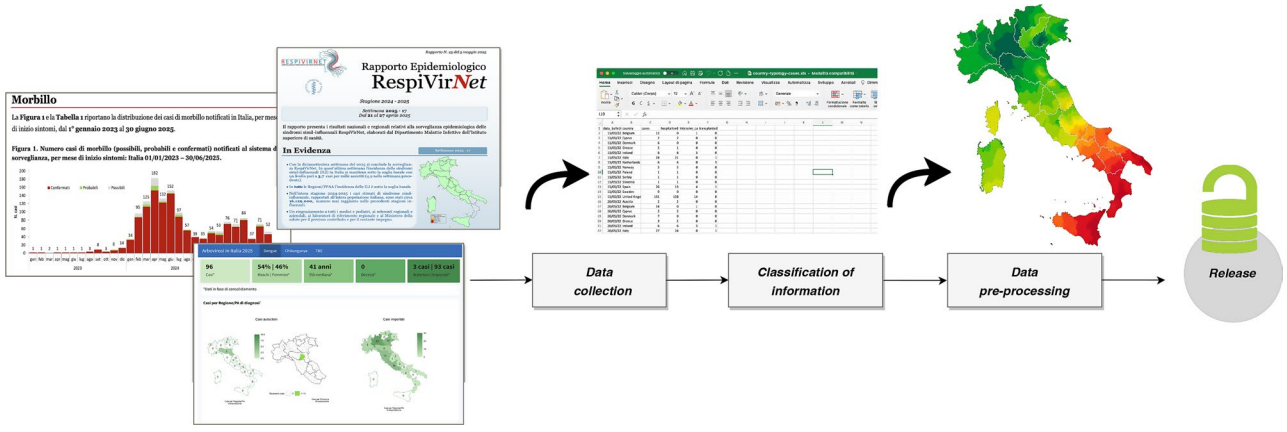

Fig. 4 Schematic representation of the data workflow in the proposed framework. The pipeline begins with the systematic acquisition of official epidemiological bulletins (from the Istituto Superiore di sanità via the EpiCentro platform), then digitally archived in a GitHub repository to ensure traceability and version control. Key epidemiological parameters are extracted through a classification phase and organized into structured excel files to support interoperability and downstream computational modelling

public health agencies, such as the periodic reports published by the Istituto Superiore di Sanità in Italy on the EpiCentro platform at https://www.epicentro.iss.it/, which are downloaded in digital format and stored in a GitHub repository to guarantee traceability and versioning. Raw datasets are obtained through a classification phase in which the main epidemiological parameters, such as number of cases, demographic and geographical distribution, are extracted and organised into structured files in Excel format to facilitate interoperability. Text data extraction from PDF bulletins was conducted using a hybrid automatic–manual approach designed to balance efficiency and accuracy, consistent with recent work on structured information extraction using large

language models for complex technical text [26]. For bulletins with a standardized tabular structure, an automatic workflow was employed, including the use of the GPT-4 (OpenAI) API orchestrated via LangChain [27] to transform unstructured data into structured JSON formats suitable for downstream processing. Automatic extraction was subjected to subsequent programmatic cross-validation and plausibility checks to mitigate errors (e.g., logical range checks, temporal consistency), a best practice in data quality assurance for structured and tabular datasets. Manual annotation was reserved for non-standardized bulletins, the presence of embedded graphics, ambiguous symbols (e.g., 'O' for zero cases), or cases where automatic confidence was low. This hybrid procedure ensured data integrity while maintaining high scalability.

In some cases, automatically extracted values may contain transcription errors, such as truncation (e.g., a regional cumulative total of 1205 cases extracted as 205) or misinterpretations of symbolic notation (e.g., letter 'O' interpreted as the digit '0'). To address such issues, a multi-level validation procedure was implemented, following established data quality assessment methodologies in epidemiological and health databases [28, 29]. Automated checks flagged suspicious records based on: (i) internal inconsistencies, such as a difference greater than 5% between the sum of sub-regional cases and the reported regional total, a commonly adopted threshold for detecting aggregation errors [28]; (ii) epidemiological plausibility violations, including demographic values outside admissible ranges (e.g., age outside 0–120 years) [30]; and (iii) recurring error patterns such as digit truncation or symbol confusion. All flagged records, together with a random sample of 10% of the automatically extracted data, were manually reviewed by comparison with the original PDF bulletin, in line with recommended practices for partial manual validation of large health datasets [31].

Quantitative values embedded in graphical elements such as line charts and histograms are digitized using external tools (e.g., WebPlotDigitizer, https://automeris.io/wpd/?v=5_2), ensuring that all available information is accurately captured in a computer-readable format. This is followed by the pre-processing phase, which comprised several structured steps. Comprehensive data cleaning was performed through (i) removal of duplicate records using unique identifiers and time-stamps; (ii) correction of transcription errors (e.g., truncated values or misrecognised "O/0" entries); (iii) harmonisation of categorical variables and diagnostic codes across bulletins; (iv) range validation for epidemiological plausibility (e.g., age restricted to 0–120 years, laboratory values checked against physiologic limits); (v) logical consistency checks (e.g., ensuring date of symptom onset did not follow hospitalization); and (vi) harmonisation of formats, units, and time zones (all dates aligned to ISO 8601, times corrected to UTC). Missingness analysis indicated an overall rate of 9% across the dataset, with variability between fields: demographic attributes (age, sex) showed $< 5\%$ missingness, clinical covariates (comorbidities, laboratory results) up to 18–20%, and geographical attributes $< 2\%$. We assumed data were Missing At Random (MAR), as missingness was related to observable patient factors (e.g., older patients more likely to lack complete laboratory panels). Variables with low missingness ($< 10\%$) were imputed by single-value strategies (median for continuous, mode for categorical). For higher rates, we used Multiple Imputation by Chained Equations (MICE), adopting predictive mean matching for continuous variables and logistic regression for categorical variables. Sensitivity analyses compared complete-case and imputed datasets to ensure stability of downstream estimates. Derived epidemiological metrics were computed using standard definitions. Incidence rates were calculated as the number of new cases meeting surveillance definitions (numerator) over the corresponding population at risk (denominator, stratified by age and sex and sourced from official statistics such as ISTAT/Eurostat). Rates were expressed per 100,000 individuals and directly age-standardised to the European Standard Population to facilitate cross-regional comparison. Temporal aggregation followed a weekly resolution, prioritising date of symptom onset when available, or laboratory confirmation date otherwise. Spatial aggregation was carried out at the NUTS-2 regional level, matching available denominator data. Confidence intervals for incidence rates were computed assuming Poisson-distributed counts. Additional derived indicators included basic reproduction number approximations (via growth rate methods) and vaccination coverage (doses administered/population eligible). During this step the data were also normalised and harmonised across different reporting formats, ensuring internal consistency and enabling integration with complementary sources such as mobility patterns [32] or climate datasets [33].

Once the pre-processing phase is completed, the curated information is aggregated into a structured epidemiological database, which serves as the main input for the modelling and simulation modules, illustrated in Fig. 5. This modular framework is designed to support modelling and simulation processes in the biomedical field, capturing the complexity of real-world diffusion phenomena. The system integrates structured data sources, network-based simulations, deep learning models and interpretability tools into a coherent and unified workflow, enabling researchers and policy makers to simulate complex dynamics, generate reliable predictions

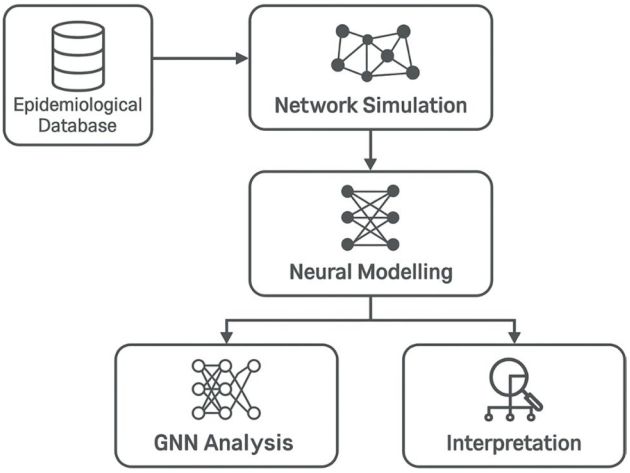

Fig. 5 Overview of the modular computational framework for diffusion modelling in complex networks. The pipeline consists of: (1) the epidemiological database, which contains real-world infection data; (2) the network Builder, responsible for constructing dynamic contact networks from demographic, spatial, or temporal attributes; (3) the diffusion simulator, integrating compartmental models and agent-based logic for scenario exploration; (4) the GNN-Based modelling module, which learns data-driven representations of diffusion using GNNs; and (5) the explainability layer, enabling the interpretation of model outputs and identification of key drivers such as superspreaders and critical edges

and interpret results to guide effective public health strategies.

Diffusion modelling and GNN explainability

The system is designed to accept inputs in standardized formats, making it interoperable with electronic health records and public health surveillance infrastructures. From these data, the framework constructs dynamic contact networks. The graphs represent social or spatial interactions between individuals or groups, with edges encoding the potential for disease transmission based on proximity, behaviour, or frequency of contact. The networks can be static or temporal and support the inclusion of weighted or probabilistic connections. This enables the modelling of diverse real-world settings such as school environments, urban communities, or cross-regional mobility networks. The construction logic is adaptable, allowing researchers to build population-specific or scenario-specific models to explore localized or large-scale epidemics.

Once the network is constructed, it serves as the substrate for simulating diffusion processes using compartmental models such as SIRV or SIRVD. These simulations operate over the network topology and can incorporate behavioural responses, vaccination strategies, and spatial mobility. The simulation engine is equipped to handle Monte Carlo experiments and parameter sweeps, supporting robust sensitivity analyses and uncertainty quantification. The ability to rapidly simulate various what-if scenarios makes the platform suitable for both prospective planning and retrospective policy evaluation.

To enhance predictive capabilities, the framework incorporates GNNs, which learn latent representations of nodes and edges based on their structural and temporal context. These models are trained to forecast future infections, classify nodes by risk category, or infer unobserved dynamics. Unlike classical machine learning models, GNNs are topology-aware, enabling them to exploit the relational structure of the data to generalize across unseen or evolving networks. This allows for more adaptive and granular predictions in situations where conditions change rapidly, such as during emerging outbreaks or policy shifts.

To identify superspreaders or critical interactions inside the network we trained a Graph Convolutional Network (GCN) [34] model to classify node states at each iteration step. The architecture comprises two graph convolutional layers followed by two fully connected layers. The first and second graph convolutional layers map the input node features to a latent representation of 32 hidden dimensions, each followed by a Rectified Linear Unit (ReLU) activation to introduce non-linearity. The resulting embeddings are subsequently processed through a fully connected layer of 16 units and a final linear classification layer, whose outputs are normalized via a softmax activation to yield class probabilities. Model optimization was carried out using the Adam algorithm with a learning rate of 0.002 over 1000 training epochs. To mitigate the effects of class imbalance, a weighted cross-entropy loss function was employed, where class weights were computed according to the empirical label distribution. Training was conducted in a full-batch manner on GPU when available, and model generalization was monitored through validation AUC scores computed at each epoch.

Recognizing the importance of explainability in clinical and policy decision-making, the system integrates XAI techniques. These include Saliency Maps and Integrated Gradients (IG), and can be extended to other interpretability techniques, such as feature attribution methods (e.g., SHAP and LIME) and counterfactual explanations. Outputs from the GNN models can thus be traced back to their structural and contextual drivers, making it possible to identify key individuals, communities, or edges responsible for propagation. In the context of infectious diseases like measles, this means identifying superspreaders, under-immunized clusters, or vulnerable demographic groups. These insights facilitate the development of targeted interventions such as focused immunization campaigns or localized containment measures.

The architecture has been tested on a case study focused on measles, a disease characterized by high infectivity and renewed global concern due to declining vaccination rates. By modelling scenarios with varying

immunization coverage, the framework supports the identification of outbreak drivers and the evaluation of containment strategies. In addition, our diffusion containment strategies were benchmarked against a baseline of four approaches: (i) no intervention, (ii) random vaccination of nodes, (iii) betweenness-based vaccination, and (iv) optimal percolation-based vaccination.

The inclusion of GNNs enables not only predictive modelling but also the identification of latent patterns in transmission pathways, offering deeper insight into how disease dynamics evolve in response to structural or behavioural factors.

Importantly, the modularity of the system allows for broad applicability beyond measles. The platform can be adapted to simulate other infectious diseases, antimicrobial resistance, behavioural health diffusion, or healthcare resource utilization. Its computational efficiency, interpretability, and integration of network-based reasoning make it a versatile tool for researchers, clinicians, and policymakers.As biomedical systems grow more interconnected and data-rich, tools that combine domain-specific modelling with machine learning and explainability are essential. This framework embodies that synthesis, offering a principled, extensible, and transparent solution for biomedical modelling and simulation. It is easily accessible via a Google Colab interface, which allows users to configure parameters and simulate customized case studies.

Results

To investigate how the structure of social interactions affects diffusion dynamics, our framework allows users to choose from a selection of synthetic network models, each capturing distinct topological features commonly observed in real-world populations. These models provide controlled environments to study how different contact patterns influence the spread and containment of infectious diseases. We begin by focusing on synthetic graphs, in line with established practices in epidemic modelling and network science, where synthetic data is used to validate methodologies under controlled and reproducible conditions. This strategy enables systematic variation of parameters and network structures before moving to the complexity of real-world health data. For example, Popper et al. [35] generated synthetic COVID-19 case data to evaluate and augment epidemiological reporting, by incorporating scenarios that cannot be captured through real-world observation; Lima et al. [36] simulated epidemic spread in a synthetic multi-agent community to test intervention strategies; and Renardy et al. [37] modelled tuberculosis transmission on synthetic populations to assess public health interventions. These works highlight how synthetic settings are commonly used as a necessary first stage in methodological validation. However, the simulation parameters used in our experiments were estimated using real data extracted from the Italian health system (measles virus, 2018–2019). Moreover, we extended our case studies to include both real contact networks and a synthetic population generated using the SynthPops tool (https://github.com/synthpops/synthpops), which simulates realistic contact patterns. These results complement our synthetic simulations and demonstrate that the proposed framework can be effectively applied to real-world healthcare scenarios, while preserving methodological transparency and reproducibility. In the following lines, the graphs employed in this study are presented, together with an overview of their main structural and statistical properties. This overview aims to provide an essential understanding of the datasets, including their size, topology, and node-level characteristics, which are relevant for interpreting the performance and behaviour of the proposed GNN model.

- Watts-Strogatz (WS) Model: Starting from a regular ring lattice where each node is linked to its k_n nearest neighbors, the WS model introduces randomness by rewiring edges with probability p. This produces small-world networks that combine short average path lengths with high clustering coefficients, features often present in social and biological systems.
- Barabási-Albert (BA) Model: Capturing the heterogeneity seen in many real-world systems, this model uses a preferential attachment rule where each new node connects to m existing nodes with probability proportional to their degree. The result is a scale-free network characterized by hub nodes and a power-law degree distribution, which influences vulnerability and resilience under epidemic scenarios.
- Random Geometric (RG) Graph: Nodes are positioned randomly in a continuous space (e.g., a unit square), and two nodes are connected if their Euclidean distance is less than a threshold r. The resulting network exhibits spatial locality and clustering, and its properties are influenced by the dimensionality of the embedding space, making it suitable for modelling geographically constrained interactions.
- Stochastic Block Model (SBM): Designed to simulate community structure, this model partitions the n nodes into k blocks, with edge probabilities defined by a $k \times k$ matrix. The parameters allow distinct levels of connectivity within and between communities, enabling the representation of modular populations and the study of intra- versus inter-group transmission.

Table 1 Network statistics and generation parameters for each graph model. The column nodes indicates the total number of vertices in the graph, while *edges* reports the number of links between them. The *clustering coefficient* measures the tendency of nodes to form tightly connected groups, and*Shortest path length* represents the average minimum number of steps needed to travel between two nodes. The column *parameters* specifies the generative settings for each model: (**i**) Barabási-Albert: the number *m* of edges each new node introduces; (ii) Watts–Strogatz: the initial number of neighbors k_n, per node and the rewiring probability *p*; (iii) Random Geometric: the connection radius *r* and the embedding dimension *dim*; (iv) Stochastic block: the number of communities *k*, the intra-community connection probability d_{in}, and the inter-community connection probability d_{out}; (**v**) synthpops: the number of nodes *n* of the synthetic population and the *location* data to use (default is Seattle metro, Washington, USA). (vi) real contact map: no parameters are needed

Network model	Nodes	Edges	Clustering Coefficient	Shortest Path Length	Parameters
Barabási–Albert	1000	3984	0.0334	3.1896	m = 4
Watts–Strogatz	1000	5000	0.5599	5.0332	$k_n = 10$, $p = 0.06$
Random Geometric	1000	31103	0.6346	4.1545	$r = 0.15$, $dim = 2$
Stochastic Block	1000	26220	0.091	2.0856	$k = 4$, $d_{in} = 0.15$, $d_{out} = 0.02$
Synthpops	964	7318	0.3593	G is not connected	$n = 1000$, location = Seattle, Washington
Real [38]	80	139	0.4914	G is not connected	-

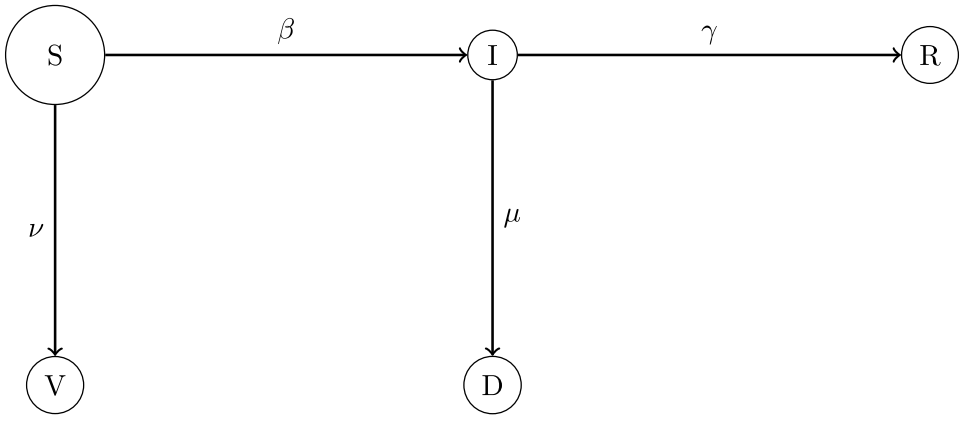

Fig. 6 Schematic representation of the SIRVD model (S = susceptible, I = infected, *R* = recovered, V = vaccinated, D = Deceased). Transitions include infection (β), recovery (γ), mortality (μ), and vaccination (ν)

- SynthPops: designed to generate synthetic populations for analysing COVID-19 (SARS-CoV-2) epidemics. It can create generic populations with different network characteristics, as well as synthetic populations that interact across different layers of a multilayer contact network. In this context, nodes represent individuals and have attributes such as age, while edges represent interactions between people, characterised by the setting in which these interactions occur (for example, household, school, or work).
- **Real contact map from** [38]: To prove the usability of our methodology in a real-world scenario, the *Contact patterns in a village in rural Malawi* dataset was used to compute the contact graph used in the spreading simulations.

Table 1 summarises the main network parameters.

SIRVD model A SIRVD model has been used to simulate the diffusion processes [20]. The SIRVD is an extension of traditional compartmental models used in epidemiology to simulate the spread of infectious diseases. It partitions the population into five compartments: Susceptible (S), Infectious (I), Recovered (R), Vaccinated (V), and Deceased (D). Transitions, as illustrated in Fig. 6, between these compartments are governed by a set of differential equations, each reflecting key epidemiological processes. Susceptible individuals (S) may become infected at a transmission rate β upon contact with infectious individuals (I). Infected individuals can either recover at rate γ, die from the disease at rate μ, or transition to the recovered (R) or deceased (D) compartments accordingly.

To model immunization effects, a vaccination rate ν transfers individuals from the susceptible group (S) to the vaccinated group (V), which is assumed to be fully immune to infection. Deceased individuals (D) no longer contribute to the transmission dynamics but are essential for quantifying the disease burden.

Vaccination is modelled through two mechanisms: a random vaccination probability ($\nu_H^r = 0.05$) and a targeted vaccination probability ($\nu_H^t = 0.20$), while in the no-vaccination scenario this probability is set to zero ($\nu_H^{no} = 0.00$). The optimal percolation vaccination

strategy was implemented with the same overall vaccination rate as the targeted strategy ($\nu_H^t = 0.20$), ensuring a consistent level of immunization across all comparative scenarios. The natural human mortality rate (μ_H) was assumed to be negligible and set to 0.0000. For vaccinated individuals, the probability of infection was set to zero ($\beta_{HV} = 0.0000$), reflecting full protection. For the SIRVD simulations, we used two years (2018–2019) of Italian measles surveillance data and obtained the following calibrated parameters: $\beta = 0.8394$, $\gamma = 0.4419$, and $\mu = 0.0010$. Each configuration was simulated five times ($N_{sims} = 5$) to ensure robustness and reproducibility of the results.

The inclusion of vaccination and mortality enhances the model's realism and applicability, allowing it to simulate various intervention strategies and public health policies.

The SIRVD model can be further enriched by implementing it on contact networks, where individuals are connected through realistic interaction structures, enabling fine-grained assessments of disease spread, superspreader identification, and tailored containment measures.

For the comprehensive SIRVD simulations, two main objectives are addressed: (I) assessing the effectiveness of *targeted vaccination strategies* informed by model explainability and (II) and disrupting *super-spreader* pathways to reduce infection. In the first case a XAI-driven vaccination protocol is implemented by determining node level influence. Edge importance scores are aggregated per node, and the top K% of influential nodes (e.g., top 40%) are identified and vaccinated with the same probability used in the baseline strategies. A new run is executed under identical conditions, and the results of XAI-driver vaccination is benchmarked against the four initial alternatives: no vaccination, random vaccination, betweenness and percolation centrality-based targeted vaccination. In the second scenario, best vaccination strategies is selected for new simulation. Then, XAI-based edge removal strategies is perfomed to remove most influent edges. In both case, the evolution of infected ("Ih") and deceased ("Dh") populations is tracked over time.

To prove the robustness of our methodology in highlight important interactions or super-spreader nodes, we repeated each simulation 5 times and reported results in terms of mean values and their standard deviation. For each graph, the complete set of vaccination strategies is simulated, running for 24 steps, each corresponding to a month; resulting GNN embeddings at a randomly selected time step are visualised; and XAI-guided interventions via edge removal and targeted vaccination are applied.

By removing 70% of the most important edges and vaccinating 30% of the total number of nodes among the ones identified by XAI (using the same vaccination probability as other strategies) in two independent scenarios, our framework proposed an effective strategy to reduce the infection peak and total deaths in XAI-guided vaccination applied to each network in SIRVD simulations.

Table 2 SIRVD XAI-based interventions in WS network. The column strategy lists the intervention method applied. *peak infected* indicates the maximum percentage of the population simultaneously infected during the outbreak. *total Dead* reports the overall mortality rate at the end of the simulation. *Duration* denotes the total number of time steps (simulation length), where each step is a month. Simulations were repeated 5 times to report mean and standard deviation values

Strategy	Peak Infected	Total Dead	Duration
No Vaccination	62.94% ± 0.18	0.32% ± 0.00091	16
Random Vaccination	61.38% ± 0.18	0.20% ± 0.00064	18
Targeted Vaccination	57.26% ± 0.17	0.22% ± 0.00071	20
Optimal Percolation Vaccination	59.78% ± 0.17	0.14% ± 0.00043	18
XAI Edge Removal (70%)	54.44% ± 0.15	0.18% ± 0.00052	20
XAI Targeted Vaccination (30%)	7.58% ± 0.025	0.04% ± 0.00019	19

Watts–Strogatz network

Table 2 reports the results obtained for the WS network, comparing them to the baseline strategies with the XAI-guided interventions. Simulations were repeated 5 times, allowing for the reporting of each value in terms of mean and standard deviation. Further details are illustrated in Figs. 7 and 8.

Figure 7 illustrates the application of GNNs and XAI in understanding the diffusion dynamics within the WS network. Figure 7 (a) shows XAI applied to simulations without vaccination, which is used to establish XAI-guided vaccination. In contrast, Fig. 7 (b) displays XAI applied to simulations with a betweenness-based vaccination strategy, identified as the best approach in the baseline, used to set-up XAI-guided edge removal.

IG is applied to demonstrate the mutual influence of nodes on the infection process. By highlighting the most influential edges, it becomes evident that certain individuals, due to their position or centrality in the network, contribute disproportionately to the infection's spread. These nodes, often referred to as super-spreaders, are critical in shaping the epidemic trajectory, and targeting them for interventions, such as vaccination or isolation, could significantly reduce the overall transmission.

Figure 8 illustrates the impact of various epidemiological intervention strategies on a WS network over a period of 24 months, highlighting the infection rates. Specifically, Fig. 8 (a) examines the effect of XAI-based

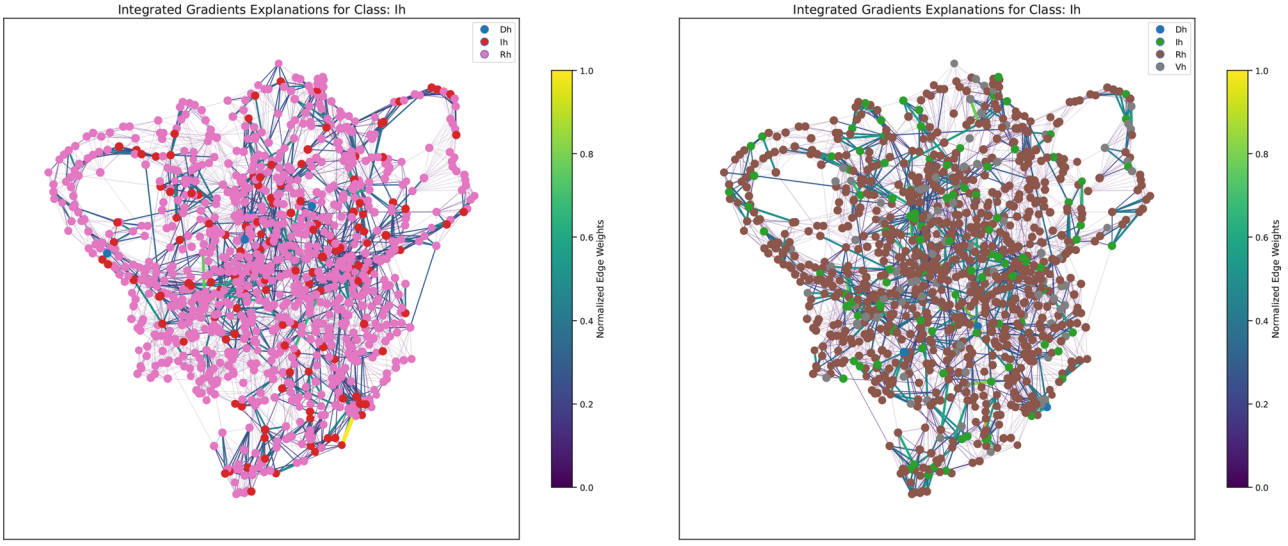

(a) IG based on no vaccination strategy (Ih class).

(b) IG based on targeted vaccination strategy (Ih class).

Fig. 7 IG insights for the WS network in the SIRVD model setting. This visualization illustrates the node-level relevance scores computed using the IG method, applied to a GNN trained to predict epidemic dynamics over a small-world topology. The WS model, known for its high clustering and short path lengths, reflects realistic patterns of social interaction. The color intensity of each edge corresponds to its attributed influence on the model's infection prediction at a selected time step. Nodes with higher integrated gradient values are identified as critical contributors to the spread of the infection. Node colors correspond to the simulation class and are indicated by abbreviations for each individual's class (sh = Susceptible humans, ih = infected humans, Rh = Recovered humans, vh = vaccinated humans)

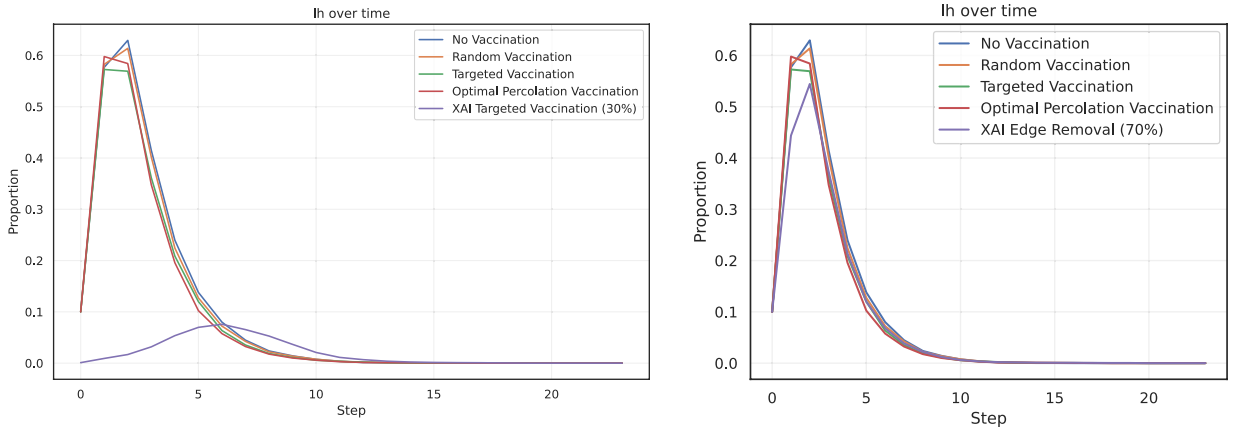

(a) XAI-based vaccination infection rate

(b) XAI-based edge removal infection rate

Fig. 8 Effect of XAI-based targeted vaccination (top 30%) and XAI-guided edge removal (70%) on infected humans(ih) over time on WS network. The mean values of all simulations (repeated five times) are compared here, with each case (baseline methods and XAI-guided intervention) distinguished by the associated colour in the legend

targeted vaccination, using the same vaccination probability as the baseline strategies to vaccinate the top 30% of all nodes identified by the XAI method as the most important. The infection rate plot demonstrates that the XAI-based vaccination is much more effective than other strategies, with significantly reduced epidemic peaks. On the other hand, Fig. 8 (b) compares XAI-based removal (top 70%) (quarantine-like case) with the baseline of four approaches, demonstrating the validity of this strategy, even though XAI-targeted vaccination has shown the best results in reducing the infection peak, as detailed in Table 2.

Barabási-Albert network

As illustrated in Table 3, both XAI-based strategies, improved the management of the spread process in the case of the BA network as well, especially in the case of XAI-guided vaccination.

In particular, the vaccination of the 30% most influential nodes drastically reduced the peak of infection

Table 3 SIRVD XAI-based interventions in BA network. The column *strategy* lists the intervention method applied. *peak infected* indicates the maximum percentage of the population simultaneously infected during the outbreak. *total Dead* reports the overall mortality rate at the end of the simulation. *Duration* denotes the total number of time steps (simulation length), where each step is a month. Simulations were repeated 5 times to report mean and standard deviation values

Strategy	Peak Infected	Total Dead	Duration
No Vaccination	74.44% ± 0.18	0.12% ± 0.00042	24
Random Vaccination	72.36% ± 0.18	0.30% ± 0.00092	17
Targeted Vaccination	63.46% ± 0.16	0.22% ± 0.00066	17
Optimal Percolation Vaccination	64.12% ± 0.16	0.28% ± 0.00087	20
XAI Edge Removal (70%)	63.34% ± 0.17	0.32% ± 0.00093	17
XAI Targeted Vaccination (30%)	30.62% ± 0.084	0.10% ± 0.00037	16

to only 30.62%, with a significant decrease in mortality to 0.084%, ending the epidemic in only 16 time steps. In comparison, the removal of the 70% most influential edges did not produce such a great result. However, it reduced the infection peak to 63.34% (slightly lower than the baseline).

Figures 9 and 10 provide a comprehensive overview of the simulation results and the integration of GNNs within the BA network setting. Figure 9 shows the contribution of nodes to the spread process, highlighting how nodes with the greatest impact on spread are ideal candidates for XAI-based interventions (vaccination or containment). As in the previous case, Fig. 9 (a) corresponds to XAI for simulations without vaccination, while Fig. 9 (b) corresponds to XAI applied to simulations with a betweenness-based vaccination strategy. Edge importance scores were computed using attribution techniques derived from the GNN model (IG in this case), and the most epidemiologically critical connections were selectively removed before simulation.

Figures 10 (a) and (b) compare the evolution of the infected population (Ih) across the simulation timeline following a targeted structural intervention. This approach aims to disrupt transmission pathways in a topologically informed manner, particularly targeting hub-to-hub connections characteristic of scale-free networks. In particular, Fig. 10 illustrates the appreciable changes in the number of infected, and the effective control of deaths, which is almost negligible with XAI-targeted strategies (0.084%), as stated earlier. The observed reduction in infection peak and acceleration of epidemic termination suggest that strategically XAI-guided vaccination can suppress contagion more effectively than baseline methods. The result highlights the utility of XAI in identifying non-obvious intervention points and supports the integration of explainability into simulation-based public health planning.

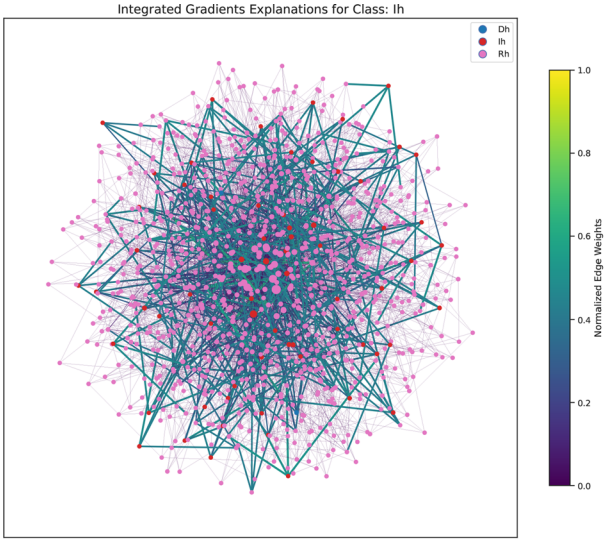

(a) IG based on no vaccination strategy (Ih class).

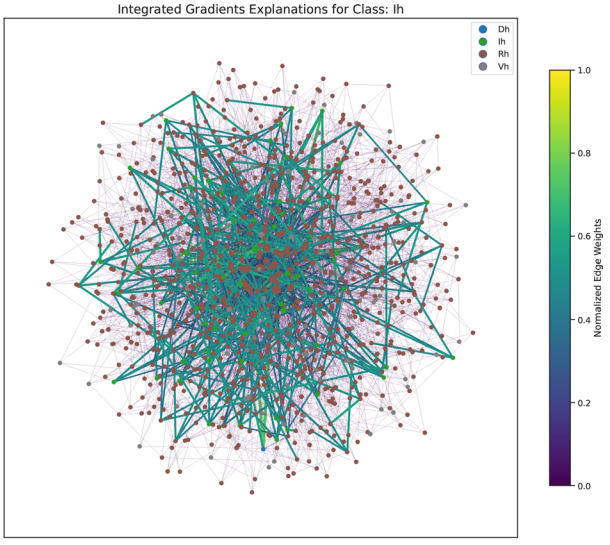

(b) IG based on targeted vaccination strategy (Ih class).

Fig. 9 IG insights for the BA network in the SIRVD model setting. This visualization illustrates the node-level relevance scores computed using the IG method, applied to a GNN trained to predict epidemic dynamics over a small-world topology. The WS model, known for its high clustering and short path lengths, reflects realistic patterns of social interaction. The color intensity of each edge corresponds to its attributed influence on the model's infection prediction at a selected time step. Nodes with higher integrated gradient values are identified as critical contributors to the spread of the infection. Node colors correspond to the simulation class and are indicated by abbreviations for each individual's class (sh = Susceptible humans, ih = infected humans, Rh = Recovered humans, vh = vaccinated humans)

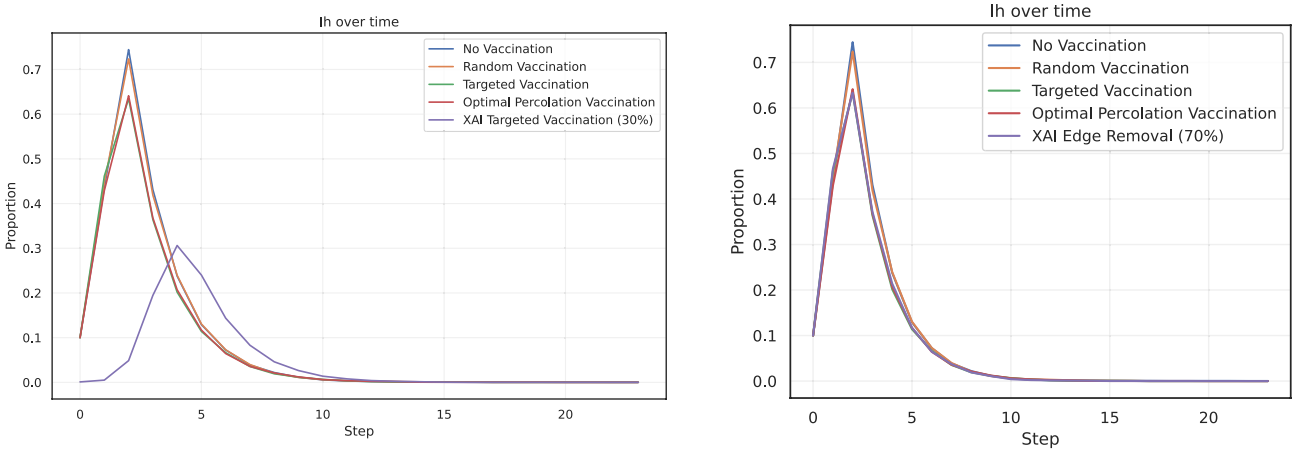

(a) XAI-based vaccination infection rate

(b) XAI-based edge removal infection rate

Fig. 10 Effect of XAI-based targeted vaccination (top 30%) and XAI-guided edge removal (70%) on infected humans(ih) over time on BA network. The mean values of all the simulations (repeated 5 times) are compared here, distinguishing each case (baseline methods and XAI-guided intervention) by the associated colour in the legend

Table 4 SIRVD XAI-based interventions in RG network. The column *strategy* lists the intervention method applied. *peak infected* indicates the maximum percentage of the population simultaneously infected during the outbreak. *total Dead* reports the overall mortality rate at the end of the simulation. *Duration* denotes the total number of time steps (simulation length), where each step is a month. Simulations were repeated 5 times to report mean and standard deviation values

Strategy	Peak Infected	Total Dead	Duration
No Vaccination	95.30% ± 0.22	0.22% ± 0.00061	17
Random Vaccination	94.98% ± 0.22	0.28% ± 0.0008	15
Targeted Vaccination	95.10% ± 0.22	0.20% ± 0.00058	15
Optimal Percolation Vaccination	94.38% ± 0.21	0.20% ± 0.00055	17
XAI Edge Removal (70%)	94.36% ± 0.21	0.16% ± 0.00043	14
XAI Targeted Vaccination (30%)	22.36% ± 0.074	0.10% ± 0.00038	18

Random geometric network

Table 4 details the XAI-guided interventions on the RG network in SIRVD simulations. In RG graphs, high clustering and spatial proximity cause infections to spread rapidly within dense local groups, producing sharp peaks. Since the degree distribution is fairly uniform, random or degree-based vaccination is less effective, leaving the epidemic curve with higher infection peaks.

Figure 11 shows the IG insights for the RG network. This visualisation highlights the node-level attributions derived from GNNs trained to predict infection dynamics in the two settings just described.

As illustrated in Fig. 12, the diffusion of infection, even if not evident as in the previous case studies, is reduced with both XAI strategies (94.36% with XAI edge removal and 22.36% with XAI vaccination); additionally, the total number of deaths is also slightly reduced.

Stochastic block network

The last synthetic simulation is made by considering a structure of network in which communities (i.e. subregions with a large number of contacts) are presents. Table 5 shows the the results on SBM network, comparing baseline simulations and XAI-based interventions.

Figure 13 illustrates node-level importance scores computed using IG-based explainability techniques applied to the GNN trained on the disease diffusion dynamics.

Figure 14 (a) shows the effect of XAI-based targeted vaccination (top 30%) on the evolution of infected humans (Ih) over time in a SBM network under the SIRVD model. Finally, Fig. 14 (b) shows the results of XAI-based edge removal (top 70%) applied on infected humans (Ih) over time on SBM network. In particular panel (b) shows the reduction of deaths in the presence of targeted vaccination and XAI based interventions.

Synthpops

{Table 6 details the XAI-guided interventions on the Synthpops Seattle, Washington network in SIRVD simulations.

Figure 15 shows the node embeddings extracted by the GNN considering separately the no vaccination scenario and the optimal percolation-guided vaccination strategy.

Figures 16 and 17 illustrate the progression of infection under the five different immunisation scenarios and the explanations. They report how removing the 70% of important edges (5122 edges removed) from the network (approximately a quarantine of the whole community) didn't have important effects on the infection peak, obtaining only a 1% of peak reduction compared to optimal percolation vaccination. On the other hand, using XAI-guided vaccination (30%, 289 vaccinated nodes)

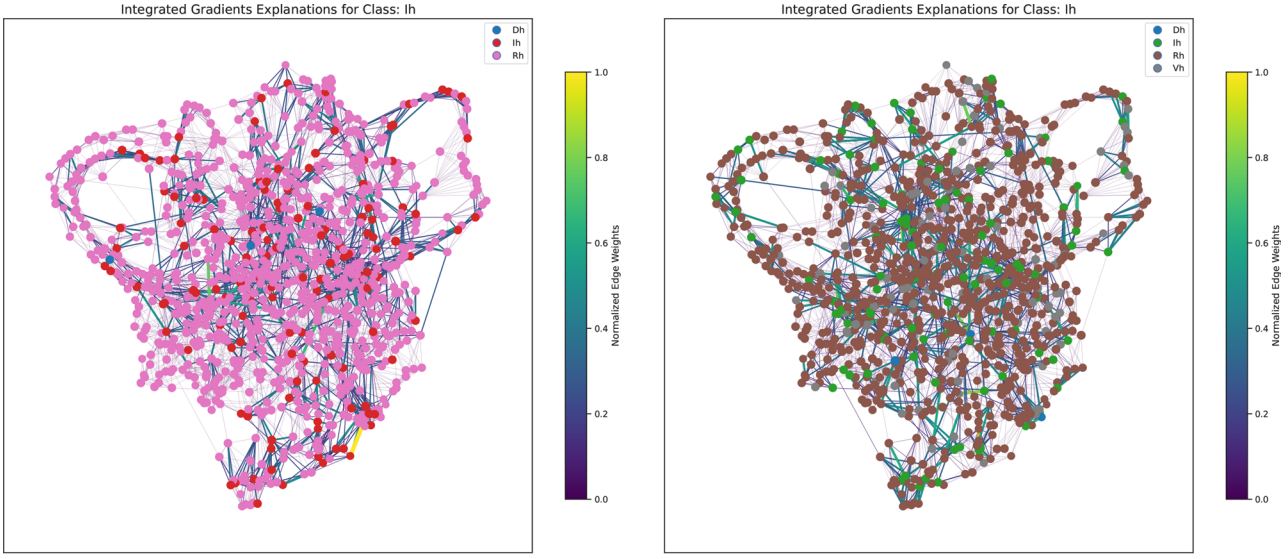

(a) IG based on no vaccination strategy (Ih class).

(b) IG based on targeted vaccination strategy (Ih class).

Fig. 11 IG insights for the RG network in the SIRVD model setting. This visualization illustrates the node-level relevance scores computed using the IG method, applied to a GNN trained to predict epidemic dynamics over a mathematical spatial network. In spatially constrained networks like RG, where connectivity is based on physical proximity, IG analysis provides a fine-grained understanding of local transmission patterns. The color intensity of each edge corresponds to its attributed influence on the model's infection prediction at a selected time step. Nodes with higher integrated gradient values are identified as critical contributors to the spread of the infection. Node colors correspond to the simulation class and are indicated by abbreviations for each individual's class (sh = Susceptible humans, ih = infected humans, Rh = Recovered humans, vh = vaccinated humans)

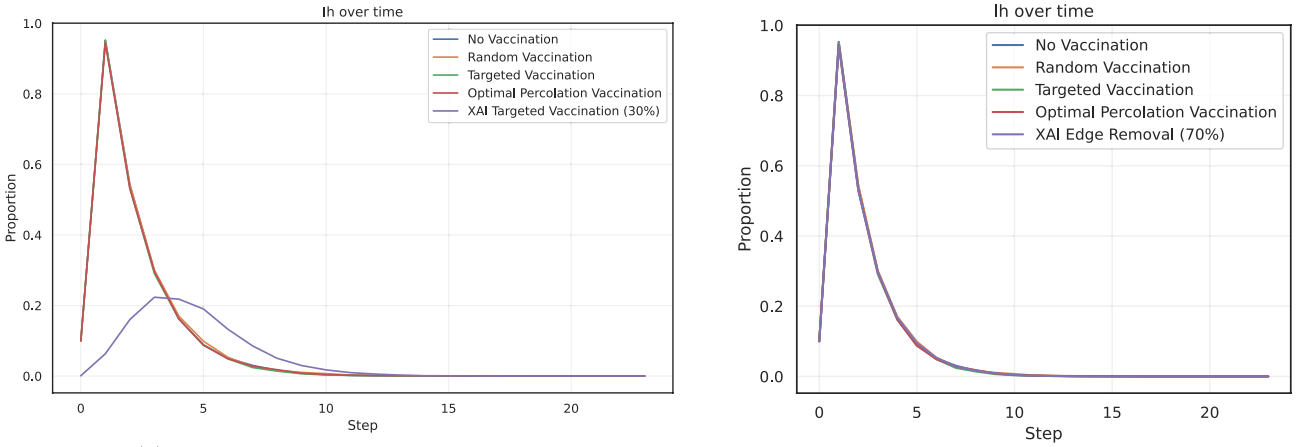

(a) XAI-based vaccination infection rate

(b) XAI-based edge removal infection rate

Fig. 12 Effect of XAI-based targeted vaccination (top 30%) and XAI-guided edge removal (70%) on infected humans(ih) over time on RG network. The mean values of all the simulations (repeated 5 times) are compared here, distinguishing each case (baseline methods and XAI-guided intervention) by the associated colour in the legend

reduced the peak to 30.73% and total deaths to 0.10% (the same target vaccination probability was used).

Real world contact map

To simulate the measles spreading on a not-synthetic real-world population we selected the dataset of daily human contact interaction dataset proposed in [38]. In this dataset, authors used wearable proximity sensors to track contact interactions in a village or rural Malawi for 26 consecutive days (from 16th December 2019 to 10th January 2020). Interactions of each day can be represented by a contact graph, were nodes are humans and one edge between two humans exists only if they interacted during the day. Both number of edges and nodes can change between days. For our analysis we are interested to the contact map with the higher number of human interactions, for this reason we selected the graph

Table 5 SIRVD XAI-based interventions in SBM network. The column *Strategy* lists the intervention method applied. *Peak infected* indicates the maximum percentage of the population simultaneously infected during the outbreak. *Total Dead* reports the overall mortality rate at the end of the simulation. *Duration* denotes the total number of time steps (simulation length), where each step is a month. Simulations were repeated 5 times to report mean and standard deviation values

Strategy	Peak Infected	Total Dead	Duration
No Vaccination	94.82% ± 0.22	0.20% ± 0.00057	19
Random Vaccination	94.68% ± 0.21	0.14% ± 0.00043	18
Targeted Vaccination	94.64% ± 0.22	0.22% ± 0.00067	18
Optimal Percolation Vaccination	94.76% ± 0.22	0.20% ± 0.00058	18
XAI Edge Removal (70%)	94.38% ± 0.21	0.26% ± 0.00073	15
XAI Targeted Vaccination (30%)	54.42% ± 0.14	0.14% ± 0.00054	17

representing the interactions of 23rd December: 80 nodes and 139 edges.

Table 7 details the XAI-guided interventions on the real contact network in SIRVD simulations. Figure 18 illustrates the progression of infection under the five different immunisation scenarios. Both report how removing the 70% of important edges (97) from the network (approximately a quarantine of the whole community) didn't have important effects on the infection peak, obtaining only a 1% of peak reduction compared to optimal percolation vaccination. On the other hand, by using our methodology to guide target vaccination, we achieved a reduction of 22.75% in the infection peak (compared to the best baseline strategy in this case, optimal percolation-guided vaccination) by vaccinating the top 30% of nodes (24 nodes) in the network. For a fair comparison, the same target vaccination probability was used. Figure 19 shows edge-level explanations extracted using integrated gradients to highlight important interactions for human measles infection in a real contact network.

Discussion

The experimental evaluation of the proposed framework demonstrates its capacity to simulate complex diffusion processes and inform strategic intervention design through a combination of network science, deep learning, and XAI. Across diverse synthetic network topologies, including Watts-Strogatz, Barabási-Albert, Random Geometric, Stochastic Block models, Synthpops and a contact map of individuals taken from a real-world scenario, our simulations consistently show that the integration of data-driven modelling with interpretability mechanisms enables not only the forecasting of epidemic dynamics but also the precise tailoring of containment strategies.

Synthetic data is emerging as a cornerstone of epidemic modelling, providing a surrogate when empirical

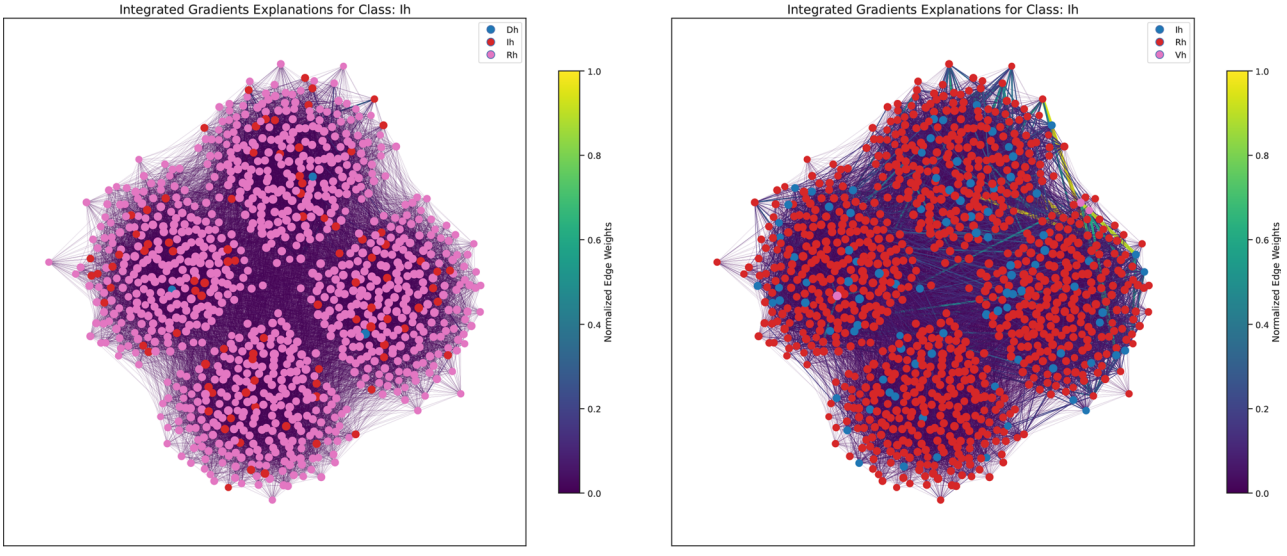

(a) IG based on no vaccination strategy (Ih class).

(b) IG based on targeted vaccination strategy (Ih class).

Fig. 13 IG insights for the SBM network in the SIRVD model setting. This visualization illustrates the node-level relevance scores computed using the IG method, applied to a GNN trained to predict epidemic dynamics. The SBM model, known for its high clustering and short path lengths, reflects realistic patterns of social interaction. The color intensity of each edge corresponds to its attributed influence on the model's infection prediction at a selected time step. Nodes with higher integrated gradient values are identified as critical contributors to the spread of the infection. Node colors correspond to the simulation class and are indicated by abbreviations for each individual's class (sh = Susceptible humans, ih = infected humans, Rh = Recovered humans, vh = vaccinated humans)

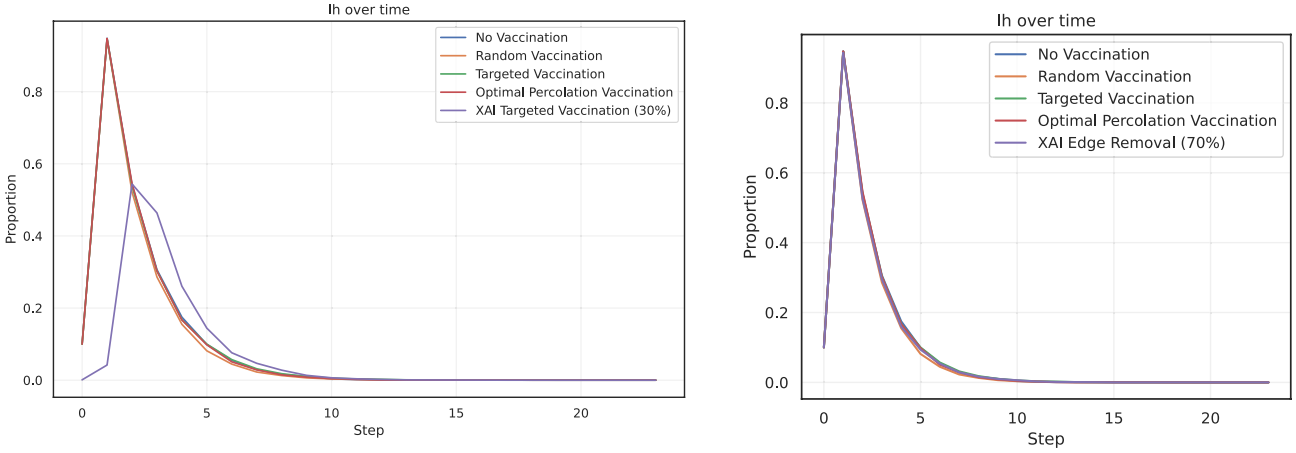

(a) XAI-based vaccination infection rate

(b) XAI-based edge removal infection rate

Fig. 14 Effect of XAI-based targeted vaccination (top 30%) and XAI-guided edge removal (70%) on infected humans(ih) over time on SBM network. The mean values of all the simulations (repeated 5 times) are compared here, distinguishing each case (baseline methods and XAI-guided intervention) by the associated colour in the legend

Table 6 SIRVD XAI-based interventions in Synthpops network. The column *Strategy* lists the intervention method applied. *Peak infected* indicates the maximum percentage of the population simultaneously infected during the outbreak. *Total Dead* reports the overall mortality rate at the end of the simulation. *Duration* denotes the total number of time steps (simulation length), where each step is a month. Simulations were repeated 5 times to report mean and standard deviation values

Strategy	Peak Infected	Total Dead	Duration
No Vaccination	62.45% (0.17)	0.25% (0.00080)	16
Random Vaccination	63.22% (0.18)	0.21% (0.00067)	18
Targeted Vaccination	60.44% (0.16)	0.25% (0.00066)	15
Optimal Percolation Vaccination	60.58% (0.16)	0.12%(0.00038)	19
XAI Edge Removal (70%)	59.44% (0.16)	0.17% (0.00051)	17
XAI Targeted Vaccination (30%)	30.73% (0.087)	0.10% (0.00037)	17

data are scarce, restricted, or incomplete [39, 40]. Population synthesis methods generate artificial datasets that preserve demographic heterogeneity, such as age structure, household composition, and spatial distribution while enabling analyses at scales that survey data cannot achieve. Beyond privacy and scalability, synthetic populations support benchmarking, allow precise control over demographic and epidemiological parameters, and enrich training corpora for machine-learning models. Recent studies demonstrate that spatially explicit synthetic populations can reproduce observed demographic statistics and that epidemic trajectories derived from PDE- or SIR-based frameworks enhance the performance of spatio-temporal graph models. Collectively, these advances show that carefully constructed synthetic data not only replicate key epidemic indicators, including peak infection timing and magnitude, but also provide a robust foundation for scenario testing, model validation, and data-driven public health planning [1].

Simulation outcomes confirm that even modest structural interventions, such as targeted vaccination informed by node-level relevance or edge removal guided by saliency attribution, can lead to substantial reductions in infection peaks and total mortality. In particular, XAI-driven approaches offer advantages over random or centrality-based heuristics by leveraging model-derived insights rather than relying solely on static network properties. This advantage becomes particularly apparent in networks with heterogeneous or modular structures, where high-risk nodes are not always intuitively or centrally located.

The use of simulation prior to real-world deployment allows health authorities to assess multiple scenarios without incurring the costs, delays, or ethical risks associated with live experimentation. This has direct implications for pandemic preparedness: policies can be pre-tested in silico, allowing for faster response times when outbreaks occur. Moreover, our results illustrate how data-driven insights can optimize vaccine allocation, avoiding resource wastage while maximizing protective effects. In the BA network, for example, XAI-informed strategies (XAI-targeted vaccination) halved the peak deaths from infections compared to no intervention or other baseline strategies.

The framework's ability to adapt across network models enhances its generalizability and usability for varied biomedical settings. It supports both prospective forecasting and retrospective policy evaluation, offering a principled infrastructure for iterative decision-making. Importantly, these simulations highlight a future in which biomedical modelling platforms act as real-time companions to epidemiological surveillance, guiding policy not only with

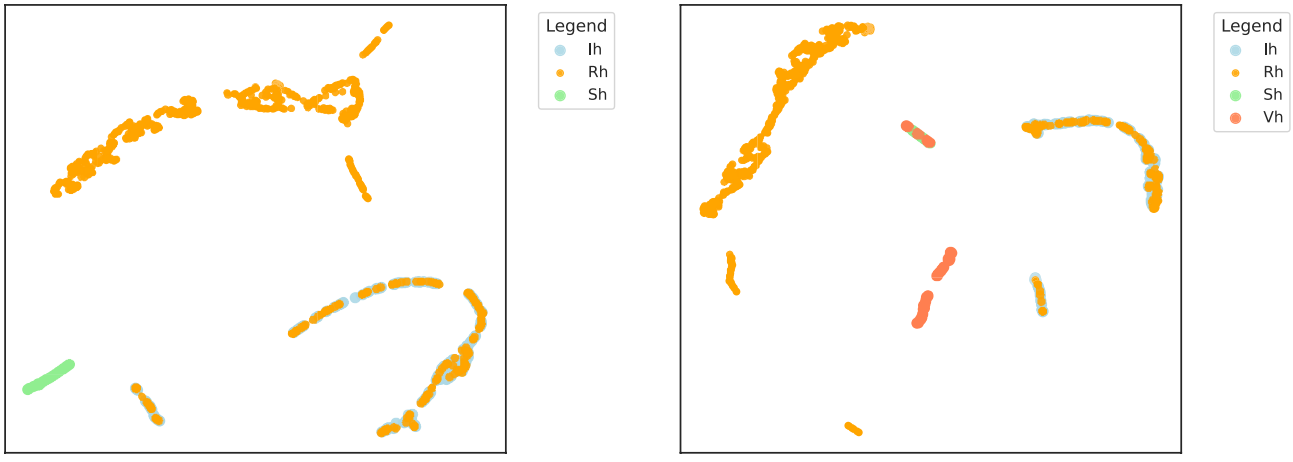

(a) Node embedding in the no vaccination scenario

(b) Node embedding in the optimal percolation-guided vaccination strategy

Fig. 15 Node embedding visualization for the Synthpop network in the SIRVD model setting. This figure represents the latent space learned by the graph neural network (GNN) during training on epidemic diffusion data. Each point corresponds to a node, projected into a two-dimensional embedding space using dimensionality reduction. The position and clustering of nodes reflect structural and dynamic similarities captured by the GNN, including community structures and temporal patterns of infection. Colors encode infection status (sh = Susceptible humans, ih = infected humans, Rh = Recovered humans, vh = vaccinated humans). This representation enables interpretable classification of high-risk nodes and supports downstream tasks such as targeted intervention and forecasting, by revealing relationships not evident from the original graph topology alone

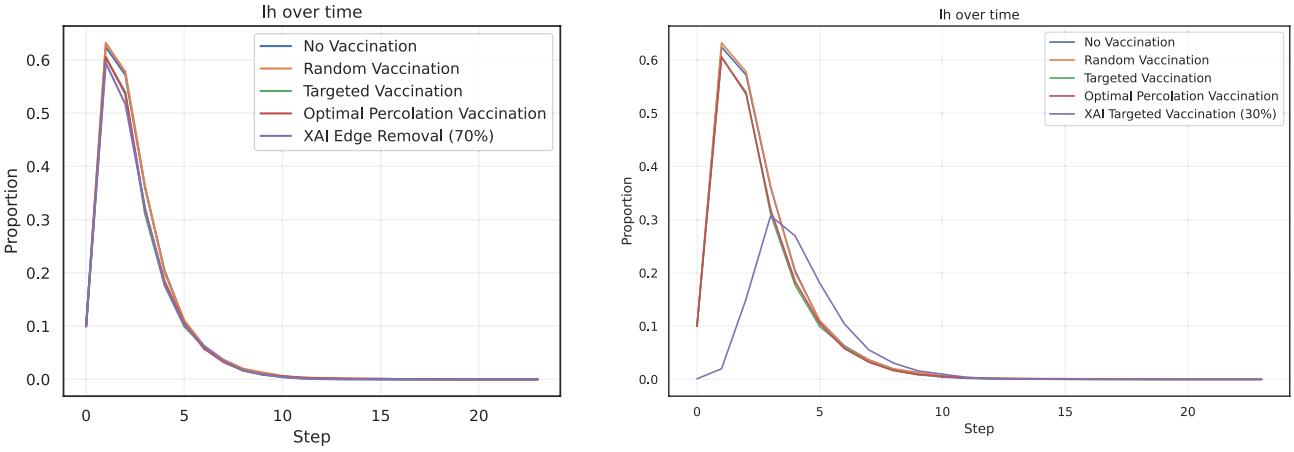

(a) XAI-based edge removal infection rate

(b) XAI-based vaccination infection rate

Fig. 16 Effect of XAI-based edge removal (top 70%) and XAI-based vaccination (top 30%) on infected humans (ih) over time on a Synthpops network. The mean values of all the simulations (repeated 5 times) are compared here, distinguishing each case (baseline methods and XAI-guided intervention) by the associated colour in the legend

predictive accuracy but also with actionable explanations. The integration of interpretability into this process ensures that interventions remain transparent and justifiable to both practitioners and the public.

In sum, the results advocate for a shift toward simulation-enhanced public health planning, where evidence-based decisions are accelerated, lives are preserved, and resources are efficiently utilized. By making these capabilities accessible and extensible, the framework addresses a critical need for proactive, intelligent, and interpretable systems in epidemic control and biomedical modelling.

Conclusion

This work presents a modular, interpretable simulation framework that bridges classical epidemiological modelling with state-of-the-art machine learning techniques to support evidence-based public health decision-making. By integrating network-based diffusion models, graph neural networks, and XAI, the platform enables users to explore complex transmission scenarios, assess intervention strategies, and extract actionable insights from simulated data.

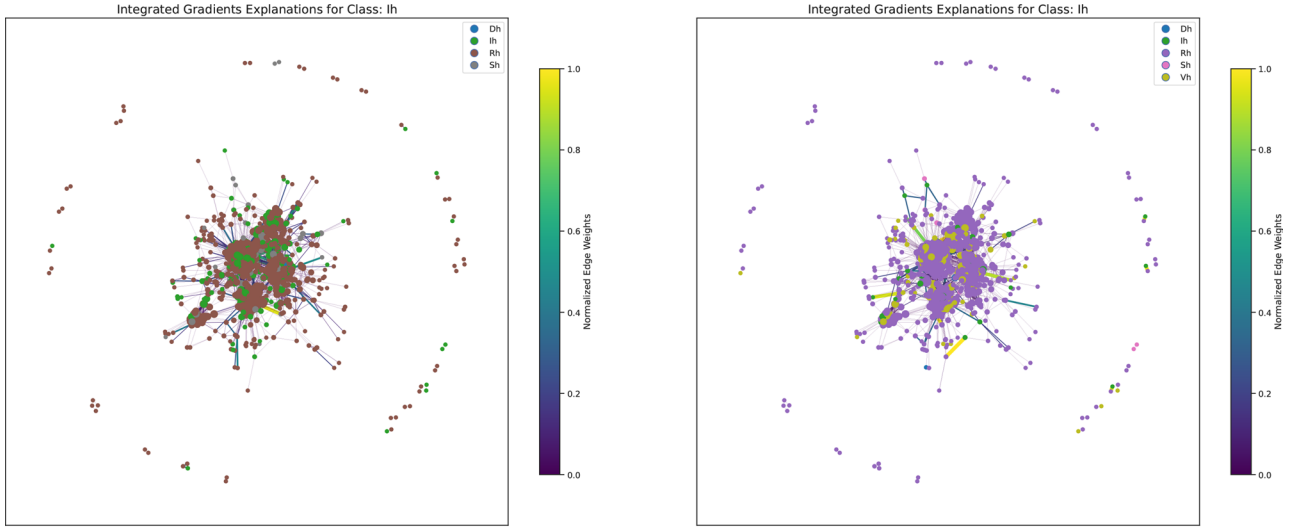

(a) IG based on no vaccination strategy (Ih class).

(b) IG based on targeted vaccination strategy (Ih class).

Fig. 17 IG insights for the Synthpops network in the SIRVD model setting. This visualization illustrates the node-level relevance scores computed using the IG method, applied to a GNN trained to predict epidemic dynamics. The color intensity of each edge corresponds to its attributed influence on the model's infection prediction at a selected time step. Nodes with higher integrated gradient values are identified as critical contributors to the spread of the infection. Node colors correspond to the simulation class and are indicated by abbreviations for each individual's class (sh = Susceptible humans, ih = infected humans, Rh = Recovered humans, vh = vaccinated humans)

Table 7 SIRVD XAI-based interventions in real-word network. The column *strategy* lists the intervention method applied. *peak infected* indicates the maximum percentage of the population simultaneously infected during the outbreak. *total Dead* reports the overall mortality rate at the end of the simulation. *Duration* denotes the total number of time steps (simulation length), where each step is a month. Simulations were repeated 5 times to report mean and standard deviation values

Strategy	Peak Infected	Total Dead	Duration
No Vaccination	39.00% (0.12)	0.0%(0.0)	19
Random Vaccination	35.25% (0.11)	0.0%(0.0)	14
Targeted Vaccination	34.00% (0.11)	0.0%(0.0)	11
Optimal Percolation Vaccination	33.25% (0.10)	0.25%(0.0005)	17
XAI Edge Removal (70%)	32.25% (0.10)	0.25% (0.0005)	17
XAI Targeted Vaccination (30%)	10.50% (0.003)	0.0% (0.0)	8

Through extensive experiments across multiple synthetic network topologies, the framework demonstrated robust predictive capacity and the ability to pinpoint structurally critical nodes and edges for targeted containment. These capabilities translate into tangible benefits in managing infectious disease outbreaks, offering a practical, cost-effective alternative to reactive public health measures. Notably, XAI-guided strategies consistently outperformed traditional heuristics, highlighting the value of interpretable machine learning in real-time epidemic control.

The potential of such a platform extends well beyond the scope of measles modelling. It sets the stage for rapid adaptation to future pandemics, antimicrobial resistance tracking, and behavioural diffusion modelling in healthcare. Ultimately, the framework supports a proactive paradigm in which simulation and interpretability enable timely, transparent, and targeted responses to public health threats.

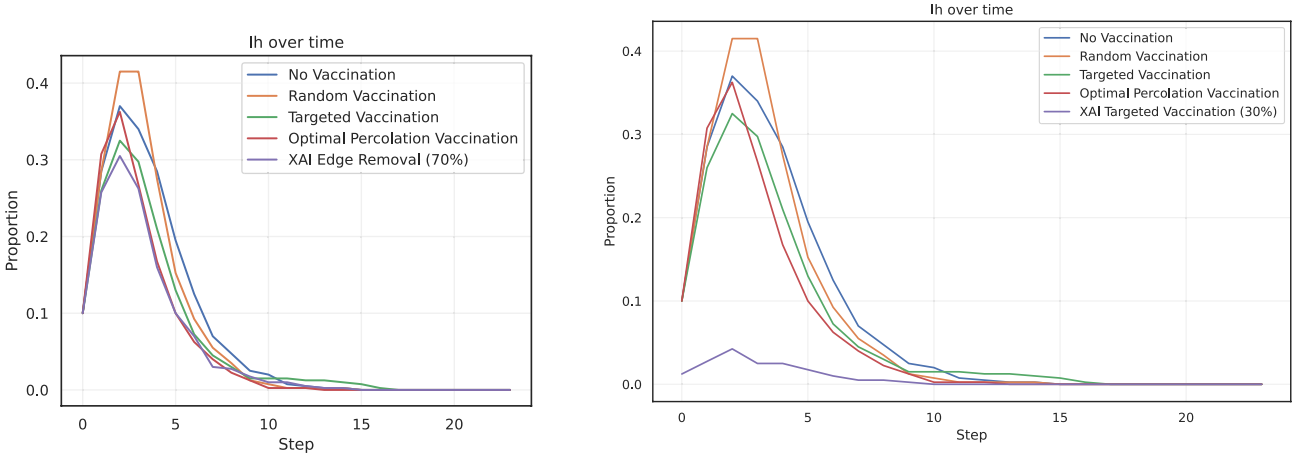

(a) XAI-based edge removal infection rate

(b) XAI-based vaccination infection rate

Fig. 18 Effect of XAI-based edge removal (top 70%) and XAI-based vaccination (top 30%) on infected humans (ih) over time on a real-world contact network. The mean values of all the simulations (repeated 5 times) are compared here, distinguishing each case (baseline methods and XAI-guided intervention) by the associated colour in the legend

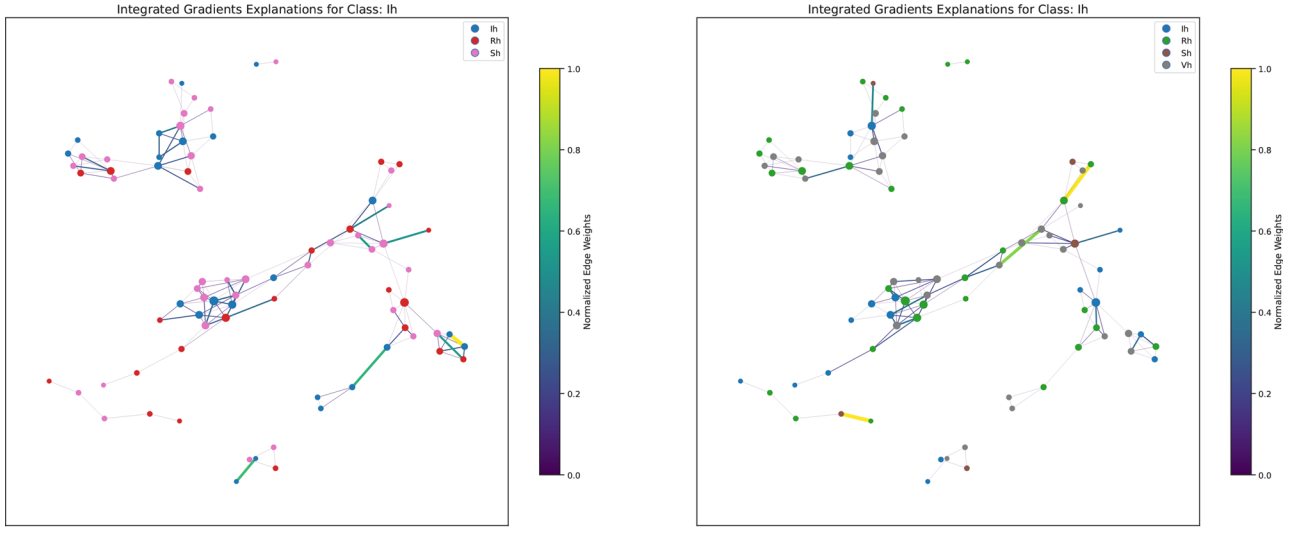

(a) IG based on no vaccination strategy (Ih class).

(b) IG based on targeted vaccination strategy (Ih class).

Fig. 19 IG insights for the real network in the SIRVD model setting. This visualization illustrates the node-level relevance scores computed using the IG method, applied to a GNN trained to predict epidemic dynamics. The color intensity of each edge corresponds to its attributed influence on the model's infection prediction at a selected time step. Nodes with higher integrated gradient values are identified as critical contributors to the spread of the infection. Node colors correspond to the simulation class and are indicated by abbreviations for each individual's class (sh = Susceptible humans, ih = infected humans, Rh = Recovered humans, vh = vaccinated humans)

Acknowledgements
We acknowledge the support of the PNRR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI, under the NRRP MUR program funded by the NextGenerationEU.

Author contributions
Author Contributions: F.B. and P.H.G. conceived the main idea and overall direction of the study. P.V. contributed to the bioinformatics aspects and guided the selection of relevant epidemiological data. M.C. and F.S. curated the mathematical modelling and simulation components, with particular focus on measles dynamics. A.D., U.L., and B.P. designed and implemented the computational framework, including the integration of network simulations, machine learning modules, and explainability tools. All authors contributed to writing, reviewed, and approved the final manuscript.

Data availability
Data are available at http://github.com/fbranda/measles.

Code availability
Code is available at https://github.com/AI4BIOCONE/xai_sirvd_measles.

Declarations

Ethics approval and consent to participate
Not applicable.

Consent for publication
All the authors read and approved the manuscript.

Relevant guidelines and regulation
Not applicable.

Competing interests
The authors declare no competing interests.

Received: 30 July 2025 / Accepted: 5 February 2026
Published online: 20 February 2026

References

1. Zitnik M, Li MM, Wells A, Glass K, Morselli Gysi D, Krishnan A, et al. Current and future directions in network biology. Bioinf Adv. 2024;4(1):099.
2. Shahrabi A, Nikpanjeh F, Hamounian A, Mohebbi H, Shekari M, Parvandi Z, Asoudeh M, Rahimi Tabar MR. Data-driven stability analysis of complex systems with higher-order interactions. Commun Phys. 2025;8(1):239.
3. Vespignani A. Modelling dynamical processes in complex socio-technical systems. Nat Phys. 2012;8(1):32–39.
4. Branda F, Giovanetti M, Romano C, Benvenuto D, Ciccozzi A, Sanna D, et al. Global measles surveillance: trends, challenges, and implications for public health interventions. Infect Disease Rep. 2024;16(2):367–79.
5. Patel MK. Progress toward regional measles elimination—worldwide, 2000–2019. MMWR Morbidity Mortal Wkly Report. 2020;69.
6. Branda F, Giovanetti M, Romano C, Ciccozzi A, Sanna D, Ciccozzi M, et al. The reemergence of measles and the urgent need for uninterrupted genetic surveillance and vaccination. Clin Microbiol Infect. 2024;30(10):1215–18.
7. Orenstein WA, Strebel PM, Papania M, Sutter RW, Bellini WJ, Cochi SL. Measles eradication: is it in our future? Am J Public Health. 2000;90(10):1521.
8. Elam-Evans LD. National, regional, state, and selected local area vaccination coverage among adolescents aged 13–17 years—United States, 2019. MMWR Morbidity Mortal Wkly Report. 2020;69.
9. Vynnycky E, White R. An introduction to infectious disease modelling. OUP oxford, ???; 2010.
10. Pastor-Satorras R, Vespignani A. Epidemic spreading in scale-free networks. Phys Rev Lett. 2001;86(14):3200.
11. Grais RF, Ferrari M, Dubray C, Bjørnstad O, Grenfell B, Djibo A, et al. Estimating transmission intensity for a measles epidemic in niamey, Niger: lessons for intervention. Trans R Soc Trop Med Hyg. 2006;100(9):867–73.
12. Alguliyev R, Aliguliyev R, Yusifov F. Graph modelling for tracking the covid-19 pandemic spread. Infect Disease Modell. 2021;6:112–22.
13. Li W, Nie Y, Li W, Chen X, Su S, Wang W. Two competing simplicial irreversible epidemics on simplicial complex. Chaos: Interdiscip J Nonlinear Sci. 2022;32(9).
14. Li W, Li J, Nie Y, Lin T, Chen Y, Liu X, et al. Infectious disease spreading modeling and containing strategy in heterogeneous population. Chaos Solitons Fractals. 2024;181:114590.
15. Guerra FM, Bolotin S, Lim G, Heffernan J, Deeks SL, Li Y, et al. The basic reproduction number (r0) of measles: a systematic review. Lancet Infect Dis. 2017;17(12):420–28.
16. Freifeld CC, Mandl KD, Reis BY, Brownstein JS. Healthmap: global infectious disease monitoring through automated classification and visualization of internet media reports. J Am Med Inf Assoc. 2008;15(2):150–57.
17. Wallinga J, Teunis P, Kretzschmar M. Using data on social contacts to estimate age-specific transmission parameters for respiratory-spread infectious agents. Am J Epidemiol. 2006;164(10):936–44.
18. Kermack WO, McKendrick AG. A contribution to the mathematical theory of epidemics. Proceedings of the royal society of london. Ser A Containing Papers Math Phys Character. 1927;115(772):700–21.
19. Guzzi PH, Petrizzelli F, Mazza T. Disease spreading modeling and analysis: a survey. Briefings Bioinf. 2022. https://doi.org/10.1093/bib/bbac230.
20. Ajelli M, Gonçalves B, Balcan D, Colizza V, Hu H, Ramasco JJ, et al. Comparing large-scale computational approaches to epidemic modeling: agent-based versus structured metapopulation models. BMC Infect Dis. 2010;10(1):1–13.
21. Koher A, Lentz HH, Hövel P, Sokolov IM. Infections on temporal networks—a matrix-based approach. PLoS One. 2016;11(4):0151209.
22. Liu Q-H, Ajelli M, Aleta A, Merler S, Moreno Y, Vespignani A. Measurability of the epidemic reproduction number in data-driven contact networks. Proc Natl Acad Sci. 2018;115(50):12680–85.
23. Colizza V, Pastor-Satorras R, Vespignani A. Reaction–diffusion processes and metapopulation models in heterogeneous networks. Nat Phys. 2007;3(4):276–82.
24. Defilippo A, Veltri P, Lió P, Guzzi PH. Leveraging graph neural networks for supporting automatic triage of patients. Sci Rep. 2024;14(1):12548.
25. Adadi A, Berrada M. Peeking inside the black-box: a survey on explainable artificial intelligence (xai). IEEE Access. 2018;6:52138–60.
26. Dagdelen J, Dunn A, Lee S, Walker N, Rosen AS, Ceder G, et al. Structured information extraction from scientific text with large language models. Nat Commun. 2024;15(1):1418.
27. Oshin M, Campos N. Learning LangChain. O'Reilly Media, Inc., ???; 2025.
28. Batini C, Cappiello C, Francalanci C, Maurino A. Methodologies for data quality assessment and improvement. ACM Comput Surv (CSUR). 2009;41(3):1–52.
29. Kahn MG, Brown JS, Chun AT, Davidson BN, Meeker D, Ryan PB, et al. Transparent reporting of data quality in distributed data networks. Egems. 2015;3(1):1052.
30. Arts DG, De Keizer NF, Scheffer G-J. Defining and improving data quality in medical registries: a literature review, case study, and generic framework. J Am Med Inf Assoc. 2002;9(6):600–11.
31. Weiskopf NG, Weng C. Methods and dimensions of electronic health record data quality assessment: enabling reuse for clinical research. J Am Med Inf Assoc. 2013;20(1):144–51.
32. Romano C, Branda F, Scarpa F, Lasinio GJ, Ciccozzi M. Monitoring measles infections using flight passenger dynamics in europe: a data-driven approach. Sci Data. 2024;11(1):1358.
33. Branda F, Nakase T, Maruotti A, Scarpa F, Ciccozzi A, Romano C, et al. Dengue virus transmission in italy: historical trends up to 2023 and a data repository into the future. Sci Data. 2024;11(1):1325.
34. Kipf TN, Welling M. Semi-supervised classification with graph convolutional networks. International Conference on Learning Representations. 2017. https://openreview.net/forum?id=SJU4ayYgl.
35. Popper N, Zechmeister M, Brunmeir D, Rippinger C, Weibrecht N, Urach C, et al. Synthetic reproduction and augmentation of covid-19 case reporting data by agent-based simulation. medRxiv, 2020;2020–11.
36. Lima CACD, Silva LA, TedescoPCDAR. Simulation of epidemic dynamics using a multi-agent model: analysis of social distancing strategies and their impacts on public health and economy. Appl Sci (2076-3417). 2024;14(19).
37. Renardy M, Kirschner DE. A framework for network-based epidemiological modeling of tuberculosis dynamics using synthetic datasets. Bull Math Biol. 2020;82(6):78.
38. Ozella L, Paolotti D, Lichand G, Rodríguez JP, Haenni S, Phuka J, et al. Using wearable proximity sensors to characterize social contact patterns in a village of rural Malawi. EPJ Data Sci. 2021;10(1):46.
39. Zhu K, Yin L, Liu K, Liu J, Shi Y, Li X, et al. Generating synthetic population for simulating the spatiotemporal dynamics of epidemics. PLoS Comput Biol. 2024;20(2):1011810.
40. Guzzi PH, Petrizzelli F, Mazza T. Disease spreading modeling and analysis: a survey. Briefings Bioinf. 2022;23(4):230.

Publisher's Note