

RESEARCH ARTICLE

Multilayer modular fusion graph attention network (MMF-GAT) for epidemic prediction

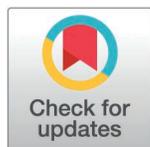
Kayo Fujimoto ^{1†*}, Lizhong Liu^{1†}, Jacky Kuo¹, Baode Gao¹, Hei Kit Chan ¹, Michelle Carr², Roger Sealy², Xi Luo¹, Justin Bahl³

1 School of Public Health, University of Texas Health Science Center at Houston, Houston, Texas, United States of America, **2** City of Houston Health Department, Houston, Texas, United States of America,

3 Institute of Bioinformatics, University of Georgia, Athens, Georgia, United States of America

† These authors are designated as co-first authors.

* kayo.fujimoto@uth.tmc.edu



Abstract

This study introduces the multilayer modular fusion graph attention network (MMF-GAT), an interpretable predictive framework that combines principles from network science and sociology to forecast infection risk. The model is based on the premise that infection risk is influenced by an individual's overlapping membership across various social groups and their participation in shared activities that function as interaction hubs. To represent this complex social fabric, the framework formalizes a multilayer social graph that is assumed to be composed of personal contact, household, and community layers via fourth-order tensor formalism. The MMF-GAT architecture employs layer-specific graph attention network (GAT) modules to preserve context-specific modularity, combined with a late-fusion mechanism that represents how individuals connect different social settings. This approach operationalizes key sociological concepts by processing each layer independently before integration. When applied to a COVID-19 surveillance dataset of 2,264 individuals from Houston, Texas, the MMF-GAT model significantly outperformed five baseline models, achieving an accuracy of 0.78, an area under the curve (AUC) of 0.90, an F1 score of 0.72, and a precision–recall area under the curve (PRAUC) of 0.89. Explainable AI (XAI) analysis identified structural features, particularly household degree, degree of personal contact, and affiliation with educational centers, as the most influential predictors. Owing to its robust predictive accuracy, the model effectively identifies high-risk individuals and settings, positioning it as a valuable tool for public health operations. This study makes a dual contribution. First, it advances multilayer network science through a novel computational architecture that preserves modularity. Second, it provides a validated, high-performance tool for public health informatics. This computational tool can support targeted interventions by optimizing contact tracing prioritization and resource allocation.

OPEN ACCESS

Citation: Fujimoto K, Liu L, Kuo J, Gao B, Chan HK, Carr M, et al. (2025) Multilayer modular fusion graph attention network (MMF-GAT) for epidemic prediction. *PLOS Complex Syst* 2(10): e0000070. <https://doi.org/10.1371/journal.pcsy.0000070>

Editor: Anjalika Nande, Johns Hopkins University Whiting School of Engineering, UNITED STATES OF AMERICA

Received: May 2, 2024

Accepted: September 1, 2025

Published: October 23, 2025

Peer Review History: PLOS recognizes the benefits of transparency in the peer review process; therefore, we enable the publication of all of the content of peer review and author responses alongside final, published articles. The editorial history of this article is available here: <https://doi.org/10.1371/journal.pcsy.0000070>

Copyright: © 2025 Fujimoto et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution,

and reproduction in any medium, provided the original author and source are credited.

Data availability statement: Data were collected through COVID-19 surveillance by the Houston Health Department, which did not require informed consent. This study was approved by the Institutional Review Committee for the Protection of Human Subjects at the University of Texas Health Science Center at Houston (HSC-SPH-20-1022). Due to ethical and legal restrictions outlined in the Memorandum of Understanding between the Houston Health Department and the University of Texas Health Science Center at Houston, the original de-identified surveillance data cannot be made publicly available. Researchers seeking access must obtain approval from the Houston Health Department's Investigative Review Committee by contacting (analysisdatarequest@houstontx.gov). More: Institutional Review Committee | Houston Health Department. The code can be accessed at the GitHub (<https://gitpapl1.uth.tmc.edu/fujimoto-lab-uthealth-hhd/mmf-gat>)<https://github.com/Fujimoto-lab-UTHealth-HHD/MMF-GAT>). Additionally, summary descriptive statistics of the dataset are provided in the online supplementary materials to support transparency and reproducibility. All methods were carried out in accordance with relevant guidelines and regulations.

Funding: This work was supported by the Centers for Disease Control and Prevention (CDC) (75D30121C10133 to JB; NU50CK000626 to JB), and partially supported by Sally W. Vernon, Ph.D. Distinguished Professorship in Social Determinants of Health, UTHealth Houston, to KF. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Author summary

We developed a new AI approach to predicting and interpreting infection risk during disease outbreaks by examining how individuals' daily social environments overlap, including personal contacts, households, and community spaces such as schools and workplaces. Unlike traditional models that treat these interactions as a simple whole, our method preserves unique structure of each social layer before integrating them. This layered approach allows us to capture the complexity of infection spread in real life. Our model was tested on public health surveillance data from COVID-19 outbreak in Houston, Texas, U.S., and was found to predict infection risk with high accuracy, outperforming existing methods. The key factors influencing risk included the size of households, the number of close contacts, and affiliation with education settings. Our findings demonstrate that infections propagate through both tight-knit groups and the key bridging roles of individuals who connect multiple social circles. This understanding offers public health officials a powerful tool to target interventions effectively, enabling effective allocation of limited resources for testing and contact tracing, consequently improving outbreak control and reducing avoidable disruption.

1. Introduction

Network models provide a structured framework for understanding disease transmission dynamics within complex social systems, moving beyond the limitations of traditional epidemiological models that assume homogeneous mixing [1,2]. By mapping social connections, network models represent localized outbreaks and superspreader events, demonstrating how infections spread across communities and informing targeted interventions.

Three widely used network models, scale-free, small-world, and modularity models, have substantially advanced our understanding of how diseases spread through social interactions. However, each has limitations in its representation of the complexity of real-world transmission dynamics. Scale-free networks emphasize the role of highly connected hubs (potential superspreaders) in driving outbreaks [3,4]; however, they often have a limited ability to account for the rapid transmission facilitated by localized clustering within closely knit groups. Small-world networks, which balance local clustering with global connectivity through shortcuts [5], may not fully reflect how relationships develop within shared social settings that influence disease transmission patterns. Modularity-based models identify communities by analyzing dense internal connections and sparse intercommunity links [6,7]. However, they rarely address cross-layer dependencies in multilayer networks. In these networks, interactions in one context (e.g., a school) can impact the spread of disease in another context (e.g., residential communities).

Most conventional network analyses rely on egocentric or sociometric one-mode networks and focus primarily on direct interactions within uniplex structures.

However, such approaches often provide an incomplete picture of disease spread in complex social systems. A more comprehensive framework can be employed to examine both direct and indirect relationships across multiple, overlapping social contexts, including social networks, shared spaces, and community activities.

This study introduces the multilayer modular fusion graph attention network (MMF-GAT), a predictive framework designed to address these gaps. Grounded in sociological theories on the duality of persons and groups [8] and the role of shared activities as *foci* for interaction [9,10], our framework positions modularity at the core of disease transmission. It organizes social interactions into three primary layers: personal contact networks (P), household coresidencies (H), and shared community spaces (C), which are further decomposed into specific site-specific sublayers (S_k) (e.g., schools). This multilayer approach defines individuals by their simultaneous membership in multiple social collectives, allowing the model to capture that overlapping affiliations and brokerage roles influence infection dynamics.

To overcome the limitations of conventional sociometric approaches, which often overlook cross-layer dynamics, this study developed a computational system by integrating graph attention networks (GATs) [11], late fusion techniques [12], and explainable artificial intelligence (XAI) [13]. This system models complex interaction patterns both within and across social layers, aiming to advance infection prediction. Its practical utility is demonstrated through analysis of COVID-19 surveillance data from Houston, Texas, U.S., providing a transparent and flexible tool to support public health efforts such as contact tracing.

2. Review of network models in epidemiology

2.1. Scale-free networks: Hubs and disease transmission

Scale-free networks are characterized by power-law degree distributions, where a small number of highly connected hubs disproportionately influence overall network dynamics [3,4,14]. These hubs are key to disease transmission, as they serve as central nodes that facilitate the rapid spread of infections [15].

Degree-based assortative mixing, namely, the tendency for nodes with similar degrees to connect, is a common feature of social networks [16]. This pattern entails a dense core of highly interconnected hubs that significantly influence disease dynamics. In networks characterized by assortative mixing, hubs tend to connect with other hubs, intensifying transmission within the network. Targeting such hubs for interventions has been shown to reduce infection rates and mitigate disease spread [17].

Scale-free network models emphasize global connectivity through hubs and often downplay the importance of local dynamics. Real-world disease transmission involves both global and localized processes. For example, rapid transmission within closely knit groups can occur due to clustering. Consequently, scale-free models may inadequately represent the dynamics between global hub-driven spread and localized outbreaks.

2.2. Small-world networks: Balancing local and global transmission

In small-world networks, real-world contact patterns are effectively represented through the integration of global connectivity with local clustering dynamics [5]. This balance is achieved through the implementation of densely connected local clusters linked by “shortcuts” that bridge otherwise distant parts of the network [18]. These shortcuts enable the rapid spread of disease across the entire network, whereas the high level of clustering within the network means that most infections still occur locally [18].

Although valuable for understanding disease spread, small-world networks may not completely account for the influence of community structure on epidemic dynamics. Research has demonstrated that community structure plays a key role in shaping infection patterns within networks [19]. Individuals who connect otherwise separate communities are shown to be important for disease transmission. Targeting these bridging individuals for immunization can substantially increase intervention effectiveness, particularly in networks with strong community structures. Combining modularity

analysis with small-world network models could provide a more comprehensive approach to understanding and managing disease transmission.

2.3. Modularity: The role of communities and containment

Modularity, also referred to as community structure [6, 7], is a key concept for understanding disease transmission in networks. It quantifies the extent to which a network is partitioned into densely connected communities with relatively sparse connections between them [6,20]. These network boundaries, defined by the limits of social groups or communities, shape disease transmission dynamics by balancing localized containment with the potential for broader disease spread. Dense intracommunity ties allow infection to spread easily within a community, acting as an incubator [21], whereas sparse inter-community links can hinder spread between communities [22].

In highly modular networks characterized by strong subgroup cohesion and fragmentation, the spread of infection across community boundaries can be hindered, potentially slowing or even containing outbreaks [23]. However, these protective effects can be offset by global shortcuts, where long-range connections bridge otherwise disconnected communities [23]. Intercommunity edges are essential for infection to reach other communities, acting as conduits for wider dissemination or enabling global spread, as seen in a spatial modular network [21,24,25].

In addition to basic dynamics, studies of the epidemic threshold indicate complexities in modular networks. Investigations into the epidemic threshold have revealed various impacts. Some studies suggest that, under otherwise identical conditions, the epidemic threshold (to first order) remains unchanged by community structure [26, 27]. However, the effects of community structure on epidemic dynamics can vary on the basis of specific network features, such as the interplay between clustering and the degree distribution, as well as the conditions of balanced communities [26,27]. Analyses of spatial models identify distinct local and global thresholds [25], and empirical studies have shown significant increases, primarily in extremely modular networks [23]. Regardless, the mesoscopic community structure is essential for predicting post-invasion dynamics [24].

An epidemiological study on HIV transmission among intravenous drug users in New York City illustrated how network structure influences disease dynamics [28]. This study identified smaller, interconnected subgroups of uninfected individuals within a larger network that included both infected and uninfected individuals. These subgroups, while not entirely isolated, had limited connections to the wider network. This limited connectivity potentially restricted HIV spread. If a new infection emerged, the outbreak was likely contained within the immediate contacts of the newly infected person in the subgroup. Longer-term HIV-infected individuals on the periphery of these islands acted as “firewalls” owing to their lower infectivity, potentially restricting the spread of HIV from newly infected individuals with high infectivity. This example illustrates how individual traits (e.g., infectivity) and network position interact to influence disease spread within a modular network.

Although this study focused on local dynamics within modular components [28], it did not explore how connections with a broader network might influence disease spread. Building on this work, small-world theory suggests that even a few long-range links can facilitate rapid transmission across communities [29], explaining how long-range connections can drive the broader spread of a disease [5,30]. More specifically, an infected individual bridging an isolated subgroup to the broader network could act as a small-world shortcut, triggering rapid transmission within previously protected communities. This implies how long-range connections can disrupt the containment provided by modularity, balancing local cohesion with global connectivity.

In addition to the influence of long-range links, assortative mixing, whether by grouping individuals with similar numbers of contacts or clustering those with shared traits or behaviors, plays a key role in influencing disease dynamics. It creates clusters within networks that experience varying levels of disease prevalence compared with the general population [18,31]. These clusters often serve as focal points for localized outbreaks and influence epidemic trajectories.

2.4. Multilayer approaches: Progress and challenges in transmission modeling

Building on foundational network models, this subsection explores multilayer frameworks developed to account for the multiscale complexity in real-world social interactions and their heterogeneous contact patterns [27]. Hierarchical meta-population models represent populations as nested systems of distinct contexts, such as local settings (e.g., hospitals, schools, or workplaces, organized within broader communities such as neighborhoods, cities, and national regions). The movement of individuals between these levels is essential for widespread pathogen diffusion [30].

Extending this concept, multilayer networks represent diverse forms of interactions as separate layers, including connections both within (intralayer) and across (interlayer) these layers [20]. For example, empirical studies have shown that age-specific social preferences lead to high infection prevalence among schoolchildren, whereas connections with individuals of different ages facilitate the wider spread of disease through the community [31]. However, conventional network investigations, which are frequently limited to single-mode structures and direct links, often fail to provide a complete representation of disease propagation. A more comprehensive framework is therefore needed to analyze direct and indirect connections across multiple, overlapping social contexts. By acknowledging the influence of contact patterns and hierarchical social structures on disease transmission, multilevel intervention strategies integrate considerations of individual relationships, local groups, and broad system-level configurations to guide effective public health responses [32].

Despite recent advances, the potential of multilayer networks to represent modular social structures in epidemiology is underexplored [20,33,34]. Prior studies have advanced the modeling of disease dynamics across network layers [22,35] but have not yet fully accounted for the dynamic and modular nature of social structures. Epidemic spread in a two-layer system with an information layer and an infection layer has been analyzed; for example, a large number of communities in the infection layer increase the propagation threshold (thus inhibiting spread), whereas a higher average degree decreases it (thus facilitating transmission) [22]. However, this model treats community structure as static within each layer, with limited consideration of how modularity evolves across layers or social contexts. Concurrent contagions (disease and opinion diffusion) across a spatial lattice and a scale-free network have similarly been examined, with a focus on centrality-based interventions that target high-degree nodes in the opinion layer to mitigate disease spread [35]. However, their focus on influence strategies underrepresents the role of structural modularity in shaping transmission patterns across layers.

To address these gaps, this study introduces the multilayer modular fusion graph attention network (MMF-GAT), a framework that positions modularity as a core component of epidemic prediction. This study makes a dual contribution. Our first contribution is to multilayer network science by adapting established tensorial formalism and providing a novel graph neural network architecture designed to preserve and integrate context-specific information. The second contribution is to operationalize public health by introducing a high-performing predictive tool.

By integrating data from personal contact (P), household coresidency (H), and community colocation (C) layers, the MMF-GAT uses layerwise attention and a late fusion mechanism to model complex interactions. As demonstrated with COVID-19 surveillance data from Houston, Texas, U.S., this approach significantly improves prediction accuracy compared with baseline models. Although sampling biases inherent in contact tracing data preclude causal claims about disease transmission, the ability of the framework to identify high-risk individuals and contexts makes it a valuable tool for guiding public health interventions such as prioritizing contact tracing and testing resources.

3. Theoretical and mathematical foundations of the MMF-GAT model

3.1. Sociological concepts of group membership and structural overlap

This study employs a multilayer modular network framework grounded in Breiger's (1974) theory of the "duality of persons and groups," which conceptualizes social structures as emerging from the interplay between individual relationships and group affiliations [8]. Individuals are shaped by the groups to which they belong and, in turn, influence the structure and

dynamics of those groups through repeated interactions. In our framework, personal networks (e.g., neighborhood-based contacts), residential settings (e.g., households), and community collectives (e.g., outbreak sites) are modeled as interrelated layers within a dynamic social system.

Sociological perspectives further clarify that social groups are defined not only by relational ties but also by shared attributes, distinct community traditions, and social contexts [36,37]. These categories establish social boundaries embedded in broader systems of resource distribution, which regulate access to opportunities and structure patterns of association [9,10,37–39]. Such boundaries can contribute to imbalances in community health outcomes by creating varied levels of exposure and sensitivity to health risks, including diseases [36,37,40–42]. In fact, social boundaries act as foundational drivers of inconsistent access to health-promoting resources [36,37,41–43]. This connection persists and influences health outcomes even as specific risk factors change over time [40].

Building on this foundation, our study posits that an individual's risk of infection is influenced by their structural position both within and across multiple layers of the social environment. Our model represents each individual with a feature set that includes their demographic profile, behavioral risk, and group affiliation. These features are designed to reflect how individuals are embedded in an overlapping modular setting, defined by households, institutions, and social ties, and how their roles within those settings contribute to disease transmission.

3.2. Social mechanisms forming modularity

To understand that infections spread through the social system, it is essential to examine the mechanisms that generate modularity, i.e., dense clusters of social ties within groups and sparser connections between them. Several sociological concepts shed light on the development of modular structures and interactions within layers (intralayer interactions).

Feld's concept of *foci* (1981, 1982) refers to shared settings or routine activities (e.g., schools, workplaces, or religious centers) that serve as focal points for repeated interaction [9,10]. These *foci* foster localized clustering of interactions and set up the structural conditions for the formation of modular communities.

Bourdieu's (1990) concept of *habitus* adds depth by emphasizing how internalized dispositions, formed through past experiences, guide behaviors within specific settings [43]. For example, hygiene practices within households or cultural norms governing physical proximity in public settings reflect habitual routines that may increase or reduce transmission risk.

Homophily, or the tendency for individuals to associate with others who share similar characteristics, contributes to the formation of modular structures [39]. It promotes clustering within groups that are socially, demographically, and behaviorally alike, which may intensify localized outbreaks but also limit immediate spread to dissimilar groups.

These mechanisms not only support the internal cohesion of modules (e.g., households or social circles) but also affect the potential for disease contamination or acceleration, depending on the scope and type of connections that link one module to another.

3.3. Cross-layer connectivity and brokerage roles

Although modularity represents the internal structure of social groupings, transmission across groups and layers is shaped by how individuals link otherwise connected segments of the network. To understand how infections cross both modular and multilayer boundaries, this study focuses on brokerage roles or specific structural positions that facilitate transmission between otherwise disconnected groups.

Gould and Fernandez (1989) provide a typology of brokerage roles on the basis of the predefined, nonoverlapping subgroup affiliations of the broker and actors being connected [44]. These roles include coordinator (within-group), itinerants (outsiders brokering within-group), gatekeepers (insiders controlling access), representatives (insiders reaching out), and liaisons (connecting members of distinct groups). This typology is especially relevant for multilayer networks,

where subgroups may exist both within and across layers (e.g., households in the residential layer and shared sites in the community layer). For example, an individual may act as a representative of a household entering a school environment (crossing residential and community layers) or serve as a gateway controlling what flows back into the household.

This brokerage framework was further extended to multilayer network contexts through the concept of “brokerage-centrality conjugated,” which emphasizes how individuals with high degree centrality (those with many connections to their immediate neighbors) are presumed to be effective as brokers across social contexts [45]. Prominent brokerage serves as a specific form of these conjugated roles, involving individuals who hold multiple brokerage positions at once through their broad, multilayer connectivity. Such positions give them disproportionate influence on epidemic trajectories by linking distinct contexts and allowing transmission to move across otherwise disconnected or weakly connected layers.

In the context of epidemic prediction, considering these structural positions allows a model such as MMF-GAT to incorporate relevant patterns of connectivity across layers. Through its late fusion strategy, MMF-GAT integrates information from each social layer in a way that reflects the underlying modular and cross-layer structures, representing not only who is connected but also how those connections span boundaries across modular and layered social systems.

3.4. Metrics for structural positions in personal and modular layers

In the personal contact network (layer P), an individual’s position is assessed via standard network analysis metrics, including degree centrality to measure the number of direct connections for each individual [46], clustering coefficients to measure the density of connections within an individual’s immediate neighborhood [47], and network exposure to measure the extent to which an individual is connected to others who are already infected [48] to determine their effects on transmission patterns. For example, a high degree centrality in layer P identifies hubs with wide-reaching ties that may accelerate pathogen spread. On the other hand, clustering coefficients quantify the local density of connections, indicating the potential for disease containment in cohesive neighbor groups.

These individual-level metrics are then contextualized within the broader multilayer framework. Affiliation in the household (layer H) and community (layer C) networks establishes modularity of the system, forming structural boundaries around family units or shared sites or their types. Although dense connections within these modules can facilitate local transmission, the sparse connections between them can limit widespread outbreaks. The social mechanism for transmission across these modular boundaries is brokerage.

Interlayer links, which represent an individual’s participation in multiple social contexts, allow them to act as brokers connecting otherwise isolated groups. Prominent brokers are identified by operationalizing the “brokerage-centrality conjugated” pattern, which targets individuals who have high degree centrality and spans multiple layers to bridge distinct social contexts [45]. These positions can facilitate disease transmission, as they create channels that cross boundaries that separate different clusters. Therefore, the interaction between personal network metrics and multilayer affiliations is important. For example, high degree centrality in the personal contact network (P) can increase an individual’s brokerage potential when they are embedded in the household (H) and community (C) layers. This provides a more comprehensive picture for understanding an individual’s risk and potential influence on epidemic dynamics.

3.5. Visual representation of multilayer modular dynamics

A multilayer modular network is operationalized via conventions from multilayer network science [33] to represent the complex organization of social interactions. This framework breaks down how individuals navigate overlapping social contexts and that their structural positions create channels for disease transmission. It lays the groundwork for the mathematical formalism presented later.

[Fig 1](#) illustrates the abstract structure of this multilayer network, featuring three primary layers (P , H , C) and seven site-type sublayers (S_k ; $k = 1, \dots, 7$). [Fig 1A](#) shows how individuals traverse these layers through overlapping affiliations,

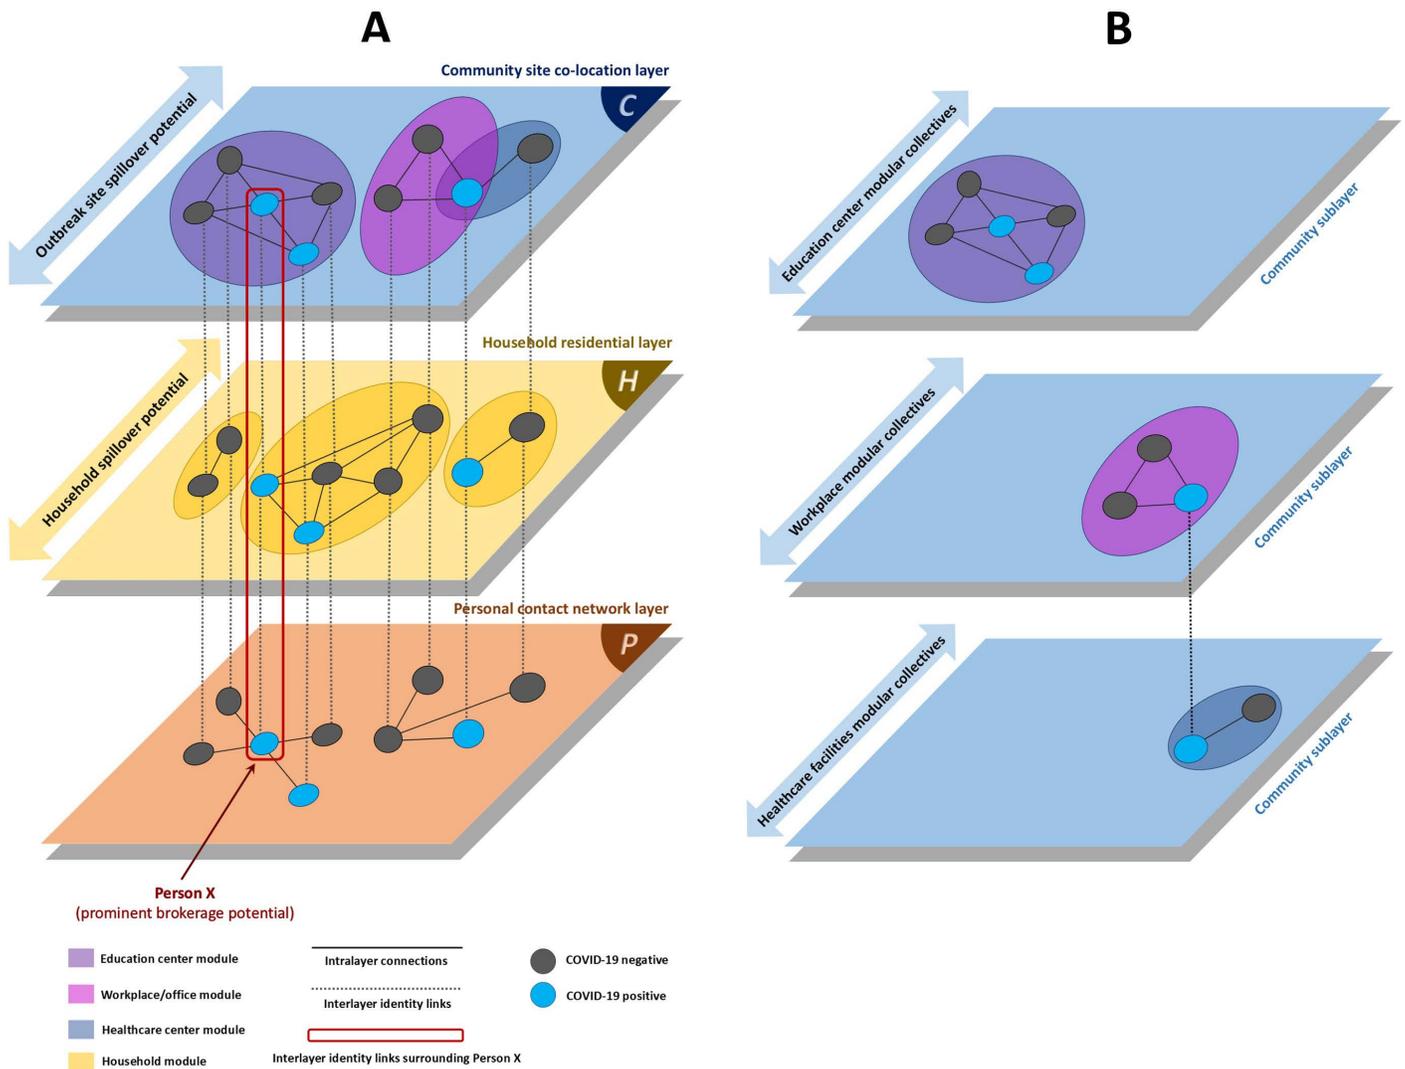


Fig 1. A. Multilayer modular network illustrating primary layers and brokerage roles in disease transmission. B. Site-type decomposition of the community layer.

<https://doi.org/10.1371/journal.pcsy.0000070.g001>

a concept grounded in the sociological theory of the duality of persons and groups [8]. Shared community sites, such as schools and workplaces within the C layer, are described as *foci* for repeated interactions and local clustering [9,10]. This setup can accelerate disease spread, especially when connected to a high-exposure personal network (P).

As illustrated in Fig 1A, “Person X” serves as an example that overlapping brokerage roles can bridge social contexts and facilitate pathogen transmission. By spanning the personal, household, and community layers, Person X’s position combines centrality and brokerage through high connectivity within and across layers. This structural position allows Person X to fulfill several key functions in disease spread, as outlined below.

3.5.1. Interlayer boundary spanner role. In Fig 1A, Person X acts as an interlayer boundary spanner. This role involves connecting distinct social settings (e.g., households and schools) and enabling pathogen movement between them. For example, Person X attends a diverse school (community layer C) and maintains friends in their personal

network (contact network layer P). If infected at school, Person X could transmit the infection to their household or friends. These spanners play a key role in bridging multilayer boundaries.

3.5.2. Intralayer modular connector role. Person X also functions as an intralayer modular connector. This role arises when interlayer boundary spanners create unintended channels that link disconnected groups within a single layer. For example, Person X lives in a high-exposure household (residential layer H) and attends a school where classmates from various backgrounds have little prior pathogen exposure. If a household member becomes infected, Person X could introduce the infection to the school, making peers vulnerable. Infected peers might then spread it to their families, causing outbreaks in previously low-exposure groups. These connections enable rapid transmission by forming shortcuts that bypass structural barriers within layers.

3.5.3. Overlapping roles and prominent brokerage. These roles illustrate how, on the basis of their network position, they can simultaneously serve as interlayer boundary spanners and intralayer modular connectors. Person X's extensive personal networks, including extracurricular school activities (P), combined with their household and school affiliations, position them with prominent brokerage potential. This prominence is characterized by high connectivity across multiple layers, increasing their ability to bridge disconnected groups within and across layers. Person X's representative (outward mediation) and gatekeeping (inward control) functions shift depending on the layer perspective. From the household (H) viewpoint, they represent household risks to the community (C). From the community (C) viewpoint, they control inflows (or gatekeepers) back to the household (H). This duality underscores the impact of these factors on transmission.

[Fig 1B](#) expands the C layer into 7 site-type sublayers (S_k), with 3 examples shown (e.g., education centers, workplaces, and healthcare facilities) corresponding to [Fig 1A](#). These sublayers are operationalized as site-type modular collectives, representing individuals who shared exposure at sites of the same type within the community layer. High degree centrality in P enhances prominent brokerage potential, which varies by site type. For example, students attending schools may exhibit personal clustering patterns that promote interlayer brokerage roles, whereas others may not.

These visualized structures and associated metrics provide the conceptual foundation for the computational framework of the MMF-GAT model, which is introduced next.

3.6. Mathematical formalism and MMF-GAT model architecture

This section introduces the computational framework developed for this study. It begins by establishing the multilayer modular network with a mathematical formalism and then specifies the architecture of the MMF-GAT model.

3.6.1. Tensorial formulation of the multilayer social graph. To formally represent the complex social system relevant to epidemic modeling, this study adopts the tensorial framework of multilayer network science described by Artime et al. [33]. In this approach, a multilayer network consists of N physical nodes (individuals), each interacting across L distinct layers, where each layer represents a unique social context such as personal contact, household coresidency, or various categories of community sites.

The network is mathematically encoded as a fourth-order adjacency tensor $M \in \mathbb{R}^{N \times L \times N \times L}$. Each element $M_{j\beta}^{i\alpha}$ specifies the presence and weight of an edge from node i in layer α to node j in layer β . Here, the Latin indices (i, j) and $i, j \in \{1, 2, \dots, N\}$ denote physical nodes, and the Greek indices (α, β) and $\alpha, \beta \in \{1, 2, \dots, L\}$ denote layers [33]. Each node i is represented by a set of state nodes (replicas), (i, α) , corresponding to its presence in each layer α . This tensorial formulation distinguishes intralayer edges ($\alpha = \beta$), which connect nodes within the same layer, from interlayer edges ($\alpha \neq \beta$), which connect nodes across different layers, preserving the modularity of the network's structure.

Following Artime et al., the tensor M is conceptually decomposed into four components: self-interactions (R), endogenous interactions (N), exogenous interactions (X), and intertwining interactions (I). Self-interactions (R) refer to ties from a node to itself within a layer ($i = j, \alpha = \beta$), endogenous interactions (N) represent ties within a layer ($i \neq j, \alpha = \beta$), exogenous interactions (X) describe ties between distinct nodes in different layers ($i \neq j, \alpha \neq \beta$), and intertwining interactions (I) connect replicas of the same node across layers ($i = j, \alpha \neq \beta$). Thus, the tensor can be decomposed as

$M = R + N + X + I$. In this study, endogenous (N) and intertwining (I) interactions are the focus, whereas self-interaction (R) and exogenous interaction (X) tensors are set to zero. This modeling assumption simplifies the interaction structure and is not a general rule from Artime et al. (2022), who include all $RNXI$ components in their comprehensive framework.

Intralayer connections are denoted by adjacency matrices A_α for each layer α , where $\alpha = \beta$. Each A_α is a symmetric binary matrix of size $N \times N$, with entries $[A_\alpha]_{ij} = M_{j\alpha}^{i\alpha} = 1$ if there is an edge between node i and node j within layer α and 0 otherwise. These matrices represent endogenous interactions within each layer. Interlayer edges, representing intertwining interactions (I), are encoded in the tensor element $M_{j\beta}^{i\alpha}$ for $\alpha \neq \beta$, representing connections between node i in layer α and node j in layer β . Although this study emphasizes intralayer processing before fusion in the MMF-GAT architecture, the tensorial framework provides the mathematical basis for both intralayer and interlayer information flow.

The network constructed for this study comprises $L = 10$ distinct layers, each corresponding to a specific social context and defined as a graph. Each layer graph G_α is associated with an intralayer adjacency matrix A_α , where α belongs to the set $\{P, H, C, S_1, S_2, \dots, S_7\}$. For each layer α , A_α is a symmetric binary matrix of size $N \times N$ ($N = 2,264$ individuals), where $[A_\alpha]_{ij} = 1$ if there is an edge from node i to node j within layer α and 0 otherwise. Each A_α thus represents the intralayer component of the fourth-order adjacency tensor for $\alpha = \beta$.

Building on this tensorial framework, both intralayer and interlayer connections inform the fusion strategies and architectural design choices in the MMF-GAT model, guiding how information is aggregated and integrated across social contexts.

Fig 2 illustrates the overall structure and flow of the MMF-GAT model, describing the input from multilayer social graphs, GAT modular processing, late fusion integration, and output predictions with key evaluation metrics.

3.6.2. Graph attention network integration for multilayer fusion. The MMF-GAT is a computational framework designed to predict infection status by learning from the relational patterns in the multilayer social graph. Its architecture translates the rank-4 tensor formalism into the graph attention network (GAT) formulation [11]. The MMF-GAT first processes each social layer individually before integrating their outputs through a late fusion mechanism.

A GAT layer learns to assign importance to different nodes within a neighborhood on a specific layer α . The initial features for all N nodes in the graph are compiled into a node feature matrix $\mathbf{X} \in \mathbb{R}^{N \times F}$, where F is the number of features

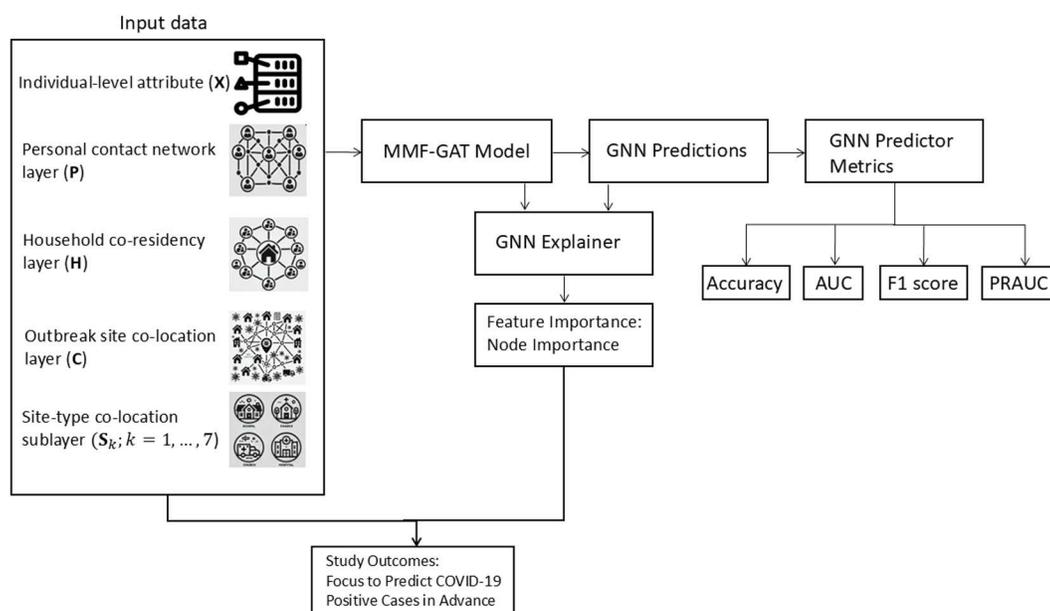


Fig 2. Workflow of the MMF-GAT architecture for COVID-19 prediction.

<https://doi.org/10.1371/journal.pcsy.0000070.g002>

per node, which distinguishes the input features from the hidden representations (h) generated by subsequent neural network layers. The attention mechanism computes unnormalized scores, $e_{ij}^{(\alpha)}$, which are then passed through a Softmax function to obtain the final attention weights, $u_{ij}^{(\alpha)}$, to be described momentarily. The updated feature vector for node i , specific to layer α , is then computed as:

$$\vec{h}_i^{(\alpha)} = \sigma \left(\sum_{j \in \mathcal{N}_i^{(\alpha)}} u_{ij}^{(\alpha)} \mathbf{W}^\alpha \vec{h}_j \right) \tag{1}$$

where \vec{h}_j is the feature vector for a neighboring node j , \mathbf{W}^α for a layer α , is a shared learnable weight matrix, σ is a nonlinear activation function (ELU), and $\mathcal{N}_i^{(\alpha)}$ is the set of neighbors of node i within layer α . This layer-specific notation illustrates how the GAT mechanism operates independently on each layer. The architecture proceeds in steps. First, it processes intralayer connections by computing the attention coefficient between node i and its neighbor j for layer α :

$$e_{ij}^{(\alpha)} = a \left(\mathbf{W}^\alpha \vec{h}_i, \mathbf{W}^\alpha \vec{h}_j \right) \tag{2}$$

These scores are normalized across the neighborhood of node i within layer α to obtain attention weights:

$$u_{ij}^{(\alpha)} = \text{softmax}_j \left(e_{ij}^{(\alpha)} \right) = \frac{\exp \left(e_{ij}^{(\alpha)} \right)}{\sum_{k \in \mathcal{N}_i^{(\alpha)}} \exp \left(e_{ik}^{(\alpha)} \right)} \tag{3}$$

where $\mathcal{N}_i^{(\alpha)}$ is the set of neighbors of node i as defined by the connections in layer α .

Second, multiple heads (with $K=8$ heads) are used to discern diverse relational patterns within each layer [49]. Thus, each head k generates a distinct embedding, $\vec{h}_{i,k}^{(\alpha)}$, and the corresponding outputs are concatenated and transformed via a learnable output weight matrix, \mathbf{W}^O , to form the final layer-specific embedding:

$$\vec{h}_i^{(\alpha)} = \text{Concat} \left(\vec{h}_{i,1}^{(\alpha)}, \vec{h}_{i,2}^{(\alpha)}, \dots, \vec{h}_{i,K}^{(\alpha)} \right) \mathbf{W}^O \tag{4}$$

Here, $\mathbf{W}^O \in \mathbb{R}^{(K \cdot d) \times d'}$, where d is the dimensionality of each head's output and where d' is the desired dimensionality of the final embedding.

Third, a deep late fusion architecture with a sparsity-optimal design based on the principles in [50] (see Model 3 in the ablation experiment) processes interlayer connections. It employs two L1-regularized fusion layers and multihop processing to integrate information, accounting for both 1-hop (direct) and 2-hop (indirect) neighborhood information. The final embeddings from each layer are then concatenated into a comprehensive vector:

$$\vec{c}_i = \text{Concat} \left(\vec{h}_i^{(1)}, \vec{h}_i^{(2)}, \dots, \vec{h}_i^{(L)} \right) \tag{5}$$

where L is the total number of layers indexed by $\alpha = 1, \dots, 10$. This unified vector, \vec{c}_i , represents an individual's composite identity, which is then fed to a final classification layer.

3.6.3. Comparative analysis of late fusion and supra-graph methods. The MMF-GAT model employs a late fusion approach that differs from the supra-graph technique [51], which flattens the multilayer tensor into a single rank-2 matrix of size $(N \cdot L) \times (N \cdot L)$, merging social contexts prematurely and obscuring layer-specific patterns [33,34]. The supra-graph

method disregards contextual differences among settings, weakening the model's ability to learn from brokerage roles conditioned on distinct modular contexts to facilitate spread. Conversely, MMF-GAT processes each layer's adjacency matrix independently with layerwise attention, safeguarding context-specific interactions before integrating embeddings through late fusion.

3.6.4. MMF-GAT design principles and hypothesis evaluation. In the MMF-GAT model, multilayer modular structures and brokerage roles are integrated through a late fusion technique in which each social layer is processed independently before its outputs are combined. Through this design, modularity within layers is preserved to reflect context-specific barriers to disease spread, whereas cross-layer bridging actions, characteristics of brokers, are simultaneously represented to account for transmission across social contexts. The model is designed to test the hypothesis that incorporating multilayer modular structures and brokerage roles produces more accurate predictions of infection status than models that omit these structural features. As such, the framework offers a sociologically grounded tool for analyzing and forecasting epidemics within complex social systems.

3.6.5. Interpretability through explainable AI (XAI) techniques. The interpretability of the MMF-GAT model was addressed by implementing the GNNExplainer method [13] via the PyTorch Geometric framework [52]. This method is designed to identify the most influential features and network connections for prediction by learning nonnegative feature masks that quantify the relative importance of each input value. Within GNNExplainer, interpretability is framed as an optimization problem that aims to maximize the mutual information between the output of the model and a selected subset of features. The output is a ranked list of feature importance scores for each prediction, which provides transparency into the factors influencing the decisions. For this study, a pragmatic top-10 cutoff was adopted for the analysis [53] and focused on interpretation. This practice is consistent with the literature, suggesting that focusing on a limited set of top-ranked features is sufficient to capture the most significant components of explanatory power [53]. This approach and the resulting explanations are compatible with the multilayer modular network framework and consistent with sociological perspectives on individuals' embedding within overlapping social contexts.

3.6.6. Computational efficiency and scalability considerations. The MMF-GAT approach is designed to prevent early aggregation and overparameterization, thus supporting the duality of individuals and groups and integrating context-specific behaviors for effective prediction. Previous supra-graph attention networks applied to health outcomes encountered challenges such as overfitting and high computational costs [50], motivating the design of MMF-GAT with independent graph aggregation and parallelized attention blocks to improve efficiency. Early stopping is incorporated as a hyperparameter option to reduce computational demands.

Although computational requirements may limit current scalability for real-time applications and data aggregation over extended periods may obscure important temporal dynamics, the MMF-GAT framework provides a theoretical foundation that could inform public health efforts, including the strategic prioritization of contact tracing and targeted allocation of testing resources during outbreaks.

3.7. Model validation through ablation experiments

This subsection describes the assessment of the impact of different fusion strategies and network depths, designing three architectural variants of the MMF-GAT. All the models operate on ten network layers representing distinct social contexts (P , H , C , and S_k for k ranging from 1–7). Each model applies to GAT blocks [11] to extract neighborhood-aware embeddings, and all use L1-regularized fusion layers for feature integration and sparsity. Fig 3 illustrates the three models that differ in depth and fusion strategy, as described below.

3.7.1. Model 1: Shallow GAT with late fusion (baseline). This baseline architecture applies a single GAT layer independently to each of the ten network layers ($\alpha = 1-10$), representing distinct social contexts. Within each layer α , the GAT processes immediate 1-hop intralayer connections as defined by the endogenous interaction (N), computes

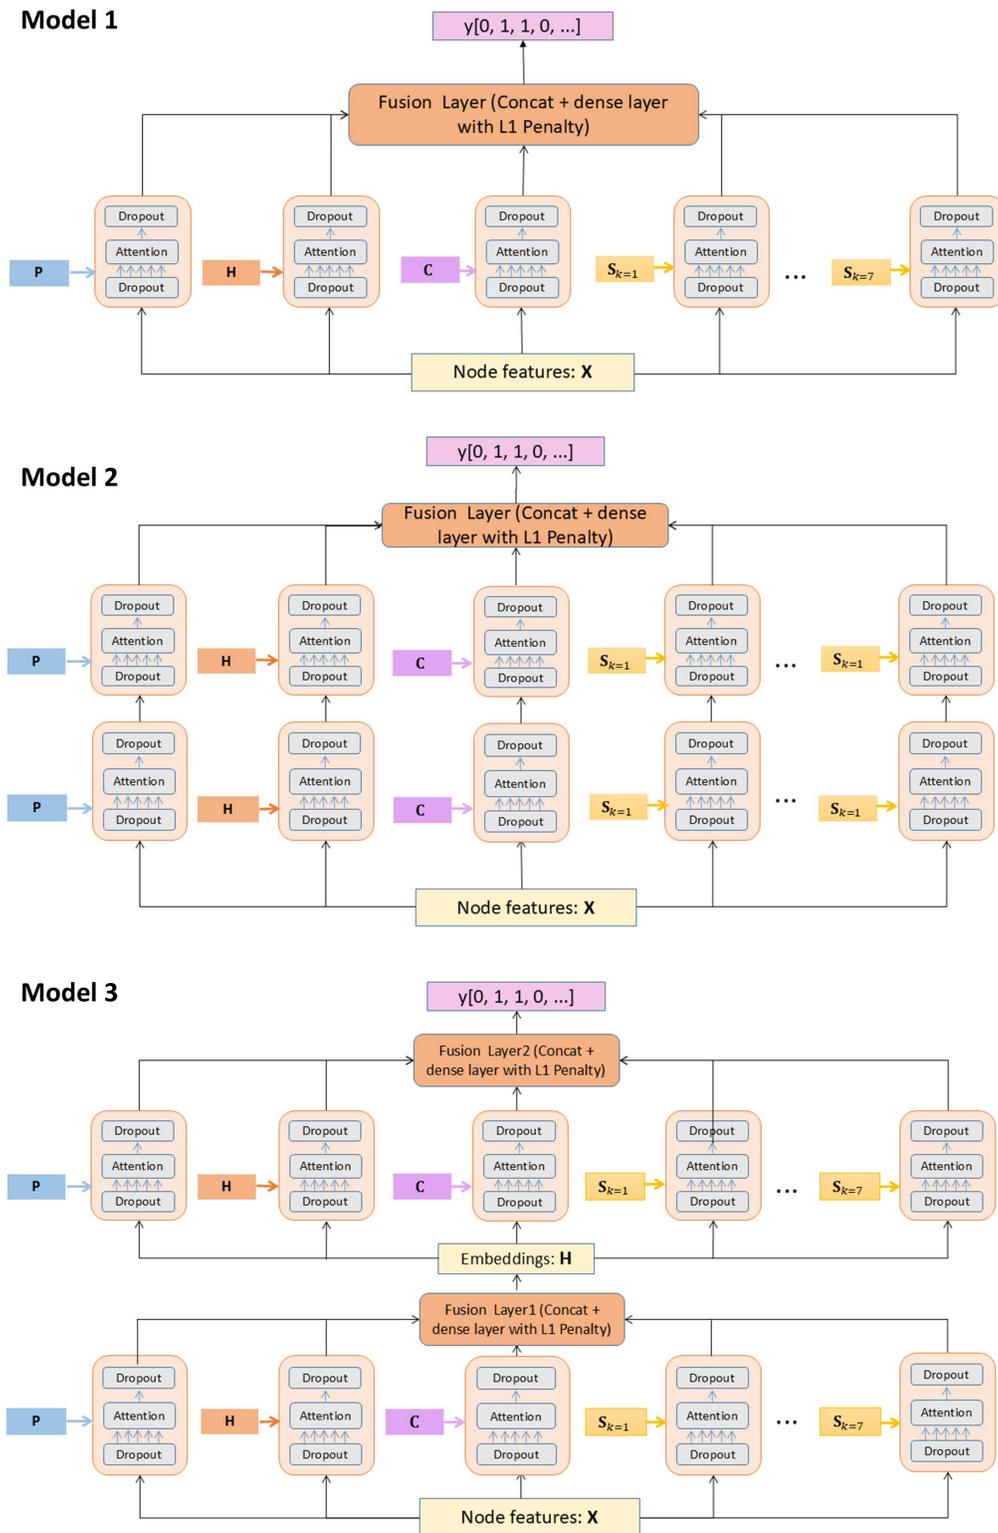


Fig 3. Ablation experiment architecture for the MMF-GAT.

<https://doi.org/10.1371/journal.pcsy.0000070.g003>

embeddings $\vec{h}_i^{(\alpha)}$ by aggregating features from the neighborhood $\mathcal{N}_i^{(\alpha)}$ via attention weights $u_{ij}^{(\alpha)}$, a weight matrix $\mathbf{W}^{(\alpha)} \in \mathbb{R}^{d \times F}$, and ELU activation, following the layer-specific update in [equation \(1\)](#). These layer-specific embeddings are then concatenated across all layers to form $\vec{c}_i = \text{Concat} \left(\vec{h}_i^{(1)}, \dots, \vec{h}_i^{(10)} \right)$, which passes through an L1-regularized dense fusion layer to produce the final embedding for classification. Compared with deeper variants, this model limits depth to one GAT layer per context, representing only direct 1-hop structures and delaying fusion until the output stage, which emphasizes basic intralayer modularity without cross-layer refinement.

3.7.2. Model 2: Deep GAT with late fusion. Building on Model 1, this variant doubles the depth by stacking two GAT layers per network layer, enabling the representation of both 1-hop and 2-hop relationships within each α ; however, it still relies on late fusion. The first GAT layer generates initial embeddings $\vec{h}_i^{(1, \alpha)}$ via attention weights $u_{ij}^{(1, \alpha)}$, a weight matrix $\mathbf{W}^{(1, \alpha)} \in \mathbb{R}^{d \times F}$, and ELU activation. These are fed into a second GAT layer to produce refined embedding $\vec{h}_i^{(2, \alpha)}$ with $u_{ij}^{(2, \alpha)}$ and $\mathbf{W}^{(2, \alpha)} \in \mathbb{R}^{d \times d}$, extending the reach to indirect neighbors. These outputs are concatenated to form $\vec{c}_i = \text{Concat} \left(\vec{h}_i^{(2, 1)}, \dots, \vec{h}_i^{(2, 10)} \right)$, which is fed into a single L1-regularized dense fusion layer before the final classification. Unlike the design of Model 1, this version supports deeper intralayer learning up to 2-hop structures, but it maintains fusion only at the end. This approach is different from double-fusion approaches in that interlayer information is not integrated midway.

3.7.3. Model 3: Deep GAT with double fusion (Intermediate + late). Model 3 advances the previous designs by introducing an additional intermediate fusion stage between the two GAT layers, resulting in a double-fusion architecture that integrates cross-layer information both earlier and later while preserving modularity. After the first GAT layer computes 1-hop embeddings $\vec{h}_i^{(1, \alpha)}$ for each node i in each layer α (using attention weights $u_{ij}^{(1, \alpha)}$, weight matrices $\mathbf{W}^{(1, \alpha)}$, and ELU activation), these layer-specific embeddings are concatenated across all $L = 10$ layers to form $\vec{c}_i = \text{Concat} \left(\vec{h}_i^{(1, 1)}, \dots, \vec{h}_i^{(1, 10)} \right)$. This unified vector \vec{c}_i is then passed through the first L1-regularized dense fusion layer to produce an intermediate fused embedding $\vec{h}_i^{\text{(fused, 1)}}$ that is then input into the second GAT layer for each layer α , which computes 2-hop embeddings $\vec{h}_i^{(2, \alpha)}$ via updated attention weights $u_{ij}^{(2, \alpha)}$ and weight matrices $\mathbf{W}^{(2, \alpha)}$. These outputs are again concatenated across all layers to form a refined $\vec{c}_i = \text{Concat} \left(\vec{h}_i^{\text{(2, 1)}}, \vec{h}_i^{\text{(2, 2)}}, \dots, \vec{h}_i^{\text{(2, 10)}} \right)$, which is subsequently passed through a second L1-regularized fusion layer to yield the final fused embedding $\vec{h}_i^{\text{(fused, 2)}}$ for classification.

This double-fusion architecture enables both early and late integration of cross-layer information, supporting modular and interpretable learning. By leveraging multihead attention mechanisms (with $K=8$ heads), the model is able to discern diverse relational patterns within and across layers and to represent dynamic interactions such as brokerage roles that span the tensor's intralayer (N) and interlayer (I) components, as defined in Section 3.6.

In all three models, the L1 penalty promotes sparse representations that improve interpretability by prioritizing key features from the multilayer graph. Model 3 was selected as the final MMF-GAT architecture on the basis of its better performance in ablation experiments (reported in [Table 1](#)), particularly in representing complicated cross-layer patterns essential for epidemic forecasting tasks.

4. Methods: Application to COVID-19 Surveillance Data

4.1. Study population and data collection

Epidemiological surveillance data for social network analysis were collected by the Houston Health Department (HHD) between March 1, 2020, and December 31, 2021. Following the receipt of a positive PCR result, index cases were

Table 1. Ablation experiment results for the MMF-GAT architectures.

Model configuration	Accuracy	AUC	F1 score	PRAUC
Model 1: Shallow GAT with late fusion	0.76*	0.89*	0.68**	0.87*
Model 2: Deep GAT with late fusion	0.76**	0.89**	0.68***	0.88**
Model 3: Deep GAT with double fusion (Intermediate + late)	0.78	0.90	0.72	0.89

Note: Model 3 applies parallel GAT blocks across all network layers and incorporates both intermediate and late L1-regularized fusion layers, achieving superior performance on all evaluation metrics. The DeLong test was used for the AUC and PRAUC. For the accuracy and F1 score, paired *t* tests across 100 cross-validation folds were conducted. **p*<0.05; ***p*<0.01; ****p*<0.001.

<https://doi.org/10.1371/journal.pcsy.0000070.t001>

interviewed by HHD surveillance investigators to document health status, personal contacts, and potential exposures at identified outbreak sites. When the manual outreach capacity was exceeded, data collection was automated via third-party contractors and supplemented by Qualtrics and telephone surveys. All data, including PCR test results, addresses, and reported contacts, were consolidated within HHD’s Houston Electronic Disease Surveillance System (HEDSS). Among individuals present in all three network layers (personal contact, household coresidency, and outbreak site colocation), a final analytical cohort of 2,264 was identified on the basis of having at least one definitive PCR test recorded during Houston’s delta-variant surge (April 4–September 4, 2021). Any inconclusive test results were excluded from this classification.

4.2. Outcome measures

The primary outcome was COVID-19 positivity, defined as one or more positive PCR results recorded between April 4 and September 4, 2021.

4.3. Network layer construction and graph definitions

The construction of the multilayer network was guided by operational definitions and data collected by HHD for epidemiological surveillance. In the early stages of the pandemic, HHD used traditional contact tracing methods (i.e., direct interviews) to gather information. When the manual outreach capacity was exceeded, this process was supplemented by automated systems, including Qualtrics and telephone surveys. The following sections detail the construction of each network layer and provide its formal graph definition with a multilayer social graph framework.

4.3.1. Personal contact layer (*P*). The personal network layer *P* aggregates close physical contact data from contact tracing interviews. HHD defined “close contact” as being within 6 feet of an infected person for more than 15 minutes. The exposure window was defined as 24–48 hours prior to symptom onset for symptomatic cases or 24–48 hours prior to a positive test for asymptomatic cases. A system was in place to manage duplicated anonymous contacts by referencing identifiers (name, date of birth, address), with manual verification by a surveillance investigator to ensure data accuracy.

The personal contact graph is formally defined as $G_P = (V, E_P)$, where *V* is the set of individuals ($N=2,264$), and an undirected intralayer edge exists if individuals *i* and *j* are identified as close contacts, meaning that edge $e(i, j)$ is a member of edge set E_P . The network is represented by a symmetric adjacency matrix A_P , which represents intralayer connections.

4.3.2. Household coresidency layer (*H*). The coresidency layer *H* indicates potential exposure through shared living spaces. Coresidency was determined programmatically by matching addresses on the basis of ZIP codes, street numbers, apartment numbers, and street names (with a Levenshtein distance similarity index greater than 0.75). A matching algorithm was developed in Python via Pandas, NumPy, RecordLinkage, and Address libraries to extract and compare address components. Each household forms a fully connected subgraph (a clique), reflecting household-level modularity.

The household coresidency graph is defined as $G_H = (V, E_H)$, where an intralayer edge exists if individuals i and j share the same residential address, meaning that edge $e(i, j)$ is a member of edge set E_H . This layer is represented by a symmetric binary adjacency matrix A_H , which represents intralayer ties.

4.3.3. Person-outbreak site colocation layer (C). Site colocation layer C represents shared exposure at community outbreak sites, which were locally defined as locations with three or more linked COVID-19 cases. The network was constructed from a bipartite graph that consists of two sets of nodes: one set V representing the set of individuals ($N=2,264$) and another set S representing the set of outbreak sites ($M=367$). Edges connect individuals to sites on the basis of affiliation. The bipartite graph, defined as $B = (V, S, E_b)$, has an edge set E_b where each edge $e(i, s)$ connects an individual i from set V to an outbreak site s from set S . This structure models affiliations before projecting them into a one-mode person-person network to represent shared exposure at common outbreak sites. The projection was performed via matrix multiplication of B with its transpose (BB^T). Since most individuals (97.7%) were affiliated with one site, the resulting value graph was binarized such that the cell values became 1 if the original value was greater than or equal to 1 and 0 otherwise, forming fully connected colocation clusters.

The person-outbreak site colocation graph is defined as $G_C = (V, E_C)$, where an intralayer edge exists if individuals i and j have an affiliation with at least one common outbreak site, meaning that edge $e(i, j)$ is a member of edge set E_C . The layer is represented by a symmetric binary adjacency matrix A_C , reflecting intralayer connections.

4.3.4. Site-type colocation sublayers (S_k). To represent more specific modular structures, the community layer was decomposed into seven sublayers on the basis of the outbreak site type. The site types were categorized via Google Maps data as follows: (1) education centers (e.g., schools, universities, daycare), (2) assisted living centers (e.g., rehabilitation, transitional care), (3) detention centers (e.g., prisons, juvenile justice centers), (4) government locations, (5) charities/homeless organizations, (6) commercial/office locations (e.g., airports, UPS stores), or (7) healthcare facilities (e.g., hospitals, clinics). For each type k , where k ranges from 1–7, a type-specific bipartite graph is defined as B^k , consisting of V , S^k , and E_b^k . Here, V is the set of individuals ($N=2,264$), S^k is the set of outbreak sites belonging to type k , and E_b^k is the set of edges connecting individuals in V to outbreak sites in S^k on the basis of documented affiliation with at least one site of that type, where each edge $e(i, s)$ links an individual i to a specific site s of type k . This structure enables the model to isolate and examine interaction patterns unique to each category of the outbreak site. These bipartite $B^k \in N \times |S^k|$ are then projected into one-mode person–person matrices via matrix multiplication of B^k with its transpose and binarized to represent shared affiliations within each site-specific sublayer S_k . Each sublayer is denoted by a corresponding symmetric binary adjacency matrix A_S^k , which represents the pairwise connections within that specific site-type category.

4.4. Structural characteristics of the network layers

To assess the structural characteristics of the multilayer network, degree distributions for personal contact (P), household coresidency (H), and overlap ($P \cap H$) were computed for the full surveillance period (March 2020–December 2021) via the entire analytic sample ($N=2,264$). Although individuals were selected on the basis of their presence in all three layers, the analysis demonstrated significant network sparsity. This sparsity is explained by the temporal constraints of the outcome variable, i.e., individuals could have a degree of zero if their contacts were not tested during the delta surge period (April–September 2021), even if connections existed at other times. The analysis revealed that 47% of individuals had a degree of zero in the personal contact network, and 62% had a degree of zero in the overlaps ($P \cap H$).

In the denser household network, 63% of individuals had a degree between one and three. More reports with a detailed breakdown of these distributions, including histograms and stratification by infection status, can be found in the online supplementary material ([S1 Section](#), [S1 Fig](#), [S2 Fig](#), [S3 Fig](#), [S4 Fig](#), [S5 Fig](#), [S1 Table](#), [S2 Table](#), [S3 Table](#), [S4 Table](#), [S5 Table](#), [S1 Text](#), [S2 Text](#), [S3 Text](#), [S4 Text](#), [S5 Text](#)). Statistical significance was evaluated using the Wilcoxon signed-rank test for differences between the P and H networks, and the Wilcoxon rank-sum tests for comparisons of overlapping degrees by COVID-19 status.

4.5. Construction of node-level features for predictive modeling

The initial feature vector, x_i , for each individual i was constructed from a combination of risk factors for COVID-19, such as age, sex, and race, converting variables to appropriate measurement scales. Age was categorized into four levels (0–17, 18–44, 45–64, 65+), sex was encoded as a binary variable for males and females, and race (excluding Hispanic ethnicity due to data limitations) was one-hot encoded into multicategory variables with levels White, Black, Asian, and Other (including multiple races). From the personal contact network (P), three network features were extracted on the basis of network theory as predictors: the number of contacts (degree centrality), the interconnectedness among neighboring nodes (clustering coefficient), and the number of COVID-19 positive contacts to measure the network exposure level. From the household network (H), two coresidency features, the number of coresidents (household degree) and the number of COVID-19 positive coresidents (household exposure), were extracted. To represent location-based modularity at the node level within the community layer (C), seven binary dummy variables were created. These variables correspond to each outbreak site type and indicate whether an individual was affiliated with at least one site of a given type, also representing membership in aggregated outbreak-type modules.

These node features served as covariates in subsequent statistical machine learning and graph-based deep learning models to predict COVID-19 positivity. On the basis of the logistic regression results, two significant interaction terms were included: [1] the interaction between degree in the P network and affiliation with education centers and [2] the interaction between the number of infected coresidents in the H network and affiliation with education center outbreak sites. The detailed process of feature engineering via logistic regression is reported in the online supplementary material ([S3 Section](#)).

4.6. Experimental design and evaluation

The network data were preprocessed to prepare adjacency matrices for analysis. Normalization was applied to the P , H , and C matrices by converting them to a sparse format and row-normalizing them to address differences in network density. The adjacency matrices A_s^k for the seven site-type specific sublayers S_k are not normalized, as their binary structure (of all entries confined to the unit interval $[0, 1]$) does not require weight scaling. To evaluate the performance of the MMF-GAT model for COVID-19 prediction, the dataset ($N=2,264$) was randomly divided into training (70%, $N=1,585$) and testing (30%, $N=679$) sets.

The performance of the MMF-GAT was compared against that of five baseline models: logistic regression (LR), random forest (RF), graph convolutional network (GCN) [54], GAT without fusion, and supra-graph GAT [51]. Models were evaluated in terms of accuracy, area under the receiver operating characteristic curve (AUC), F1 score, and precision-recall AUC (PRAUC). The statistical significance of the performance differences was assessed via the DeLong test for AUC and PRAUC and paired t tests for accuracy and F1 scores across cross-validation folds. The negative log-likelihood function was used as the loss function, which is mathematically equivalent to optimizing the cross-entropy loss for a multiclass classification problem where the final layer uses a log softmax activation function.

The hyperparameters for the GCN and GAT baseline models were optimized via grid search. For the GCN, the following settings were used: learning rate = 0.001, weight decay = 0.005, dropout = 0.5, hidden dimension = 64, and early stopping patience = 50. For the GAT, the configuration included a learning rate = 0.001, weight decay = 5×10^{-5} , dropout = 0.3, hidden dimension = 64, number of attention heads = 12, and attention slope (α) = 0.2. For the proposed GAT fusion model, the optimal hyperparameters were determined as follows: training epochs = 500; learning rate = 0.0005; weight decay = 5×10^{-5} ; dropout = 0.3; hidden dimensions (layers 1 and 2) = 16; fusion layer dimension = 4; number of attention heads = 8; attention slope (α) = 0.2; L1 regularization coefficient (λ) = 0.0001; and early stopping patience = 50.

5. Results

This section presents the findings from the application of the multilayer modular fusion graph attention network (MMF-GAT) model to COVID-19 surveillance data from Houston, Texas. The results demonstrate the predictive capability of our proposed model for infection status within a multilayer social framework, covering descriptive network visualizations, logistic regression outcomes, ablation experiments, model performance evaluations, and analysis of feature importance.

5.1. Descriptive statistics

Among the 2,264 individuals, 797 (35%) tested positive for COVID-19, and 1,467 (65%) tested negative. There were significant differences in infection rates across age groups: the youngest group had a 40% infection rate, whereas the 18–44 age group had a significantly lower infection rate of 30% ($p=0.0005$), and the 65+ group had an even lower infection rate of 22% ($p=0.016$). The “Other” racial/ethnic category also had a significantly lower infection rate of 30% ($p=0.038$). No statistically significant differences were observed by sex.

In the personal contact network (P), t tests revealed that infected individuals had a significantly greater average degree (1.47 vs. 0.71; $p<0.0001$) and more infected contacts (0.88 vs. 0.32; $p<0.001$). Similarly, in the household network (H), infected individuals had a higher household degree (0.99 vs. 0.69; $p<0.0001$) and greater exposure to infected coresidents (0.69 vs. 0.16; $p<0.001$). With respect to outbreak site affiliations, settings such as detention centers (16%; $p=0.024$) and assisted living facilities (20%; $p=0.001$) had significantly lower infection rates. In contrast, commercial offices (53% infected; $p=0.036$) and government facilities (56% infected; $p=0.020$) presented higher infection rates. However, the small sample sizes for these latter two settings ($N=38$ and $N=32$) warrant caution in interpretation. Descriptive statistics for the study sample across all features can be found in the online supplementary material ([S2 Section](#), [S6 Table](#), [S6 Text](#)).

5.2. Visualization of the multilayered network structure

The multilayer networks were visualized via the *igraph* package in R [55]. [Fig 4](#) visually demonstrates how integrating multiple social contexts refines the network’s structure. Starting with the personal contact network (P) in Panel A, distinct social clusters are visible. Adding household coresident data (H) in Panel B increases network density within these clusters by introducing additional ties. Incorporating outbreak site colocation information (C) in Panel C creates new connections that bridge previously separate clusters, enhancing connectivity across the network. Panel D further categorizes these colocation ties by outbreak site type (S_k), showing that education centers (purple nodes) and assisted living facilities (red nodes) are the predominant settings forming these bridging links. These layers demonstrate how household affiliations strengthen local cohesion, whereas community affiliations promote broader integration through institutional settings.

5.3. Results of important network features from logistic regression analysis

A logistic regression analysis was conducted to identify key predictors and inform the feature set for the MMF-GAT model. The main model (Model 1) revealed that a higher degree in the personal contact network (P) was associated with increased odds of infection ($OR = 1.72$, $p<0.001$), as was a greater number of infected coresidents ($OR = 5.57$, $p<0.001$). Conversely, the 18–44 ($OR = 0.64$, $p<0.01$) and 65+ ($OR = 0.39$, $p<0.05$) age groups had lower odds of infection than did the 0–17 reference group. Affiliation with certain outbreak sites, such as education centers ($OR = 0.51$, $p<0.01$) and detention centers ($OR = 0.31$, $p<0.05$), was also associated with lower odds of infection. The second model (Model 2) incorporated significant interaction terms, demonstrating the combined effects of personal contact degree with education center affiliation ($P \times S_1$) ($OR = 1.42$, $p<0.05$) and the number of infected coresidents with education center affiliation ($H \times S_1$) ($OR = 2.74$, $p<0.05$). The results can be found in the online supplementary material ([S7 Table](#), [S7 Text](#)).

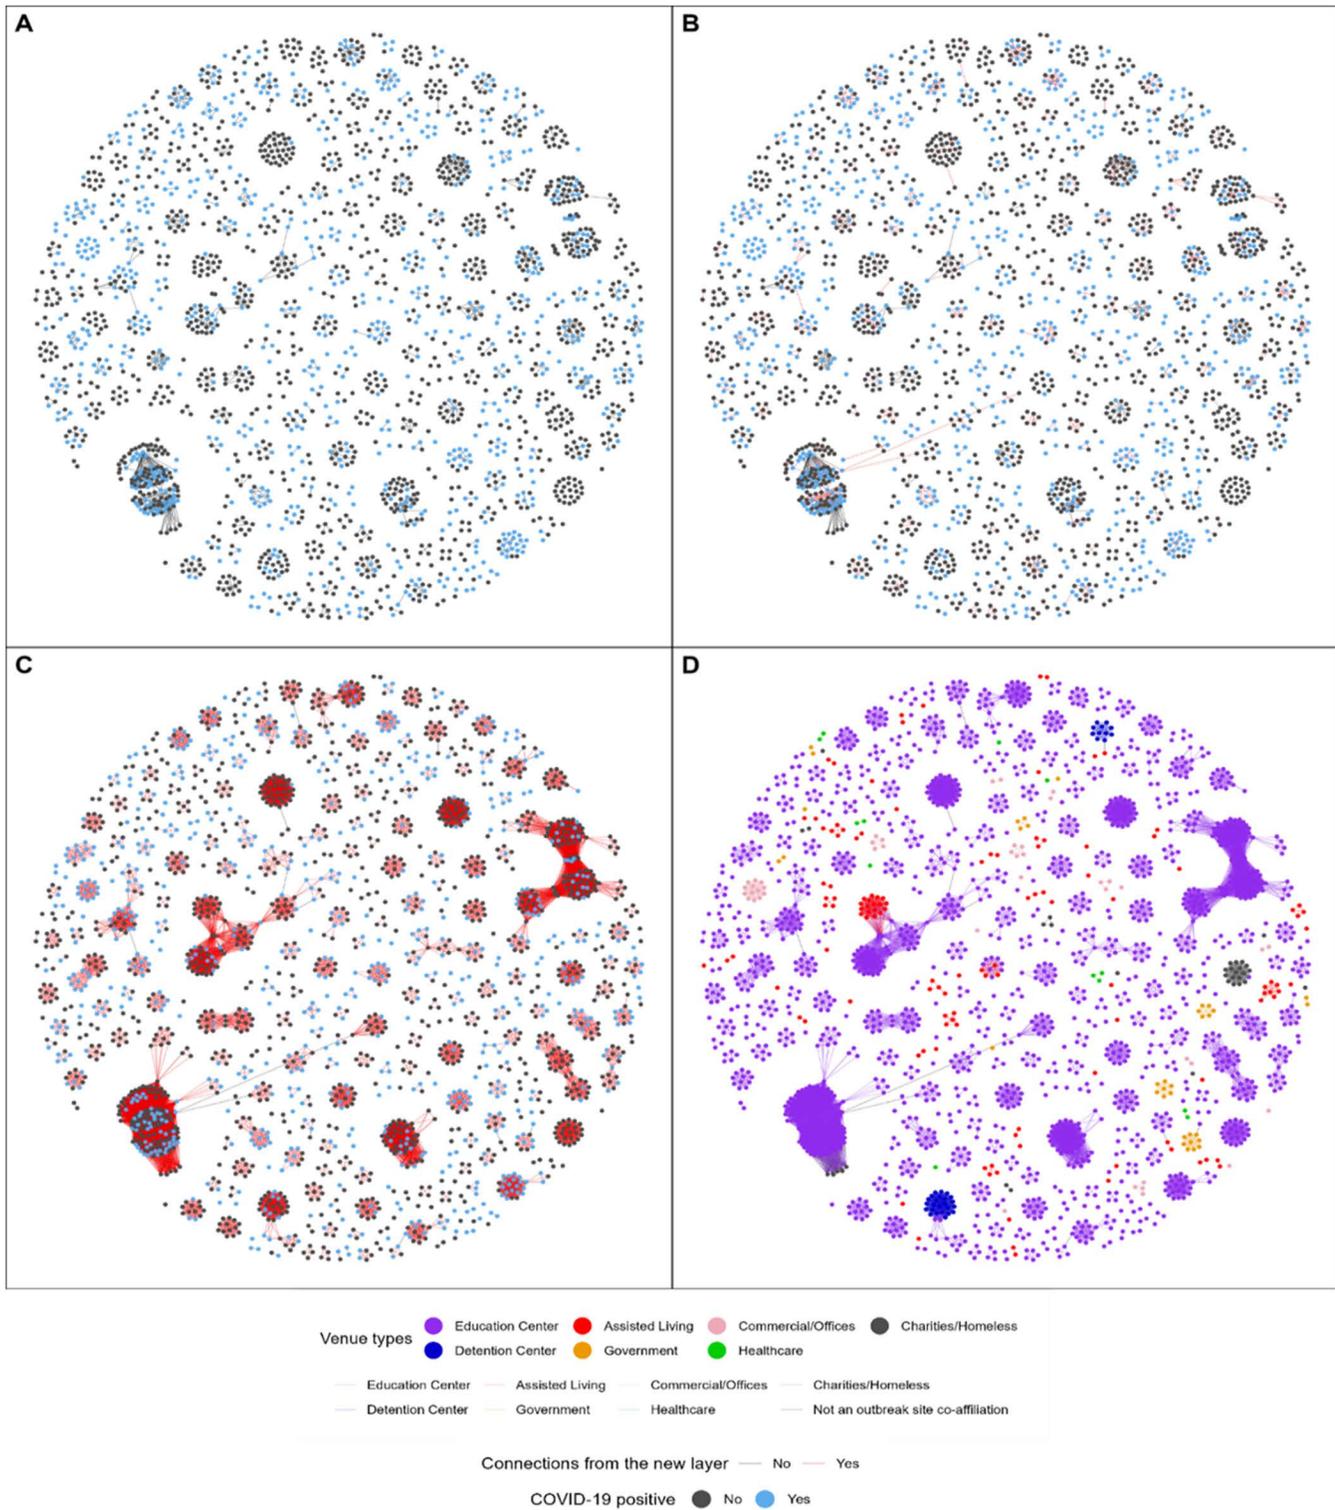


Fig 4. Progressive multilayer network and transmission channels.

<https://doi.org/10.1371/journal.pcsy.0000070.g004>

5.4. Performance outcomes of the ablation experiment

As presented in [Table 1](#), Model 3 consistently outperformed the other configurations across multiple metrics, achieving an accuracy of 0.78, an AUC of 0.90, an F1 score of 0.72, and a PRAUC of 0.89. These results support the superiority of its multiphase processing approach, which preserves the modular organization of social contexts such as households, outbreak sites, and specific site types by processing each layer independently before integration. This method also enables the model to detect patterns associated with bridging roles across networks through its multihop strategy.

Model 3 was selected as the optimal architecture because of its ability to balance comprehensive integration with computational focus on significant features. By employing two L1-regularized fusion layers, it effectively consolidates embeddings, emphasizing the most relevant social interactions for prediction. Although this approach increases computational demands, it addresses limitations in multilayer network frameworks for epidemic modeling by preserving the unique structures of community contexts [\[20,22\]](#). This design improves predictive accuracy and offers a computational framework to better comprehend the complex social dynamics driving disease transmission.

5.5. Performance evaluation of the MMF-GAT model

[Table 2](#) compares the performance of the MMF-GAT model against baseline models, including logistic regression (LR), random forest (RF), graph convolutional network (GCN), GAT without fusion, and supra-graph GAT. The MMF-GAT model ($P+H+C+S_k$) achieved superior results, with an accuracy of 0.78, an AUC of 0.90, an F1 score of 0.72, and a PRAUC of 0.89. The performance slightly decreased in 100-fold cross-validation (accuracy=0.77, AUC=0.83, PRAUC=0.76), perhaps due to dataset imbalances and variability in test set sizes. Statistical tests (DeLong for AUC and PRAUC; paired t tests for accuracy and F1 score) confirmed the model's significant improvement over baselines ($p<0.05$ to $p<0.001$). Adding household (H) and outbreak site (C) layers incrementally enhanced performance, with notable gains in the F1 score and AUC, indicating the value of multilayer data integration.

5.6. XAI feature analysis

[Table 3](#) presents the distribution of COVID-19 positive individuals located at various outbreak site types, illustrating the prevalence of disease transmission at specific locations. Education centers accounted for the highest frequency of collocation events, representing 87.0% of all individuals, 87.4% of empirical positive cases, and 92.1% of predicted positive cases. Smaller outbreak sites, such as assisted living facilities (4.6%), commercial offices (1.7%), and detention centers (1.7%), had lower representation but still contributed to transmission dynamics.

[Fig 5](#). illustrates the relative importance of features identified by the GNNExplainer applied to the MMF-GAT model.

Table 2. Model performance and node classification.

Model structures	Accuracy	AUC	F1 score	PRAUC
Logistic regression (Model 2)	0.67***	0.76***	0.53***	0.69***
Random forest	0.68***	0.72***	0.56***	0.69***
GCN Model (P)	0.68***	0.79***	0.58***	0.78***
GAT Model (P)	0.72**	0.85**	0.61***	0.82**
MMF-GAT model ($P+H$)	0.73***	0.84**	0.65***	0.81***
MMF-GAT model ($P+H+C$)	0.75**	0.87**	0.68**	0.83***
Supra-graph GAT model ($P+H+C+S_k$)	0.69***	0.67***	0.58***	0.74***
MMF-GAT model ($P+H+C+S_k$)	0.78	0.90	0.72	0.89

Note: Statistical significance was assessed against the MMF-GAT model. The DeLong test was used for the AUC and PRAUC. For the accuracy and F1 score, paired t tests across 100 cross-validation folds were conducted. * $p<0.05$; ** $p<0.01$; *** $p<0.001$.

<https://doi.org/10.1371/journal.pcsy.0000070.t002>

Table 3. Empirical and MMF-GAT predicted COVID-19 positive cases by outbreak site type.

All individuals			Empirical positive cases		MMF-GAT predicted cases	
Outbreak site type	Count	Percentage	Count	Percentage	Count	Percentage
Education center	1,899	87.0%	663	87.4%	385	92.1%
Assisted living	101	4.6%	20	2.6%	5	1.2%
Commercial/offices	38	1.7%	20	2.6%	4	1.0%
Detention center	37	1.7%	6	0.8%	0	0.0%
Government	42	1.5%	18	2.4%	10	2.4%
Charities/homeless	68	3.1%	27	3.6%	12	2.9%
Healthcare	9	0.4%	5	0.7%	2	0.5%

<https://doi.org/10.1371/journal.pcsy.0000070.t003>

Ranked Feature Importance Scores

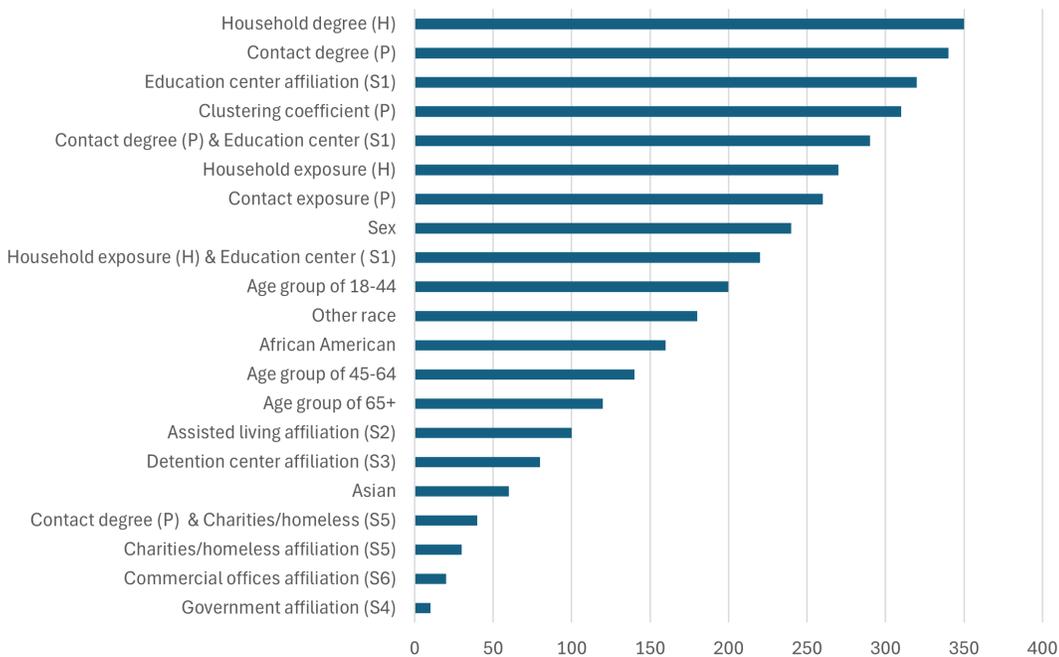


Fig 5. Ranked feature importance scores estimated by the MMF-GAT model.

<https://doi.org/10.1371/journal.pcsy.0000070.g005>

Among the 21 features considered, the top 10 relative feature importance scores ranged from 10–350. The distribution of these scores had a mean of 178, a median of 180, and an interquartile range from 40 (25th percentile) to 310 (75th percentile). Household degree (*H*) emerged as the most influential predictor, followed by personal contact degree (*P*), education center affiliation (*S*₁), and the clustering coefficient (*P*). The interaction effect between the degree of personal contact and education center affiliation (*P* × *S*₁) was also ranked highly. Additional important features included exposure to infected coresidents (*H*), exposure to infected personal contacts (*P*), sex, interaction between infected coresidents and education center affiliation (*H* × *S*₁), and 18–44 years of age. These findings illustrate how infection prediction depends on a combination of network structures, institutional contexts, and demographic factors, reflecting the multilayer modular nature of social interactions.

6. Discussion

This study introduced the multilayer modular fusion graph attention network (MMF-GAT), a framework designed not to uncover universal laws of disease transmission but to serve as a powerful predictive tool for operational public health. Architecture, which processes each social layer independently before fusion, is a deliberate choice grounded in the sociological principles of modularity and brokerage. By preserving the distinct interaction patterns within personal networks and household and community sites, the MMF-GAT avoids the premature flattening of the social context common in supra-graph methods. Our results confirm that this design improves predictive performance, providing a more accurate assessment of individual infection risk.

The XAI analysis identified household degree as the most influential predictor. The apparent contradiction between descriptive statistics (where a higher degree correlated with infection) and the adjusted logistic regression model (where it was negatively associated) suggests a complex relationship. After controlling for known exposures such as infected coresidents, the model indicates higher transmission intensity in smaller, more crowded households where isolation is difficult, which is consistent with an observed tendency reported in other studies [56]. However, this interpretation remains speculative and could be influenced by artifacts in the address-matching algorithm.

The second most important feature, personal contact degree, must be interpreted with extreme caution. Its high importance is likely an artifact of the contact tracing-based sampling design, which overrepresents individuals who have had more contact with positive cases. This structural bias means that the association is a feature of the data collection method rather than a generalizable epidemiological finding.

Affiliations with education centers ranked third, underlying their role as institutional hubs for sustained interaction. This is consistent with evaluation studies on school closures, which emphasize that schools are high-risk sites for transmission due to dense and repeated contacts. For example, closures have been shown to reduce community incidence among school-age children during periods of high transmission, as they disrupt interaction patterns and limit exposure opportunities [57]. These findings explain the elevated infection risk for individuals affiliated with education centers in our model. Targeting specific modular structures within schools, such as grades, classrooms, courses, or extracurricular activities, could further minimize risks through tailored interventions. Our results support alternatives to broad closure, such as regular testing, which can effectively mitigate transmission to preserve educational continuity.

Interaction terms, particularly between personal contact degree and education center affiliation, illustrate how an individual's risk is impacted by their connections across multiple social layers. Higher personal contact in school settings can exacerbate exposure, leading to large variability in transmission where some individuals disproportionately drive the spread [57]. This underlines the need for targeted strategies that address both direct contacts and institutional affiliations to effectively mitigate spread.

However, the predictive power of network features is not absolute. For example, the interaction between infected coresidents and education center affiliation ranked ninth, below the demographic variable of sex (ranked eighth). This finding suggests that individual attributes can sometimes carry greater predictive weights than can specific cross-layer structural metrics.

This study has several limitations. First, the primary limitation originates from the sampling bias in contact tracing data, which confines the findings to predictive applications. Second, potential confounding factors may also pose a concern. For example, affiliation with a high-density setting such as schools may correlate with age and increased contact reporting. Third, the computational demands of the model could hinder real-time scalability, although aggregating network data over two years might mask important temporal dynamics. Fourth, the presence of individuals with a degree of zero, which reflects the temporal filtering of the outcome variable rather than actual social isolation, may limit the interpretation of network-based transmission dynamics. Finally, although GNExplainer produces a useful ranking of feature importance, the absence of a predefined statistical threshold for significance limits its precision in distinguishing key features from

noise. Future studies could improve the overall transparency of the model by exploring advanced methods to distinguish truly influential features from background noise, with the potential for improved interpretability by integrating graph-based techniques with signed attributions that measure the directional impact of features [58].

7. Conclusion

This study successfully developed and validated the multilayer modular fusion graph attention network (MMF-GAT), a framework that integrates personal, residential, and community network data to predict the COVID-19 infection status. The MMF-GAT demonstrated strong predictive performance, significantly outperforming multiple baseline models via an architecture designed to computationally preserve the modular and overlapping nature of human social systems. The primary contribution of this study is its practical utility for public health. Although the built-in biases in contact tracing data limit the model from being used to infer causal mechanisms of disease transmission, its ability to accurately identify individuals and contexts associated with a higher predicted risk makes it a valuable operational tool. By using this framework, public health agencies can enhance situational awareness and strategically allocate limited resources, such as by prioritizing contact tracing efforts or targeting testing campaigns. This study provides a validated computational method and aims to make methodological contributions to the fields of network science and public health informatics.

Supporting information

S1 Table. Descriptive Statistics of the Degree Distribution Stratified by COVID-19 Status. Descriptive statistics (mean, standard deviation, median, minimum, maximum) for degree in personal contact network (P), household coresidency network (H), and overlapping personal-household network ($P \cap H$), stratified by COVID-19 infection status (positive, negative, and overall).

(PDF)

S2 Table. Proportion of Degree in Personal Contact Network (P). Proportion (%) of individuals at each degree in the personal contact network (P), for the total study sample.

(PDF)

S3 Table. Proportion of Degree in Household Network (H). Proportion (%) of individuals with each degree in the household coresidency network (H), for the total study sample.

(PDF)

S4 Table. Proportion of Overlapping Degree in $P \cap H$. Proportion (%) of individuals with each degree in the overlapping personal contact and household network ($P \cap H$), for the total study sample.

(PDF)

S5 Table. Wilcoxon Signed-Rank and Rank Sum Test Results for Network Degrees. Wilcoxon signed-rank test for differences between the degree distribution of P and H networks; Wilcoxon rank-sum test for overlapping degree in $P \cap H$ by COVID-19 status. Statistical significance indicated at $p < 0.05$, 0.01 , and 0.001 .

(PDF)

S1 Fig. Histogram of P Network Degrees Stratified by COVID-19 Infection Status. The histogram illustrates the distribution of degrees in the persona network (P), stratified by individuals' COVID-19 infection status.

(PDF)

S2 Fig. Histogram of H Network Degrees Stratified by COVID-19 Infection Status. The histogram illustrates the distribution of degrees in the household coresidency network (H), stratified by individuals' COVID-19 infection status.

(PDF)

S3 Fig. Histogram of $P \cap H$ Network Degrees Stratified by COVID-19 Infection Status. The histogram illustrates the distribution in the overlapping network ($P \cap H$), stratified by COVID-19 infection status.

(PDF)

S4 Fig. Histogram of P Network Degrees (Overall). The histogram illustrates the distribution of the personal network (P) across the total study population.

(PDF)

S5 Fig. Histogram of H Network Degrees (Overall). The histogram illustrates the distribution of the household co-residency network (H) across the total study population.

(PDF)

S6 Table. Descriptive Statistics of the Study Sample. Summary of demographic, network, and outbreak site attributes for COVID-19 positive, negative, and overall study population, with p values for group comparisons.

(PDF)

S7 Table. Logistic Regression Results for COVID-19 Positivity. Logistic regression odds ratios and confidence intervals for sociodemographic, network, and outbreak site predictors of COVID-19 positivity, as well as significant interaction terms used for feature engineering.

(PDF)

S1 Text. Descriptive Statistics of the Degree Distribution Stratified by COVID-19 Status.

(PDF)

S2 Text. Proportion of Degree in Personal Contact Network (P).

(PDF)

S3 Text. Proportion of Degree in Household Network (H).

(PDF)

S4 Text. Proportion of Overlapping Degree in $P \cap H$.

(PDF)

S5 Text. Wilcoxon Signed-Rank and Sum Test Results for Network Degrees.

(PDF)

S6 Text. Descriptive Statistics for Study Sample.

(PDF)

S7 Text. Logistic Regression Results for COVID-19 Positivity.

(PDF)

S1 Section. Results from the Degree Distribution Analysis.

(PDF)

S2 Section. Descriptive Statistics for Study Sample.

(PDF)

S3 Section. Process of Feature Engineering via Logistic Regression with the Results.

(PDF)

S1 Striking Image. Visualization of multilayered networks using Multilayer Modular Fusion Graph Attention Networks (MMF-GAT) to predict epidemic dynamics. Image generated with the assistance of Nano Banana (an AI generator); finalized and curated by the authors. Licensed under Creative Commons Attribution (CC BY) License. (TIFF)

Acknowledgments

We acknowledge staff from the Houston Health Department for collecting the data. During the preparation of this work, the author(s) used artificial intelligence (AI)-assisted technologies (such as large language models) to improve language, structure, and readability. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

Author contributions

Conceptualization: Kayo Fujimoto.

Data curation: Lizhong Liu, Jacky Kuo, Hei Kit Chan.

Formal analysis: Lizhong Liu, Jacky Kuo.

Funding acquisition: Justin Bahl, Kayo Fujimoto.

Investigation: Kayo Fujimoto, Jacky Kuo, Michelle Carr, Roger Sealy.

Methodology: Kayo Fujimoto, Lizhong Liu, Jacky Kuo, Baode Gao, Xi Luo.

Project administration: Roger Sealy, Justin Bahl.

Resources: Michelle Carr, Roger Sealy.

Supervision: Kayo Fujimoto, Xi Luo, Justin Bahl.

Validation: Baode Gao, Hei Kit Chan.

Writing – original draft: Kayo Fujimoto, Lizhong Liu, Jacky Kuo.

References

1. Eames KTD, Keeling MJ. Modeling dynamic and network heterogeneities in the spread of sexually transmitted diseases. *Proc Natl Acad Sci U S A*. 2002;99(20):13330–5. <https://doi.org/10.1073/pnas.202244299> PMID: [12271127](https://pubmed.ncbi.nlm.nih.gov/12271127/)
2. Keeling MJ, Rohani P. Modeling infectious diseases in humans and animals. Princeton Univ Press. 2008.
3. Barabási AL. Scale-free networks: a decade and beyond. *Science*. 2009;325(5939):412–3. <https://doi.org/10.1126/science.1173299> PMID: [19628854](https://pubmed.ncbi.nlm.nih.gov/19628854/)
4. Barabási AL, Albert R. Emergence of scaling in random networks. *Science*. 1999;286(5439):509–12. <https://doi.org/10.1126/science.286.5439.509> PMID: [10521342](https://pubmed.ncbi.nlm.nih.gov/10521342/)
5. Watts DJ, Strogatz SH. Collective dynamics of “small-world” networks. *Nature*. 1998;393(6684):440–2. <https://doi.org/10.1038/30918> PMID: [9623998](https://pubmed.ncbi.nlm.nih.gov/9623998/)
6. Newman MEJ. Modularity and community structure in networks. *Proc Natl Acad Sci U S A*. 2006;103(23):8577–82. <https://doi.org/10.1073/pnas.0601602103> PMID: [16723398](https://pubmed.ncbi.nlm.nih.gov/16723398/)
7. Newman MEJ, Girvan M. Finding and evaluating community structure in networks. *Phys Rev E Stat Nonlin Soft Matter Phys*. 2004;69(2 Pt 2):026113. <https://doi.org/10.1103/PhysRevE.69.026113> PMID: [14995526](https://pubmed.ncbi.nlm.nih.gov/14995526/)
8. Breiger RL. The duality of persons and groups. *Soc Forces*. 1974;53(2):181–90.
9. Feld SL. The focused organization of social ties. *Am J Sociol*. 1981;86(5):1015–35.
10. Feld SL. Social structural determinants of similarity among associates. *Am Sociol Rev*. 1982;47:797–801.
11. Veličković P, Cucurull G, Casanova A, Romero A, Liò P, Bengio Y. In: Proceedings of the International Conference on Learning Representations (ICLR); 2018 Apr 30-May 3; Vancouver, BC, Canada. Available from: <https://openreview.net/forum?id=rJXMpikCZ>.

12. Sleeman WCI, Kapoor R, Ghosh P. Multimodal classification: Current landscape, taxonomy and future directions. *ACM Comput Surv.* 2022;55(7):1–31.
13. Ying Z, Bourgeois D, You J, Zitnik M, Leskovec J. GNNExplainer: Generating explanations for graph neural networks. In: *Advances in Neural Information Processing System 32 (NeurIPS 2019)*; 2019 Dec 8–14, Vancouver, BC, Canada. Available from: https://proceedings.neurips.cc/paper_files/paper/2019/file/d80b7040b773199015de6d3b4293c8ff-Paper.pdf.
14. Barabási AL, Albert R, Jeong H. Scale-free characteristics of random networks: the topology of the worldwide web. *Physica A.* 2000;281(1–4):69–77.
15. Pastor-Satorras R, Vespignani A. Epidemic spreading in scale-free networks. *Phys Rev Lett.* 2001;86(14):3200–3. <https://doi.org/10.1103/PhysRevLett.86.3200> PMID: [11290142](https://pubmed.ncbi.nlm.nih.gov/11290142/)
16. Newman MEJ. Assortative mixing in networks. *Phys Rev Lett.* 2002;89(20):208701. <https://doi.org/10.1103/PhysRevLett.89.208701> PMID: [12443515](https://pubmed.ncbi.nlm.nih.gov/12443515/)
17. Herrmann HA, Schwartz J-M. Why COVID-19 models should incorporate the network of social interactions. *Phys Biol.* 2020;17(6):065008. <https://doi.org/10.1088/1478-3975/aba8ec> PMID: [32702678](https://pubmed.ncbi.nlm.nih.gov/32702678/)
18. Keeling MJ, Eames KTD. Networks and epidemic models. *J R Soc Interface.* 2005;2(4):295–307. <https://doi.org/10.1098/rsif.2005.0051> PMID: [16849187](https://pubmed.ncbi.nlm.nih.gov/16849187/)
19. Salathé M, Jones JH. Dynamics and control of diseases in networks with community structure. *PLoS Comput Biol.* 2010;6(4):e1000736. <https://doi.org/10.1371/journal.pcbi.1000736> PMID: [20386735](https://pubmed.ncbi.nlm.nih.gov/20386735/)
20. Kinsley AC, Rossi G, Silk MJ, VanderWaal K. Multilayer and multiplex networks: an introduction to their use in veterinary epidemiology. *Front Vet Sci.* 2020;7:596. <https://doi.org/10.3389/fvets.2020.00596> PMID: [33088828](https://pubmed.ncbi.nlm.nih.gov/33088828/)
21. Nematzadeh A, Ferrara E, Flammini A, Ahn Y-Y. Optimal network modularity for information diffusion. *Phys Rev Lett.* 2014;113(8):088701. <https://doi.org/10.1103/PhysRevLett.113.088701> PMID: [25192129](https://pubmed.ncbi.nlm.nih.gov/25192129/)
22. Ma J, Wang P. Epidemic spreading on multilayer community networks. *Phys Lett A.* 2025;532:130199.
23. Sah P, Leu ST, Cross PC, Hudson PJ, Bansal S. Unraveling the disease consequences and mechanisms of modular structure in animal social networks. *Proc Natl Acad Sci U S A.* 2017;114(16):4165–70. <https://doi.org/10.1073/pnas.1613616114> PMID: [28373567](https://pubmed.ncbi.nlm.nih.gov/28373567/)
24. Stegehuis C, van der Hofstad R, van Leeuwen JSH. Epidemic spreading on complex networks with community structures. *Sci Rep.* 2016;6:29748. <https://doi.org/10.1038/srep29748> PMID: [27440176](https://pubmed.ncbi.nlm.nih.gov/27440176/)
25. Gross B, Havlin S. Epidemic spreading and control strategies in spatial modular network. *Appl Netw Sci.* 2020;5(1):95. <https://doi.org/10.1007/s41109-020-00337-4> PMID: [33263074](https://pubmed.ncbi.nlm.nih.gov/33263074/)
26. Landry NW, Restrepo JG. Opinion disparity in hypergraphs with community structure. *Phys Rev E.* 2023;108(3–1):034311. <https://doi.org/10.1103/PhysRevE.108.034311> PMID: [37849151](https://pubmed.ncbi.nlm.nih.gov/37849151/)
27. Volz EM, Miller JC, Galvani A, Ancel Meyers L. Effects of heterogeneous and clustered contact patterns on infectious disease dynamics. *PLoS Comput Biol.* 2011;7(6):e1002042. <https://doi.org/10.1371/journal.pcbi.1002042> PMID: [21673864](https://pubmed.ncbi.nlm.nih.gov/21673864/)
28. Friedman SR, Kottiri BJ, Neaigus A, Curtis R, Vermund SH, Des Jarlais DC. Network-related mechanisms may help explain long-term HIV-1 seroprevalence levels that remain high but do not approach population-group saturation. *Am J Epidemiol.* 2000;152(10):913–22. <https://doi.org/10.1093/aje/152.10.913> PMID: [11092433](https://pubmed.ncbi.nlm.nih.gov/11092433/)
29. Newman ME, Watts DJ. Scaling and percolation in the small-world network model. *Phys Rev E Stat Phys Plasmas Fluids Relat Interdiscip Topics.* 1999;60(6 Pt B):7332–42. <https://doi.org/10.1103/physreve.60.7332> PMID: [11970678](https://pubmed.ncbi.nlm.nih.gov/11970678/)
30. Watts DJ, Muhamad R, Medina DC, Dodds PS. Multiscale, resurgent epidemics in a hierarchical metapopulation model. *Proc Natl Acad Sci U S A.* 2005;102(32):11157–62.
31. Mossong J, Hens N, Jit M, Beutels P, Auranen K, Mikolajczyk R, et al. Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLoS Med.* 2008;5(3):e74. <https://doi.org/10.1371/journal.pmed.0050074> PMID: [18366252](https://pubmed.ncbi.nlm.nih.gov/18366252/)
32. Robins G, Lusher D, Broccatelli C, Bright D, Gallagher C, Karkavandi MA, et al. Multilevel network interventions: goals, actions, and outcomes. *Soc Networks.* 2023;72:108–20. <https://doi.org/10.1016/j.socnet.2022.09.005> PMID: [36188126](https://pubmed.ncbi.nlm.nih.gov/36188126/)
33. Artime O, Benigni B, Bertagnolli G, d'Andrea V, Gallotti R, Ghavasieh A. *Multilayer network science: from cells to societies.* Cambridge Univ Press. 2022.
34. Kivelä M, Arenas A, Barthelemy M, Gleeson JP, Moreno Y, Porter MA. Multilayer networks. *J Complex Netw.* 2014;2(3):203–71.
35. Fügenschuh M, Fu F. Overcoming vaccine hesitancy by multiplex social network targeting: an analysis of targeting algorithms and implications. *Appl Netw Sci.* 2023;8(1):67. <https://doi.org/10.1007/s41109-023-00595-y> PMID: [37745797](https://pubmed.ncbi.nlm.nih.gov/37745797/)
36. Bourdieu P. The forms of capital. In: Richardson J, editor. *Handbook of theory and research for the sociology of education.* New York: Greenwood; 1986. p. 241–258.
37. Lamont M, Molnár V. The study of boundaries in the social sciences. *Annu Rev Sociol.* 2002;28(1):167–95.
38. Bourdieu P. *Distinction: a social critique of the judgment of taste.* Cambridge (MA): Harvard University Press; 1984.
39. McPherson M, Smith-Lovin L, Cook JM. Birds of a feather: homophily in social networks. *Annual Review of Sociol.* 2001;27(1):415–44.

40. Link BG, Phelan J. Social conditions as fundamental causes of disease. *J Health Soc Behav.* 1995;35:80–94.
41. Pescosolido BA. Beyond rational choice: the social dynamics of how people seek help. *Am J Sociol.* 1992;97(4):1096–138.
42. Williams DR, Collins C. US socioeconomic and racial differences in health: patterns and explanations. *Annu Rev Sociol.* 1995;21(1):349–86.
43. Bourdieu P. *The logic of practice.* Stanford (CA): Stanford Univ Press; 1990.
44. Gould RV, Fernandez RM. Structures of mediation: a formal approach to brokerage in transaction networks. *Sociol Methodol.* 1989;89–126.
45. Fujimoto K, Hallmark CJ, Mauldin RL, Kuo J, Smith C, Del Vecchio N, et al. Brokerage-centrality conjugates for multilevel organizational field networks: toward a blockchain implementation to enhance coordination of healthcare delivery. In: Weber MS, Yanovitzky I, editors. *Networks, knowledge brokers, and the public policymaking process.* Cham (Switzerland): Palgrave Macmillan; 2021. p. 265–314.
46. Freeman LC. Centrality in social networks conceptual clarification. *Soc Netw.* 1978;1(3):215–39.
47. Watts DJ. Networks, dynamics, and the small-world phenomenon. *Am J Sociol.* 1999;105(2):493–527.
48. Valente TW. Network models and methods for studying the diffusion of innovations. *Models and methods in social network analysis.* 2005. p. 98–116.
49. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I. Attention is all you need. *Adv Neural Inf Process Syst.* 2017;30(1):5998–6008.
50. Shanthamallu US, Thiagarajan JJ, Song H, Spanias A. GrAMME: Semisupervised learning using multilayered graph attention models. *IEEE Trans Neural Netw Learn Syst.* 2020;31(10):3977–88. <https://doi.org/10.1109/TNNLS.2019.2948797> PMID: [31725400](https://pubmed.ncbi.nlm.nih.gov/31725400/)
51. Xiang Y, Fujimoto K, Li F, Wang Q, Del Vecchio N, Schneider J, et al. Identifying influential neighbors in social networks and venue affiliations among young MSM: a data science approach to predict HIV infection. *AIDS.* 2021;35(Suppl 1):S65–73. <https://doi.org/10.1097/QAD.0000000000002784> PMID: [33306549](https://pubmed.ncbi.nlm.nih.gov/33306549/)
52. Fey M, Lenssen JE. Fast graph representation learning with PyTorch Geometric. *arXiv preprint.* 2019. <https://doi.org/10.48550/arXiv.1903.02428>
53. Yu E, Du J, Xiang Y, Hu X, Feng J, Luo X, et al. Explainable artificial intelligence and domain adaptation for predicting HIV infection with graph neural networks. *Ann Med.* 2024;56(1):2407063. <https://doi.org/10.1080/07853890.2024.2407063> PMID: [39417227](https://pubmed.ncbi.nlm.nih.gov/39417227/)
54. Zhang S, Tong H, Xu J, Maciejewski R. Graph convolutional networks: a comprehensive review. *Comput Soc Netw.* 2019;6(1):11. <https://doi.org/10.1186/s40649-019-0069-y> PMID: [37915858](https://pubmed.ncbi.nlm.nih.gov/37915858/)
55. Csardi G, Nepusz T, Traag V, Horvát S, Zanini F, Noom D, Müller K. *igraph: network analysis and visualization in R [computer program].* R package version 1.5.1; 2023. <https://doi:10.5281/zenodo.8240644>
56. Chung MK, Hart B, Santillana M, Patel CJ. Pediatric and young adult household transmission of the initial waves of SARS-CoV-2 in the United States: administrative claims study. *J Med Internet Res.* 2024;26:e44249. <https://doi.org/10.2196/44249> PMID: [37967280](https://pubmed.ncbi.nlm.nih.gov/37967280/)
57. Colosi E, Bassignana G, Contreras DA, Poirier C, Boëlle P-Y, Cauchemez S, et al. Screening and vaccination against COVID-19 to minimise school closure: a modelling study. *Lancet Infect Dis.* 2022;22(7):977–89. [https://doi.org/10.1016/S1473-3099\(22\)00138-4](https://doi.org/10.1016/S1473-3099(22)00138-4) PMID: [35378075](https://pubmed.ncbi.nlm.nih.gov/35378075/)
58. Lundberg SM, Erion G, Chen H, DeGrave A, Prutkin JM, Nair B, et al. From local explanations to global understanding with explainable AI for trees. *Nat Mach Intell.* 2020;2(1):56–67. <https://doi.org/10.1038/s42256-019-0138-9> PMID: [32607472](https://pubmed.ncbi.nlm.nih.gov/32607472/)